Gigamon®

GigaVUE Fabric Management Guide

Product Version: 6.14

Document Version: 1.0

(See Change Notes for document updates.)

Copyright © 2026 Gigamon Inc. All rights reserved.

Information in this document is subject to change without notice. The software described in this document is furnished under a license agreement or nondisclosure agreement. No part of this publication may be reproduced, transcribed, translated into any language, stored in a retrieval system, or transmitted in any form or any means without the written permission of Gigamon Inc.

Trademark Attributions

Gigamon and the Gigamon logo are trademarks of Gigamon in the United States and/or other countries. Gigamon trademarks can be found at www.gigamon.com/legal-trademarks. All other trademarks are the trademarks of their respective owners.

Gigamon Inc.
3300 Olcott Street
Santa Clara, CA 95054
408.831.4000

Change Notes

When a document is updated, the document version number on the cover page will indicate a new version and will provide a link to this Change Notes table, which will describe the updates.

Product Version	Document Version	Date Updated	Change Notes
6.14	1.0	07/10/2026	The original release of this document with 6.14.00 GA.

Contents

GigaVUE Fabric Management Guide **1**
Change Notes 3
Contents 4
Fabric Management **54**
GigaVUE-FM **54**
Supported Nodes **55**
Device and Configuration Management **56**
How to enable Web Server for Node Management **58**
Supported Browsers for GigaVUE-FM 58
GigaVUE-FM Features and Benefits **58**
Get Started with GigaVUE-FM **60**
Log In to GigaVUE-FM **61**
Log Out of GigaVUE-FM **61**
GigaVUE-FM **62**
Homepage 62
Left Navigation Pane 62
Top Navigation Bar 63
Page-Level Header 64
GigaVUE-FM GUI Navigation 64
Footer 64
Theme Settings 65

Configure a Custom Banner **65**
Quick Views **66**
Recently Viewed **66**
Find Pages **67**
Return to the Dashboard **67**
Table View Customization **67**
Notifications Panel **68**
Long-term Notifications **69**
Notification Type Icons **70**
Notification Banners **70**
Notification Banners for Disk Space Utilization 71
How to Add the GigaVUE-FM Instance Name **73**
How to Search in GigaVUE-FM **74**
Performing a Search **75**
Search Examples **77**
Search Maps 77
Example 1: Search for a Map by Alias 77
Example 2: Search for a Map with an IP Address 78
Example 3: Search for a Map with a MAC Address 80
Example 4: Search for Maps with Down Ports 80
Search for Roles and Users 81
Example 1: Search for Monitor Role 81
Example 2: Search for a User 83
Search Ports 84
Example 1: Search for Down Ports 84
Example 2: Search for Port Details of Devices Managed by GigaVUE-FM 85
Example 3: Search for Port and Linked Resources Using a Port Alias 86
Filtering Search Results **86**
How to Apply Filters **87**
Dashboard **88**
Physical and Virtual Dashboard **88**
Overview of the Physical and Virtual Dashboard **89**
Physical Dashboard Profiles **89**
Physical Dashboard Quick Views **91**
Telemetry Insights Dashboard **91**
Prerequisites 92

Executive Health Summary 92
Operational Summary 92
Recommended actions 93
View the Telemetry Insight dashboard 94
Rules and Notes 94
Physical Dashboard Widgets 94
Highest Traffic 95
Lowest Traffic 98
Traffic Comparison By Tags 99
Most Utilized Traffic 103
Least Utilized Traffic 106
Inventory 107
Nodes by Model 107
Nodes by Software Version 108
Status Summary 110
Nodes' Status Summary 110
Port Link Status Summary 111
Unhealthy Maps 112
Unhealthy Flows 113
Port Drops and Errors 113
Unhealthy Fabric Maps 116
Audit Logs 116
Events 119

FM Health Dashboard **122**
Overview of GigaVUE-FM Health Dashboard **123**
CPU Utilization **124**
Memory Utilization **124**
Storage Utilization **125**
Alarm Thresholds and Notifications **126**
GigaVUE-FM Reference Materials **127**
Disk Size on GigaVUE-FM **128**
Increase Disk Size on a New or Existing GigaVUE-FM Installation on KVM **128**
How to Clean up Disk Space on a GigaVUE-FM Instance **131**
Data Transfer Rate Units **132**
Open Ports in GigaVUE-FM **132**
Open Ports in HC Series Devices **136**
Health Status **138**
How GigaVUE-FM Computes Health Status **138**
Node Health Status **140**
Port Health Status **142**
Map Health Status **144**
GigaSMART Map Health Status **145**
GigaSMART Group Health Status 145
vPort Health Status 147
GigaSMART Operations Health Status 147
IP Interfaces Health Status 147
Flow Health Status **148**
Priority Map Set Health **149**
Flow Health Computation **150**
GigaVUE-FM APIs **151**
Supported Cloud Environments **151**
Analytics **153**
Rules, Notes, and Limitations for Analytics 153
GigaVUE-FM Statistics and Data Roll-up 155
Control Filters in Analytics Dashboards 155

Get Started with Analytics UI **156**
Work with the Analytics User Interface **159**
Share 159
Clone 160
Reporting 160
Set as Default 160
Copy Dashboard Path 160
Auto Refresh Tags in Visualizations 161
Filter Data Using Tags in Control Filters 161
Search Data 164
Filter Data in Visualizations **164**
Visualizations **167**
View System Visualization 168
Create Custom Visualizations **168**
Create a Visualization 168
Clone a Visualization 171
Visualizations - Example Work Flows **171**
TSVB Chart Displaying Traffic Trend 171
Pie Chart Displaying Alarms Summary 173
Metric Displaying Card Count 175
Bar Chart for Alarms by Severity 176
Dashboards **178**
Clone Dashboard **179**
Create New Dashboard **180**
Edit Dashboard **183**
Reports **184**
Discover **186**
Filter 188
Save Search 188
Find Data 189
Default Dashboards **189**
Dashboard **204**
Dashboards for Volume-based Licenses Usage **205**
Virtual Inventory Statistics and Cloud Applications Dashboard **208**
Rules and Notes 214

FHA Dashboards for 5G Applications **214**
Overall 5G Apps Dashboard 215
Detailed 5G Apps Dashboard 216
NokiaSCP Statistics 218
OracleSCP Statistics 219
EricssonSCP Statistics 219
NokiaHEP3 Statistics 220
Message Statistics 220
Protocol Statistics 221
Ericsson vTAP Statistics 221
Packet Statistics 221
TCP Statistics 222
Ericsson Packet Processing Statistics 223
GigaSMART Mobility Session and Flow Filtering Dashboards **224**
How to Access the Dashboards 224
Rules and Notes 227
GigaSMART Inline TLS/SSL Dashboards **227**
Basic Dashboards 228
Advanced Dashboards 234
System Requirements 234
Configure Advanced dashboard 235
Rules and Notes 235
OpenSearch Indices and Fields Used in Analytics **238**
Deprecated Fields in Analytics **242**
Analytics Appendix **243**
Ports 243
Maps 244
Map Rule 245
Nodes and Clusters **247**
GigaVUE Nodes and Clusters **247**
GigaVUE® HC Series and TA Series Overview **247**
Notes on TA Series Nodes 251
About Cluster **251**
Overview of Seed Node **252**
Manage GigaVUE® Nodes and Clusters **253**
Configure Physical Nodes **253**
Node Control Options 255

Add New Physical Node or Cluster to GigaVUE-FM **257**
Add Nodes Manually 257
Add Nodes From an Excel Spreadsheet 259
Device Level Tagging **260**
NAT Behind Nodes **261**
Prerequisites 261
Add NAT Behind Nodes 261
Limitations 262
Cluster Discovery Behavior **262**
ARP/NDP Timer Settings **262**
Change the ARP/NDP Timer Settings 263
Enable Gratuitous ARP on Management Interface **263**
SNMPv3 Support **264**
Enable SNMPv3 on Nodes 264
Create and Manage Clusters **265**
About Cluster **265**
Cluster Node Limit **266**
Cluster Scaling 267
Cluster Topologies 268
Separate Paths for Cluster Control and Stack Traffic **268**
About Cluster Roles **270**
Setting a Node's Priority in the Leader Election Process 270
About the "Unknown" Cluster Role 271
Sample Cluster Control Connections **272**
Sample Cluster Control Configurations 272
Zeroconf for Cluster Management Ports 272
Keep Cluster Management Ports Connected! 272
Sample Stack-Link Configurations **273**
Creating Clusters: A Roadmap **274**
IPv6 Based Clustering **276**
Rules, Notes, and Limitations 276
Create IPv6 based Cluster 277
Switch IPv4 based Cluster to IPv6 based Cluster 277
Rules and Recommendations for Nodes and Clusters **279**
Using Command Line Interface for Managing Clusters 281

GigaVUE TA Series and GigaVUE-HC3 Clustering Recommendations **284**
Cluster Rules **284**
Best Practices for OOB Clusters with IGMP Snooping **285**
About IGMP Snooping in a Cluster 285
Allow IGMP Traffic 286
Enable an IGMP Querier 286
E-Tag Clustering **287**
Stacking Mode 288
Switching Stacking Mode 289
Notes 290
New 6.0 290
Cluster Safe and Limited Modes **290**
Safe Mode **291**
Limited Mode **291**
Support for Cluster Types **292**
Create Clusters **292**
Regular Cluster Formation Workflow **292**
Deployment Checklist 292
Create Regular Cluster Formation 293
Customize Stack Links 295
Edit a Cluster 297
Prerequisites 298
Delete a Node from a Cluster 300
How to Change the Leader Preference of a Device 300
Edit Cluster **300**
Prerequisites: 302
Inband Cluster Management **302**
Inband Cluster Management Topologies **302**
Inband Cluster Management Stack Ports **303**
Inband Cluster Management Stack Ports Example 303

Inband Cluster Management Configuration Flow Chart 304
Inband Cluster Management Configuration 306
Enable Cluster Management for GigaVUE TA Series Nodes 306
Add Nodes to a Cluster 306
Remove Nodes from a Cluster 307
Edit Cluster Parameters 308
Check Cluster Status 309
Cluster Management Events 309
Audit Logs 310
Export Nodes and Clusters 310
Problems with SCP? 311
Events 311
Alarms 312
Audit Logs 312
Search for Specific Nodes Using Keywords 312
Search for Ports on a GigaVUE Node 313
Overview Page 315
Systems Information 315
Ports Information 316
Traffic 317
Workflows 317
Overview of Workflows 317
How to Use Workflows 318
Chassis Table View 320
Live Graphing 321
Safe and Limited Modes 321
Safe Mode 322
Limited Mode 323
Enable SNMP Trap for Safe Mode and Limited Mode 323
Collect Information for Technical Support 324
Manage Not-reachable Nodes in Cluster 324
Rules, Notes, and Limitations 326
Remove Offline Chassis 326

Alarms and Health Status of Not-reachable Nodes 327
Upgrade Cluster with Not-reachable Nodes 328
Backup and Restore of Cluster with Not-reachable Nodes 329
Multi-Path Leaf and Spine 330
Introduction to Multi-Path Leaf and Spine 330
Path Protection 332
Leaf Node Failure 332
Restoration 332
Affected Time 332
Stack Link Failure on Leaf (TAP Connected) 333
Restoration 333
Affected Time 333
Spine Node Failure 333
Restoration 334
Affected Time 334
Stack Link Failure on Leaf (Tool Connected) 334
Configuration Overview 335
Notes and Considerations 336
Leaf-Spine Cluster Deployment 337
Deployment Checklist 337
Pre-deployment checklist 337
Formation Scenario 337
Leaf-Spine Cluster Formation Workflow 338
Create a Leaf-Spine Cluster 338
Cluster Configuration 340
Stack Links configuration 340
Edit a Cluster 342
Prerequisites 343
Delete a Node from a Cluster 346
How to Change the Leader Preference of a Device 347

Spine to Spine and Leaf **349**
Introduction to Spine to Spine and Leaf **349**
Configuration Overview **351**
Notes and Considerations **353**
Configuration of Spine to Spine and Leaf Architecture **353**
Limitations **353**
Leaf-Spine Cluster Deployment **354**
Deployment Checklist **354**
Documentation 354
Pre-deployment checklist 354
Formation Scenario **355**
Fabric Statistics **355**
About Fabric Statistics **356**
Display Fabric Statistics for All Ports **356**
Display Fabric Statistics for a Single Port **359**
Port Quick View **359**
Export Fabric Statistics **360**
Filter Fabric Statistics **360**
Clear Fabric Statistics Counters 362
Enable Discovery Protocols **362**
Enable Gigamon Discovery on Chassis **363**
Enable LLDP and CDP on Chassis **363**
Enable LLDP, CDP, and Gigamon Discovery on Ports **364**
Limitations of Discovery **365**
Topology Visualization **365**
Edit Topologies 371

Configure Network Devices 380
Configure Connections 382
Configure Tools 384
Flows 386
About Flows 387
View Flows 390
View the Flow Summary and Statistics 391
View Maps and Ports 394
View Total Ports 395
View Total Unhealthy Ports 395
View Total Maps 396
View Unhealthy Maps 396
Filter Flows 396
How to Change the Flow Layout 397
How to Update Flows 397
View Events 398
Set Notifications 398
Limitations of Flows 399
Device Logs and Event Notifications 399
Stream Device Logs to GigaVUE-FM 400
Cluster Behavior 400
Standardized Logs 400
Device Log Categories 401
Device Log Message Types 401
Device Logging Levels 402
Device Logging Processes 403
View Device Logs 404
Arrange Columns in the Logs View 406
VLAN Resources 407
Device Log Host Servers 408
Add an External Logging Host Server to a Node 409
Host Server Options 410

Edit Host Server Settings **410**
Storage Management for Device Logs **411**
Access Storage Management **411**
Manage Device Log Output **412**
Traffic Filtering **414**
Ports and GigaStreams **415**
About Ports **415**
About Network and Tool Ports **415**
Network (Ingress) Ports Defined 416
Tool (Egress) Ports Defined 416
Ports on GigaVUE® TA Series Traffic Aggregator Nodes 417
Hybrid Ports 419
Stack Ports 420
Inline Network Ports 420
Inline Tool Ports 420
Circuit Ports 420
GigaSMART Engine Ports 421
Port Lists **421**
Port Aliases **421**
Work with Hybrid Ports **422**
Port Filters **424**
Port Filter—Rules and Notes 424
Port-Filter Maximums 425
How to Apply Port Filters 426
View Port Filter Statistics 427
View Filter Resources for a Slot 428
Status of Line Cards/Nodes and Ports **429**
How to Check Port Status with Ports Page 429
How to Check Port Status with Chassis Page 430
Managing Ports **430**
Ports **431**
All Ports 431
Port Quick View 434
Port List Filter 434
Quick Port Editor 437
Configure Ports 438
Ports Discovery 440
Statistics 441

Port Groups **442**
All Port Groups 442
Create Port Groups 443
Edit Port Group 443
Clone Port Group 444
Port Group Statistics 445
Port Pairs **445**
Create Port Pair 446
Tool Mirrors **447**
Create Tool Mirror 448
Edit Tool Mirror Description 448
Clone Tool Mirror 448
Stack Links **449**
IP Interfaces **450**
Configure IP Interface 451
Prerequisites 451
IP Interface Statistics 452
Circuit Tunnels **452**
Port Discovery **453**
Port Discovery with LLDP and CDP **453**
Enable Port Discovery **455**
Limits of Discovery Information **455**
Port Discovery Support **456**
Port Discovery for a Cluster 456
Port Discovery Supported for SNMP 456
Ingress and Egress VLAN **456**
About Ingress Port VLAN Tagging **457**
Ingress Port VLAN Tagging 458
Adding VLAN Tags 458
Deleting VLAN Tags 459
Using VLAN Tags in Maps **459**
Ingress Port VLAN Tag Limitations **460**
Second Level Maps 460
Double-Tagged Packets 460
IP Interfaces 460
Local Tool Port Ingress VLAN Tag 461
Configure Egress Port VLAN Stripping **461**
Enable Egress Port VLAN Stripping 461

Disable Egress Port VLAN Stripping 462
Display Egress Port VLAN Stripping 462
Egress Port VLAN Stripping Limitations 462
How to Use Both Ingress Tagging and Egress Stripping 463
How to Use GigaStream 463
GigaStreams 464
Accessing the GigaStream page 464
About GigaStream 465
Regular GigaStream 466
Controlled GigaStream 466
Regular GigaStream 466
Regular Tool GigaStream 467
Regular Stack GigaStream 468
Configure Regular GigaStream 469
Edit Regular GigaStream 470
Edit Regular Stack GigaStream 471
Traffic Distribution Across Regular GigaStream 471
Regular GigaStream Failover Protection—Resiliency 472
Regular Circuit GigaStream 472
Controlled GigaStream 473
Notes and Considerations for Controlled GigaStream 475
Controlled GigaStream Configuration 476
Edit Controlled GigaStream 478
Traffic Distribution Across Controlled GigaStream 479
Failover and Controlled GigaStream 481
Advanced Hashing 482
How to Change Advanced Hash Criteria 482
Advanced Hash Settings 482
Advanced Hash Examples 484
Hashing Behavior 486
Notes and Considerations for Advanced Hashing 488
Advanced Hashing with MPLS 488
Advanced Hashing with GTP TEID 490
Packet Distribution and the Advanced Hash Algorithm 492
Weighted GigaStream 492
GigaStream Rules and Maximums 493
Maximum Ports per GigaStream 494

Port Statistics and Counters **495**
Display Port Statistics **495**
Display Port Statistics for a Single Port 495
Display Statistics for All Ports 497
How to Clear Traffic Counters **499**
Header Stripping **499**
About VXLAN Header Stripping **499**
VXLAN Tunnel Decapsulation Versus VXLAN Header Stripping 501
VXLAN Header Stripping – Rules, Notes, and Limitations 501
Configure VXLAN Header Stripping **503**
Configure Ageing Interval for VXLAN Header Stripping **504**
View Header Stripping Statistics **504**
About MPLS Header Stripping **505**
Advanced MPLS Header Stripping for GigaVUE-TA400 506
MPLS Header Stripping – Rules, Notes, and Limitations 507
Configure MPLS Header Stripping **509**
Create a New Map 509
Configure MPLS Labels on a Chassis **512**
Enable Header Stripping Protocol on Ports **512**
Tunnels **515**
About Circuit-ID Tunnels **516**
Circuit-ID Tunnels—Rules and Notes **517**
Circuit-ID Tunnel Encapsulation **518**
Circuit-ID Tunnel Decapsulation **518**
About Layer 2 Generic Routing Encapsulation (L2GRE) Tunnels **519**
L2GRE Tunnel Configuration—Rules and Notes 520
Limitation 521
Configure L2GRE Tunnel to Encapsulate Traffic **521**
Configure L2GRE Tunnel to Decapsulate Traffic **522**
Orchestrated Workflow Configuration of L2GRE Tunnels **523**
About Virtual Extensible LAN (VXLAN) Tunnels **525**
VXLAN Tunnel Configuration—Rules and Notes 525

Configure VXLAN Tunnel to Encapsulate Traffic **527**
Configure VXLAN Tunnel to Decapsulate Traffic **528**
Orchestrated Workflow Configuration of VXLAN Tunnels **529**
Create Tunnel **530**
Create VXLAN / L2GRE Group **531**
Configure L2GRE / VXLAN Identifier **532**
Configure L2GRE / VXLAN Identifier for a Chassis 532
Configure L2GRE / VXLAN Identifier for a Network Port 533
View VXLAN / L2GRE ID Statistics **533**
Tunnel Monitoring **535**
Rules, Notes, and Limitations 536
View Tunnel Monitoring 537
Packet Capture (PCAP) **539**
Rules, Notes, and Limitations 540
Configure PCAP Profile 542
View PCAP 545
Delete PCAP 545
View PCAP Files 545
Configure PCAP on IP Interface 546
View PCAP on IP Interface 547
Delete PCAP on IP Interface 547
Flow Mapping® **547**
About Flow Mapping® **548**
Flow Mapping® Overview **548**
Get Started with Flow Mapping® **549**
Check Status of Nodes and Ports 550
Designated Port Types **550**
About Shared Collectors 550
No Map Statistics for Shared-Collector Only 551
Shared Collector Configuration 551
About Map-passall Maps 552
Map-Passall and Regular Byrule Map 553
Map-Passall and Shared-Collector Only 553
Map-passalls Configuration 553
Map-passalls 553
Define Map Source Port Lists 553
Share Network Ports Between Maps 554

Share Tool Ports Between Maps 554
Map Priority 554
Adjust Map Priority in GigaVUE-FM 555
Flow Map Syntax and Construction 555
Map Types 555
Define Map Source Port Lists for First Level Maps 556
Null Port in Maps 559
Map Subtypes 559
Map Type and Subtype Modification 560
Backwards Compatibility 560
Minimum Requirements for Map Creation 560
Map Rules 560
Other Types of Map Rules for GigaSMART Operations 561
IPv4/IPv6 and Map Rules 561
Set Map Rules for TCP Control Bits 563
How to Use Bit Count Netmasks 565
How to Combine Rules and Rule Logic 567
How to Mix Pass and Drop Rules 567
Configuring Port Criteria and Bi-Directional Rules in By Rule Maps 568
Work with User-Defined Pattern Match Rules 569
Inner Header and MPLS Header Filtering 574
How to Handle Q-in-Q Packets in Maps 577
Comparison of Q-in-Q Tagging 578
Priority Tagged Packets 579
Flow Mapping® on Inner VLAN Tags 579
Inner VLAN Limitation 580
Work with Map-Passalls and Port Mirroring in GigaVUE-FM 581
Syntax for Maps-passalls and Port Mirroring 581
Rules for Map-Passalls and Port Mirroring 581
View and Delete Map-passalls 582
Port Access and Map Sharing 582
Port-based Access Levels 582
How to share Maps 583
Map Examples 584
Example: How to Create a Simple Map 585
Example: How to Handle Overlaps when Sending VLANs and Subnets to Different Tools 586

Manage Maps **588**
Map Views **588**
List View 588
Map Topology View 589
Quickview 590
Free Layout 591
Reset to Default Layer 592
Manage Maps **592**
Create a new map 593
Clone Map 596
Edit Maps 596
Delete Maps 597
Create Map Groups 597
Description to Map Rules 598
Error Messages 598
Edit Map Rule Description 598
How to Use the Quick Editor for Pass and Drop Rules 599
How to Use the Quick Editor to Add Port Numbers 599
How to Use the Quick Port Editor to Add IP Address 601
Map Templates **601**
Create Map Templates 601
Edit Map Templates 602
Map Template Quick View 602
Manage Map Rule Resources **602**
Template Groups 602
Add Tags to Map Rules **605**
Notes 606
Map Rule Statistics Dashboard 606
Filter Templates **607**
Filter Template Qualifiers and Defaults 608
Filter Template Configuration 609
Filter Template Limits 609
How to Understand Map Filter Resources 609
Filter Template Rules and Recommendations 610
Filter Template Best Practices 610
Filter Templates in a Cluster 611
Filter Templates Formulas 612
Troubleshoot Custom Filter Template Issues 614
Review Map Statistics with Map Rule Counters **616**
Viewing Map Statistics with the Statistics Page 617

Viewing Map Statistics with Quick View 617
Flow Mapping® FAQ 617
How Many Map Rules are Supported? 617
How Many Rules Can Each Map Have? 622
How Many Maps Can Run at Once? 623
What Criteria can be Filtered in Q-in-Q Packets? 623
How Many Maps Can Share a Network Port? 623
How Many Network Ports and Tool Ports Can Be in a Map? 623
Are Port-Filters Supported? 623
Does Flow Mapping® Support Passalls? 624
Does Flow Mapping® Support port-pairs? 624
Does Flow Mapping® Support UDA Pattern Matches? 625
Can a GigaStream Act as a Shared Collector? 625
What Are the GigaStream Maximums? 625
What order are the map rules displayed in "show running config"? 626
Active Visibility 626
MAC Address Rewrite 629
Configuring MAC Address Re-write 630
License 631
Limitations 631
IP Address Rewrite 632
Configuring IP Address Re-write 633
License 634
Limitations 634
VLAN Manipulation 634
Overview 635
Rule-Level Add VLAN Manipulation 635
Rule-Level Modify VLAN Manipulation 636
Rule-Level Strip VLAN Manipulation 637
Map-Level VLAN Manipulation 638
Limitations 641
Monitor Port Utilization 641
Port Utilization Availability by Port Type 642
Set Port Utilization Thresholds 642
Utilization Alarm/SNMP Trap Generation 643
Configure Alarm Buffer Thresholds 644
Set Alarm Buffer Thresholds 645
Microburst 646
Best Practices to Improve Burst Tolerance 647

Logical Grouping of Ports in GigaVUE-HC3, GigaVUE-TA100, and GigaVUE-TA200 647
Flexible Inline Arrangements 648
Flexible Inline Arrangements 649
Flexible Inline Arrangement vs Inline Bypass Solution 650
Supported Platforms 651
Flexible Inline Arrangement License 652
Software Version 652
Flexible Inline Solution Supported in Clustered Nodes 652
Limitations 654
Functionalities Not Supported by Flexible Inline Arrangement 654
Benefits of Flexible Inline Arrangements 654
About Flexible Inline Maps 655
Types of Flexible Inline Maps 656
Configure Flexible Inline Maps 657
Flexible Inline Arrangements—Rules and Notes 657
Visualize the Flexible Inline Arrangements Canvas 658
Configure Flexible Inline Flows 659
Troubleshoot Flexible Inline Flows 659
Status 659
Rules and Notes 660
Statistics 661
HSM Statistics 662
ICAP Show Stats 663
Troubleshoot 665

Example: Troubleshoot Traffic Issues Between Side A and Side B **666**
View the Forwarding States of Inline Networks **667**
Import and Export Flexible Inline Solution **674**
Backup and Restore Flexible Inline Flows **675**
Inline Bypass Solutions **676**
About Inline Bypass Solutions **676**
Introduction to Inline Bypass Solutions **677**
Capabilities of Inline Bypass Solutions **677**
Logical Bypass and Physical Bypass **678**
Configure Inline Bypass Solutions **680**
Inline Networks **681**
Unprotected Inline Network 681
Protected Inline Network **682**
Mix of Protected and Unprotected 683
Network Port Link Status Propagation Parameter 684
Traffic Path Parameter 684
Physical Bypass Parameter 687
Redundancy Profile Parameter 687
Display Current State of Inline Bypass Solution 688
How to Use SNMP Polling to Obtain Inline Network State 692
SNMP Notification of Forwarding State Change 693
How to Use Syslog to Obtain Inline Network State 693
Inline Network Groups **693**
Inline Tool Sharing 695
Configurable VLAN Tagging 695
Add VLAN Tag 696
Tools in Bridge Mode 696
Inline Tools **697**
Inline Tool Failover Action 698
Inline Tool Failover Action with Inline Flow Mapping® 699
Inline Tool Recovery Mode 699
Inline Tool Sharing Mode 701
How to Use SNMP Polling to Obtain Inline Tool State 701
Heartbeats **702**
Rules and Notes 702
Heartbeat Profiles 703
Standard Heartbeat 704

Standard or Custom Heartbeat Packet 704
Detect Inline Tool Failure 705
Negative Heartbeat Profiles 705
Heartbeat Support Between GigaVUE Nodes 706
Troubleshoot Heartbeat Failures 708
Configure a Heartbeat Profile 708
View Heartbeat Profile Statistics 709
Heartbeat Status after System Reload 710
Inline Tool Groups 710
Inline Tool Group Failover Action 713
Inline Tool Group Spare Inline Tool 714
Symmetrical and Asymmetrical Hashing 714
Asymmetrical Hashing Restrictions 716
Resilient Weighted Hashing 717
Inline Serial Tools 717
Inline Serial Tools Global Failover Action 720
Inline Tool Series Local Failover Action 722
Inline Tool Series Per-Direction Order 724
Associate Inline Networks with Inline Tools Using Inline Maps 726
Inline Map Passall 727
Inline Map 727
Inline Map Shared Collector 728
Inline Maps to Individual Members of an Inline Tool Group 729
Map, Inline Tool, and Inline Tool Group Configuration Restrictions 730
Inline Tool Failures and Failover Actions 731
Maps That May Lead to Selective Traffic Drops 732
Out-of-Band (OOB) Map 732
Symmetric and Asymmetric Maps 734
Configuration Steps 734
Configuration Step Details 737
Configure Inline Network Ports 737
Configure Inline Network (Unprotected) 737
Configure Inline Network Group 738
Create Heartbeat Profile 738
Create Negative Heartbeat Profile 739
Configure Inline Tool Ports 740
Configure Inline Tool 741
Create Inline Tool Group 742
Configure Inline Tool Series 743
Configure When GigaVUE® HC Series Modules are Operationally Up 744

Avoid Over subscription 744
Configure Gigamon Resiliency for Inline Protection 744
How to Cable GigaVUE Nodes 745
How to Handle Recovery 746
Both Nodes Go Down and Only Secondary Comes Up 746
How to Cable GigaVUE Nodes 746
Configure GRIP Solution Software 746
Limitation for Suspended Role 748
Configure Synchronization 748
Display Redundancy Control State 749
How to Use Suspended Role for Maintenance 749
Upgrade Procedure Recommendations 750
Rules and Notes 750
Limitations 751
Troubleshoot 751
Signaling Ports Down 751
Traffic outage in Inline Tool 751
Network Traffic Outage 752
FAQs 752
Migrate Inline Bypass(Classic) Solution from version 6.5 to 6.12 753
Migration Scenarios 753
Migration Considerations 753
Migrate Classic Inline Bypass Solution using GigaVUE-OS CLI 754
Migrate Classic Inline Bypass Solution using GigaVUE-FM 755
Inline Bypass Solution Examples 756
Example 1: Unprotected Inline Bypass with an Inline Tool Group 757
Example 2: Unprotected Inline Bypass with Default Heartbeat 759
Example 3: Protected Inline Bypass Using Combo Modules 760
Example 4: Gigamon Resiliency for Inline Protection 763
Configure Primary Role GigaVUE-HC2 763
Configure Secondary Role GigaVUE-HC2 765

Timestamps **767**
About Timestamps **768**
Using Timestamps **769**
Why PTP? **769**
Synchronization of the PTP and Local System Clock 771
Using PTP to Timestamp Packets **772**
Configuring PTP Globally **772**
Enabling PTP on a Network Port **773**
Enabling Timestamp on a Port **774**
Viewing PTP Details **775**
Viewing the PTP Configuration on a Port **776**
Viewing the PTP Clock Details **778**
Viewing the PTP Foreign Source Details **782**
Viewing PTP Statistics **783**
PTP and Timestamp—Rules and Notes **784**
Fabric Maps **785**
Supported Topologies **785**
Multihop Topology 786
Supported Multihop Topology 787
Unsupported Multihop Topology 787
Leaf and Spine Topology 788
Spine Source Map Mode 789
Dual Multipath Leaf and Spine Topology 789
Supported Failovers 790
Scenario 1 791
Scenario 2 792
UnSupported Failovers 792
Scenario 1 793
Scenario 2 794
Fabric Maps Prerequisites **794**
Notes on Circuit Ports and Circuit GigaStream in Fabric Maps **795**
Notes for Circuit ID 796
797
Create Links between Clusters **797**
Create Manual Links 797
Create Gigamon Discovery Protocol (GDP) Based Links 798

Create Fabric Maps **799**
Edit and Delete Fabric Maps **802**
Edit Fabric Maps **803**
Edit Fabric Map Component 803
Delete Fabric Maps **804**
Prioritize Fabric Maps **805**
Configure Fabric Map Settings **806**
Share Circuit ID Resource 806
Spine Source Map Mode 806
Limitations of Spine Source Map Mode 807
Fabric Maps Statistics **808**
Display Fabric Map Statistics **808**
Display Fabric Map Details **808**
Filter Fabric Maps List View **809**
Fabric Port Group **810**
How Fabric Port Groups Work? 810
Rules and Notes **811**
Create Fabric Port Group **812**
Configure Fabric Map Using Fabric Port Group 813
Troubleshooting **814**
How to Troubleshoot Partial Success Errors **817**
Example 1: 817
Config Audit **819**
Limitations of Fabric Maps **820**
Backup and Restore Fabric Maps and Orchestrated Configurations **821**
Orchestrated Configurations **821**
About Intent Based Orchestrated Configurations 822
Benefits of Orchestrated Configurations 823
Orchestrated Configuration: Examples 824
Priority-Free Policies 824
Overlapping Sources in Policies 827
Supported Topologies 828
Rules and Notes for Orchestrated Configurations 828
Create Orchestrated Policies 829
Prerequisites 829
How to Create a Policy 830

Tagging and RBAC Support 832
Drop Rules 833
Import and Export Orchestrated Policies 835
Rules and Notes 837
Import and Export Orchestrated Configuration 837
Egress Filters for Additional Filtering Capabilities 838
Traffic Policy 840
Map Components in a Traffic Policy 840
How Traffic Policy Choose Map Types 842
Map Priorities in Traffic Policy 843
Priority Options 844
How Traffic Policy Prioritize Maps 844
How Map Priority Affects Traffic Flow 844
Best Practices for Map Priority 845
Source Overlap Behavior in Traffic Policy 845
Traffic Policy Life Cycle 845
Best Practices for Traffic Policy 846
Important Notes 846
Limitations 846
Explore Traffic Policy Interface 847
Overview 847
Dashboard 848
Map Details Grid 849
Menu 853
View Migration Results 853
Filters 854
Settings 855
Actions 855
Config Audit 856
Resubmit All 856
Edit 856
Delete 856
Clone 857
Deploy 857
Undeploy 857
View Details 857
View Statistics 859
View Troubleshoot 860
Resources Tab 860
Fabric Port Groups 860
Tool Templates 860

Create and Manage Fabric Port Group **861**
Create Fabric Port Group 861
Edit Fabric Port Groups 862
Delete Fabric Port Groups 862
Export Fabric Port Groups 862
Configure Traffic Policy using Fabric Port Groups 862
Configure a Virtual Port 863
Configure a First-Level Map: 863
Configure a Second-Level Map: 864
Manage Rulesets across Traffic Policies **866**
Copy Rules from a Source Policy 866
Paste Rules into a Target Policy 867
Manage Rulesets across Traffic Policies — Rules and Notes 868
Configure Traffic Policy Without GigaSMART Operations **869**
Configure VLAN Manipulation 871
Map-Level VLAN Manipulation 872
Configure Traffic Policy for Traffic Intelligence Solutions **874**
GigaSMART Adaptive Packet Filtering (APF) **874**
Identify Social Security Numbers in User-Level Transactions **875**
Configure a GigaSMART Group 876
Configure a Virtual Port 876
Configure Traffic Policy for a First-Level Map 877
Configure Traffic Policy for a Second-Level Map 878
Mask Social Security Numbers **880**
Configure a GigaSMART Group 881
Configure a Virtual Port 881
Configure Traffic Policy for a First-Level Map 882
Configure Traffic Policy for a Second-Level Map 883
Filter on Fiber Channel over Ethernet (FCOE) Traffic **885**
Configure a GigaSMART Group 885
Configure a Virtual Port 886
Configure Traffic Policy for a First-Level Map 886
Configure Traffic Policy for a Second-Level Map 888
Filter on Fiber Subscriber Device IP (User-Endpoint IP or UE-IP) **889**
Configure a GigaSMART Group 890

Configure a Virtual Port 891
Configure Traffic Policy for First-level Maps 891
Configure Traffic Policy for Second-Level Maps 893
Filter on Inner Layer 2-4 Parameters for Unrecognized Headers 895
Configure a GigaSMART Group 895
Configure a Virtual Port 896
Configure Traffic Policy for a First-Level Map 896
Configure Traffic Policy for Second-Level Maps 897
GTP Tunnel ID-Based Filtering 899
Configure a GigaSMART Group 901
Configure a Virtual Port 901
Configure Traffic Policy for a First-Level Map 902
Configure Traffic Policy for Second-Level Maps 903
ERSPAN Tunneling 905
Configure a GigaSMART Group 906
Configure a Virtual Port 906
Configure Traffic Policy for a First-Level Map 907
Configure Traffic Policy for a Second-Level Map 908
Distribute Traffic Based on Inner IP Addresses and Inner TCP Port values 910
Configure a GigaSMART Group 911
Configure a Virtual Port 912
Configure Traffic Policy for a First-Level Map 912
Configure Traffic Policy for Second-Level Maps 913
MPLS Label-Based Filtering 916
Configure a GigaSMART Group 917
Configure a Virtual Port 917
Configure Traffic Policy for a First-Level Map 918
Configure Traffic Policy for Second-Level Maps 919
To create a second-level map to filter based on MPLS label 921
Conditional Header Stripping 921
Configure a GigaSMART Group 922
Create a Virtual Port 923
Configure Traffic Policy for a First-Level Map 923
Configure Traffic Policy for Second-Level Maps 924
Facilitate Overlapping Rules 926
Configure a GigaSMART Group 927
Configure a Virtual Port 928
Configure Traffic Policy for a First-Level Map 928

Configure Traffic Policy for Second-Level Maps 930
Combine APF with GigaSMART Operations 931
Configure a GigaSMART Group 933
Configure a Virtual Port 933
Configure Traffic Policy for a First-Level Map 933
Configure Traffic Policy for Second-Level Maps 935
GigaSMART Load Balancing 937
Stateful Load Balancing 937
Configure a Port Group 938
Configure a GigaSMART Group 939
Configure a Virtual Port 939
Configure a Traffic Policy for First-Level Map 940
Configure a Traffic Policy for Second-Level Map 941
Configure a Port Group 943
Configure a GigaSMART Group 944
Configure a Virtual Port 944
Configure Traffic Policy for a First-Level Map 944
Configure Traffic Policy for a Second-Level Map 946
Stateless Load Balancing 948
Configure a Port Group 948
Configure a GigaSMART Group 949
Configure a Traffic Policy for Regular Map 949
Configure a Port Group 951
Configure a GigaSMART Group 951
Configure Traffic Policy for a First-Level Map 952
Configure a Port Group 954
Configure a GigaSMART Group 954
Configure a Virtual Port 954
Configure Traffic Policy for a First-Level Map 955
Configure Traffic Policy for a Second-Level Map 956
Enhanced Load Balancing 958
GigaSMART De-Duplication 958
Configure Traffic Policy 959
GigaSMART Masking 962
Configure a GigaSMART Group 963
Configure a Traffic Policy for Masking 963
GigaSMART NetFlow Generation 965
Configure a GigaSMART Group 966
Configure a NetFlow Exporter 966
Configure an IP Interface 967

Configure a NetFlow Record 968
Configure a NetFlow Monitor 969
Add the NetFlow Monitor to GigaSMART Group 970
Configure a Virtual Port 970
Configure Traffic Policy for a First-Level Map 971
Configure Traffic Policy for a Second-Level Map 972
GigaSMART Packet Slicing 973
Configure a Traffic Policy for Packet Slicing. 974
GigaSMART Tunneling Operations 976
GigaSMART IP Encapsulation/Decapsulation 976
Example: GigaSMART Encapsulation/Decapsulation (GigaVUE V Series 977
Configure Sending End of Tunnel: GigaVUE V Series map rules in Sydney 977
Configure Receiving End of Tunnel: GigaVUE HC Series with GigaSMART in Melbourne 978
Configure a GigaSMART Group 978
Configure an IP Interface 978
Configure a Traffic Policy 979
GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation 981
Layer 2 GRE Header Length 982
GigaSMART L2GRE Tunnel Encapsulation 983
Configure a GigaSMART Group 983
Configure an IP Interface 984
Configure a Traffic Policy 984
GigaSMART L2GRE IPv6 Tunnel Encap/Decap with Load Balancing 986
Configure a Port Group 987
Configure GigaSMART Group 987
Configure an IP Interface 988
Configure a Traffic Policy 988
GigaSMART L2GRE Tunnel Encapsulation Stateless Load Balancing 990
Configure a GigaSMART Group 990
Configure an IP Interface 991
Configure a Tunnel Endpoint 991
Configure a Port Group 992
Configure a Traffic Policy 992
GigaSMART L2GRE Tunnel Decapsulation 994
Configure a GigaSMART Group 995

Configure an IP Interface 995
Configure a Traffic Policy 996
GigaSMART L2GRE Tunnel Encapsulation Load Balancing 997
Create IP Interface 998
Configure a Tunnel Endpoint 999
Configure a Port Group 999
Configure Traffic Policy 1000
GigaSMART VXLAN Tunnel Decapsulation 1002
Configure a GigaSMART Group 1002
Configure an IP Interface 1003
Configure a Traffic Policy 1003
GigaSMART ERSPAN Tunnel Decapsulation 1005
Configure an IP Interface 1006
Configure a Port Group 1006
Configure a GigaSMART Group 1006
Configure a Traffic Policy for Regular Map. 1007
GigaSMART Custom Tunnel Decapsulation 1008
Configure a GigaSMART Group 1009
Configure an IP Interface 1010
Configure a Traffic Policy 1010
Configure Traffic Policy for Application Intelligence Solutions 1012
How Traffic Policy Choose Deployment Type for Application Intelligence Solutions 1013
Application Intelligence – Rules and Notes 1014
Create and Manage Custom Tool Template 1015
Create Custom Tool Template 1015
Edit Tool Template 1018
Delete Tool Template 1018
Clone Tool Template 1018
Export Tool Template 1018
Import Tool Template 1019
Application Intelligence Field References 1019
Application Intelligence Session 1019
Application Filtering Intelligence (AFI) 1020
Application Metadata Intelligence (AMI) 1022

Use Case: Reduce Tool Overload **1025**
Use Case: Enhance Context for Incident Response **1028**
Use Case: Maintain Visibility Whilst Filtering Traffic **1032**
Configure the Traffic Policy for Subscriber Intelligence Solutions **1037**
GigaSMART GTP and CUPS Correlation **1038**
Identifying High-Value and/or Roaming Subscribers Based on IMSIs **1039**
Configure a GigaSMART Group 1040
Configure a Virtual Port 1041
Configure a First-Level Map 1041
Configure Second-Level Maps 1043
Identifying GTP Versions **1044**
Configure a GigaSMART Group 1046
Configure a Virtual Port 1046
Configure First-Level Maps 1047
Configure Second-Level Maps 1048
Same Subscriber Filter on Different Versions **1050**
Configure a Virtual Port 1052
Configure Port Groups 1052
Configure an Ingress First-Level Map 1052
Configure Traffic Policy for Second-Level Maps 1054
Same Subscriber Filter on Different Interfaces **1056**
Configure a GigaSMART Group 1057
Configure a Virtual Port 1057
Configure Port Groups 1058
Configure First-Level Maps 1058
Configure Second-Level Maps 1060
EPC Filtering in Scenario 1 ..... **1062**
Configure a GigaSMART Group 1062
Configure a Virtual Port 1063
Configure a Port Group 1063
Configure a First-Level Map 1063
Configure a Second-Level Map 1065
EPC Filtering in Scenario 2 ..... **1067**
Configure a Virtual Port 1068
Configure a Port Group 1068
Configure a First-Level Map 1068
Configure a Second-Level Map 1070

GigaSMART Rotational Sampling 1072
Configure a GigaSMART Group 1072
Configure a Virtual Port 1073
Configure a Second-Level Map 1073
GigaSMART 4G RAN Correlation 1075
Configure Port Groups 1076
Create a GigaSMART Group 1076
Configure a Second-Level Map 1077
GigaSMART SIP/RTP Correlation 1079
SIP/RTP Load Balancing Example 1079
Configure a Port Group 1079
Configure a GigaSMART Group 1080
Configure a Virtual Port 1080
Create a SIP Whitelist 1081
Associate the SIP Whitelist to GigaSMART Group 1081
Configure Traffic Policy 1082
Create First-Level Maps 1082
Create Second-Level Maps 1083
GTP Overlap Flow Sampling 1085
Configure GigaStream 1087
Configure Port Groups 1087
Create a Virtual Port 1088
Create a GTP Whitelist 1089
Associate the GTP Whitelist to GigaSMART Group 1089
Configure Traffic Policy 1090
Create a First-Level Map 1090
Create Second-Level Maps 1091
GigaSMART GTP Whitelisting and GTP Flow Sampling 1094
GTP Whitelisting 1094
Configure a Port Group 1095
Configure a GigaSMART Group 1096
Create a Virtual Port 1096
Create a GTP Whitelist 1097
Associate the GTP Whitelist to GigaSMART Group 1097
Configure Traffic Policy 1097
Create First-Level Maps 1097
Create a Second-Level Map 1099
GTP Whitelisting with Multiple Maps 1101
Configure a Port Group 1102
Create a Virtual Port 1103
Create a GTP Whitelist 1103

Associate GTP Whitelist to GigaSMART Group 1104
Configure Traffic Policy 1104
Create First-Level Maps 1104
Create Second-Level Maps 1106
GTP Flow Sampling 1108
Configure a Port Group 1109
Create a GigaSMART Group 1109
Create a Virtual Port 1110
Create a GTP Whitelist 1110
Associate the GTP Whitelist to GigaSMART Group 1111
Configure Traffic Policy 1111
Create First-level Maps 1111
Create a Second-level Map 1113
GTP Whitelisting, GTP Flow Sampling, and Load Balancing 1115
Configure a Port Group 1116
Create a Virtual Port 1117
Create a GTP Whitelist 1118
Associate the GTP Whitelist to GigaSMART Group 1118
Configure Traffic Policy 1119
Create First-level Maps 1119
Create a Second-level Map 1120
GTP Flow Sampling with Multiple Maps 1123
Configure a Port Group 1126
Create a Virtual Port 1127
Create a GTP Whitelist 1127
Associate the GTP Whitelist to GigaSMART Group 1128
Configure Traffic Policy 1128
Create First-Level Maps 1128
Create Second-Level Maps 1130
APN for GTP Forward Listing and GTP Flow Sampling 1134
Configure a Port Group 1136
Create a Virtual Port 1137
Create a GTP Whitelist 1137
Associate GTP Whitelist to GigaSMART Group 1138
Configure Traffic Policy 1138
Create First-level Maps 1138
Create a Second-level Map 1140
APN for GTP Whitelisting, APN and QCI for GTP Flow Sampling 1142
Configure a Port Group 1143

Create a Virtual Port 1144
Create a GTP Whitelist 1145
Associate GTP Whitelist to GigaSMART Group 1145
Configure Traffic Policy 1145
Create First-level Maps 1145
Create a Second-level Map 1147
4G/5G Traffic Monitoring using UPN 1149
Configure a Port Group 1151
Configure a GigaSMART Group 1151
Create a GTP Whitelist 1152
Associate GTP Whitelist to GigaSMART Group 1152
Configure a Virtual Port 1153
Configure Traffic Policy 1153
Create First-level Maps 1153
Create Second-Level Maps 1155
GigaSMART 5G CUPS 1157
Configure a Port Group 1159
Configure a GigaSMART Group 1159
Configure an IP Interface 1160
Configure TCP/IP Host 1160
Create a GTP Whitelist 1161
Associate GTP Whitelist to GigaSMART Group 1161
Configure a Virtual Port 1162
Configure Traffic Policy 1162
Create First-level Maps 1162
Create Second-Level Maps 1165
GTP Scaling 1167
GTP Engine Grouping Configuration with 2 Engine Ports 1168
Create a GigaSMART Group 1169
Create a Virtual Port 1169
Create First-level Maps 1169
Create a Second-level Map 1171
GTP Engine Grouping Configuration with 4 Engine Ports 1173
Configure GigaStream 1175
Configure a Stack Link 1176
Configure Port Groups 1176
Create a GigaSMART Group 1176
Create a Virtual Port 1177
Create a GTP Whitelist 1177

Associate the GTP Whitelist to GigaSMART Group 1178
Configure Traffic Policy 1178
Create First-level Maps 1178
Create a Second-Level Map 1180
Migrate Maps to Traffic Policy 1183
Dry Run before Migration 1183
Migrating Traffic Intelligence and Subscriber Intelligence Solutions 1184
Migrating Flow Maps 1184
Migrating Fabric Maps 1185
Migrating Application Intelligence Solutions 1186
View Migration Results 1186
Configure Map Settings 1187
Share Circuit ID Resource 1188
Spine Source Map Mode 1189
Limitations of Spine Source Map Mode 1189
Configuration Sync 1189
Troubleshoot Traffic Policy Issues 1190
View the Troubleshooting Page 1190
Troubleshooting Statistics and Health Indicators 1191
Policy Level Health 1191
Component Level Health 1191
FAQ Traffic Policy 1192
GigaSMART® 1194
GigaSMART Operations 1194
About GigaSMART® Applications 1195
Quick Glance- How to Configure a GigaSMART Application 1195
Application Intelligence Solutions 1197
Subscriber Intelligence 1198
Installation and Configuration of Subscriber Intelligence Solution using Ansible 1198
System Requirements 1199
Installation and Configuration of Gigamon Ansible Module 1199
Deployment Report 1200
Check Mode 1200
Reapplying Golden Payload 1201
Rules and Notes 1201
Configuration of Non-CUPS using Ansible 1201
Prerequisite 1206

Site 1206
cpNode 1206
upNode 1207
5GPolicy 1207
LTEPolicy 1207
Deployment of Non-CUPS Solution 1207
GigaSMART GTP and CUPS Correlation 1208
Filtering on Subscriber IDs and Version 1209
Session Correlation 1212
Supported Interfaces 1213
Conditional S10 Support 1215
GTP Session Timeout 1215
Priorities for Flow Rules and Maps 1216
GTP Correlation Configuration Examples 1219
GigaSMART Rotational Sampling 1238
Entire session sampling for Mobile subscribers 1239
Dynamic Session Sampling for Fixed Wireless Access Subscribers 1240
Supported Platforms Compatibility 1242
Configuring Rotational Sampling in GigaVUE-FM 1242
Viewing the Configuration of Rotational Sampling for non-CUPS LTE and UPN 1245
GigaSMART 4G RAN Correlation 1247
Rules and Notes 1248
GigaSMART SIP/RTP Correlation 1248
CallerID Tracking 1250
Support for SIP, RTP, and RTCP 1250
SIP/RTP Correlation Engine 1250
Configure SIP/RTP Correlation Engine 1251
SIP Whitelist 1252
RTP Flow Sampling 1252
Support for Sessions 1252
Support for IPv4 and IPv6 1253
Support for Content Masking 1253
Behaviors of Some SIP Methods 1254
SIP Whitelisting in a Cluster 1254
Support for NAT 1254
Not Supported by SIP/RTP Correlation 1255
Display SIP/RTP Reports 1255
Display SIP Map Statistics 1256
SIP/RTP Support for Tool Throttling 1256
Admission Control 1257

SIP/RTP Examples **1258**
SIP/RTP Load Balancing Example 1258
SIP/RTP Minimum Configuration Example 1258
GigaSMART IP FlowVUE **1266**
Configure FlowVUE 1267
FlowVUE Configuration Examples 1268
Display FlowVUE Statistics 1268
FlowVUE_Examples **1268**
GigaSMART GTP Whitelisting and GTP Flow Sampling **1274**
GTP Whitelisting 1275
GTP Flow Sampling 1281
GTP Subscriber Aware Random Sampling 1286
GigaSMART GTP Whitelisting and GTP Flow Sampling Examples **1287**
4G/5G Traffic Monitoring using UPN **1313**
Configure 4G/5G Traffic Monitoring using UPN 1315
PFCP Messages 1315
PFCP Load Balancing 1316
Map Configuration 1316
Stateful Session Recovery 1317
5G Stateful Session Recovery 1318
Configure 5G Stateful Session Recovery 1319
View Backup and Restore Information 1319
Configure GTP Overlap Mapping 1320
Configuration Considerations 1321
About Standalone UPN Flow Sampling Map Mode and Port Groups 1321
Standalone UPN Flow Sampling Map Priority 1323
Virtual Port Configuration in Standalone UPN Mode 1324
Overlap Support 1324
Standalone UPN **1324**
PFCP (Packet Forwarding Control Protocol) 1324
PFCP Messages 1325
PFCP Load Balancing 1325
Map Configuration 1326
Stateful Session Recovery 1326
5G Stateful Session Recovery 1327
Configure 5G Stateful Session Recovery 1328
View Backup and Restore Information 1329
Configure GTP Overlap Mapping 1330
Configuration Considerations 1331

About Standalone UPN Flow Sampling Map Mode and Port Groups 1331
Standalone UPN Flow Sampling Map Priority 1333
Virtual Port Configuration in Standalone UPN Mode 1334
Overlap Support 1334
GigaSMART Rotational Sampling Support 1334
Interface Filtering and APN/DNN Filtering 1334
Interface Filtering 1334
APN (Access Point Name)/DNN (Data Network Name) Filtering 1336
Custom Interface Selection 1337
GigaSMART 5G CUPS 1339
Overview 1339
Supported Platforms 1342
5G Correlation 1342
5G Load Balancing 1343
Configure 5G Load Balancing 1343
5G RAN Correlation 1344
Configure 5G RAN Correlation 1345
5G Network Slice Correlation 1345
Configure 5G Network Slice Correlation 1347
5QI Correlation 1347
Rules and Notes 1347
5G Flow Sampling and Filtering 1347
About Flow Sampling Rules and Maps 1348
GTP Overlap Flow Sampling for 4G and 5G 1351
Configure GTP Overlap Mapping 1351
Configuration Considerations 1351
Overlap Map Statistics 1361
5G Stateful Session Recovery 1361
Configure 5G Stateful Session Recovery 1362
View Backup and Restore information 1362
Remove Backup files 1364
5G Whitelisting 1365
Configure Whitelist Maps 1365
Change Priority of Whitelist Maps 1366
Delete Whitelist Maps 1366
User Plane Node Traffic Monitoring 1368
Rules and Notes 1369
Configure Stand-Alone User Plane Node Traffic Monitoring 1369

Control Plane Metadata **1370**
Configuration of 3G/4G Control Plane Metadata using Ansible **1372**
Creating Inventory Directory 1372
Creating fmInfo.yml 1372
Creating ansible_inputs.json 1372
Creating 3G/4G control plane metadata inventory file 1373
Creating host_vars directory 1375
Creating host_vars files 1376
Prerequisite 1376
Site 1377
cpNode 1377
upNode 1377
5GPolicy 1377
LTEPolicy 1377
Deployment of 3G/4G Control Plane Metadata Solution 1377
Set up of additional variable for Single GigaVUE-FM instance 1377
Execute the Playbook 1377
Configuration of 5G Control Plane Metadata using Ansible **1378**
Creating Inventory Directory 1378
Creating fmInfo.yml 1379
Creating ansible_inputs.json 1379
Creating 5G control plane metadata inventory file 1379
Creating host_vars directory 1382
Creating host_vars files 1383
Prerequisite 1383
Site 1384
cpNode 1384
upNode 1384
5GPolicy 1384
LTEPolicy 1384
Deployment of 3G/4G Control Plane Metadata Solution 1384
Set up of additional variable for Single GigaVUE-FM instance 1384
Execute the Playbook 1384
CPN-UPN Communication for Support of RAN and Network Slice Attributes **1385**
CPN-UPN Communication Configuration **1387**
Configure CPN-UPN Communication using Ansible 1387
Configure CPN-UPN Communication using CLI 1388

Configure CPN-UPN Communication Solution using Ansible **1389**
Create Inventory Directory 1389
Create fmInfo.yml 1389
Create ansible_inputs.json File 1389
Create an Inventory File for CPN-UPN Communication 1390
Create host_vars directory 1391
Create host_vars files 1391
Deploy CPN-UPN Communication Solution in Ansible 1397
View CPN-UPN Communication Solution in GigaVUE-FM 1398
Remove CPN-UPN Communication 1399
Map Rules for CPN-UPN Communication **1399**
Flow Sample Map Rules for CPN-UPN Communication 1399
Rules and Notes: 1399
Forward List Map Rules for CPN-UPN Communication 1402
Rules and Notes: 1402
View CUPS Communication Dashboard **1404**
Upgrade Standalone UPN to CPN-UPN Communication Solution **1407**
Rollback from CPN-UPN Communication Solution to Standalone UPN **1408**
Quick Rollback from CPN-UPN Communication Solution to Standalone UPN **1409**
Monitoring of Subscriber Intelligence Solutions **1409**
Monitoring CUPS Solution 1409
Health Status of Solution 1410
Monitor Session and Tunnel Utilization 1412
Configure SNMP Traps for Subscriber Intelligence Solution **1412**
Display Flow Ops Reports **1414**
Flow Ops Report - Field Reference 1416
Export Flow SIP Session Reports 1419
Export Flow Filtering Reports 1420
Generate Delta Reports 1420
GTP Overlap Flow Sampling Maps **1423**
Configure GTP Overlap Mapping 1425
Configuration Considerations 1425

GTP Overlap Flow Sampling Maps Example 1431
This section contains: 1431
Example 1: GTP Overlap Mode 1431
GTP Stateful Session Recovery 1436
Configure GTP Persistence 1437
GTP Scaling 1440
GigaSMART Cards in GigaVUE-OS Devices 1440
GTP Engine Grouping 1441
Passing GTP Control Traffic 1442
Configure GTP Engine Grouping 1443
Display Statistics 1443
GTP Engine Grouping Configuration Examples 1443
GigaSMART TLS/SSL Decryption for Inline Tools 1453
GigaSMART TLS/SSL Decryption for Out-of-Band Tools 1454
GigaSMART Passive TLS/SSL Decryption 1454
About Passive TLS/SSL Decryption 1455
Supported Protocols, Algorithms, and Ciphers 1458
Limitations 1460
Licensing 1460
Create and Reset TLS/SSL Keychain Passwords 1460
Work with Keys and Services 1462
Set Up Key Store Certificate Management for Passive TLS/SSL 1462
Upload TLS/SSL Private Keys 1463
Delete TLS/SSL Key 1463
Display Key Store Details for Passive TLS/SSL 1464
Create TLS/SSL Service 1465
ECODES for Troubleshooting Passive TLS/SSL Decryption 1466
View Passive TLS/SSL Decryption Flow Ops Report 1467
View Certificate Statistics 1468
View Service Statistics 1468
View Error Statistics 1468
Configuring Passive TLS/SSL Decryption Examples - Command Line Reference 1469
Example 1: TLS/SSL Decryption with a Regular Map 1469
Example 2: TLS/SSL Decryption with De-duplication 1470
Other Usage Examples 1471
Thales-Luna Network HSM for TLS/SSL Decryption for Out-of-Band Tools(Passive) 1471
Thales-Luna HSM for TLS/SSL Decryption for Out-of-Band Tools—Rules and Notes 1472

Configure HSM for TLS/SSL Decryption for Out-of-Band Tools 1473
Configure HSM Group 1473
Configure Set Key Handler 1474
Configure Passive TLS/SSL Network Access 1476
Use RFS to Manage Encrypted Keys 1477
Configure a GigaSMART Group 1478
Create a GigaSMART Operation (GSOP) 1479
Create TLS/SSL Keychain Password 1480
Upload TLS/SSL Private Keys 1481
Configure TLS/SSL Service 1482
Configure Maps 1483
Traffic Intelligence 1486
GigaSMART Adaptive Packet Filtering (APF) 1487
Feature Overview 1487
Implement APF Through the UI 1488
Content-based Filtering 1489
Encapsulation Awareness 1493
Pattern Matching on Gen 2 GigaSMART modules 1494
Cross-Packet Pattern Matching 1500
Pattern Matching on Gen 3 GigaSMART modules 1502
Regex Masking on Gen 3 GigaSMART modules 1504
Configuring a Regex Masking on Gen 3 GigaSMART modules 1504
Feature Parity with Gen 2 GigaSMART module and Gen 3 GigaSMART module 1507
Map Statistics 1507
Display APF Statistics 1507
Adaptive Packet Filtering Examples 1507
GigaSMART Load Balancing 1549
Stateful Load Balancing 1550
Examples 1554
Stateless Load Balancing 1559
Examples 1562
Enhanced Load Balancing 1568
GigaSMART De-Duplication 1570
Feature Overview 1571
De-duplication Configuration Steps 1572
Configure GigaSMART Parameters for Packet De-duplication 1572
Example – GigaSMART De-duplication 1573
Display De-duplication Statistics 1575
GigaSMART Traffic Performance Enhancement 1575
Flow Masking 1575

GigaSMART Encapsulated Traffic Performance Enhancement 1576
GigaSMART Header Stripping 1579
Header Stripping Illustrated 1583
Example – Stripping MPLS Headers and Adding a VLAN ID 1583
Example – Stripping Layer 3 GRE IP Encapsulated Packets 1584
Example – Stripping Layer 2 GRE Ethernet Encapsulated Packets 1585
Display Header Stripping Statistics 1586
Example – FM6000 Timestamping 1586
Generic Header Stripping 1588
Custom 1590
PPPoE 1592
Cisco LISP 1593
L2 MPLS 1595
VXLAN 1596
TRILL 1598
Avaya SPB 1599
GigaSMART Header Addition 1601
GigaSMART Masking 1602
Examples – GigaSMART Masking 1604
GigaSMART NetFlow Generation 1606
NetFlow Generation Components 1609
Network Ports 1610
Map(s) 1610
GigaSMART Group 1610
GSOP 1611
NetFlow Records 1611
NetFlow Monitors 1611
Sampled NetFlow Data 1612
NetFlow Exporters 1612
IP Interface with Tool Ports 1613
Enhancements to NetFlow 1613
IPFIX Extension: HTTP Response Code 1617
Display Exporter Statistics 1636
Display Monitor Statistics 1637
Display IP Interfaces Statistics 1637
NetFlow Generation Configuration Modification and Removal 1637
Modify NetFlow Generation Monitor Configuration 1637
Modify NetFlow Generation Record Configuration 1638
Remove NetFlow Generation Configuration 1639
V5 Fixed Record Template 1640
NetFlow Generation Match/Key and Collect/Non-Key Elements 1641

Match/Key Syntax 1642
Collect/Non-Key Syntax 1645
NetFlow Cacheless export using TCP protocol 1653
Rules and Notes 1653
Configure Netflow Generation 1654
GigaSMART Packet Slicing 1662
Example – GigaSMART Packet Slicing 1663
Display Slicing Statistics 1664
GigaSMART Advanced Flow Slicing 1664
Limitations 1664
Create Advanced Flow Slicing Profile 1665
Example-GigaSMART Advanced Slicing 1666
Display Slicing Statistics 1666
GigaSMART SSL Decryption 1667
PCAPng Application 1667
Tunneling Operations 1668
GigaSMART IP Encapsulation/Decapsulation (GigaSMART Tunnel) 1669
Configure Sending End of Tunnel: GigaVUE V Series map rules in Sydney 1671
Configure Receiving End of Tunnel: GigaVUE HC Series with GigaSMART in Melbourne 1672
Configure IP Interfaces for GigaSMART Tunnels 1673
About IP Interface Centralization 1674
Configure GigaSMART IP Encapsulation/Decapsulation 1675
Configure Sending End of Tunnel: GigaVUE-HC1 in Reno 1678
Configure Receiving End of Tunnel: GigaVUE-HC3 with GigaSMART in San Francisco 1680
GigaSMART IP Encapsulation (GigaSMART Tunnel) 1681
GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation 1682
Layer 2 GRE Header Length 1683
Load Balancing to Multiple Destinations 1684
L2GRE IPv4 Encapsulation/Decapsulation 1684
L2GRE IPv6 Encapsulation/Decapsulation 1685
L2GRE Tunnel Termination 1686
Configure L2GRE Tunnel Encapsulation and Decapsulation 1687
Display L2GRE IPV4/IPv6 Tunnel Statistics 1687
Display L2GRE Tunnel Encapsulation Statistics 1687
Display L2GRE Tunnel Decapsulation Statistics 1688

Configure GigaSMART Operation for Layer 2 GRE 1688
Example 1 – GigaSMART L2GRE Tunnel Encapsulation 1689
Example 2 – GigaSMART L2GRE Tunnel Encap Stateful LB 1690
Example 3 – GigaSMART L2GRE Tunnel Encap Stateless LB 1693
Example 4 – GigaSMART L2GRE Tunnel Decapsulation 1694
Example 5 – GigaSMART L2GRE IPv6 Tunnel Encamp/Decap with Load-Balancing 1696
Orchestrated Workflow Configuration of GigaSMART L2GRE Tunnels 1702
GigaSMART VXLAN Tunnel Encapsulation 1703
VXLAN Header Length 1704
VXLAN IPv4 Encapsulation/Decapsulation 1704
VXLAN IPv6 Encapsulation/Decapsulation 1706
Display VXLAN IPV4/IPv6 Tunnel Statistics 1706
Display VXLAN Tunnel Encapsulation Statistics 1707
Orchestrated Workflow Configuration of GigaSMART VXLAN Tunnels 1707
IP Fragmentation and Reassembly on L2GRE and GMIP Tunnels 1709
IP Fragmentation on Encapsulation 1709
IP Reassembly on Decapsulation 1710
Notes and Considerations 1710
GigaSMART ERSPAN Tunnel Decapsulation 1711
ERSPAN Type III 1712
Configure GigaSMART Operations for ERSPAN 1714
ERSPAN Tunnel Header Removal 1714
ERSPAN Type III Tunnel Header Removal 1716
Display ERSPAN Statistics 1718
Orchestrated Workflow Configuration of GigaSMART ERSPAN Tunnels 1719
GigaSMART VXLAN Tunnel Decapsulation 1720
GigaSMART Custom Tunnel Decapsulation 1725
Custom Tunnel Decapsulation Configuration Example 1726
GigaSMART TCP tunnel 1729
Configuration 1729
Supported Devices 1729
Secure Tunnels 1729
Supported Platforms 1731
Configure Secure Tunnels 1731

Edit Secure Tunnel 1734
Delete Secure Tunnel 1734
Export Secure Tunnels 1734
View Tunnel Statistics 1735
Limitations 1735
Work with GigaSMART Operations 1735
GigaSMART Licensing 1736
Access GigaSMART from GigaVUE-FM 1737
Quick Glance- How to Configure a GigaSMART Application 1737
Create GigaSMART Operations – A Summary 1739
Groups of GigaSMART Engine Ports 1740
Reload GigaSMART engine ports in a group 1742
How to Use GigaSMART Operations – Example 1743
Engine Watchdog Timer in GigaSMART 1745
Tunnel Health Checks 1746
Configure Hashing 1748
GigaSMART Rules and Tips 1749
Virtual Ports 1750
Virtual Port Rules 1750
Create Virtual Port 1751
Multiple Virtual Ports for First Level Map 1754
Multiple Virtual Port Rules 1756
Multiple Virtual Port with Other GigaSMART Applications 1757
Virtual Port Statistics 1760
Differences in GigaSMART Nomenclature Between the CLI and GigaVUE-FM 1761
GigaSMART Operations in Clusters 1762
How to Combine GigaSMART Operations 1764
How to Read GigaSMART Operations Table 1765
Work with GigaSMART Operation Combinations in GigaVUE-FM 1765
Supported GigaSMART Operations 1766
Supported GigaSMART Operations on GigaVUE HC Series 1766

Supported GigaSMART operations on GigaVUE V Series Node 1770
Order of GigaSMART Operations 1773
View GigaSMART Statistics 1773
Definitions of GigaSMART Statistics 1774
NetFlow Monitor Statistics Definitions 1774
NetFlow Exporter Statistics Definitions 1775
IP Interfaces Statistics Definitions 1776
TLS/SSL Application Statistics Definitions 1778
ASF Statistics Definitions 1779
GigaSMART Group Statistics Definitions 1780
GigaSMART Group Flow Ops Report Statistics Definitions 1781
Flow Ops Report Statistics Definitions for FlowVUE 1781
Flow Ops Report Statistics Definitions for GTP 1782
Flow Ops Report Statistics Definitions for GTP Overlap 1786
Flow Ops Report Definitions for SIP/RTP Correlation 1787
Flow Ops Report Statistics for Passive TLS/SSL Decryption 1789
GigaSMART Operations Statistics Definitions 1790
Passive TLS/SSL Decryption Statistics Definitions 1791
Inline TLS/SSL Decryption Statistics Definitions 1791
De-duplication Statistics Definitions 1791
ERSPAN Statistics Definitions 1792
Tunnel Decapsulation Statistics Definitions 1793
Tunnel Encapsulation Statistics Definitions 1795
APF Statistics Definitions 1797
ASF Statistics Definitions 1797
Masking Statistics Definitions 1798
Slicing Statistics Definitions 1798
Header Stripping Statistics Definitions 1799
Generic Header Stripping Statistics Definitions 1799
Trailer Statistics Definitions 1800
FlowVUE Statistics Definitions 1800
NetFlow Statistics Definitions 1801
Display GigaSMART Application Resource Usage 1801
Overview of GigaSMART Application Resources 1802
Resources for Buffer ASF 1803
Reload GigaSMART Line Card or Module 1803
GigaSMART CPU Utilization Statistics 1804
Display GigaSMART CPU Utilization 1805
Configure Threshold 1806
Configure Threshold Crossing Notification 1806

How to Use GigaSMART Trailers **1806**
Source ID Field 1807
Remove GigaSMART Trailers 1808
Multiple GigaSMART Trailers and Cascade Connections 1810
Interpret GigaSMART Trailer 1811
GigaSMART Trailer Reference 1811
GigaSMART Trailer Format 1811
GigaSMART Trailer Format 1811
GigaSMART Trailer TLVs 1812
Format of the GIGAMON_SRCID TLV 1813
GigaSMART Logs **1820**
Create GS Log file **1820**
Delete Log File **1821**
Upgrade GigaSMART Cards **1821**
Upgrade GigaSMART card using External Image Server 1821
Upload image to external image server 1822
Add the external server to GigaVUE-FM 1822
Upgrade GigaSMART card using External Image 1822
Upgrade GigaSMART card using Internal Image Server 1823
Download Image 1823
Upload the image file to GigaVUE-FM 1823
Upgrade GigaSMART card using Internal image 1824
Additional Sources of Information **1825**
Documentation **1825**
How to Download Software and Release Notes from My Gigamon 1828
Documentation Feedback **1828**
Contact Technical Support **1829**
Contact Sales **1830**
Premium Support **1830**
The VÜE Community **1830**
Glossary **1831**

Fabric Management

GigaVUE-FM—also called GigaVUE-FM fabric manager—provides overall management for small or large Gigamon Deep Observability Pipeline deployments. This section describes how to set up and use GigaVUE-FM, and reviews a broad range of management and configuration options.

Learn about:

- how GigaVUE-FM presents a centralized fabric management platform (GigaVUE-FM, GigaVUE Nodes and Clusters)
- how the Core Intelligence provided by GigaVUE-OS and managed through GigaVUE-FM can be used to optimize your network traffic flow (Traffic Filtering)
- how to configure GigaSMART® operations to provide network visibility and increase the efficiency of network monitoring, security, and performance tools. (GigaSMART®)
- how to identify and monitor application usage and filter traffic accordingly (Application Intelligence)
- how the Gigamon Visibility and Analytics Fabric is made possible through GigaVUE-FM (Fabric Maps)
- how to view the reporting options that GigaVUE-FM offers for the whole fabric (Dashboard)
- how to deploy GigaVUE Cloud Suite solutions (Virtual and Cloud Environments)

NOTE: Refer to the *GigaVUE-FM Installation and Upgrade Guide* to get GigaVUE-FM installed and running in your network environment. Then, turn to later chapters for information on using product features.

GigaVUE-FM

GigaVUE-FM fabric manager is a web-based fabric management interface that provides high-level visibility and management of both the physical and virtual traffic visibility nodes that form the Gigamon Traffic Visibility Fabric™. GigaVUE-FM can manage the following types of traffic visibility nodes:

- **Physical GigaVUE Nodes** – GigaVUE-FM manages GigaVUE TA Series, and HC Series nodes, allowing for a unified workspace, while also providing an easy-to-use wizard-based approach for configuring Flow Mapping®® and GigaSMART® traffic policies. For a list of GigaVUE TA Series, and HC Series nodes supported for management in this release, refer to Supported Nodes .

- **Virtual GigaVUE Nodes** – GigaVUE-FM also extends the GigaVUE feature set into the virtual space by allowing for the discovery, configuration, and management of the new GigaVUE-VM virtual traffic visibility node. GigaVUE-VM provides powerful Flow Mapping® technology for the traffic flowing between virtual machines, allowing you to distribute cloud-based traffic to physical tool ports in the visibility fabric.

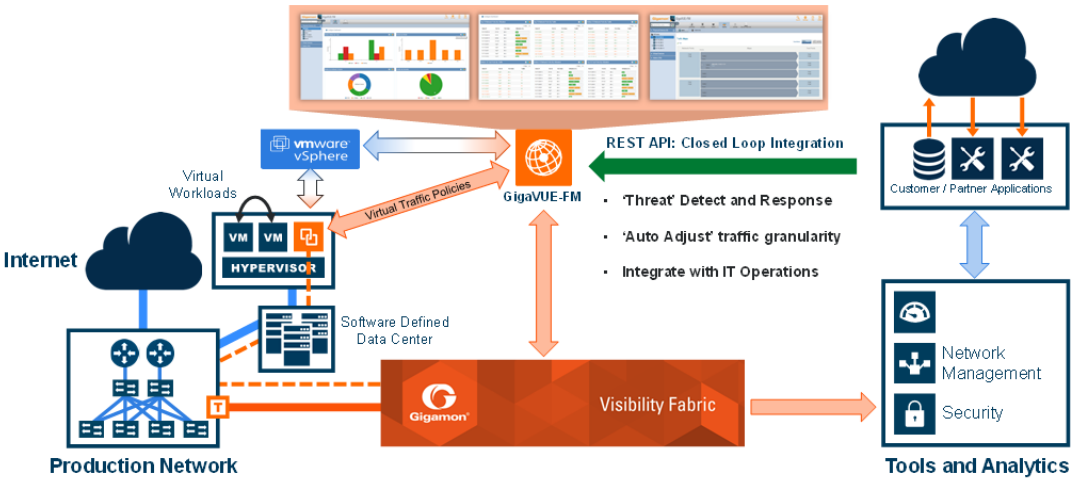


Figure 1 *GigaVUE-FM Overview*

Topics

- Supported Nodes
- GigaVUE-FM Features and Benefits
- Get Started with GigaVUE-FM
- Dashboard

Supported Nodes

GigaVUE-FM provides supports the following GigaVUE Nodes

- GigaVUE TA Series
- HC Series

NOTE: After upgrading to software version 5.15, GigaVUE-FM will no longer support to manage GigaVUE G Series Nodes and all the existing GigaVUE G Series Nodes in GigaVUE-FM will be removed.

GigaVUE-FM provides support for GigaVUE TA Series, and HC Series nodes through the administration of physical nodes (Device Management and Configuration Management) feature.

Topics:

- Device and Configuration Management

- How to enable Web Server for Node Management

Device and Configuration Management

During the initial start up of GigaVUE-FM, it performs a discovery on GigaVUE® HC Series and TA Series nodes listed in the table below. It supports the versions listed in the following table. GigaVUE-FM allows you to perform configuration tasks on these devices.

Table 1: GigaVUE-FM Managed GigaVUE Nodes Software Version ***GigaVUE*** **HC Series Nodes**

GigaVUE Series	Node	Compatible GigaVUE-OS versions
GigaVUE HC Series Nodes	GigaVUE-HC1	v5.4.xx, v5.5.xx, v5.6.xx, v5.7.xx, v5.8.xx, v5.9.xx, v5.10.xx, v5.11.xx, v5.12.xx, v5.13.xx, v5.14.xx, v5.15.xx, v5.16.xx, v6.0.xx, v6.1.xx, v6.2.xx, v6.3.xx, v6.4.xx, v6.5.xx, v6.6.xx, v6.7xx, v6.8xx, v6.9xx, v6.10xx, v6.11.xx, v6.12.xx, v6.13.xx, v6.14.xx
	GigaVUE-HC1-Plus	v6.0.xx, v6.1.xx, v6.2.xx ,v6.3.xx, v6.4.xx, v6.5.xx, v6.6.xx, v6.7xx, v6.8xx, v6.9xx, v6.10xx, v6.11.xx, v6.12.xx, v6.13.xx, v6.14.xx
	GigaVUE-HC3 CCv2	v5.4.xx, v5.5.xx, v5.6.xx, v5.7.xx, v5.8.xx, v5.9.xx, v5.10.xx, v5.11.xx, v5.12.xx, v5.13.xx, v5.14.xx, v5.15.xx, v5.16.xx, v6.0.xx, v6.1.xx, v6.2.xx ,v6.3.xx, v6.4.xx, v6.5.xx, v6.6.xx, v6.7xx, v6.8xx, v6.9xx, v6.10xx, v6.11.xx, v6.12.xx, v6.13.xx, v6.14.xx
	GigaVUE-HCT	v6.5.xx, v6.6.xx, v6.7xx, v6.8xx, v6.9xx, v6.10xx, v6.11.xx, v6.12.xx, v6.13.xx, v6.14.xx
	GigaVUE-HC2 CCv1 (End of Sale)	v5.4.xx, v5.5.xx, v5.6.xx, v5.7.xx, v5.8.xx, v5.9.xx, v5.10.xx, v5.11.xx, v5.12.xx, v5.13.xx, v5.14.xx, v5.15.xx, v5.16.xx, v6.0.xx, v6.1.xx, v6.2.xx, v6.3.xx, v6.4.xx, v6.5.xx
	GigaVUE-HC2 CCv2 (End of Sale)	v5.4.xx, v5.5.xx, v5.6.xx, v5.7.xx, v5.8.xx, v5.9.xx, v5.10.xx, v5.11.xx, v5.12.xx, v5.13.xx, v5.14.xx, v5.15.xx, v5.16.xx, v6.0.xx, v6.1.xx, v6.2.xx, v6.3.xx, v6.4.xx, v6.5.xx
	GigaVUE-HC3 CCv1 (End of Support)	v5.4.xx, v5.5.xx, v5.6.xx, v5.7.xx, v5.8.xx, v5.9.xx, v5.10.xx, v5.11.xx, v5.12.xx, v5.13.xx, v5.14.xx, v5.15.xx, v5.16.xx

Table 2: GigaVUE-FM Managed GigaVUE Nodes Software Version ***GigaVUE*** **TA Series Nodes**

GigaVUE Series	Node	Compatible GigaVUE-OS versions
GigaVUE TA Series Nodes	GigaVUE-TA25	v5.11.xx, v5.12.xx, v5.13.xx, v5.14.xx, v5.15.xx, v5.16.xx,v6.0.xx, v6.1.xx, v6.2.xx ,v6.3.xx, v6.4.xx, v6.5.xx, v6.6.xx, v6.7xx, v6.8xx, v6.9xx, v6.10xx, v6.11.xx, v6.12.xx, v6.13.xx, v6.14.xx
	GigaVUE-TA25E	v6.1.xx, v6.2.xx ,v6.3.xx, v6.4.xx, v6.5.xx, v6.6.xx, v6.7xx, v6.8xx, v6.9xx, v6.10xx, v6.11.xx, v6.12.xx, v6.13.xx, v6.14.xx
	GigaVUE-TA200	v5.4.xx, v5.5.xx, v5.6.xx, v5.7.xx, v5.8.xx, v5.9.xx, v5.10.xx, v5.11.xx, v5.12.xx, v5.13.xx, v5.14.xx, v5.15.xx, v5.16.xx, v6.0.xx, v6.1.xx, v6.2.xx, v6.3.xx, v6.4.xx, v6.5.xx, v6.6.xx, v6.7xx, v6.8xx, v6.9xx, v6.10xx, v6.11.xx, v6.12.xx, v6.13.xx, v6.14.xx
	GigaVUE-TA200E	v6.1.xx, v6.2.xx, v6.3.xx, v6.4.xx, v6.5.xx, v 6.6.xx, v6.7.xx, v6.8xx, v6.9xx, v6.10xx, v6.11.xx, v6.12.xx, v6.13.xx, v6.14.xx
	GigaVUE-TA400	v5.14xx, v5.15xx, v5.16.xx, v6.0.xx, v6.1.xx, v6.2.xx, v6.3.xx, v6.4.xx, v6.5.xx, v6.6.xx, v6.7.xx, v6.8xx, v6.9xx, v6.10xx, v6.11.xx, v6.12.xx, v6.13.xx, v6.14.xx
	GigaVUE-TA400E	v6.11.xx, v6.12.xx, v6.13.xx, v6.14.xx
	GigaVUE-TA10 (End of Sale)	v5.4.xx, v5.5.xx, v5.6.xx, v5.7.xx, v5.8.xx, v5.9.xx, v5.10.xx, v5.11.xx, v5.12.xx, v5.13.xx, v5.14.xx, v5.15.xx, v5.16.xx, v6.0.xx, v6.1.xx
	GigaVUE-TA40 (End of Sale)	v5.4.xx, v5.5.xx, v5.6.xx, v5.7.xx, v5.8.xx, v5.9.xx, v5.10.xx, v5.11.xx, v5.12.xx, v5.13.xx, v5.14.xx, v5.15.xx, v5.16.xx, v6.0.xx, v6.1.xx
	GigaVUE-TA100 (End of Sale)	v5.4.xx, v5.5.xx, v5.6.xx, v5.7.xx, v5.8.xx, v5.9.xx, v5.10.xx, v5.11.xx, v5.12.xx, v5.13.xx, v5.14.xx, v5.15.xx, v5.16.xx, v6.0.xx, v6.1.xx, v6.2.xx, v6.3.xx, v6.4.xx, v6.5.xx
	GigaVUE-TA100-CXP (End of Sale)	v5.4.xx, v5.5.xx, v5.6.xx, v5.7.xx, v5.8.xx, v5.9.xx, v5.10.xx, v5.11.xx, v5.12.xx, v5.13.xx, v5.14.xx, v5.15.xx, v5.16.xx, v6.0.xx, v6.1.xx
	GigaVUE-TA1 (End of Support)	v5.4.xx, v5.5.xx, v5.6.xx, v5.7.xx

For more information on GigaVUE-OS that has reached EOS, refer to Summary of Software release Status for GigaVUE-OS GigaVUE-FM and Cloud-Suite.

NOTE: Although GigaVUE-FM may recognize earlier versions of GigaVUE® TA and HC Series nodes, the versions listed in the above table are the officially supported versions. Earlier versions are not managed by GigaVUE-FM.

How to enable Web Server for Node Management

GigaVUE-FM can only discover and manage nodes with their web servers enabled and operating on the default HTTP port of 80. The request from Port 80 is immediately redirected to port 443 (HTTPS) for secure connections over SSL.

Supported Browsers for GigaVUE-FM

GigaVUE-FM supports the following browsers:

OS	Browser	Minimum Browser Version
Windows®	Mozilla Firefox™	143.0.1
	Microsoft Edge	140.0.3485.94
	Google® Chrome®	140.0.7339.208
Mac OSX	Google® Chrome®	140.0.7339.214
	Mozilla Firefox™	143.0.1
	Apple® Safari®	26.0.1
	Microsoft Edge	140.0.3485.94
Linux	Google® Chrome®	137.0.7151.27
	Mozilla Firefox™	138.0.4

GigaVUE-FM Features and Benefits

The GigaVUE-FM is a web-based management interface that provides unified access, centralized administration, and high-level visibility for all GigaVUE traffic visibility nodes in the enterprise or data center, allowing a global perspective which is not possible from individual nodes.

The following table summarizes the major benefits of GigaVUE-FM:

Table 3: Features and Benefits of GigaVUE-FM

Benefit	Descriptions
Centralized Management and Control	Provides centralized management, monitoring, and configuration of the physical and virtual traffic policies for the Visibility Fabric, allowing administrators to map and direct network traffic to the tools and analytics infrastructure.
Programmable APIs for Software Defined Visibility	REST APIs that can be used by the traffic monitoring or IT operations management tools to perform various tasks, such as • Improve security through better network detection, reaction, and response by automating NetFlow generation and SSL decryption so that current security appliances are not overtaxed when performing deep packet inspection. • Program the Visibility Fabric flow maps when security threats are detected. • Discover the Visibility Fabric nodes for inventory and status collection. • Performing common tasks, such as provisioning and ticketing of network port configurations. • Programmatically create, update, or delete port properties, including port-type, admin state, speed, and others. • Programmatically create, update, or delete traffic maps and GigaSMART operations. **Note**: GigaVUE-FM can handle only 150 concurrent (write - POST/PUT/PATCH/DELETE) API requests at any given time.
Fabric-wide reporting	Summarized and customizable dashboards for inventory, node or cluster status, events, audit trail, and Top-N/Bottom-N port/map usage with options to export and schedule HTML or PDF reports for off-line viewing.
Advanced Monitoring	Proactively monitor and troubleshoot hot spots in your Visibility Fabric: • Top-N, Bottom-N Network/Tool Port and Map usage widgets in the dashboard • Global search across the Visibility Fabric for quick access to monitoring hot spots • Audit trail of user operations for enterprise security compliance • Historical trend analysis (1 hour, 1 day, 1 week, 1 month) for port and traffic policies • Quick Views for easy access to Visibility Fabric details (node, port, traffic policies)
Scheduling capabilities	Initiates version updates to one or many fabric nodes to streamline software rollouts in an automated fashion.
Backup and Restore Capabilities	Supports configuration backup and restore across multiple visibility nodes to quickly back-out changes if required due to errors or change control requirements. **Note**: Restoration of backed up configuration fails in G-Series. Alternatively you can download the file from GigaVUE-FM and use the GigVUE-OS-CLI to restore the configuration.

Benefit	Descriptions
Improved Operational Efficiencies	Minimizes resources required to configure, manage, and monitor multiple visibility nodes: • Create/Update/Delete port properties including port-type, admin state, speed, and others • Create/Update/Delete traffic maps and GigaSMART operations
Near-Real Time Config Status	Provides Near-Real Time (NRT) status of the following: • **Configuration changes performed using GigaVUE-FM APIs**: Immediately reflected in the GigaVUE-FM GUI. • **State changes occurring in the device (for example, oper up/down changes happening in the device)**: Immediately reflected in the GigaVUE-FM GUI. • **Configuration changes made in the devices (for example, attributes such as 'alias' changed by the user in the device through CLI or API):** Get reflected in GigaVUE-FM only after the next config sync cycle, which could be few minutes. **Note**: ▪ You must enable SNMP traps in GigaVUE-FM for the NRT status updates to get reflected. However, SNMP traps are enabled by default in GigaVUE-FM. Refer to the "SNMP Traps" section in the *GigaVUE Administration Guide* for more details. ▪ To avoid flooding of alarms, devices implement SNMP throttling. Refer to the "SNMP Throttling" section in the *GigaVUE Administration Guide*.
Scalability	Provides a reliable and stable environment for managing a large number of physical devices without any impact to the performance. • **Vertical Scale**: Allows GigaVUE-FM to manage approximately 1000 Nodes. • **Horizontal Scale**: Allows GigaVUE-FM to manage approximately 3000 Nodes. Horizontal scale support is applicable only in GigaVUE-FM High Availability mode. Horizontal scale support involves service distribution to the standby nodes.

Get Started with GigaVUE-FM

This section provides an overview of the GigaVUE-FM interface. It also provides information about table customization and search features available in GigaVUE-FM.

It includes the following major sections:

- Log In to GigaVUE-FM
- Log Out of GigaVUE-FM
- GigaVUE-FM
- Configure a Custom Banner
- Quick Views
- Return to the Dashboard
- Table View Customization

- Notifications Panel
- How to Add the GigaVUE-FM Instance Name
- How to Search in GigaVUE-FM

Log In to GigaVUE-FM

The GigaVUE-FM login page provides information about the security policy login banner beside the username and password fields. The login banner is customizable. For more information about configuring a custom banner, refer to Configure a Custom Banner.

GigaVUE-FM is preconfigured with one user with the fm_super_admin role assigned (username - admin, password - admin123A!!). The default password (admin123A!! on the admin account must be changed to a non-default password (as it is no longer allowed to have the default password).

- If GigaVUE-FM is deployed inside AWS or OpenStack then, use the **Instance ID** as the default password.
- If you try to access the GigaVUE-FM internal page URLs (for example Port page *<fmipaddress>/app/#/node/10.115.32.12/ports/ports*) without logging in to GigaVUE-FM, then after logging in, you will be redirected to the Dashboard page and not the specific page that you tried to access.
- When there is a change in the IP address of GigaVUE-FM, the cms server must be restarted. Only when the cms server is restarted, the new IP address will be updated in all the relevant places.
- Log in to GigaVUE-FM once the page appears, as the session is valid for about 15 minutes.

Log Out of GigaVUE-FM

To log out of GigaVUE-FM, click on the user-profile drop-down ( ) icon on the top right of GigaVUE-FM GUI and select **Logout**.

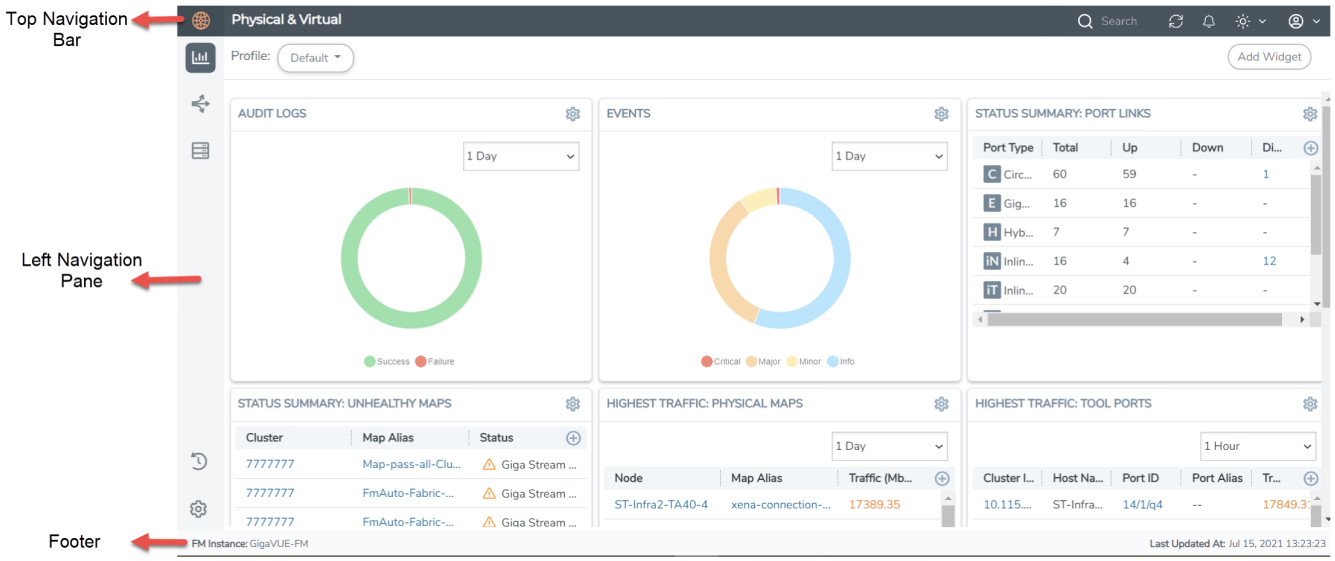
NOTE: You will be automatically logged out of GigaVUE-FM after a period of inactivity (based on the auto-logout time configured). To configure the auto-logout time, refer to the "Preferences" section in the GigaVUE Administration Guide.

GigaVUE-FM

Homepage

When you first log in to GigaVUE-FM, the Dashboard - Physical & Virtual page is displayed by default as shown in the following figure.

NOTE: You can add widget in GigaVUE-FM only if you have an active license.



Left Navigation Pane

The GigaVUE-FM landing page has a left navigation pane that expands into a floating pane which navigates to the following menus :

- **Dashboards** : Consists of the physical and virtual dashboards, the health monitor dashboards, and the Analytics dashboards. This page also includes the alarms, the events, and the audit logs pages.
- **Traffic** : Consists of the fabric solutions that the users must configure to monitor the flow of traffic.
- **Inventory** : Consists of the physical and virtual resources which the users must configure before configuring the traffic flow and solutions.
- **Recently Viewed** : Displays the list of recently viewed pages. Refer to Recently Viewed Section for detailed information.

- **Settings** : Consists of administrative, authentication, system resources that needs to be configured by the user.

Top Navigation Bar

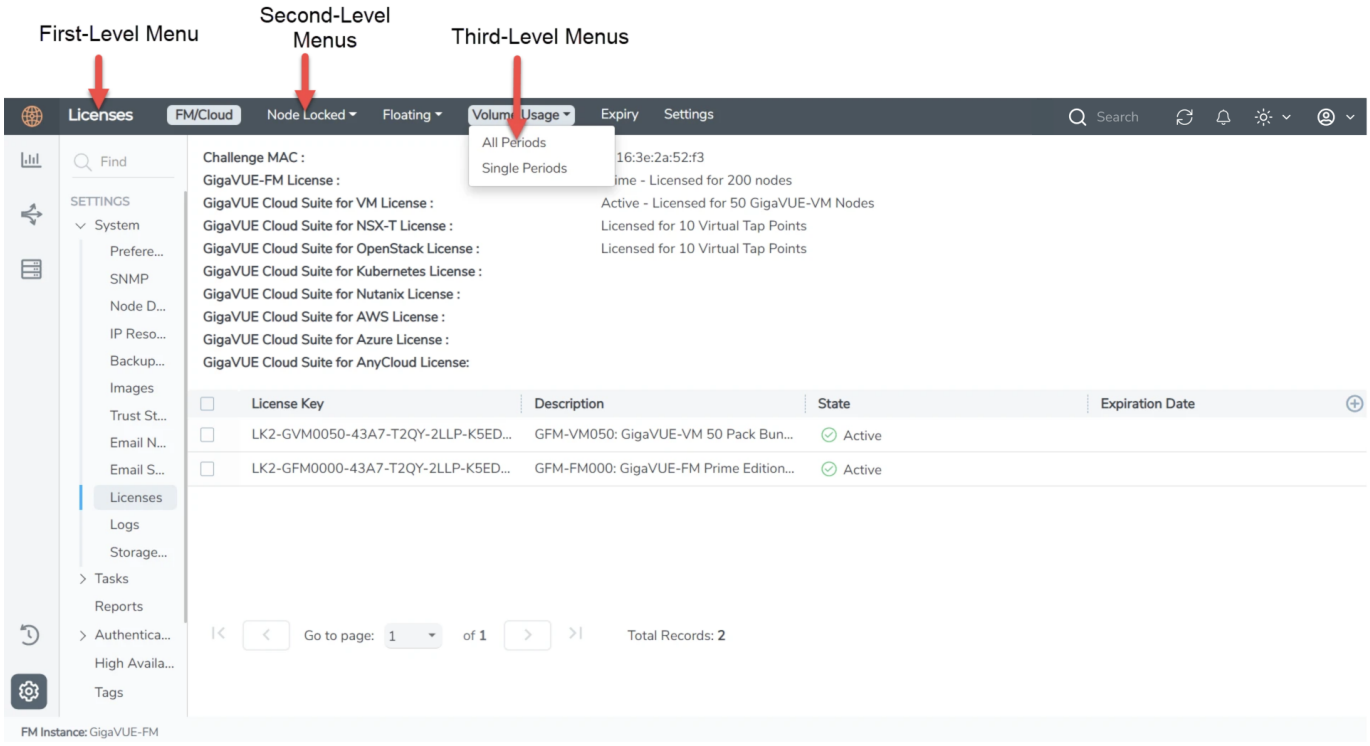
The top navigation bar contains:

- Page-level headers
- Search
- Refresh
- Rediscover - You can view the rediscover icon on the top navigation bar only when you click a node or a cluster ID (Go to **Inventory** > **Physical** > **Nodes**). Click the icon to pull the latest configuration from the GigaVUE-OS device. This works same as the Rediscover option in the Physical Nodes page. For more information refer to Cluster Discovery Behavior.
- Theme - For information on the Theme icon refer to Theme Settings.
- Profile options- Help option is available under . To create keyboard shortcuts to navigate to different menu pages, click on and select Keyboard Shortcuts.

NOTE: The Save Configuration option is available in the Node Overview, Flow Maps, and Active Visibility pages.

Page-Level Header

The page-level headers of the GigaVUE-FM instance are displayed as shown in the figure.



- The first-level menu is displayed as the Main Header in the top navigation bar.
- The second-level menus are displayed next to the first-level menu.

  NOTE: The second-level menus that overflow in the top navigation bar are displayed as a drop-down with an option to expand or collapse.

- The third-level menus are displayed as drop-down under the second-level menus.

GigaVUE-FM GUI Navigation

Use the navigation sidebar and the appropriate page-level headers to navigate to the various GUI pages. Depending on the user role and access rights of the user, the fields and buttons in the individual pages may either be enabled or disabled. Mandatory fields in the GUI are pages are notified appropriately.

Footer

The footer of GUI displays the GigaVUE-FM instance name, Node Synchronized time when accessing pages related to nodes, and NRT time stamp.

NOTE: On Windows Server 2019, some layout elements, such as the bottom bar, may not display correctly due to system-level display behavior. Additionally, certain screens could have minor issues with font rendering, spacing, or resolution scaling.

Theme Settings

The theme setting options allow you to choose among light, dark, and system themes. Once the theme is set, the settings remain the same, even when you log out of GigaVUE-FM, close or reopen the browser, or open a new tab for GigaVUE-FM. It is also independent of the browser settings.

The following table shows the options available when you click on the theme settings icon:

S.No	Options	Description
1	Light theme	UI changes to light color. It is the default theme.
2	Dark theme	UI changes to Dark color.
3	System	UI applies the theme based on the Operating System theme settings.

NOTE: When you cannot view the theme changes applied in Fabric Health Analytics (FHA), clear the browser cache and reload to view the updated changes.

Configure a Custom Banner

It is recommended to configure a pre-login banner which states the security policy of your company or organization. The banner appears on the login screen before the users log into GigaVUE-FM.

Only the users with fm_admin and fm_super_admin role assigned can view and configure the custom banner.

To configure the custom banner:

1. Go to **Settings** and select **Systems** > **Preferences**.

2. Click **Edit** and the Edit Preferences page appears.
3. Enter the custom banner message in the **Login infobox** text box.
4. Click **Save**.

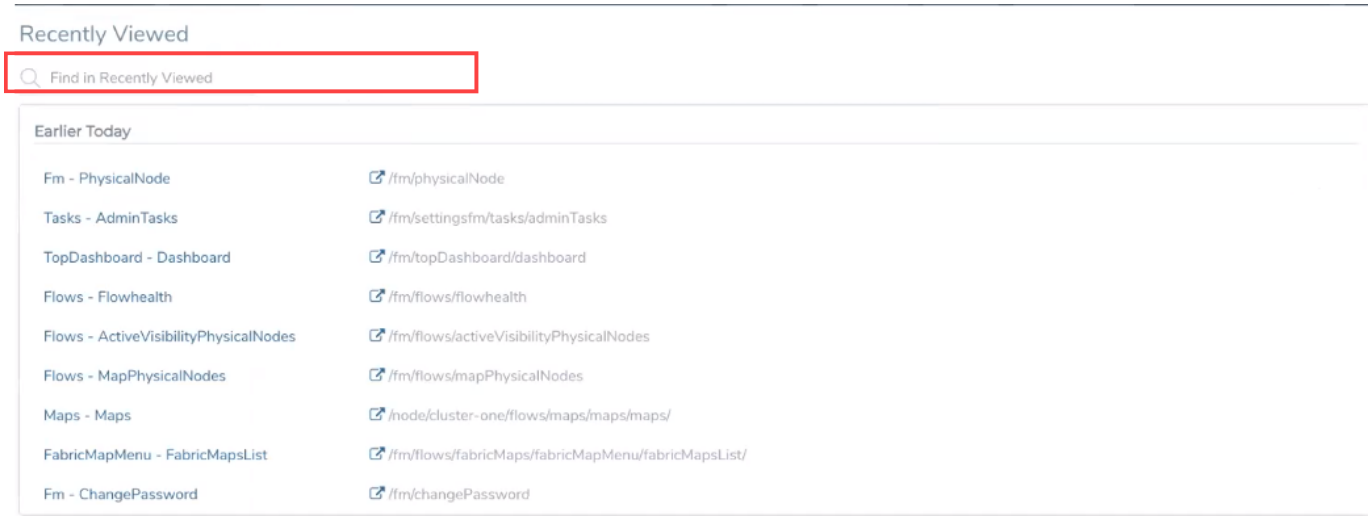
Quick Views

A quick view provides easy access to Visibility Fabric details such as nodes, ports, and traffic policies. In GigaVUE-FM, you can click on items such as port ID, map alias, port error counts, and so on, and get detailed information about the selected item.

Recently Viewed

Click on  icon on the left navigation pane to view the list of recently viewed pages.

- You can view up to last 100 pages that you have visited.
- The pages are listed in chronological order with the most recent entry listed at the top.
- Click on a link to navigate to that page.
- Use the 'Find in Recently Viewed' to filter the required pages.

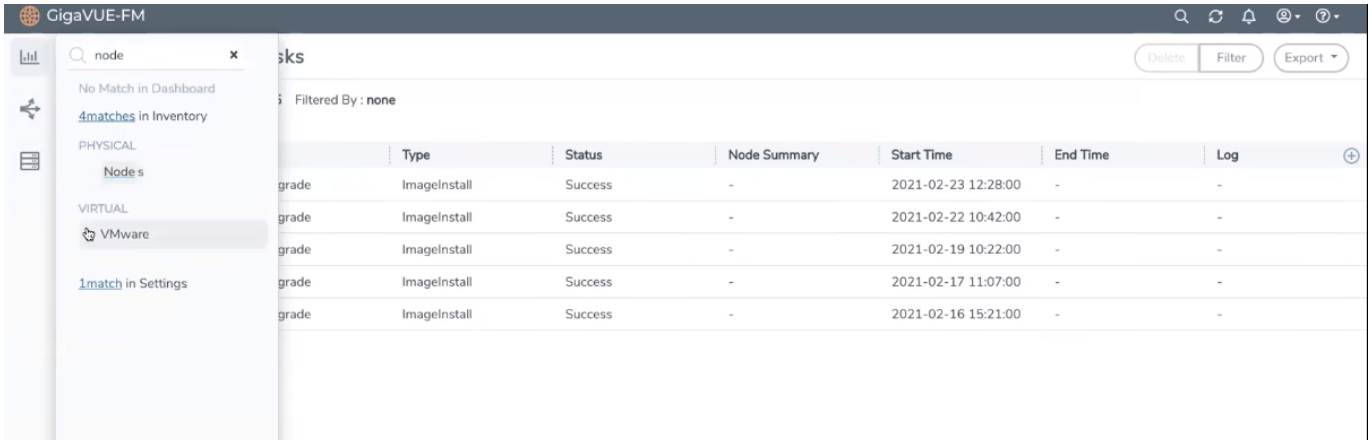


Find Pages

The Find option in the expanding tool bar allows you to search and quickly navigate to the pages without navigating through each of the menus and submenus.

- Enter the name of the page that you want to search (fully or in part). The number of matching entries under each of the menu and submenu category is displayed.
- Click the link under the category to navigate to the respective page.

For example, if you search for the term 'node', then the corresponding matching entries in the menus and submenus are displayed.



.

Return to the Dashboard

At any time, to return back to the Dashboard, click on the  icon on the top left of GigaVUE-FM GUI. Refer to Dashboard. By default, the Physical & Virtual Dashboard page is displayed.

Table View Customization

GigaVUE-FM enables you to customize the appearance of tables. You can choose the columns you want to show and hide in the table. You can also choose the order in which you want to view the columns in the table.

To customize the columns:

1. Click the '+' icon on the top-right edge of the table.

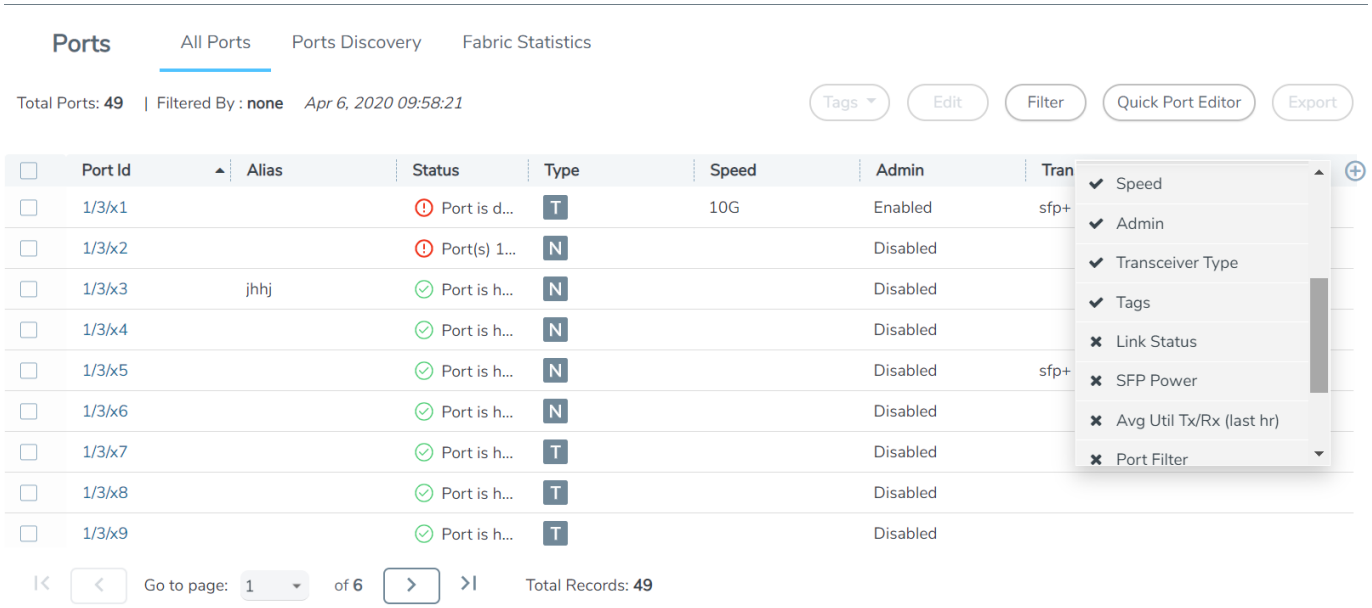


Figure 2 *Table menu to configure columns*

2. Click on a column name to change the show/hide setting. A check mark indicates the columns to show and an X indicates the columns to hide.
3. To rearrange the columns in the table, select a column heading and drag it to the new location. Your customizations are automatically saved.
4. Click on 'Reset columns to default' to reset the columns to the default view.

NOTE: The customized column settings are preserved for the user profile. When you logout and log back in, the tables display the same customized columns.

The pagination option on the bottom-right corner of the page allows you to scroll through long lists of data that span across multiple pages. You can also jump to a specific page by clicking the page number. Each page can show up to 100 rows of data per view.

Click the export option to export the tables either in CSV file format or in XLSX format. You can either export all the records or export only the selected records.

Notifications Panel

The Notifications icon in the top navigation bar displays a number if there are any system alerts. Immediate alerts appear on the left side of the page as individual pop-ups.

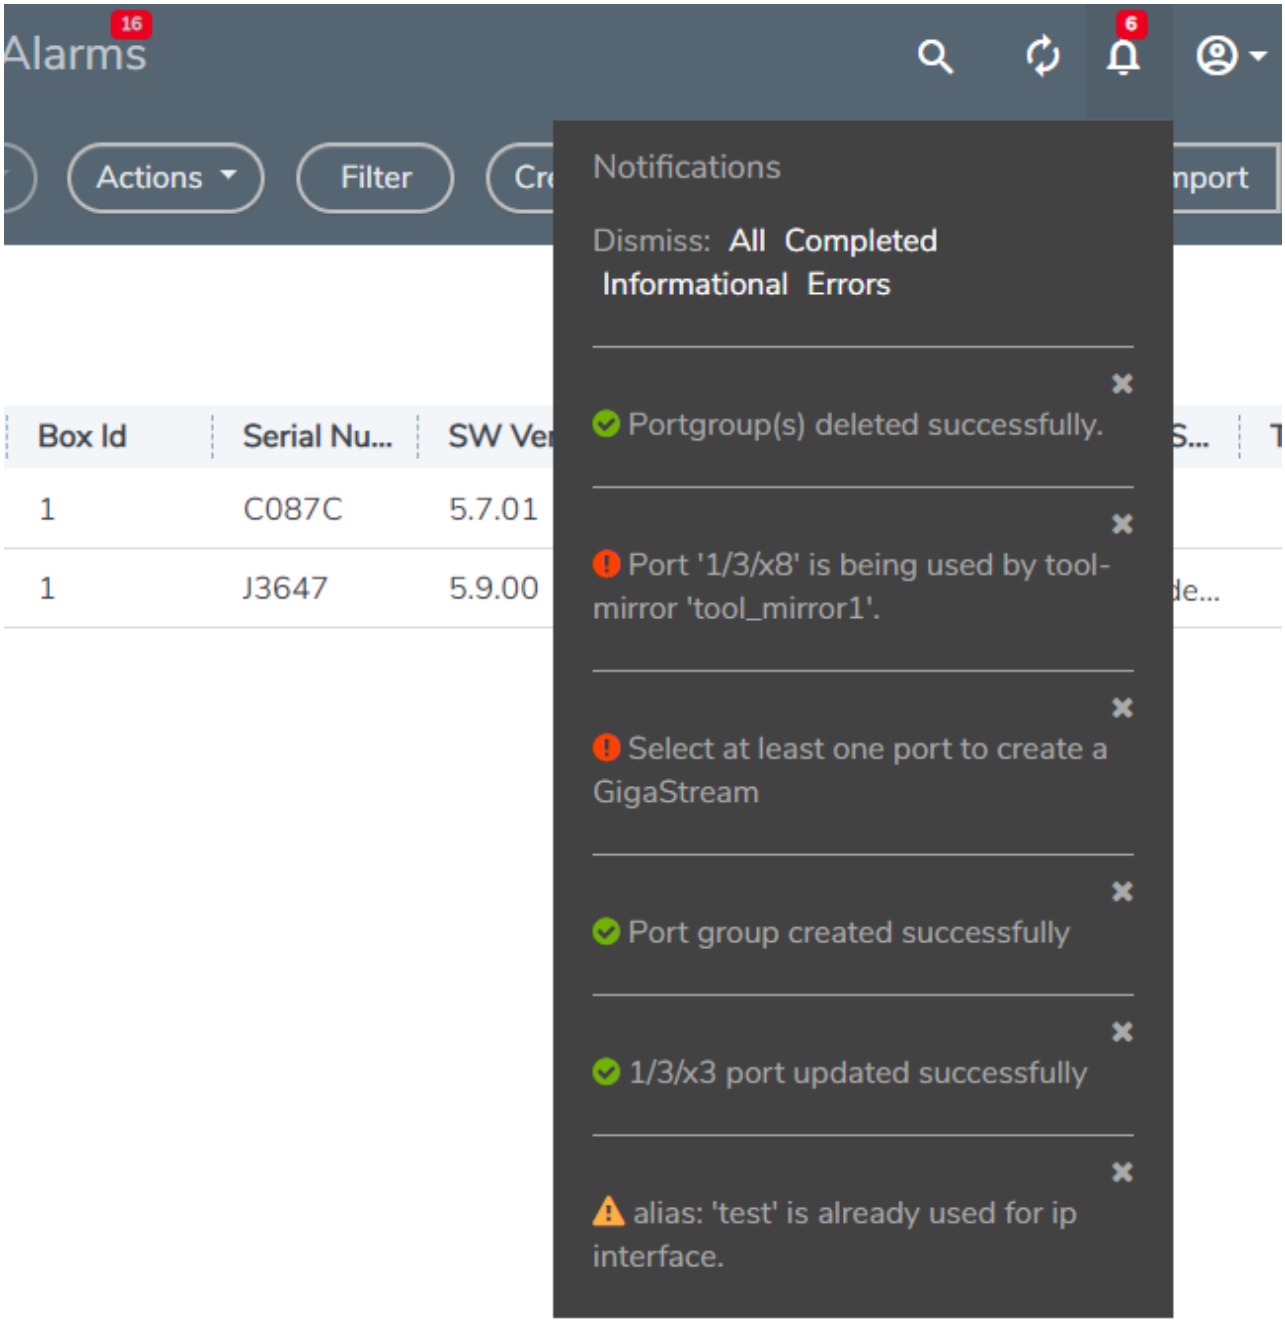


Figure 3 *Pop-up messages*

Long-term Notifications

Click on the Notifications icon to display a Notifications Panel listing long-term alert messages. If the list is long, a scroll bar appears on the left so you can scroll through the list.

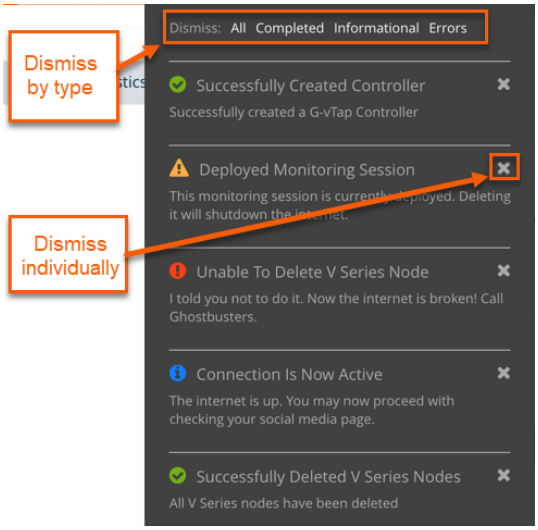


Figure 4 *Long-term messages*

Messages are updated as more information is received from packets. You can dismiss the individual messages or by type.

Notification Type Icons

Notification messages include a summary and a notification-type icon indicating the severity level of the alert. Some notifications have titles as well.

	Process Completed (green circle with a check mark)
	Warning (yellow triangle with an exclamation point)
	Error (red circle with an exclamation point)
	Information (blue circle with an “i”)
	Alert being processed (gray spinner)

Notification Banners

Notification banners are displayed in the GigaVUE-FM GUI to alert users of critical or warning situations that require immediate user attention, such as the following:

- when disk space exceeds a predefined threshold limit.
- when the licenses are in expired state
- when Volume-based License (VBL) usage exceeds the daily usage allowance limit

Refer to the GigaVUE Licensing guide for banners related to licenses.

NOTE: You must be a read-only user to view the notification banners.

Notification Banners for Disk Space Utilization

The disk space in GigaVUE-FM is used to store statistical data and syslogs. External log files, sysdumps, and device image files also take up disk space. It is important to monitor the disk space, as increased usage beyond a threshold level results in system slowness, system instability, and UI inaccessibility. If appropriate action is not taken GigaVUE-FM will stop functioning.

GigaVUE-FM monitors the disk space utilization of the /config directory. Whenever the disk space usage exceeds a pre-defined threshold level, GigaVUE-FM performs the following steps sequentially:

1. Cleans the statistical data.
2. Disables syslog and stats collection services if the disk usage is still above the threshold. Corresponding alarms are triggered and displayed in the Alarms page. Email notifications are also sent to notify about the critical disk usage levels (if configured).
3. Sends SNMP traps

To alert the users immediately about the critical disk usage levels, starting from software version 6.2.00, notification banners are displayed when disk space usage level exceeds a specific threshold limit. Notification banners are displayed in the following cases:

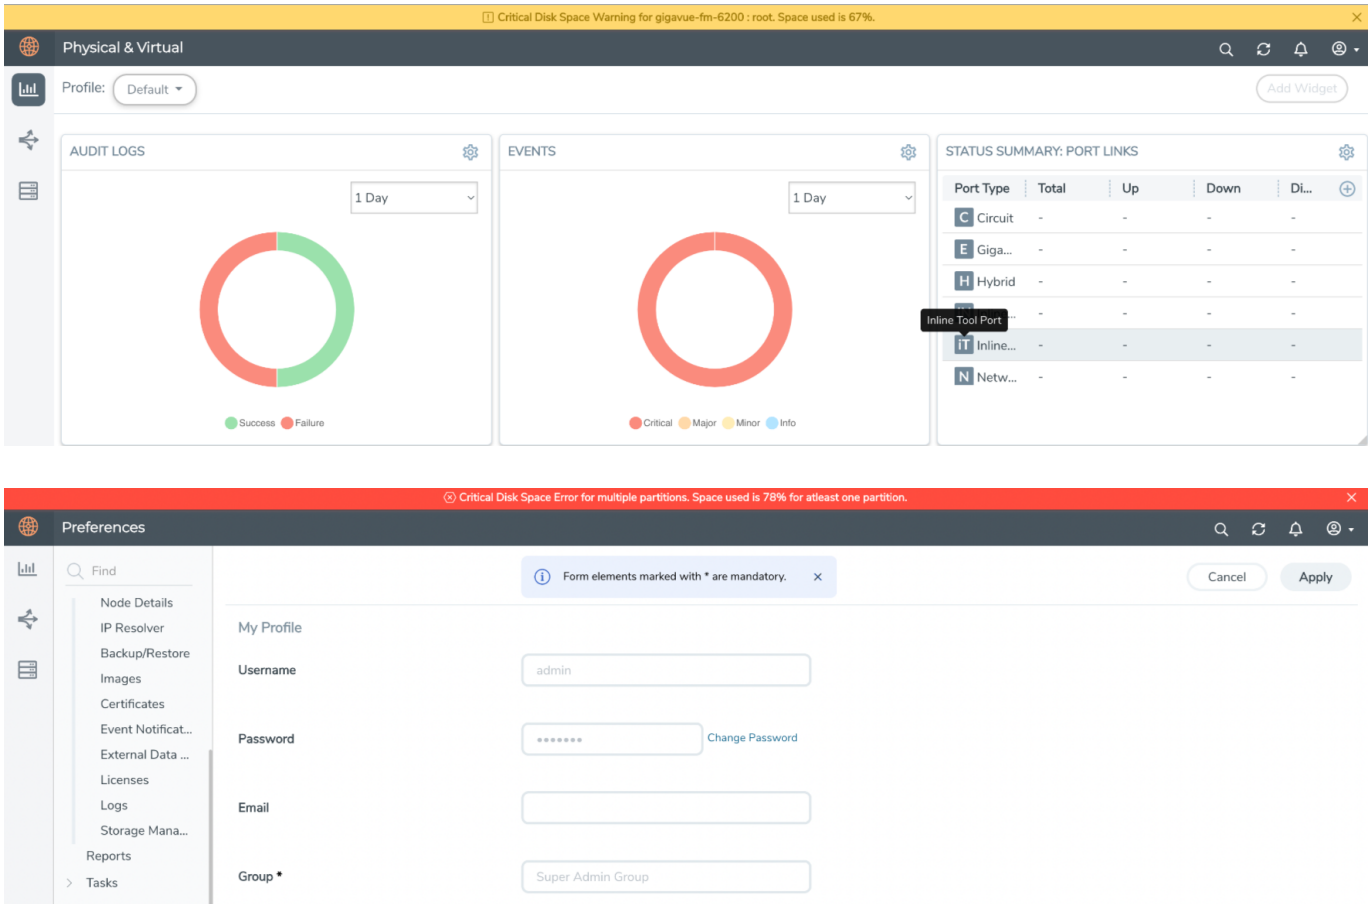
- When you login to a GigaVUE-FM instance.
- Near Real-time notifications also get displayed at the time when the disk usage exceeds the configured threshold value.

NOTE: Notification banners for disk space utilization alerts users about critical disk usage levels in all three partitions (/config, /var, /) of GigaVUE-FM.

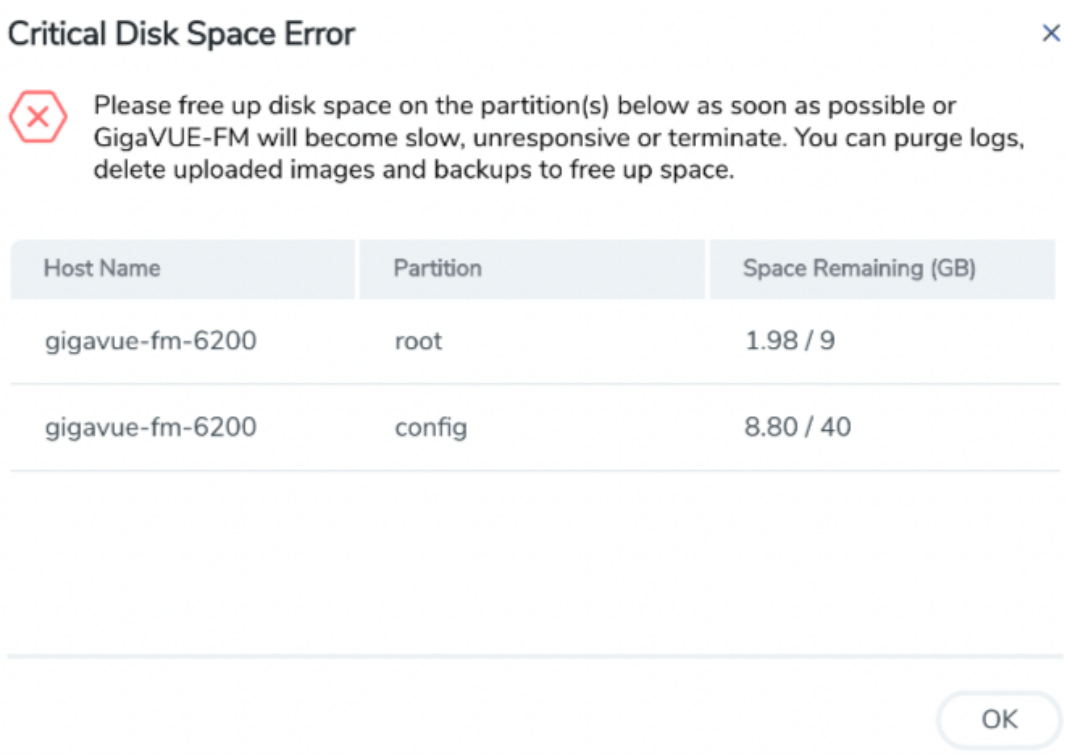
Depending on the severity, the color of the banner varies:

- **Warning - Yellow-colored banner**: If disk space utilization has reached 60%.
- **Critical - Red-colored banner**: If disk space utilization has exceeded 75%

If the severity of the disk space utilization varies across the partitions, then the color of the banner is determined based on the highest severity level.



If you click on the notification banner ribbon, a pop-up with details about the disk space utilization is displayed. See the following example:



You may experience a slight delay while the notification banner pops-up and clears.

Disk utilization monitoring is also applicable in dynamic GigaVUE-FM High Availability setups. The details of the GigaVUE-FM HA host names are displayed in the pop-up.

NOTE: For information about how to free-up disk space refer to How to Clean up Disk Space on a GigaVUE-FM Instance for details.

How to Add the GigaVUE-FM Instance Name

The default name of the GigaVUE-FM instance is GigaVUE-FM, and it is displayed on the footer.

When you have multiple GigaVUE-FM Instances running in your system, it becomes difficult to differentiate the instances and switch between tabs.

To customize the GigaVUE-FM instance name:

1. Go to **Settings** and select **System** > **Preferences**.
2. Click **Edit**.

3. In the **FM Instance Name** box, enter a name for the GigaVUE-FM instance.
4. Click **Save**.

   The customized GigaVUE-FM instance name is displayed in the footer.

How to Search in GigaVUE-FM

When searching for items, GigaVUE-FM performs a full search across multiple categories and displays the categories as part of the results. A Filter option is also available, which displays as a quick view, making it possible to quickly refine the search results.

GigaVUE-FM has an global search feature that allows you to search for information in GigaVUE-FM as well all the devices and device configurations managed by GigaVUE-FM. Essentially, if it is part of the GigaVUE-FM UI or managed by GigaVUE-FM, you can search for items based on keywords because almost all items are indexed for global search. GigaVUE-FM enhances port alias search, allowing you to enter full or partial aliases to instantly find the corresponding port and all directly associated resources including maps, GigaStreams, inline components, port groups, and interfaces consolidated in a single result set for faster troubleshooting and impact analysis. Refer to the Example 3: Search for Port and Linked Resources Using a Port Alias.

NOTE: The clusterId is a string, and sorting is performed lexicographically (dictionary order), not by numeric IP order.
For example:
Searching the clusterId "192.168.2.200" vs "192.168.2.79".
192.168.2.200 comes before 192.168.2.79 in ascending order because "200" is compared as text, and "2" comes before "7".
The clusterId is stored as a string since it needs to accommodate both IP addresses and domain names.

The search feature allows you to search for items across the following:

- Maps
- Roles and Users
- GigaSMART
  - GigaSMART Operations
  - GigaSMART Groups
  - Virtual Ports
  - NetFlow/IPFIX Generation
  - SSL Decryption
  - GTP Whitelists
  - Application Session Filtering (ASF)

- Ports
  - Port Groups
  - Port Pairs
  - Tool Mirrors
  - Stack Links
  - Tunnel Endpoints
  - IP Interfaces
  - Circuit Tunnels
  - GigaStreams
  - Port Alias
- Chassis and port inventory
- Node Clusters
- VMs
- IP, DNS, and MAC address

The following are not currently searchable: audit logs, events, NetFlow data, RBAC, IP ranges, or statistics.

Categories are another important component of Elastic Search. When you provide a keyword, the search system displays the matching category or categories related to the keyword search. This provides an automatic filtering of the keyword, helping to narrow your search. Some of the general categories are:

- Cluster
- GigaSMART
- Inline Bypass
- Maps
- NSX-V
- Ports
- Users
- VMware

You can refine the search categories by using the Filter feature (refer to Filtering Search Results) to further narrow the search results. For example, if the search keyword falls into the GigaSMART category, you can narrow the search further to NetFlow Records.

Performing a Search

To find an item in GigaVUE-FM, do the following:

1. Click the **Search** icon in the GigaVUE-FM top menu to open the Keyword field.

2. Type a keyword in the text field.

   As you enter the keyword, the system displays the categories in which the keyword appears and the total matches in that category along with the specific instances. For example, as shown in Figure 5GigaVUE-FM Search, you start to enter an IP address. The search shows that 10.115 occurs in the Cluster category 29 times, Ports 1 time, and IP Interfaces 2 times.

NOTE: You can type up to 128 characters in the **Keyword** field. This is because you cannot create a component with alias more than 128 characters (Map aliases, Port aliases and other such aliases).
The Search results start appearing when you type a minimum of three characters in the search pane.

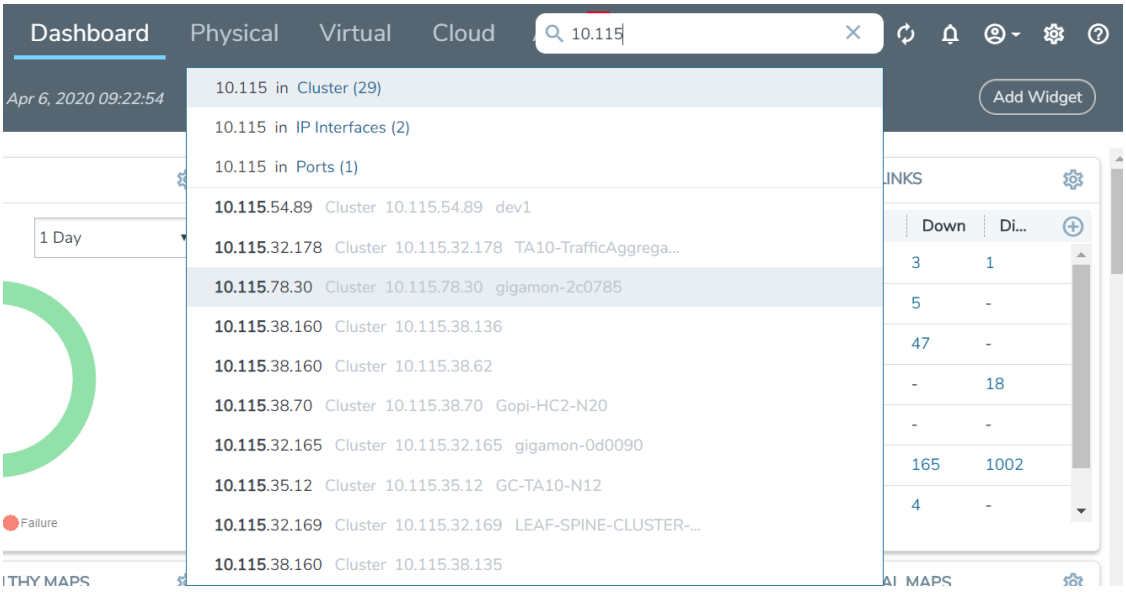


Figure 5 *GigaVUE-FM Search*

3. Once you are done entering the keyword, you can scroll through the list by using the up and down keys on the keyboard. Select an item by pressing the Enter key.
   - Selecting a category, displays all the results in that category, using the category as a filter. You can further refine the results by clicking the **Filter** button. For details on filtering, refer to Filtering Search Results.
   - Selecting a specific result opens the page for that item.

     If the search result is a cluster or standalone node, clicking the results takes you to the Overview page of that node and the Keywords field displays the node's ID as shown in the following figure. When the node ID is displayed in the Keywords field, it indicates that the scope of searches is narrowed to the current node or cluster.

NOTE: The clusterId is a string (keyword), sorting is lexicographical (dictionary order), not numeric IP order.
For example: "192.168.2.200" vs "192.168.2.79" "2" (in 200) comes before "7 (in 79)" → so 200 comes before 79

Search Examples

This section provides few examples to show how to use the global search feature. The examples cover the following:

- Search Maps
- Search for Roles and Users
- Search Ports

Search Maps

You can search for maps based on the map alias, IP address associated with maps, MAC address, port status and so on. Click on the **Filter by Cluster** icon in the search results page to refine the search results based on the cluster ID.

This section provides several examples of searching Maps:

- Example 1: Search for a Map by Alias
- Example 2: Search for a Map with an IP Address
- Example 3: Search for a Map with a MAC Address
- Example 4: Search for Maps with Down Ports

Example 1: Search for a Map by Alias

In this example, you are looking for a map where you remember the map's alias but are not sure which node or cluster it is on.

1. Click the Search icon in the GigaVUE-FM header to open the Keyword field and start typing the map alias. In this example, Up.

   As you type, the search displays the categories and items that match the keyword. Figure 6 Searching by Map Alias shows the search results and you can see that the map with Alias Up is on node 10.60.94.73 without typing the entire string.



Figure 6 *Searching by Map Alias*

2. In the search results, click on Up-1.

   GigaVUE-FM opens the Map page as shown in Figure 7Map Page.

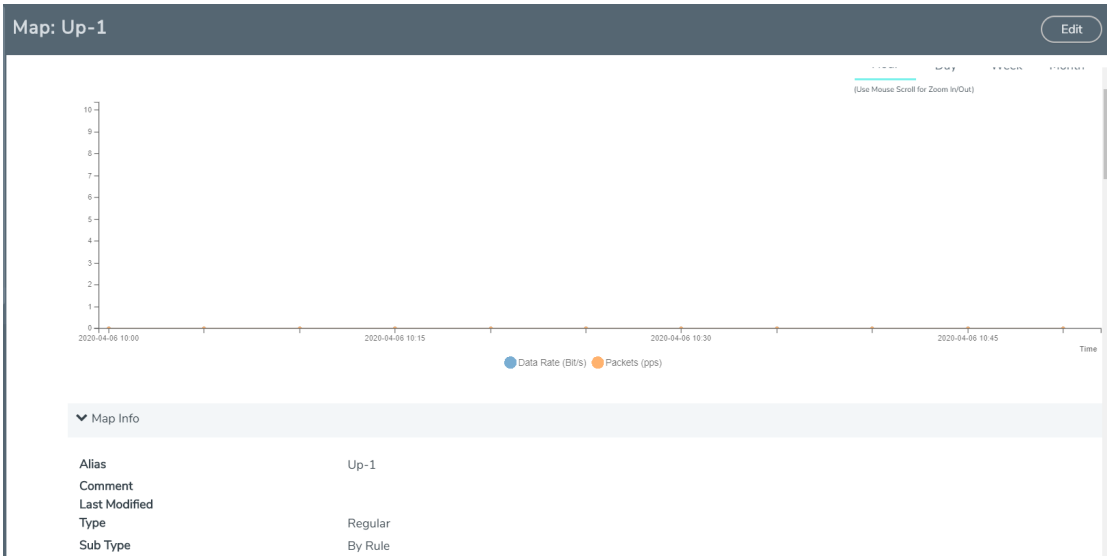


Figure 7 *Map Page*

Example 2: Search for a Map with an IP Address

In this example, you are searching for a map or maps that contain a specific IP address.

1. Click the Search icon in the GigaVUE-FM header to open the Keyword field and start typing an IP address. In this example, the IP address is 10.60.94.73.

   As you type the search displays the categories and items that match the keyword. Figure 8Search for IP Address shows the search results.

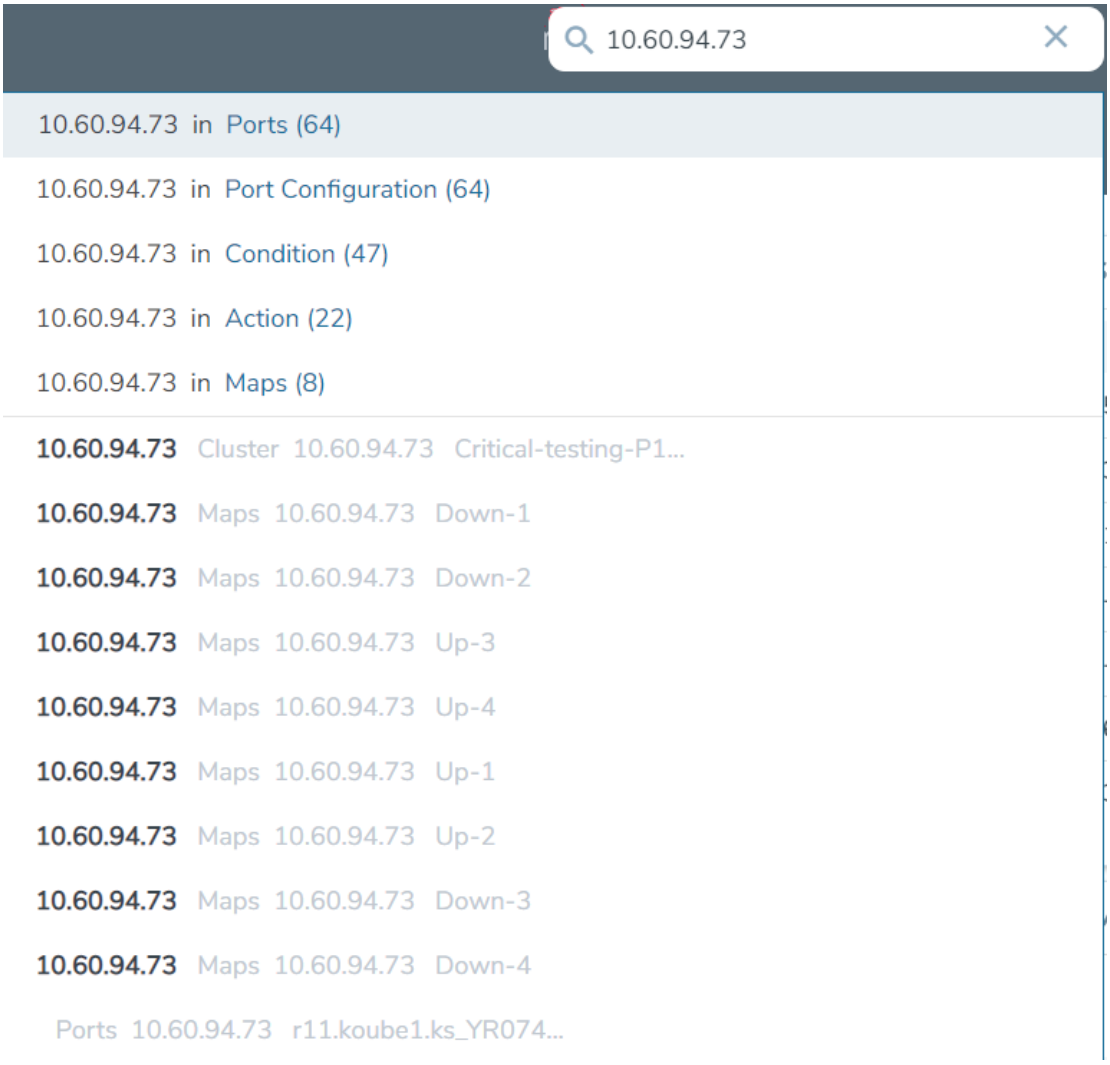
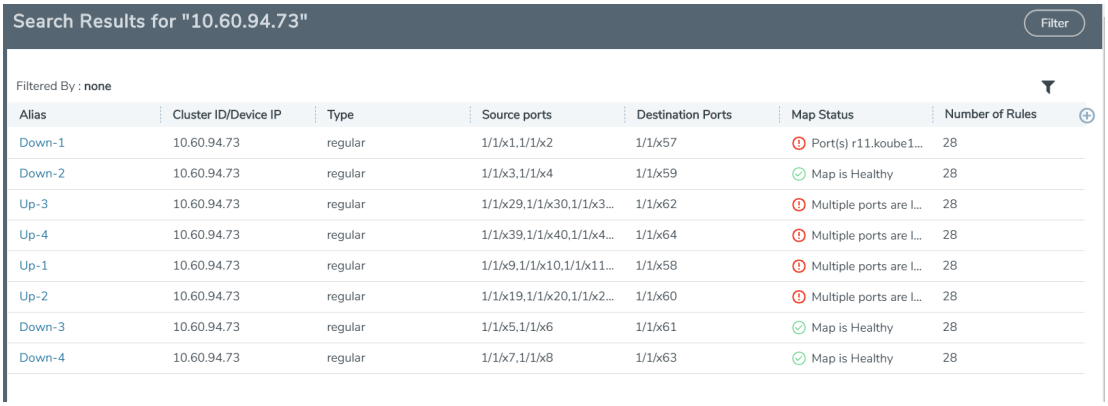


Figure 8 *Search for IP Address*

2. Click on Maps in the categories section of the search results.

   The Filter page opens. Figure 9Search Result Filter Page shows the results with maps the that contain the IP address. In this scenario, the item of interest is Down-1, so you click on the cluster name.



Alias	Cluster ID/Device IP	Type	Source ports	Destination Ports	Map Status	Number of Rules
Down-1	10.60.94.73	regular	1/1/x1,1/1/x2	1/1/x57	Port(s) r11.koube1...	28
Down-2	10.60.94.73	regular	1/1/x3,1/1/x4	1/1/x59	Map is Healthy	28
Up-3	10.60.94.73	regular	1/1/x29,1/1/x30,1/1/x3...	1/1/x62	Multiple ports are l...	28
Up-4	10.60.94.73	regular	1/1/x39,1/1/x40,1/1/x4...	1/1/x64	Multiple ports are l...	28
Up-1	10.60.94.73	regular	1/1/x9,1/1/x10,1/1/x11...	1/1/x58	Multiple ports are l...	28
Up-2	10.60.94.73	regular	1/1/x19,1/1/x20,1/1/x2...	1/1/x60	Multiple ports are l...	28
Down-3	10.60.94.73	regular	1/1/x5,1/1/x6	1/1/x61	Map is Healthy	28
Down-4	10.60.94.73	regular	1/1/x7,1/1/x8	1/1/x63	Map is Healthy	28

Figure 9 *Search Result Filter Page*

A Map quick view opens, showing the information for the map.

Example 3: Search for a Map with a MAC Address

In this example, you are searching for a map or maps that contain a specific MAC address.

1. Click the **Search** icon in the GigaVUE-FM header to open the Keyword field and start typing an IP address. In this example, the IP address is 11:11:11:11:11.

   After entering the MAC address in the Keyword field, only one map is found as shown in Figure 10Search for MAC Address.



Figure 10 *Search for MAC Address*

2. Click on the result with the map named MacSrcMap to view the map details.

Example 4: Search for Maps with Down Ports

In this example, you are looking for maps that have a port that is in the “down” state.

1. Click the Search icon in the GigaVUE-FM header to open the Keyword field and type down.

   As shown in Figure 11Search Results for down Keyword, the keyword occurs in several categories.

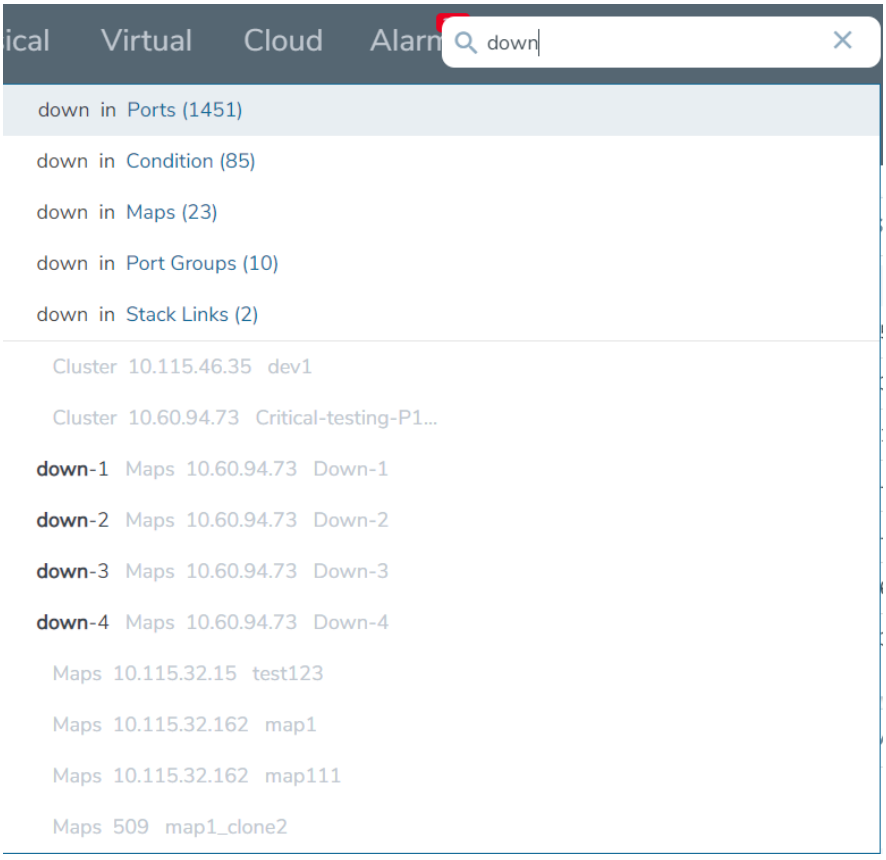


Figure 11 *Search Results for down Keyword*

2. Because you are searching for Maps, you click on the Maps category.

Search for Roles and Users

This section provides examples of searching for a role and for a user:

- Example 1: Search for Monitor Role
- Example 2: Search for a User

Example 1: Search for Monitor Role

In this example, you are looking for where the fm_User role is applied.

1. Click the Search icon in the GigaVUE-FM header to open the Keyword field and type Monitor.

   As shown in Figure 12Categories for Monitor., the keyword occurs in the several categories.

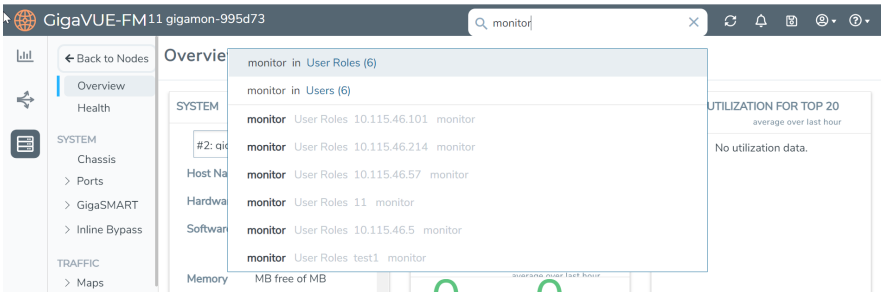


Figure 12 *Categories for Monitor.*

2. Select the User Roles category. The search results shows the nodes where the Monitor role is used.

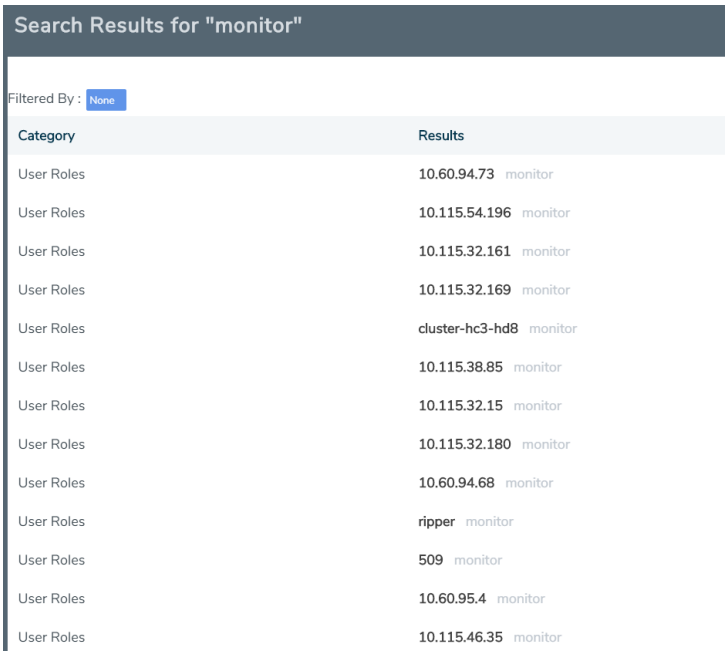


Figure 13 *Search Results for Monitor.*

3. From the search results, drill down further by selecting one of the results.

   GigaVUE-FM takes you to the User Setup page for the selected cluster as shown in Figure 14Role From Search Result and indicates in the Keywords field that further searches are restricted to the cluster.

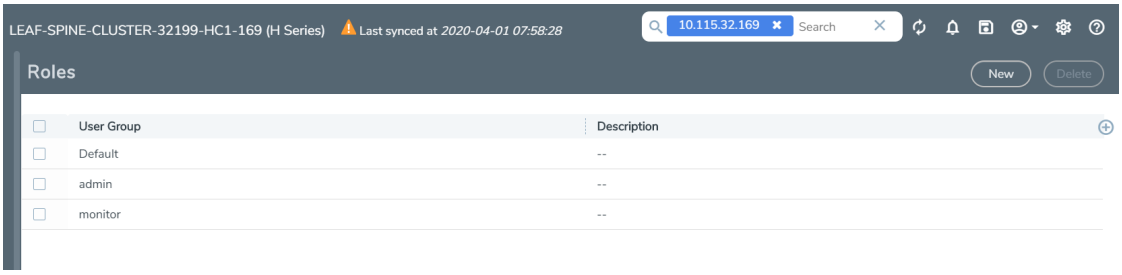
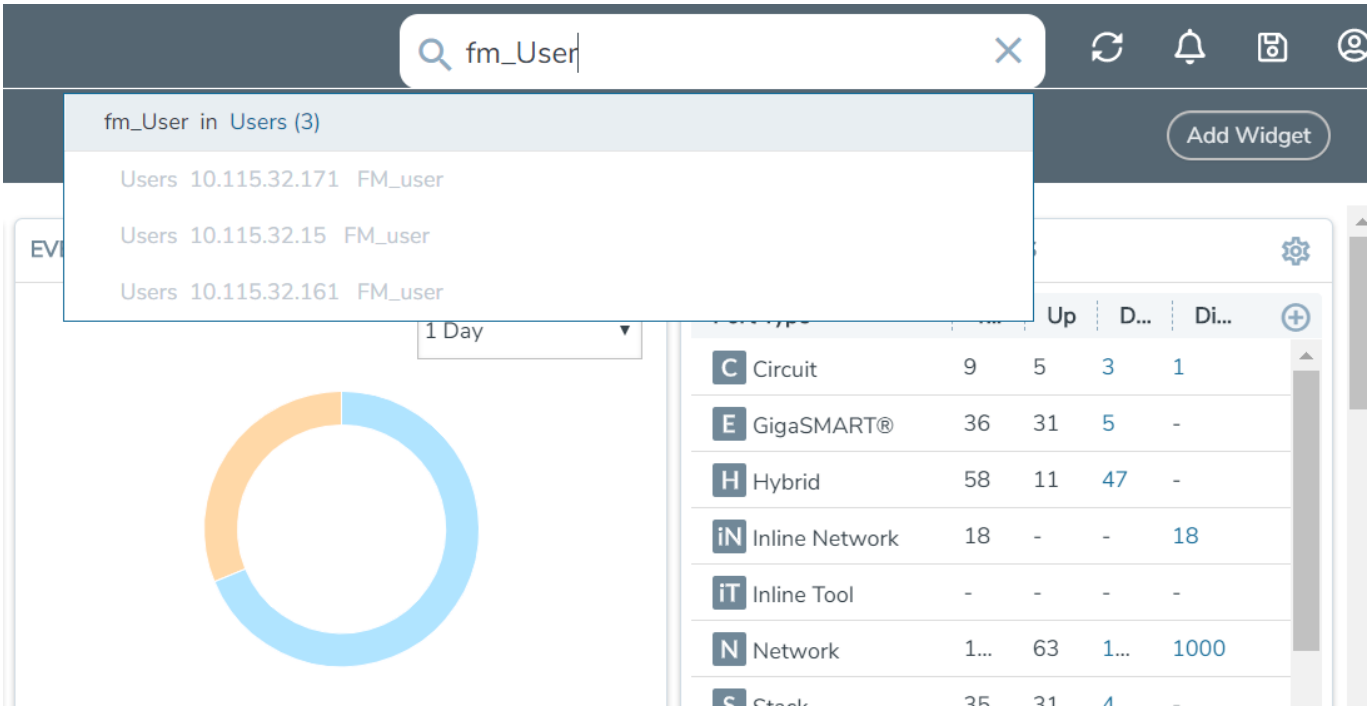


Figure 14 *Role From Search Result*

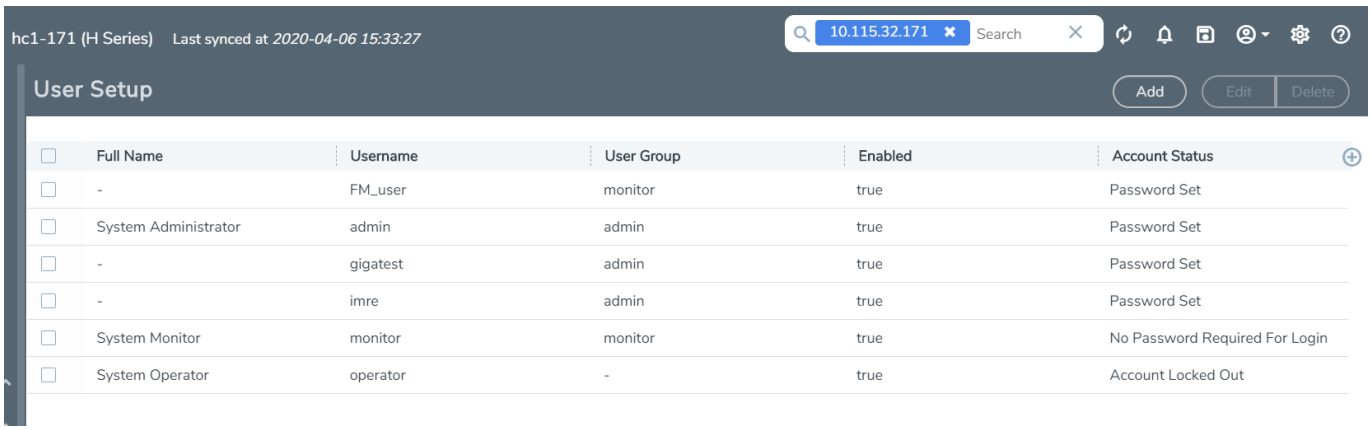
Example 2: Search for a User

In this example, you are looking for a specific user.

1. Click the Search icon in the GigaVUE-FM header to open the Keyword field and type the user name for the user. In this example, the user is *fm_User*.



2. Click on *fm_User* in the categories and items list. The User page on the node with the *fm_User* opens a shown in the following figure. The Keywords field also indicates further searches are restricted to the current node.



Search Ports

NOTE: You can search for ports based on the port id, port alias, cluster ID/Device IP and so on. You can also view the neighboring ports information from the port search results page. Click on the **Filter by Cluster** icon in the search results page to refine the search results based on the cluster ID.

This section provides few examples related to searching for ports:

- Example 1: Search for Down Ports
- Example 2: Search for Port Details of Devices Managed by GigaVUE-FM
- Example 3: Search for Port and Linked Resources Using a Port Alias

Example 1: Search for Down Ports

In this example, you are searching for a particular port by its ID. This example also shows how to combine keywords.

1. Click the Search icon in the GigaVUE-FM header to open the Keyword field and type the port ID 1/1/x5 followed by the keyword *down*.

   As shown in Figure 15Categories Returned for Port ID and Down, the 1/1/x5 and down occur 60 times in the Port category.

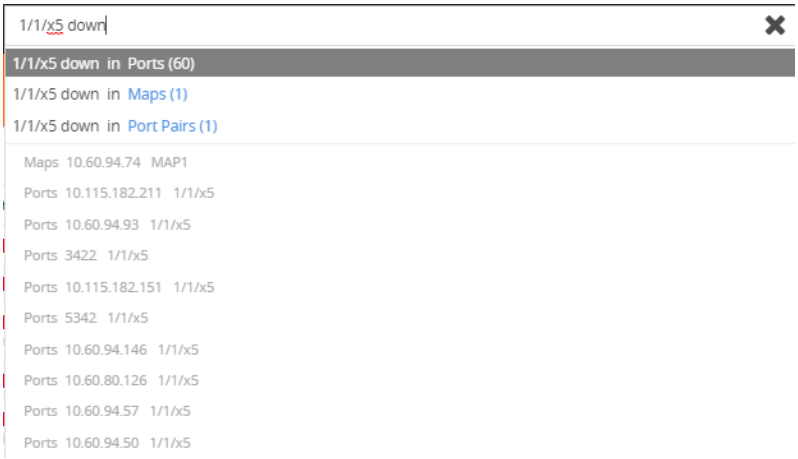


Figure 15 *Categories Returned for Port ID and Down*

2. Click on the Ports category to view the results.
3. Click on an item search results to see the port information. A Port quick view displays for the selected port as shown in Figure 16Port Quick View for Search Results.

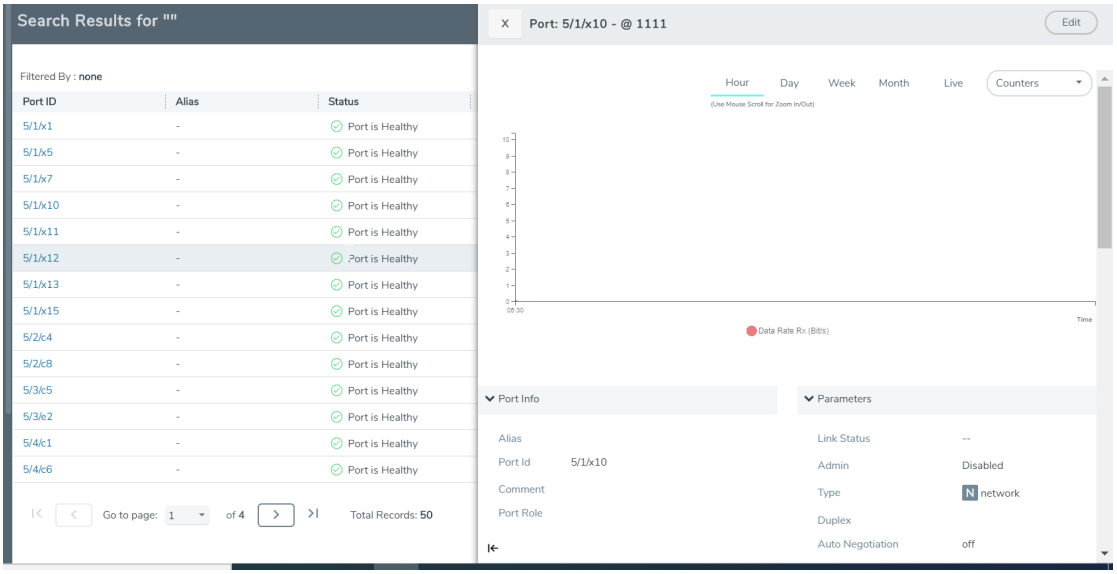


Figure 16 *Port Quick View for Search Results*

Example 2: Search for Port Details of Devices Managed by GigaVUE-FM

In this example, you want to retrieve port information from all the devices managed by GigaVUE-FM and that are up. This example also shows the use of a non-alphabetic character as the keyword.

1. Click the Search icon in the GigaVUE-FM header to open the Keyword field and type the back-slash character (/) and up.

   As shown in Figure 17Categories Returned for Port Details of All Devices and up, the search returns results in several categories.

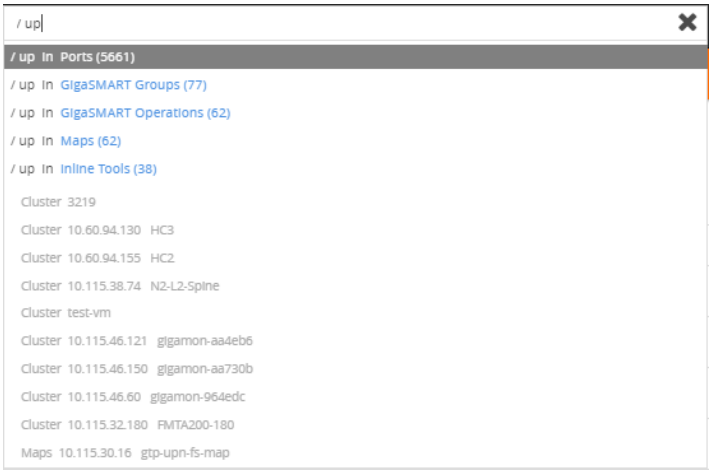


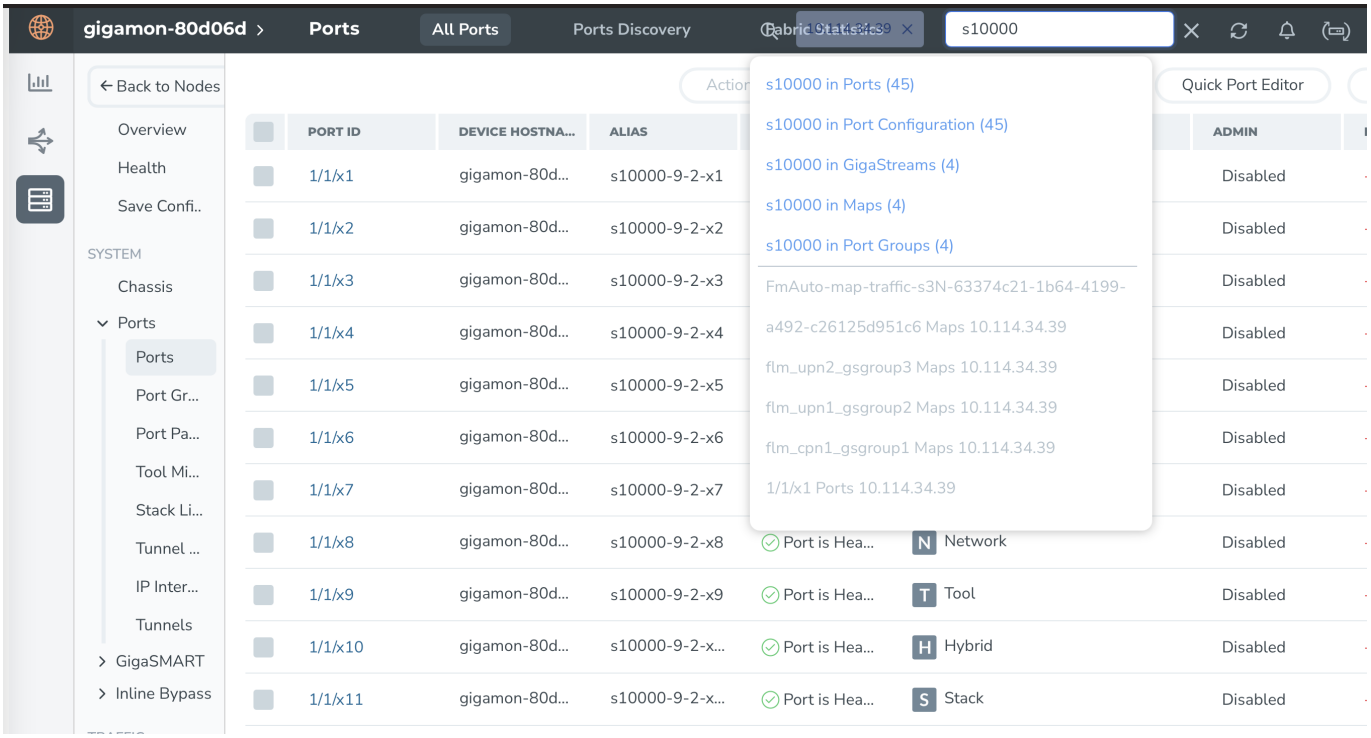
Figure 17 *Categories Returned for Port Details of All Devices and up*

2. Click on the Ports category to view the results.
3. Click on an item search results to see the port information.

Example 3: Search for Port and Linked Resources Using a Port Alias

In this example, to search using a port alias, perform the following steps:

1. Log in to GigaVUE-FM.
2. Click **Inventory > Physical > Nodes**.
3. Select the node.
4. Select **System > Ports > Ports**.
5. In the global search, searching by port alias returns the corresponding port and all directly associated resources, including maps, GigaStreams, inline components, port groups, and interfaces.



Filtering Search Results

When selecting a search result for a category, GigaVUE-FM opens a page that lists the results for that category.

To change the filter, do the following:

1. Click **Filter**.

   The Filter quick view displays.
2. Add or remove filters by selecting items from the Filter quick view.

How to Apply Filters

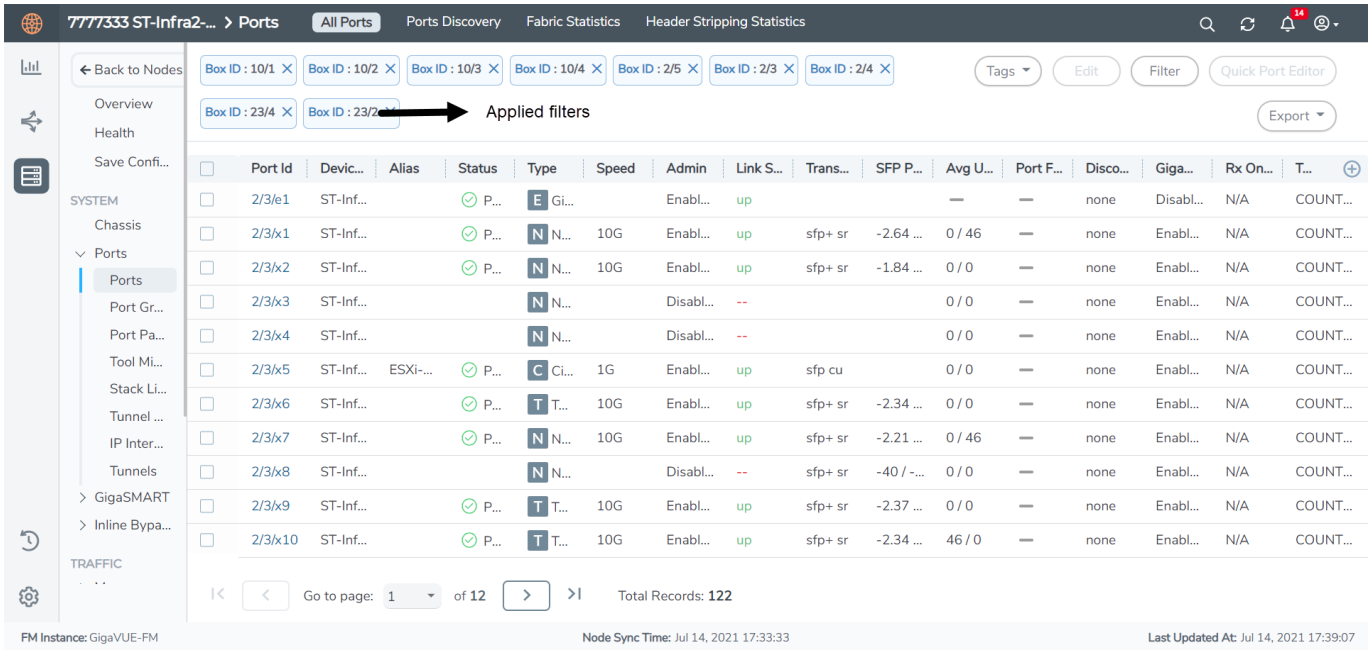
The filter functionality allows you to search and narrow down the options you want to display on a particular page.

To use the filter functionality, do the following:

1. Click on the **Filter** button.
2. The **Filter quick view** dialog is displayed.
3. Specify the parameters to be filtered.

   The filter selection appears above the list for reference. To remove a particular filter, click on the '**x**' icon next to the filtered item.
4. Click the '**x**' icon to exit the Filter quick view dialog.

The following figure shows how the applied filters are displayed on the GigaVUE-FM instance page:



Dashboard

This section describes the Dashboards available in GigaVUE-FM. Refer to the following table for details:

Left Navigation Pane	Description	Reference
Overview **Physical and Virtual**	Provides a quick visual overview of the traffic, health status, inventory and audit logs of the physical and virtual nodes managed by GigaVUE-FM.	Overview of the Physical and Virtual Dashboard
System		
Alarms	Lists the alarms triggered in GigaVUE-FM	Overview of Alarms
Audit Logs	captures audit logs for all users connected to GigaVUE-FM	
Events		
FM Health	Provides an overview of GigaVUE-FM health.	FM Health Dashboard
Analytics		Analytics

Physical and Virtual Dashboard

This chapter describes the dashboards that provide information about the physical nodes, ports, port links, maps, GigaSMART, audit logs, and events on a single page.

This chapter covers the following topics:

- Overview of the Physical and Virtual Dashboard
- Physical Dashboard Profiles
- Physical Dashboard Quick Views
- Physical Dashboard Widgets

Overview of the Physical and Virtual Dashboard

The Physical and Virtual Dashboard is a central location to monitor all the physical and virtual nodes and clusters that are managed by GigaVUE-FM. The widgets in the dashboard provides a quick visual overview of inventory and events, GigaSMART traffic, highest and lowest traffic by maps and ports, traffic comparison by tags, most and least utilized traffic, health status, and audit logs.

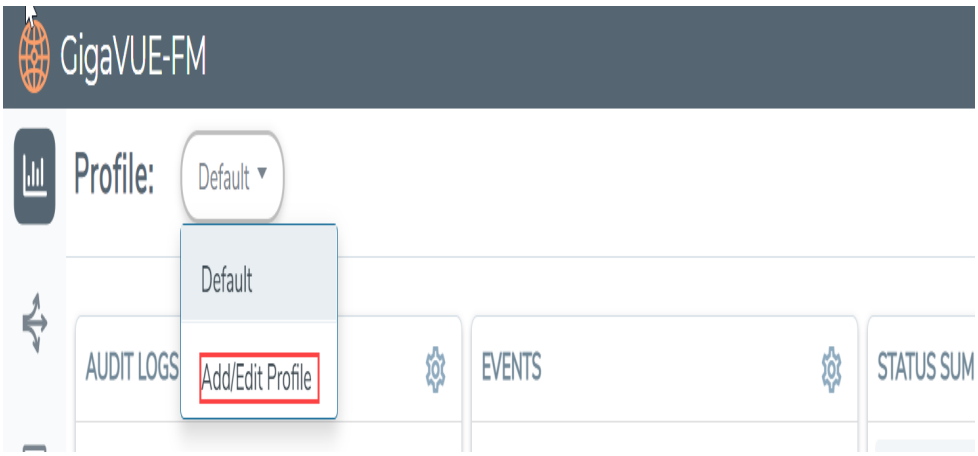
NOTE: You must scroll the individual widgets horizontally to view the invisible information.

Physical Dashboard Profiles

The Physical Dashboard displays a number of default widgets when you first log in. They are displayed with the profile labeled as **Default**. You can create multiple profiles and choose the widgets to be displayed in each profile based on the data you want to proactively monitor and troubleshoot in your Visibility Fabric.

To create a new profile

1. On the left navigation pane, click on  and from the Physical & Virtual dashboard page, select the profile drop-down and click **Add/Edit Profile**.



2. In the **Add/Edit Profile**... box, enter the name of the new profile and click **Enter**.

   The new profile name is displayed under Profiles. The new profile page is displayed.

3. (Optional) Click the Edit icon and select **Set as Login Profile** if you want the new profile to display as your default Physical Dashboard.

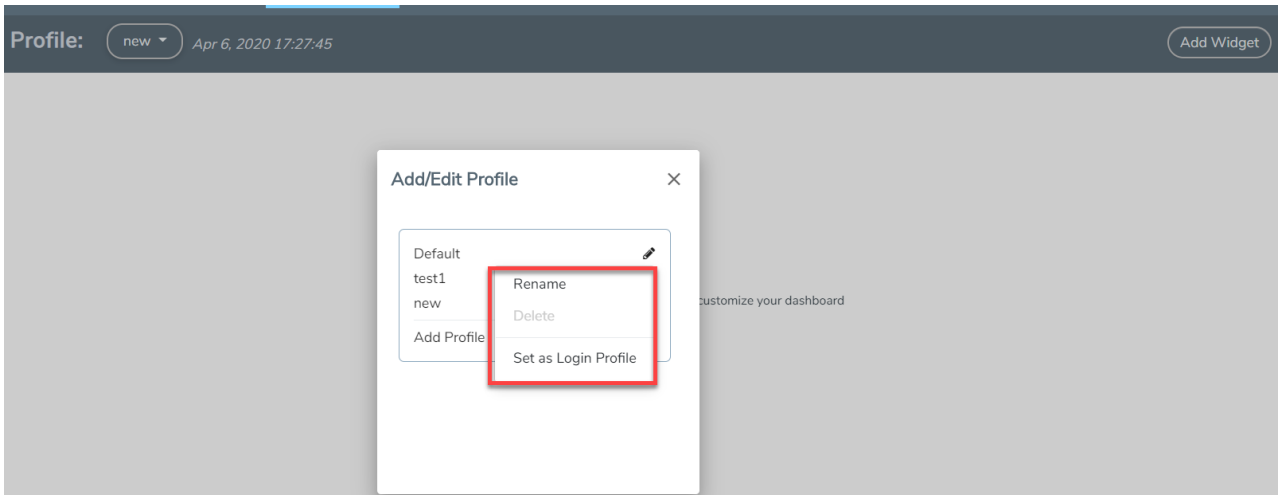


Figure 18 *Profile Settings*

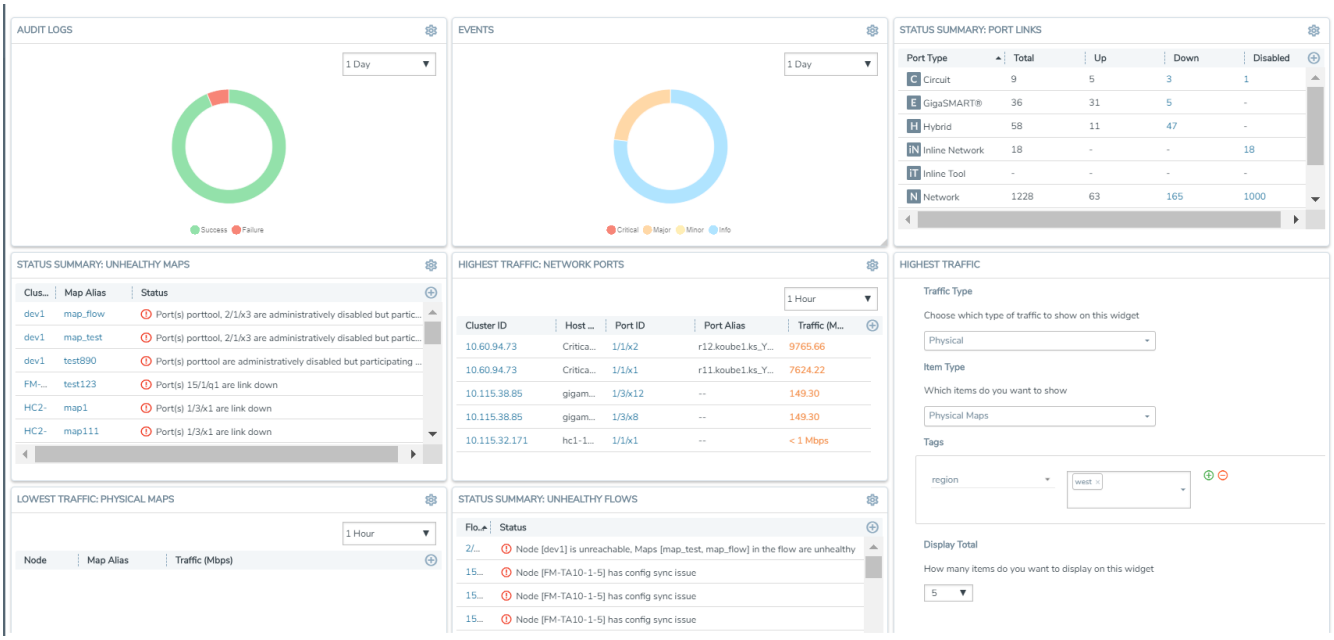
To change the profile name, click **Rename**, edit the name, and press **Enter**.

To delete the profile, click **Delete**. If you delete a default profile, then the initial default profile is automatically set as the login profile unless you actively select another one. Once deleted, there is no option to recover those dashboards.

NOTE: When GigaVUE-FM is upgraded, the profiles created in the previous version are not retained in the latest version.

Keep in mind the following regarding the widgets on the dashboard.

- Widget and trending data is available based on the GigaVUE-FM license purchased. For the base package, the data is not stored for more than 1 day. The prime package users can select any option including 1 month.
- Individual widgets can be resized and saved as part of the profile. Each widget can expand in both horizontal and vertical planes. The other widgets self-adjust when the widgets are manipulated.
- The widgets can also be dragged and dropped to different section of the page. Refer to Physical Dashboard Profiles .
- The data points can be viewed when the mouse is hovered over the graph as shown in Physical Dashboard Profiles .
- The widgets such as Unhealthy Maps opens a quick view when clicked on the cluster info or the map alias. The quick view shows more details relating to that specific map.
- The trending information can also be changed for each widget on the same dashboard.
- The port and map health status changes on the device reflect instantly on the screen.
- The color-coded legends are available at the bottom of each widget.



Note: If the percentage displayed in a pie-chart is negligible or less , then it would be difficult to click on the pie-chart arc and view the details .

Physical Dashboard Quick Views

When reviewing the widgets available on the Physical & Virtual dashboard, clicking on the options in the widgets takes you to the details page relating to the information for that node. For example, on the Nodes by Model or Software Version widget, you click on the node and it takes you to the Physical Nodes page.

For more information about Physical Nodes, refer to Manage GigaVUE® Nodes and Clusters.

Telemetry Insights Dashboard

The Telemetry Insight dashboard gives you a single, actionable view of the health, usage, licensing risks, and security posture of your GigaVUE-FM, along with recommended actions.

The Telemetry Insights dashboard has three major sections:

- Executive Health Summary
- Operational Summary
- Recommended actions

Prerequisites

- Enable GigaVUE-FM Connected Experience to populate the Telemetry Insight Dashboard. If you do not enable it, the dashboard shows no data, and GigaVUE-FM does not compute the insight data points.
- Enable SNMP traps on the devices and configure them to send traps to GigaVUE-FM.
- Enable the following trap categories:
  - Packet drops
  - Packet errors
  - Port link status
  - High and low utilization

Executive Health Summary

Once Connected Experience is enabled in GigaVUE-FM, the dashboard is populated with subsequent data at 5 AM UTC.

In the **Executive Health Summary** section, GigaVUE-FM shows a point-in-time health status of:

- Physical nodes such as GigaVUE ® HC Series and TA.
- Virtual nodes such as V Series nodes, and V Series-proxies.
- Universal Cloud Tap (UCT) such as UCT-V Controller, UCT-V, UCT-C, UCT-C Controller.

This section lets you quickly see whether your physical, virtual, and cloud tap components are healthy or need attention.

Operational Summary

In the **Operational Summary** section, GigaVUE-FM highlights the information about:

- **Compliance and Security**: GigaVUE-FM shows operational risks and inefficiencies over the last 7 days, which will help you to understand the performance issues and configuration problems without drilling into multiple device pages. This section helps you to quickly assess whether GigaVUE-FM and its access controls are configured according to internal policies. It also summarizes configuration elements that affect security, including:
  - FIPS Status- Whether GigaVUE-FM is running in FIPS mode.
  - Common Criteria Compliant- GigaVUE-FM checks settings such as FIPS, Crypto Mode, Network Audit Log, CRL Validation, and Lock Users, and ensures they are compliant.
  - Authentication mode (for example, basic authentication vs third-party).
  - API rate-limit exceed counts.
  - About to Expire tokens -Tokens that are expiring (for example, expiring within 7 days).

- **Licenses**: This section helps you to identify where you may soon lose support or entitlements so you can renew or migrate in time. You can view information about:
  - Installed and expired licenses.
  - License SKUs that are approaching End-of-Sale (EOS) or End-of-Life (EOL).
  - License SKUs that are already past EOS or EOL.
  - Volume-based license (VBL) over-usage in the last 7 days.
- **Aggregated Physical Nodes** - This section shows view of health, performance, and configuration-related signals across physical inventory managed by GigaVUE-FM, helping to identify operational risks, underutilized resources, and nodes that need attention. You can infer the following information:
  - Ports with High Packet Drops - Highlights physical ports experiencing sustained packet drops.
  - Admin Enabled Ports with no Traffic - Highlights ports that are administratively enabled but not carrying traffic.
  - Ports With High Utilization - Tracks ports with high utilization.
  - Unused Flow Maps - Shows configured flow maps that are currently not in use.
  - Physical Nodes with Older than LTS Version - Shows physical nodes running software versions older than the supported LTS baseline.
- **Aggregated Cloud Fabrics** - This section summarizes the health of V Series nodes, UCT-V agents. It combines the health signals for those cloud environments into a single cloud-level view so that you can quickly confirm whether there are problems in any of the managed cloud locations. When the aggregated cloud fabrics tile shows issues, you can use the rest of the dashboard and drill-down views to identify which specific cloud nodes or fabrics are affected. You can infer the following information:
  - GigaVUE V Series added and deleted - Shows the count of V Series nodes added to or removed from the environment during the reporting window.
  - UCT-V Inventory Added - Displays newly added UCT-V instances in the environment.
  - Virtual Nodes with outdated Software Version - Highlights virtual nodes running versions behind the supported software baseline.
  - GigaVUE V Series with Status Flaps - Highlights V Series nodes whose status is changing frequently between healthy and unhealthy states.

Recommended actions

In the recommended actions section, you can view the lists of actionable items that are derived from the other sections, such as:

- Expired Licenses
- About to expire License
- End Of Sale/End Of Life SKU
- Unreachable Nodes (Node Healthy) Physical
- Token about to expiry
- Basic Authentication
- GigaVUE-FM Storage Space (/Config /var /root)

This section helps you to decide what to fix first. For example:

- Issues marked with this symbol must be treated as Priority 1 (P1).
- Issues marked this with symbol must be treated as Priority 2 (P2).
- Issues marked with this symbol must be treated as Priority 3 (P3).

View the Telemetry Insight dashboard

To view the Telemetry Insights Dashboard, in GigaVUE-FM, go to Dashboard > **System** > **Telemetry**. The Telemetry Insights dashboard page appears.

Rules and Notes

- The Telemetry Insight Dashboard shows data only the next day after GigaVUE-FM Connected Experience is enabled. If you do not enable it, the dashboard appears but stays empty.
- Only Super Admin users can view the dashboard. This prevents incomplete views caused by RBAC limits in this release.
- If SNMP traps are not enabled or do not reach GigaVUE-FM, trap-based sections will be empty or partly filled. GigaVUE-FM still shows sections based on inventory and configuration.
- Dashboard time windows are relative (for example, last 7 days). Absolute dates are not supported.
- Configuration-change insights show only the changes made through GigaVUE-FM (GUI, REST API, or other GigaVUE-FM based methods). GigaVUE-FM does not capture changes made directly on devices, such as through the device CLI. This may cause lower reported configuration activity.

Physical Dashboard Widgets

This section describes the widgets that can be created and viewed on the Physical Dashboard.

- Highest Traffic
- Lowest Traffic
- Traffic Comparison By Tags
- Most Utilized Traffic
- Least Utilized Traffic
- Inventory
- Status Summary

The default profile displays the following widgets:

- Highest Traffic: Network Ports, Tool Ports, and Physical Maps
- Status Summary: Unhealthy Maps and Port Links
- Audit Logs
- Events

You can customize the widgets by modifying the physical dashboard profiles. Refer to Physical Dashboard Profiles for more information.

Highest Traffic

The Highest Traffic widget can be created for the following:

- Physical
  - Physical maps
  - Fabric maps
  - Network ports
  - Tool ports
  - Stack ports
  - Hybrid ports
  - Inline network ports
  - Inline tool ports
- GigaSMART
  - GigaSMART groups
  - GigaSMART operations

You can create as many Highest Traffic widgets as you want listing up to 5, 10, 15, 20, 50, or 100 items in each widget.

The highest traffic is measured in megabytes per second (Mbps). You can specify the period over which the amount of traffic must be calculated. The period can be 1 hour, 1 day, 1 week, or 1 month.

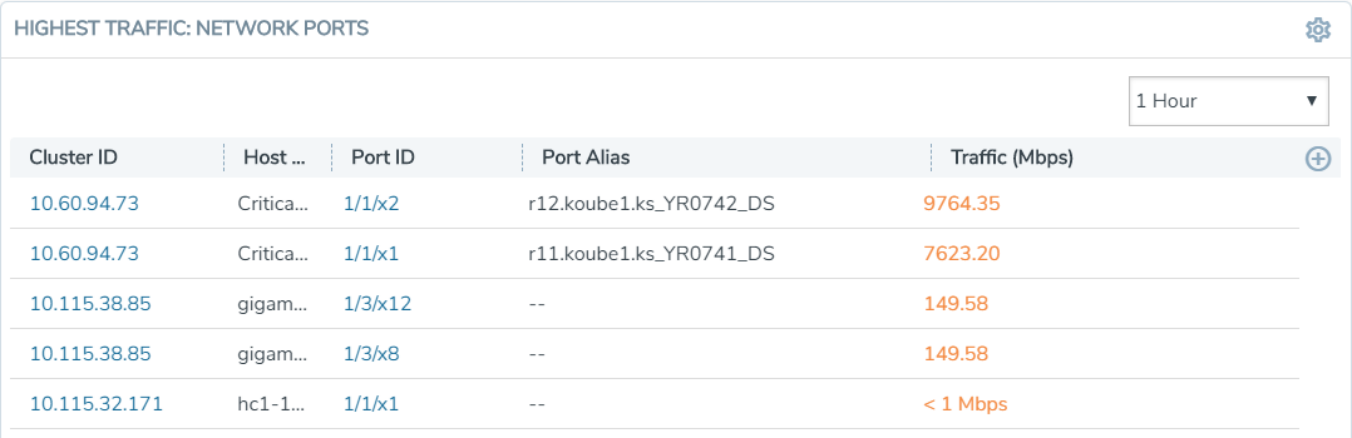


Figure 19 *Highest Traffic: Example*

The physical maps are listed by the node ID, map alias, and the traffic in Mbps.

The ports are listed by the node on which they are used and the port alias. You can create the Highest Traffic widget for the following ports:

- Network ports
- Tool ports
- Stack ports
- Hybrid ports
- Inline network ports
- Inline tool ports

The highest traffic for GigaSMART operations or GigaSMART group can be displayed as shown in Figure 20Highest Traffic GigaSMART.

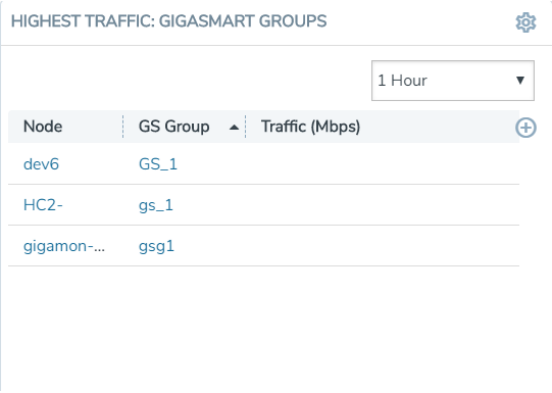


Figure 20 *Highest Traffic GigaSMART*

To configure the Highest Traffic widget:

1. On the left navigation pane, click on  and from the Physical & Virtual dashboard page, select the profile in which you want to add the widget.

2. Click **Add New Widget**. The Add New Widget window is displayed. Refer to Figure 21Add New Widget.

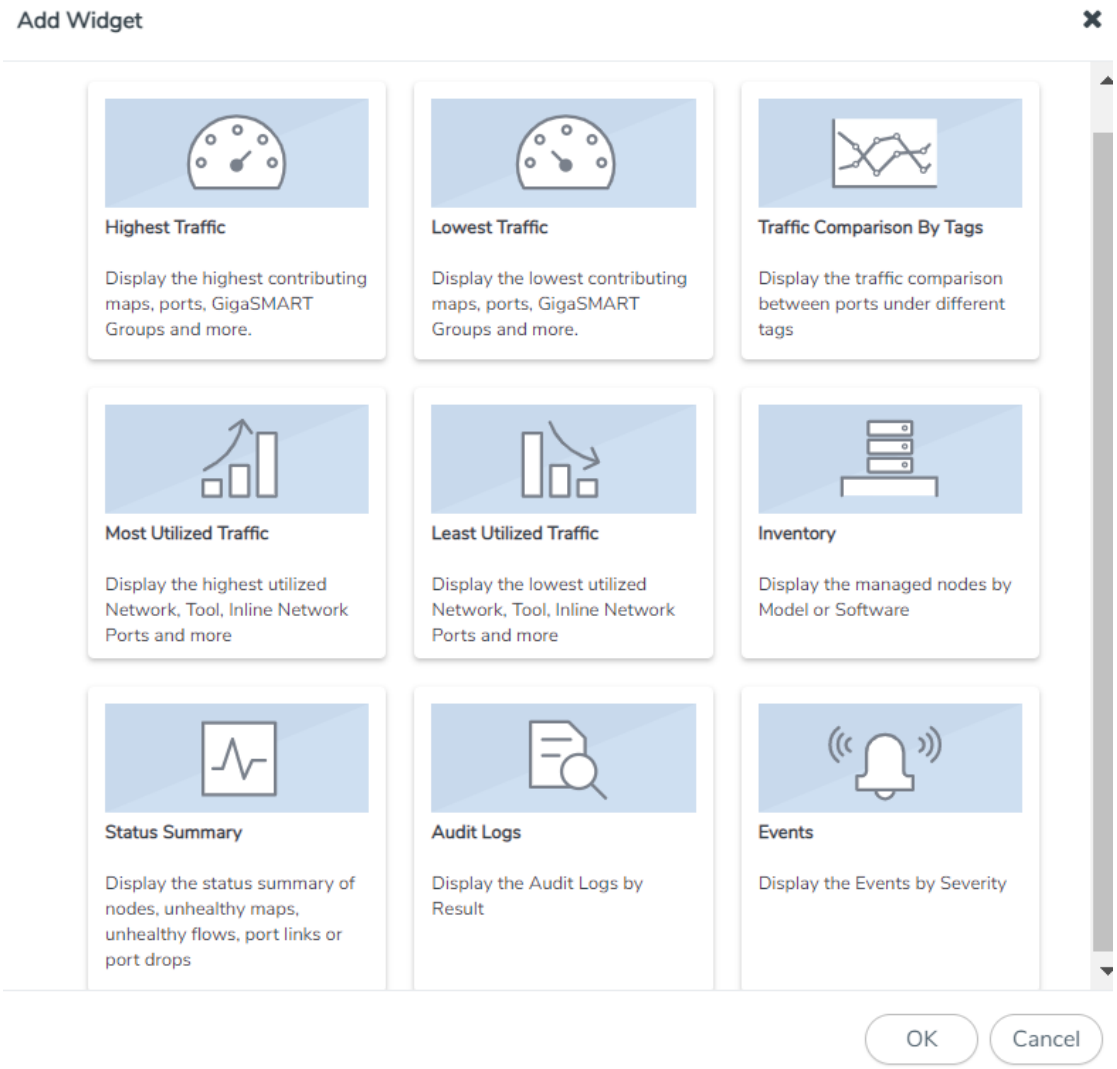


Figure 21 *Add New Widget*

3. In the Add New Widget window, select **Highest Traffic** and click **OK**. The Highest Traffic configuration window is displayed. Refer to Figure 22Highest Traffic Configuration.

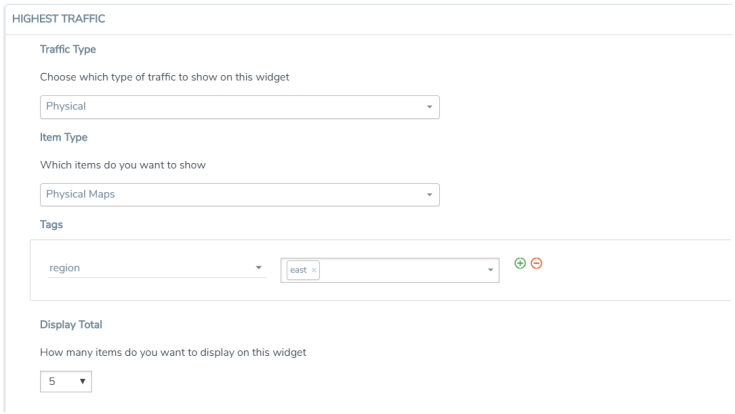


Figure 22 *Highest Traffic Configuration*

4. From the **Traffic Type** drop-down list, select one of the following traffic types:
   - Physical—Allows you to view the physical maps and ports that contribute to the highest traffic distribution.
   - GigaSMART—Allows you to view the virtual ports, GigaSMART groups, and GigaSMART operations that contribute to the highest traffic distribution.
5. From the **Item Type** drop-down list, select the item you want to view. The options displayed are based on the traffic type you selected in step 5.
6. Select the required tag key and tag value combination (for example: tag key is 'Site' and tag value is 'East') for which the highest traffic distribution must be displayed. This step is optional.
7. From the **Display Total** drop-down list, select the number of items to be displayed. By default, the number of items selected for display is 5.
8. Click **OK**.

Lowest Traffic

The Lowest Traffic widget lists the physical maps, flow maps, ports, and GigaSMART that contribute to the lowest traffic within a specified time. You can create as many Lowest Traffic widgets as you want listing up to 5, 10, 15, 20, 50, or 100 items in each widget.

The traffic flowing through a port or a map rule is measured in megabytes per second (Mbps). You can specify the period over which the amount of traffic is calculated. The period can be 1 hour, 1 day, 1 week, or 1 month.

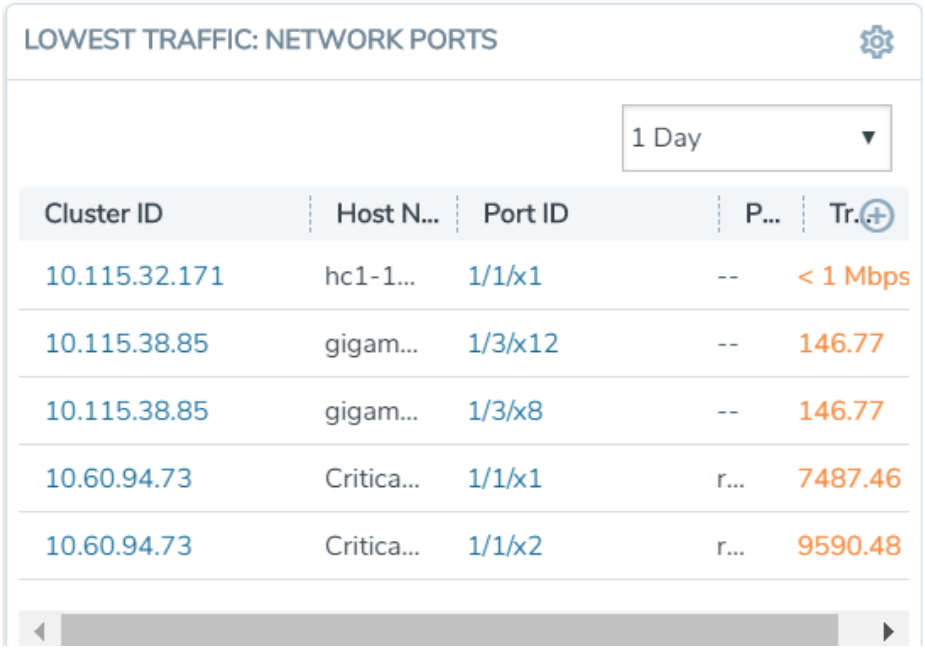


Figure 23 *Lowest Traffic*

The Lowest Traffic widget is configured exactly the same way as the Highest Traffic widget. To configure the Lowest Traffic widget, refer to the configuration steps provided in Highest Traffic . In **step 4**, select **Lowest Traffic** and click **OK**. The Lowest Traffic configuration window is displayed.

Traffic Comparison By Tags

The Traffic Comparison By Tags widget allows you to compare the aggregated traffic flowing through the list of ports associated to tags. You can choose to view up to four traffic comparisons in a single widget. You can create as many Traffic Comparison By Tags widgets as necessary in the selected profile and provide a customized name for each widget. The customized name helps you to differentiate multiple traffic comparison widgets in a single profile.

In this example, there is traffic flowing from GigaVUE-TA10 to GigaVUE-HC3. You can group the tool ports in GigaVUE-TA10 and create a tag as T A_TOOL. Then, you can group the network ports in GigaVUE-HC3 and create a tag as HC3-NETWORK. Refer to Figure 24Example for Traffic Comparison By Tags Widget.

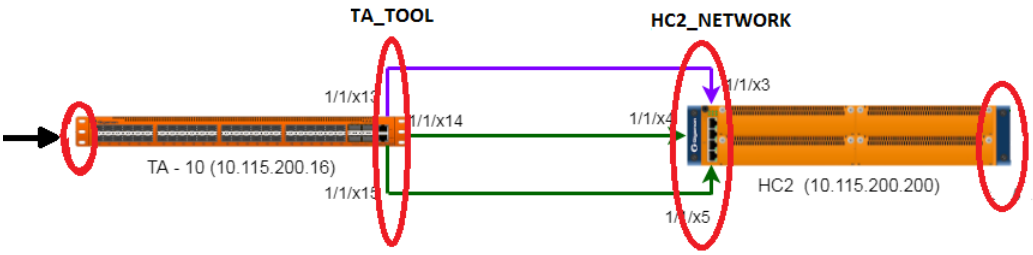


Figure 24 *Example for* Traffic Comparison By Tags Widget

Using the Traffic Comparison By Tags widget, you can compare the egress traffic passing through the ports associated with T A_TOOL with the ingress traffic passing through the ports associated with HC3-NETWORK, and quickly analyze if there is any packet loss associated. Refer to Figure 25Traffic Comparison By Tags

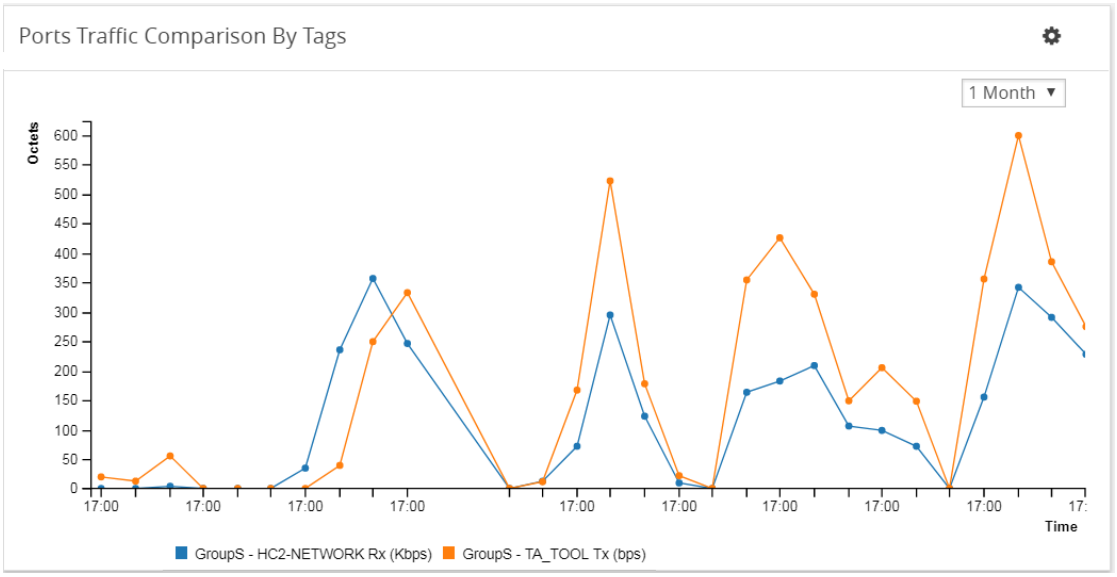


Figure 25 *Traffic Comparison By Tags*

The Traffic Comparison By Tags widget also allows you to choose just the egress traffic passing through the ports associated with T A_TOOL and view the graph.

The following statistics can be viewed for physical ports and GigaSMART:

Traffic Type	Statistics
Physical Ports	Data Rate Packet Rate Packet Errors Packet Discards Packet Drops

Traffic Type	Statistics
	Port Utilization
GigaSMART	Data Rate Packet Rate Packet Drops Packet Errors Packet Buffer Packet Terminated

The aggregated traffic comparison is displayed as a graph. You can choose to display the data over a day, an hour, a week, or a month. However, when you select a week or a month, the time period is not persisted. The data is defaulted to 1 day when you navigate away from the Physical Dashboards page and then return to the page. Click the notification icon at the top of the window and view the alarms and notifications displayed (refer to Figure 26Notifications):

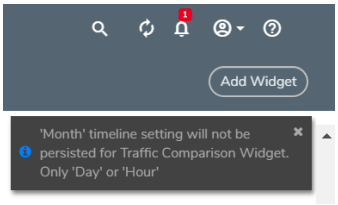


Figure 26 *Notifications*

Hovering the mouse over the lines in the graph displays the tag name, traffic direction, and traffic flow (Mbps).

To configure the Traffic Comparison By Tags widget:

1. On the left navigation pane, click on  and from the Physical & Virtual dashboard page, select the profile in which you want to add the widget.
2. Click **Add New Widget**. The Add New Widget window is displayed. Refer to Figure 21Add New Widget.

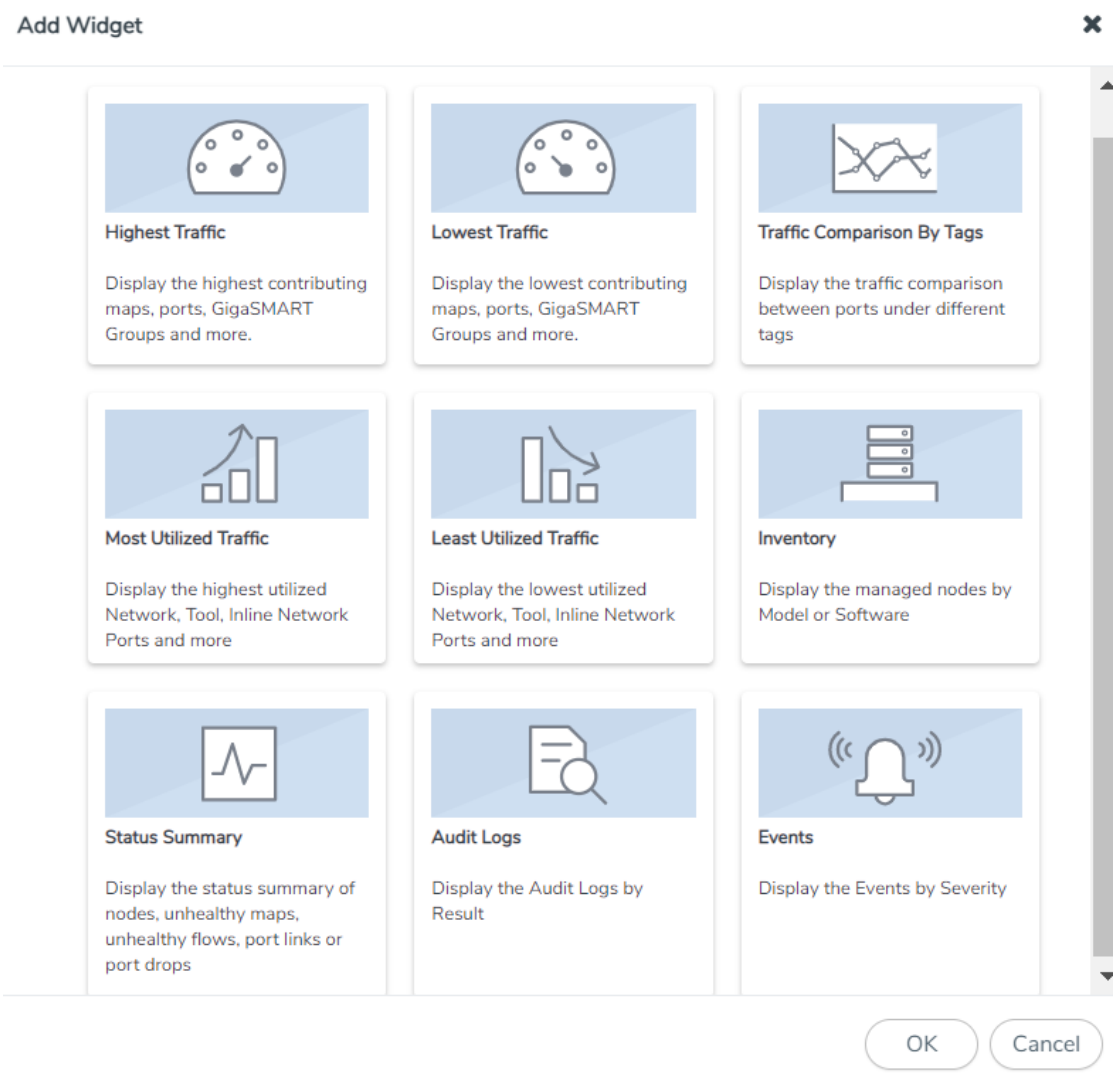


Figure 27 *Add New Widget*

3. In the Add New Widget window, select **Traffic Comparison By Tags** widget and click **OK**. The Traffic Comparison by Tags configuration window is displayed. Refer to Figure 28Traffic Comparison By Tags Configuration.

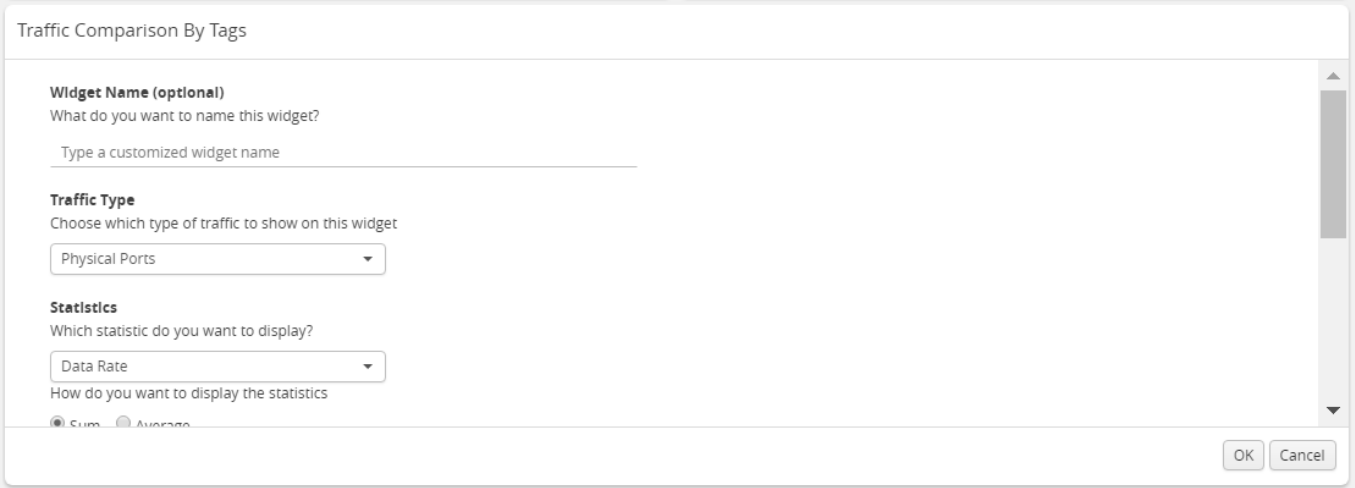


Figure 28 *Traffic Comparison By Tags Configuration*

4. (Optional) In the **Widget Name** box, enter a customized name for the widget. Customized name helps to differentiate multiple traffic comparison widgets in the same profile.
5. From the **Traffic Type** drop-down list, select one of the following traffic types:
   - Physical Ports
   - GigaSMART
6. From the **Statistics** drop-down list, select the type of statistic to view in the comparison graph.
7. Select **Sum** or **Average** to determine the way to display the statistics.
8. In Tag Items, select two or more tags to compare.
   a. For Traffic 1, select the tag name and tag value from the drop-down lists.
   b. Select **Ingress (Rx)** or **Egress (Tx)** to determine the traffic direction.
   c. Repeat step a and step b to select the next traffic for comparison.
9. Click **OK**.

Most Utilized Traffic

The Most Utilized Traffic widget allows you to view the ports with highest percentage utilization. The highest percentage utilization is displayed over the selected period. The period can be 1 hour, 1 day, 1 week, or 1 month to view the utilization percentage.

The Most Utilized Traffic widget lists the ports with the cluster ID, port Id, port alias, and the utilization percentage.

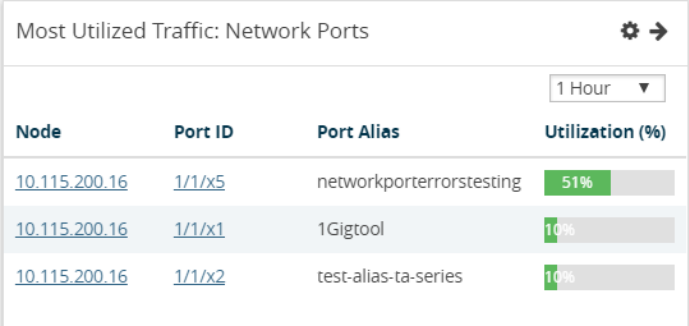


Figure 29 *Most Utilized Traffic Widget*

To configure the Most Utilized Traffic widget:

1. On the left navigation pane, click on  and from the Physical & Virtual dashboard page, select the profile in which you want to add the widget.
2. Click **Add New Widget**. The Add New Widget window is displayed. Refer to Figure 30Add New Widget.

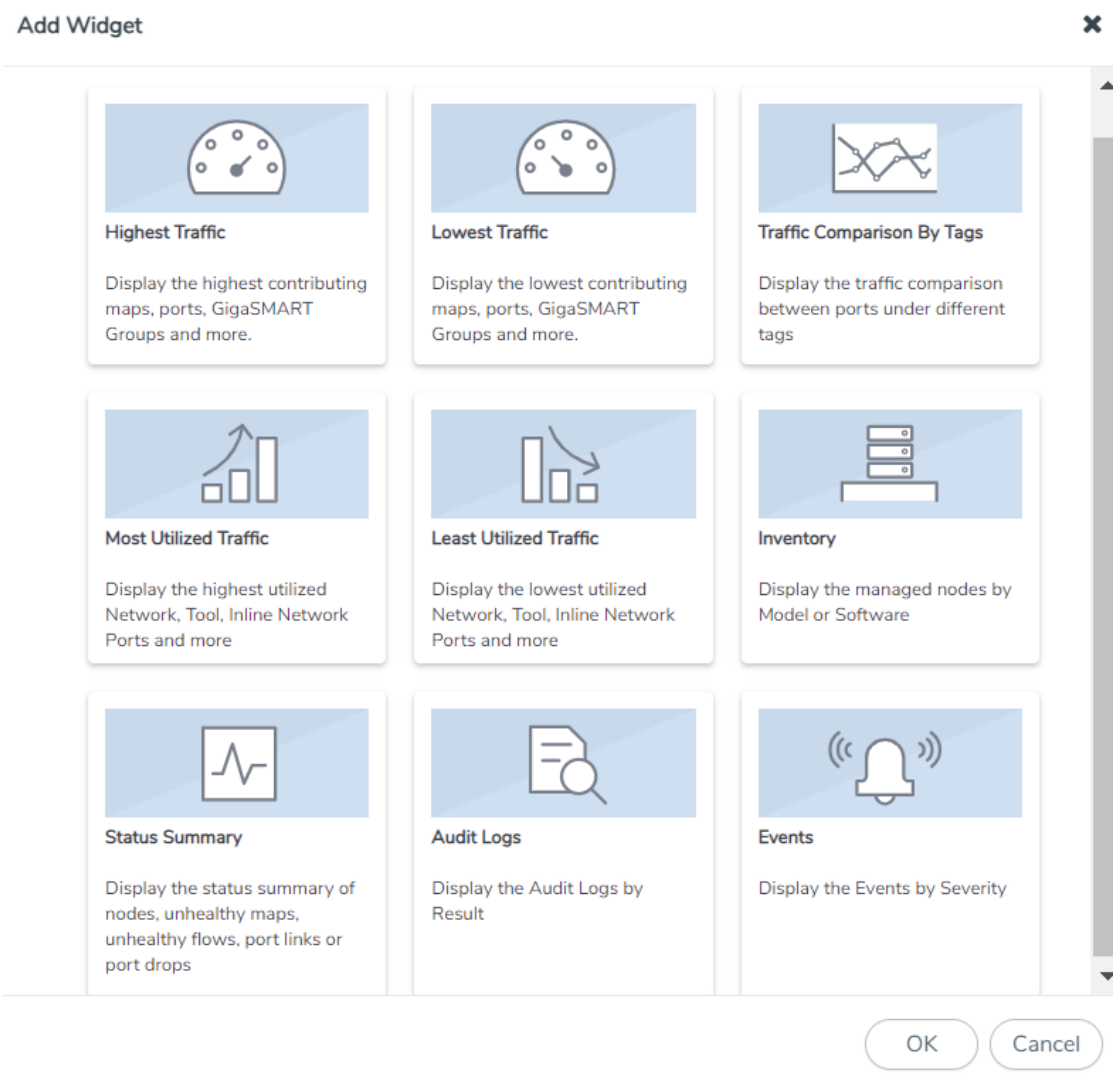


Figure 30 *Add New Widget*

3. In the Add New Widget window, select **Most Utilized Traffic** and click **OK**. The Most Utilized Traffic configuration window is displayed. Refer to Figure 31Most Utilized Traffic Configuration.

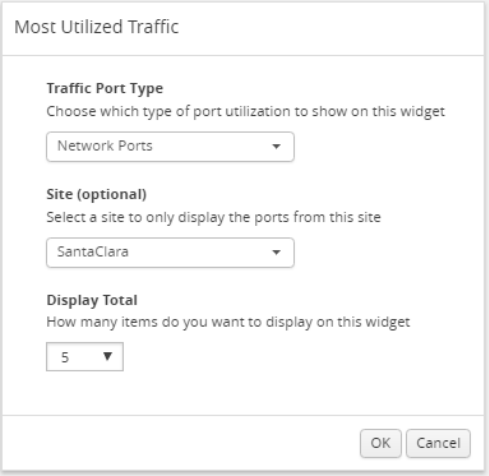


Figure 31 *Most Utilized Traffic Configuration*

4. From the **Traffic Port Type** drop-down list, select one of the following port types:
   - Network Ports
   - Tool Ports
   - Stack Ports
   - Hybrid Ports
   - Inline Network Ports
   - Inline Tool Ports
5. Select the required tag key and tag value combination (for example: tag key is 'Site' and tag value is 'East') for which the most utilized traffic configuration must be displayed. This step is optional.
6. From the **Display Total** drop-down list, select the number of items to be displayed. By default, the number of items selected for display is 5.
7. Click **OK**.

Least Utilized Traffic

The Least Utilized Traffic widget allows you to view the lowest percentage utilization for all the ports. The lowest percentage utilization is displayed over the selected period, You can choose 1 hour, 1 day, 1 week, or 1 month to view the utilization percentage.

The Least Utilized Traffic widget lists the ports with the cluster ID, host name, port number, port alias, and the utilization percentage.

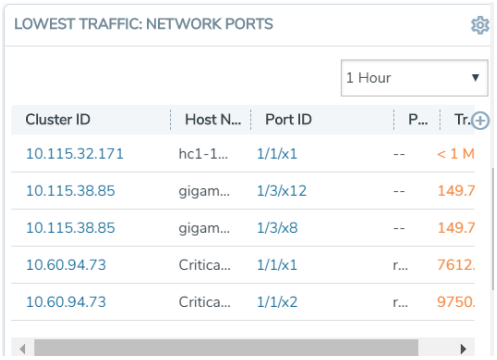


Figure 32 *Least Utilized Traffic*

The Least Utilized Traffic widget is configured exactly the same way as the Most Utilized Traffic widget. To configure the Least Utilized Traffic widget, refer to the configuration steps provided in Most Utilized Traffic. In **step 4**, select **Least Utilized Traffic** and click **OK**.

Inventory

The Inventory widget provides information about the physical nodes by model and software.

Nodes by Model

The Nodes by Model widget displays the number of nodes managed by the current instance of GigaVUE-FM as a bar graph. Each bar in the graph indicates the number of each device model managed. Hovering the mouse over a bar in the graph displays the model name and the total number. Figure 33Nodes by Model shows a Node by Model widget displaying six different nodes managed by GigaVUE-FM. Hover the mouse over the bar to view the number of devices in each node.

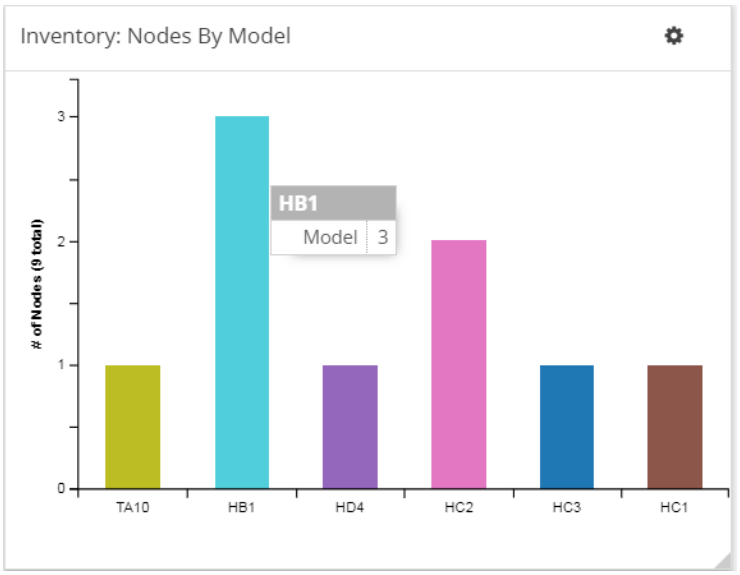


Figure 33 *Nodes by Model*

Nodes by Software Version

The Nodes by Software Version widget presents a graph that helps you to quickly view the software versions of the nodes that GigaVUE-FM is managing and the total percentage of each version. Each software version is assigned a color in the graph, which is specified by the legend. Hovering the mouse over an area in the graph displays the total number of software version used as a percentage. In Figure 34Nodes by Software Version, the Nodes by Software Version widget shows that there are 7 instances of version 5.1 and 2 instances of version 5.0, which is 22 percent of the total versions installed.

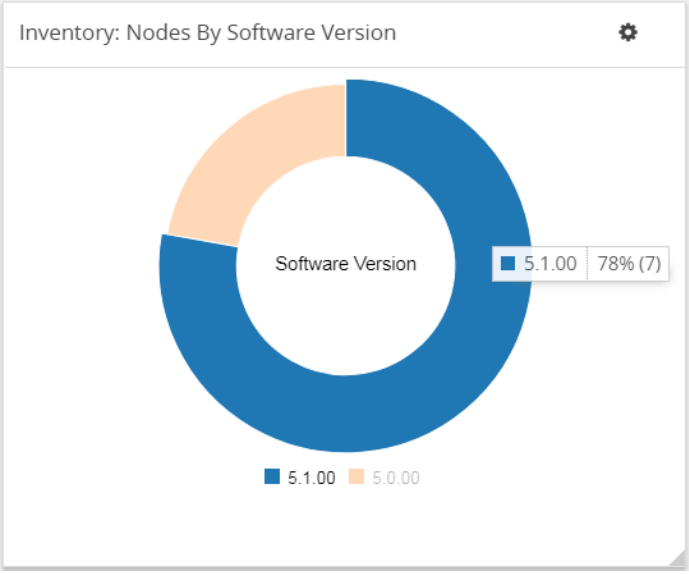


Figure 34 *Nodes by Software Version*

To configure the Inventory widget:

1. On the left navigation pane, click on  and from the Physical & Virtual dashboard page, select the profile in which you want to add the widget.
2. Click **Add New Widget**. The Add New Widget window is displayed. Refer to Figure 35Add New Widget.

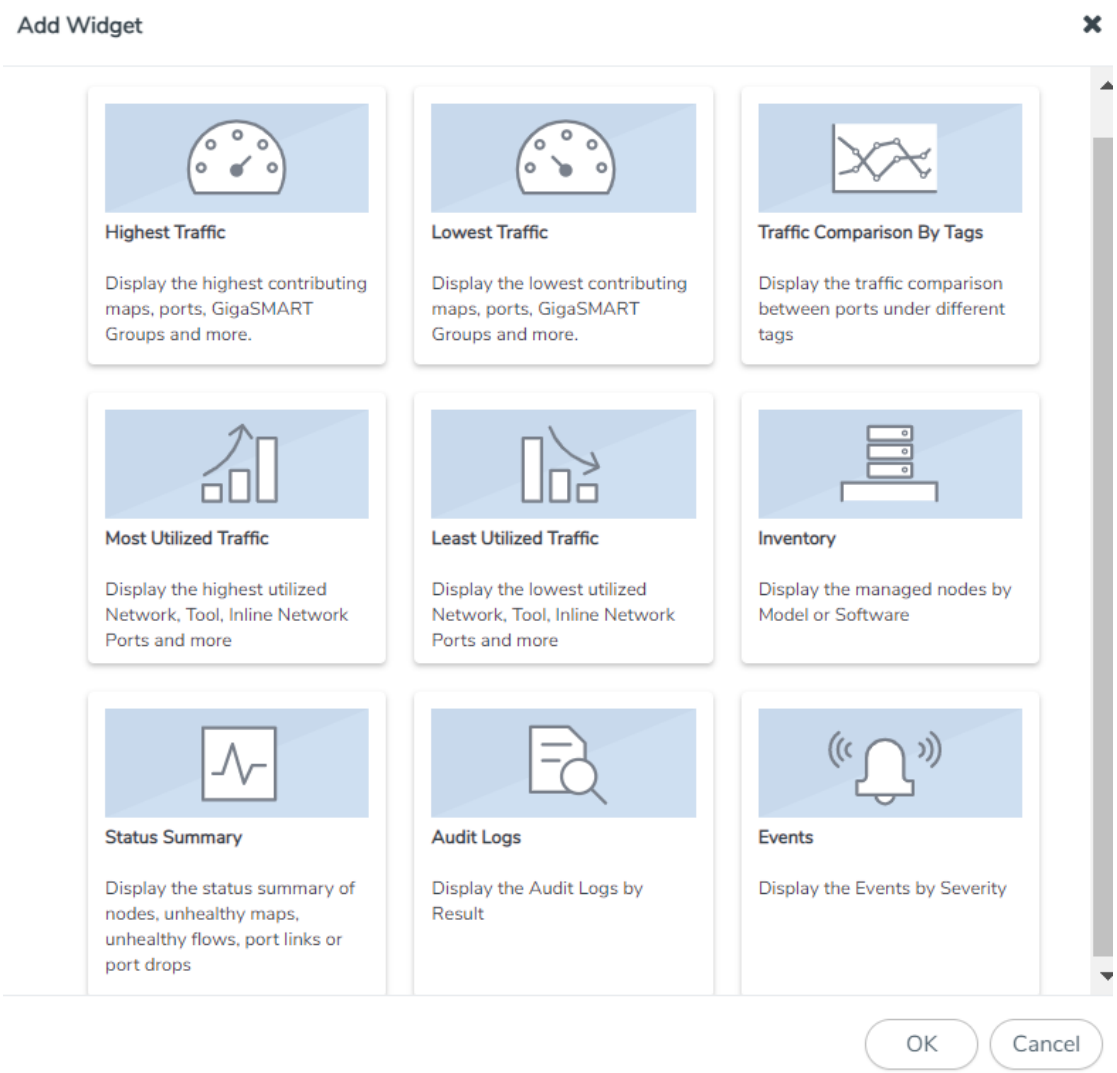


Figure 35 *Add New Widget*

3. In the Add New Widget window, select **Inventory** and click **OK**. The Inventory configuration window is displayed. Refer to Figure 31Most Utilized Traffic Configuration.

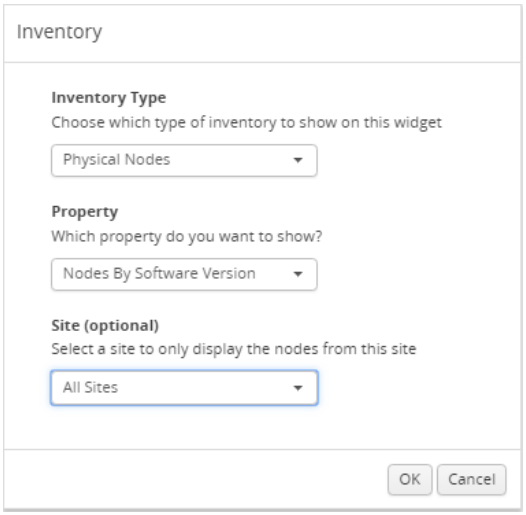


Figure 36 *Inventory Configuration*

4. From the **Inventory Type** drop-down list, select the Physical Nodes.
5. From the **Property** drop-down list, select one of the following:
   - Nodes by Model
   - Nodes by Software Version
6. Select the required tag key and tag value combination (for example: tag key is 'Site' and tag value is 'East') for which the inventory type details must be displayed. This step is optional.
7. Click **OK**.

Status Summary

Refer to the following section for the Status Summary widget details:

- Nodes' Status Summary
- Port Link Status Summary
- Unhealthy Maps
- Unhealthy Flows
- Port Drops and Errors
- Unhealthy Fabric Maps

Nodes' Status Summary

The nodes' status summary widget presents a graph that allows you to quickly view the current status of the physical nodes that GigaVUE-FM is managing and the number of nodes in a particular status, which is indicated by a color in the graph. The possible statuses

are:

- Normal (green)
- Warning (yellow)
- Error (orange)
- Critical (red)

For information about how the device status is computed, refer to Node Health Status.

Hovering the mouse over an area in the graph displays the percentage of nodes in that status. In Figure 37Nodes' Status Widget, the widget shows that there are 5 nodes in Normal status, 4 nodes in Warning status, and 5 nodes in Critical status. There are no nodes in Error status.

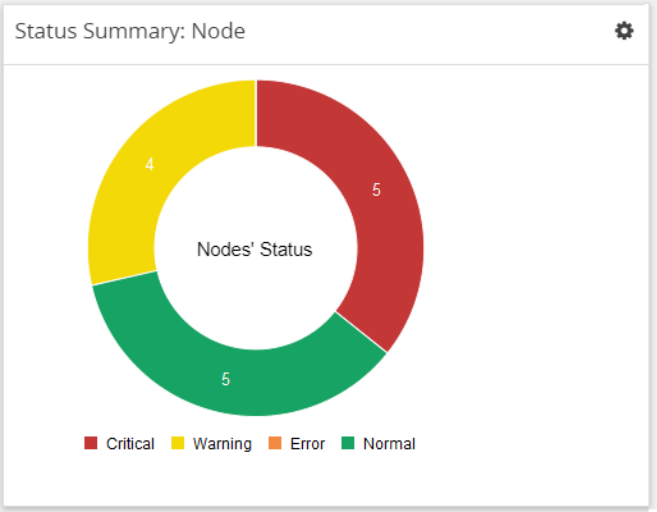


Figure 37 *Nodes' Status Widget*

Port Link Status Summary

The Port Link status summary widget allows you to view the current link status of all the ports available in the physical nodes currently managed by GigaVUE-FM. Optionally, the Port Link status can be displayed for the required combination of tag key and tag value (for example: tag key is 'Site' and tag value is 'East'). When a particular site tag is selected, the port link status of all the ports available in the nodes associated to that site are displayed in the Port Link Status Summary dashboard.

Refer to Figure 38Status Summary: Port Links Widget for Status Summary: Port Links dashboard.

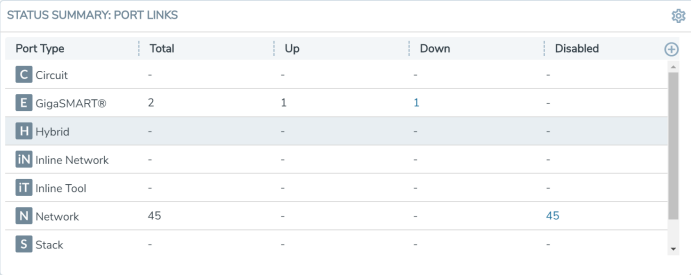


Figure 38 *Status Summary: Port Links Widget*

The Status Summary: Port Links widget lists the following:

- Ports type
- Total number of ports in each type
- Total number of ports in the up, down, or disabled state

Click the numbers in the down or disabled column. A quick view provides detailed information with the cluster ID, device host name, port ID, and port alias of all the ports in the down or disabled state. If the port status is down, the quick view also provides information about the time since when the port has been in down state. The down time is displayed in minutes, hours, days, or months.

Click the port ID link for a detailed view of the packet errors, packet drops, data rate transmitted or received, packet transmitted or received, and so on occurring on an hourly, daily, weekly, or monthly basis. You can also view the related maps, transceiver type, speed, and other detailed information about the port.

NOTE: All gateway ports on GigaVUE TA Series nodes are tool ports.

Unhealthy Maps

The Unhealthy Maps status summary widget lists the maps that are in unhealthy state. The health of a map is determined by the health status of its associated components such as ports, port groups, port pairs, GigaStream, tool port, GigaSMART group, tunneled port, virtual port, inline network, inline tool, inline tool group, inline serial tool group, inline network group, and GigaSMART operations. If the status of any one of the component is down, the corresponding map is also considered unhealthy.

The Unhealthy Maps widget shows the cluster ID, map alias, and current status of the unhealthy map. The possible statuses are:

- Critical (red)
- Warning (amber)
- Unknown (gray)

The health status of a map is shown as gray when the traffic health is still being computed. The status will be updated eventually.

Click on the ID to go directly to the node. Click on the map alias to display the quick view for the map. Hovering the mouse over the status bubble for the map displays the port or ports related to the map that is in an unhealthy state.

Unhealthy Flows

The Unhealthy Flows status summary widget lists the flows that are in unhealthy state. The health of a flow is determined by the health status of the pass-all maps and the priority maps involved in the flow.

A priority map group consists of one or more maps configured with the same source ports. The health of a priority map group is determined by the aggregated health of the constituted maps. The health of the maps is determined by its associated components such as ports, port groups, port pairs, GigaStream, and so on. If any one of the maps in the priority map group is unhealthy, the corresponding priority map group is also considered unhealthy. But, the overall health status of a flow is determined by the aggregated health of the maps that are involved in the flow.

The Unhealthy Flows widget shows the names of the flows that are in unhealthy state and the names of the maps that are unhealthy in the flow. Click the Flow Name to open the flow view page.

For more information about Flows, refer to Flows.

Port Drops and Errors

The Port Drops and Errors status summary widget helps in identifying the ports with packet drops or packet errors in the network. When a particular site value is selected, the status summary widget lists the port types associated to that site and the number of ports having packet drops, transmitting errors, or receiving errors in the site. You can view the number of ports with packet drops or packet errors occurring on a daily or an hourly basis.

NOTE: To view the unhealthy ports for GigaSMART, the GigaVUE-OS node must have Software version 5.0.

To view detailed information about the port drops and errors, click the number in the Pkt Drops, Rx Errors, or Tx Errors column. A quick view displays the cluster ID, host name, port ID, port alias, and the number of packet drops or errors for the list of unhealthy ports in the port

type. If there are too many ports, click the Filter icon and filter the ports based on the cluster ID, host name, port ID, or port alias. To clear the filters, click **Clear Filters** in the filter dialog box.

In the Unhealthy Ports quick view, click the port ID for a detailed view of the type of packet errors or packet drops occurring on a daily or an hourly basis. You can also view the related maps, transceiver type, speed, and other detailed information about the port, which helps to investigate the reason for the packet drops or packet errors. To return to the Ports quick view, click **Back**.

To configure the Inventory widget:

1. On the left navigation pane, click on  and from the Physical & Virtual dashboard page, select the profile in which you want to add the widget.
2. Click **Add New Widget**. The Add New Widget window is displayed. Refer to Figure 21Add New Widget.

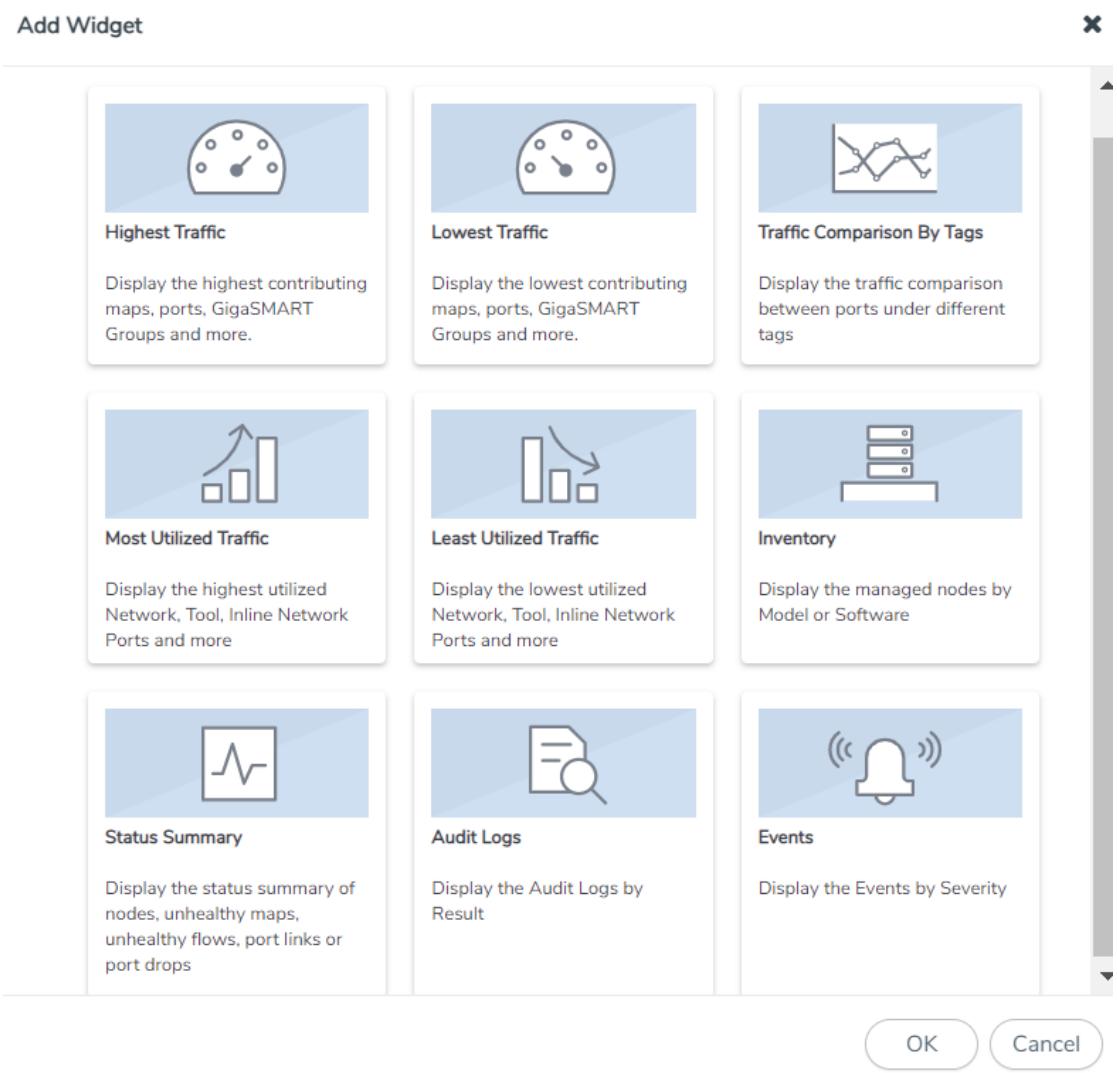


Figure 39 *Add New Widget*

3. In the Add New Widget window, select **Status Summary** and click **OK**. The Status Summary configuration window is displayed.

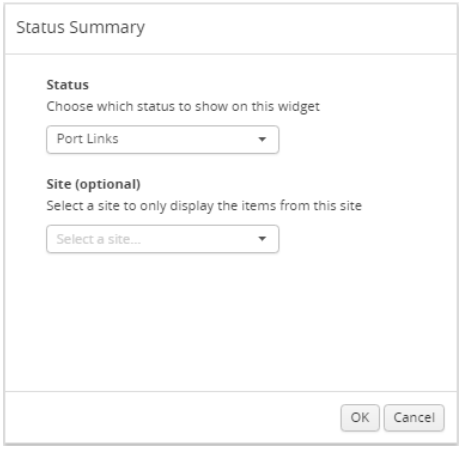


Figure 40 *Status Summary Configuration*

4. From the **Status** drop-down list, select one of the following:
    - Node—For information, refer to Nodes' Status Summary.
    - Port Links—For information, refer to Port Link Status Summary.
    - Unhealthy Maps—For information, refer to Unhealthy Maps.
    - Port Drops & Errors—For information, refer to Port Drops and Errors.
5. Select the required tag key and tag value combination for which the highest traffic distribution must be displayed. This step is optional.
6. Click **OK**.

Unhealthy Fabric Maps

The unhealthy flow maps widget displays the list of unhealthy flow maps.

The Unhealthy Fabric Maps status summary widget lists the maps that are in unhealthy state. The health of a fabric map is determined by the health status of its associated components such as maps, ports, port groups, port pairs, GigaStream, tool port, GigaSMART group, virtual port and GigaSMART operations. If the status of any one of the component is down, the corresponding fabric map is also considered unhealthy.

The Unhealthy Maps widget shows the fabric map alias and the current status of the unhealthy fabric map. The possible statuses are:

- Critical (red)
- Warning (amber)
- Unknown (gray): The health status of a map is shown as gray when the traffic health is still being computed. The status will be updated eventually.

Audit Logs

The Audit Logs widget shows the audit logs of successful and failed events. Optionally, the audit logs can be displayed for a specified site. When a particular site tag is selected, the audit logs pertaining to the clusters and nodes associated to that site are displayed in the Audit Logs dashboard.

The Audit Logs widget presents a graph that allows you to quickly view the number of logs in successful or failure state. In the graph, the state is indicated by color. The possible log results are:

- Success (green)
- Failure (red)

Hovering the mouse over an area in the graph displays the percentage of audit logs in that result. You can also specify the audit log statuses that have occurred over the past hour, day, week, or month. Figure 41Audit Logs by Result shows the audit log results over each of the time periods.

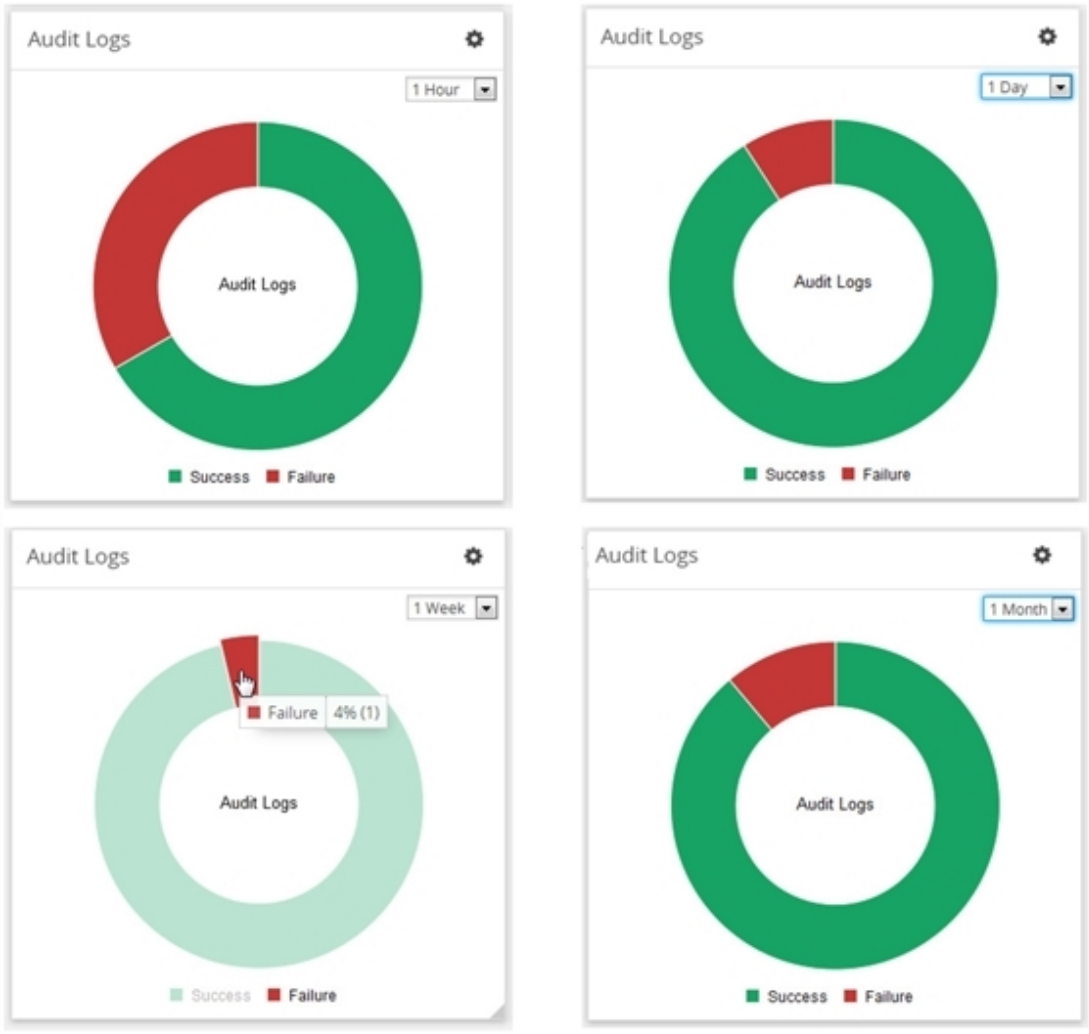


Figure 41 *Audit Logs by Result*

In this example, the Audit Logs widget shows that there is a log with the failure status that has occurred in the last hour. When you go to the audit logs page, you can see the entries, which matches with the information displayed in Audit Logs widget: two successes and one failure due to an incorrect log in.

To configure the Audit Logs widget:

1. On the left navigation pane, click on  and form the Physical & Virtual dashboard page, select the profile in which you want to add the widget.
2. Click **Add New Widget**. The Add New Widget window is displayed. Refer to Figure 21Add New Widget.

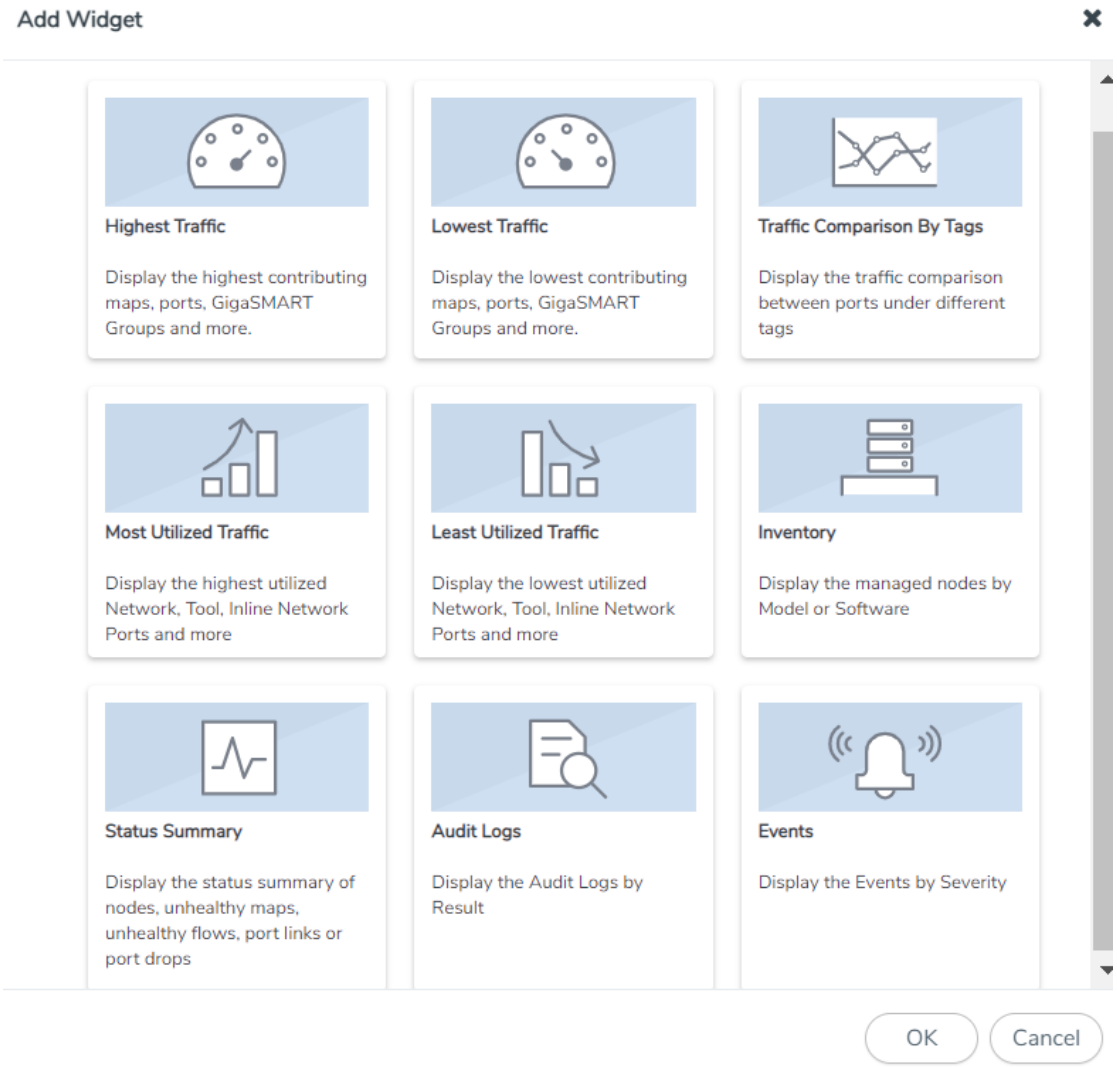


Figure 42 *Add New Widget*

3. In the Add New Widget window, select **Audit Logs** and click **OK**. The Audit Logs configuration window is displayed. Refer to Figure 31Most Utilized Traffic Configuration.

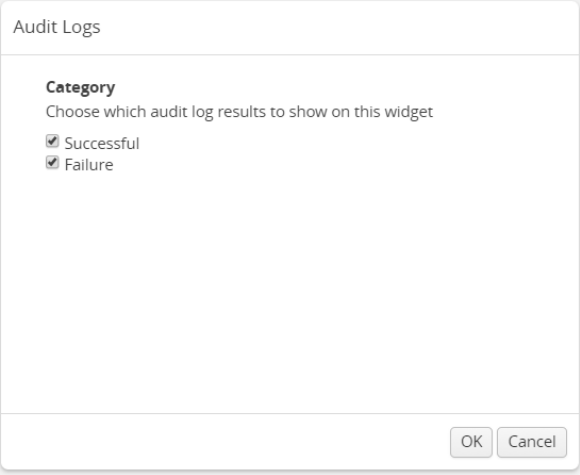


Figure 43 *Status Summary Configuration*

4. Choose the category of audit log results you want to view:
   - Successful
   - Failure
5. Click **OK**.

Events

The Events widget presents a graph that shows the number of events that have occurred within a particular severity level, which is indicated by a color in the graph. Optionally, the events can be displayed for a specified site. When a particular site tag is selected, only the events pertaining to the clusters and standalone nodes associated to that site are displayed in the Events dashboard.

The possible severity levels are:

- Information (blue)
- Major (orange)
- Minor (yellow)
- Critical (red)

Hovering the mouse over an area in the graph displays the percentage and the number of events that have occurred within the selected severity level. You can also select the time period to view the number of events that have occurred over the past hour, day, week, or month.

Figure 44Events Widget shows the number of events that have occurred in the past week in each severity level for all sites. If you want more detail about the events, select **Events** in the Physical page.

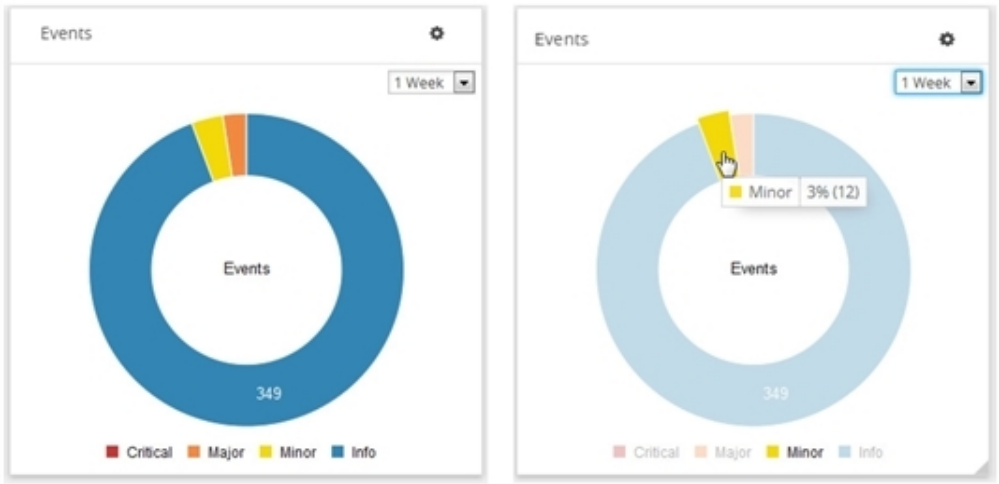


Figure 44 *Events Widget*

Figure 44Events Widget shows the number of events that have occurred in the past week in each severity level for the tag Santa Clara site.

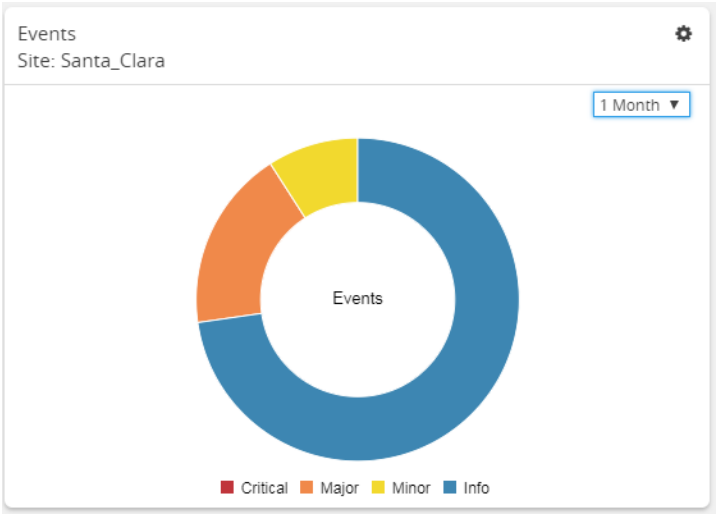


Figure 45 *Events Widget for Santa Clara*

To configure the Events widget:

1. On the left navigation pane, click on  and from t he Physical & Virtual dashboard page, select the profile in which you want to add the widget.
2. Click **Add New Widget**. The Add New Widget window is displayed. Refer to Figure 21Add New Widget.

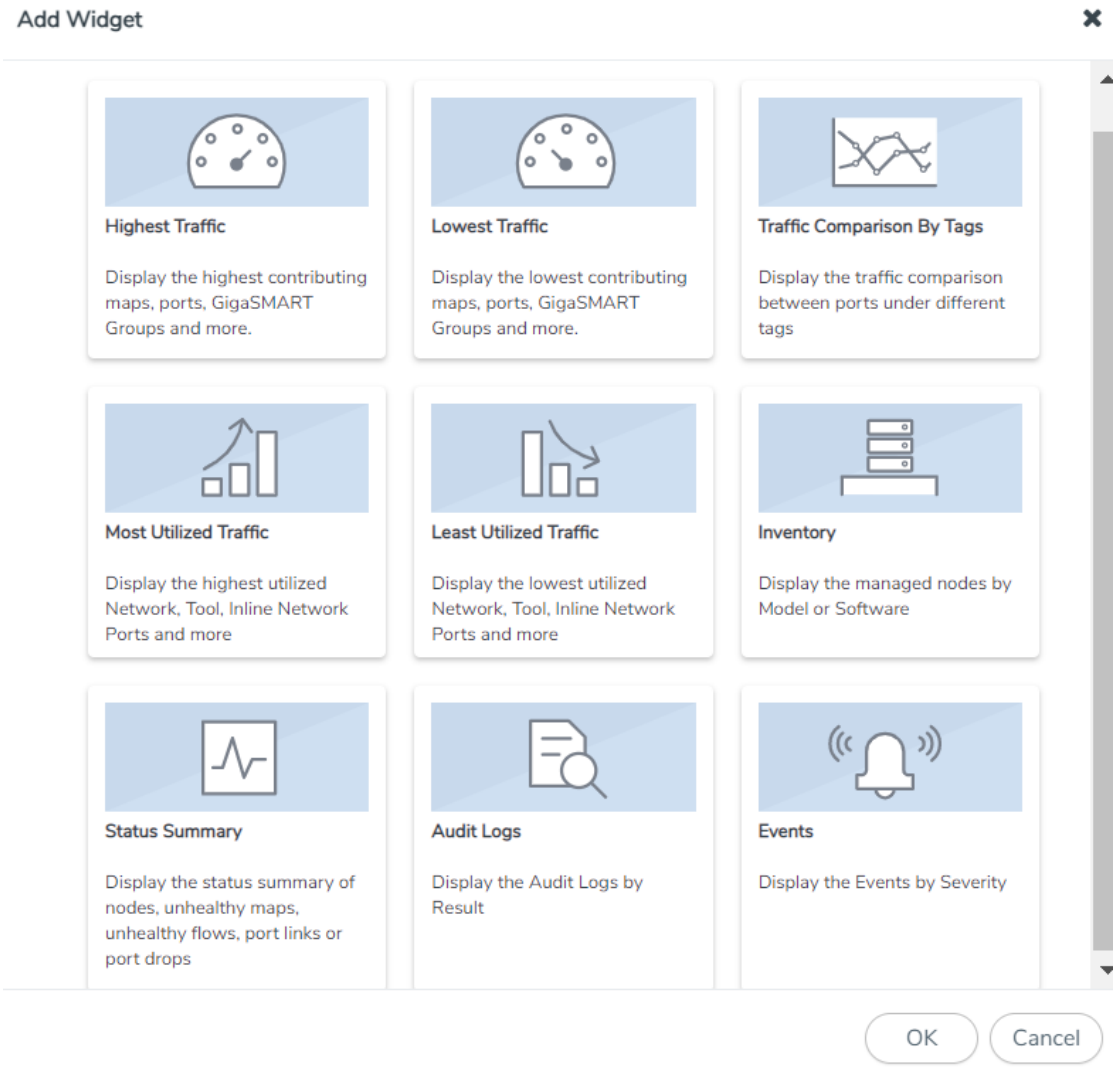


Figure 46 *Add New Widget*

3. In the Add New Widget window, select **Events** and click **OK**. The Events configuration window is displayed. Refer to Figure 31Most Utilized Traffic Configuration.

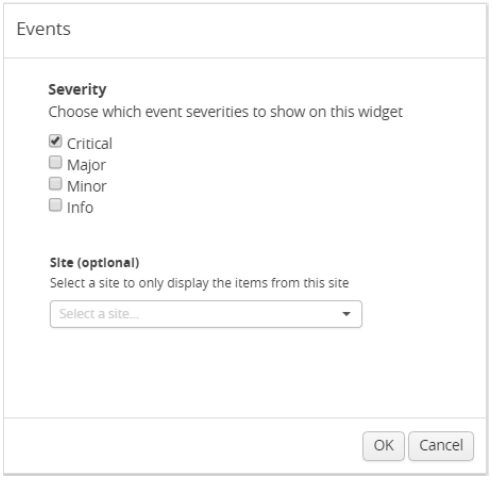


Figure 47 *Status Summary Configuration*

4. Choose the event you want to view in the widget:
    - Critical
    - Major
    - Minor
    - Info
5. Select the required tag key and tag value combination for which the events must be displayed. This step is optional.
6. Click **OK**.

FM Health Dashboard

This chapter describes the Health Monitor Dashboard of GigaVUE-FM.

This chapter covers the following topics:

- Overview of GigaVUE-FM Health Dashboard
- Alarm Thresholds and Notifications

You can access the Health Monitor Dashboard and view the current system performances such as CPU, Memory and Disk Usage without being authenticated in to GigaVUE-FM. Type *<fmip>/fmHealth* in your browser to view the Health Monitor Dashboard.

Overview of GigaVUE-FM Health Dashboard

GigaVUE-FM is the central management appliance for the visibility fabric. Therefore, knowing its current health is important in order to maximize the availability of the appliance. The Health Monitor dashboard provides health information about GigaVUE-FM and makes it possible to do the following:

- Detect problems with GigaVUE-FM so that they can be responded to in a timely fashion.
- Provide alerts about issues that could impact the performance, such as CPU or disk over-utilization.

The Health Monitor provides the following monitors:

- CPU utilization
- Memory utilization
- Disk utilization

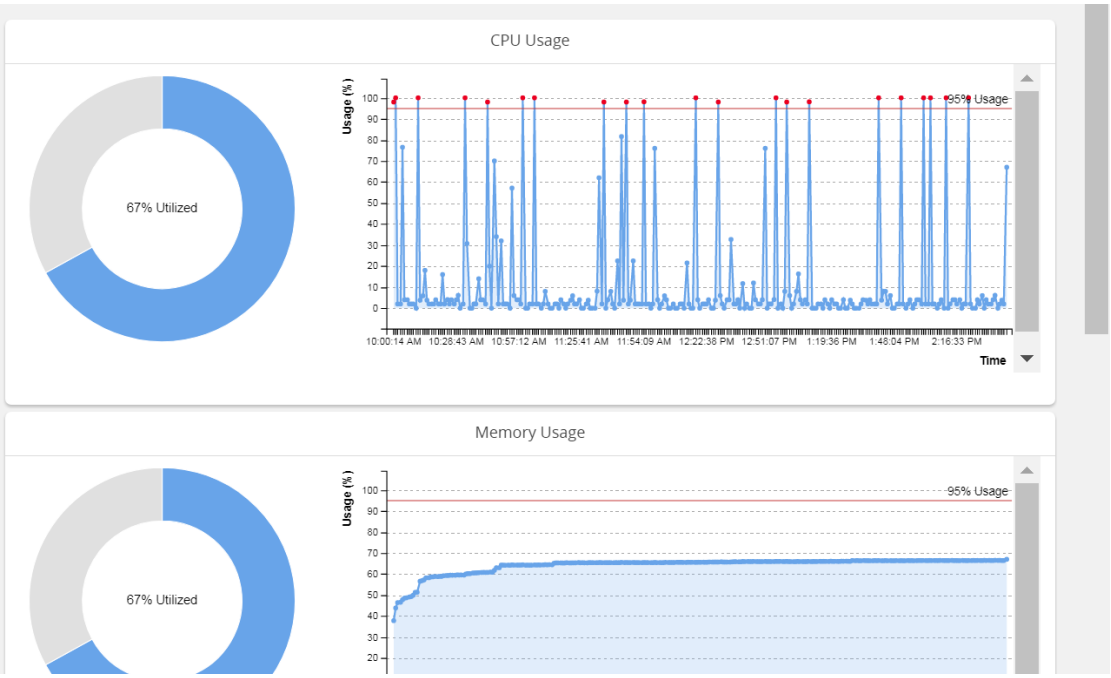


Figure 48 *Health Monitor Dashboards*

Note: If the percentage displayed in a pie-chart is negligible or less, then it would be difficult to click on the pie-chart arc and view the details.

CPU Utilization

The CPU Utilization Monitor displays overall CPU usage over time, providing information about peak CPU usage. This indicates whether there is sufficient CPU processing power for the currently deployed GigaVUE-FM appliance deployment. For example, peak CPU usage above a high-utilization mark of 90 for a long period for more than 30 minutes could indicate that the CPU power of the server is not adequate for GigaVUE-FM to manage the size of the deployed visibility fabric.

The CPU Utilization Monitor displays utilization as donut and time charts. The donut chart shows the percentage of utilized and available CPU. The time chart shows utilization at specific intervals. By clicking on a point in the time chart, you can see the utilization at a specific point in time. In Figure 49CPU Utilization Monitor, the CPU utilization at 10:49:55 AM is 1.0 percent.

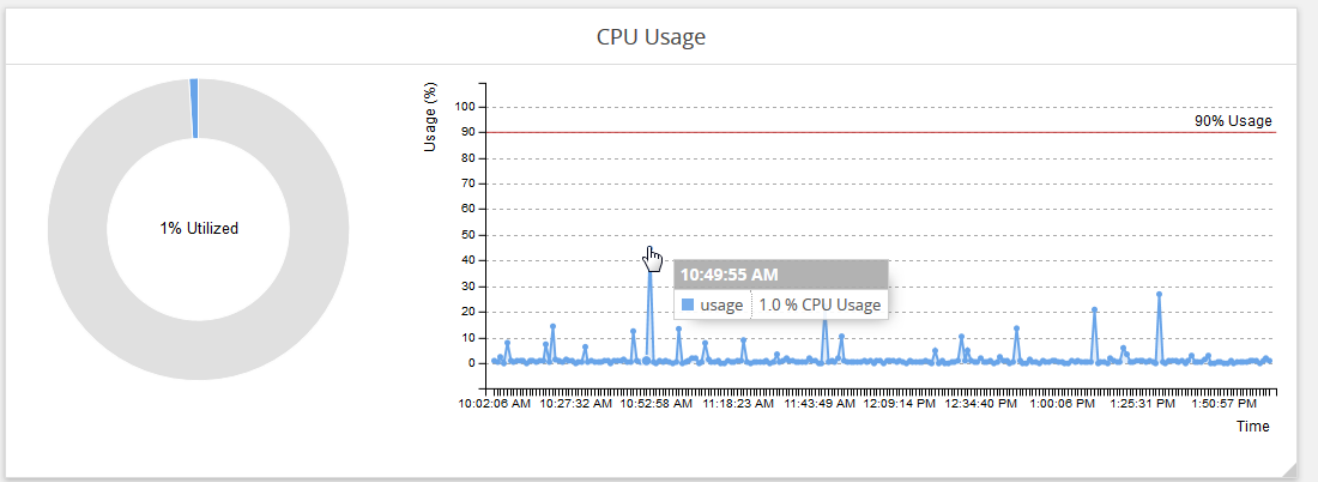


Figure 49 *CPU Utilization Monitor*

Memory Utilization

The Memory Utilization Monitor displays overall memory usage over time, providing information about peak memory usage. This indicates whether there is sufficient memory to handle the size of the visibility fabric managed by GigaVUE-FM. For example, memory usage above a high-utilization mark over a period for more than 30 minutes could indicate that the amount of memory supplied to GigaVUE-FM is insufficient.

The Memory Utilization Monitor displays utilization as dough nut and time charts. The dough nut chart shows the percentage of utilized and available memory. The time chart

shows utilization as specific intervals. By clicking on a point in the time chart, you can see the utilization at a specific point in time. In the following figure, the memory utilization at 11:48:49 AM is 17 percent.

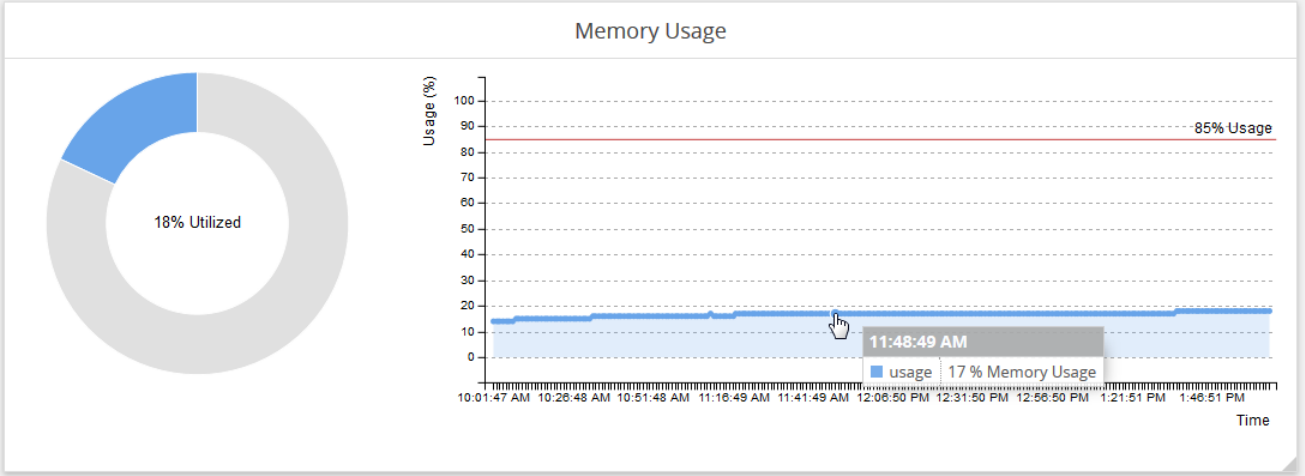


Figure 50 *Memory Utilization Monitor*

NOTE: When GigaVUE-FM starts up, the Memory Utilization Monitor displays around 60-65%. The utilization slowly increases up to 85% with the back end operations running and the memory gets stabilized at this point. A spike in memory is also observed when syslogs are more than usual.

Storage Utilization

The Storage Utilization Monitor displays disk usage levels over time for individual partitions, providing information about peak disk usage for GigaVUE-FM logs. This provides information that can help prevent outages due to disk out-of-space issues.

The Storage Utilization Monitor displays utilization for GigaVUE-FM logs. The bar charts show the percentage of disk utilization in the partitions for GigaVUE-FM logs. The time chart shows utilization as specific intervals for both partition. By clicking on a point in the time chart, you can see the utilization at a specific point in time. In Memory Utilization, the disk usage at 11:16:01 AM for GigaVUE-FM logs is 18.5 percent and the disk usage for GigaVUE-FM data is 12%.

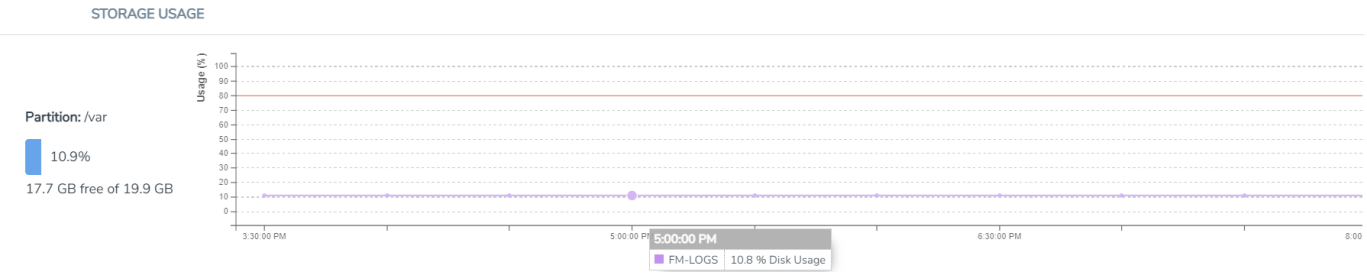


Figure 51 *Storage Utilization Monitor*

Alarm Thresholds and Notifications

For CPU, memory, and disk utilization monitoring, an alarm is triggered in the Alarms page if the following threshold levels are exceeded:

- CPU Utilization - 80%
- Memory Utilization- 90%
- File System (/var) - 80%
- File System (/config) - 80%

GigaVUE-FM Reference Materials

This section provides additional information useful for GigaVUE-FM.

Topics:

- Disk Size on GigaVUE-FM
- Data Transfer Rate Units
- Open Ports in GigaVUE-FM
- Open Ports in HC Series Devices
- Health Status
- GigaVUE-FM APIs

Disk Size on GigaVUE-FM

This section describes how to increase the data storage space available on GigaVUE-FM. It also explains how to clear the space in the /var directory.

- After installing GigaVUE-FM, the size of /config can be increased by increasing the size of the disk used for /config and then rebooting GigaVUE-FM.
- Increase Disk Size on a New or Existing GigaVUE-FM Installation on KVM describes increasing the disk size for a new or existing GigaVUE-FM installation as well as for an upgrade from previous releases.

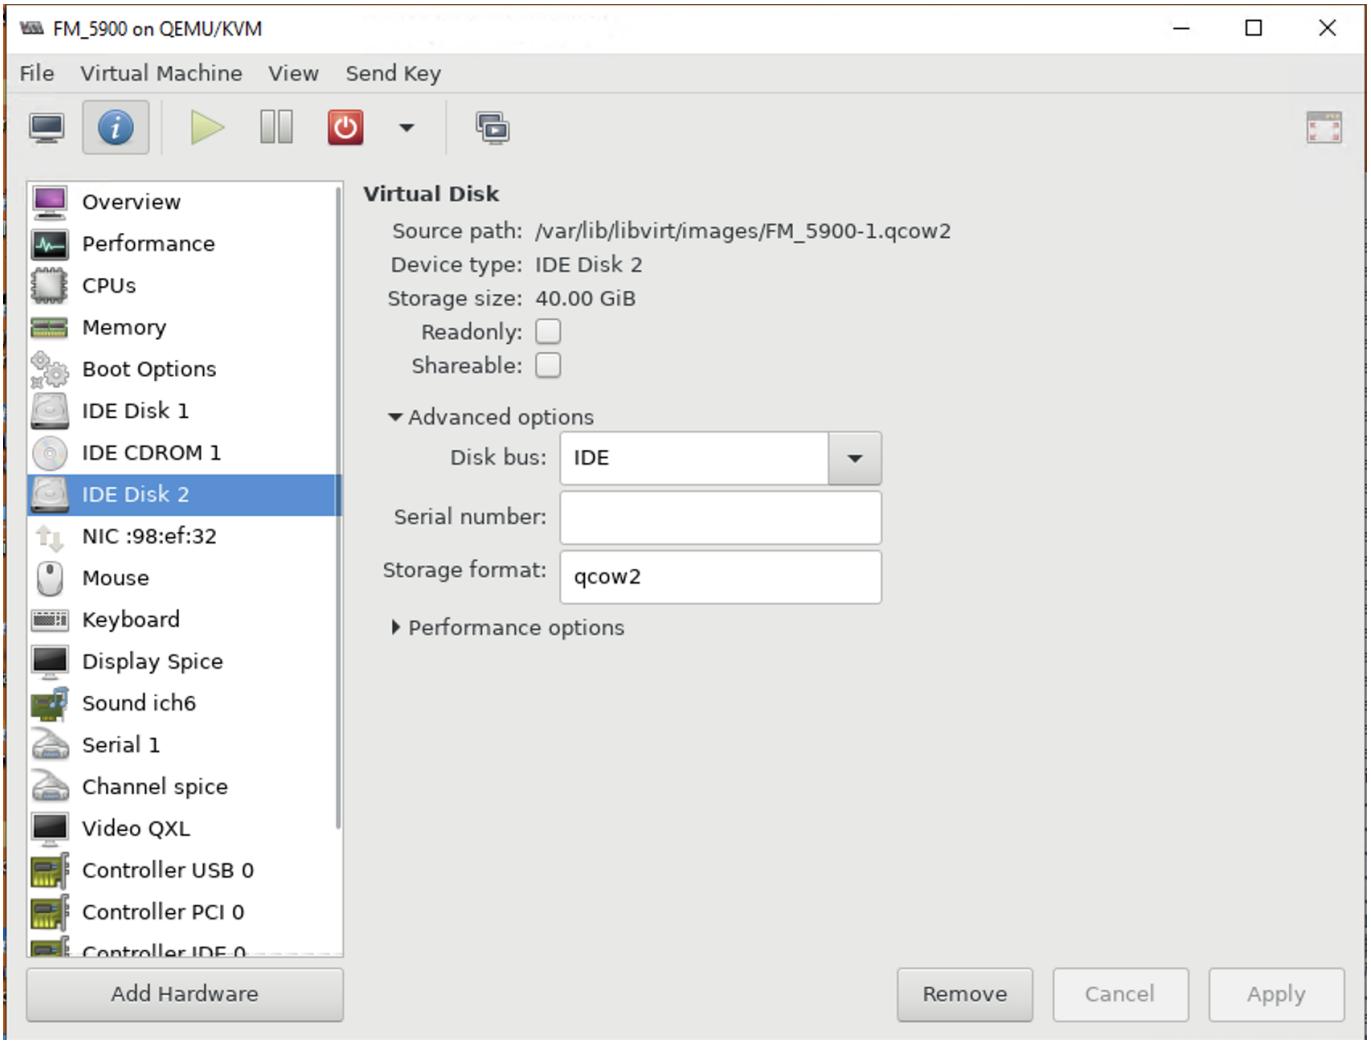
  Note: This procedure only applies to installations on KVM. Ensure that you apply this procedure only to GigaVUE-FM version 5.8 and later.

Increase Disk Size on a New or Existing GigaVUE-FM Installation on KVM

Note: This procedure only applies to installations on KVM. Ensure that you apply this procedure only to GigaVUE-FM version 5.8 and later.

To increase disk size on a new or existing GigaVUE-FM installation, do the following:

1. Shutdown the GigaVUE-FM system.
2. Open the Virtual Machine Manager, and then go to **IDE Disk 2**.



3. Open the console, and run the following command to increase the disk size:

   /var/lib/libvirt/images# qemu-img resize <instance _name>-1.qcow2 +20G

   The disk size is increased.

4. Start-up the GigaVUE-FM system.
5. Login to the GigaVUE-FM system using CLI, and then run the **df -h** command to verify the disk size.

```
[admin@unconfigured-gigavue-fm ~]$ df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        7.9G     0  7.9G   0% /dev
tmpfs           7.9G     0  7.9G   0% /dev/shm
tmpfs           7.9G  680K  7.9G   1% /run
tmpfs           7.9G     0  7.9G   0% /sys/fs/cgroup
/dev/sda6       9.5G  3.1G  6.5G  32% /
/dev/sda5        19G  2.2G   17G  12% /var
/dev/sdb         60G  4.6G   56G   8% /config
/dev/sda1       677M   41M  637M   6% /boot
tmpfs           1.6G     0  1.6G   0% /run/user/1000
[admin@unconfigured-gigavue-fm ~]$
```

How to Clean up Disk Space on a GigaVUE-FM Instance

The /var directory can sometimes run out of storage space and result in performance degradation of GigaVUE-FM. The best practice is to periodically check the storage space (refer to the "*Storage Management*" section in the *GigaVUE Administration Guide*) and clear up the disk space to avoid system misbehavior.

The disk space on a GigaVUE-FM Instance can be cleared up in multiple ways:

- Purging the statistics older than a certain date. For more information, refer to the Storage Management section in the *GigaVUE Administration Guide*.
- Removing the unused internal images. For more information, refer to the Internal Image Files section in the *GigaVUE Administration Guide*.
- Deleting the logs. For more information, refer to the Delete a Log File section in the *GigaVUE Administration Guide*.

Data Transfer Rate Units

Data Transfer Rate is measured as multiples of unit bits per second (bit/s) or as bytes per second (B/s). The transfer rates shown on the dashboard and in other places are measured as decimal multiples of bits. The following table shows the units of data transfer:

Table 1: Decimal Multiple of Data Transfer Rate in Bits

Data Rate	Symbol	Rate
Kilobit per second	kbps	1000 bits per second
Megabit per second	Mbps	1,000,000 bits per second
Gigabit per second	Gbps	1,000,000,000 bits per second

NOTE: 8 kilobits = 1 kilobyte

Table 2: Decimal Multiple of Data Transfer Rate in Bytes

Data Rate	Symbol	Rate
Kilobyte per second	kBps	8000 bits per second
Megabyte per second	MBps	8,000,000 bits per second
Gigabyte per second	GBps	8,000,000,000 bits per second
Terabyte per second	TBps	8,000,000,000,000 bits per second

Open Ports in GigaVUE-FM

GigaVUE-FM Open Ports

The following table provides information about the ports:

Inbound

Protocol	Port Number	Service	Source CIDR	Purpose
TCP	22	SSH	Administrator Subnet	Allows CLI access to user-initiated management and diagnostics.
TCP	443	HTTPS	Administrator Subnet /	Allows GigaVUE-FM to accept Management connection using REST API from user and Gigamon devices.

Protocol	Port Number	Service	Source CIDR	Purpose
			GigaVUE-OS / Cloud Fabrics	Allows users to access GigaVUE-FM UI securely through an HTTPS connection.
TCP	514	Syslog	GigaVUE-OS Node	Allows GigaVUE-OS node to send syslog message to GigaVUE-FM over TCP.
UDP	514	Syslog	GigaVUE-OS Node	Allows GigaVUE-OS node to send syslog message to GigaVUE-FM over UDP.
UDP	162	SNMP	GigaVUE-OS Node	Allows GigaVUE-OS node to send SNMP events to GigaVUE-FM over UDP.
UDP	2056	FluentD	GigaVUE-OS Node / GigaVUE V Series	Allows GigaVUE-OS / GigaVUE V Series nodes to send Application Intelligence monitoring reports.
UDP	2096	FluentD	GigaVUE-OS Node	Allows GigaVUE-OS node to send Basic Inline SSL Session Stats.
UDP	2097	FluentD	GigaVUE-OS Node	Allows GigaVUE-OS node to send Advanced Inline SSL Session Stats.
TCP	5671	RabbitMq	Cloud Fabric Nodes	Allows Cloud Fabric nodes to send health events, solution status, statistics, and other notifications.
TCP	9600	StepCA	Cloud Fabric Nodes	Allows Cloud Fabric nodes to configure and renew the certificates.

Note: For FMHA, it is essential to open all the previously mentioned ports from GigaVUE-OS to every GigaVUE-FM node.

Outbound

Protocol	Port Number	Service	Source CIDR	Purpose
TCP	389	LDAP	GigaVUE-FM	Allows GigaVUE-FM to reach the LDAP server for authentication. Required only if the LDAP is configured for FM user authentication.
TCP	636	LDAP SSL	GigaVUE-FM	Allows GigaVUE-FM to reach the LDAP server over SSL for authentication. Required only if the LDAP is configured for FM user authentication.
UDP	1812	RADIUS (RFC 2865)	GigaVUE-FM	Allows GigaVUE-FM to reach the RADIUS server for authentication. Required only if the RADIUS is configured for FM user authentication.
TCP	49	TACACS	GigaVUE-FM	Allows GigaVUE-FM to reach the TACACS server for authentication. Required only if the TACACS is configured for FM user authentication and TACACS uses TCP..

Protocol	Port Number	Service	Source CIDR	Purpose
UDP	53	DNS	GigaVUE-FM	Allows GigaVUE-FM to reach the DNS server for name resolution.
UDP	68	DHCP	GigaVUE-FM	Allows GigaVUE-FM to reach the DHCP server for network configuration.
UDP	123	NTP	GigaVUE-FM	Allows GigaVUE-FM to reach the NTP server for time synchronization.

Open Ports for Communication between members of GigaVUE-FM High Availability Cluster

The following table lists the ports that must be open for communication between the members of GigaVUE-FM High Availability cluster:

NOTE: These ports cannot be accessed by standalone GigaVUE-FM instances.

Bidirectional

Protocol	Port Number	Service	Source CIDR	Purpose
TCP	443	HTTPS	GigaVUE-FM	REST API communication between HA members.
TCP	8300	Consul	GigaVUE-FM	RPC communication between Consul members.
TCP	8301	Consul	GigaVUE-FM	Heartbeat and Gossip between Consul members.
UDP	8301	Consul	GigaVUE-FM	Heartbeat and Gossip between Consul members
TCP	8302	Consul	GigaVUE-FM	Heartbeat and Gossip between Consul members over WAN.
UDP	8302	Consul	GigaVUE-FM	Heartbeat and Gossip between Consul members over WAN.
TCP	27017	MongoDB	GigaVUE-FM	Used for data replication across HA members and data access through GigaVUE-FMCLI.

Protocol	Port Number	Service	Source CIDR	Purpose
TCP	9300	OpenSearch	GigaVUE-FM	Used for data replication across cluster members.
TCP	30865	CSync2	GigaVUE-FM	Used for the Synchronization of files / directories across HA members. For example, Image files during GigaVUE-FM HA Upgrade.
TCP	24224	FluentD	GigaVUE-FM	Used for receiving / forwarding the packets from / to other HA members.
UDP	24224	FluentD	GigaVUE-FM	Used for receiving / forwarding the packets from / to other HA members.
UDP	4500	IPSec Tunnel	GigaVUE-FM	Used for encrypted communication between HA members.
	500			
Protocol 50 and Protocol 51				

Open Ports in HC Series Devices

Open Ports in Embedded Devices

Direction	Protocol	Port Number	Service	Description
Bidirectional	TCP	22	SSH	Used to access the device.
Bidirectional	TCP	80	HTTP	Used for any HTTP connection.
Bidirectional	TCP	443	HTTPS	Used for any HTTPS connection.
Bidirectional	UDP	111	NFS	Used for Internal Communication with gigasmart cards.
Bidirectional	UDP	662	NFS	Used for Internal Communication with gigasmart cards.
Bidirectional	UDP	890	NFS	Used for Internal Communication with gigasmart cards.
Bidirectional	UDP	892	NFS	Used for Internal Communication with gigasmart cards.

Direction	Protocol	Port Number	Service	Description
Bidirectional	DP	2020	NFS	Used for Internal Communication with gigasmart cards.
Bidirectional	UDP	2049	NFS	Used for Internal Communication with gigasmart cards.
Bidirectional	UDP	4045	NFS	Used for Internal Communication with gigasmart cards.
Bidirectional	TCP	6379	Redis	Used by REDIS for cluster communication.
Bidirectional	TCP	6381	Redis	Used by REDIS for cluster communication.
Bidirectional	TCP	16381	Redis	Used by REDIS for cluster communication.

Note: The following ports are blocked by firewall internally (and no security issues have been observed). You cannot access the device using these ports:

- 111
- 662
- 890
- 892
- 2020
- 2049
- 4045
- 6379
- 6381
- 16389

Health Status

GigaVUE-FM computes the health status of the devices by collecting inventory information and also information related to physical and logical components. The health of the devices is recomputed periodically based on the confic sync updates.

This appendix provides the health status information of the following components:

- Node Health Status
- Port Health Status
- Map Health Status
- GigaSMART Map Health Status
- Flow Health Status

How GigaVUE-FM Computes Health Status

GigaVUE-FM determines the health status of the clusters based on the events received and information collected from the devices. In a scaled environment, GigaVUE-FM manages a large number of devices and computes the health state in an optimized way:

Prioritization of Tasks

GigaVUE-FM receives events such as state changing and traffic related SNMP traps. It also periodically collects configuration changes through the config refresh cycle. GigaVUE-FM prioritizes these tasks in the following order:

- Config Update/Delete (performed through GigaVUE-FM APIs)
- Node Addition/Deletion (performed through GigaVUE-FM APIs)
- SNMP Traps/RMQ Events
- User Triggered Config Sync
- Background Config Sync

The SNMP traps are in turn prioritized in the following order:

- Utilization Trap
- Packet Drop Trap

- Packet Error Trap
- Module State Change Trap
- Trap Link State Change Trap
- Physical Component State Changing Trap

NOTE: A background config refresh collects inventory information for the components and submits its task to GigaVUE-FM to process the health state of the components. If GigaVUE-FM receives a state changing trap (link down, link up) simultaneously for the same component, then based on priority, the trap is processed first and the config refresh health computation on the component will happen next. This leads to old state information of the component to be replaced with the new state information, which will be corrected in the next config refresh cycle.

Health Queue Threshold

GigaVUE-FM maintains a health queue threshold to throttle the number of tasks submitted to it. This ensures that GigaVUE-FM is not oversubscribed with a huge number of tasks. The health queue threshold is applicable only to the following tasks and is derived based on the system profile:

- Config Refresh
- SNMP Traps

Backoff Mechanism

GigaVUE-FM implements a 'Backoff Mechanism' which ensures that after a particular threshold limit is reached, background tasks such as config refresh and SNMP traps will back off from starting its process of performing rediscovery and config refresh on clusters. For example, if config refresh threshold (health queue) is reached, then the subsequent config refresh for the clusters will be backed off. GigaVUE-FM maintains counters for the clusters that were backed off when they reached the threshold limit.

NOTE: If a user is forcefully rediscovering the cluster and if the health queue is busy, then the request will be ignored with the following message 'Backing Off the Attempt as Health Engine is busy'.

Node Health Status

The status of a node is determined by the health status of the following components:

- Ports
- Cards
- Fan Trays
- Power Modules
- Memory utilization
- CPU utilization

The health of a port and a fan depends on the health status of its associated components. For example, the health of a card depends on the port health. If more than 50% of the ports in a card are up and the operational status of the card is also up, then the card is determined as healthy (green). Similarly, the health of a fan depends on the operational status of the fan tray.

NOTE: GigaVUE-FM computes the health status of the node based on FanChange and PowerChange traps and the same is reflected in GigaVUE-FM GUI. The card health status is computed based on ModuleChange trap and the same is reflected in GigaVUE-FM GUI.

A node is considered unhealthy if:

- at least 50% of the cards are down
- at least 50% of the ports in a card are down
- at least 1 power module is down
- the average memory usage over the past one hour is more than 70%
- the CPU load per core is more than 50% overloaded
- the Different Power Supply Modules [PWR-HC1P-01 and PWR-HC1P-02] are used (Applicable only for GigaVUE-HC1-Plus and GigaVUE-HC3)

The change in the health status of a node is indicated in Events.

The cluster health is determined by the health status of the devices associated to the cluster.

The health status of a node is indicated by the following colors:

Color	Health Status
Green	Up (connected, healthy)
Amber	Warning This state is displayed when the operational status of the card is up and 50% of the associated ports are up.
Red	Down (disconnected), unreachable
Gray	Unknown This state is displayed when newly added nodes are yet to be discovered by GigaVUE-FM.

GigaVUE-FM determines the health state of the ports and devices based on the following SNMP traps, which are enabled by default when the GigaVUE-FM fabric manager initializes:

- Packet Drop
- GigaSMART Packet Drop
- Packet Error
- Port Utilization
- Low Port Utilization
- GigaSMART Port Utilization
- GigaSMART Port Low Utilization
- System Memory Threshold
- Process Memory Threshold

GigaVUE-FM determines the health status of the devices based on the receipt of the traps mentioned above:

- If GigaVUE-FM receives any of the traps mentioned above: Port state is set to yellow or red.
- If the traps are not received within the configured interval specified in the SNMP Throttling Page: Port state is set to green.

NOTE: You must configure the throttling interval in the SNMP Throttling page for the traps mentioned above. Otherwise, GigaVUE-FM cannot determine the health status of the ports.

Port Health Status

The health of a port is determined by multiple factors:

- Operational status
- Packet drops (Optional)
- Packet errors (Optional)

GigaVUE-FM computes the health status of a port and its associated logical components such as Map when a port link changes. When a port flaps, GigaVUE-FM computes the health status and the same is reflected immediately in GUI.

Click  on the top navigation bar. On the left navigation pane, select **System > Traffic Health Thresholds**. The packet drops and errors are enabled by default for computing the health status of a port. For information about setting traffic health thresholds, refer to the *"Traffic Health Thresholds"* section in the *GigaVUE Administration Guide*.

In this example, the Port Packet Drops and Port Packet Errors are enabled. The threshold value is set to 15000 packets over a time interval of 15 min. Refer to the following table to view how the port health status is calculated.

Color	Health Status	Operational Status	Packet Drops over 15 min	Packet Errors over 15 min
Green	Up (healthy)	Up	< 15000	< 15000
Red	Down (unhealthy)	Down	< 15000	< 15000
Red	Down (unhealthy)	Up	< 15000	> 15000
Red	Down (unhealthy)	Up	> 15000	> 15000
Red	Down (unhealthy)	Up	> 15000	> 15000

Inline Networks

The health of an inline network port depends on the forwarding state of the inline networks. GigaVUE-FM checks the forwarding state every 5 min. Refer to the following table to view how the health status of the inline network port is calculated.

Color	Health Status	Forwarding State
Green	Up (healthy)	Normal
Red	Down (unhealthy)	Failure-introduced Drop
Red	Down (unhealthy)	Network Ports Forced down
Red	Down (unhealthy)	DISCONNECTED
Red	Down (unhealthy)	ABNORMAL
Amber	Warning	Failure - Introduced Bypass
Amber	Warning	Forced Bypass with Monitoring
Amber	Warning	Disabled
Amber	Warning	Forced Bypass

Map Health Status

The health of a map is determined by the health status of its associated components such as ports, port groups, port pairs, GigaStream, tool port, GigaSMART group, tunneled port, virtual port, inline network, inline tool, inline tool group, inline serial tool group, inline network group, and GigaSMART operations. If the status of any one of the component is down, the corresponding map is considered unhealthy.

If a user creates/edits a Map through GigaVUE-FM, then those changes are reflected immediately to other users who are viewing the Map List View.

The health status of a map is indicated by the following colors:

Color	Health Status
Green	Up (healthy)
Amber	Warning This state is displayed when one or more ports associated to the map are unhealthy.
Red	Critical (unhealthy)
Gray	Unknown

NOTE: If you hover your mouse over the **Map Status** field, the health of the stack GigaStream or stack port (configured in the stack link) or the source and destination ports is displayed in a tooltip (if ports are unhealthy). You can derive the map health status based on the details displayed in the tooltip.

GigaSMART Map Health Status

The health of a GigaSMART map is determined by the health of the following components:

- All GigaSMART engine ports in the GigaSMART group (gsgroup)
- Virtual ports (vport) associated with the GigaSMART group
- GigaSMART operations
- IP Interfaces

Refer to the following table to view the health status of a map with GigaSMART:

Color	Health Status	GigaSMART Group	vPort	GigaSMART Operations	IP Interfaces
Green	Up	Up	Up	Up	Up
Red	Down	Up	Down	Up	Up
Red	Down	Up	Up	Down	Up
Red	Down	Up	Up	Up	Down
Red	Down	Down	Down	Down	Down

GigaSMART Group Health Status

The health of a GigaSMART group (gsgroup) depends on the aggregated health of the associated GigaSMART engine ports. The following components contribute to the health of the GigaSMART engine ports:

- **Operational Status**—The operational status of the associated GigaSMART engine ports. If the operational status of any GigaSMART engine port in the GigaSMART group is down, then the GigaSMART group becomes unhealthy.
- **GigaSMART Engine Port Packet Correlation**—The percentage (%) of packet correlation seen in a GigaSMART engine port. The GigaSMART engine packet correlation is calculated based on the following factors:
  - the cumulative number of packets coming into a GigaSMART group
  - the cumulative number of packets going out of a GigaSMART interface
  - the cumulative number of packets dropped at a GigaSMART operation for a map

  If the number of packets going out of a GigaSMART interface exceeds the threshold set in **Administration** > **System** > **Traffic Health Thresholds**, the GigaSMART engine port becomes unhealthy.

- **GigaSMART Engine Port Packet Drops**—The cumulative number of packets dropped due to over subscription of a GigaSMART engine port. If the number of GigaSMART engine port packet drops exceed the threshold set in **Administration** > **System** > **Traffic Health Thresholds**, then the GigaSMART engine port becomes unhealthy.
- **GigaSMART Engine Port Packet Errors**—The cumulative number of packet errors coming into a GigaSMART engine port. If the number of GigaSMART engine port packet errors exceed the threshold set in **Administration** > **System** > **Traffic Health Thresholds**, then the GigaSMART engine port in the GigaSMART group becomes unhealthy.

For information about setting traffic health thresholds, refer to the *"Traffic Health Thresholds"* section in the *GigaVUE Administration Guide*. All threshold types are enabled by default for computing the health status of the GigaSMART engine ports.

In this example, the thresholds are enabled and the threshold values are set as follows:

Type	Threshold Value	Interval
GigaSMART engine port packet correlation	50	15
GigaSMART engine port packet drops	15000	15
GigaSMART engine port packet errors	15000	15

Refer to the following table to view how the GigaSMART engine port health status is calculated.

Color	Health Status	Operational Status	Packet Correlation	Packet Drops over 15 min	Packet Errors over 15 min
Green	Up (healthy)	Up	< 50	< 15000	< 15000
Red	Down (unhealthy)	Down	-	-	-
Red	Down (unhealthy)	Up	> 50	< 15000	< 15000
Red	Down (unhealthy)	Up	< 50	> 15000	< 15000
Red	Down (unhealthy)	Up	< 50	< 15000	> 15000

The GigaSMART group health status is determined by the aggregated health of the GigaSMART engine ports. Refer to the following table for computing the health of a GigaSMART group:

Color	GigaSMART Group Health	GigaSMART Engine Ports
Green	Up	Up
Amber	Warning	Some engine ports are down
Red	Down	All engine ports are down

vPort Health Status

GigaSMART virtual port is used as an aggregation point for traffic directed to second level maps. Second level maps include an Adaptive Packet Filtering component or a GTP rule.

A vPort is healthy when the GigaSMART group associated with the vPort is healthy. If a gsgroup is unhealthy and the vPort is healthy, this indicates that the vPort is not participating in the maps.

GigaSMART Operations Health Status

GigaSMART Operation consists of one or more advanced processing applications.

A GigaSMART operation (gsop) is healthy when the GigaSMART group associated with the gsop is healthy. If a gsgroup is unhealthy and the gsop is healthy, this indicates that the gsop is not participating in the maps.

IP Interfaces Health Status

IP interfaces are used for GigaSMART encapsulation and decapsulation operations on both network ports and tool ports.

An IP interface is healthy when the GigaSMART group associated with the IP interface is healthy. If a gsgroup is unhealthy and the IP interface is healthy, this indicates that the IP interface is not participating in the maps.

For information on how to calculate the map health, refer to Map Health Status.

Flow Health Status

The health of a flow is determined by the aggregated health of the maps in the flow. The factors that determine the health of a flow is as follows:

- Health of the priority maps in the flow
- Health of the maps that constitute the priority map set
- Gigamon Discovery or Manual links

Priority Map Set Health

A priority map set consists of more than one map configured with the same source ports in priority order. The health of such map is determined by the aggregated health of the constituted maps. For information about how map health is computed, refer to Map Health Status.

The following table provides a summary of the health status of a map chain:

Color	Priority Map Set Health Status	Maps in the Map Chain
Green	Healthy	All maps are healthy
Amber	Warning	One or many maps in warning state
Red	Unhealthy	One or many maps in unhealthy state

Flow Health Computation

Starting with software version 5.4, the current throughput percentage used in determining the health of a tool device would is be determined as follows:

Current Throughput %	Tool Health	Reasons
<85	Green	All maps are healthy
>85 && <100	Yellow	>85 && <100 Yellow Tool device is experiencing throughput between 85% to 100%
>100	Red	One or many maps in unhealthy state

Based on the tool device health, flow health for the flows having tool device would be computed as follows:

Existing Flow Health	Tool Health	New Flow health	Reasons
Red	Red/Yellow/Green	Red	Existing Flow Health reasons + tool health reasons
Yellow	Red	Red	Existing Flow Health reasons + tool health reasons
Yellow	Yellow	Yellow	Existing Flow Health reasons + tool health reasons
Yellow	Green	Yellow	Existing Flow Health reasons
Green		Red	Tool health

Existing Flow Health	Tool Health	New Flow health	Reasons
			reasons
Green	Warning	One or many maps in warning state	Tool health reasons
Green	Green	Green	--

GigaVUE-FM APIs

GigaVUE-FM APIs are designed with the Representational State Transfer (REST) architecture, which provides a well structured architecture for performing query, update, and delete functions in a programmatic manner. GigaVUE-FMAPIs are implemented based on Open API 3.0 specification (formerly known as Swagger) and conform to required standards.

NOTE: You can access the GigaVUE-FM APIs without being authenticated in to GigaVUE-FM. Type <fm-ip/apiref/apiref.html> in your browser to view the API Reference page.

Click Here for the Online API Reference.

Supported Cloud Environments

The GigaVUE Cloud Suite solution for the various cloud platforms offers network traffic visibility in the respective cloud platforms. The GigaVUE Cloud Suite solutions acquire, optimize and distribute selected traffic to security and monitoring tools. The below table lists the available cloud suites and their respective guides.

Cloud Platform	Guides
Public Cloud	
AWS	GigaVUE Cloud Suite Deployment Guide - AWS GigaVUE Cloud Suite for AWS–GigaVUE V Series 1 Guide
Azure	GigaVUE Cloud Suite Deployment Guide - Azure GigaVUE Cloud Suite for Azure–GigaVUE V Series 1 Guide

Cloud Platform	Guides
Private Cloud	
OpenStack	GigaVUE Cloud Suite for OpenStack–GigaVUE V Series 1 Guide GigaVUE Cloud Suite Deployment Guide - OpenStack
VMware	GigaVUE Cloud Suite Deployment Guide - VMware GigaVUE Cloud Suite for VMware—GigaVUE-VM Guide
Container Visibility	
Kubernetes	GigaVUE Cloud Suite for Kubernetes Container Visibility Configuration Guide
Gigamon Containerized Broker	Gigamon Containerized Broker Deployment Guide
Universal Container Tap	Universal Cloud TAP - Container Deployment Guide
Other Cloud Platforms	
Third party Orchestration	GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration
Nutanix	GigaVUE Cloud Suite for Nutanix Guide—GigaVUE-VM Guide

The guides provide instructions on configuring the GigaVUE Cloud components and setting up traffic monitoring sessions for the respective Cloud platform.

Analytics

Analytics in GigaVUE-FM is a standalone service that provides data visualization capabilities for the following entities in GigaVUE-FM:

- Physical resources, specifically the nodes and the ports
- Virtual resources
- Alarm Management services
- GigaVUE-FM CPU, Memory and Disk Storage services

Using Analytics[1] you can create visual elements such as charts that are embedded as visualizations. The visualizations are grouped together in dashboards. You can also create search objects using Analytics. Dashboards, Visualizations and Search Objects are called Analytics objects.

Analytics provides the following advantages:

- Real-time data for visualization as the required data is taken from GigaVUE-FM.
- Data to be analyzed and visualized is fetched based on the access control rights of the user.
- Trending Analysis: Visualize the trends in the traffic using pre-defined widgets.
- Capacity Planning: Utilization of resources based on health and inventory summary data.
- Generation of reports based on the available data.

Rules, Notes, and Limitations for Analytics

- All GigaVUE-FM users can create, edit and delete Analytics objects[2]. However, you can perform these operations only on the objects created by you. You cannot delete or edit system-defined objects such as dashboards and visualizations.
- You can only view the Analytics objects created by other users, but you cannot edit them.
- There is no limit on the number of dashboards per user.
- The data available in dashboard is controlled by Role Based Access Control. That is, the data fetched depends on the accessibility rights of the user based on the user role and user-defined tags.

[1]Analytics uses the OpenSearch front-end application to visualize and analyze the data in the OpenSearch database of GigaVUE-FM.
[2]Analytics objects include Dashboards, Visualizations and Saved Search Objects.

- The GigaVUE-FM backup/restore operation preserves both default and custom dashboards and visualization. This allows the dashboards and visualizations created in the earlier software versions to be restored in software version 5.13.00. This enables rapid resumption of GigaVUE-FM services and also provides the users the ability to operate and monitor the visualizations in GigaVUE-FM.
- Statistical dashboards display data in the following two tabs.
  - **Metric**: Displays maximum and average values of statistical counters for a specific time period, example Rx Ports Aggregated Max Traffic, Rx Ports Aggregated Average Drop Traffic.
  - **Trends**: Displays trend of statistical counters for a specific time period, example Rx Port Statistics (bps), Tx Port Statistics (bps), Tool GigaStream Maximum Rx Rate (bps) by member port Alias.
- The Inbound and Outbound Port Statistics dashboards display metric values for the top 1000 port elements in visualizations (as the total count it set to 1000).
- By default, all statistical dashboards display data based on a pre-selected cluster ID to avoid performance issues in a scaled environment.
- When browsing the Analytics page with 5 minute granularity, you can observe visual gaps in the displayed results. This is because stats collection happen in GigaVUE-FM with 5 minute granularity and a small delay of few seconds in stats collection can result in such gaps over time.
- For custom dashboards, enable the resource type filter to generate faster search results.
- For rate-based visualizations in the default system dashboards and cloned system dashboards, the axis-min setting (under panel option) is set to zero (0). An empty graph is therefore displayed when the data point values are zero. However, for visualizations in the dashboards created using the **Create Dashboard** option, you must manually change the axis-min setting to '0' for an empty graph to be displayed.
- For gauge-based visualizations "no data to display" message is received in case of the following scenarios:
  - GigaVUE-FM or the device is down during a particular time interval.
  - No traffic in the device and the axis-min panel option is set to value > 0.
- For rate-based visualizations, "no data to display" message is received in case of the following scenarios:
  - With *Drop Last Bucket* set to 'Yes' and time interval range < 15 minutes or time range has less than three data points.
  - With *Drop Last Bucket* set to 'No' and time interval range < 10 minutes or time range has less than two data points.
  - When a large time range is selected, the chart uses a larger data interval for rendering. If the selected range has insufficient data points, the visualization may appear empty. Select a smaller time range to improve trend visibility.

GigaVUE-FM Statistics and Data Roll-up

Until software version 6.0.00, the statistical dashboards of the physical and logical resources display data collected for a period of 30 days with 5 minute granularity (considered as the base index). Starting from software version 6.1.00, the statistical dashboards display data collected for a period of 35 days.

Resource	Related Statistical Dashboards
Ports	• Inbound Port Statistics • Outbound Port Statistics
Maps	Map Statistics
Nodes	Inventory Statistics
GigaStreams	• Circuit and Stack GigaStream Statistics • Tool GigaStream Statistics
GigaSMART	GS Group Statistics

With Data Rollup, GigaVUE-FM allows you to collect, summarize and display historical data up to 120 days with hourly granularity. Refer to the following notes:

- Starting from software version 6.1.00, Rollup is enabled by default. To disable Data Rollup, navigate to the Storage Management page. Refer to the GigaVUE Administration Guide for details.
- Rollup is applicable only for ports, maps, and map rules. You can search and analyze port statistics for a duration of up to 120 days. GigaVUE-FM collects statistical data from the day you upgrade to software version 6.1.00 or higher. Complete 120 days of data will be visible only from 121st day (and not backwards from the day GigaVUE-FM was upgraded).

  NOTE: Associated Map Rule Visualization will only show the latest map rule count of 35 days.

- On upgrading to software version 6.1.00, ensure to have increased disk space for data roll-up. Refer to the Recommended Resource Requirements for Scaled Environments section in the GigaVUE-FM Installation Guide for detailed information.

Control Filters in Analytics Dashboards

You can use control filters in your dashboards to filter and display the data you want to explore:

- To filter data in the visualizations based on the required tag keys and tag values, clone the required statistical dashboard and edit the *Sample - Tags* option in the Control Filters. Refer to Filter Data Using Tags in Control Filters

- Starting from software version 6.2.00, Cluster ID is a mandatory control filter field in the statistical dashboards. When cloning the dashboards and visualizations, you must use at least one of the following control filter fields to filter the data:
  - Cluster ID
  - Tag
- To optimize the performance of GigaVUE-FM, starting from software version 6.3.00, the data listed in the following control filter fields in Analytics dashboards is fetched from fminventory index (instead of fmstats index):
  - Cluster ID
  - Tags
- As data is fetched from fminventory index, you cannot view the details of devices that are no longer managed by GigaVUE-FM (as they will not be available in the fminventory index).

  You can disable this behavior by logging in to the GigaVUE-FM CLI and performing the following steps:

  1. Log in to the command line prompt:
     ```
     sudo su -
     ```
  2. Navigate to the folder where the .json file is located
     ```
     cd /var/lib/fabricHealth/conf
     ```
  3. Open the file using vi editor.
     ```
     vi custom_settings.json
     ```
  4. Set the **`load_legacy`** attribute value to "true".
  5. Restart Analytics service.
     ```
     systemctl restart fabricHealth.service
     ```

Reference Topics

Refer to the following sections for details:

- Get Started with Analytics UI
- Work with the Analytics User Interface

Get Started with Analytics UI

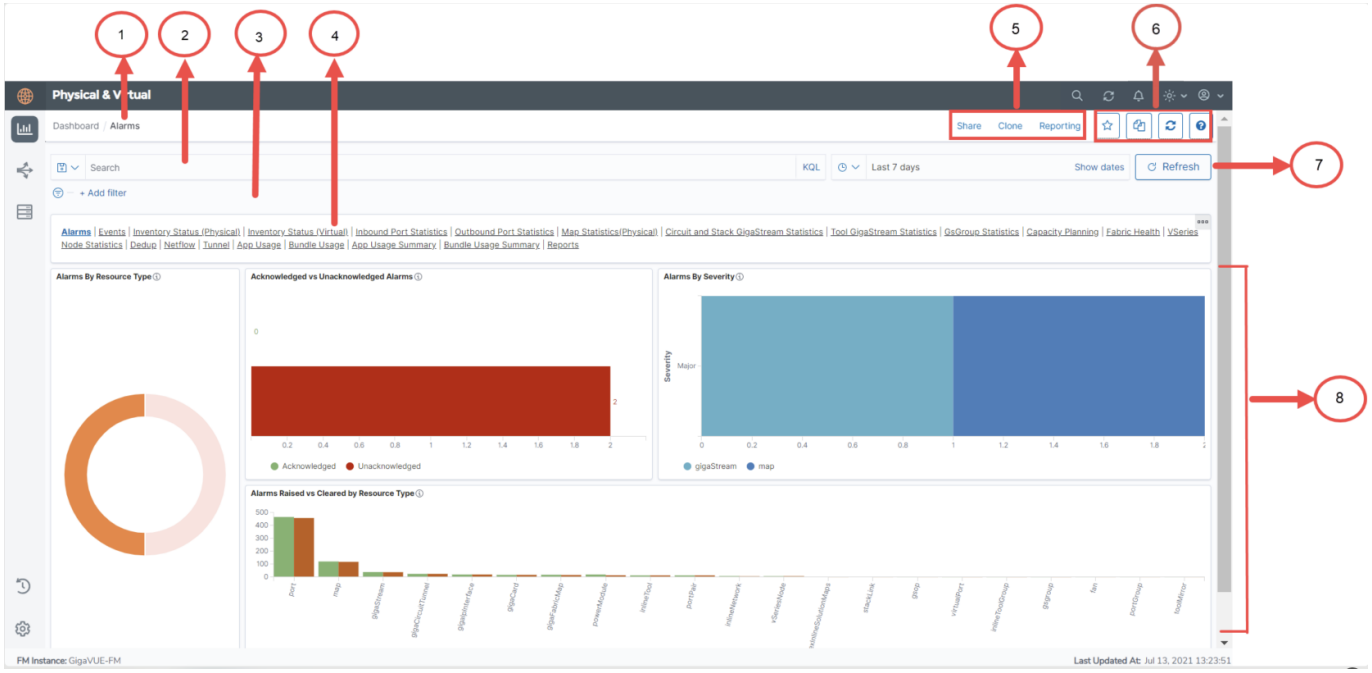
The Analytics option is listed under the Dashboards. To access the Analytics UI:

1. On the left navigation pane, click on  .
2. Select **Analytics** . The following options are listed:
   - **Dashboards**: Refer to the Dashboards section for details.
   - **Visualization**: Refer to the Visualizations section for details.
   - **Discover**: Refer to the Discover section for details.
   - **Reports**: Refer to the Reports section for details.

GigaVUE-FM allows you to view the Fabric Health Analytics (FHA) Dashboard with data points having granularity levels for less than 48 hours in Trend Line Dashboards.

GigaVUE-FM provides granularity to all the custom FHA Dashboards or Visualizations.

The following figure shows the available options in the Analytics Dashboard page.



Refer to the following table for details:

S.No	Description	
1	Name of the Dashboard: Example Alarms.	Dashboards
2	Search box	Filter Data in Visualizations
3	Add Filter	In the Dashboard, refrain from using multiple filters in the global filter for better results.
4	Dashboard Navigation bar	Copy Dashboard Path
5,6	Working with the GUI	Work with the Analytics User Interface
7	Refresh	Use to manually refresh visualizations
8	Visualizations	Visualizations

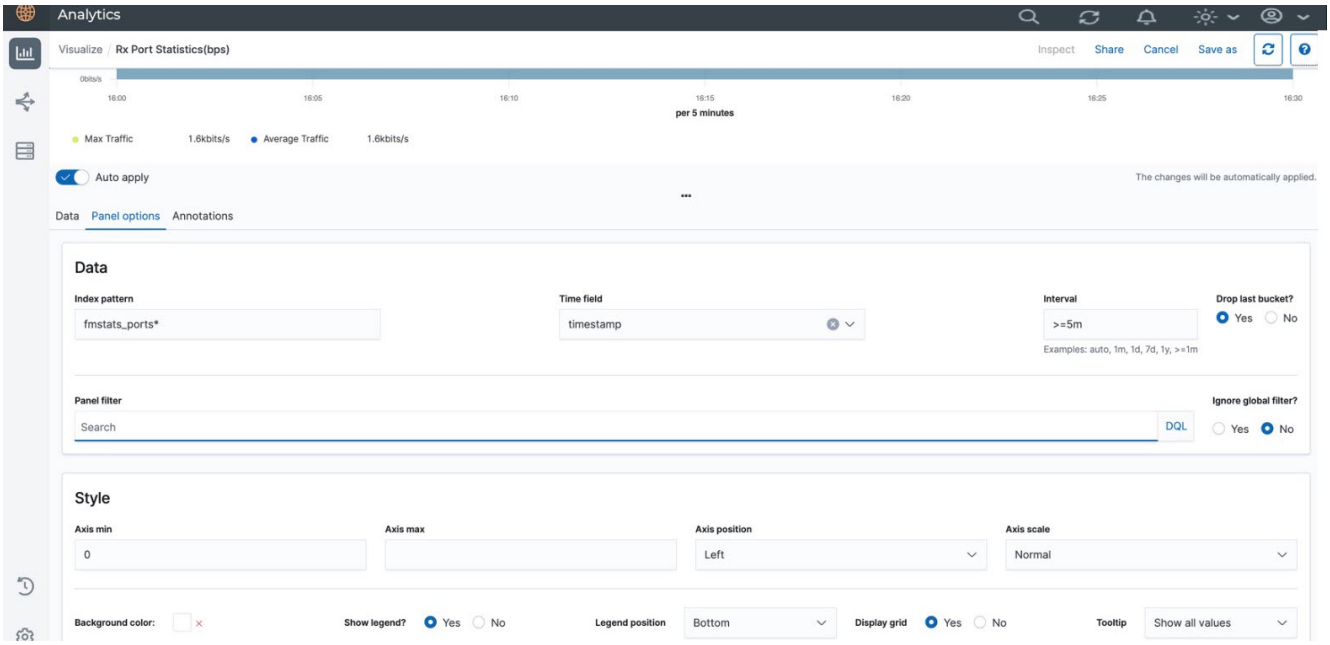
NOTE: When you cannot view the theme changes applied in Fabric Health Analytics (FHA), clear the browser cache and reload to view the updated changes.

The following table provides the granularity level at which the data is shown in the FHA Dashboard:

Interval	Granularity level at which data is shown
Up to 14 hours	5 minute granularity
15 to 48 hours	10 minute granularity

For intervals > 48 hours, 12 data points are displayed, is an expected behavior.

NOTE: If the analytics are available for < 30 minutes, it is recommended to select 5-minute granularity.



Work with the Analytics User Interface

GigaVUE-FM allows you to share the dashboard, clone the dashboard, generate the report, copy the dashboard, search data, etc., To know more about the user interface, refer to following sections:

Share

Use to share Dashboard pages with other users. To share a dashboard:

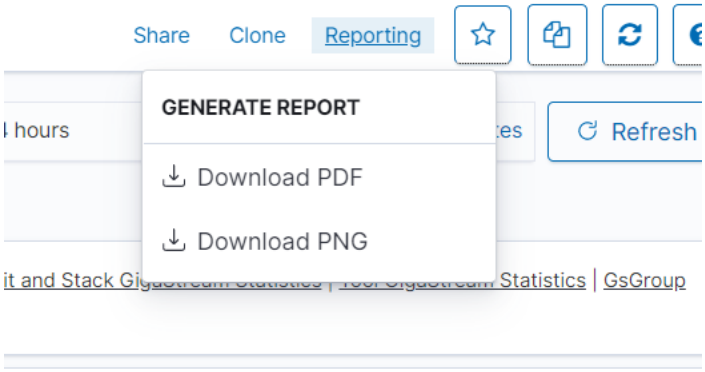
1. Navigate to the Dashboard page that you want to share.
2. From the top navigation bar, click **Share**.
3. The following options are available:
   - **Embed Code**: Use to share the dashboard in an iFrame Embed URL.
   - **Permalinks**: Use to share the permalink URL of the dashboard page.

Clone

Use to clone the system dashboard pages. Refer to the Clone Dashboard section for the details.

Reporting

Use to generate the report in PDF or in PNG format that can be downloaded instantly. The generated report is also listed in the Reports page. Refer to the Reports section for details.



Set as Default

Use to change the default dashboard page:

- Navigate to the specific dashboard page.
- Click to set this dashboard page as default.

Copy Dashboard Path

The GigaVUE-FM Analytics page allows you to add links in custom dashboard pages for navigating from one dashboard page to another dashboard page without going to the listing page. To do this:

- Navigate to the custom dashboard page on which you need to add the link.
- Click on  icon to copy the relative path.
- Create a new markdown visualization or Edit an existing markdown visualization.
- Paste the link in *[Title](RelativePath)* format. For example ***[Alarms](#/dashboard/fha-alarms)***.
- Add the markdown visualization to any of the custom dashboards on which the links of other dashboards needs to be added.

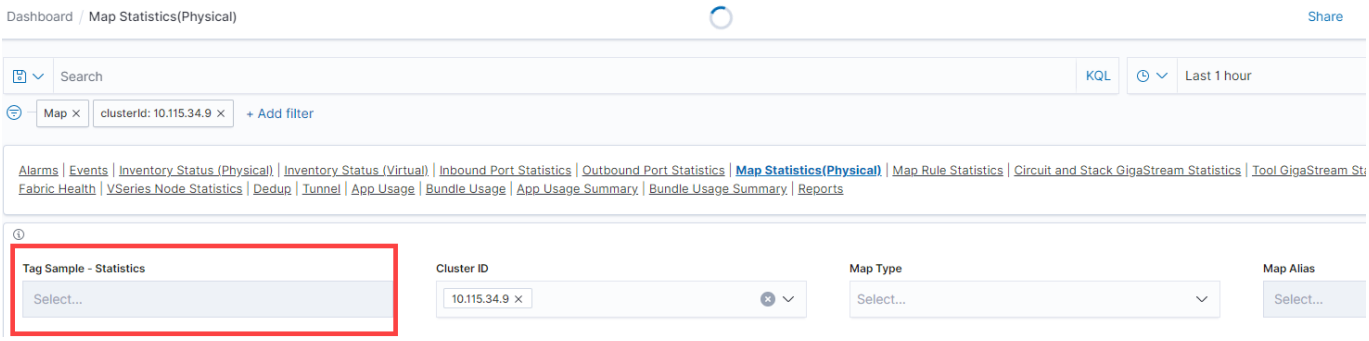
Auto Refresh Tags in Visualizations 

In Fabric Health Analytics, the data is saved as index patterns in the OpenSearch database and is fetched into the various fields in the visualizations. However, when you add tag ids dynamically in GigaVUE-FM and associate the tags to the various resources, the tag ids are not automatically refreshed in the visualizations. To add the new tags in visualizations:

1. Create new tag key and tag values from the tags page. Refer to the Create User-defined Tag section in the GigaVUE Administration Guide.
2. Associate the resources to the tag keys and tag values. For example, to add tag key and tag values to the physical nodes, refer to the Add New Physical Node or Cluster to GigaVUE-FM
3. Click on  Refresh Index Pattern icon for the newly added tags to get reflected in the visualization filters in the system and custom visualizations.

Filter Data Using Tags in Control Filters

Customize the **Tag Sample - Statistics** option available in some of the Statistical Dashboards to filter the data based on the required tags.



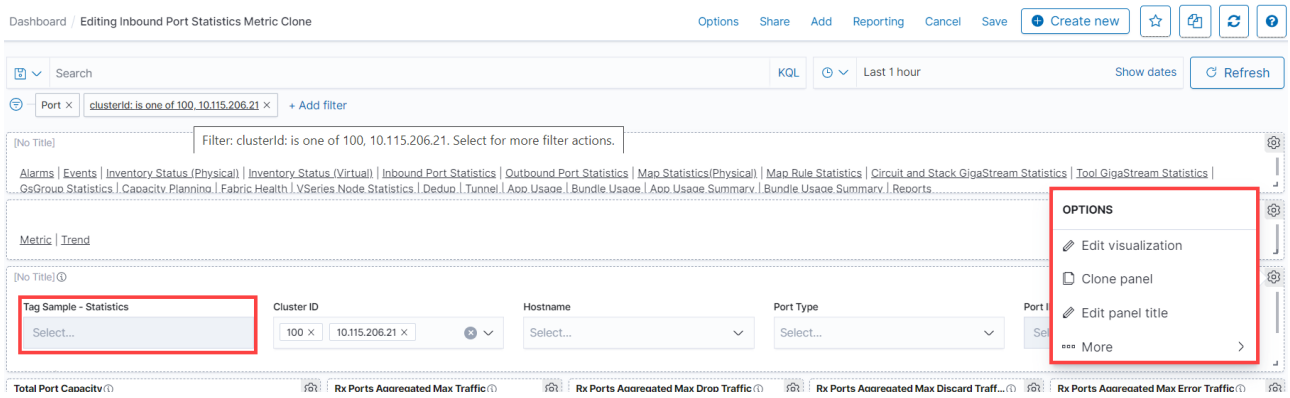
Pre-requisites

- Ensure to add the required tag keys and tag values to the tags page, and associate the tag values to the resources.
- Use the  Refresh Index Pattern icon for the newly added tags to get reflected in the Analytics page.

Consider a scenario in which you want to filter the inventory details in GigaVUE-FM based on a newly created tag key called SITE:

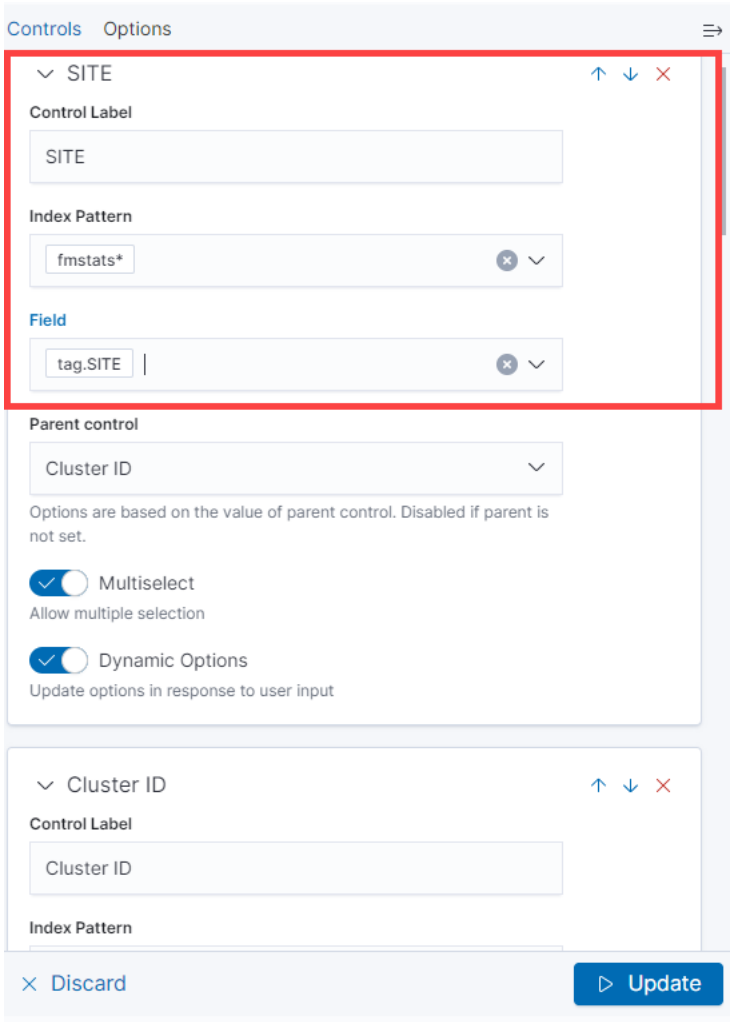
To do this:

1. Clone the required statistical dashboard.
2. Click **Edit.**
3. Scroll to the Control Visualization panel. Click **Options** and select **Edit Visualization**.



4. In the **Control** tab, configure the following. This is for filtering the data based on the tag value SITE.

Control Label	*Site*
Index Pattern	*fmstats**
Field	*tag.site*



5. Configure the existing tags as required:

Parent Control	*Use to configure a specific field as a parent based on which the other fields are filtered*
Multiselect	*Use to select multiple variables within a field*
Dynamic Options	*Use to update the dashboards and visualizations dynamically based on this criteria.*

6. In the **Option** tab, configure the following:

Update Kibana filters on each change	
Use time filter	*To update the data based on the time filter configured in the dashboard.*
Pin filters for all applications	

7. Click **Update** to update the changes.

8. Customize the panel title or remove it, as required. Move the control visualization to the top.

You can also create a new control visualization using the steps described above.

Search Data

To search your data:

- Enter the search criteria in the Query bar.
- Press Enter or click Update/Refresh button to submit the request.
- Click the Saved Queries icon to save the current query.

You can use Kibana's standard query language (KQL).

NOTE: When you submit a search request, the histogram, Documents table, Fields list and all the widgets in the dashboard are updated to reflect the search results.

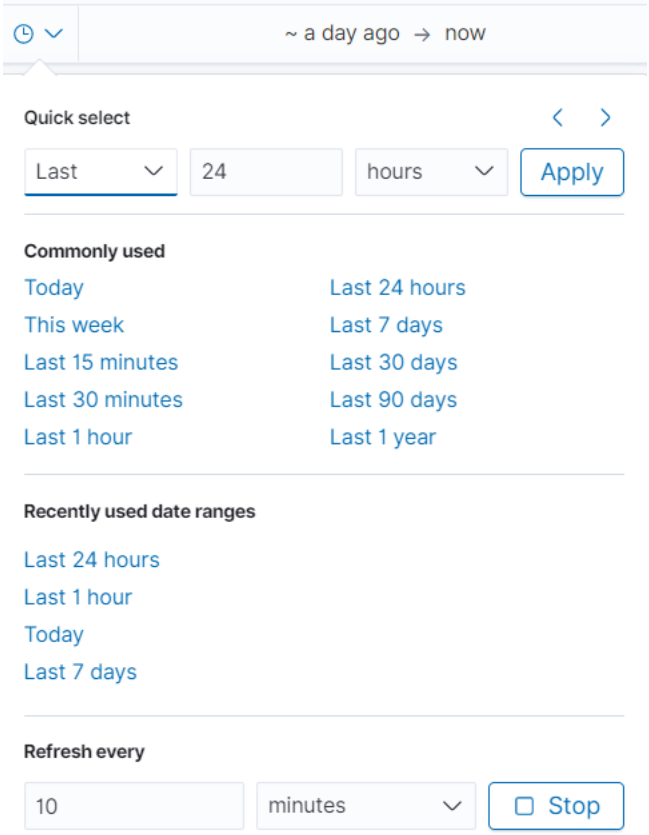
Filter Data in Visualizations

Fabric Health Analytics (FHA) provides various options to filter your data:

- **Time Filter:** Use time filter to retrieve search results for a specific time period.

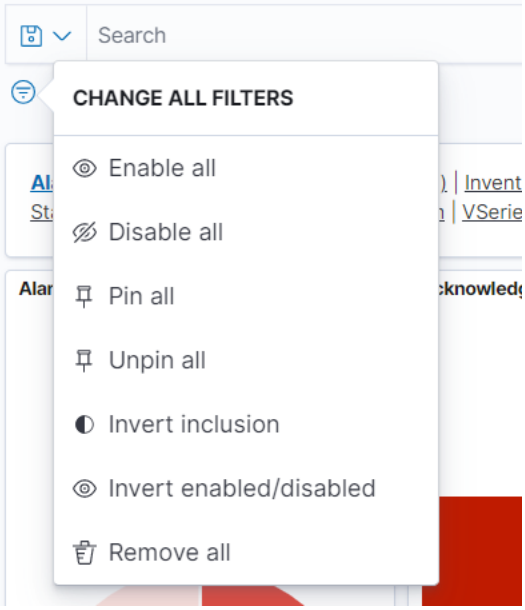
  Use the Refresh option to refresh the dashboards for the selected time interval. It is recommended to configure a longer time interval.

  Use Quick Select option to change the time interval.



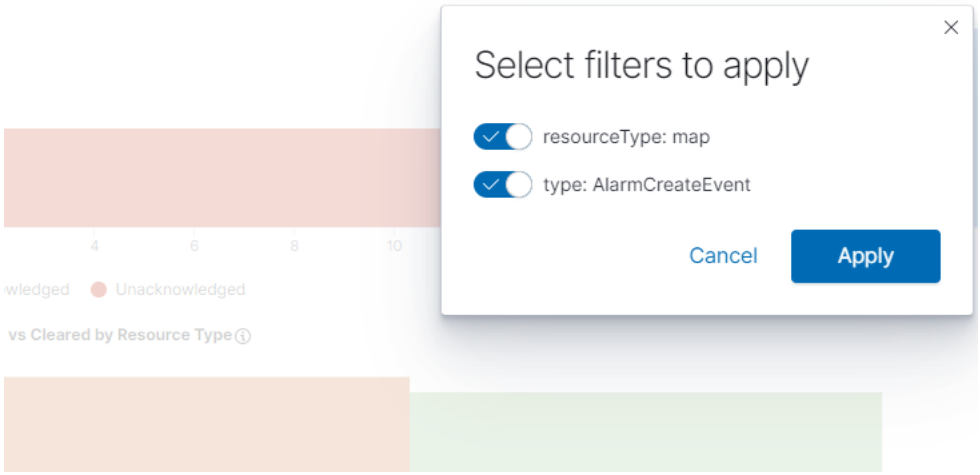
  Use the CHANGE ALL FILTERS to configure the following options:

  - **Enable all**: Enables all saved filters
  - **Disable all**: Disables all saved filters
  - **Pin all**: Filter is applied to all the dashboards
  - **Unpin all**: Filter is no longer applied to all the dashboards
  - **Invert Inclusion**: Included filters will be inverted.
  - **Invert enabled/disabled**: Enable and disable options are inverted
  - **Remove all**: Removes all filters

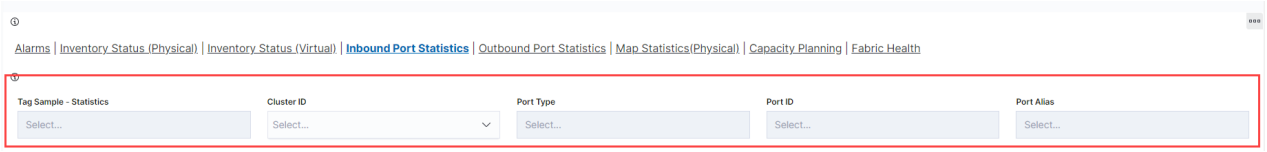


Filtering in visualizations:

- Select and drag an area of the visualization for a specific time interval. All the visualizations in the dashboard get updated for that time interval. The time interval also gets updated in the Quick Select.
- Double click on an area in the visualization and select the required filters to apply.



- **Control Visualizations:** Use Control Visualizations to filter the data based on tags. For example, in the Inbound Port Statistics visualization, you can filter the data based on the cluster id, port number, port id or port alias. Refer to the following sections:
  - Auto Refresh Tags in Visualizations
  - Filter Data Using Tags in Control VisualizationsFilter Data Using Tags in Control Filters



.

Visualizations

Visualization refers to the visual representation of data in various forms such as pie charts, time series graphs and other visual elements.

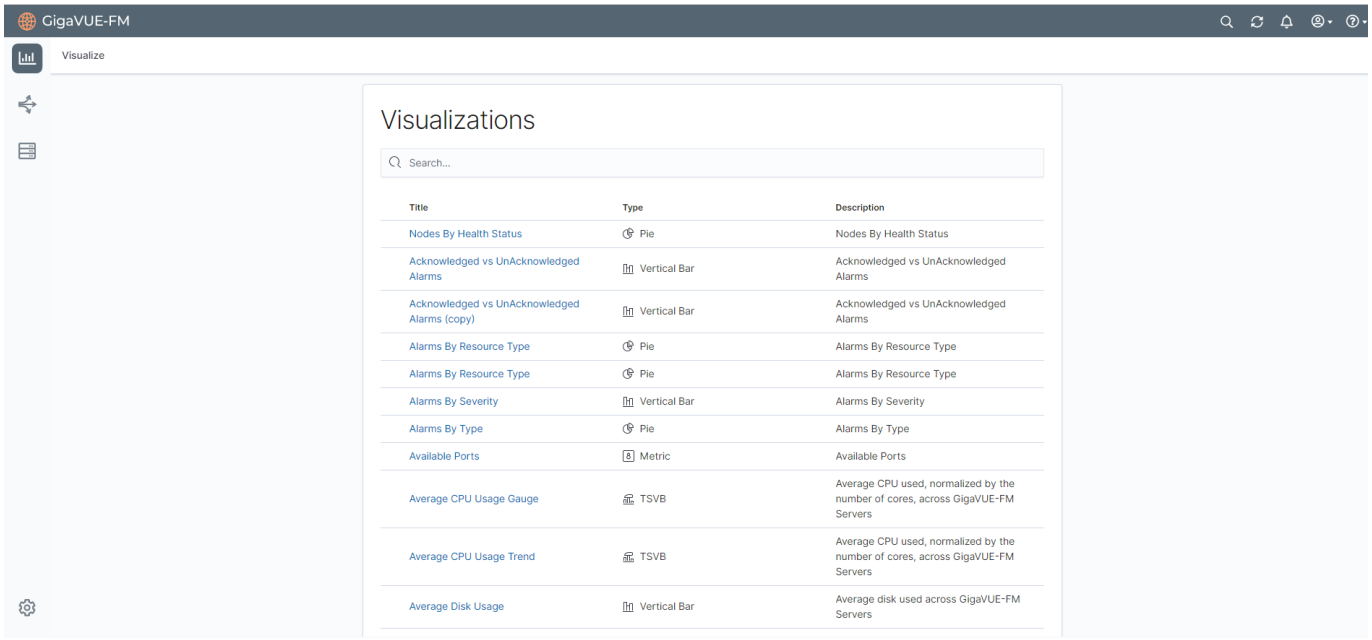
Using Analytics you can:

- View System Visualizations
- Create Custom Visualizations

View System Visualization

System visualizations are pre-defined visualizations that are available by default in GigaVUE-FM. To view the system visualization:

1. Click the Dashboard icon on the left navigation pane.
2. Select **Analytics > Visualizations**.
3. Choose the required system visualizations.



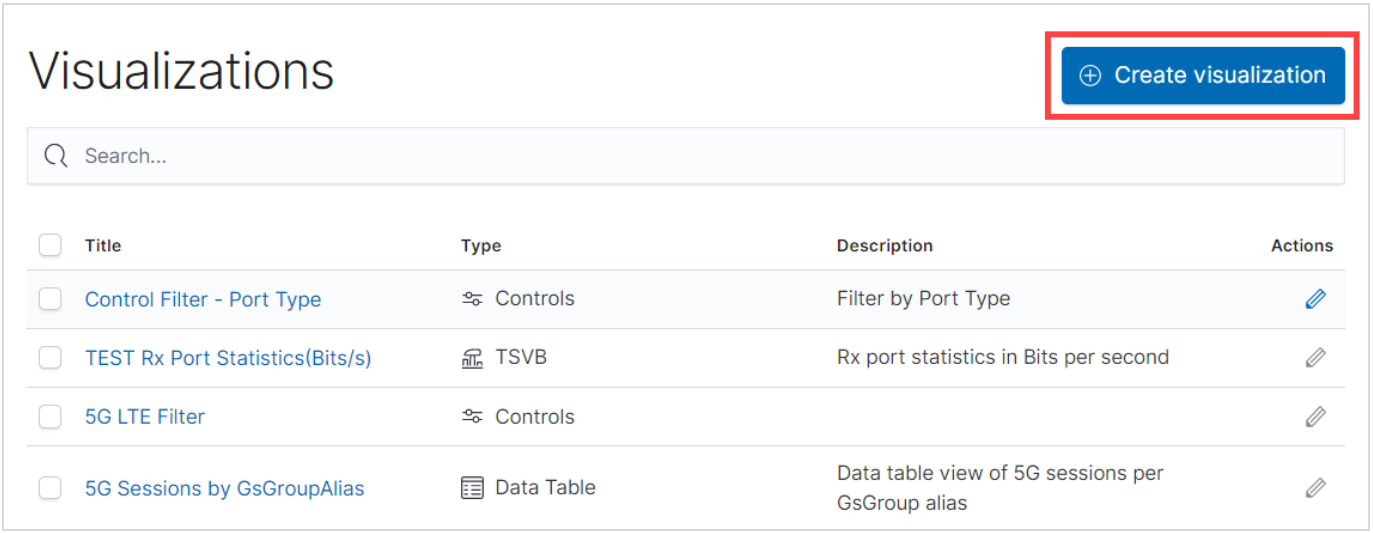
Create Custom Visualizations

You can create custom visualizations by cloning the existing system visualizations or creating a new visualization.

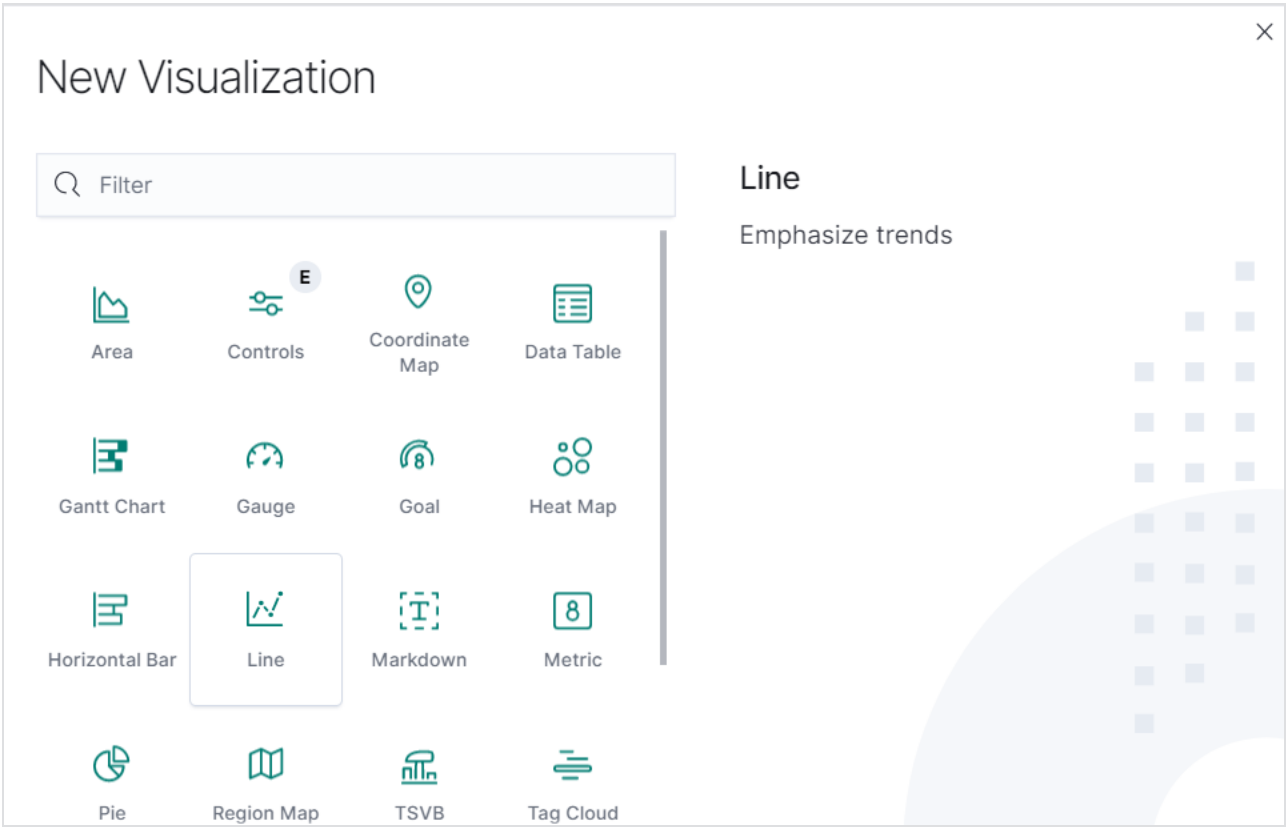
Create a Visualization

To create a new visualization:

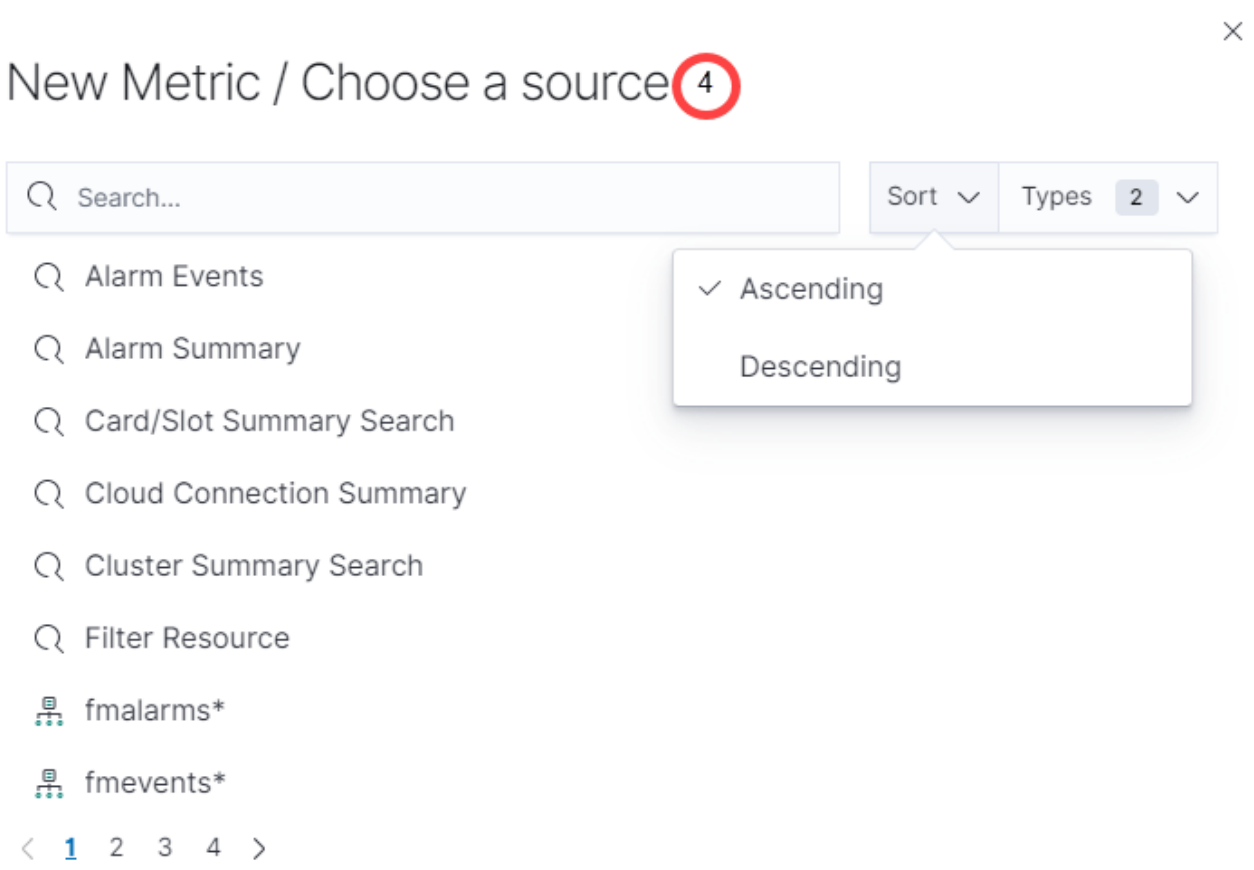
1. On the left navigation pane, click on  . Select **Analytics > Visualization**.
2. Click **Create Visualization**.



3. In the New Visualization page, select the required **Visualization Type.**



4. Select the data source for the visualization.



5. Enter the required details for the type of visualization selected. Refer to the table below for more details.
6. Click **Save**.
7. In the **Save Visualization** dialog enter the Title and Description for the visualization and click **Save**.

Type of Visualization	Description
Metric	Displays a single number for the selected aggregation.
Data table	Displays the raw data of a composed aggregation.
Pie Chart	Display each source's contribution to a total.
TSVB	Combines an infinite number of aggregations and pipeline aggregations to display complex data in a meaningful way
Line Chart, Area Chart, Horizontal and Vertical Bar charts	Compares different series in X/Y charts.
Heat maps	Shade cells within a matrix.
Markdown widget	Display free-form information or instructions.
Goal and Gauge	Displays a gauge
Coordinate map	Associate the results of an aggregation with geographic locations.
Region map	Thematic maps where a shape's color intensity corresponds to a metric's value.

NOTE:
1. When creating TSVB visualizations using derivative aggregation, you have the option to utilize the 'ignore' keyword in the unit field. This can be particularly beneficial when working with counter metrics and aiming to display only the difference instead of the rate per time unit in the graphical representation. Due to OpenSearch's limitation, the visualization's maximum search results will be limited to 10,000 entries. It is recommended to use filters for refined search results.
2. In the OpenSearch dashboards, you can only change the bucket priority by dragging fields if your screen resolution is set to 100 percent.

Clone a Visualization

To clone a visualization:

1. On the left navigation pane, click on  . Select **Analytics > Visualization**.
2. Select the visualization for which you need to create a clone.
3. Make the required changes.
4. Click **Save As**.
5. Enter a name for the new Visualization.
6. Click **Save** As. The new visualization will be added to the list page.

Visualizations - Example Work Flows

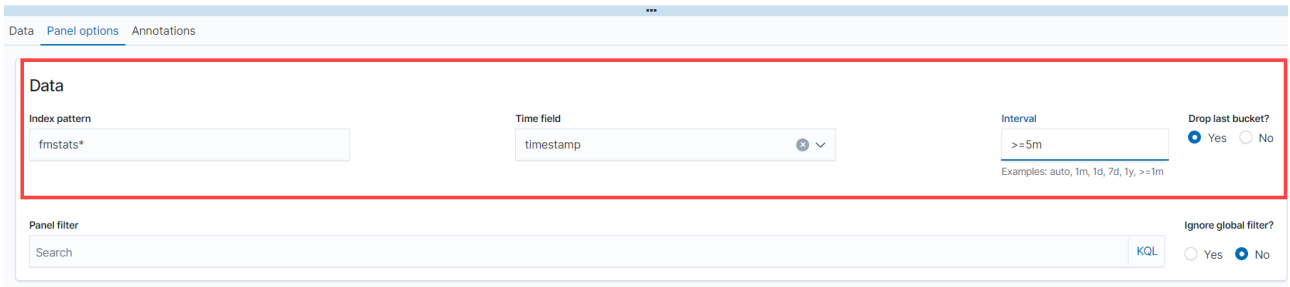
This section includes examples for configuring the visualizations:

- TSVB Chart Displaying Traffic Trend
- Pie Chart Displaying Alarms Summary
- Metric Displaying Card Count
- Bar Chart for Alarms by Severity

TSVB Chart Displaying Traffic Trend

To create a TSVB chart that shows traffic trend for every five minutes:

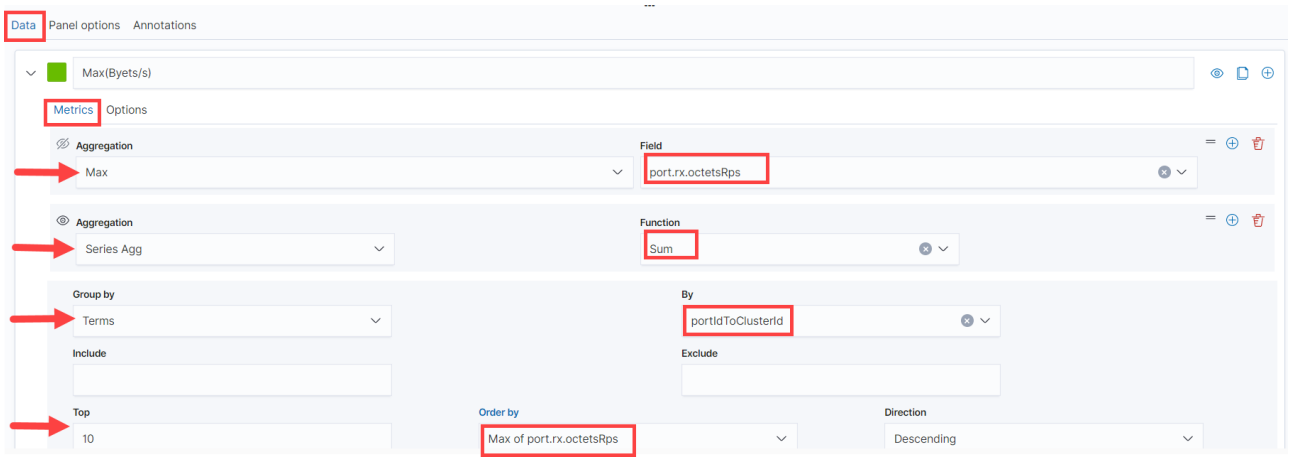
1. Click **Create Visualization**. In the New Visualization page, select **TSVB.**
2. Click **Panel Options.**



3. Select or enter the following details under Data:

Index Pattern	*fmstats**
Time field	*Timestamp*
Interval	*Must be >=5 minutes*
Drop last bucket	*Must be checked*

4. Click **Data** and Select **Metrics.**

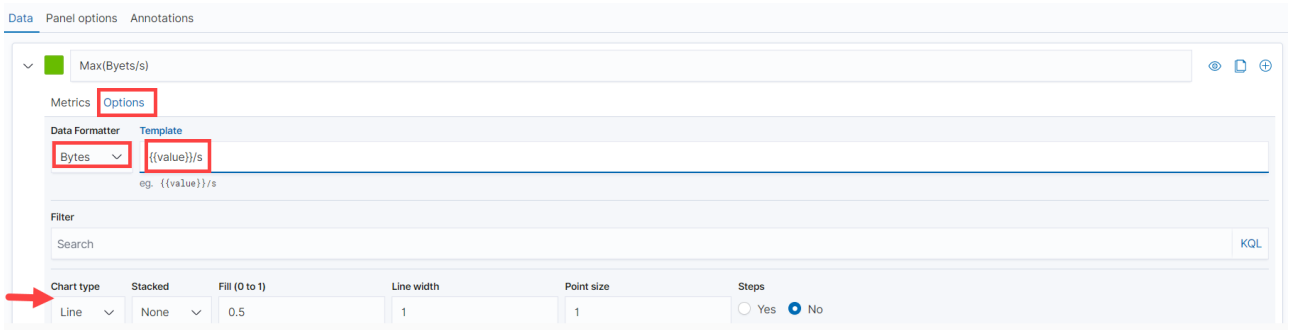


Select or enter the following details:

Label: Configure the label as Max Rate (Bytes/s)	
Aggregation	*Max*
Field	*port.rx.octets.Rps*
Create another aggregation	

Aggregation	*Series Agg*
Function	*Sum*
Group By	*Terms*
By	*PortIdToClusterId*
Top	*10*
Order by	*Max of port.rx.octets.Rps*

5. Select **Options.**



Data Formatter	*Bytes*
Template	*Values/s*

5. Click **Save** to save the visualization.

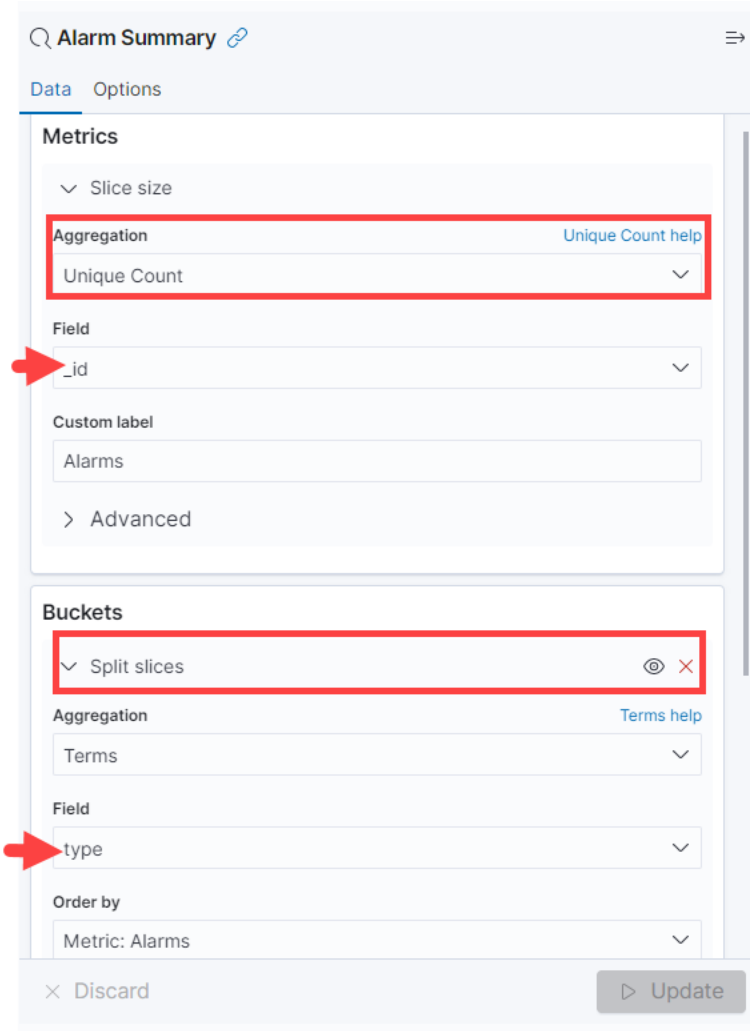
NOTE: Use the chart type option allows you to configure the chart.

Pie Chart Displaying Alarms Summary

To create a pie chart that shows alarm summary:

1. Click **Create Visualization**. In the New Visualization page, select Pie**.**
2. Select the data source for the visualization. It can be either an index pattern or a saved search object. *In this example, Alarm Summary is selected as the data source.*

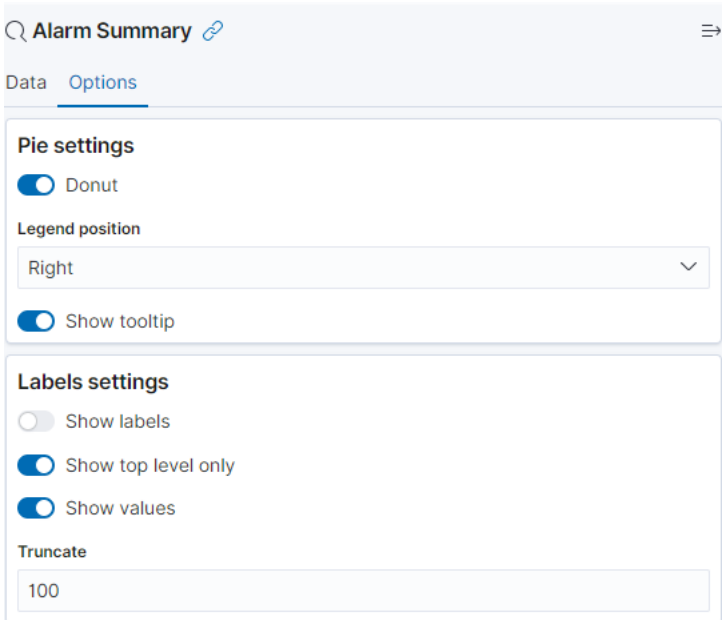
3. Click **Data.** Configure the Metrics and Buckets as shown in the following figure.



- **Aggregation function**: Select Unique Count. This returns the number of unique values in a field.
- **Field values**: _Id and Type can be changed as per your requirements
  - **Field value _Id**: The field that you want to visualize (in the example, id is being used because it is unique for each node)
  - **Field value Type**: The Field that you want to use to split the pie chart (in the example, type is being used to slice the chart)

4. Click **Update** to update the visualization.

NOTE: Use the **Options** tab to configure the required visual effects such as configuring the pie chart as a donut, adjusting the position of the legend, and so on.

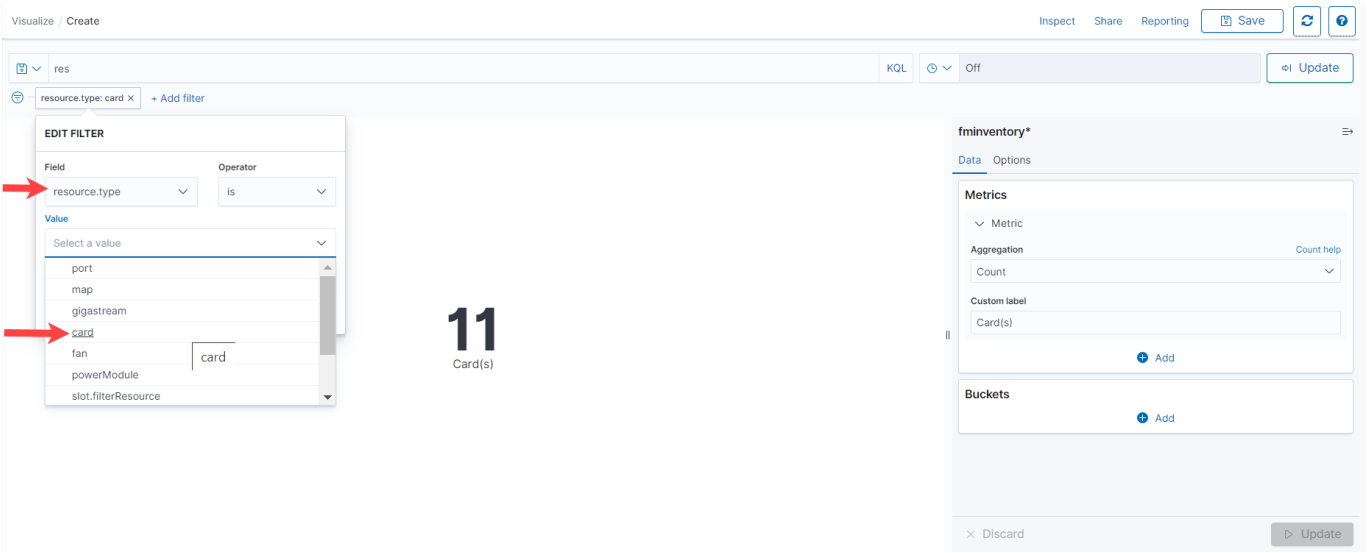


Metric Displaying Card Count

To create a Metric visualization that shows the number of cards:

1. Click **Create Visualization**. In the New Visualization page, select **Metric.**
2. Select the data source for the visualization. It can be either an index pattern or a saved search object.
3. Configure the **Aggregation** as **Count**.
4. In the Add Filer option, configure the following:
   - **Filter: *resource.type***
   - **Operator: is**
   - **Value: *Card***
5. Click **Update.**

Refer to the following image:

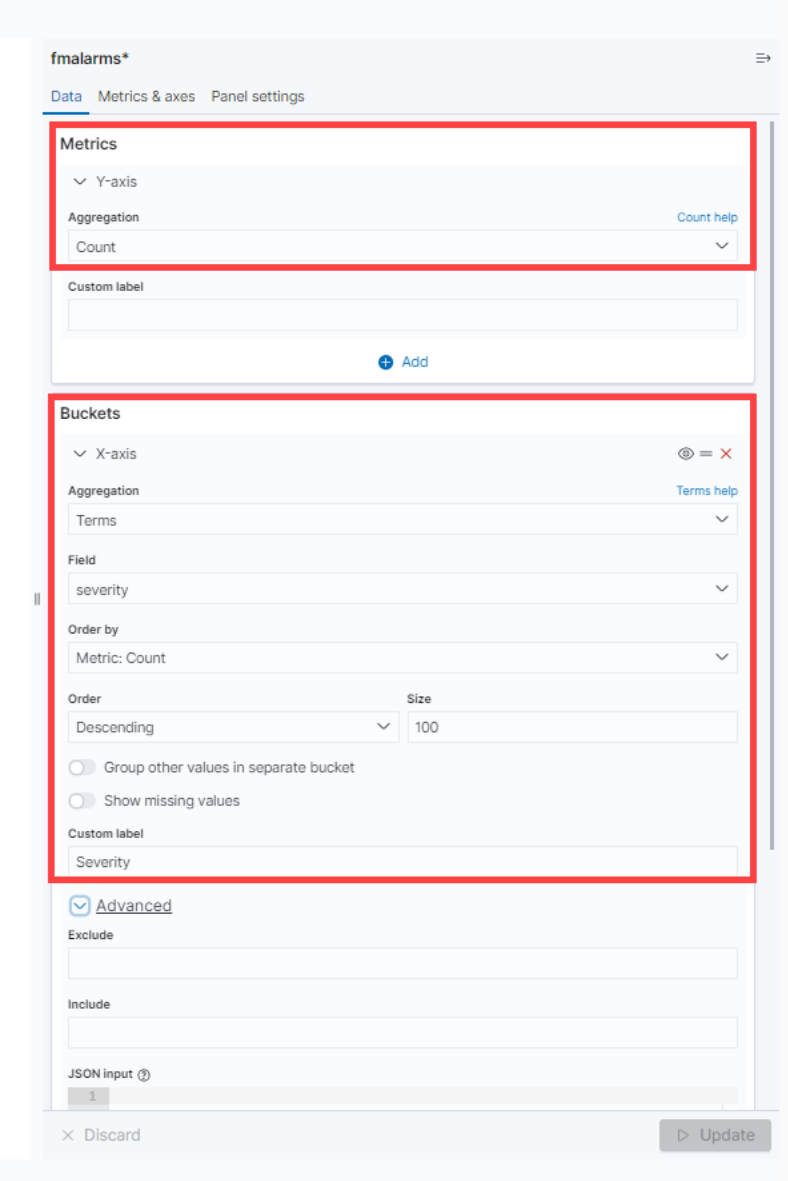


NOTE: Use the **Options** tab to configure the required visual effects such as adjusting the font size.

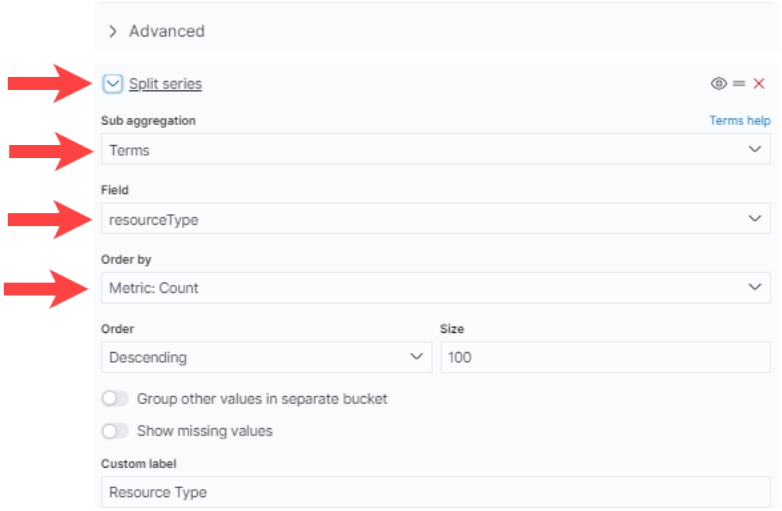
Bar Chart for Alarms by Severity

To create a bar chart that shows the number of alarms based on severity:

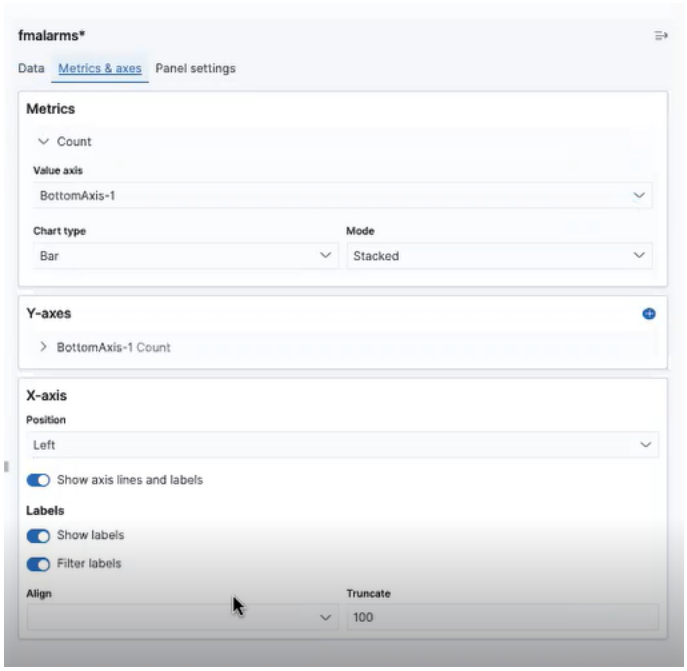
1. Click **Create Visualization**. In the New Visualization page, select **Vertical Bar.**
2. Select the data source for the visualization. It can be either an index pattern or a saved search object. In this example, *fmalarms* is selected as the data source.
3. Click **Data.** Configure the **Metrics** and **Buckets** as shown in the following figure.
   - Metrics
     - *Aggregation: Count*
   - Bucket
     - *Aggregation: Terms*
     - *Field: Severity*
     - *Order by: Metric Count*



4. Under **Advanced**, select **Split series** and configure the following:
   - *Sub aggregation: Terms*
   - *Field: resourceType*
   - *Order by: Metric Count*



5. Under Metrics & Axes, configure the following:



6. Click **Update**.

Dashboards

A dashboard is a collection of visualizations. Click on **Analytics** > **Dashboard**. The Alarms dashboard page appears.

Click the **Dashboards** menu on the top navigation bar of the Analytics page to view all the dashboards.

From the Dashboards page, you can:

- View Default Dashboards
- Create Custom Dashboards by cloning existing system dashboards or creating new dashboards.

View Default Dashboards

The Default dashboards is the list of pre-defined dashboards created in GigaVUE-FM. Refer to the Default Dashboards section for the list of default dashboards and the associated visualizations.

Notes:

- You cannot edit or delete the default dashboards. However, you can create your own personalized dashboards as per your requirements.
- GigaVUE-FM displays default static description text for each of the default dashboards.

You can perform the following operations:

Share	Use to share the Dashboard page. The following options are available: **Embed Code**: To share the code either as a snapshot or saved object. **Permalinks**: To copy the permalink of the dashboard page.
Clone	Use to clone the default dashboard page. Refer to the Clone Dashboard for more details
Reporting	Use to generate the report in PDF or in PNG format.

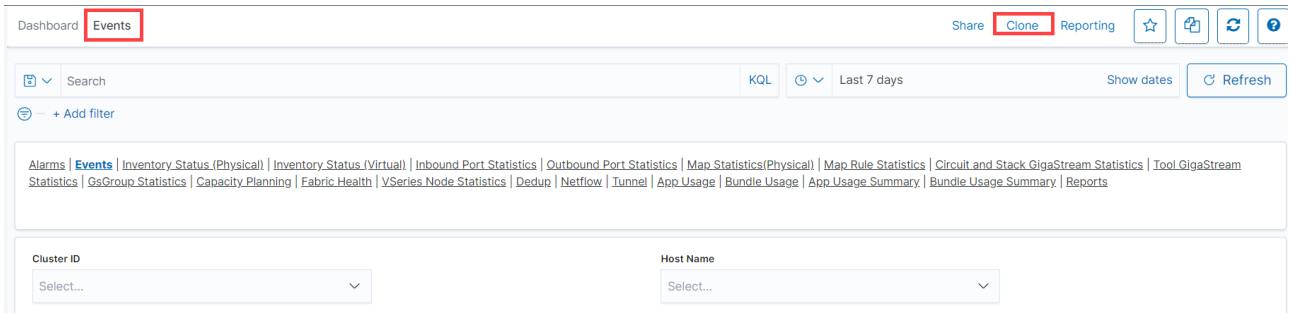
Clone Dashboard

GigaVUE-FM does not allow you to edit the default dashboards. However, you can clone the system dashboard and make changes to the new dashboard.

To clone a dashboard:

1. On the left navigation pane, click on  . Select **Analytics** > **Dashboards**.
2. Navigate to the dashboard page for which you need to create a clone.
3. Click the **Clone** button on the submenu bar.
4. Enter a name for the new dashboard and click **Confirm Clone.** The new dashboard is

created.

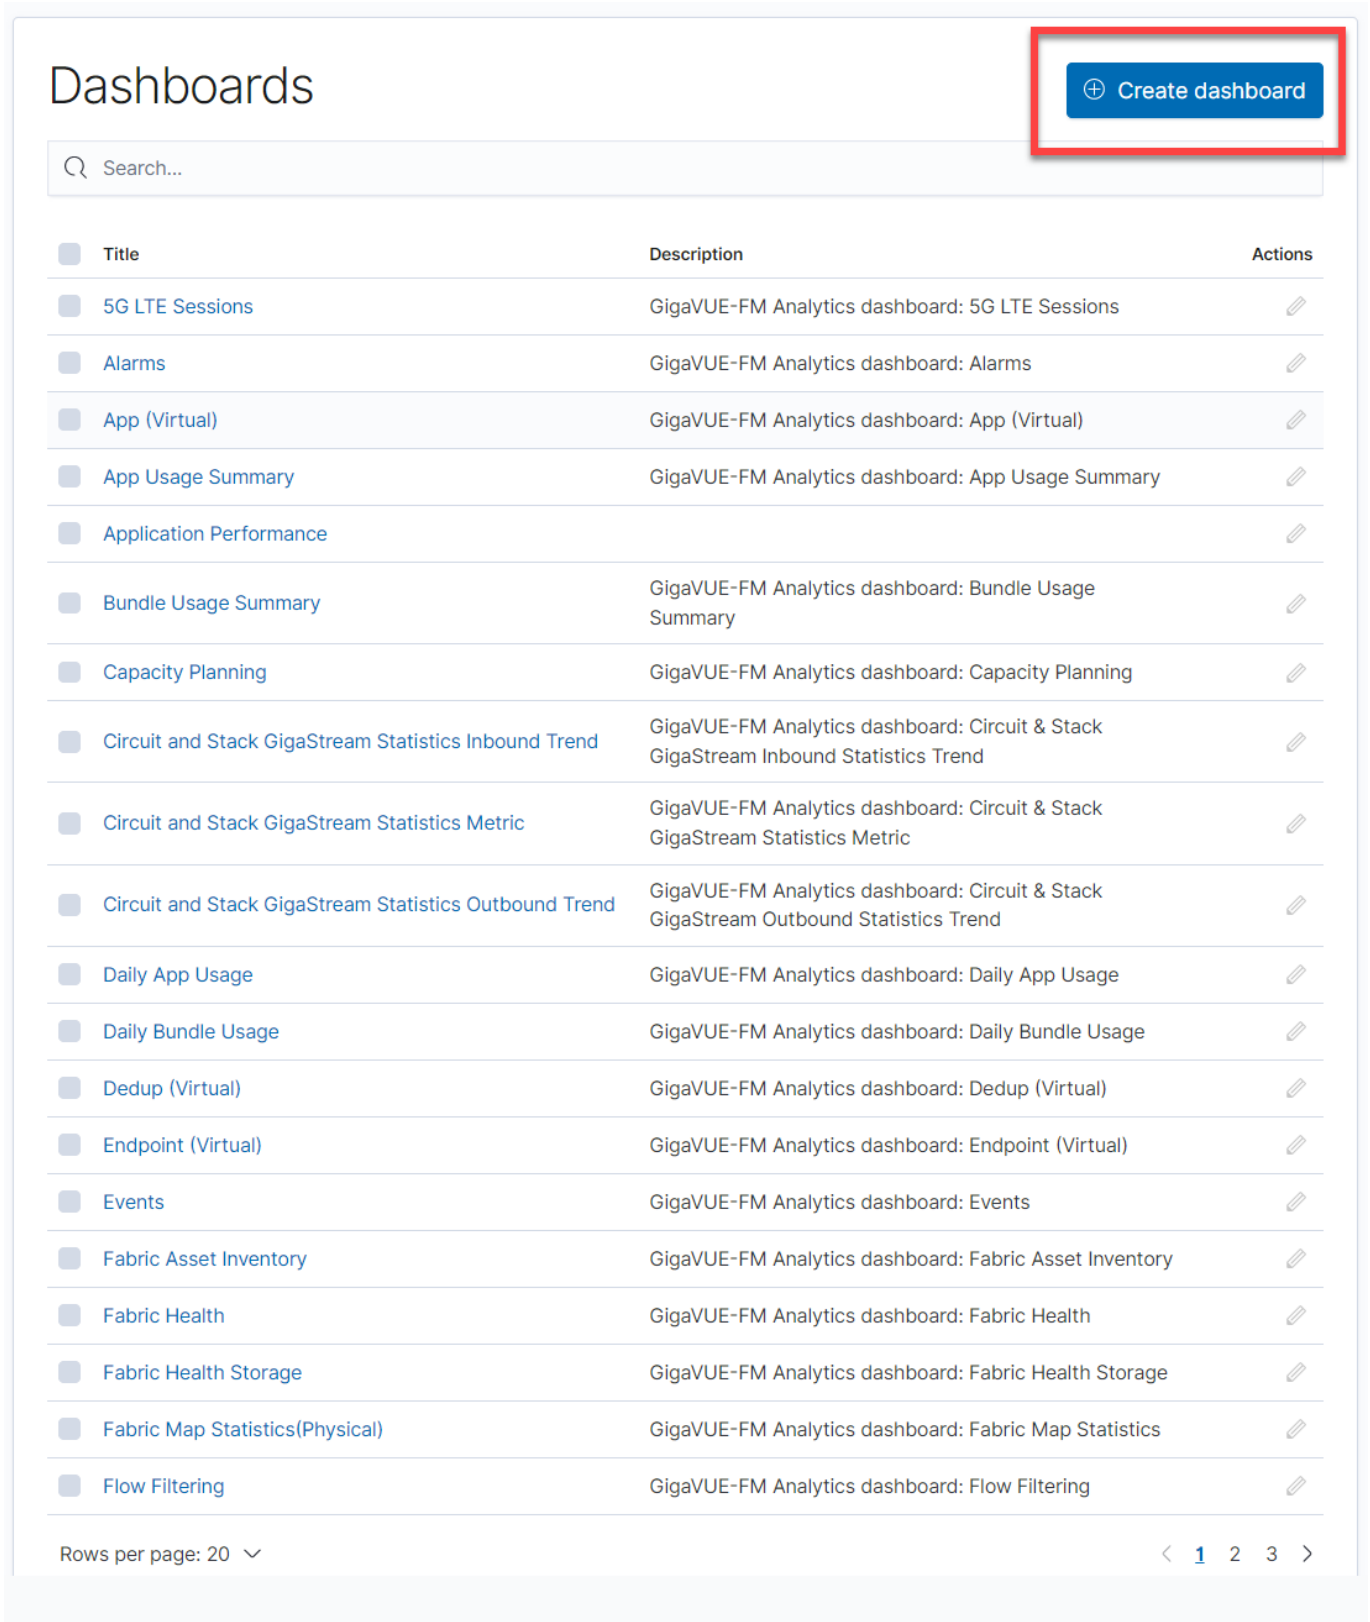


Refer to the Edit Dashboard section for details on editing the dashboard.

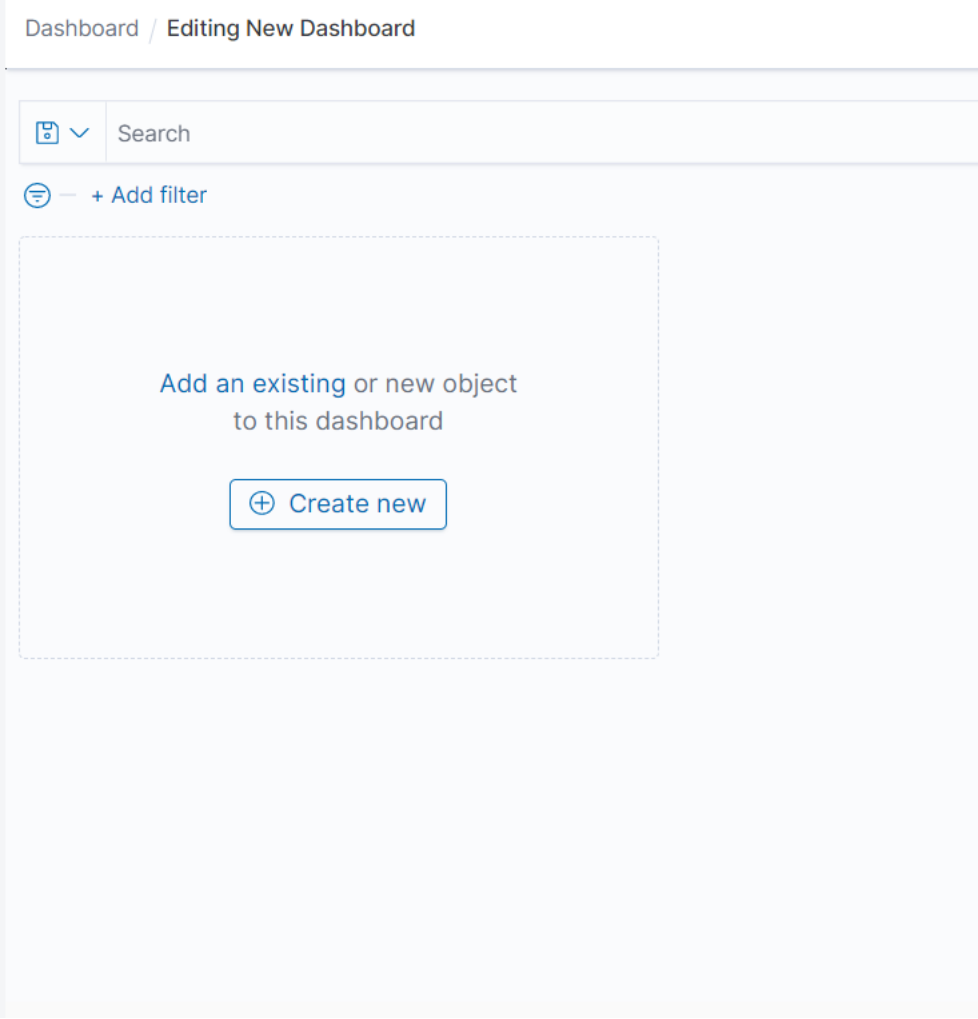
Create New Dashboard

To create a new dashboard:

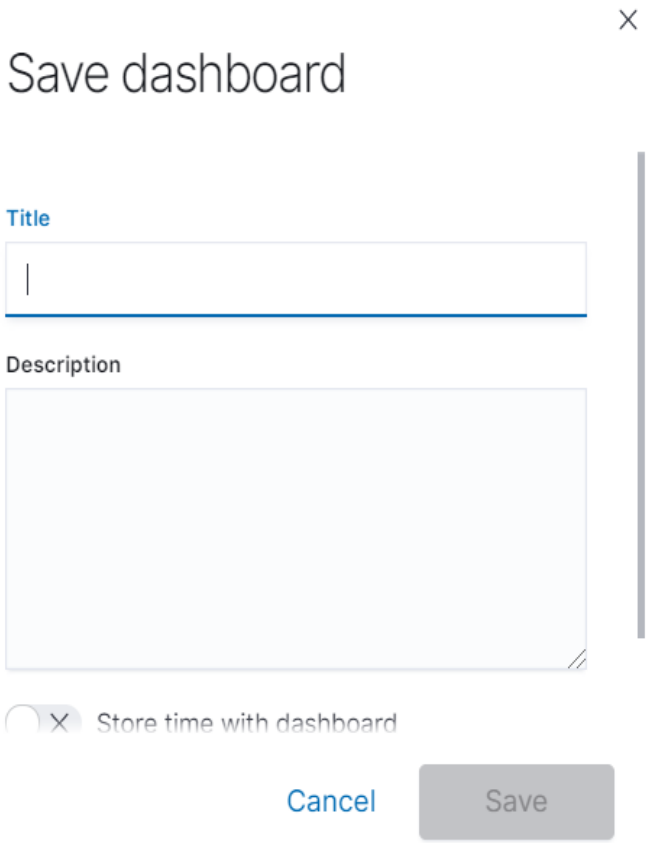
1. Go to Dashboards -> **Analytics** -> **Dashboards**.
2. Click the **Dashboards** menu on the top navigation bar of the Analytics page.
3. Click **Create Dashboard**.


Dashboards
Create dashboard
Search...
Title
Description
Actions
5G LTE Sessions
GigaVUE-FM Analytics dashboard: 5G LTE Sessions
Alarms
GigaVUE-FM Analytics dashboard: Alarms
App (Virtual)
GigaVUE-FM Analytics dashboard: App (Virtual)
App Usage Summary
GigaVUE-FM Analytics dashboard: App Usage Summary
Application Performance
Bundle Usage Summary
GigaVUE-FM Analytics dashboard: Bundle Usage Summary
Capacity Planning
GigaVUE-FM Analytics dashboard: Capacity Planning
Circuit and Stack GigaStream Statistics Inbound Trend
GigaVUE-FM Analytics dashboard: Circuit & Stack GigaStream Inbound Statistics Trend
Circuit and Stack GigaStream Statistics Metric
GigaVUE-FM Analytics dashboard: Circuit & Stack GigaStream Statistics Metric
Circuit and Stack GigaStream Statistics Outbound Trend
GigaVUE-FM Analytics dashboard: Circuit & Stack GigaStream Outbound Statistics Trend
Daily App Usage
GigaVUE-FM Analytics dashboard: Daily App Usage
Daily Bundle Usage
GigaVUE-FM Analytics dashboard: Daily Bundle Usage
Dedup (Virtual)
GigaVUE-FM Analytics dashboard: Dedup (Virtual)
Endpoint (Virtual)
GigaVUE-FM Analytics dashboard: Endpoint (Virtual)
Events
GigaVUE-FM Analytics dashboard: Events
Fabric Asset Inventory
GigaVUE-FM Analytics dashboard: Fabric Asset Inventory
Fabric Health
GigaVUE-FM Analytics dashboard: Fabric Health
Fabric Health Storage
GigaVUE-FM Analytics dashboard: Fabric Health Storage
Fabric Map Statistics(Physical)
GigaVUE-FM Analytics dashboard: Fabric Map Statistics
Flow Filtering
GigaVUE-FM Analytics dashboard: Flow Filtering
Rows per page: 20
1 2 3

4. In the Editing New Dashboard page, you can:
   a. Add an existing visualization. Click **Add an Existing** link.
   b. Create a new object. Click **Create New.** For instructions, refer to Create Custom Visualizations



5. Click **Save**.
6. In the **Save dashboard** dialog box, enter the **Title** and **Description** for the dashboard and click **Save**.



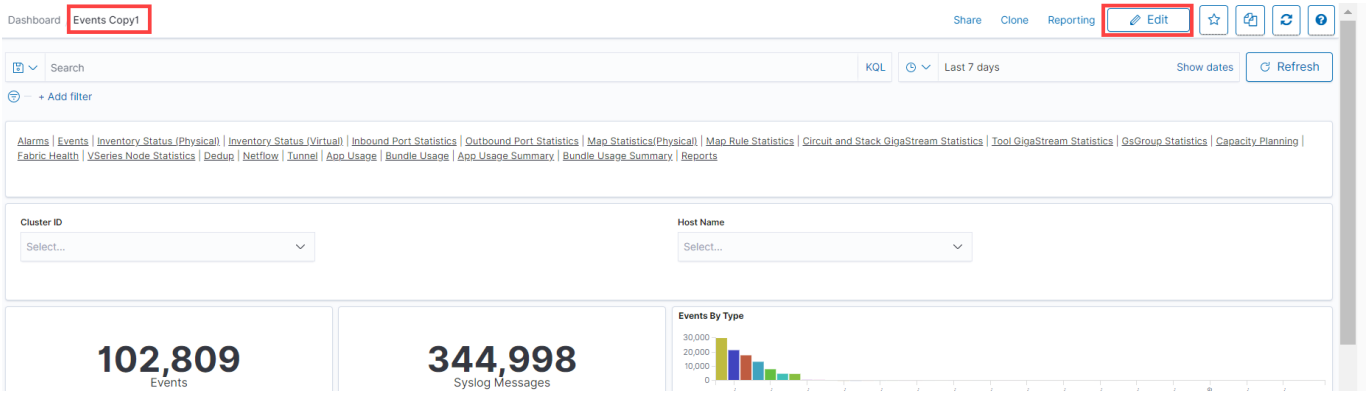
To make further changes to the dashboard, refer to the Edit Dashboard section.

Edit Dashboard

Edit the new dashboard page to suit your requirements. From the cloned dashboard page, click **Edit** to perform the following operations:

Option	Description
Options	Use to set the following options: • Use margins between panels • Show panel titles

Add	Use to Add Panels to the dashboard
Create New	Use to create a new visualization
⚙	Use to edit the following: • **Edit Visualization**: Edit the required visualizations. • **Clone Panel**: Clone the panel. • **Edit Panel Title**: Edit the panel title. • **Maximize Panel**: Maximize the panel. • **Replace Panel** • **Delete from Dashboard**: Delete visualizations that you no longer need on the new dashboard



Click **Save** to save the changes to the dashboard. In the **Save Dashboard** dialog box, use the toggle option to save the changes to a new dashboard. Click **Cancel** to discard the changes.

Reports

The Reports option allows you to download the data in the dashboards and visualizations in PDF or PNG format.

NOTE: You can download data in .csv format from the Discover page.[1]

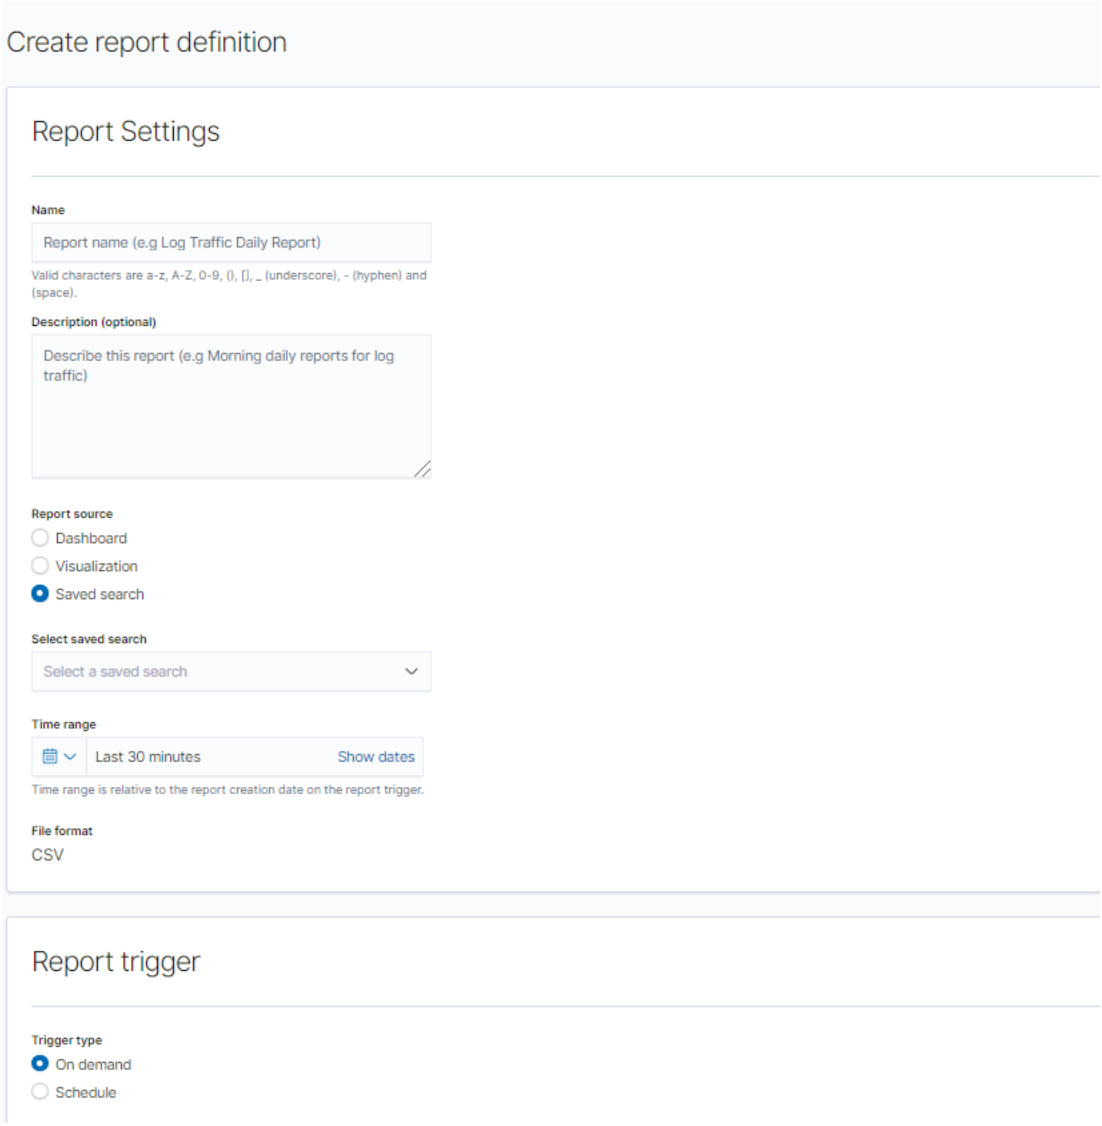
To download the reports

1. On the left navigation pane, click on  .
2. Select **Analytics** and click **Reports.** The list of reports is displayed**.** It can be either **On Demand** or **Schedule**.
3. Click on a report to view the details and download the report.

[1]The CSV file will have empty fields if the field values are zero in the report.

The **Report Definitions** option allows you to schedule automatic generation of reports. To create report definition:

1. Click **Create**.



2. Select or enter the following details under **Report Settings**:

Field	Description
Name	Name of the report.
Description	Description for the report.
Report Source	The source from which the report is generated. It can be Dashboard, Visualization or a Saved Search object. **Note**: When downloading reports of any dashboard, the parent

	dashboard link of the report will not get highlighted in the report (only the child dashboard link - such as metric/trend will not be highlighted).
Select	Select your Dashboard, Visualization or the Saved Search object, accordingly. Dashboards and Visualizations will be downloaded in PNG or PDF format. Saved search objects will be downloaded in CSV format.
Time Range	Select the time range for your report.
File Format	Select the required file format.
Header and Footer	Add a header or footer for the report. Headers and footers are only available for dashboard or visualization reports.

3. Select or enter the following details under **Report trigger**:
   - **Trigger type**
     - On Demand
     - Schedule
   - **Request time**
     - **Recurring**: Select **Frequency** and the **Request time**.
     - **Cron-based**: Select the Custom cron expression.
     - Select the required time zone.

     NOTE: GigVUE-FM allows you to create a report with recurring or cron-based options but does not allow you to edit these options after creating the report. In such cases, you must delete the report and create a new report using the required option.
4. Click **Create.**

Discover

The Discover page allows you to view and explore your data. This page consists of the following sections:

- **Add filters**: Use to create queries and filters. Click **Add filters** to add a filter. You can also use time filters along with the filter created. Use the saved filters in dashboards and visualizations. Refer to the Discover section for details.
- **Date Histogram**: Displays how data is ingested over time.
- **Documents**: Displays the documents. Expand the documents to view more details.
- **Field list on the left**: Displays fields available in the data. Click on a field to view the most common values.

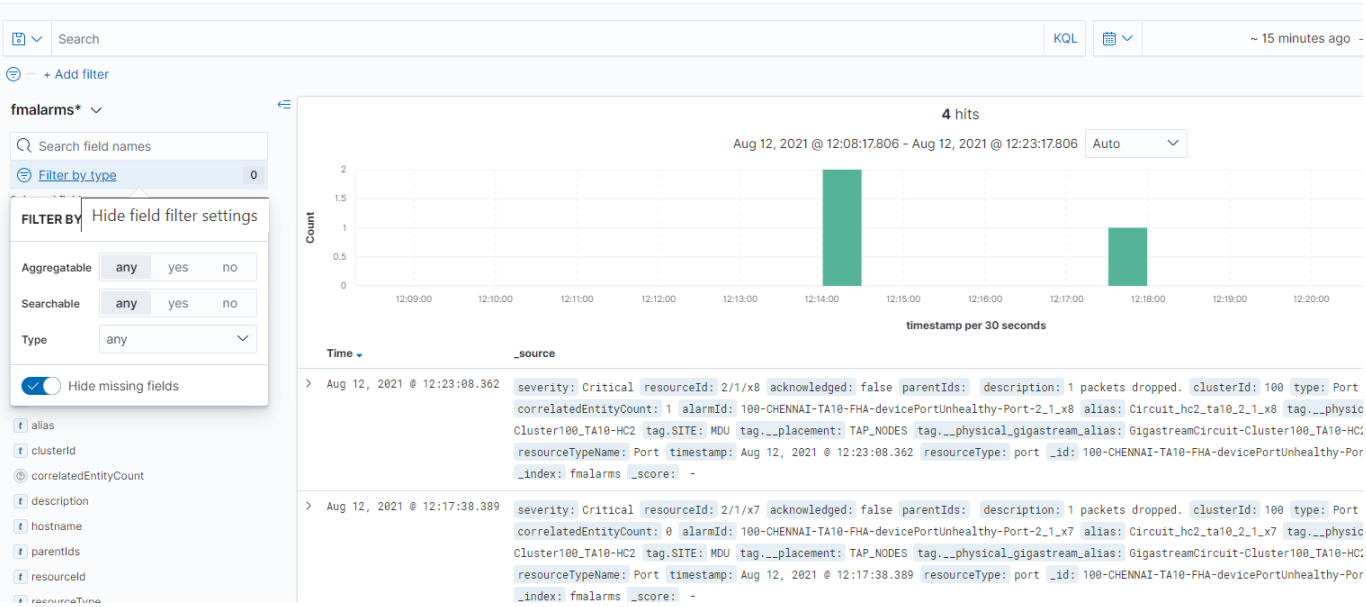
The **Filter by Type** option allows you to filter the data based on the following criteria:

Aggregatable Select Yes to extract summaries from matching documents. For example, count is a type of aggregation.

Searchable Select Yes to filter the data based on specific conditions. For example, filter the data for the last 24 hours.

Type Field type. Allowable values are:

- String
- Number
- _Source
- Date



You can perform the following operations from the Discover page:

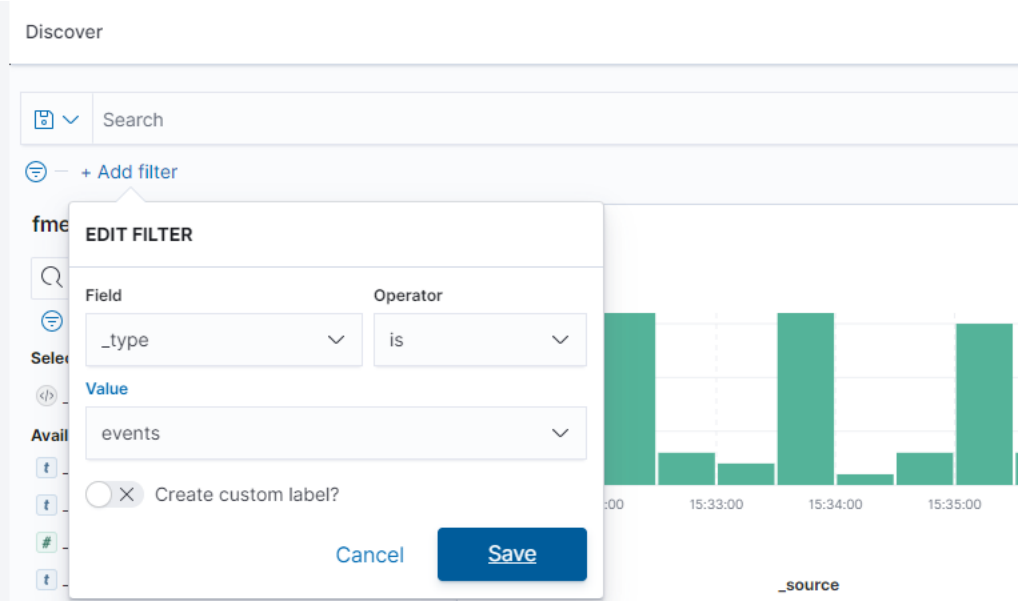
New	Use to create a new saved search object.
Save	Use to save your search and use it later. You can also generate a CSV report or use the saved search object in dashboards and visualizations. Refer to Save Search section for details.
Reporting	Use to generate and download the report in CSV format.
Open	Use to open the list of existing saved search objects.
Share	Use to share the saved search object to other users.
Inspect	Use to view details such as number of hits, index pattern, index pattern id, request and response details.

NOTE: Currently, the Discover page in GigaVUE-FM does not support the search and filtering option on RollUp Index..

Filter

To add a filter:

1. Click **Add filter**. The **Edit Filter** pop-up appears.
2. Select the required field, the operator, and the value.
3. Click **Save** to save the filter.
4. You can create custom label for the filter.



Save Search

Use the **Save Search** option to save queries, filters, and current view of the Discover page, such as the columns selected in the Document table, the sort order and also the index pattern. To save a search:

1. Create a search criteria that you want to reuse, click **Save** in the toolbar.
2. Enter a name for the search and click **Save**.
3. Use the saved search objects in the dashboards and visualizations by selecting the search objects using the Add from library option.

Find Data

Use the Discover page to find the data you need to analyze. You can also specify the time range in which to view that data:

1. On the left navigation pane, click on  .
2. Select **Analytics** and click **Discover.**
3. Select the index patterns for which you want to find the data. For example, *fmalarms*.
4. Adjust the time range to view the data for the required time range.

   NOTE: The range selection is based on the default time field in your data. If the data does not have a time field, the range selection is not available.
5. To view the count of data for a given time in the specified range, click and drag the mouse over the histogram.

   NOTE: After you upgrade or install GigaVUE-FM, when you query the *fmstats_port* index, additional device statistical data from *fmstats* index is also combined and displayed. This behavior is observed only on day 1 after installation or upgrade of GigaVUE-FM, and will not be noticed from the subsequent day onwards.

Default Dashboards

The following table lists the various default dashboard pages. The default dashboards are categorized as follows:

- System Dashboards
- Physical Dashboards
- Cloud Dashboards
- Hybrid Dashboards
- Inline TLS/SSL Dashboards

NOTE: Few dashboard pages have Control Visualizations that help you narrow down the data displayed in the visualizations based on the selected criteria.

Table 1: System Dashboards.

Dashboard	Details	Visualizations
System Dashboards		
Alarms	Displays data related to Alarms. Alarms is the default dashboard page. Clicking on a legend in the following visualizations in the Alarms dashboard navigates you to the Alarms page. The alarms are listed based on the filters in the Analytics page: • Alarms By Resource Type • Acknowledged vs Unacknowledged Alarms • Alarms by Severity • Unsuppressed vs Suppressed Alarms by Resource Type Use the following fields in control visualizations to filter the Syslog data: • **Tag Sample - Alarms** • **Cluster ID**	*Alarms by Resource Type* *Alarms by Severity* *Acknowledged vs. Unacknowledged Alarms* *Alarms Raised vs Cleared by Resource Type* *Suppressed Alarms Count* *Suppressed Alarms by Nodes* *Unsuppressed vs Suppressed Alarms by Resource Type*
Events	Displays data related to Events. Clicking on a legend in the following visualizations in the Events dashboard will navigate you to the Events page. The events will be listed based on the filters in the Analytics page: • Events by Severity • Events by Type • Top 10 Event Contributors Use the following fields in control visualizations to filter the Events data: • **Cluster ID** • **Host Name**	*Events by Type* *Events by Severity* *Top 10 Event Contributors*
Syslog	Displays data related to Syslog. Use the following fields in control visualizations to filter the Syslog data: • **Syslog Source** • **Host Name** • **Severity**	*Syslog by Severity* *Top 10 Syslog Contributors*

Dashboard	Details	Visualizations
Fabric Health	The Fabric Health tab consists of the following two dashboards: • CPU and Memory • Storage	
	CPU and Memory The CPU and Memory dashboard displays metric and trend visualizations for both threshold and maximum values. **Note**: After upgrading to software version 5.16.00, the line that depicts the threshold values (blue line) in the following visualizations will not be available for historic time ranges. This is because historical data for threshold is not collected in GigaVUE-FM prior to the upgrade: • Average Memory Usage Trend • Max Memory Usage Trend Use the **Server Name** control visualization to view the memory and storage metric and trend for the particular server.	*Average CPU Usage Gauge* *Max CPU Usage Gauge* *Average CPU Usage Trend* *Max CPU Usage Trend* *Average Memory Usage Gauge* *Max Memory Usage Gauge* *Average Memory Usage Trend* *Max Memory Usage Trend* *Used Vs Total System Memory*
	Storage The Storage dashboard displays metric and trend visualizations for both threshold and maximum values. **Note**: After upgrading to software version 5.16.00, the following visualizations will not be available for historical time ranges. This is because historical data for these visualizations is not collected in GigaVUE-FM prior to the upgrade: • Average Dis Usage (/) • Max Disk Usage (/) • Average Disk Usage (/) Trend • Max Disk Usage (/) Trend Use the Server Control Filter to view the memory and storage metric and trend for the particular server.	*Average Disk Usage (/var)* *Max Disk Usage (/var)* *Average Disk Usage (/var) Trend* *Max Disk Usage (/var) Trend* *Average Disk Usage (/config)* *Max Disk Usage (/config)* *Average Disk Usage (/config) Trend* *Max Disk Usage (/config) Trend* *Average Disk Usage (/)* *Max Disk Usage (/)*

Dashboard	Details	Visualizations
		Average Disk Usage (/) Trend
		Max Disk Usage (/) Trend
Physical Dashboards		
Inventory Status	Displays status of the physical resources. The following metrics are displayed at the top: • Number of clusters • Number of standalone nodes • Number of Nodes • Number of Ports • Number of Cards Use the **Tag Sample - Inventory** control visualizations to filter the data.	*Nodes by model and software version* *Port by type and health* *Card by type and health*
Inbound Port Statistics	Displays statistics of the receiving (Rx) ports in packets per second, bits per second. The dashboards are categorized into: • Metric • Trend The metric tab displays the following visualizations: • Total Port Capacity • Rx Ports Aggregated Average Traffic • Rx Ports Aggregated Average Drop Traffic • Rx Ports Aggregated Average Discard Traffic • Rx Ports Aggregated Average Error Traffic Use the following control visualizations to filter the data: • **Tag Sample - Statistics** • **Cluster-ID** • **Host Name** • **Port Type** • **Port ID** • **Port Alias** With Rollup enabled, you can view statistics for a period of 120 days on hourly granularity.	*Top Rx ports by Max Rate (bps)* *Top Rx ports by Average Rate (bps)* *Rx Port Statistics (bps)* *Rx Port Statistics(pps)* *Top Rx ports by Max Traffic Trend (bps)* *Top Rx ports by Max Traffic Trend (pps)* *Top Rx ports by Average Traffic Trend (bps)* *Top Rx ports by Average Traffic Trend (pps)* *Rx Drop Rate(pps)* *Rx Discard Rate(pps)* *Rx Error Rate(pps)*
Outbound Port	Displays statistics of the transmitting (Tx) ports in packets per	*Top Tx ports by Max*

Dashboard	Details	Visualizations
Statistics	second, bits per second. The dashboards are categorized into: • Metric • Trend The metric tab displays the following visualizations: • Total Port Capacity • Tx Ports Aggregated Average Traffic • Tx Ports Aggregated Average Drop Traffic • Tx Ports Aggregated Average Discard Traffic • Tx Ports Aggregated Average Error Traffic • Tx Ports Aggregated Max Traffic Use the following control filters to filter the statistics: • **Tag Sample - Statistics** • **Cluster-ID** • **Host Name** • **Port Type** • **Port ID** • **Port Alias** You can view statistics for a period of 120 days on a hourly granularity.	*Rate (bps)* *Top Tx ports by Average Rate (bps)* *Tx Port Statistics (bps)* *Tx Port Statistics(pps)* *Top Tx ports by Max Traffic Trend (bps)* *Top Tx ports by Max Traffic Trend (pps)* *Top Tx ports by Average Traffic Trend (bps)* *Top Tx ports by Average Traffic Trend (pps)* *Tx Drop Rate(pps)* *Tx Discard Rate(pps)* *Tx Error Rate (pps)*
Map Statistics	Displays statistics of the Maps. The metric tab displays the following visualizations: • Maps Aggregated Average Traffic (bps) Maps Aggregated Average Traffic (pps) Use the following control visualizations to filter the data: • **Tag Sample - Statistics** • **Cluster-ID** • **Map Type** • **Map Alias**	*Top Maps by Avg Rate (bps)* *Top Maps by Avg Rate (pps)* *Map Statistics (bps)* *Map Statistics(pps)* *Top Map Average Traffic Trend (bps)* *Top Map Average Traffic Trend (pps)*
Map Rule Statistics	Displays statistical data related to the map rules and the associated traffic.	*Refer to Map Rule Statistics Dashboard for more details.*
Map Traffic Statistics	Displays statistics about traffic that has passed/dropped through the fabric to the destination(s). The dashboard is categorized into:	*Map Average/ Pass/Drop Traffic (bps)*

Dashboard	Details	Visualizations
	• bps • pps Use the following control visualizations to filter the data: • Tag • Cluster-Id (default control filter) • Map Type • Map Alias	*Map Average /Pass/Drop Traffic (pps)*
Fabric Map Statistics	Displays statistical data related to Fabric Maps. Use the following control visualizations to filter the data: • Tag Sample - Statistics • Fabric Map • Cluster ID to Destination .	*Fabric Map Destination Tool Traffic (bps)* *Fabric Map Destination Tool Traffic (pps)* *Top Fabric Map Max Tool Traffic Trend (bps)* *Top Fabric Map Max Tool Traffic Trend (pps)*
Circuit and Stack GigaStream Statistics	Displays the statistics of the circuit and stack GigaStream. The dashboards are categorized into: • Metric • Inbound Trend • Outbound Trend The metric tab displays the following visualizations: • Circuit and Stack GigaStream Total Capacity (bps) • Circuit and Stack GigaStream Total Capacity by member Ports • Circuit and Stack GigaStream Tx Aggregated Max • Circuit and Stack GigaStream Tx Aggregated Average • Circuit and Stack GigaStream Tx Aggregated Average Drop • Circuit and Stack GigaStream Tx Aggregated Average Discard • Circuit and Stack GigaStream Tx Aggregated Average Error • Circuit and Stack GigaStream Rx Aggregated Max • Circuit and Stack GigaStream Rx Aggregated Average • Circuit and Stack GigaStream Rx Aggregated Average Drop • Circuit and Stack GigaStream Rx Aggregated Average Discard • Circuit and Stack GigaStream Rx Aggregated Average Error Use the following control visualizations to filter the data: • Sample Tag	*Circuit and Stack GigaStream Tx Traffic Rate (bps)*

Dashboard	Details	Visualizations
	• Cluster ID • Host Name • GigaStream	

Dashboard	Details	Visualizations
		Circuit and Stack GigaStream Tx Traffic Rate(pps)
		Circuit and Stack GigaStream Maximum Tx Rate (bps) by member port Alias
		Circuit and Stack GigaStream Maximum Tx Rate (pps) by member port Alias
		Circuit and Stack GigaStream Average Tx Rate (bps) by member port Alias
		Circuit and Stack GigaStream Average Tx Rate (pps) by member port Alias
		Circuit and Stack GigaStream Tx packets Drop Rate (pps)
		Circuit and Stack GigaStream Tx packets Discard Rate (pps)
		Circuit and Stack GigaStream Tx packets Error Rate (pps)
		Circuit and Stack GigaStream Rx Traffic Rate (bps)
		Circuit and Stack GigaStream Rx Traffic Rate (pps)
		Circuit and Stack GigaStream Maximum Rx Rate (bps) by member port Alias
		Circuit and Stack

Dashboard	Details	Visualizations
		GigaStream Maximum Rx Rate (pps) by member port Alias
		Circuit and Stack GigaStream Average Rx Rate (bps) by member port Alias
		Circuit and Stack GigaStream Average Rx Rate (pps) by member port Alias
		Circuit and Stack GigaStream Rx packets Drop Rate (pps)
		Circuit and Stack GigaStream Rx packets Discard Rate (pps)
Tool GigaStream Statistics	Displays the statistics of the Tool GigaStream. The dashboards are categorized into: • Metric • Trend The metric tab displays the following visualizations: • Tool GigaStream Total Capacity ((bps) • Tool GigaStream Total Capacity by member Ports • Tool GigaStream Tx Aggregated Average (bps) • Tool GigaStream Tx Aggregated Average Drop Traffic Rate (pps) • Tool GigaStream Tx Aggregated Average Discard Traffic Rate (pps) • Tool GigaStream Tx Aggregated Average Error Traffic Rate (pps) • Tool GigaStream Tx Ports Aggregated Max (bps) Use the following control visualizations to filter the data: • Sample Tag • Cluster ID • Host Name • GigaStream	*Tool GigaStream Tx Traffic Rate (bps)*
		Tool GigaStream Tx Traffic Rate(pps)
		Tool GigaStream Maximum Tx Rate (bps) by member port Alias
		Tool GigaStream Average Tx Rate(pps) by member port Alias
		Tool GigaStream Average Tx Rate (bps) by member port Alias
		Tool GigaStream Maximum Tx Rate (pps) by member port Alias
		Tool GigaStream Tx packets Discard Rate (pps)
		Tool GigaStream Tx packets Drop Rate

Dashboard	Details	Visualizations
		(pps)
GsGroup Statistics	Displays the GigaSMART group statistics: • GsGroup Total Capacity (bps) • GsGroup Total Capacity by member Ports Use the following control visualizations to filter the data: • Sample Tag • Cluster ID • Host Name • GSGroup	*GsGroup Max Rx Rate (bps)*
		GsGroup Max Rx Rate (pps)
		GsGroup Average Rx Rate (bps)
		GsGroup Average Rx Rate (pps)
		GsGroup Max Rx Rate (bps) by GsEngine Member ports
		GsGroup Max Rx Rate (pps) by GsEngine Member ports
		GsGroup Avg Rx Rate (bps) by GsEngine Member ports
		GsGroup Avg Rx Rate (pps) by GsEngine Member ports
		GsGroup Drop Rate (pps)
		GsGroup Drop Rate by Member Port in Percentage
		GsGroup Packet Buffer Utilization in Percentage
		GsGroup CPU Utilization in Percentage
		Gsgroup Packet buffer and CPU utilization Alert
Fabric Asset Inventory	Displays details about the devices managed by GigaVUE-FM: • Devices • Cards • Fans • Power • Power Modules	

Dashboard	Details	Visualizations
	• Port SFPs Click the **Export: Formatted** link option to download the data in each of the visualizations in CSV format. You can also use the **Reporting** option in the top menu to download the report in PDF or PNG format. Use the following control visualizations to filter the data: • Tag Sample - Inventory • Cluster ID • Host Name	
Capacity Planning	Consists of the following dashboards: • Port • Filter Resources Use the following control visualizations to filter the data: • Tag Sample - Inventory • Cluster ID • Host Name • Port Type • Port ID • Port Alias	
	Port	*Port Capacity Distribution*
		Port Capacity Distribution - Cluster Id
		Port Summary
	Filter Resources	*Filter Resource - Map Rule*
		Filter Resource - App Filter
		Filter Resource - Port Filter
Fault Collector	The Fault Collector tab consists of the following two dashboards: • Port Flapping • GigaSMART Core Crash Use the following Control Visualizations to filter the data: • Sample Tag	

Dashboard	Details	Visualizations
	• Cluster ID • Host Name • Port ID • Port Alias	
	Port Flap Displays visualizations related to port flapping.	• *Port Flap Events* • *Port Flapping Histogram* • *Port Link State Changes* • *Port Flaps by SFP Part Number* • *Port Flaps by SFP Vendor Name* • *Port Flapping Report*
	GigaSMART Core Crash Displays visualizations related to GigaSMART application crash.	• *GigaSMART App Core Crash Events* • *GSApp Crashes Histogram* • *GSApp Crashes by Cluster ID* • *GSApp Crashes Report*
Reports	Provides launch point for various reports. Click the **Reports** link to view the reports. You can perform the following operations: • Download the data tables in CSV file format. • Export data using **Reporting** → **Generate PDF/PNG**. **NOTE:** The CSV files display only a maximum of 10,000 records. If the number of records exceed 10,000 the additional records do not get displayed in the report.	• *Inventory Reports* • *Performance and Utilization Reports* • *Fabric Asset Inventory (Physical) Reports*
Fault Collector	The Fault Collector tab consists of the following two dashboards: • Port Flapping • GigaSMART Core Crash Use the following Control Visualizations to filter the data: • Sample Tag • Cluster ID	

Dashboard	Details	Visualizations
	• Host Name • Port ID • Port Alias	
Cloud Dashboards	Displays the statistics of virtual resources.	***Refer to*** *Virtual Inventory Statistics and Cloud Applications Dashboard*
Hybrid Dashboards	Displays statistics of hybrid solutions . This tab has the following dashboard: • Secure Tunnels.	
	Secure Tunnels Displays visualizations related to Secure tunnels. Use the following control visualizations to filter data: • Environment • Server / Remote/ Listener Port • Server /Remote / Listener IP	• *Average TLS Tunnel Decap Packets* • *Average TLS Tunnel Encap Packets* • *Average TLS Tunnel Decap Errors* • *Average TLS Tunnel Encap Errors* • *Average TLS Tunnel Encap Packets per IP and port* • *Average TLS Connection Stats*
Inline SSL Dashboards	Displays statistics for Inline TLS/SSL sessions. This tab has the following dashboards: • Session Overall • Session Engine Overview • Traffic Insights • Session Insights • Session Table	

Dashboard	Details	Visualizations
	• Engine Diagnostics	
Session Overall	Displays overall count of TLS/SSL sessions intercepted by the node including decrypted, non-decrypted and non-SSL sessions. Use the following control visualizations to filter data: • Host Name • GsGroup Alias • GsEngine ID	• *Total Intercepted Sessions* • *Sessions Trend* • *Average Decryption Rate* • *Client TLS Version Trend* • *Server TLS Version Trend* • *Policy Based Intercepted Sessions* • *Intercepted Sessions By Policy Rules*
Session Engine Overview	Displays Inline TLS/SSL Session statistics per engine. Use the following control visualizations to filter data: • Host Name • GsGroup Alias • GsEngine ID	• *Session Rate Per Engine* • *Average Decryption Rate Per Engine* • *Average CPS Per Engine* • *Engine Metric Table* • *Average CPU Per Engine*
Traffic Insights	Displays the traffic details for each Inline TLS/SSL Session. Use the following control visualizations to filter data: • Host Name • GsGroup Alias • GsEngine ID	• *Client Throughput(bps)* • *Server Throughput(bps)* • *Overall Volume (Bytes)* • *Overall Decrypted Volume(Bytes)* • *Average & Peak value of CPU & CPS* • *Average CPU Per Engine* • *Max CPU Per Engine* • *Max CPS Per*

Dashboard	Details	Visualizations
		Engine • *CPU Trend Per Engine* • *CPS Trend Per Engine* • *CPS Trend & CPU Trend Correlation* • *Throughput Trend On Network* • *Throughput Trend On Tool*
Session Insights	Displays detailed information about the TLS/SSL session. Use the following control visualizations to filter data: • Host Name • GsEngine ID • Source IP • Destination IP	• *Decryption Status* • *SSL Mode* • *SSL State* • *Policy Match By Rules* • *TLS Version* • *Top URLs (Max 10)* • *Top URL Categories (Max 10)* • *Top Ciphers (Max 10)* • *Certificates By Type*

Dashboard	Details	Visualizations
Session Table	Displays the entire details of a TLS /SSL session in a tabular format. Use the following control visualizations to filter data: • Host Name • GsEngine ID • URL • Source IP • Destination IP • URL Category	• *Session Table* • *Session Table with Policy Debug*
Engine Diagnostics	Engine Diagnostics tab provides troubleshooting information of Inline TLS/SSL sessions intercepted by the engine. Use the following control visualizations to filter data: • Host Name • GsGroup Alias • GsEngine ID	• *Certificates verified in Cache* • *SSL Alerts*

Dashboard

When you first login to GigaVUE-FM, the Dashboard - Physical & Virtual page is displayed by default. You can navigate to Health Monitor Dashboards from the left navigation pane.

Table 2: Summary of Dashboard Navigation provides descriptions of the top navigation bar and the left navigation pane in GigaVUE-FM.

Table 2: Summary of Dashboard Navigation

Top Navigation Bar	Left Navigation Pane
Dashboard	Physical & Virtual —Displays the informational widgets for the physical and Virtual nodes.
	Health Monitor—Displays information about GigaVUE-FM such as CPU usage, amount of storage used and available.
Profiles	Allows you to create and view the profiles. A profile allows you to create a customized dashboard to monitor the physical and virtual nodes.

Dashboards for Volume-based Licenses Usage

Licensed GigaSMART applications, when running on a GigaVUE V Series node, generate usage statistics. In the Volume-based Licensing (VBL) scheme, a license entitles specific applications on your V Series nodes to use a specified amount of total data volume over the term of the license. GigaVUE-FM tracks the total amount of data processed by the various licensed applications and provides visibility into the actual amount of data, each licensed application is using on each node, and track the overuse if any.

In cloud environment:

- when a monitoring session is created and deployed, you can only use applications that are licensed at that point.
- When a license expires, you will be notified, along with a list of monitoring sessions that would be affected in the near future.
- When a license finally expires (and has not been renewed yet), the monitoring sessions using the corresponding license will be undeployed, but not deleted from the database.
- When a license is later renewed or newly imported, such undeployed monitoring sessions will be redeployed.

Using the Volume-based License Application Usage and Bundle Usage dashboards, you can plan for better utilization of the licenses. These dashboards work on the principles of Analytics and are listed together with other dashboards. These VBL dashboards include both summary and daily dashboard pages.

- **Summary usage dashboard**: Displays summary for each period of VBL usage
- **Daily usage dashboard:** Displays more detailed view (down to the granularity of one day) about the app and bundle usages.

NOTE: Clicking on a bar chart on the App or Bundle Usage Summary dashboards does not display any further information. To view the originally displayed visualization if clicked inadvertently as mentioned above, navigate to a different dashboard and return to the original dashboard.

To access the dashboards:

1. Go to  -> **Analytics -> Dashboards.**
2. Click on the required dashboard to view the visualizations.

Table 3:

Dashboard	Visualizations	Description
Daily App Usage	Licensed App Allowance vs Usage	Displays details about the daily usage of each of the licensed application against the allowance provided and the overage. By default, it is shown for a 90-day time period. However, you can change the interval to the required time period.
	Aggregate Summary	Displays the following aggregation statistics: • Highest Daily Usage • Average Daily Usage • 95th Percentile Daily Usage[1]. • Highest Daily Overage • Average Daily Overage • Average Daily Allowance
Daily Bundle Usage	VBL Bundle Usage	Displays details about the usage of bundled license against the allowance and the overage. The bundle can be coreVUE or netVUE.

[1]The 95th percentile daily usage for any day is calculated as follows: the daily usage for the trailing 90 days, including and up to the current day, are sorted in ascending order and the usage at the 95th percentile (near the high end) is reported as the 95th percentile usage for the day. In the daily usage widgets, the aggregate statistics uses the maximum of the 95th percentile usage for the days selected as per the Time Filter. The 95th percentile usage statistic allows the user to disregard exceptionally high values of usage (might have occurred due to extraordinary conditions) which do not represent normal high values.

Dashboard	Visualizations	Description
	Aggregate Summary	Displays the following aggregation statistics: • Highest Daily Usage • Average Daily Usage • 95th Percentile Daily Usage • Highest Daily Overage • Average Daily Overage • Average Daily Allowance
App Usage Summary	• Usage (all applications) • Overage (all applications) • Summary per period (Incoming traffic) • Data Usage vs Overage (Incoming traffic)	The **Usage Period** drop-down option at the top allows you to choose the period for which you want to view the usage details. The duration of each period is 3 months. The following visualizations are displayed for the selected period: • **Usage (all applications)**: Displays breakdown of usage of all your applications • **Overage all applications**: Displays breakdown of overage of all your applications • **Summary per period (Incoming traffic)**: Displays a tabular view of the license usage/overage summary for the selected period, considering the incoming data traffic (before being processed by the Gigamon applications), • **Data Usage vs Overage (Incoming traffic):** Displays a bar chart of the license usage vs. overage summary for the selected period. If you do not select the Usage Period, the aggregation of all periods' data is displayed in the top visualizations, and the summary for each of the periods is displayed in the bottom visualizations in the dashboard.
Bundle Usage Summary	• Summary per Period • Bundle Usage vs Overage (Incoming traffic)	The Bundles drop-down allows you to choose the bundle for which you want to see the usage details. • **Summary per Period:** Displays a tabular view of the bundle usage summary for each period. Days with overage within each period are displayed within this tabular view. The tabular view also includes the total licensed data allowance for the days already elapsed in each period. • **Bundle Usage vs Overage (Incoming traffic):** Displays a bar chart of the license usage vs. overage summary for the selected bundle.

Dashboard	Visualizations	Description
		If you do not select any bundle, the summary views include the sum of metric values for all bundles that were active during a period.

Refer to the Analytics section for details on how to clone a dashboard, create a new visualization, and other detailed information.

Virtual Inventory Statistics and Cloud Applications Dashboard

Analytics dashboards allow users to monitor the physical and virtual environment and detect anomalous behavior and plan accordingly. In case of virtual environment, Analytics support is available for the following cloud platforms:

- AWS
- OpenStack
- VMware ESXi
- Azure

To access the dashboards:

1. Go to  -> **Analytics -> Dashboards.**
2. Click on the required dashboard to view the visualizations.

The following table lists the various virtual dashboards:

Dashboard	Displays	*Visualizations*	Displays
Inventory Status (Virtual)	Statistical details of the virtual inventory based on the platform and the health status. You can view the following metric details at the top of the dashboard: • Number of Monitoring Sessions • Number of V Series Nodes • Number of Connections • Number of GCB Nodes You can filter the visualizations based on the following control filters: • Platform • Health Status	*V Series Node Status by Platform*	Number of healthy and unhealthy GigaVUE V Series Nodes for each of the supported cloud platforms.
		Monitoring Session Status by Platform	Number of healthy and unhealthy monitoring sessions for each of the supported cloud platforms
		Connection Status by Platform	Number of healthy and unhealthy connections for each of the supported cloud platforms
		GCB Node Status by Platform	Number of healthy and unhealthy GCB nodes for each of the supported cloud platforms
V Series Node Statistics	Displays the Statistics of the V Series node such as the CPU usage, trend of the receiving and transmitting packets of the V Series node. You can filter the visualizations based on the following control filters: • Platform • Connection • V Series Node	*V Series Node Maximum CPU Usage Trend*	Line chart that displays maximum CPU usage trend of theGigaVUE V Series Node in 5 minutes interval, for the past one hour. **Note**: The maximum CPU Usage trend refers to the CPU usage for service cores only. Small form factor GigaVUE V Series Nodes do not have service cores, therefore the CPU usage is reported as 0.

Dashboard	Displays	*Visualizations*	Displays
		V Series Node with Most CPU Usage For Past 5 minutes	Line chart that displays Maximum CPU usage of the GigaVUE V Series Node for the past 5 minutes. **Note**: You cannot use the time based filter options to filter and visualize the data.
		V Series Node Rx Trend	Receiving trend of theGigaVUE V Series Node in 5 minutes interval, for the past one hour.
		V Series Network Interfaces with Most Rx for Past 5 mins	Total packets received by each of the V Series network interface for the past 5 minutes. **Note**: You cannot use the time based filter options to filter and visualize the data.
		V Series Node Tunnel Rx Packets/Errors	Displays the reception of packet at the Tunnel RX. This is the input toGigaVUE V Series Node, Grouping by tunnel identifier comprising {monDomain, conn, GigaVUE V Series Node, tunnelName}, before aggregation.
		V Series Node Tunnel Tx Packets/Errors	TX is for output tunnels from VSN. GigaVUE V Series Node Tunnel Tx

Dashboard	Displays	*Visualizations*	Displays
			Packets/Errors
Dedup	Displays visualizations related to De-duplication application. You can filter the visualizations based on the following control filters: • Platform • Connection • V Series Node	*Dedup Packets Detected/Dedup Packets Overload*	Statistics of the total de-duplicated packets received (ipV4Dup, ipV6Dup and nonIPDup) against the de-duplication application overload.
		Dedup Packets Detected/Dedup Packets Overload Percentage	Percentage of the de-duplicated packets received against the de-duplication application overload.
		Total Traffic In/Out Dedup	Total incoming traffic against total outgoing traffic
Tunnel (Virtual)	Displays visualizations related to the tunneled traffic in both bytes as well as the number of packets. You can select the following control filters, based on which the visualizations will get updated: • **Monitoring session**: Select the required monitoring session. The cloud platform, monitoring domain and connection within the monitoring domain that is used by the V Series node are shown in square brackets, comma-separated, after the name, to distinguish the whole path to it. • **V Series node**: Management IP of the V Series node. Choose the required V Series node from the drop-down. • **Tunnel:** Select any of the tunnels shown in the Tunnel drop-down. The direction for each tunnel is shown with the prefix in or out.	*Tunnel Bytes*	Displays received tunnel traffic vs transmitted tunnel traffic, in bytes. • For input tunnel, transmitted traffic is displayed as zero. • For output tunnel, received traffic is displayed as zero.

Dashboard	Displays	*Visualizations*	Displays
	The following statistics are displayed for the tunnel: • Received Bytes • Transmitted Bytes • Received Packets • Transmitted Packets • Received Errored Packets • Received Dropped Packets • Transmitted Errored Packets • Transmitted Dropped Packets		
		Tunnel Packets	Displays packet-level statistics for input and output tunnels that are part of a monitoring session.
App (Virtual)	Displays Byte and packet level statistics for the applications for the chosen monitoring session on the selected V Series node. You can select the following control filters, based on which the visualizations will get updated: • **Monitoring session** • **V Series node** • **Application**: Select the required application. By default, the visualizations displayed includes all the applications. By default, the following statistics are displayed: • Received Bytes • Transmitted Bytes • Received Packets • Transmitted Packets • Errored Packets • Dropped Packets	*App Bytes*	Displays received traffic vs transmitted traffic, in Bytes.

Dashboard	Displays	*Visualizations*	Displays
		App Packets	Displays received traffic vs transmitted traffic, as the number of packets.
End Point (Virtual)	Displays Byte and packet level statistics for the un-tunneled traffic deployed on the GigaVUE V Series Node. The following statistics that are shown for Endpoint (Virtual): • Received Bytes • Transmitted Bytes • Received Packets • Transmitted Packets • Received Errored Packets • Received Dropped Packets • Transmitted Errored Packets • Transmitted Dropped Packets The endpoint drop-down shows *<V Series Node Management IP address : Network Interface>* for each endpoint. You can select the following control filters, based on which the visualizations will get updated: • **Monitoring session** • **V Series node** • **Endpoint:** Management IP of the V Series node followed by the Network Interface (NIC)	*Endpoint Bytes*	Displays received traffic vs transmitted traffic, in Bytes.
		Endpoint Packets	Displays received traffic vs transmitted traffic, as the number of packets.

NOTE: The Tunnel (Virtual), App (Virtual) and Endpoint (Virtual) dashboards do not show data from the previous releases if the *Monitoring Session [Platform : Domain : Connection]* dashboard filter is applied. This is because, this filter relies on the new attributes in the Opensearch database, which are available only from software version 5.14.00 and beyond.

Rules and Notes

- You cannot edit or delete these default dashboards. However, you can clone the dashboards and visualizations. Refer to the Clone Dashboard section for more details.
- Use the Time Filter option to select the required time interval for which you need to view the visualization.
- Refer to the Analytics section for details on how to create a new dashboard, clone a dashboard, create a new visualization, and other information about the Discover page and Reports page.

FHA Dashboards for 5G Applications

Analytics dashboards allow users to monitor the physical and virtual environment and detect anomalous behavior and plan accordingly.

5G Applications dashboard provides comprehensive insights into the performance and statistics of GVHTTP2, 5G-Cloud, and 5G-SBI applications. Users can monitor Transaction Metrics, Queue Information, Request Metrics, URI Validation, and VXLAN Network Traffic using the 5G Apps dashboard. Analytics support is available for the following cloud platforms:

- VMware ESXi
- OpenStack
- Third Party Orchestration

To access the dashboards:

1. Go to  -> **Analytics -> Dashboards.**
2. Navigate to **System Dashboards ->5G Apps Dashboards**. Click on the **Overall 5G Apps** or **Detailed 5G Apps** to view the visualizations.

Refer to the following sections to view the various 5G Apps visualizations:

- Overall 5G Apps Dashboard
- Detailed 5G Apps Dashboard
- NokiaSCP Statistics
- OracleSCP Statistics
- EricssonSCP Statistics
- NokiaHEP3 Statistics
- Ericsson vTAP Statistics

- Ericsson Packet Processing Statistics

Overall 5G Apps Dashboard

Displays the overall visualization details of 5G Apps. To access the dashboards:

1. Select the required 5G-Cloud Monitoring Session Name and the GVHTTP2 Monitoring Session Name from the drop-down lists.
2. Select the required GigaVUE V Series Node names from the drop-down list.

Table 1: Overview of 5G Apps Dashboard – Metrics, Visualizations, and Transaction Details

Visualizations	Details
GVHTTP2 Transactions	Displays the cumulative count of Total HTTP2 transactions that occurred in the selected GVHTTP2 Servers, along with a breakdown of successful and failed transactions.
5G-Cloud Transactions	Displays the Total, Partial, and Failed Transactions that occurred in 5G-Cloud. For Nokia SCP, Total and Partial Transactions are not supported.
GVHTTP2 Queue	Displays the cumulative number of packets placed in or removed from the RX queue and the number of packets attempting to enter the RX Queue when it is full, and RX packets dropped due to queue limitations or excessive message size.
GVHTTP2 Stream Failures	Displays the cumulative count of total number of packet where headers or data frames were received but could not be processed because stream size exceeds maximum supported stream size, also total occurrences where a new HTTP2 streams which are timed out on the GVHTTP2 servers.
GVHTTP2 SSL Connection / Socket / nghttp2 library failures	Displays the cumulative count of total SSL connection failures and socket failures in either mTLS or clear text mode, and the total nghttp2 library errors in the GVHTTP2 servers.
GVHTTP2 Tx Socket Errors	Displays the cumulative count of the total number of packets that failed to transmit to the TX Tunnel.
GVHTTP2 Request Types	Displays the cumulative count of total HTTP2 requests received by the GVHTTP2 servers, categorized into GET, POST, PUT, and other types of requests.
GVHTTP2 URI's	Displays the cumulative count of the number of HTTP2 requests with Valid and Invalid paths for the current GVHTTP2 mode.
GVHTTP2 Tx Packets	Displays the count of the total number of packets successfully transmitted to the TX Tunnel. This visualization is not intended for comparative analysis.
5G-Cloud Rx Packets	Displays the cumulative count of the total received packets from the GVHTTP2 application. This visualization is not intended for comparative analysis.
5G-Cloud Rx Drops	Displays the total number of messages/packets drops due to queue

Visualizations	Details
	full and VNI ID mismatch.
5G-Cloud Transaction Flow	Displays the total number of transaction flows allocated and active, the number of times a new Transaction Flow was retrieved, and the number of times Transaction Flows were inserted.
5G-Cloud Transaction Flow Error	Displays the total number of errors encountered while attempting to retrieve new Transaction Flows, insert Transaction Flows into the active table, and age out Transaction Flows.
5G-Cloud TCP Flow	Displays the total number of TCP flows allocated and active, the total number of times a new TCP Flow was retrieved, and the total number of times TCP Flows were inserted into the Active TCP Flow table during the interval.
5G-Cloud TCP Flow Error	Displays the total number of failures when attempting to retrieve new TCP flows, errors when inserting TCP flows into the active table, and instances of TCP flows that have aged out.
5G-Cloud Tool Synthesized Requests & Responses	Displays the cumulative count of Mirrored Transaction HTTP/2 requests and responses synthesized towards the tool port output, including retransmissions.
5G-Cloud Tool Synthesized Packets	Displays the total cumulative count of Mirrored Transaction packets synthesized toward the tool port. This visualization is not intended for comparative analysis.
5G-Cloud Tx Packets	Displays the cumulative count of total transmitted packets sent to the tool. This visualization is not intended for comparative analysis.

Detailed 5G Apps Dashboard

Displays the detailed visualizations details of 5G Apps. To access the dashboards:

1. Select the required 5G-Cloud Monitoring Session Name and the GVHTTP2 Monitoring Session Name from the drop-down lists.
2. Select the required GigaVUE V Series Node names from the drop-down list.

Table 2: Detailed 5G Apps Dashboard – Visualization and Metrics Overview

Visualizations	Details
Total Transactions occurred per GVHTTP2	Displays the total number of HTTP2 transactions that occurred per GVHTTP2 Server.
Total Transactions Successful per GVHTTP2	Displays the total number of Successful HTTP2 transactions that occurred per GVHTTP2 Server.
Total Transactions Failed per GVHTTP2	Displays the total number of Failed HTTP2 transactions that occurred per GVHTTP2 Server.
5G-Cloud Transactions	Displays the Total, Partial, and Failed Transactions that occurred in 5G-Cloud. For Nokia SCP, Successful and Partial transactions are not applicable.

Visualizations	Details
Packets Queued In per GVHTTP2	Displays the number of packets in RX queue.
Packets Queued Out per GVHTTP2	Displays the number of packets dequeued from the RX queue.
Packets Queued Full per GVHTTP2	Displays the number of packets dropped due to RX queue being full.
Rx Packet Drops per GVHTTP2	Displays the number of RX packets dropped due to queue limitations.
Rx packets Oversized per GVHTTP2	Displays the number of RX packets that were dropped due to the greater message size.
Stream Header Overflows per GVHTTP2	Displays the total number of packets where headers were received but could not be processed because stream size exceeds maximum supported stream size on the GVHTTP2 server.
Stream Buffer Overflows per GVHTTP2s	Displays the total number of packets where data frames were received but could not be processed because stream size exceeds maximum supported stream size on the GVHTTP2 server.
Stream Timeouts per GVHTTP2	Displays the total occurrences where HTTP2 streams are timed out on the GVHTTP2 server.
Total SSL Connection Failures per GVHTTP2	Displays the total SSL connection failures in the GVHTTP2 server.
Total Socket Failures per GVHTTP2	Displays the total socket failures in mTLS or clear text mode in the GVHTTP2 server.
Total nghttp2 library errors per GVHTTP2	Displays the total nghttp2 library errors in the GVHTTP2 server.
TX Socket Send Errors per GVHTTP2	Displays the total number of packets that failed to transmit to the TX Tunnel.
TX Socket Size Errors per GVHTTP2	Displays the total number of packets that failed to transmit to the TX Tunnel due to the larger packet size.
Total Requests per GVHTTP2	Displays the total number of HTTP2 requests received by the GVHTTP2 Server.
GET Requests per GVHTTP2	Displays the total number of HTTP2 GET requests received by the GVHTTP2 Server.
POST Requests per GVHTTP2	Displays the total number of HTTP2 POST requests received by the GVHTTP2 Server.
PUT Requests per GVHTTP2	Displays the total number of HTTP2 PUT requests received by the GVHTTP2 Server.
Other Requests per GVHTTP2	Displays the total number of HTTP2 Other requests received by the GVHTTP2 Server.
Total URI's per GVHTTP2	Displays the number of URI's.
Requests with Valid URI's per GVHTTP2	Displays the number of HTTP2 requests with Valid paths.
Requests with Invalid URI's per GVHTTP2	Displays the number of HTTP2 requests with Invalid paths.
TX Packets per GVHTTP2	Displays the total number of packets successfully transmitted to the VXLAN TX Tunnel per GVHTTP2. This visualization is not intended for comparative analysis.

Visualizations	Details
5G-Cloud Rx Packets	Displays the cumulative count of the total packets received from the GVHTTP2 application. This visualization is not intended for comparative analysis.
5G-Cloud Rx Drops	Displays the total number of messages/packets drops due to queue full and VNI ID mismatch.
5G-Cloud Transaction Flow	Displays the total number of transaction flows allocated and active, the number of times a new Transaction Flow was retrieved, and the number of times Transaction Flows were inserted.
5G- Cloud Transaction Flow Error	Displays the total number of errors encountered while attempting to retrieve new Transaction Flows, insert Transaction Flows into the active table, and age out Transaction Flows.
5G-Cloud TCP Flow	Displays the total number of TCP flows allocated and active, the total number of times a new TCP Flow was retrieved, and the total number of times TCP Flows were inserted into the Active TCP Flow table during the interval.
5G-Cloud TCP Flow Error	Displays the total number of failures when attempting to retrieve new TCP flows, errors when inserting TCP flows into the active table, and instances of TCP flows aging out from the Active TCP Flow table during the interval.
5G-Cloud Tool Synthesized Requests & Responses	Displays the cumulative count of Mirrored Transaction HTTP/2 requests and responses synthesized towards the tool port output, including retransmissions.
5G-Cloud Tool Synthesized Packets	Displays the total cumulative count of Mirrored Transaction packets synthesized toward the tool port. This visualization is not intended for comparative analysis.
5G-Cloud Tx Packets	Displays the cumulative count of total transmitted packets sent to the tool. This visualization is not intended for comparative analysis.

NokiaSCP Statistics

Displays the visualization details of Nokia SCP statistics. To access the dashboards:

- Select the required 5G-Cloud Monitoring Session Name and the GigaVUE V Series Node names from the drop-down lists.

Table 3: NokiaSCP Statistics – Overview of Message Parsing and Processing Metrics

Visualizations	Details
Message Parsing	Displays the total number of messages received from the SCP and indicates whether they were parsed successfully.
Parsing Errors	Displays the errors encountered while parsing the message/frame in Nokia SCP.

Visualizations	Details
Message Processing	Displays the total number of messages received from the SCP that failed to process.
Processing Errors	Displays the errors encountered while processing the message/frame in Nokia SCP.

OracleSCP Statistics

Displays the visualization details of Oracle SCP statistics. To access the dashboards:

- Select the required 5G-Cloud Monitoring Session Name and the GigaVUE V Series Node names from the drop-down lists.

Table 4: Overview of OracleSCP message parsing and processing metrics, including error statistics.

Visualizations	Details
Message Parsing	Displays the total number of messages received from the SCP and indicates whether they were parsed successfully.
Parsing Errors	Displays the errors encountered while parsing the message/frame in Oracle SCP.
Message Processing	Displays the total number of messages received from the SCP that failed to process.
Processing Errors	Displays the errors encountered while processing the message/frame in Oracle SCP.

EricssonSCP Statistics

Displays the visualization details of Ericsson SCP statistics. To access the dashboards:

1. Select the required 5G-Cloud Monitoring Session Name and the DPDK MonitoringSession Name from the drop-down lists.
2. Select the required GigaVUE V Series Node names from the drop-down list.

Table 5: Ericsson SCP packet and stream statistics overview.

Visualizations	Details
Rx packets from Ericsson SCP	Displays the packet statistics received at Ericsson SCP, showing only the packets forwarded to the 5G-Cloud App
Tx packets to 5g Cloud	Displays the packet statistics transmitted to the 5G-Cloud application.
5g Cloud App Packets	Displays the packet statistics for the 5G-Cloud application, including packets received from Ericsson SCP and packets transmitted to Tools.
HTTP2 Monitored Streams	Displays the statistics of concurrent HTTP2 monitored stream flows.
HTTP2 Monitored Stream Errors	Displays statistics of concurrent HTTP2 Monitored Streams during

Visualizations	Details
	error and age-out conditions.
TCP Monitored Flows	Displays statistics of concurrent TCP monitored flows.
TCP Monitored Flow Errors	Displays statistics of concurrent TCP Monitored Flows during error, age-out, and TCP reassembly timeout conditions.

NokiaHEP3 Statistics

Nokia HEP3 Statistics dashboard displays the following visualizations. To access the dashboards:

- Select the required 5G-Cloud Monitoring Session Name and the GigaVUE V Series Node names from the drop-down lists.

Message Statistics

Table 6: Displays the visualization details of Nokia HEP3 message statistics

Visualizations	Details
HEP3 Message Overall Statistics	Displays the overall statistics of HEP3 messages received at the 5G-Cloud application.
HEP3 Rx Traffic Statistics	Displays the count of HEP3 bits received each second at the 5G-Cloud application.
HEP3 Message Statistics	Displays the count of HEP3 messages received at the 5G-Cloud application.
Missing Mandatory HEP3 Chunks Statistics	Displays the statistics of missing mandatory HEP3 chunks, detailing the total number of messages that lacked these required chunks, along with individual counts for each specific missing chunk.
Available Mandatory HEP3 Chunks Statistics	Presents the statistics for the mandatory HEP3 chunks that are available. This includes the overall number of messages containing all the required chunks, along with counters for each specific mandatory chunk. In each message, only one of the IPv4 or IPv6 chunks will be included, never both.
Tx Traffic Rate in Bits from 5G-Cloud to Tool	Displays the total number of bits transmitted per second from the 5G-Cloud application to the tools or probes.
Tx Traffic Rate in Packets from 5G-Cloud to Tool	Displays the total number of packets transmitted per second from the 5G-Cloud application to the tools or probes.
HEP3 FQDN Queries Statistics	Displays the statistics for FQDN queries conducted within the 5G-Cloud application during the synthesis of HEP3 messages.

Visualizations	Details
HEP3 FQDN Interface Name Statistics	Displays the statistics related to FQDN queries for interface names within the 5G-Cloud application during the synthesis of HEP3 messages.
Service Map Queries Statistics	Displays Service Map lookup query statistics generated in the 5G-Cloud Application during HEP3 message synthesis.
Service Map Action Statistics	Displays statistics for actions performed on Service Map queries in the 5G-Cloud Application during HEP3 message synthesis.

Protocol Statistics

Table 7: Displays the visualization details of Nokia HEP3 protocol statistics

Visualizations	Details
Transport Protocol Statistics	Displays the identified transport protocol statistics in the HEP3 payload.
Application Protocol Statistics	Displays the identified application protocol statistics in the HEP3 payload.
Outbound TCP Flow Statistics	Displays the Outbound TCP flow statistics data from the 5G-Cloud application to the tool/probes.
Outbound TCP Flow Error Statistics	Displays the Outbound TCP flow error statistics data from the 5G-Cloud application to the tool/probes.

Ericsson vTAP Statistics

To access the dashboards:

- Select the required EVTAP Monitoring Session Name and GigaVUE V Series Node Name from the drop-down lists to view the visualizations and metric details. Ericsson vTAP Statistics dashboard displays the following visualizations:

Packet Statistics

Table 8: Displays the visualization details of Ericsson vTAP packet statistics

Visualizations	Details
Rx Packets from EVTAP	The view displays aggregated packet statistics collected from the Ericsson VTAP interface.
5g SBI App Packets	The view displays the Rx and Tx packet statistics for the 5G-SBI application.
Tx Packets	The view displays the transmitted packet statistics.
5g SBI App Bits	The view displays the Rx and Tx bit statistics for the 5G-SBI application.
Layer 7 HTTP Ingress TEP Statistics	The view displays the packet statistics received at ingress TEPs that

Visualizations	Details
	handle Layer 7 HTTP traffic.
Layer 7 HTTP Egress TEP Statistics	The view displays the packet statistics transmitted from egress TEPs that handle Layer 7 HTTP traffic.
Native HTTP Ingress TEP Statistics	The view displays the packet statistics received at ingress TEPs that handle native HTTP traffic.
Native HTTP Egress TEP Statistics	The view displays the packet statistics transmitted from egress TEPs that handle native HTTP traffic.
Non SBI Ingress TEP Statistics	The view displays the packet statistics received at ingress TEPs that handle non-SBI traffic.
Non SBI Egress TEP Statistics	The view displays the packet statistics transmitted from egress TEPs that handle non-SBI traffic.

TCP Statistics

Table 9: Displays the visualization details of Ericsson vTAP TCP statistics

Visualizations	Details
TCP Flows	The view displays the TCP flow statistics for the 5G-SBI application.
TCP Flow States	The view displays the statistics related to TCP flow states for the 5G-SBI application.
HTTP2 Transaction Flows	The view displays the HTTP/2 transaction flow statistics for the 5G-SBI application.
HTTP2 Transaction Flow States	The view displays the statistics related to HTTP/2 transaction flows for the 5G-SBI application.
Synthesized Statistics	The view displays the statistics related to message synthesis in the 5G-SBI application.
Error Statistics	The view displays the message synthesis error statistics for the 5G-SBI application.
Messages Statistics	The view displays the message statistics for the 5G-SBI application.
Messages Error Statistics	The view displays the message error statistics for the 5G-SBI application.
NF ID Look up	The view displays the NF Instance ID lookup statistics for the 5G-SBI application.
FQDN Look up	The view displays the FQDN lookup statistics for the 5G-SBI application.

Ericsson Packet Processing Statistics

Table 10: Displays the visualization details of Ericsson Packet Processing Statistics

Visualizations	Details
Rx Packets	The view displays the Packet statistics received at EVP Application. Rx Packets -- Average number of packets received per second (RX side).
Tx Packets	The view displays the Packet statistics transmitted from EVP Application.
Rx Bits	The view displays the Bits statistics received at EVP Application.
Tx Bits	The view displays the bits transmitted from EVP application.
Rx Statistics - Others	The view displays other RX statistics for the EVP Application: dropped, errored, and bypassed packets. All statistics show the average rate of change per second.
Tx Statistics - Others	The view displays other TX statistics for the EVP Application: dropped and errored packets. All statistics show the average rate of change per second..
Flow Statistics	The view displays UDP flow statistics for the EVP Application: total, active, free, retrieved, and inserted active UDP flows. All statistics except Flows show the average rate of change.
Flow Error Statistics	The view displays UDP flow error statistics for the EVP Application: flow retrieval, free flow, and flow allocation errors. All statistics show the average rate of change.
Bucket Sort Packet Stats	The view displays bucket sort packet statistics for the EVP Application: total, sent, and successfully sorted packets. All statistics show the average rate of change per second.
Bucket Sort Bit Stats	The view displays bucket sort bit statistics for the EVP Application: bits sent from the bucket to the TX queue. Statistics show the average rate of change per second.
Bucket Sort Statistics-Others	The view displays other bucket sort statistics for the EVP Application: bypassed, discarded, too new, and too old packets. All statistics show the average rate of change per second.
Packet Sort Statistics	The view displays packet sort statistics for the EVP Application: successfully inserted and total inserted packets. All statistics show the average rate of change per second.
Bucket Sort Failure Statistics	The view displays bucket sort failure statistics for the EVP Application: list full, access, allocation, packet list full, unknown, and send failures. All statistics show the average rate of change per second.

GigaSMART Mobility Session and Flow Filtering Dashboards

GigaSMART mobility session and flow filtering dashboards display flow filtering summary reports and statistics pertaining to the mobility solutions on an hourly, daily, weekly, or monthly interval. You can export the data from the dashboards page, and create your own flow ops visualizations for your required use cases.

How to Access the Dashboards

To access the Mobility Session and Flow Filtering dashboards:

- **From the Dashboard menu:** Go to -> **Analytics -> Dashboards** -> **5G LTE Sessions**
- **From the Traffic menu**: Go to -> **Physical** -> **Orchestrated Flows** -> **Mobility** -> **Dashboard**
- **From the Inventory menu:** Go to **-> Physical -> Nodes ->GigaSMART Groups -> Report -> Flow Filtering**

The following dashboards are displayed:

5G LTE Sessions

The following table displays visualizations for 5G and LTE Session statistics

Table 4: Default Mobility Dashboards.

Visualizations	**Details**
Mobility 5G Session Statistics	Displays maximum aggregation of 5G numSessions, numSessions available and session capacity for all the GigaSMART groups pertaining to the mobility solutions against the time stamp.
Mobility LTE Session Statistics	Displays maximum aggregation of LTE numSessions, numSessions available and session capacity for all the GigaSMART groups pertaining to the mobility solutions against the time stamp.
5G Sessions by GSGroup Alias	Displays the maximum aggregation of all 5G statistics parameters for each GigaSMART group pertaining to the mobility solutions. The data columns are sortable and the user can find top N values by sorting across each field.
LTE Session by GSGroup Alias	Displays the Max aggregation of all LTE statistics parameters for each GigaSMART group pertaining to the mobility solutions. The data columns are sortable and the user can find top N values by sorting across each field.
Mobility Session No SFFP Match	Displays Max aggregation of LTE and 5G num of SFFP No Match for all the GigaSMART groups pertaining to the mobility solutions against time stamp.

Flow Filtering

The following table displays visualizations related to flow filtering statistics

Table 5: Default Mobility Dashboards.

Visualizations	Details
Flow Filtering Statistics	Displays Max aggregation of controlTunnels, controlUserTunnels, controlOnlyTunnels and pendingSession for all the GigaSMART groups pertaining to the mobility solutions against timestamp.
Flow Filtering Statistics per GsGroup	Displays Max aggregation of of controlTunnels, controlUserTunnels , controlOnlyTunnels and pendingSession for each GigaSMART group pertaining to the mobility solutions. The data columns are sortable and the user can find top N values by sorting across each field.

Flow Filtering Interface Statistics

The follwing table displays visualizations related to flow filtering interface statistics

Table 6: Default Mobility Dashboards.

Visualizations	Details
Flow Filtering Interface Statistics - Packets	Displays Max aggregation of rxPkts, txPkts and droppedPkts for all the interface types in all GigaSMART group pertaining to the mobility solutions against timestamp.
Flow Filtering Interface Statistics - Bytes	Displays Max aggregation of rxBytes, txBytes and droppedBytes for all the interface types in all GigaSMART groups pertaining to the mobility solutions against timestamp.
Flow Filtering Interface Statistics - GTPC Packets	Displays Max aggregation of GTPC packets for all the interface types in all GigaSMART groups pertaining to the mobility solutions against timestamp.
Flow Filtering Interface Statistics - Packets Percentage	Displays Percentage for sum of tx and sum of dropped packets to sum of rx packets for all the interface types in all GigaSMART group pertaining to the mobility solutions against timestamp.
Flow Filtering Interface Statistics per Interface type	Displays Max aggregation of all Flow Filtering Interface stats for each interface type in each GigaSMART group pertaining to the mobility solutions. The data columns are sortable and the user can find top N values by sorting across each field.

The Flow Filtering Correlation Statistics

The following table displays visualizations related to flow filtering correlation statistics.

Table 7: Default Mobility Dashboards.

Visualizations	Details
Flow Filtering Correlation Statistics - Control	Displays Max aggregation of control correlations statistics for all GigaSMART group pertaining to the mobility solutions against timestamp.
Flow Filtering Correlation Statistics - User	Displays Max aggregation of user correlations statistics for all GigaSMART group pertaining to the mobility solutions against timestamp.

Visualizations	Details
Flow Filtering Control Correlation Statistics per Control Message	Displays the Max aggregation of all Flow Filtering Correlation User stats for each GigaSMART group pertaining to the mobility solutions.

Flow Filtering Correlation Statistics

The following table displays visualizations related to flow filtering correlation statistics

Table 8: Default Mobility Dashboards.

Visualizations	Details
Flow Filtering Correlation Statistics - Control	Displays Max aggregation of control correlations statistics for all GigaSMART group pertaining to the mobility solutions against timestamp.
Flow Filtering Correlation Statistics - User	Displays Max aggregation of user correlations statistics for all GigaSMART group pertaining to the mobility solutions against timestamp.
Flow Filtering Control Correlation Statistics per Control Message	Displays the Max aggregation of all Flow Filtering Correlation User stats for each GigaSMART group pertaining to the mobility solutions.

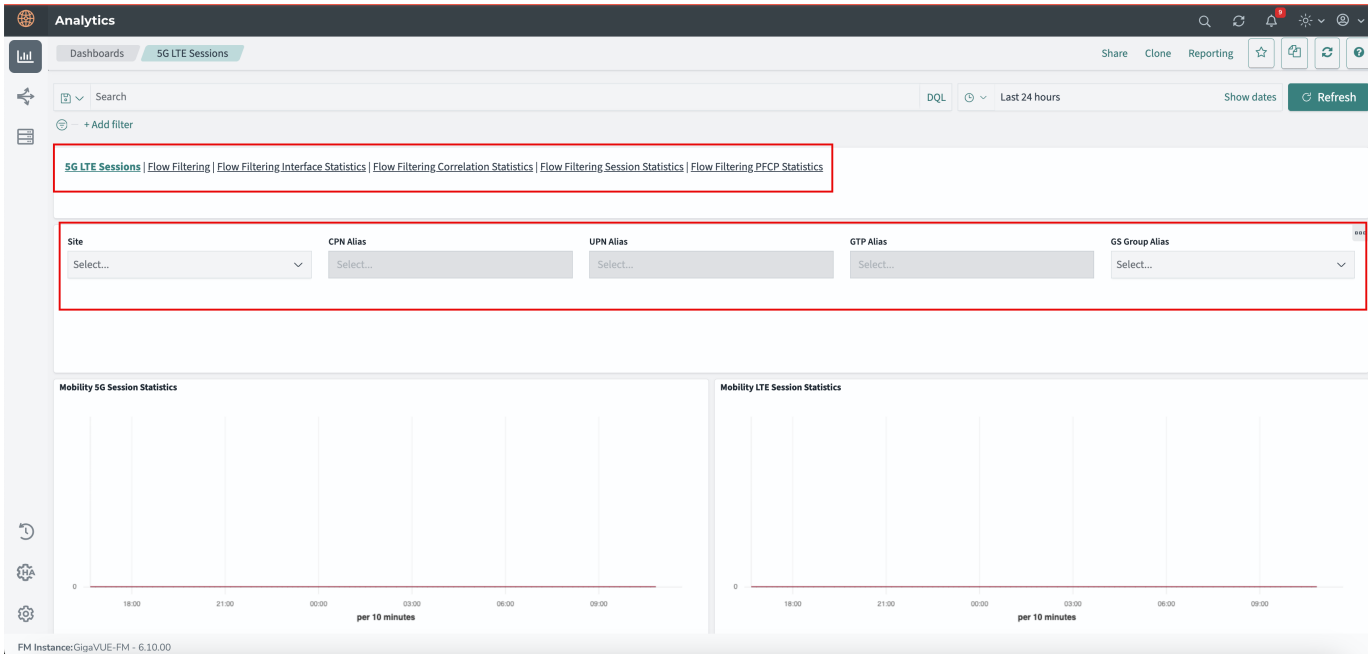
Flow Filtering Session Statistics

The following table displays visualizations related to flow filtering session statistics

Table 9: Default Mobility Dashboards.

Visualizations	Details
Flow Filtering Session Statistics - Sessions	Max aggregation of sessions for all GigaSMART group pertaining to the mobility solutions against timestamp.
Flow Filtering Session Statistics - Tunnels	Max aggregation of tunnels for all GigaSMART group pertaining to the mobility solutions against timestamp.
Flow Filtering Session Statistics per Interface type	Displays Max aggregation of all Flow Filtering session and tunnel stats for each interface type in each GigaSMART group pertaining to the mobility solutions. The data columns are sortable and the user can find top N values by sorting across each field.
GTP Threshold limit utilization Alerts	Provides an alert when the threshold has crossed the configured limit for Session usage, Tunnel usage, CPU Utilization and l Packet Buffer Utilization.
Gsgroup CPU usage in percentage	Displays percentage of CPU utilization in all GigaSMART groups pertaining to the mobility solution

Refer to the following screen shot.



Rules and Notes

- You cannot edit or delete these default dashboards. However, you can clone the dashboards and visualizations.
- You cannot visualize CLI configurations on these default dashboards.
- Flow Filtering Dashboards are not supported for 5G flow-ops report.
- You can use the following control visualizations to filter and visualize the data based on the following criteria:
  - Site
  - CPN Alias
  - UPN Alias
  - GTP Alias
  - GS Group Alias

GigaSMART Inline TLS/SSL Dashboards

GigaSMART Inline TLS/SSL Dashboards offer insights into session performance, network capacity, traffic decryption, compliance analysis, and historical data. Monitoring decryption statuses and anomalies helps organizations enhance security.

These dashboards provide real-time alerts and detailed reports for network security administrators to maintain data integrity and security compliance. It allows you to visualize the information with GigaVUE-FM. These dashboards are supported only for Gen 3 GigaSMART card platforms.

A few of the use case scenarios where the Inline TLS/SSL Dashboard could detect and manage anomalies:

- Alert administrators when a TLS handshake involves certificates signed with insecure hash algorithms.
- Alerts can be triggered when CBC mode is used, especially in older versions of TLS (for example, TLS 1.0 and TLS 1.1), advising an upgrade to more secure cipher modes like GCM (Galois/Counter Mode).
- Identify and report the use of certificates with weak signatures in the network traffic, facilitating a swift response to enhance security.
- Automatic detection and reporting of expired certificates help maintain continuous security compliance and trust.
- Monitoring and analyzing trends in decryption success and failure rates can pinpoint disruptions or anomalies in encrypted traffic handling.
- Ensure only approved cryptographic standards are used and generate compliance reports for auditing purposes.

To access the dashboard:

1. Go to  -> **Analytics -> Dashboards.**
2. Click on the required dashboard to view the visualizations.

Inline TLS/SSL Dashboard can be categorized into two types:

- Basic Dashboards
- Advanced Dashboards

Basic Dashboards

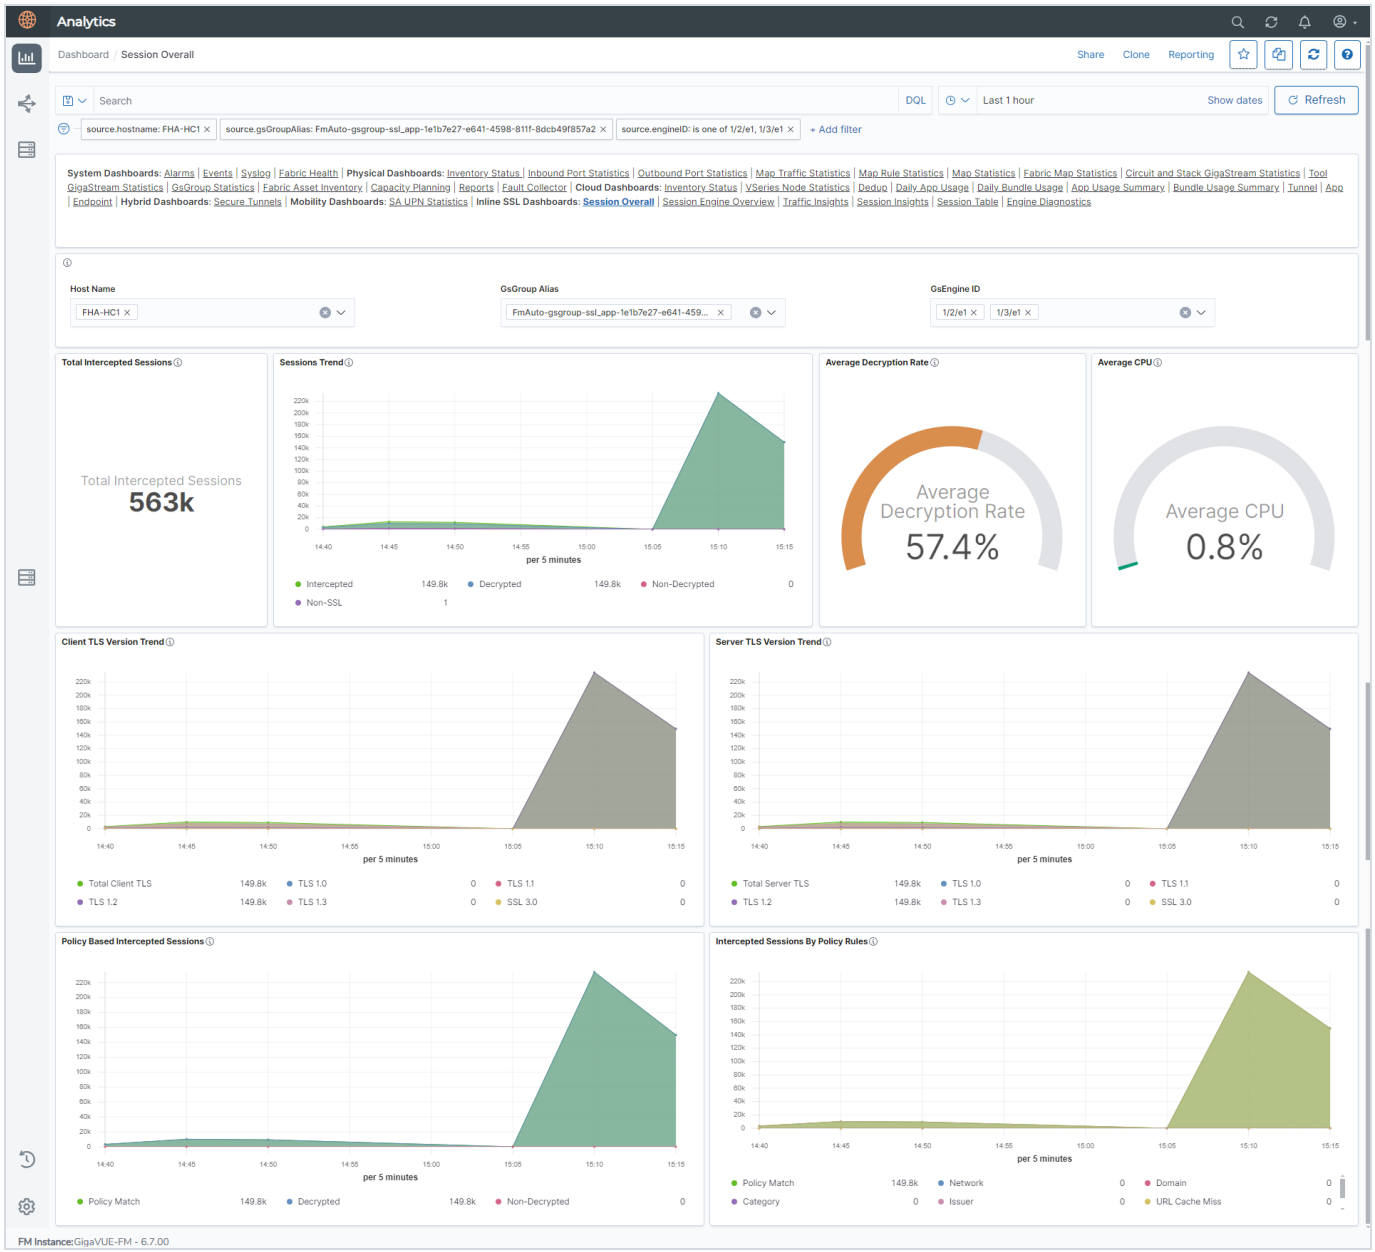
The Basic dashboards are available by default and provides an overall information on the session. You can use the below control filters and specify the time period to visualize and filter the dashboard information:

- Host Name
- GigaSMART Groups Alias (GSGroup Alias)
- GigaSMART Engine ID (GSEngine ID)

The following table displays visualizations on the overall details of encrypted traffic.

Table 10: Session Overall Dashboard

Visualizations	Details
Total Intercepted Sessions	Displays overall count of intercepted sessions by the node over time period. This page does not display per engine unless specified by a filter.
Sessions Trend	Displays the trend of all Inline TLS/SSL session that has been received over a specified time period and per specified GigaSMART engines. The trend included the visualization of the Intercepted/Decrypted/Non-SSL Sessions.
Average Decryption Rate	Displays the average rate of Inline TLS/SSL sessions that have been decrypted over the specified time period..
Average CPU	Displays an average CPU utilization of all engines in Session Overall Page.
Client TLS Version Trend	Displays an overview of the incoming traffic's TLS version of the incoming data.
Server TLS Version Trend	Provides an insight into the TLS version distribution at the server side.
Policy based Intercepted Session	Displays the trend of decryption status of Inline TLS/SSL session based policy.
Intercepted Sessions By Policy Rules	Displays the trend of Inline TLS/SSL session based on Policy rules such as; Domain ,Category, Issuer, URL Cache Miss, Network and Default. **Note**: The no. of sessions that gets matched to a Network Policy Rule will not be displayed in the Total Intercepted Session widget.

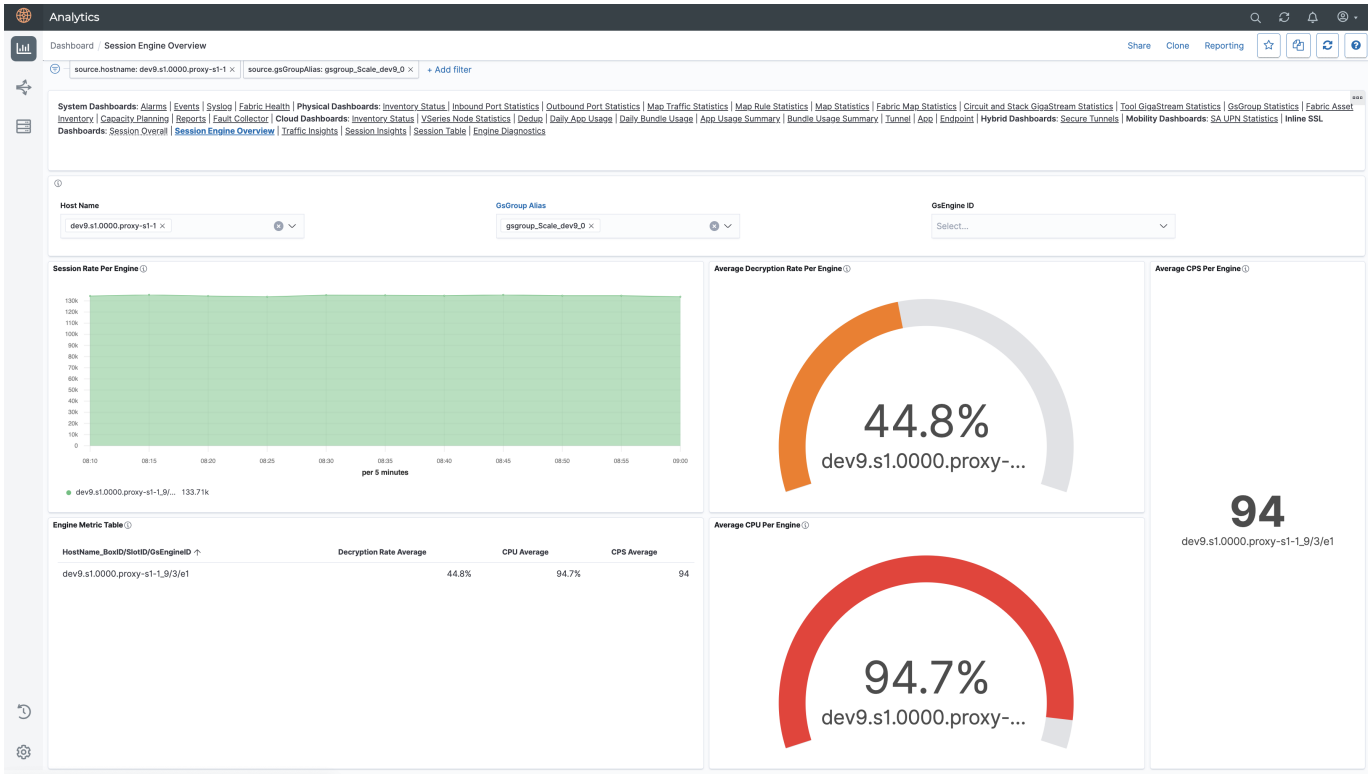


The following table displays visualizations related to Inline TLS/SSL Sessions per Engine

Table 11: Session Engine Overview Dashboard

Visualizations	Details
Sessions Rate per Engine	Displays the rate at which Inline TLS/SSL sessions are intercepted per engine.
Average Decryption Rate per Engine	Displays the average rate of sessions that got decrypted per engine.

Visualizations	Details
Average CPS per Engine	Displays the average Connections per Second (CPS) performance metric per engine.
Average CPU per Engine	Displays the average CPU utilization per engine.
Engine Metric Table	Displays the Decryption rate per engine, average CPS and average CPU rate in a tabular format. The details are displayed as Host Name/Engine ID. For example; FHA-HC1 (Host Name)_1/3/e11(GSEngine ID)

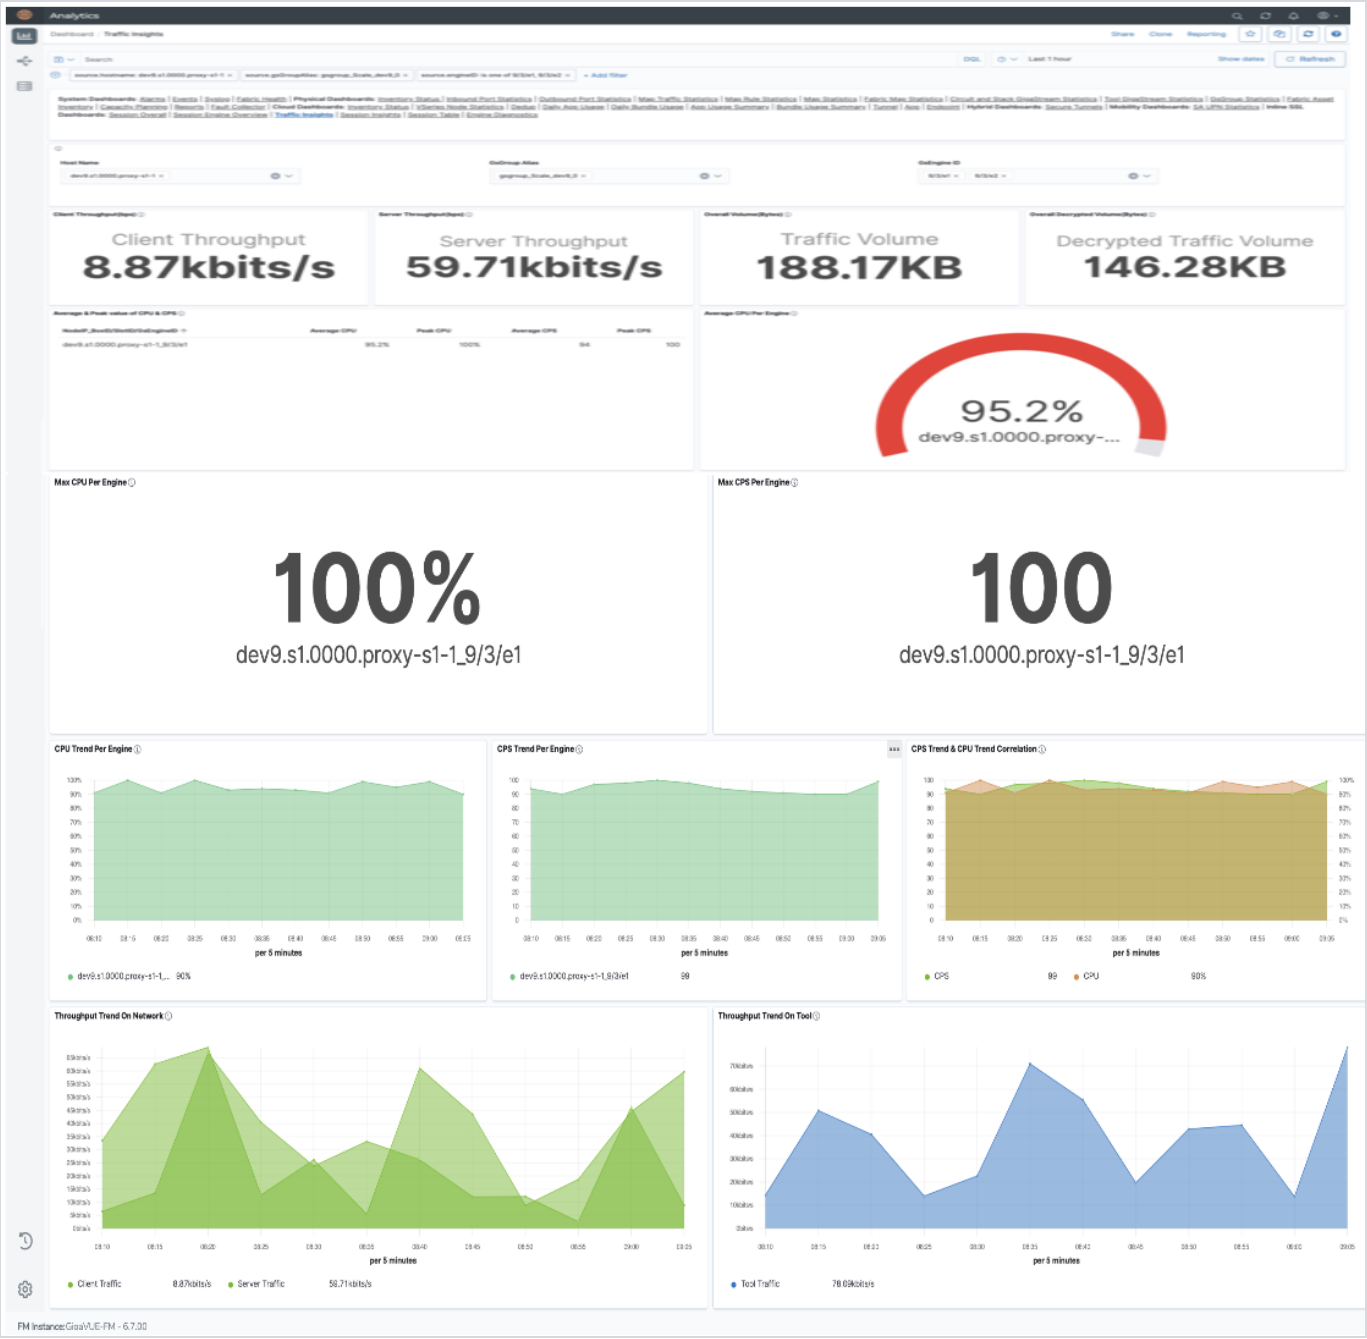


The following table displays visualizations related to the traffic that is handled with a Inline TLS/SSL sessions

Table 12: Traffic Insights Dashboard

Visualizations	Details
Client and Server Throughput (bps)	Displays the traffic throughput that is received from the client and the throughput that is handled at the server side. This throughput is displayed in bits per second (bps)value.
Overall Volume(Bytes)	Displays the volume of traffic that is being handled in

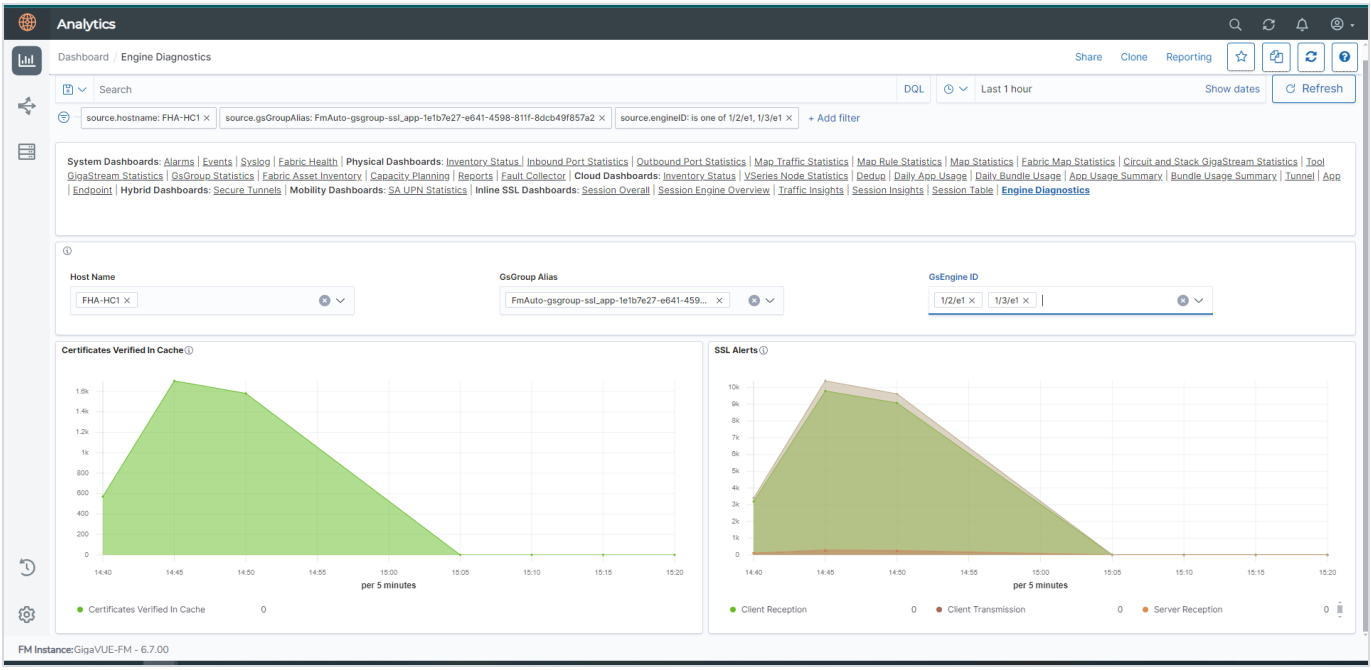
Visualizations	Details
	Bytes. This takes into account both TCP and SSL sessions.
Overall Decrypted Volume (Bytes)	Displays the overall decrypted volume of all engines unless filtered by engine ID control filter in Bytes unit
Average CPU Per Engine	Displays the average CPU performance per engine
Max CPU Per Engine	Displays the maximum CPU utilization that was observed per engine. This is static rate and is not displayed based on a time frame.
Max CPS Per Engine	Displays the maximum Connection Per Second (CPS) rate that was observed per engine. This is static rate and is not displayed based on a time frame.
Average & Peak value of CPU & CPS	Displays the average and peak values of CPU and CPS observed per engine in a tabular format.
CPU Trend per Engine	Displays a trend of CPU utilization that was achieved over a time period per engine.
CPS Trend per Engine	Displays a trend of the Connections per Second that was achieved over a time period per engine.
CPS Trend & CPU Trend Correlation	Displays a correlation between the CPU and CPS trend of the engine within a time period.
Throughput Trend on Network	Display the throughput trend of traffic that was received from both client and server side.
Throughput Trend on Tool	Displays the throughput trend of traffic that was received on the Tool.



The following displays the certificates and SSL alerts related to a GigaSMART engine

Table 13: Engine Diagnostics Dashboard

Visualizations	Details
Certificates verified in Cache	Displays the number of certificates that were verified in cache over a time period
SSL Alerts	Displays the number of SSL alerts that were received both from client and server.



Advanced Dashboards

Advanced Dashboards are available only if you enable it while configuring your Inline TLS/SSL Decryption session.

System Requirements

The system requirements for utilizing Inline TLS/SSL Advanced Dashboards are as shown below.

Requirements	Support up to 100 Devices (GigaVUE-FM Standalone)	Support up to 100 Devices (GigaVUE-FM HA Mode)
Memory	128GB	128GB
Virtual CPU	Minimum 12 CPU **NOTE:** It is recommended to have 16 CPU for continuous traffic with maximum supported limit of 18k sessions/second for three Advanced Statistics enabled GigaSMART engines.	Minimum 12 CPU **NOTE:** It is recommended to have 16 CPU for continuous traffic with maximum supported limit of 36k sessions/second for six Advanced Statistics enabled GigaSMART engines.
Disk Space	Refer to "*Large Configuration*" category under "*Virtual Computing Resource Requirement in Scaled Environments*" section in GigaVUE-FM Installation and Upgrade Guide for disk space details.	Refer to "*Large Configuration*" category under "*Virtual Computing Resource Requirement in Scaled Environments*" section in GigaVUE-FM Installation and Upgrade Guide for disk space details.
Virtual Network Interface	1	1
Number of GigaVUE-FM nodes	1	3

Configure Advanced dashboard

To configure advanced dashboards:

- Go to, **Traffic** **>Configuration Canvas > Select the device> Inline SSL APP**.
- Enable the toggle option **Advanced Session Statistics**.

Rules and Notes

Keep in mind the following rules and notes when using the Advanced Dashboard:

- Advanced Dashboard data will be retained for 24 hours.
- For a standalone GigaVUE-FM node, the Advanced Dashboard is available for a maximum of three GigaSMART engines.
- In a GigaVUE-FM High Availability group with three GigaVUE-FM nodes, a maximum of six GigaSMART engines will be supported.
- Configure NTP time sync or ensure that your device and GigaVUE-FM are synchronized with the date and time zone.

You can use the below control filters and specify the time period to visualize and filter the dashboard information:

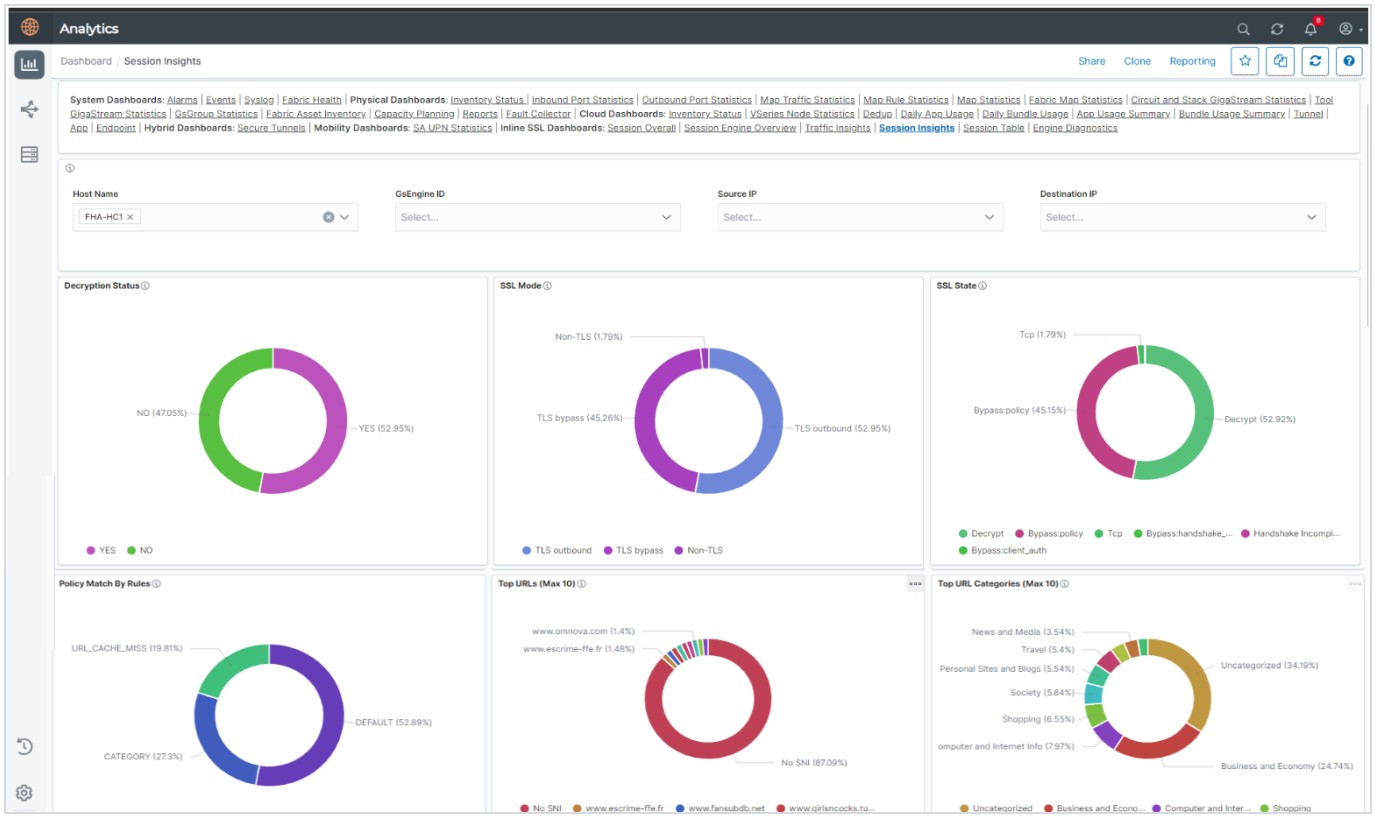
- Host Name
- GigaSMART Engine ID (GSEngine ID)

- URL (Only for Session Table Dashboard)
- Source IP
- Destination IP
- URL Category (Only For Session Table Dashboard)

The following table displays visualizations on the details of an Inline TLS/SSL session.

Table 14: Session Insight Dashboard

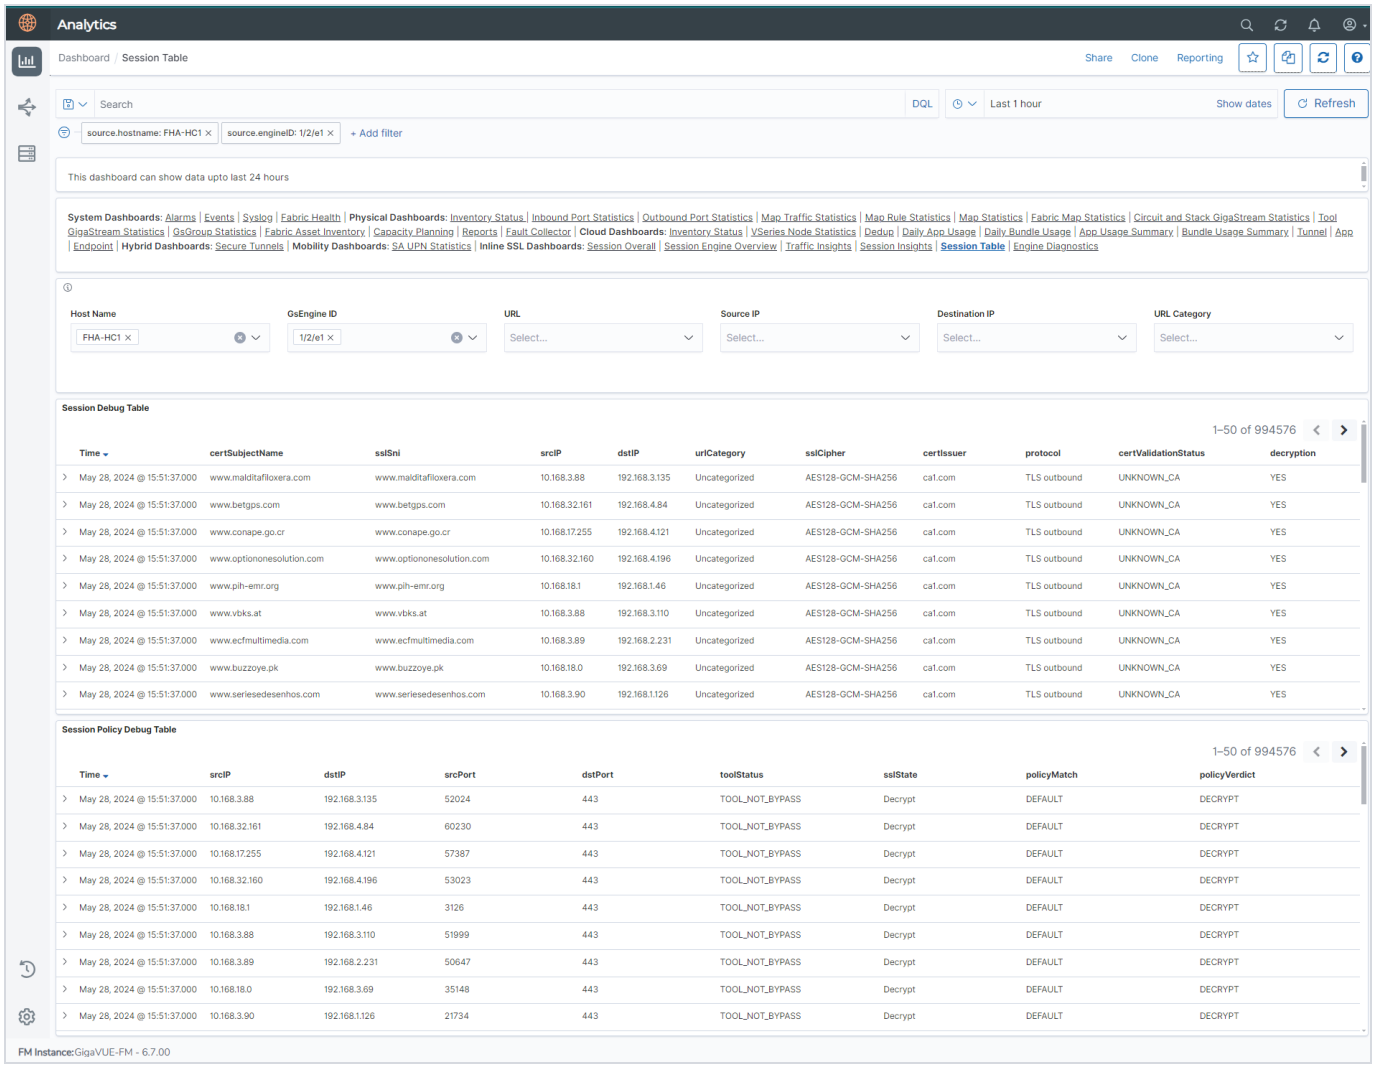
Visualizations	Details
Decryption Status	Displays the number of Inline TLS/SSL sessions that were decrypted and not decrypted.
SSL Mode	Displays the distribution of TLS/SSL Session modes. The modes are as follows: • TLS/SSL Outbound- : Sessions decrypted due to ISSL inbound deployment. • TLS/SSL Inbound-Sessions decrypted due to ISSL outbound deployment. • TLS/SSL Bypass- The session mode that is neither inbound or outbound. • Non-SSL - TCP sessions that are not an TLS/SSL session.
SSL State	Displays the distribution of TLS/SSL Session statuses.
Policy Match By Rules	Provides an insight into the TLS/SSL session that matches the Policy Rules. **Note**: The Policy Rule CATEGORY indicates the URL category.
TLS Version	Displays the TLS version of the sessions. **Note**: The counter “Bypass/Error” denotes sessions that were not able to determine the TLS version.
Top URLS (Max 10)	Displays the top 10 URLs that were accessed during the Inline TLS/SSL Session.
Top URL Category (10 Max)	Displays the Category of top 10 URLs accessed in Inline TLS sessions **Note**: 'Uncategorized' signifies SNIs that could not be categorized or Non TLS sessions. **Note**: 'Unknown' signifies TLS Bypass and IP address based URLs.
Top Ciphers (Max 10)	Displays the top 10 Ciphers that performed the Inline TLLS/SSL Decryption.
Certificates by Type	Displays the certificates received are valid or non-valid.



The following table displays visualizations related to Inline TLS/SSL Sessions per Engine in a tabular format.

Table 15: Session Table Dashboard.

Visualizations	Dashboard
Session Debug Table	Displays the entire Sessions Debug details throughout the system that has enabled Advanced Session Statistics. Each field can be added or removed as a customized filter option by using ⊕ ⊖ button.
Session Policy Debug Table	Displays the entire Sessions Policy Debug details throughout the system that has enabled Advanced Session Statistics. It points out to the policy rules that got matched or the policy verdict of Decryption or non decryption. Each field can be added or removed as a customized filter option by using ⊕ ⊖ button.



OpenSearch Indices and Fields Used in Analytics

GigaVUE-FM uses the following OpenSearch indices to store data[1]. These OpenSearch indices store different sets of data based on usability. The document mappings of the indices are also different.

- ***fmstats****: For storing statistics information of entities such as maps, gsgrops, gsops, etc (except port)
- ***fmstats_ports****: For storing statistics information of the port

[1]These indices were introduced during different time frames for various purposes. The data models of these indices were designed as per the need during the time of development.

- **fminventory**: For storing the inventory and assets.
- **fmevents**: For storing events
- **fmalarms**: For storing alarms.

The following table lists the new keys introduced in software version 6.0:

Table 16:

New Field	Applicable Indices	Type	Description
resource.id.entityid	• **fmalarms** • **fmevents** • **fmstats:** Applicable for the following document types: ◦ **vport stats** ◦ **gsop stats** ◦ **gtap port stats**	text	Represents the entity ID. Holds value exactly similar to **resourceId** in the existing alarms/events document. If the **resourceId** does not exist in the document, the **resource.id.entityId** also will not exist. Example: For alarms related to port, port ID will be the **resourceId**, and **resource.id.entityId**.
resource.id.deviceId	• **fmalarms** • **fmevents** • **fmstats_ports** • **fmstats:** Applicable for the following document types: • **gsgroup stats** • **gsgroup port rate core rate stats** • **vport stats** • **Map and MapRule stats** • **gsop stats**	text	Represents the **deviceId**. That is, hostname. If the hostname does not exist in the document, the **resource.id.deviceId** also will not exist in the document
resource.id.clusterId	• **fmalarms**	text	Represents the **clusterId**. If the **clusterId** does not exist in the document, then **resource.id.clusterId** also will not exist in the document

New Field	Applicable Indices	Type	Description

	• **fmstats:** Applicable for the following document types:
	  • **vport stats**
	  • **gsop stats**		

resoure.type	• **fmalarms**
	• **fmevents**
	• **fmstats:** Applicable for the following document types:
	  • **vport stats**
	  • **gsop stats**
		text	Represents the resource type (port, map, device, gsgroup)

resource.id.alias	• **fmalarms**
	• **fmevents**
	• **fminventory**
	• **fmstats_ports**
	• **fmstats:** Applicable for the following document types:
	  • **vport stats**
	  • **gsop stats**
	  • **Map and Map Rule stats**
		text	Represents the alias name of the associated entity of the alarm. If the alias does not exist in the document, then **resource.id.alias** will not exist in the document.

resource.name	• **fminventory**
	• **fmstats**
	• **fmstats:** Applicable for the following document types:
	  • **vport stats**
	  • **gsop stats**
			Represents the key used to uniquely identify a particular entity across the system. For example, for ports, ClusterID, and PortId can be combined and used to uniquely identify the entity. In case of map, ClusterID, and map alias name are used. The value will be in the form of **ClusterId__ EntityId**

port.alias	**fmstats_ports**		Represents the alias name of the port's alias.
port.dir	**fmstats_ports**		Represents the direction of the port.

New Field	Applicable Indices	Type	Description
gsGroup.alias	• **fmstats:** Applicable for the following document types: • **gsgroup stats** • **gsgroup port rate, and core rate stats documents** • **vport stats** • **gsop stats**		Represents the alias name of the gsgroup alias
map.alias	• **fmstats:** Applicable for the following document types: • **Map and MapRule stats** • **gsop stats**		Represent the alias name of the map alias
gsVport.alias	• **fmstats:** Applicable for the following document types: • **vport stats** • **gsop stats**		Represents the alias name of the vport alias

New Field	Applicable Indices	Type	Description
cluster.mode	**fminventory**		Represents the mode of the device with respect to the cluster. Applicable values are as follows: • leader • standby • standalone • normal
linkSpeedInBits	• **fmstats_ports**		Represents the link speed in bits per second. **Note**: It is recommended to use the **linkSpeedInBits** field in custom visualizations (instead of **linkSpeed**)

Deprecated Fields in Analytics

The following table lists the fields that will be deprecated in the future software versions. It is recommended to use the replaced fields instead.

Table 17:

Deprecated Field	Replaced Field
Ports	
portIdToClusterId	*resource.name*
portId	*resource.id.entityId*
clusterId	*resource.id.clusterId*
objectType	*resoure.type*
alias	*resource.id.alias*
Maps	
clusterId	*resource.id.clusterId*
objectType	*resoure.type*
alias	*resource.id.alias*

Deprecated Field	Replaced Field
Map Rule	
clusterId	*resource.id.clusterId*
objectType	*resoure.type*
alias	*resource.id.alias*

Analytics Appendix

The below section describes the Analytics components and the specific fields .

Ports

The below section lists the dimensions and metrics collected for Ports .

Dimensions

Fields	Description
hostname	The hostname of the device.
portType	The type of Port.
resource.name	This is a combination of cluster id and port Id (entity ID).
resource.type	This shows the type of statistics that is being analyzed.
resource.id.entityId	The Port Id.
resource.id.clusterId	The Cluster Id.

Metrics

The metrics is further categorized to Gauge Metrics and Counter Metrics

Gauge Metrics

For Rolled Up data the Gauge metrics stores average , maximum, minimum and sum value.

Fields	Description
port.tx.utilztn	The transmission packets total utilization.
port.tx.packetRps	The transmission packets by rate per second
port.tx.bitsRps	The transmission packet bits by rate per second.
port.rx.utilztn	The received packets total utilization.

Fields	Description
port.rx.sfpPowerMin	The minimum SFP power.
port.rx.sfpPowerMax	The maximum SFP power.
port.rx.bitsRps	The Received packet bits rate per second.

Counter Metrics

For Rolled Up data the counter metrics stores only maximum value.

Fields	Description
port.tx.packetDiscard	The transmission packets that were discarded.
port.tx.octets	The transmission octets that passed via port.
port.tx.packetErr	The transmission packets that encountered error at the port.
port.tx.packetDrop	The transmission packets that were dropped.
port.tx.packets	The transmission packets that were passed
port.rx.packetDiscard	The received packets that were discarded
port.rx.octets	The received octets that were passed
port.rx.packetErr	The received packets that encountered error at the port.
port.rx.packetDrop	The received packets that were dropped.
port.rx.packets	The total received packets.
linkSpeedInBits	The link speed of port in bits.

Maps

The below section lists the dimensions and metrics collected for Maps.

Dimensions

Fields	Description
cluster Id	The cluster Id of the device.
host	The host of the device
objectType	The type of statistics.
alias	The alias assigned to the map.
mapType	The type of map.
mapSubtype	The sub type of the map.
ruleOrder	The order in which the rule is configured.
resource.name	This is a combination of host, alias and rule order.
resource.type	The type of statistics.

Fields	Description
resource.id.deviceId	The device id.
resource.id.alias	The alias assigned to the map.
resource.id.entityId	The entity Id of the map.
resource.id.clusterId	The cluster Id of the map.
map.alias	The alias assigned to the map.

Metrics

The metrics is further categorized to Gauge Metrics and Counter Metrics . Here we have described the counter metrics.

Counter Metrics

For Rolled Up data the counter metrics stores only maximum value.

Fields	Fields
map.octets	The total octets processed
map.octets_rule_pass	The number of octets passing the rule.
map.octets_rule_drop	The number of octets dropped due to the rule.
map.packets	The total packets.
map.packets_rule_pass	The packets that passed the matching rule.
map.packets_rule_drop	The packets That were dropped due to not matching the rule.

Map Rule

The below section lists the dimensions and metrics collected for Map Rules.

Dimensions

Fields	Fields
clusterId	The cluster Id of the device.
host	The host name of the device.
objectType	The type of statistics.
alias	The alias assigned to the map rule.
ruleType	The type of map rule.
ruleUniqueId	The unique ID assigned to the map rule.
ruleOrder	The order of map rule.
resource.name	This is a combination cluster id, alias and rule order.

Fields	Fields
resource.type	The type of statistics.
resource.id.deviceId	The device Id.
resource.id.alias	The alias assigned to the map rule.
resource.id.entityId	The entity Id of the map rule.
resource.id.clusterId	The cluster Id of the device.
map.alias	The alias assigned to the map rule.
mapRule.ruleUniqueId	The unique ID assigned to the map rule.

Metrics

The metrics is further categorized to Gauge Metrics and Counter Metrics . Here we have described the counter metrics.

Counter Metrics

For Rolled Up data the counter metrics stores only maximum value.

Fields	Fields
mapRule.octets	The total octets processed through map rule.
mapRule.packets	The total packets processed through map rule.
mapRule.octets_rule_drop	The total octets dropped due to the map rule.
map.accepted	The total octets accepted.
mapRule.acceptedBytes	The total bytes accepted.
mapRule.acceptedPkts	The total packets accepted.
mapRule.matched	The total octets matched with map rule.
mapRule.matchedBytes	The total bytes matched with map rule.
mapRule.matchedPkts	The total packets matched with map rule.
mapRule.rejected	The total octets that were rejected due to not matching with map rule.
mapRule.rejectedPkts	Total packets rejected not matching rule.
mapRule.rejectedBytes	Total bytes rejected not matching rule.

Nodes and Clusters

This section introduces the GigaVUE HC Series and GigaVUE TA Series of the GigaVUE Traffic Visibility nodes. It also describes node and fabric management activities that you can perform on the physical nodes' embedded GigaVUE-OS from the GigaVUE-FM GUI.

Topics:

- GigaVUE Nodes and Clusters
- Manage GigaVUE® Nodes and Clusters
- Multi-Path Leaf and Spine
- Spine to Spine and Leaf
- Fabric Statistics
- Topology Visualization
- Flows
- Device Logs and Event Notifications

NOTE: For device backup/restore details refer to the *GigaVUE Administration Guide*.

GigaVUE Nodes and Clusters

This section introduces the GigaVUE HC Series and GigaVUE TA Series of GigaVUE Traffic Visibility nodes. It also describes their features and functions of the GigaVUE family of nodes. It includes the following major sections:

- GigaVUE® HC Series and TA Series Overview
- About Cluster

GigaVUE® HC Series and TA Series Overview

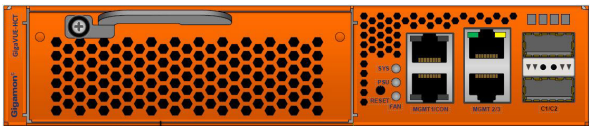
The GigaVUE HC Series delivers performance and intelligence in each of its Deep Observability Pipeline nodes, with port density and speeds that scale to your needs, from 1Gb to 100Gb. With an intuitive Web-based interface (GigaVUE-FM) and a powerful GigaVUE-OS, the Deep Observability Pipeline is able to replicate, filter, and selectively forward network traffic to monitoring, management, and security tools.

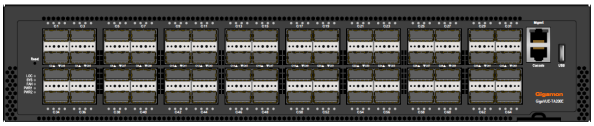
The GigaVUE HC Series and GigaVUE TA Series include the following models that run GigaVUE-OS:

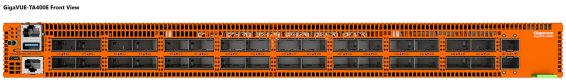
- GigaVUE-HC1
- GigaVUE-HC3
- GigaVUE-HC1-Plus
- GigaVUE-HCT
- GigaVUE-TA25
- GigaVUE-TA25E
- GigaVUE-TA100
- GigaVUE-TA200
- GigaVUE-TA200E
- GigaVUE-TA400
- GigaVUE-TA400E

NOTE: This document describes how to configure and operate the GigaVUE-OS for GigaVUE HC Series and GigaVUE TA Series nodes.

GigaVUE-HC1	• 1RU Footprint • Two Module Slots (Bays) • Dedicated Cluster Management Port • Supports all GigaVUE-HC1 Modules • Cluster with GigaVUE HC Series and GigaVUE TA Series Nodes • All ports, excluding BPS ports, of same type and speed can be used to create GigaStream.	
GigaVUE-HC3	• 3RU Footprint • Four Module Slots (Bays) • Internal Control Card • Extension Board • Dedicated Cluster Management Port • Supports all GigaVUE-HC3 Modules • Cluster with GigaVUE HC Series and GigaVUE TA Series Nodes. • All ports, excluding BPS ports, of same type and speed can be used to create GigaStream.	

GigaVUE-HC1-Plus	• 1RU Footprint • Two Module Slots (Bays) and one Fixed Base Module • 4 x 100Gb / 40Gb, 8 x 25Gb, 10Gb, 1Gb connectivity • Supports GigaSMART applications with a fixed rear GigaSMART Module. • Supports Flex Inline in unprotected ports with 10Gb/40Gb/25Gb/100Gb/4 x 10Gb/4x25Gb speed Nodes. • Cluster with GigaVUE HC Series and GigaVUE TA Series Nodes • All ports, excluding BPS ports, of same type and speed can be used to create GigaStream.	
GigaVUE-HCT	• 1RU and half rack width system • One Module Slot (Bay) and one Fixed Base Module • 2 x 100Gb/40G bconnectivity • Supports Flex Inline in unprotected ports with 40Gb/100Gb/4x10G/4x25G speed Nodes. • Cluster with GigaVUE HC Series and GigaVUE TA Series Nodes • All ports, excluding BPS ports, of the same type and speed can be used to create GigaStream.	
GigaVUE-TA25	• 1 RU Footprint • 1Gb/10Gb/48 x 25Gb/ ports and 8 x 100Gb/40Gb ports, dual hot-pluggable power supplies (AC/DC), four rear hot swappable fan modules, two console ports, and a 10Mb/100Mb/1Gb management port • Optional patch or breakout panel support • Cluster with GigaVUE HC Series and GigaVUE TA Series Nodes	

GigaVUE-TA25E	• 1 RU Footprint • 48 x 25Gb/10Gb/1Gb ports and 8 x 100Gb/40Gb port, dual hot-pluggable power supplies (AC/DC), four rear hot swappable fan modules, two console ports, and a 10Mb/100Mb/1Gb management port • Optional patch or breakout panel support • Cluster with GigaVUE HC Series and GigaVUE TA Series Nodes.	
GigaVUE-TA100	• 1RU Footprint • 32 x 100Gb/40Gb ports, hot swappable fan modules • Optional patch or breakout panel support • Cluster with GigaVUE HC Series and GigaVUE TA Series Nodes.	
GigaVUE-TA200	• 2RU Footprint • 64 x 100Gb/40Gb ports, hot swappable fan modules • Optional patch or breakout panel support • Cluster with GigaVUE HC Series and GigaVUE TA Series Nodes.	
GigaVUE-TA200E	• 2RU Footprint • 64 x QSFP28 ports (40G/100G), dual power supply modules (AC/DC) and hot swappable fan modules • UART RS232 Console port (RJ45) and Management port (RJ45).	
GigaVUE-TA400	• 1RU Footprint • 32 x 400Gb /100Gb/ 40Gb QSFP-DD/ QSFP28/QSFP+ ports, dual hot-pluggable power supplies (AC/DC), seven rear hot swappable fan modules, console port and a 10M/100M/1G management port. • Cluster with GigaVUE HC Series and GigaVUE TA Series Nodes.	

GigaVUE-TA400E	▪ 2RU Footprint ▪ 32 40Gb/100Gb/400Gb QSFP-DD/ QSFP28/QSFP+ ports, dual hot-pluggable power supplies (AC/DC), six hot-swappable fan modules, console port and a 10M/100M/1G management port. ▪ Breakout panel support ▪ Console port (RJ45) and Management port (RJ45). ▪ In addition provides 2 x 10Gb/1Gb SFP+/SFP ports ▪ Cluster with GigaVUE HC Series nodes and GigaVUE TA Series Nodes.	

Notes on TA Series Nodes

- On the GigaVUE-TA100, only the first 16 out of 32 100Gb ports are enabled. Two port licenses are available to enable an additional 8 or 16 ports to expand from 16 ports to 24 ports or from 16 ports to 24 ports and then to 32 ports.
- The ports on the GigaVUE-TA100 can be used as network, tool, or hybrid ports.
- For more information about the GigaVUE TA Series nodes, refer to the *GigaVUE TA Series Hardware Installation Guide*.

For adding a physical node to GigaVUE-FM, refer to Configure Physical Nodes.

About Cluster

You can use GigaVUE-FM to create a cluster. Cluster is created from standalone nodes that are currently managed by GigaVUE-FM. The type of cluster that can be created in this software version is out-of-band.

Any GigaVUE® HC Series and GigaVUE® TA Series nodes can be a part of a cluster. However, a GigaVUE® TA Series node cannot be a leader. It can only join a cluster with other GigaVUE HC Series nodes.

In addition to creating a new cluster, you can also manage an existing cluster through GigaVUE-FM. You can add nodes to an existing cluster and remove nodes from an existing cluster.

When a new cluster is created, the nodes joining the cluster must be standalone nodes. If a node is initially part of another cluster, it must be removed from that cluster so that it becomes a standalone node, before it can be added to the new cluster.

Refer to the following notes and considerations for all nodes in a cluster:

- must be currently managed by GigaVUE-FM
- must be running software version 5.1.00, at a minimum
- must be running the same software version and build version
- must be reachable by GigaVUE-FM, that is, it must be online and not have any authentication failures
- GigaVUE TA Series nodes that need an Advanced Features License, must be licensed
- member nodes joining the cluster must have preference less than the leader node
- From GigaVUE-FM 5.7 and above , a node which is prior 5.4.00 will show its installed "Cluster License" as “Advanced Features License”.

For information on clustering concepts, refer to Manage GigaVUE® Nodes and Clusters.

Overview of Seed Node

The seed node concept in GigaVUE-FM is different from the cluster leader role.

When a cluster is created, one of the nodes that you have selected for inclusion in the cluster will be deemed as the seed node. The seed node will be used to start the formation of a cluster and will be determined by the cluster leader preference settings of the nodes selected for the cluster.

You can override the seed node selected by GigaVUE-FM. However, the seed node must be a node that has the ability to become the leader.

Initially, the seed node is the source of the configuration information for the other nodes in the cluster. However, the cluster still consists of a leader, a standby node, as well as normal nodes. With the addition of more nodes to the cluster, a new cluster leader may be desired. If the desired leader is different from the seed node, the leader will then become the source of the configuration information for the other nodes in the cluster.

For creating and managing clusters using GigaVUE-FM, refer to Create and Manage Clusters.

Manage GigaVUE® Nodes and Clusters

This section describes how to add and manage GigaVUE® HC Series and GigaVUE® TA Series nodes on a GigaVUE-FM:

- Configure Physical Nodes describes the process to add, configure and manage GigaVUE nodes through GigaVUE-FM.
- Create and Manage Clusters describes the process to create a cluster using the wizard, add nodes to a cluster, remove nodes from a cluster, edit cluster parameters, and add stack links.
- Upgrade Software on a GigaVUE Node or a Cluster from GigaVUE-FM describes the process to upgrade standalone nodes and clusters through GigaVUE-FM.
- Search for Specific Nodes Using Keywords provides information about each of the standalone nodes and clusters, including a visual indication of each nodes status.
- Overview Page provides the information on each node connected to the GigaVUE-FM. This section covers the following: Systems Information, Ports Information, and Traffic.
- Workflows describes how to use the workflow wizards to create four different types of maps. With the wizards, you can create the following types of maps:
  - Out-of-band maps
  - Inline maps
  - Basic out-of-band GigaSMART maps
  - Advanced out-of-band GigaSMART maps
- Chassis Table View describes the Chassis Table View when managing a node with GigaVUE-FM.
- Safe and Limited Modes describes Safe Mode and Limited Mode.
- IPv6 based Clustering Support
- Rules and Recommendations for Nodes and Clusters

Topics:

Configure Physical Nodes

The Physical Nodes page displays a list of physical nodes and clusters. It provides information about a device's cluster ID, role, model, connection status, device status, and many other details.

To access physical nodes attached to an instance of GigaVUE-FM, log into GigaVUE-FM. On the left navigation pane, go to **Inventory** and under **Physical**, select **Nodes** to view all the physical nodes and clusters managed by GigaVUE-FM.

The Physical Nodes page displays the following information:

Field	Description
Cluster Id	The name of the cluster.
Host Name	The host name of the box.
Task Status	The status of the upgrade process. When the upgrade process is in progress, the task status displays the number of steps completed successfully out of the total number of steps to be completed. For example, upgrade: step (2/5) Image Fetch Complete. Once the upgrade process is complete, the upgrade status is displayed as Upgrade Success or Upgrade Failure. **Note**: To clear the task status of the upgrade process, select the required cluster IDs and choose **Actions > Clear Task status** option.
Node Address	The IP address of the physical node. **Note**: When assigning an IP address to a device using DHCP, the IP address may change after a reload or device upgrade, resulting in lost connectivity. Previous entries are retained in the node, and stale entries are retained in the GigaVUE-FM database even though the configuration has been deleted from the GigaVUE-FM end. To resolve this issue, Map the IP address to the Mac address of the node in the DHCP server so that GigaVUE-FM won't end up in this state due to the IP address change.
Role	The role of the node in the cluster. The role of the node can be one of Leader, Standalone, Member, or Standby.
Model	The type of the GigaVUE HC Series model.
Box Id	The box Identifier of the node.
Serial Number	Serial number of the device.
SW Version	The version number of GigaVUE-FM.
Licensed	The status of the physical node or Advanced Features license.
Attempted Sync Time	The time and date when the physical node attempted synchronizing with GigaVUE-FM.
Successful Sync Time	The time and date when the physical node was last synchronized with GigaVUE-FM.
Health	The current health status of the GigaVUE node or cluster. To know about how the device health status is computed, refer to Node Health Status.

Field	Description
	Note: You can monitor the health status of the device by enabling the SNMP notifications. For more information on configuring the email notifications, refer to the *"Notifications"* section in the *GigaVUE Administration Guide*.
Alarm	Alarm triggered by the device.
Tag	The tag or site name and value associated to the physical node or cluster. The tag names associated to the physical node or cluster are displayed as separate columns. Under the tag or site name, the respective tag or site value is displayed.
Uboot Version	Uboot version of the device.

NOTE: The columns in the Physical Nodes page can be customized based on the type of content you want to view in the table. For customizing the columns, refer to Table View Customization.

Changes made to the cluster through the CLI are reflected in GigaVUE-FM when it synchronizes with the cluster, which is typically every 5 minutes.

If the latest configuration data is not retrieved from the cluster for more than 30 minutes, a warning is displayed in the cluster Overview page indicating the last time GigaVUE-FM successfully synchronized with the cluster.

To view the last synchronized status, click the cluster and view the status at the top of the Overview page.

Node Control Options

Use the following buttons in the Nodes page to perform specific actions:

Button	Description	Reference
Tags	The Tags **drop-down** menu button allows you to perform the following actions: • **Add**: Use to **Add** tags to standalone nodes and clusters. • **Delete**: Use to **Remove** tags from standalone nodes and clusters • **Device Level Tagging**: Use to associate tags to devices. Refer to the Device Level Tagging section for more details.	Refer to the Tags section in the *GigaVUE Administration Guide* for details on how to create, edit, delete tags.

Button	Description	Reference
	• **Export Tags for Selected**: Use to **Export** tags for the selected nodes. • **Export Tag Resources for Selected**: Use to Export tag resources for the selected nodes. **Note**: To create tags you must be a **fm_ super_admin** user.	
Actions	The **Actions** drop-down menu option allows you to perform the following actions:	
	• **Edit**: Use to **Edit** the selected standalone node or cluster.	
	• **Image Upgrade**: Use to upgrade the image for the selected node or cluster.	Refer to Upgrade Software on a GigaVUE Node or a Cluster from GigaVUE-FM
	• **Backup**: Use to backup the nodes and cluster. • **Restore**: Use to restore the nodes and cluster.	Refer to the Backup and Restore section for more details.
	• **Reboot**: Use to reboot the nodes and cluster.	Refer to the Reboot the Nodes section for more details.
	• **Suppress Alarms:** Use to suppress alarms. • **Stop Alarm Suppression**: Use to stop suppression of alarms.	Refer to Suppressed Alarms section for more details.
	• Clear Task Status: Use to clear the Task status.	Refer to Admin Tasks for more details.
	• **Rediscover:** Use to rediscover the node.	
	• **Disconnect Node**: Use to disconnect the node from GigaVUE-FM.	
Filter	The **Filter** button allows you to filter the nodes.	
Create Cluster	Use to **Create Cluster**.	Refer to Create and Manage Clusters section for more details.
Add	Use to **Add** a physical node to GigaVUE-FM.	Refer to Add New Physical Node or Cluster to GigaVUE-FM.
Delete	Use to **Delete** a node from GigaVUE-FM.	

Button	Description	Reference
Import	Use to **Import** nodes to GigaVUE-FM.	Refer to Add Nodes From an Excel Spreadsheet section for more details.
Export	Use to **Export** the nodes from GigaVUE-FM.	Refer to Export Nodes and Clusters section for more details.

Add New Physical Node or Cluster to GigaVUE-FM

You can add physical nodes and clusters to GigaVUE-FM either manually or by importing an Excel spreadsheet. However, before adding a new physical node, ensure that the node credentials are added under the **System** > **Node Details**.

The node credentials (username/password) used to add a node in GigaVUE-FM must be configured in the node and must have the necessary privileges for read and write operation. If you add a node with read-only privileges, it will impact the device operations performed from GigaVUE-FM.

NOTE: In a cluster configuration, the Normal nodes are seen as Member nodes in GigaVUE-FM.

Add Nodes Manually

To add physical nodes manually, do the following:

1. On the left navigation pane, go to **Inventory** and select **Physical** .
2. On the Physical Nodes page, click **Add**. The Add Physical Node page displays.
3. Select **Add Manually**.
4. Enter or select the information as shown in the following table:

Field	Description
Node	DNS name or IP address of the node. Click + to add additional nodes or - to remove a node. **Note**: You must add the nodes by their FQDN or hostname rather than IP, especially in case of NAT nodes. This ensures that GigaVUE-FM and device communication is not interrupted when there is a change in the IP address of

Field	Description
	the node. When adding clusters, use the VIP of the cluster to add. If VIP is not configured, they you must add the leaders FQDN or IP, which is the primary address of the cluster. **Port** The HTTPS port number of the device, if it has been changed using CLI. **Note**: If you do not provide the HTTPS port number, then the default value 443 will be used for device communication.
Secondary IP Address	IP address (nodeAddress) of the node that is used to connect to the cluster, in case the primary IP address is not reachable. This node will be the one which becomes the leader if the leader fails. **Note**: Secondary IP address is applicable only for clusters configured without VIP. You cannot add a secondary IP address for standalone nodes. You can edit an existing cluster and add the secondary IP address. In case both the primary and the secondary IP address is not reachable the cluster is marked as unreachable.
Custom Login Credentials	Check box to enable custom login credentials. If enabled enter the following details: ▪ User Name ▪ Password
SNMP Version	The SNMP version to be used to register GigaVUE-FM as a SNMP target on the node.
Node behind NAT	Check box to indicate if the node being added is behind NAT. Refer to Add NAT Behind Nodes for details.
Tags	Select the tag key and the tag value to which the node must be associated to. **Note**: You can only view tags that are permitted for your role. Refer to the "Tags" section in the *GigaVUE Administration Guide* for more details.

5. Click **Submit** to add the node or nodes to the list of physical nodes GigaVUE-FM is managing.

For standalone nodes, both node IP (Device IP) and cluster ID is the node Address (nodeAddress) provided by the user while adding the node.

NOTE: In GigaVUE-FM version 5.5.01, the Node IP was the actual IP and the Cluster ID was either the DNS ID (if the IP address was resolved) or the Node IP (if the IP address was not resolved).

Add Nodes From an Excel Spreadsheet

You can add nodes to GigaVUE-FM by uploading an Excel spreadsheet that contains a list of the physical nodes that you want GigaVUE-FM to manage. You can create the spreadsheet or use a spreadsheet from a previous export of the nodes. From GigaVUE-FM 6.7 release onwards, the format of the spreadsheet must have at least a column with the node addresses and a separate column for the HTTPS port number (if the port number of the node has been changed using CLI) and the header Node Address.

- If you do not provide the HTTPS port number in the Excel spreadsheet, then the default value 443 will be used for device communication.
- For software version greater than 5.11.00, Node_IP should be changed to Node Address to avoid any errors while importing nodes using Excel spreadsheet.

Host Name	Node Address	Model	Box Id	HTTPS Port Number	Serial Number	SW Version	Licensed	Attempted Sync Time	Successful Sync Time	Alarms
gigamon-a0fa62	10.114.34.113	HC1	1	443	HFA62	6.7.00	Yes	2024-06-20 11:23:12	2024-06-20 11:23:12	0
gigamon-a046fc	10.114.34.196	HC3	7	443	J46FC	6.1.00	Yes	2024-06-20 11:23:12	2024-06-20 11:23:12	0
MSWMNAJ9C01339	10.114.84.141	HC3	2	443	JBEE9	6.5.00	Yes	2024-06-20 11:23:11	2024-06-20 11:23:11	2
gigamon-951a34	10.114.178.10	HC3	3	443	J1A34	6.7.00_Beta	Yes	2024-06-20 11:23:10		0
gigamon-a0cd12	dev1.fmtest.lab	HC3	4	8443	JCD12	6.6.00	Yes	2024-06-20 11:23:11	2024-06-20 11:23:11	0
gigamon-a08802	10.114.43.207	HC1	7	443	H8802	6.6.00	Yes	2024-06-20 11:23:14	2024-06-20 11:23:14	0
gigamon-a00e52	10.114.43.79	HC3	4	443	J0E52	6.6.00	Yes	2024-06-20 11:23:14	2024-06-20 11:23:14	0
gigamon-a0e448	10.114.42.213	HC1	3	443	HE449	6.6.00	Yes	2024-06-20 11:23:14	2024-06-20 11:23:14	0

Figure 1 *Node List Spreadsheet for Import/Export*

NOTE: The format of the spreadsheet is changed in GigaVUE-FM 6.7. To import a spreadsheet created by GigaVUE-FM prior to GigaVUE-FM 6.7, you can modify the previous spreadsheet or do a new export after upgrading to the current version of GigaVUE-FM. An example of a spreadsheet is shown in the figure.

You can also import nodes to GigaVUE-FM by uploading an Excel spreadsheet that contains only the Node Address, username, and password of the nodes that you want to import. An example of such spreadsheet is shown below:

Cluster ID	username	password
10.114.34.119	admin	admin123A!!

To add physical nodes by importing from a spreadsheet, do the following:

1. On the left navigation pane, go to **Inventory** select **Physical** .
2. On the Physical Nodes page, click **Import**.

   The Add Physical Node page displays with the **Import from Excel** option automatically selected. The page displays an area for selecting or dragging and dropping a file.
3. Drop an Excel file onto the page or click **Select File** to upload the file.

   The page displays the list of IP address.
4. Select the nodes to add to GigaVUE-FM.
5. Click Submit to add the node or nodes to the list of Physical Nodes GigaVUE-FM is managing.

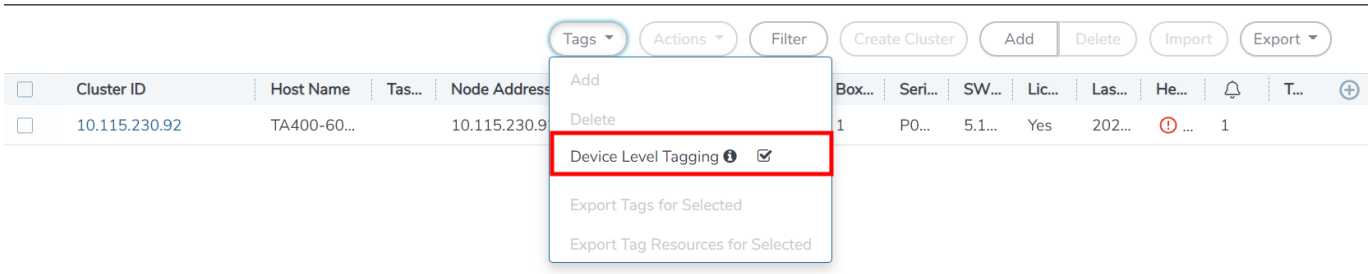
NOTE: If the nodes/clusters that are being imported in to GigaVUE-FM are associated to tags, you must ensure to add the tags to that GigaVUE-FM instance before importing the nodes. Otherwise, the nodes will not be associated to the tags. Refer to the GigaVUE Administration Guide for more details.

Device Level Tagging

Each of the individual nodes in a cluster can be associated with a tag key and a tag value. To assign a tag key to a device:

1. On the left navigation pane, go to **Inventory** go to **Physical** > **Nodes**.
2. Click on **Tags** and enable the **Device Level Tagging** option.
3. Select the required node(s). Click **Tags** and select **Add**. The **Add Tags to Resources** page appears.
4. Select the required tag key and the tag value to which the nodes must be tagged.
5. Click **OK**.

NOTE: Only aggregation tags can be configured to the devices.



The nodes in a cluster are now associated to tags which can be used for creating the required topology.

NAT Behind Nodes

GigaVUE-FM can manage nodes and clusters that use Network Address Translation (NAT) for IP address conversion. GigaVUE-FM identifies the nodes behind NAT based on the **Node Behind NAT** option which you must select at the time of adding the nodes. Refer to Add Nodes Manually.

Prerequisites

In order for GigaVUE-FM to manage the devices and clusters that are behind NAT:

- GigaVUE-FM must be running software version 5.7 or higher.
- Devices must be running software version 5.3 or higher.
- Devices must be assigned NAT IP and Fully Qualified Domain Names.
- Domain name server and default search domain must be configured in GigaVUE-FM.

Add NAT Behind Nodes

When managing clusters behind NAT, GigaVUE-FM does not know the NAT IP of the members in the cluster. You cannot also configure the NAT IP of the members. GigaVUE-FM only identifies the private IP of the nodes through the API. Therefore, for device specific operation, GigaVUE-FM uses host names to connect to the nodes. You must therefore configure the domain name server and default search domain in GigaVUE-FM. Refer to the *"NAT Configuration"* section in the *GigaVUE Administration Guide*. If you do not configure the domain name server and search domains, then it may lead to failure in node specific operations.

NOTE: When adding nodes, ensure the following:

- Select the **Node behind NAT** option for a node that is behind NAT.
- Do not select the **Node behind NAT** option for a normal node.

If you fail to perform the above operations, then you may experience inconsistency in managing the nodes.

Refer to the *"NAT Configuration"* section in the *GigaVUE Administration Guide* for details.

Limitations

Though GigaVUE-FM can manage nodes behind NAT from software version 5.3, SNMP traps, Syslog messages and event notifications will not be pushed from the node to GigaVUE-FM. This is because GigaVUE-FM registers its eth0 address as target address on the nodes to receive the notifications but the nodes behind NAT will not be able to reach GigaVUE-FM by its eth0 IP. These limitations are not applicable for node version 5.7 and above.

NOTE: Refer to the *"NAT Configuration"* section in the *GigaVUE Administration Guide* for details.

Cluster Discovery Behavior

GigaVUE-FM does not detect individual nodes that were part of a cluster, if the cluster was dismantled using the CLI. GigaVUE-FM always reaches the cluster by its virtual IP address (if configured). When a cluster is dismantled through the CLI, the virtual IP address of the cluster will no longer be available and the cluster is therefore marked as unreachable. There is no change in detecting node additions, removals, or membership changes performed from CLI. Refer to the Rules and Recommendations for Nodes and Clusters for details on Cluster Configuration through CLI.

Note: When you reboot a device in cluster , the device status appears as **Green** ,whereas the connection status of the device appears as **Connection Failed**, and the role of the device appears as **Standalone**. In the subsequent discovery cycle , the rebooted devices appear in the cluster with their **Role, Device Status** ,and **Connection Status.**

ARP/NDP Timer Settings

The Address Resolution Protocol (ARP) or the Neighbor Discovery Protocol (NDP) timer specifies the aging time on the IP interface. The ARP timer is used for IPv4 addresses and the NDP timer is used for IPv6 addresses. The timer is configurable from 3 to 30 seconds. The default is 30 seconds. When an IP interface is configured, ARP/NDP requests are sent to the IP interface to find the gateway MAC address. When Tunnel encapsulation GSOP Map is configured with destination tool in local network, ARP/NDP requests are sent to the IP interface to find the tool MAC address. In response, the gateway sends an ARP/NDP reply. The control card tries to match the IP address of the IP interface with the IP address of the ARP/NDP message received. If a match is found, the ARP/NDP status changes to resolved (otherwise the ARP/NDP status is not resolved).

Once the ARP/NDP status is resolved, the ARP/NDP timer of the IP interface controls the interval at which an ARP/NDP request is sent to the gateway to detect if the gateway is reachable or not. Above is applicable to detect tool's reachability part of local network that are configured as part of tunnel encapsulation GSOP Map.

You must enable the IP Gateway Status SNMP trap to send SNMP notifications when the ARP/NDP status changes. To enable SNMP notifications, refer to Enable or Disable Events for SNMP Notifications.

Change the ARP/NDP Timer Settings

The default ARP/NDP timer value is 30 seconds. To change the timer setting at the node-level:

1. On the left navigation pane, go to **Inventory** select **Nodes**.
2. In the Physical Nodes page, select the node for which you want to change the ARP/NDP timer setting.
3. Go to **Settings > Global Settings > ARP/NDP**.
4. Click **Settings**.
5. In the ARP/NDP Settings page, choose the required **ARP Refresh Time Interval** or the **NDP Refresh Time Interval** in seconds.
6. Click **OK**.

The ARP Entries table and the IPv6 Neighbor Entries table dynamically refresh to display information such as the IP address and Hardware address mapping, aging, state, and interface details.

Click **Clear > Clear ARP Entries** or **Clear > Clear IPv6 Neighbor Entries** to remove the entries from the tables.

Enable Gratuitous ARP on Management Interface

A Gratuitous ARP is an Address Resolution Protocol response that is sent without an ARP request. The Gratuitous ARP response is sent as a broadcast in the network. It is a mechanism by which a device can announce or update its IP address to MAC address mapping. This helps the neighboring devices to discover the device that has sent the Gratuitous ARP response and forward subsequent packets to the device without any delay.

You can enable the Gratuitous ARP on the management interface of a GigaVUE HC Series or GigaVUE TA Series device. Ensure that the management interface is configured with static IP address. A Gratuitous ARP response is sent to the network when:

- the device boots up.
- the IP address of the device changes.
- the link state of the management interface comes up.

NOTE: You cannot enable Gratuitous ARP on the management interface that is configured with DHCP.

To enable Gratuitous ARP on a management interface:

1. On the left navigation pane, go to **Inventory** and select **Nodes**.
2. Click on a node on which you want to enable Gratuitous ARP. The Overview page of the node appears.
3. In the left navigation pane, select **Settings > Interface > Protocol Configuration**.
4. Select the box ID on which you want to enable Gratuitous ARP, and then click **Edit**.
5. Select the **Gratuitous ARP** check box.
6. Click **Apply**. The Gratuitous ARP is enabled on the management interface of the device.

SNMPv3 Support

Starting with software version 5.4 GigaVUE-FM creates an SNMPv3 user during the upgrade process from version 5.3. The SNMPv3 user is created with same authentication and privacy settings previous held.

NOTE: SNMPv3 support is only available on GigaVUE-FM software version 5.4 user and SNMPv3 users created by GigaVUE-FM cannot be modified by users.

Enable SNMPv3 on Nodes

You can enable SNMPv3 from the node addition page, If node version 5.4 or higher and SNMPv3 is selected, the GigaVUE-FM registers itself as a SNMPv3 trap receiver. If the SNMPv3 is chosen and the node version is 5.3 or below then, FM registers itself as a SNMPv2 trap receiver. When this occurs GigaVUE-FM reports an event that GigaVUE-FM does not support SNMPv3 on the node with 5.3 or lesser version.

Enabling SNMPv3 During Upgrade

1. Go to **Settings** and select **System> Credentials.** The Credentials page displays a listing of the nodes and SNMP versions.
2. Select a node to upgrade, and click **Edit**.
3. Select SNMP version - SMNPv3
4. Click **Save**.

Enabling SNMPv3 on a New Node

1. Select **Physical>Nodes> Add.** The Add Physical Node displays.
2. Enter **Node Name**.
3. Click **Enable**.
4. Enter **Username** and **Password**.
5. SNMP Version: SNMPv3.
6. Click **Submit**.

Create and Manage Clusters

This section describes the GigaVUE-FM clustering. Refer to the following sections for details:

- About Cluster
- Create Clusters
- Support for Cluster Types
- Regular Cluster Formation Workflow
- Edit Cluster
- Add Nodes to a Cluster
- Remove Nodes from a Cluster
- Edit Cluster Parameters
- Check Cluster Status

About Cluster

A cluster consists of multiple GigaVUE-OS nodes operating as a unified fabric such that packets entering the cluster on one node can be sent to a destination port on any other node. You set up packet distribution using the standard box ID/slot ID/port ID format, allowing maps to distribute traffic to any port in the cluster.

The following figure illustrates a sample cluster of GigaVUE-OS nodes.

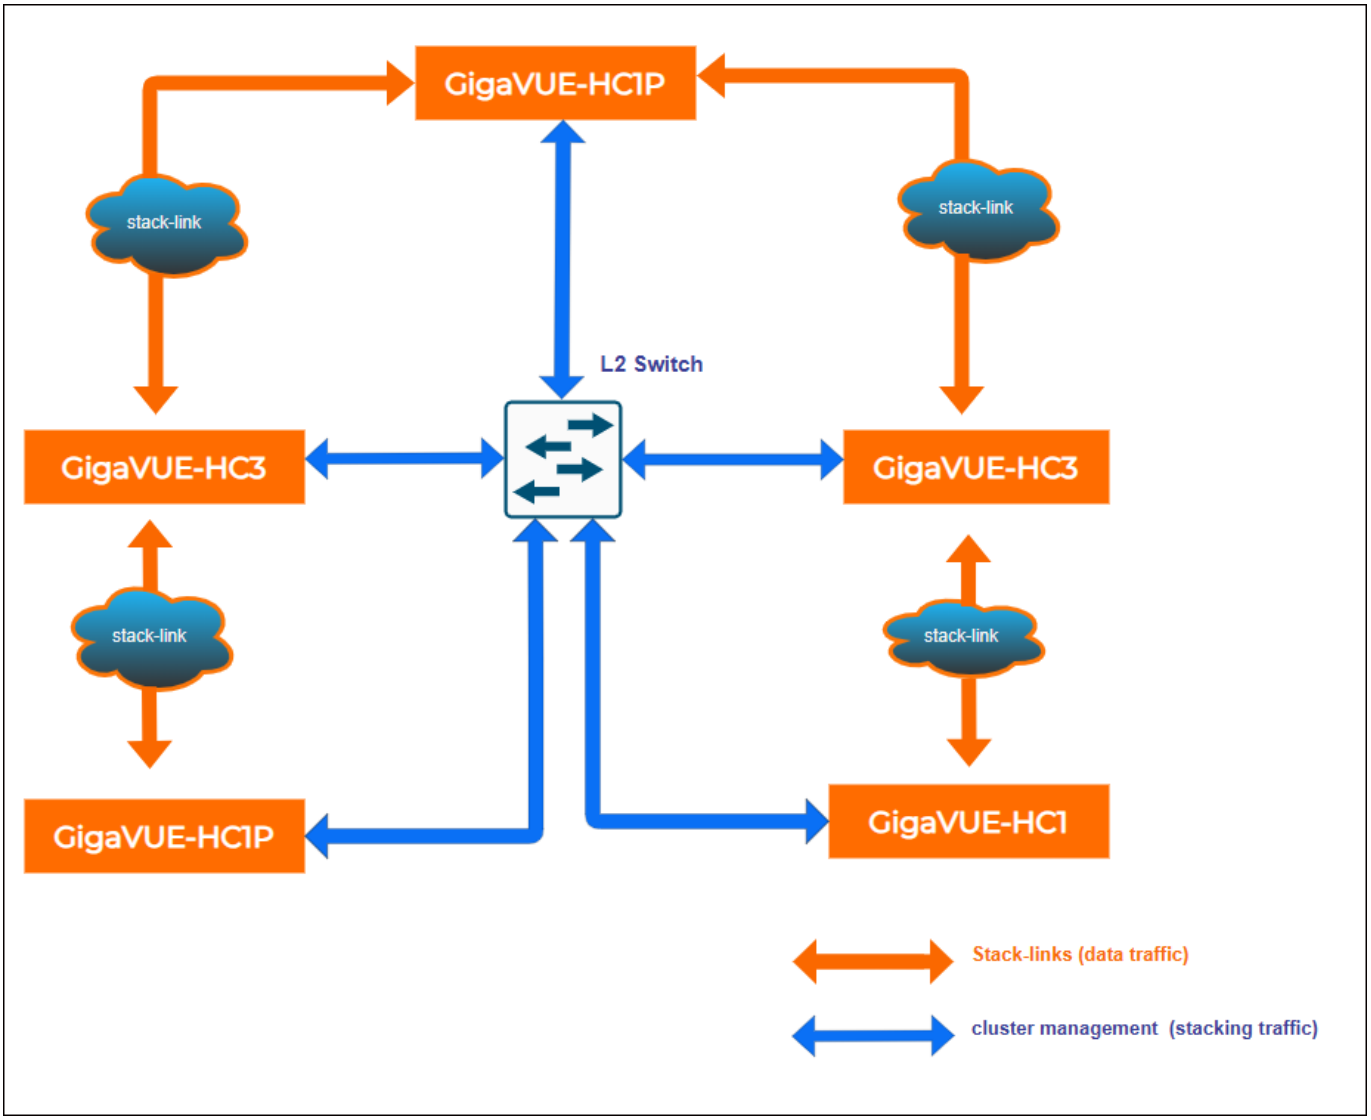


Figure 2 *Sample Cluster*

Cluster Node Limit

Any GigaVUE-OS nodes can join other nodes in a cluster. The GigaVUE-HC1, GigaVUE-HC3, and GigaVUE TA Series nodes, including Certified Traffic Aggregation White Box (white box), can be included in the same cluster. Starting in software version 4.5, the maximum number of nodes supported in a cluster is 32. For details, refer to Cluster Scaling.

In addition, there is another independent limit, which is for the maximum number of line cards supported in a cluster, across all of the nodes in the cluster. This limit is determined by cost units. Cost units are based on the total number of line cards, line card types, and chassis type. Cost units measure the resources that a node needs in a cluster. The higher the cost unit, the more resources are needed to manage the node.

The following table has examples of line card and chassis types, and their cost units.

Line Card Type or Chassis	Cost Unit
GigaVUE-HC1-Plus	1
GigaVUE-HCT	1
GigaVUE-HC1 Chassis	1
GigaVUE-HC3 Chassis	1
GigaVUE-HC3 Chassis with Control Card version 2 (HC3 CCv2)	1
GigaVUE-TA25	1
GigaVUE-TA25E	1
GigaVUE-TA100	1
GigaVUE-TA200	1
GigaVUE-TA200E	1
GigaVUE-TA400	1
Certified Traffic Aggregation White Box	1

For example, if a GigaVUE-HC3 has one PRT-HC3-C08Q16 line card, one SMT-HC3-C08 GigaSMART line card, and one BPS-HC3-Q35F2G + PRT-HC3-X24 line cards, the cost units are 4 + 1 = 5.

The cost unit maximum for the various nodes in the cluster is 255.

Therefore, the largest cluster supported is determined either by the maximum number of nodes (32) or by the cost unit maximum, whichever is reached first.

Cluster Scaling

The maximum number of nodes and map rules supported in a cluster is as follows:

Table 1: Maximum Number of Nodes and Map Rules Supported in a Cluster

When a Cluster is Configured with:	Number of Nodes	Maximum Map Rules
Out-of-Band Cluster Management	32	38000
Inband Cluster Management	16	38000

The maximum number of map rules supported in a cluster apply to all nodes in the cluster including GigaVUE® HC Series nodes: GigaVUE-HC3, GigaVUE-HC1, and GigaVUE TA Series nodes: GigaVUE-TA100, and GigaVUE-TA200.

Cluster Topologies

The following cluster topologies are supported:

- star
- daisy-chain
- tree

Separate Paths for Cluster Control and Stack Traffic

There are two types of clustering: out-of-band and inband.
The nodes in a cluster are constantly communicating with one another, exchanging heartbeats to check on one another's status, exchanging configuration information so that changes made on the leader are propagated to other nodes, and making changes to cluster roles based on changes in status.
GigaVUE-OS separates cluster control traffic from the actual flow of packets from ingress ports on one node to egress ports on another for the two types of clustering, as follows:

Out-of-Band Clustering on Mgmt Ports	With out-of-band clustering, cluster control traffic is carried out-of-band on its own network as follows: • GigaVUE-HC3 nodes use the Mgmt port (eth0) . • GigaVUE-HC1 nodes use the Mgmt port (eth0). • GigaVUE-HC1-Plus nodes use the Mgmt port (eth0). • GigaVUE-HCT nodes use the Mgmt port (eth0). • GigaVUE-TA25 nodes use the Mgmt port (eth0). • GigaVUE-TA25E nodes use the Mgmt port (eth0). • GigaVUE-TA100 nodes use the Mgmt port (eth0). • GigaVUE-TA200 nodes use the Mgmt port (eth0). • GigaVUE-TA200E nodes use the Mgmt port (eth0). • GigaVUE-TA400 nodes use the Mgmt port (eth0). **Cluster Management port(s)**. Using the cluster management port(s) lets you route cluster control traffic over a separate network from the network used to access the Mgmt port. This prevents cluster control traffic from overloading the traffic used to access the Mgmt port. **Mgmt Port (eth0)**. You can also select to use the standard Mgmt port for cluster control traffic. In this implementation, cluster control traffic uses the Mgmt port's Ethernet connection.
Inband Clustering on Stack-Links	Stack-links are used to create a stacking connection between two GigaVUE nodes in a cluster. Stack-links carry traffic from network ports on one node to tool (or GigaSMART) ports on a destination node. With inband clustering, cluster control traffic is carried inband through the stack-link. Stack-links can be constructed out of individual stack ports, for example, a 40Gb port on a PRT-HC0-Q06 or stack GigaStream. You decide which to use with the **gigastream** and **ports** arguments in the **stack-link** command. For example, the following command creates a stack-link between the q1 40Gb port on box 1/slot 1 and the q1 port on box 2/slot 7: **(config) #stack-link alias biglink between ports 1/1/q1 and 2/7/q1 comment "40Gb Stack"** Stack links are supported at speeds of 10Gb, 40Gb, and 100Gb. Refer to the *Hardware Installation Guide* for each GigaVUE node for information on stack link support. Keep in mind that because of the 10Gb port density offered by GigaVUE-OS nodes, using only one 10Gb port for a stack-link could cause a serious bottleneck. A stack GigaStream dramatically increases the bandwidth available for stack-links, letting you connect GigaVUE-OS nodes in a cluster and still take advantage of the 10Gb port density. Alternatively, nodes with 40Gb or 100Gb ports can take advantage of their high bandwidth for stack-links. When using stack GigaStream for stack-links, you must create a stack GigaStream on each side of the stack-link and each side must consist of the same number of ports running at the same speed.

About Cluster Roles

Communication with a GigaVUE-OS cluster is accomplished using a leader virtual IP address assigned to the cluster as a whole. Physically, the virtual IP address resolves to only a single leader at any one time. However, the leader role on the GigaVUE-OS node is not statically assigned to a single node. Instead, any node (except GigaVUE TA Series and the nodes residing on a different management subnet) in the cluster can take on the leader role if the situation requires it (for example, if both the leader and the current standby nodes go down).

When a new node becomes the leader, it takes ownership of the virtual IP address used for leader access to the cluster. Because all of the nodes in the cluster share the same database, this transition takes place seamlessly, ensuring that the cluster survives the failure of one or more nodes.

The virtual IP address is assigned to the primary control card in the configuration jump-start wizard:

Step xx: Cluster mgmt virtual IP address and masklen? [0.0.0.0/0]

Each node in the cluster is performing one of the following roles at any given time:

- **Leader** – This node has possession of the cluster's virtual IP (VIP) address and takes responsibility for dispatching commands to the entire cluster.
- **Standby** – This node takes over the leader role in the event of a failure on the node currently holding the role.
- **Normal** – These nodes perform normal GigaVUE operations with minimal cluster responsibilities. However, they, too, have a complete copy of the cluster's database. When a leader fails and standby is promoted to be the new leader, an election process takes place automatically between all normal nodes, ensuring that a new standby is found.

Setting a Node's Priority in the Leader Election Process

Clusters of GigaVUE-OS nodes perform a leader election in the following situations:

- Cluster reload
- Leader or standby node failure

In either of these cases, a new node is selected to perform the necessary role(s). You can set the **cluster leader preference** for each individual node in the cluster to specify how likely the node is to claim a leader/standby role. Higher values are more likely to claim the leader/standby role; lower values are less likely.

Use preference settings from 10 to 100 for leader, standby, and normal roles. Use preference settings from 1 to 9 for normal nodes that are excluded from taking the leader or standby role.

In software version 4.5, the preference cannot be set to zero. A node with a preference of 0 in an earlier software version will be changed to 1 after upgrading to 4.5 or higher.

GigaVUE-OS sets defaults for the **preference** argument based on the type of control card in use. If you choose to change a node's **preference** setting, it is generally preferable to set higher priorities for nodes with more processing power. GigaVUE-HC3 node provides the most processing power, followed by GigaVUE-HC1 nodes, followed by GigaVUE TA Series nodes.

> **NOTE:** All GigaVUE TA Series nodes including the white box, will automatically be added to a cluster with preference set to 1 because any Traffic Aggregator can never take the role of, or be eligible to be, the leader.

Note: The Clustering Daemon (Clusterd) restarts ,if the no card slot 1/4 down force command is executed after performing a cluster reload.

In addition, in an event of a cluster reboot, any GigaVUE TA Series node in a cluster may show as standby for a couple of minutes while the cluster is coming up from the reboot cycle. However, once the cluster is up and running, none of the GigaVUE TA Series nodes can be a standby.

About the "Unknown" Cluster Role

In addition to the standard roles in About Cluster Roles, the system may occasionally report a node operating with an **unknown** cluster role. A node with an unknown cluster role is no longer being actively managed by the leader.

When a node that was formerly part of a cluster transitions to an **unknown** cluster role, its database will typically be out of synchronization with the leader's. You can restore the node to the cluster by using the **reset factory keep-all-config** command, followed by a reboot, and running **configuration jump-start** to rejoin the cluster with a clean local database.

> **NOTE:** If a leader gets unknown cluster role, do not perform any configuration on the leader node as database sync may not happen.

Sample Cluster Control Connections

The GigaVUE-OS provides a flexible approach to cluster control traffic, allowing you to route it over cluster management or Mgmt ports. The ports available and their eth**x** designations vary by control card version and node type, as summarized in the following table:

Control Card/Node Type	Possible Cluster Control Ports	Deployment Models
HCCv2 Control Card	**Mgmt (eth0)**	▪ Mgmt (eth0) and L2 switch
GigaVUE-HC3 Node	**Mgmt (eth0)**	▪ Mgmt (eth0) and L2 switch
GigaVUE-HC1 Node	**Mgmt (eth0)**	▪ Mgmt (eth0) and L2 switch
GigaVUE TA Series Nodes	**Mgmt (eth0)**	▪ Mgmt (eth0) and L2 switch

Sample Cluster Control Configurations

Nodes in the same cluster must use the same cluster interface. For example, if there is a GigaVUE-TA100 in the cluster, all nodes in the cluster must use eth0.

Zeroconf for Cluster Management Ports

By default, cluster management ports use zero configuration networking (zeroconf) to establish networking settings. This eases configuration when establishing clusters using the cluster management port(s).

Keep Cluster Management Ports Connected!

IMPORTANT: Clusters implemented using the cluster management ports for cluster control traffic must ensure that the cluster management ports of all nodes in the cluster are connected at all times. This prevents a situation where multiple leaders claim the management VIP address, resulting in the inability to connect to it at all.

Sample Stack-Link Configurations

This section illustrates some sample configurations for the data-carrying stack-links in a cluster, including a star configuration and a daisy-chain using a GigaStream. You can see a combination of star and daisy-chain in About Cluster.

IMPORTANT: Ensure that you do not cable the stack-links in a loop. Use a star or daisy-chain configuration, as follows:

Star Configuration

Use a GigaVUE-OS node as the hub in a star configuration. This makes it easy to create a star configuration that maximizes traffic distribution efficiency. With a star configuration, no destination is further than two hops away. Note that the following image only shows the stack-link connections and not the cluster control connections from the control cards.

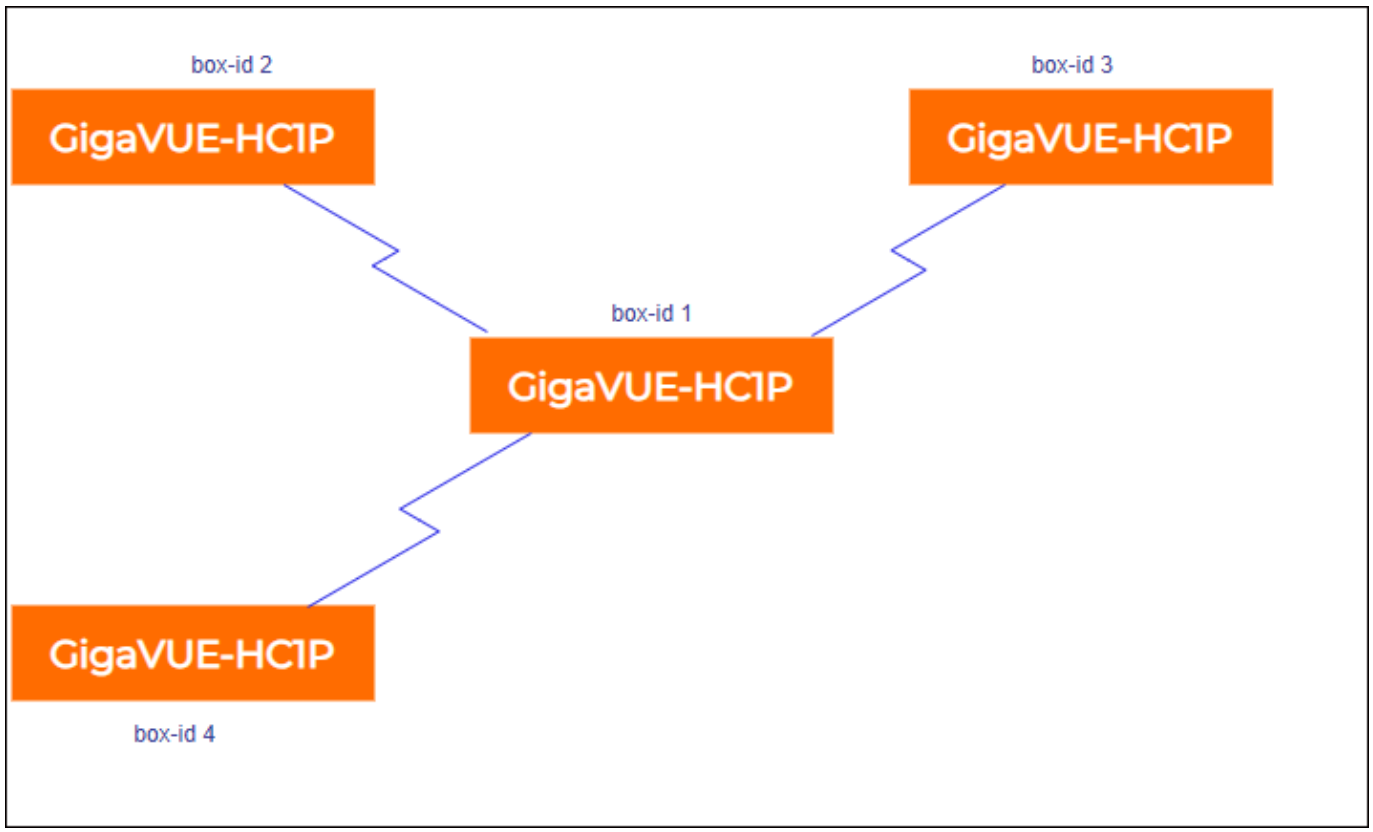


Figure 3 *Star Configuration*

Daisy-Chain Configuration Using GigaStream for Stack-Link

You can connect two GigaVUE-OS nodes together in a daisy-chain, for example, using any 10Gb line card port. Since ,there can be up to 96 10Gb ports on a single GigaVUE-HC3 node, the stack-link needs enough bandwidth to handle expected cross-node traffic volume. Create a stack GigaStream out of up to 24x10Gb (PRT-HC3-X24) or 8x100Gb (PRT-HC3-C08Q16) or 16x40Gb (PRT-HC3-C08Q16) ports to handle expected cross-node traffic loads.



Figure 4 *Daisy-Chain Configuration using GigaStream for Stack-link*

Creating Clusters: A Roadmap

Setting up a cluster consists of the major steps shown in the following figure:

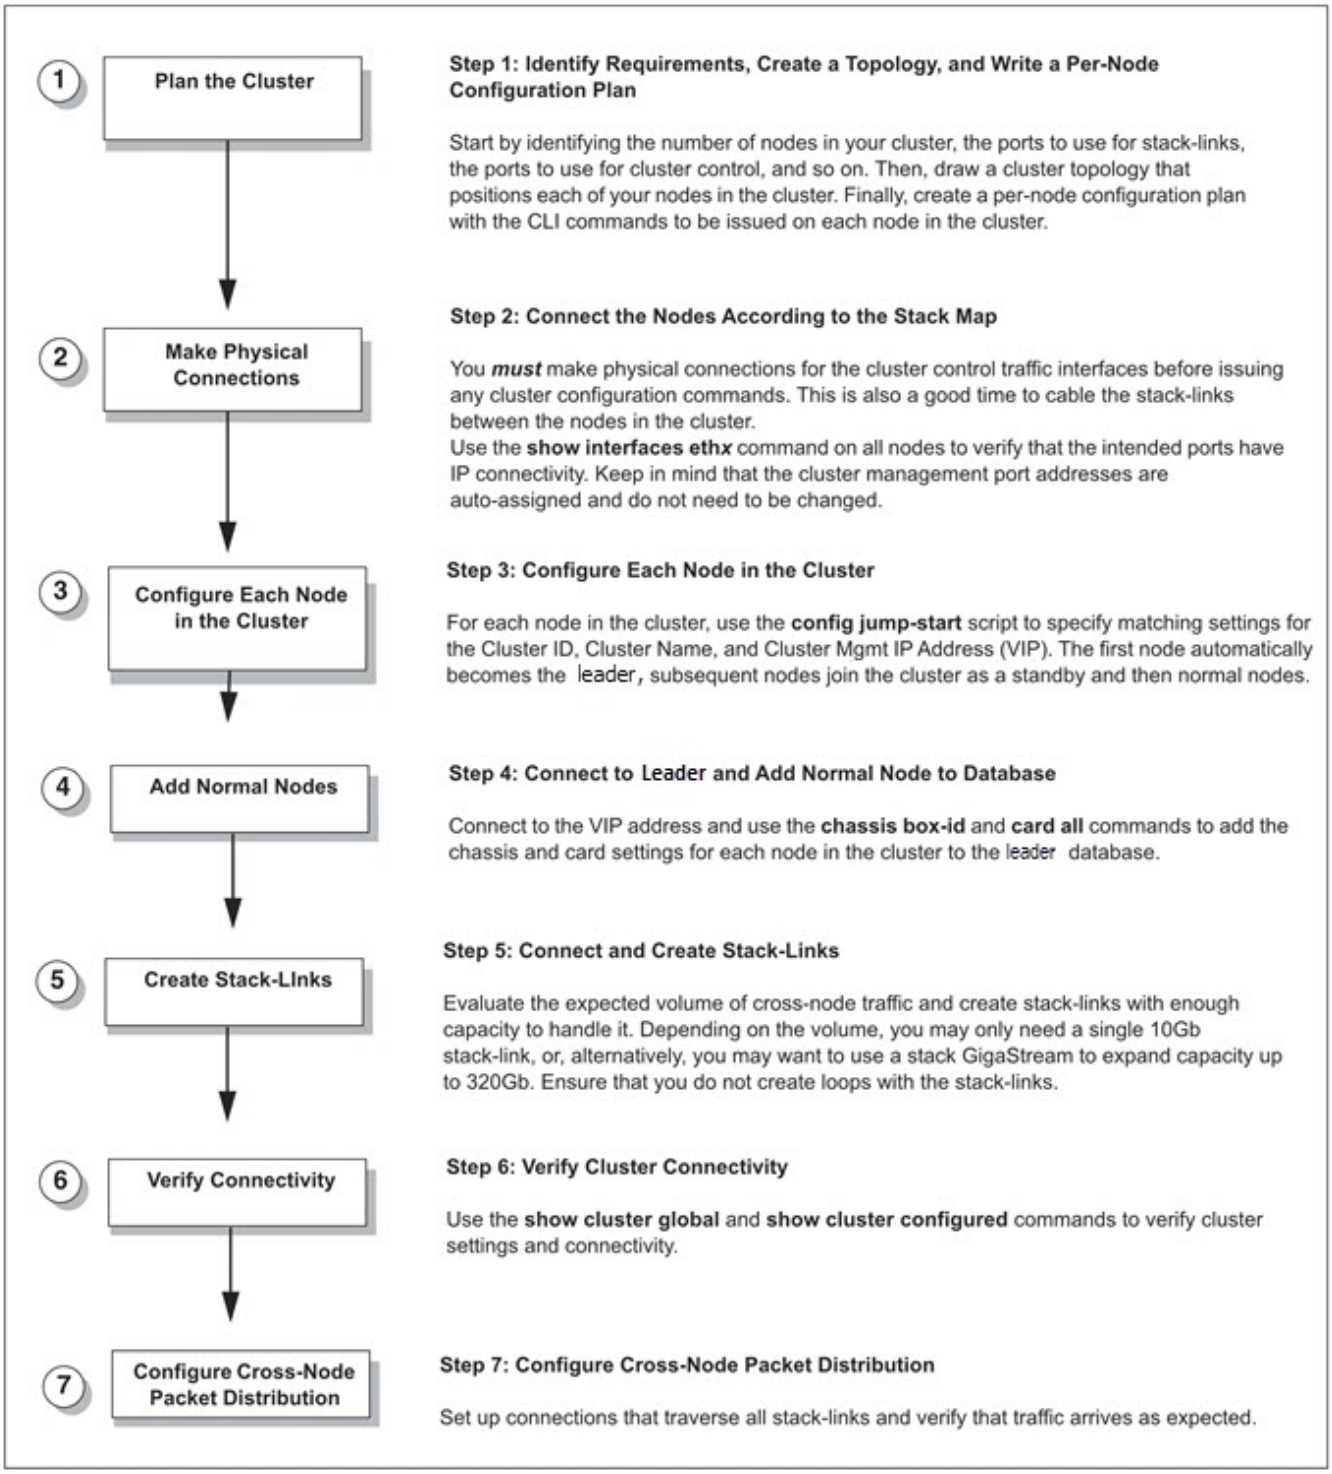


Figure 5 *Creating Clusters: Major Steps*

IPv6 Based Clustering

Until software version 6.1.00, you can create new clusters or manage an existing cluster only using the IPv4 protocol. Starting from software version 6.2.00, support for IPv6 based clustering is available based on which you can create new clusters and manage existing clusters using IPv6 protocol.You can configure the cluster topologies and perform cluster-related operations using IPv6 stacking. IPv6-based clustering is applicable for both GigaVUE HC series devices and GigaVUE TA Series devices.

Rules, Notes, and Limitations

Refer to the following rules, notes, and limitations:

- IPv6 based clustering is available only for the following types of clusters:
  - Physical device out-of-band clustering: Auto-discovery of IPv6 cluster is supported.
  - Layer 3 out-of-band clustering: Manual discovery of the IPv6 cluster using the primary IP address, secondary IP address and cluster VIP.
- IPv6 support for Inband clustering is not supported.
- You can create new IPv6 based clusters from both GigaVUE-FM and the CLI. However, switching an existing IPv4 based cluster to IPv6 based cluster can be performed only from the CLI. Refer to Switch IPv4 based Cluster to IPv6 based Cluster for details.
- Both auto-discovery and manual discovery of the cluster devices are supported in IPv6 based clustering. You can configure the following addresses as IPv6 addresses for manual discovery:
  - Primary Address
  - Secondary Address
- Regular OOB clusters and Layer 3 OOB clusters can be created and managed using either all IPv4 addresses or all IPv6 addresses. You cannot have both IPv4 and IPv6 address within the same cluster. However, for Layer 3 OOB clusters, Virtual IP address can be configured as either IPv4 or IPv6 address, irrespective of the cluster control protocol.
- When adding nodes to existing cluster, you cannot add an IPv6 node to an existing IPv4 cluster and vice-versa. Also, when creating clusters, select the cluster protocol based on the IP address of the node that is managed in GigaVUE-FM. For example, if a node is managed with IPv4 address , then when creating a cluster , cluster protocol must be IPv4.
- The leader node should be configured with a higher preference than the standby nodes to ensure failover management.
- When creating a IPv6 cluster from GigaVUE-FM, it is recommended to use 0.0.0.0 as the VIP address.

Create IPv6 based Cluster

You can configure regular clusters and leaf-spine clusters by selecting the required protocol as either IPv4 or IPv6. Refer to the following topics for details:

- Regular Cluster Formation Workflow
- Leaf-Spine Cluster Formation Workflow

Switch IPv4 based Cluster to IPv6 based Cluster

GigaVUE-FM does not allow you to directly switch an IPv4 based cluster to IPv6 based cluster and vice-versa. You must use the GigaVUE-OS CLI to perform the switching operation. Refer to the following table:

S.No	From	Steps
1	GigaVUE-FM GUI	Disconnect the device from GigaVUE-FM. **1.** Go to **Inventory** > **Physical** > **Nodes**. **2.** Select the required cluster. **3.** From the **Actions** drop-down, select **Disconnect Node**. The device gets disconnected from GigaVUE-FM. Config-sync updates will not take place and API updates will be blocked. **Note**: Do not perform any configurations in the device after disconnecting the device from GigaVUE-FM.
2	CLI	Execute the following command from the CLI prompt of the leader node: **Seattle [1010: leader] config # cluster ip-ver-switch diagnostics** The diagnostics command establishes a server-client relationship between the leader and the member nodes of the cluster and checks if all the devices in the cluster are reachable and if there are any firewall or routing issues. Fix the issues and re-execute the command.
3		Execute the following command in the CLI prompt of the leader node to change the communication protocol version from IPv4 to IPv6 and vice-versa. **Seattle [1010: leader] config # cluster ip-ver-switch** **Note**: Before executing this command ensure that configurations are synced and all unsaved changes are saved. This command checks the following:

		• if all the devices have the proper IPv4 or IPv6 addresses to switch over. • If the devices are stable. If any of the devices in the cluster is in Safe mode, this command will be aborted. This command reloads the cluster. **NOTE:** In case of Layer3 OOB nodes, reconfigure the primary and secondary IP address of the leader node in all the member nodes of the cluster.
4	GigaVUE-FM GUI	1. Go to **Inventory** > **Physical** > **Nodes**. 2. Select the required cluster. 3. From the **Actions** drop-down, select **Edit Cluster**. 4. Click **Change Protocol / VIP** option. 5. Enter the following details: • From the **Protocol** drop down, Select IPV6 • VIP (mandatory) • Primary IP (mandatory for L3 OOB Clusters) • Secondary IP 6. Click **OK**. **Note**: Before connecting the Node in FM, if the VIP address is modified in device , you must add the new VIP address and credentials in **Node Details** page(Refer to Node Details section for adding VIP address and its credentials).
5	GigaVUE-FM GUI	Reconnect the device back to GigaVUE-FM. 1. Go to **Inventory** > **Physical** > **Nodes**. 2. Select the required cluster. 3. From the **Actions** drop-down, select **Connect Node**.

Refer to the following sections the GigaVUE-OS CLI Reference guide for configurations related to IPv6 based clustering:

S.No	Details in the GigaVUE-OS CLI Reference Guide	Reference
1	The cluster CLI command is updated for the following: • To configure protocol as ipv4 or ipv6 protocol • To verify if the devices in the cluster are reachable by the cluster leader • To switch from IPv4 to IPv6 or vice-versa	cluster

2	Example to remove a device from IPv4 cluster and add the same device to another IPv6 cluster.	How to Remove a Node from an IPv4 Cluster and Add it to an IPv6 Cluster and Vice-versa
3	For the Layer 3 Out of Band clustering, the primary IP address and secondary IP address of the cluster is cleared. After the protocol is switched, you must reconfigure the primary IP address and secondary IP address based on the switched IP.	Clustering a Node Using Layer 3 Out-of-Band Manual Discovery

Rules and Recommendations for Nodes and Clusters

Keep in mind the following rules and recommendations when managing nodes and clusters:

- If you add a node as a standalone node or as a cluster manually in GigaVUE-FM, the node will remain in that state unless you delete the standalone node/cluster. Refer to Using Command Line Interface for Managing Clusters for rules on adding and deleting the nodes using CLI.
- You can add a cluster to GigaVUE-FM using the cluster VIP or the IP address of the leader.
- The state of the cluster is decided based on the cluster API response from the device.
- Before joining an existing GigaVUE node to a cluster, it is recommended to use the **reset factory only-traffic** command to clean up existing traffic-related configuration. For example, in a cluster there is one leader and the other nodes are GigaVUE-TA nodes. When the leader is removed from the cluster, the GigaVUE-TA nodes moves to the unknown state. If another leader joins the cluster with a different database, the GigaVUE-TA nodes that are existing in the cluster will remain in the unknown state.
- Remove all physical loops before enabling the cluster. An accurate cluster topology will help with this. The GigaVUE-OS node automatically detects and prevents configurations that would cause loops, but it is best to avoid them in the first place.
- Star configurations offer the most efficient use of bandwidth. In general, use one GigaVUE-OS node at the hub of your star and then connect spokes off of that.
- GigaVUE-FM does not discover devices when both the IP interface address and GigaVUE-FM IP are on the same network.
- Create stack-links with enough capacity to match expected cross-node traffic. For example, you can use a 24x10Gb (PRT-HC0-X24) or 6x40Gb (PRT-HC0-Q06) for the GigaVUE-HC3 node.

- Configure only the stack type ports that you will use in the stack-link configuration. Loops can be created if stack type ports are configured but then not used in a stack-link.
- The first node added to the cluster becomes the leader. This is important when creating a new cluster using an existing, already-configured node and a new node. If you want to keep the configuration on your existing node, use it as the first node in the cluster. This way, the existing node becomes the leader and the new node inherits its configuration, preserving your existing settings. GigaVUE TA Series nodes are an exception since they cannot be the leader.
- When joining a new node to an existing cluster, give the new node a lower precedence than the leader. Once the database has synchronized to the existing leader, you can increase the precedence to make the newly joined node the leader, if that is required.
- You cannot have more than one leader in a cluster.
- For Inband Clustering, you must make physical connections for the cluster control traffic interfaces before issuing any cluster configuration commands. Because the first node added to a cluster becomes the leader, configuring cluster settings before physically connecting the cluster control network results in a situation with multiple leaders attempting to connect to the same cluster.

  NOTE: Merging clusters is not supported.
- For Layer 3 Out of Band clustering, the leader and the standby nodes must reside within the same subnet and must have the Gigamon Discovery Protocol (GDP) enabled for auto discovery.
- GigaVUE TA Series and Certified Traffic Aggregation White Box nodes in a cluster can have tool, network, hybrid, and stack ports.
- A GigaVUE TA Series node cannot be a leader. It can only join a cluster with other node types, such as GigaVUE-HC1, GigaVUE-HC3, GigaVUE-HC1-Plus, or GigaVUE-HCT.
- A GigaVUE TA Series node cannot be a standby node either. If the cluster has one leader and all other nodes are GigaVUE TA Series nodes, the cluster will not have a standby.
- Since a GigaVUE TA Series node can never be a leader or a standby in a cluster, a database restore is not possible. The best option is a text restore that has the information of the other nodes in the cluster removed from the text backup of the GigaVUE TA Series.
- If GigaVUE-HC1 and GigaVUE-HC3 devices are member nodes of a scaled cluster of approximately 10K map rules, then the following sequence of steps should be avoided:

  a. Removal of nodes from the cluster.
  b. Reload the removed node with 10K map rules.

If the above steps are done, the nodes may get locked out and you will not be able to log in to the devices. To remove a GigaVUE-HC1 or GigaVUE-HC3 node from a cluster and use it as a standalone node:

- Remove the node out of cluster.
- Erase the configuration using the "no traffic all" command on the standalone node.
- Reload the node.

RBAC and Tag Control on Nodes and Clusters:

- A user belonging to the Infrastructure Management resource category can create clusters in GigaVUE-FM using the nodes that he has access to, depending on the tag keys and the tag values assigned to him. The tag keys and the tag values of the user will be applied on the nodes and the cluster. Similarly, when the user removes a node from the cluster, the removed node gets the tags of the user.

- If a user has tag key and tag value set to ALL (as in the case of a super admin or admin user), then the tag key and the tag value applied to the nodes in the cluster depends on the tag value of the leader of the cluster. Similarly, if a node is removed from the cluster, then the tag value of the leader is applied to the removed node.

  Example:

  Consider the following:

  - **tag key:** Location
  - **tag values:** California, Texas, Washington and New York
  - User configured with tag key Location and has access to all tag values (ALL). Also, the user has access to the following devices.
  - dev1 - 10.115.46.11 - Tag-> Location: Texas
  - dev2 - 10.115.46.12 - Tag -> Location: California

  If the user creates a cluster C1 with dev1 and dev2, and dev1 as a seed node, as the user has access to all tag values of the tag Location, the tag value is derived from the seed device. In this example, the seed device is dev1 which has tag value Texas.

  The created cluster C1 will have the tag key: Location with value Texas.

- For nodes with single-valued tags, if you change the tag value of the node, a confirmation message pops-up. Based on the response, the tag value of the node is changed. Refer to the Create User-defined Tag section in the GigaVUE Administration Guide.

Using Command Line Interface for Managing Clusters

The following table provides information on the behavioral changes observed in GigaVUE-FM when the nodes are managed from CLI.

Process	Until Software Version 5.10.00	From Software Version 5.11.00
Add Nodes to Cluster	• **If the node is already managed by GigaVUE-FM:** The node will be added to the cluster and stack mode will be updated. • **If the node is not already managed by GigaVUE-FM:** The node will be added to GigaVUE-FM.	• **If the node is already managed by GigaVUE-FM:** Config sync will fail and is notified with appropriate messages. You must first delete the node from GigaVUE-FM and then add the node to the cluster. For example, consider the following nodes and clusters managed by GigaVUE-FM: • Stand alone node A • Cluster C with nodes B and C. To add the standalone node A to cluster C, you must first delete the node A from GigaVUE-FM and then add the node to cluster C. • **If the node is not already managed by GigaVUE-FM:** The node will be added to GigaVUE-FM.
Remove nodes from cluster	The node removed will be managed as a standalone node in GigaVUE-FM.	The node removed from the cluster will also be removed from GigaVUE-FM. If you want GigaVUE-FM to manage the node as a standalone device, you must add the node again. Refer to Add New Physical Node or Cluster to GigaVUE-FM section.
Edit Cluster ID	The new cluster id will get updated during the next config sync cycle.	The next config sync will fail because of the change in cluster ID. You must remove the cluster and add it again, so that the cluster is rediscovered with the new cluster ID.
Moving nodes between clusters	The node is removed from the existing cluster and added to the new cluster.	The node will be moved to the new cluster only if it is removed from the existing cluster. If the node is not removed from the existing cluster, config sync will fail until the node is removed from existing cluster.This could happen when the config sync for the cluster from which the

Process	Until Software Version 5.10.00	From Software Version 5.11.00
		device is moved out is triggered first. For example, consider the following nodes and clusters managed by GigaVUE-FM: **Cluster C1:** Nodes A, B, C **Cluster C2**: Nodes E, F, G If node E is moved from C2 to C1 through CLI, then following will be the behavior in the next config sync: If cluster C1 first completes the config sync: • Config sync will fail. This is because a device which is already in C2 is now being claimed in C1. • As soon as cluster C2 completes the config sync, node E will be removed from C2. • The next config sync will succeed for cluster C1. If cluster C2 first completes the config sync: • Node E will be removed from C2. • When config sync for C1 is completed, node E will be added to C1. If you keep moving the nodes from C2 to C1, then all but the last node will get moved. To move the last node you must first delete the cluster C2. In the subsequent config sync, node G will get added to cluster C1.

NOTE: If a node in a cluster does not report during config sync, GigaVUE-FM will remove the node, and an alarm with severity status 'Critical' is triggered (*Alarm description: Device is not reported*). If the node reports back to the same cluster, then the alarm will be cleared. However, if you acknowledge or delete the alarm and the node does not report back (if the node has been removed or added to another cluster) with in an hour after the alarm has been acknowledged, then GigaVUE-FM will clear the alarm.

GigaVUE TA Series and GigaVUE-HC3 Clustering Recommendations

The following recommendations are for GigaVUE-TA Series and GigaVUE-HC3 nodes in a cluster:

- When a GigaVUETA Series or GigaVUE-HC3 is connected to a node of a different type, ports may not become operationally up until the stack-links are created between the stack ports. To ensure the ports become operationally up:
  - Configure the specified ports as stack ports.
  - Configure the stack-link between the stack ports.

Cluster Rules

Clusters must adhere to the following rules:

Rule
All GigaVUE-OS nodes in a cluster must run the same version of the GigaVUE-OS software, including the major and minor version numbers.
Each GigaVUE-OS node in a given cluster must share the same **Cluster ID**, **Cluster Name**, and **Cluster Mgmt IP Address.** You can configure these settings in the **config jump-start** script, or, alternatively, use separate **cluster** commands to set them. When adding a node present on the same IP subnet to an existing cluster, so long as you specify the cluster ID correctly, the cluster Mgmt IP address (VIP) will be synchronized from the leader automatically.
Cluster management ports must be on the same IP subnet.
Each GigaVUE-OS node in a cluster must have its own unique box ID. The box ID is assigned to a chassis from the leader with the **chassis box-id <box ID> serial-num <serial number>** command (). Keep in mind that if you are using GigaSMART trailers to identify ingress ports, only box IDs from 1-64, inclusive, are supported.

Rule
You can only connect optical-to-optical stack-links. Stack-links must be at least 10Gb. In addition, they must use the same transceiver types, such as LR-to-LR, or SR-to-SR.
Use a stack-link between different types of GigaVUE-OS nodes so long as the medium, speed, and number of ports involved is the same on both sides.

Best Practices for OOB Clusters with IGMP Snooping

The following are best practices for out-of-band (OOB) clusters if Internet Group Management Protocol (IGMP) snooping is enabled in the cluster.

Clustering relies on the IGMP protocol to discover peer nodes and to communicate with them. Switches often have IGMP snooping enabled by default, which will filter IGMP packets from ports that do not have periodic IGMP membership reports. This can cause IGMP packet drops in out-of-band clusters.

Refer to About IGMP Snooping in a Cluster for more information. Also refer to the following best practices:

- allow Internet Group Management Protocol (IGMP) traffic by using an IP filter chain. Refer to Allow IGMP Traffic.
- enable an IGMP querier. Refer to Enable an IGMP Querier.

These best practices result in the following:

- hostnames being properly displayed in CLI commands that display cluster information such as **show cluster global brief**
- nodes joining clusters faster, especially nodes that are not capable of becoming a leader, such as GigaVUE TA Series nodes
- no multiple leaders being created in an out-of-band cluster. This can occur when a node that is capable of becoming a leader is not able to see the current leader and hence elects itself as a leader.

About IGMP Snooping in a Cluster

IGMP snooping is a networking feature that monitors IGMP membership reports received from different ports on a networking switch and learns the ports to which multicast groups belong. When a port stops sending membership reports about a multicast group, the switch will stop forwarding the group's traffic to the port.

An IGMP querier is a router (or switch) feature that periodically queries the network for multicast group interests. If a node on the network belongs to a certain multicast group, it responds to the queries, the router then records or refreshes its record of the node's interest in the traffic for the group, and the router forwards traffic to the network towards the node. The switches on the network with IGMP snooping enabled also learn from the responses and maintain their records about the nodes' interests in groups and forward traffic accordingly.

Hostnames are detected using Multicast Domain Name System (mDNS) packets, which are in multicast group 224.0.0.251.

An IP filter is a chain of rules for the treatment of packets. Refer to the *"Using IP Filter Chains for Security"* section in the *GigaVUE-OS CLI Reference Guide*.

Allow IGMP Traffic

If IP filtering is enabled (and IGMP snooping is enabled):

- Verify that IGMP traffic is allowed.
- For example, issue the following CLI commands:

  (config) # ip filter chain INPUT rule append tail target ACCEPT dup-delete protocol igmp

  (config) # ipv6 filter chain INPUT rule append tail target ACCEPT dup-delete protocol igmpv6
- Verify that mDNS traffic is allowed.

  If IGMP snooping is disabled, you do not need to allow IGMP traffic. However, you must allow UDP multicast traffic that targets 224.0.0.251. For example, issue the following CLI command:

  (config) # ip filter chain INPUT rule append tail target ACCEPT dup-delete dest-addr 224.0.0.251 /32

  where:

  dest-addr specifies the multicast group

Enable an IGMP Querier

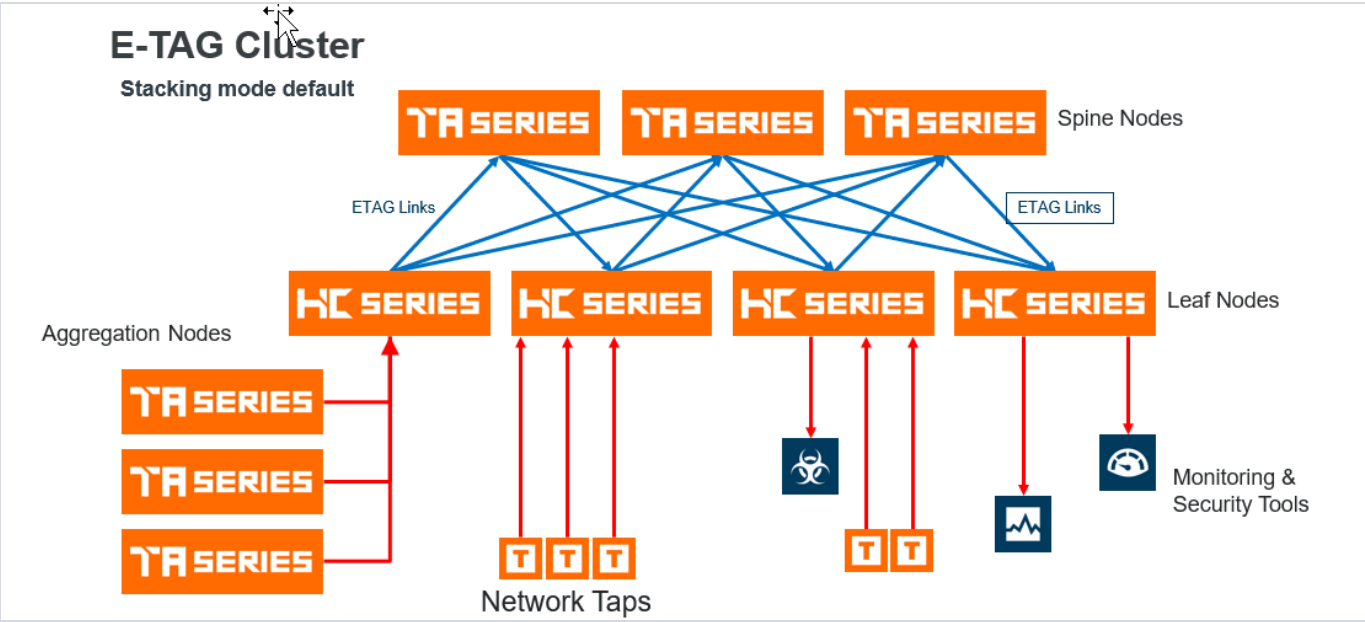
If IGMP snooping is enabled:

- Check if there is an IGMP querier on the cluster network. The querier periodically sends queries that trigger the nodes in the cluster to send IGMP membership reports. For example, use a sniffer tool to verify if there is an IGMP querier on the network, such as Wireshark.
- IGMP snooping and IGMP snooping querier settings vary by networking switch. Refer to the respective documentation for how to configure them on your device.

When IGMP traffic is allowed and an IGMP querier is enabled in the network, the switches in the network will be refreshed through the IGMP membership reports.

E-Tag Clustering

The 5.16 release introduces a new E-Tag based stacking mode to cluster with GigaVUE-TA400 or GigaVUE-HCT. All legacy cluster topologies (leaf-spine, tree, etc.) with OOB cluster management are supported with E-tag stacking mode. Refer to the example below.



The new E-TAG stacking mode is supported on all the platforms listed below.

- GigaVUE-HC3 CCv1 & CCv2
- GigaVUE-HC1
- GigaVUE-HC1-Plus
- GigaVUE-HCT
- GigaVUE-TA25
- GigaVUE-TA25E
- GigaVUE-TA100
- GigaVUE-TA200
- GigaVUE-TA200E
- GigaVUE-TA400
- GigaVUE-TA400E
- DELL S4112F-ON

The BiDi optics on stack ports are also supported in E-Tag stacking mode.

NOTE: The GigaVUE-TA400 and GigaVUE-HCT platforms support only Default (E-Tag) stacking mode and do not support Legacy stacking mode. All other platforms support both Legacy and Default (E-Tag) stacking modes.

Stacking Mode

The E-Tag Clustering introduces a user selectable stacking mode. The stacking modes are as follows:

- ***Legacy***: This option selects Legacy stacking mode.
- ***Default***: This option selects E-Tag stacking mode.

To configure the default stacking mode or the legacy stacking mode through GigaVUE-OS - CLI use the following commands :

(config) # system stacking-mode legacy: Selects the Legacy Stacking mode

(config) # no system stacking-mode legacy: Selects the Default Stacking mode.

Legacy Stacking Mode	E-Tag (Default) Stacking Mode
(config) # system stacking-mode legacy **(config) # show system stacking-mode** **Stacking Mode configuration:** **Mode : legacy-mode**	**(config) # no system stacking-mode legacy** **(config) # show system stacking-mode** **Stacking Mode configuration:** **Mode : default-mode**

NOTE: After user confirmation the system stacking-mode legacy command immediately resets the traffic configuration and initiates a cluster reload of all nodes. After the cluster is up, the configuration saved in the backup file must be applied manually to restore the traffic configuration.

For example :

```
[cluster: leader] (config) # no system stacking-mode legacy
! WARNING: Changing stacking mode will automatically
! - Take backup of config in file stacking_mode_config_backup.txt
! - Reset factory only traffic config
! - RELOAD the cluster
! - User must manually apply stacking_mode_config_backup.txt after bootup
Confirm stacking mode change? [no] yes
Configuration saved to database 'initial'
System shutdown initiated -- logging off.
```

After the cluster is up and user is logged back in, apply the saved configuration in the backup file:

[cluster: leader] (config) # configuration text file stacking_mode_config_backup.txt apply fail-continue

Refer to the system section in GigaVUE-OS CLI Reference Guide to know more.

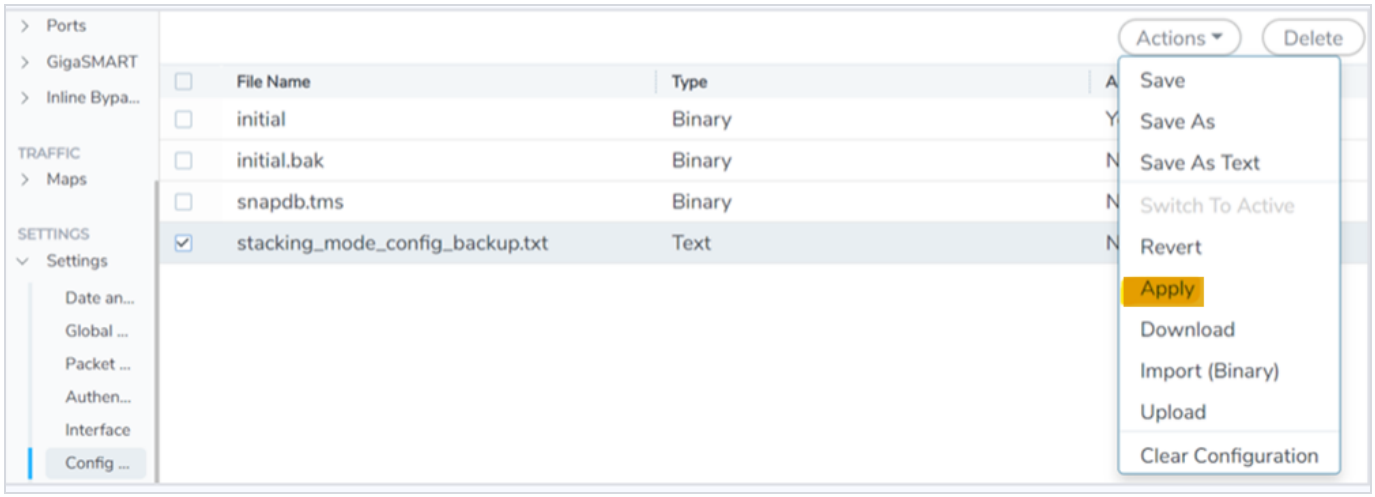
To configure from GigaVUE-FM use the stacking mode drop-down options. Refer to the Regular Cluster Formation Workflow,Edit Cluster Parameters,Leaf-Spine Cluster Formation Workflow.

Switching Stacking Mode

The Stacking mode can be switched from Legacy to Default and vice versa. Switching the stacking mode will reload the cluster and maintain a backup file of the last applied configuration. A Confirmation window appears advising that a backup file is generated with the updated traffic configuration. The configuration saved in the backup file must be applied manually to restore the traffic configuration after the cluster is up. The command used to apply the backup configuration is as follows:

[cluster: leader] (config) # configuration text file stacking_mode_config_backup.txt apply fail-continue

Refer below, to apply the backup configuration file from GigaVUE-FM.



Switching the stacking mode will not be allowed if the cluster includes an unsupported node.

Refer to Regular Cluster Formation Workflow,Edit Cluster Parameters,Leaf-Spine Cluster Formation Workflowto know more.

Notes

1. The L3 and L4 fields hashing of double tag Q-in-Q traffic is not supported over stack links in E-Tag mode on non GigaVUE-TA400 platforms. This limitation is not applicable for GigaVUE-TA400.
2. Common hash settings for all stack and non-stack GigaStream in E-Tag mode. Fabric advanced-hash command is not applicable to stack GigaStream in default stacking mode but applicable to GigaSMART groups on GigaVUE® HC Series.
3. In an E-tag cluster having MPLS Header Stripping enabled on GigaVUE-TA25, GigaVUE-TA25E and GigaVUE-HC1-Plus source/network ports, the customer outer vlan tag gets stripped along with MPLS labels specifically for tagged MPLS customer traffic.

New 6.0

1. Circuit ports as destination in GigaSMART first level maps (L2 circuit encapsulation tunnels with first level map) are not supported in E-Tag mode in 5.16 release. This is supported from release version 6.0.
2. The L2GRE and VXLAN tunnel encapsulation and decapsulation are not supported in E-Tag mode in 5.16 release. This is supported from release version 6.0.
3. The MPLS Header Stripping is not supported on GigaVUE-TA25 in the E-Tag mode in 5.16 release.It is supported on GigaVUE-TA25 from release version 6.0. On GigaVUE-TA25, the outer customer VLAN tag will also be stripped along with MPLS header for the VLAN tagged MPLS packets.
4. MPLS Header Stripping in E-Tag mode in 5.16 release on GigaVUE-HC3, GigaVUE-TA200, GigaVUE-TA100 has one limitation of 4 byte VLAN tag added when header strip enabled ports are mapped to tool ports within the same box/node specifically. This limitation is not applicable from 6.0 release version.
5. VXLAN Header Stripping in E-Tag mode in 5.16 release on platforms other than GigaVUE-TA400, GigaVUE-TA25 has limitation of 4 byte VLAN tag added when header strip enabled ports are mapped to tool ports within the same box/node specifically. This limitation is not applicable from 6.0 release version.

Cluster Safe and Limited Modes

Safe and limited modes in cluster safeguard critical provisioning errors for both standalone nodes and nodes in a cluster.

During provisioning operations such as configuring a map, in rare scenarios there can be unrecoverable system errors that can potentially put the cluster, clustered nodes, or standalone nodes into unsafe or unstable states. Once in such a state, additional operations

or configuration changes can cause the node to crash, the cluster to deform, or data traffic to be impacted. For example, due to a node attempting to rejoin a cluster, a chassis can end up in a reboot loop. In previous software versions, there was no way to prevent entering the loop.

These modes provide notification, stop further operations from being performed, and give you time to troubleshoot and plan the recovery of the cluster, the clustered node, or the standalone node.

Two modes are supported. The first is called safe mode and is triggered when the node detects unrecoverable errors, but the existing flow maps are not impacted. The second is called limited mode and is triggered when the node detects continuous system reboots. In this mode, the node will become standalone and only basic configuration will be allowed.

Safe Mode

A node enters safe mode when there are unrecoverable errors. Any node in a cluster can enter this mode. The purpose of this mode is to detect system configuration failures early and avoid future failures, such as system crashes.

Examples of unrecoverable errors are when there are inconsistencies between the system and the running configuration or when the cluster configuration did not merge properly with the existing configuration.

As part of merge error recovery, nodes joining a cluster are automatically restarted so the merge error can be fixed. If the restart cannot correct the merge error, the node will enter safe mode.

Another example is that a GigaVUE TA Series node could enter safe mode when unlicensed cluster ports are used in an offline configured map. (It is recommended to use only licensed ports in map configurations.)

A node will automatically enter safe mode.

To recover from safe mode, reload the node. If safe mode persists, contact Gigamon Technical Support.

Limited Mode

A node automatically enters limited mode when it detects repeated system crashes.

Limited mode is triggered when there are three (3) failures/system crashes within 15 minutes. In limited mode, the cluster configuration is ignored. No cluster configuration or GigaVUE-OS configuration is accepted when the node is in limited mode.

When limited mode has been detected, collect information and report it to Gigamon Technical Support.

Support for Cluster Types

The GigaVUE-FM workflow supports only out-of-band clusters; not inband clusters. To create and manage an inband cluster, refer to the *GigaVUE-OS CLI Reference Guide*.

Create Clusters

GigaVUE-FM supports workflow-based configurations for forming clusters. Starting from software version 6.2.00, you can create only IPv6 clusters from GigaVUE-FM.

- Refer to Regular Cluster Formation Workflow for instructions on how to use the regular cluster formation workflow.
- Refer to Leaf-Spine Cluster Formation Workflow for how to use the leaf-spine cluster workflow

Regular Cluster Formation Workflow

Gigamon's Cluster formation can be done for maximum of 64 devices.

GigaVUE-FM supports workflow-based configurations for forming a cluster. This workflow walks through the required steps to form a complete cluster for a regular cluster.

NOTE: Refer to Leaf-Spine Cluster Formation Workflow for how to use the Leaf-Spine Cluster workflow

Deployment Checklist

Before forming a Cluster, it is strongly recommended that you familiarize yourself with the relevant documentation and review the following deployment checklist:

- Gigamon Fabric Management must be upgraded to GigaVUE-FM 5.3.00 or later.

- Gigamon device must be upgraded to GigaVUE-OS 5.2.00 or later
- Advanced Features License must be installed on GigaVUE-TA devices.
- Physical connection must be established to create stack links.
- Devices must be physically connected to create links among devices from GigaVUE-FM.

Create Regular Cluster Formation

To create a cluster:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**..
2. Click **Create Cluster** and select Normal Cluster.

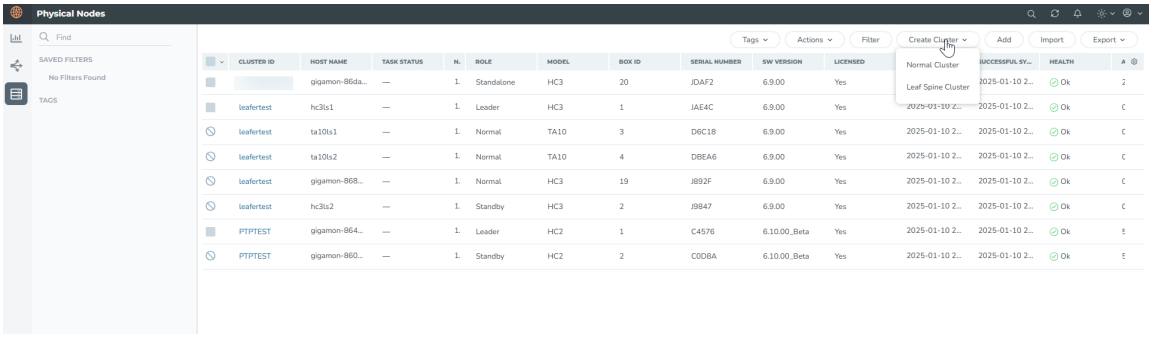
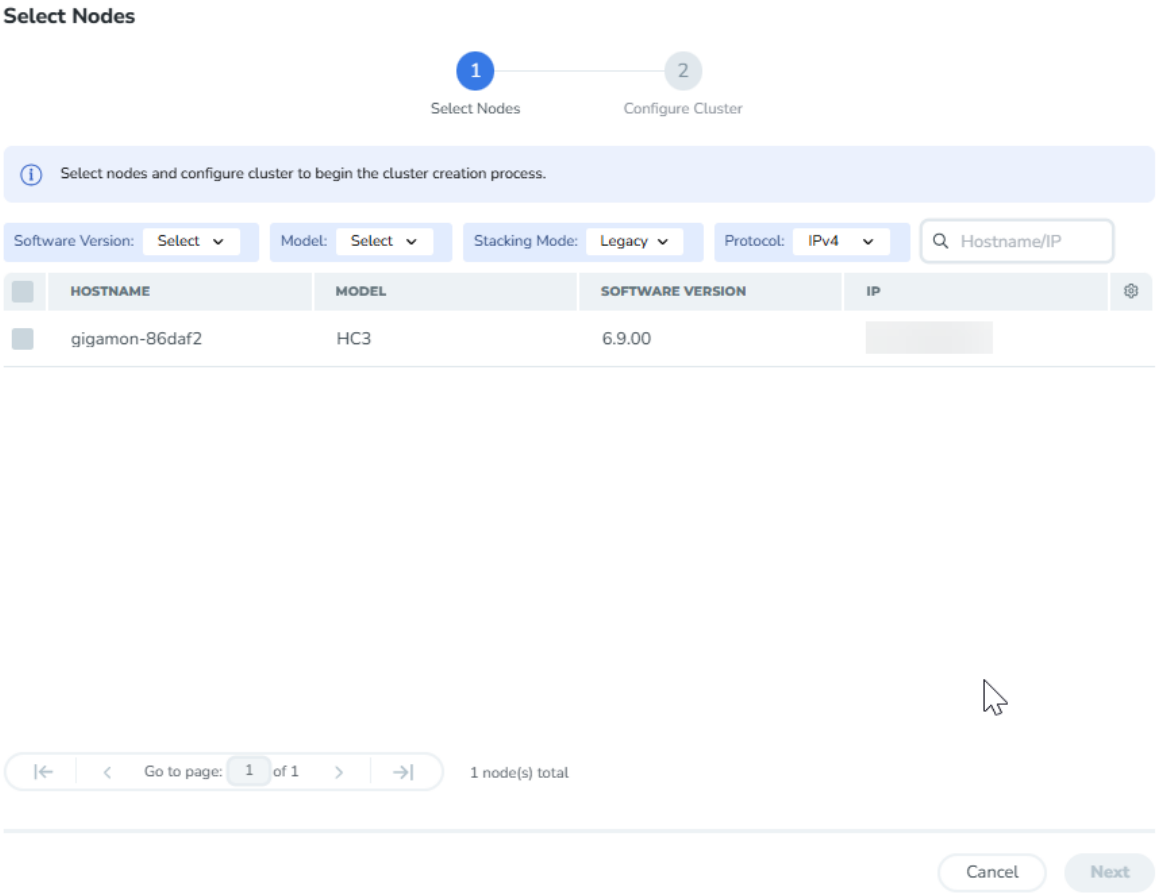


Figure 6 *Create Cluster*

3. Filter the nodes based on the Software Version, Model, Stacking Mode, Hostname or IP address, and the Protocol through which the nodes communicate with each other.

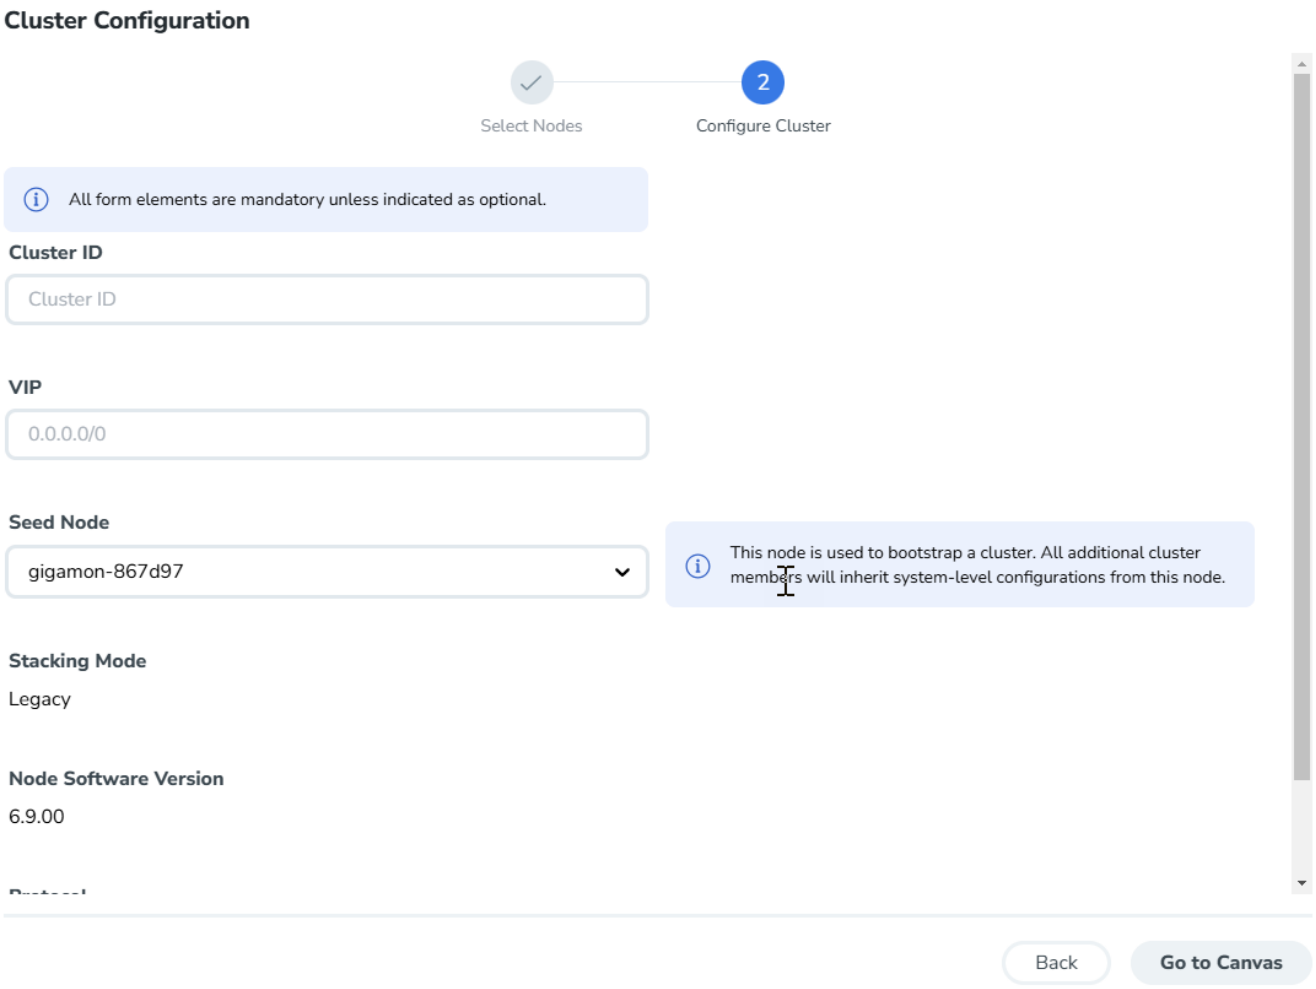
Notes:

- GigaVUE-TA and GigaVUE-HCT devices are not supported in Legacy stacking mode.



4. Select the nodes to include in this cluster and click **Next**. (Click a device to select it; click it again to deselect it).
5. In the Cluster Configuration window, enter a valid **Cluster ID** and Virtual IP (**VIP**) and select the leader in the **Seed Node** list.

   NOTE: The leader cluster preferences in GigaVUE-FM determine which of the nodes will be the default seed node. GigaVUE-TA devices cannot be a leader.

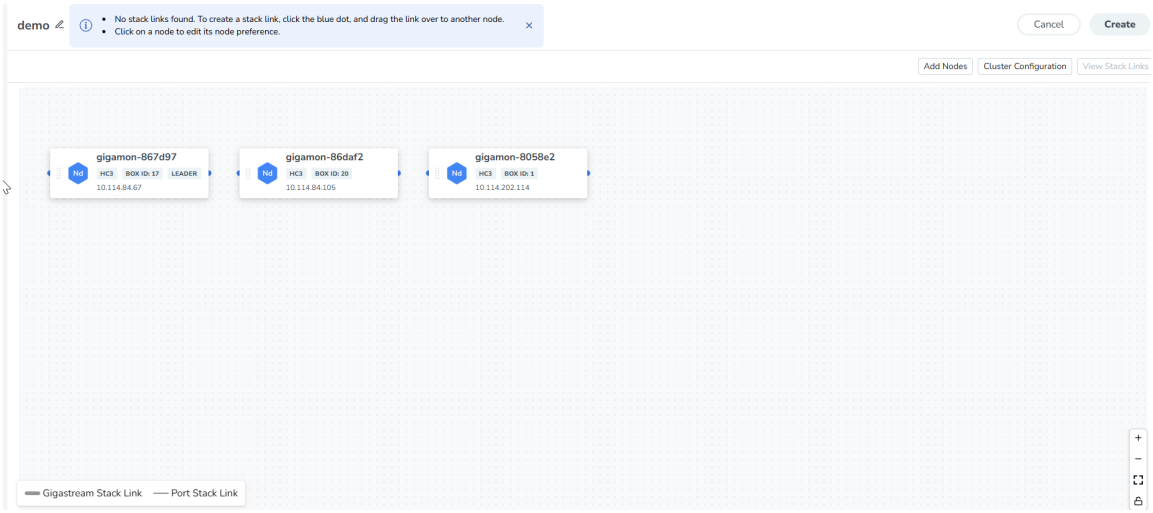


6. After completing the Cluster Configuration details, click **Go to Canvas**.

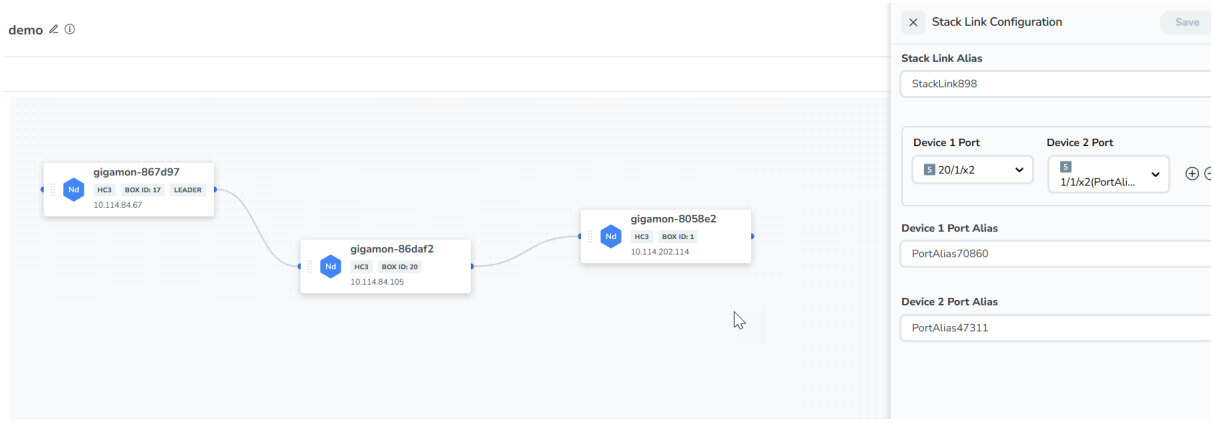
   NOTE: Use the **Back** button to return to the Select Devices page and revise the device selection for this cluster.

Customize Stack Links

Finally, customize the stack links to define how the nodes should be connected.



7. Connect any two devices to create a stack link between those two devices. Click the tip of the node and drag your cursor to the second node tip to create a link. After you create the link, a dotted line illustrates the connection.



8. Configure the formed links in the **Stack Links** Configuration panel:
   a. Select ports in each device that are compatible, for example: x-x ports, x-q ports, q-c ports, and x-c ports.
   b. Select one or more ports from each device to create a GigaStream stack link.
   c. Click **Save**.

   GigaVUE-FM generates the aliases for each Port stack link and GigaStream stack link. You can edit these aliases.

NOTE: While creating cluster, click **View Stack Links** on the right corner of the canvas to view the details of the stack links of the cluster.

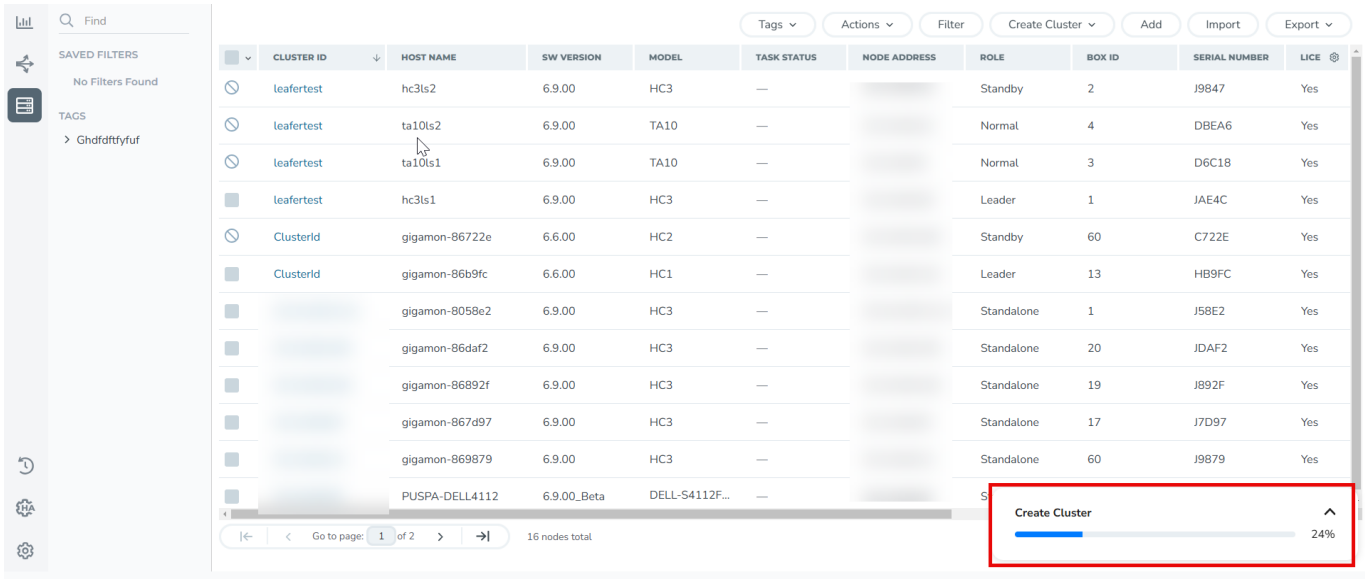
9. Click **Create** to start the cluster creation process.

The Create Cluster progress window in the lower right corner of the page shows the status of every node as it joins the cluster. It takes a few minutes for the cluster to form. The cluster creation process involves the following steps:

- Cluster[clusterName] Creation Successful followed by Seed device
- Verifying Nodes[Will display HostName of all devices]
- Adding Node[HostName] to cluster [clusterName]
- Node[HostName] successfully joined to the cluster.
- Configuring cards for cluster[clusterName]
- Rediscovering cluster[clusterName]
- Configuring ports for the cluster[clusterName].
- Configuring ports will display the status of each stack link and GigaStream whether the creation is successful or not.

When the cluster formation process is complete, you can view the completion message in the progress window.

NOTE: Refer to Check Cluster Status for Events.



10. Click **Go to Cluster** in the progress window to view the cluster overview.

Edit a Cluster

The Edit cluster option supports only the following operations to the existing cluster:

- Multiple devices can be added to the existing cluster in a single update operation.
- Stack links can be created only from the new device which is added into the cluster wizard.
- Leader preferences can be changed for each device through edit cluster option.
- Stack link alias and GigaStream alias can be edited for newly created links.
- Stacking mode can be changed from Legacy to Default and vice versa.

Notes:

- You cannot change the stacking mode and also add or remove a device simultaneously. Only one task can be performed at a time.

NOTE:

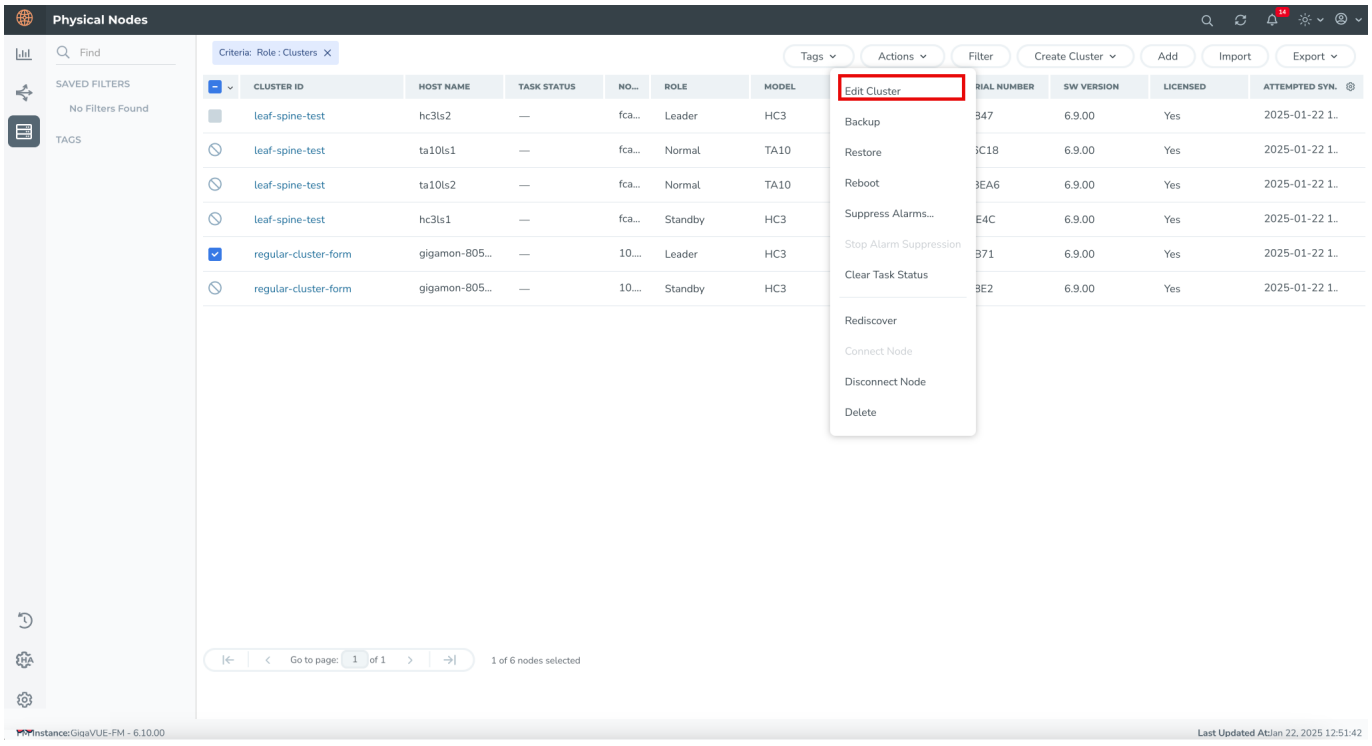
- No option to remove the existing stack links through cluster canvas.
- No option to create links in existing devices.
- Addition and deletion of devices in a single update operation should not be appreciated.
- No option to edit the existing stack link alias and GigaStream alias.

Prerequisites

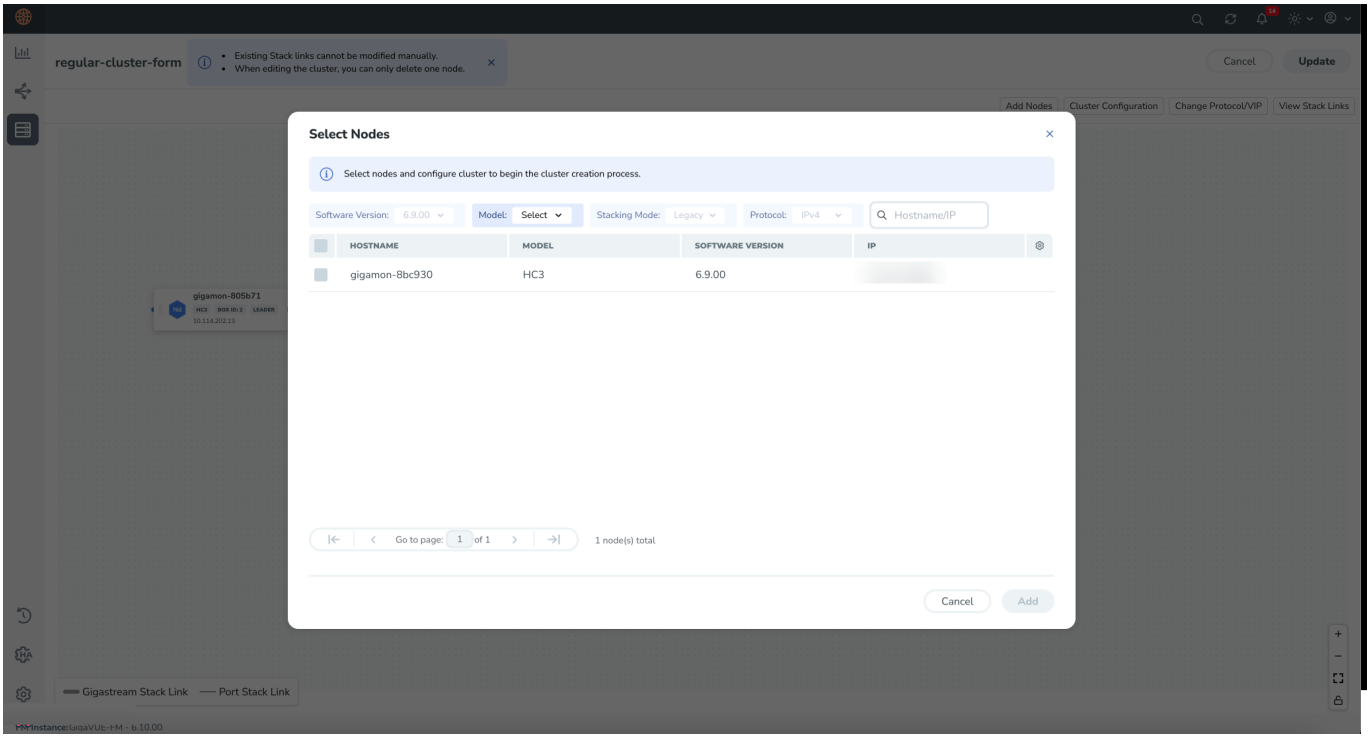
Standalone devices that have maps cannot be added to cluster if ports used in maps overwrites with the selected ports in stack link table.

This workflow describes how to add a node to a existing cluster.

1. Select a cluster and choose **Actions > Edit cluster**.



2. The Edit Canvas displays the existing stack link configuration details in the cluster canvas. Click **Add Nodes** to view the standalone devices.
3. Select the required devices from the table in the **Select Nodes** window and click **Add**.



4. In the canvas, draw the links between the newly added device. (**NOTE**: no new link is created for the existing device.)
5. Configure the stack link details in the stack link quick view.
6. Click **Save**.
7. Change the stacking mode as required.
8. Click **Update** to initiate the update process. A confirmation window appears advising that a backup file is generated with the updated traffic configurations. The configurations saved in the backup file must be applied manually to restore the traffic configurations.
9. Click **OK** to run the cluster update.

   When the cluster update operation starts, a notification window appears at the right corner of the GigaVUE-FM window to show the status progression of each node, card, GigaStream and stack link.

   When the cluster update operation is complete, a notification window confirms the completion of the cluster updates.
10. Click **Go to Cluster** to go to view the cluster overview.

    The created GigaStreams appears in the device Port Groups page, and the created stack links appear in the device Stack Links page.

Delete a Node from a Cluster

This workflow describes how to remove a device node from a cluster.

NOTE: Only one device can be removed from the cluster per update operation.

The device should not contain any map configurations in a cluster. Those devices cannot be removed until the maps are present.

1. Select a cluster and choose **Actions > Edit cluster**. Only one device can be deleted from the canvas.
2. To remove a device, click the device to be removed from the canvas and click **Remove Node**.

   The removed device will be deleted from canvas.
3. Click **Update** to initiate the cluster-update operation.

   The Manage Cluster notification window shows the progress of nodes being removed from the cluster. When the device is successfully removed from the cluster, a notification window confirms the successful deletion of the cluster.
4. Click **Go To Cluster** to go into device overview page and see the cluster details.

How to Change the Leader Preference of a Device

This workflow describes how to change the device's leader preference.

1. Select a cluster and choose **Edit cluster** under Actions.
2. To set the leader preference for a device, click the device and update the **Leader Preference** details in the **Node Configuration** quick view that appears in the side pane.
3. Update the leader preference text box and click **Update** to proceed.

   The Update Cluster notification appears to show the progress of the cluster update. When the process is complete, **Cluster updation completed**, message appears in the notification window.
4. Click **Go To Cluster** to go to the device overview page and see the cluster details.

Edit Cluster

The **Edit Cluster** action supports the following types of changes:

- Leader preferences can be changed for each device.
- Multiple devices can be added to the existing cluster in a single update operation.

- Stack links can be created only from a new device that is being added.
- Stack link alias and GigaStream alias can be edited for newly created links.
- Nodes can be removed from an existing cluster, one at a time.
- To enable the Cluster VIP change, the credentials for the new VIP must be added on the Node Details page in GigaVUE-FM, for a seamless connection of GigaVUE-FM with the cluster.
- You can edit the tag id and tag values associated to nodes and clusters:
  - To edit a configured tag ID and tag value of a node/cluster, you must be an admin user with read and write access.
  - If you replace a single-valued tag value associated to a node, a confirmation message pops-up before replacing the tag value.
- Stacking mode can be changed from Legacy to Default and vice versa.

Notes:

- You cannot change the stacking mode and also add or remove a device simultaneously. Only one task can be performed at a time.

The following options are not supported by the Edit Cluster action:

- There is no option to remove existing stack links through the cluster wizard.
- There is no option to create links in existing devices.
- Addition and deletion of devices in a single update operation is not supported and is not recommended. If you attempt to add and delete devices in a single update operation, you may get unexpected results.
- There is no option to edit the existing stack link alias and GigaStream alias.
- There is no option to rename a cluster directly from GigaVUE-FM. If you rename a cluster using the GigaVUE-OS CLI, the rename does not reflect in the GigaVUE-FM. To rename a cluster that is managed from GigaVUE-FM:
  1. Remove the cluster from GigaVUE-FM.
  2. Take a backup/log session. Run the **`show run`** and **`show diag detail`** commands from the leader node.
  3. Disable cluster on all nodes using the command **`no cluster enable`**. Start with the member nodes and disable the leader node at the end.
  4. From leader node, in configure mode, run the command **`cluster id <new-cluster-id>`** to change the cluster name.
  5. Enable the cluster on all nodes. Start with the leader by running the command **`cluster enable`**. Leader node will push the new cluster name to all the nodes.
  6. Once the cluster is up and running, add the cluster back to GigaVUE-FM.

Prerequisites:

You must clear all configurations on a node before adding it to a cluster.

Edit options:

- Add Nodes to a Cluster
- Remove Nodes from a Cluster
- Edit Cluster Parameters

Inband Cluster Management

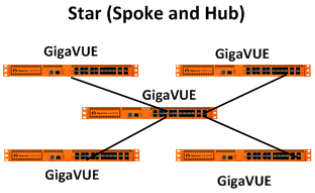
Inband Cluster Management simplifies traditional network management and maintenance by creating a virtual device to manage multiple physical nodes. This simplified approach makes it possible to oversee large networks by defining policies that span across multiple devices. The Inband Cluster Management feature is designed to reduce operational cost and extend coverage by eliminating a dedicated management network.

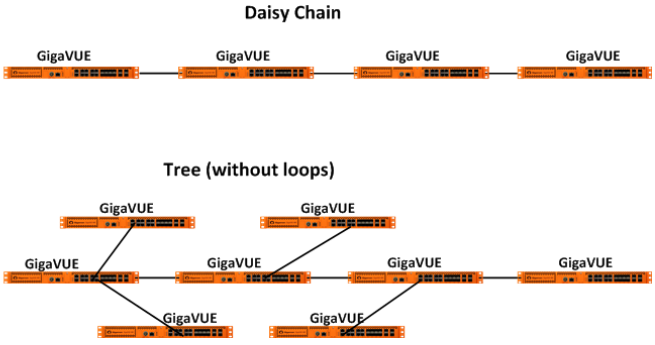
Inband Cluster Management is supported on all GigaVUE-OS nodes except GigaVUE-TA400, GigaVUE-TA400E, and GigaVUE-HCT.

Inband Cluster Management Topologies

The benefits of Inband Cluster Management are to eliminate the Layer-2 Ethernet network and create a virtual management network through the data path where the data traffic is flowing.

Inband Cluster Management supports multiple topologies that include:





NOTE: Subsets or aggregations of these topologies may be created; however, it is important not to create a loop within these specified topologies.

Loops are typically created in the following scenarios:

- **Two Node Loops** occur between two nodes in a cluster forming two or more stack links and the stack links are not contained in one GigaStream.
- **Multi-Node Loops** occur when multiple nodes form a cluster whereby a link connects between node A and node B, another link connects between node B and node C, and yet another link connects between node C and node A.

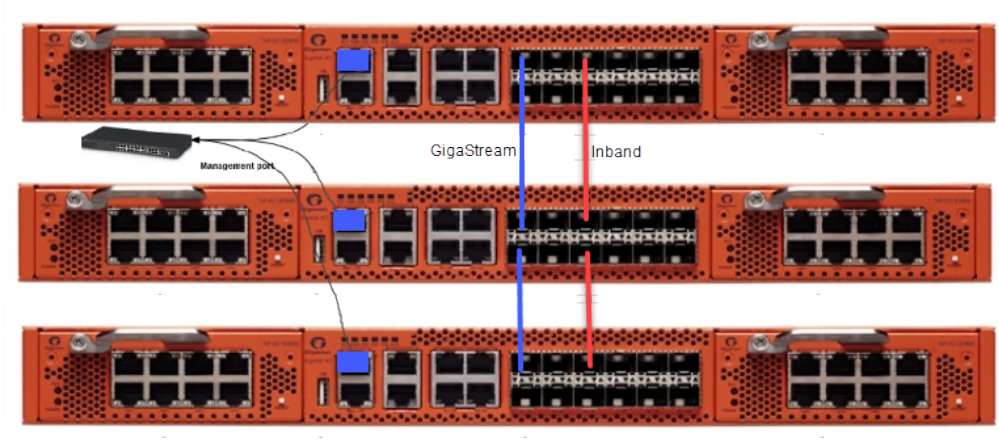
Inband Cluster Management Stack Ports

A common Inband Cluster Management topology is configured between the Layer 2 device's Ethernet management port to a GigaVUE-OS node using a stack port configuration.

Two or more GigaVUE-OS nodes may be directly connected in a one-to-many relationship between physical connections. GigaVUE-OS nodes may also be indirectly connected if there is a path of stack ports between the nodes.

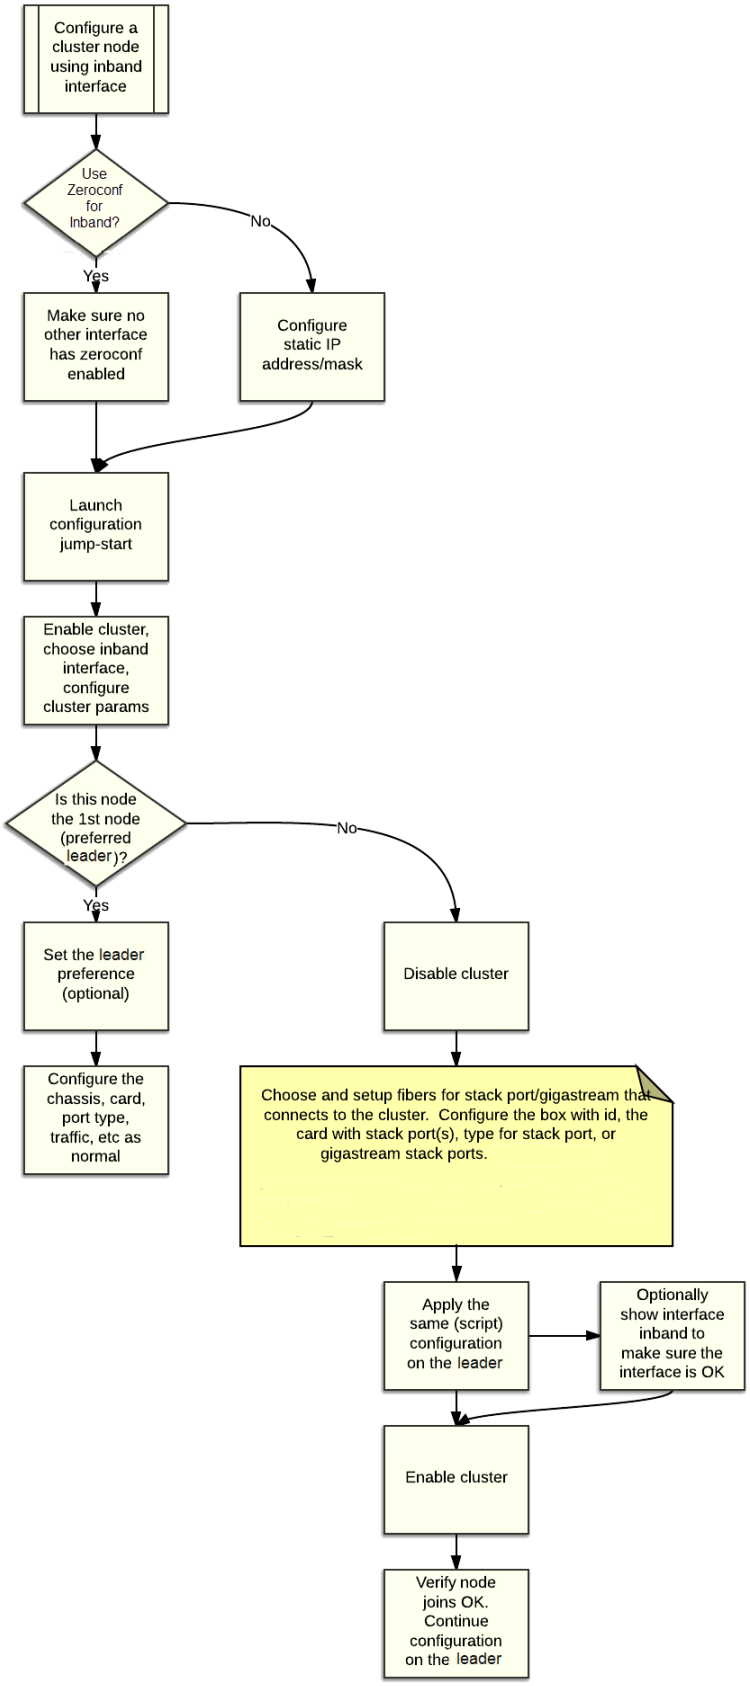
Inband Cluster Management Stack Ports Example

Inband Cluster Management Stack PortsThe following figure visually depicts how Inband Cluster Management uses the grouping of stack ports to connect between GigaVUE-HC1 nodes.



NOTE: Ensure that there is a physical connection between the stack ports of the two nodes that are being added to the Inband cluster.

Inband Cluster Management Configuration Flow Chart


Configure a cluster node using inband interface
Use Zeroconf for Inband?
No
Yes
Make sure no other interface has zeroconf enabled
Configure static IP address/mask
Launch configuration jump-start
Enable cluster, choose inband interface, configure cluster params
Is this node the 1st node (preferred leader)?
No
Yes
Set the leader preference (optional)
Disable cluster
Configure the chassis, card, port type, traffic, etc as normal
Choose and setup fibers for stack port/gigastream that connects to the cluster. Configure the box with id, the card with stack port(s), type for stack port, or gigastream stack ports.
Apply the same (script) configuration on the leader
Optionally show interface inband to make sure the interface is OK
Enable cluster
Verify node joins OK. Continue configuration on the leader

Inband Cluster Management Configuration

An interface called "Inband" has been created upon boot-up to ensure backward compatibility with an existing clustered infrastructure. This interface has similar properties and characteristics of a typical Ethernet interface such as eth0.

NOTE: Ensure that there is a physical connection between the stack ports of the two nodes that are being added to the Inband cluster.

Enable Cluster Management for GigaVUE TA Series Nodes

To enable clustering, GigaVUE TA Series nodes require an Advanced Features License. This license can be obtained by contacting Gigamon Sales team. In order to obtain the license for a Gigamon node, have the node serial number available. All licenses are tied to the serial number and cannot be moved.

For licensing the GigaVUE-OS on a white box, users can access the GigaVUE-OS licensing portal and obtain the license key online. In order to generate the license, the following are required: the serial number of the white box, digital footprint, and Gigamon Installation Key (GIK).

Add Nodes to a Cluster

You can manage an existing cluster through GigaVUE-FM by adding nodes to it. The nodes must be standalone nodes that are currently managed by GigaVUE-FM.

When a node joins an existing cluster, all of its existing traffic configuration, including maps, will be replaced by the configuration of the leader.

- If you add a node to a cluster using CLI commands, then the node will be added to GigaVUE-FM only if is not already managed by GigaVUE-FM. If the node is already managed by GigaVUE-FM, then config sync operation will fail. Refer to the Rules and Recommendations for Nodes and Clusters section for more details.

- When a new node is added to an existing cluster, if the leader is in secure cryptography mode, then the node joining the cluster will also be changed to secure cryptography mode.

The following is an example of adding nodes to an existing cluster using GigaVUE-FM.

1. From the left navigation pane, go to **Inventory > Physical > Nodes**.
2. Select a cluster and choose **Actions > Edit Cluster**. The Edit Cluster - Canvas appears showing the existing stack link configuration details in the cluster wizard. Standalone devices appear in the Devices pane.
3. Drag the required devices from the Devices pane into the Edit Cluster canvas under the leaf or spine container.
4. Connect the newly added devices to other devices to create stack links.

   Click the tip of the node and drag your cursor to the second node tip to create a link. After you create the link, a dotted line will illustrate the connection

   NOTE: No new link is created for existing devices; they need to be added manually.
5. Configure the stack link details in the stack link table and click **Save**.
6. Click **Update** to update the configuration.

   A Confirmation window appears advising that all traffic configurations will be erased on newly added or removed nodes.
7. Click **OK** to continue.
8. The Manage Cluster update notification window appears showing the status of each update activity on the nodes, cards, GigaStreams and stack links.
9. After the cluster update operation completes, a “Manage Cluster Completed” message appears.
10. Click **Go to Cluster** to view the updated cluster overview

Remove Nodes from a Cluster

You can manage an existing cluster through GigaVUE-FM by removing nodes from it. After a node is deleted from a cluster, it will become a standalone node. FM will continue to manage it.

For nodes leaving a cluster, the username and password of the admin account on the cluster will be used for managing the node after it has been removed from the cluster.

NOTE: If you remove a node from a cluster using CLI commands, then the node is removed from GigaVUE-FM. If you want the node to be managed by the same GigaVUE-FM instance, you must add the node as a new device by providing the credentials.

To remove nodes of an existing cluster using GigaVUE-FM:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**..
2. Select a cluster and choose **Actions > Edit Cluster**.

   The Edit Cluster - Canvas appears showing the existing stack link configuration details in the cluster wizard. Standalone devices appear in the Devices pane.

NOTE: Only one device can be deleted from canvas per operation. It can be any device from the cluster.

3. Right-click the device to be removed from the canvas and click **Remove**.

   Right click to remove
4. The removed device is deleted from the canvas.
5. Click **Update** to update the configuration.

   A Confirmation window appears advising that all traffic configurations will be erased on newly added or removed nodes.
6. Click **Ok** to continue.
7. The Manage Cluster update notification window appears showing the status of each update activity on the nodes, cards, GigaStreams and stack links.
8. After the cluster update operation completes, a “Manage Cluster Completed” message appears.
9. Click **Go to Cluster** to view the updated cluster overview.

Edit Cluster Parameters

When editing a cluster node, you can only edit the cluster leader preference. You can only change the cluster leader preference on one node at a time.

For the leader preference, higher values are more likely to claim the Leader and Standby roles; lower values are less likely.

To edit leader preferences:

1. To set the leader preference for devices, right-click the required device and click the **Edit Details** options button.

2. The Device configuration quick view should appear on the right. Edit the Leader Preference in the text box.
3. Click **OK**.

NOTE: Most of the fields are read-only. You can change the cluster leader preference, as desired. Use preference settings from 10 to 100 for leader, standby, and member roles. Use preference settings from 1 to 9 for normal nodes that are excluded from taking the leader or standby role. GigaVUE TA Series nodes always have a preference of 1.

4. After saving your changes to the nodes, click **Update** to apply the changes to the cluster.

   A Confirmation window appears advising that a backup file is generated with the updated traffic configurations. The configurations saved in the backup file must be applied manually to restore the traffic configurations after the cluster is up.
5. Click **Ok** to continue.
6. The Manage Cluster update notification window appears showing the status of each update activity on the nodes, cards, GigaStreams and stack links.
7. After the cluster update operation completes, a “Manage Cluster Completed” message appears.
8. Click **Go to Cluster** to view the updated cluster overview

Check Cluster Status

When a cluster is being created, you can check the status through cluster management events or audit log entries. Refer to the following sections:

- Cluster Management Events
- Audit Logs

Cluster Management Events

On the **Events** page, the following event types indicate the progress of the cluster as it is being formed:

- ClusterCreationStarted
- ClusterCreationCompleted
- ClusterCreationFailed

The following events indicate the status of nodes added to or removed from the cluster:

- NodeJoinedToCluster

- NodeFailedToJoinCluster
- NodeRemovedFromCluster
- NodeFailedToRemoveFromCluster

Figure 7 *Node Joined to Cluster Event*

Audit Logs

The following audit logs indicate the cluster management actions triggered by users from GigaVUE-FM:

- User <username> created cluster <clustername>
- User added device <device IP> to cluster <clustername>
- User removed device <device IP> from cluster <clustername>

Export Nodes and Clusters

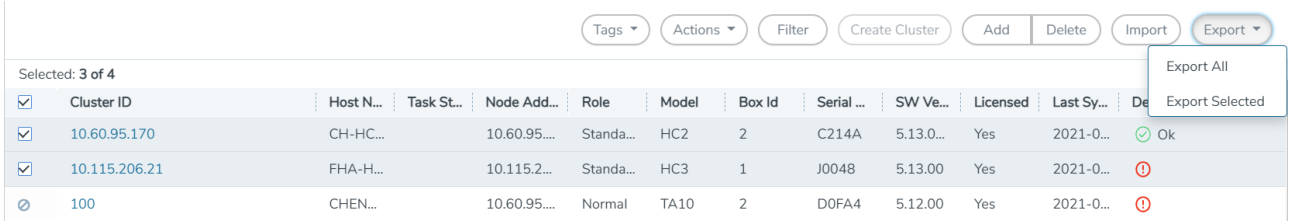
To export the nodes and clusters:

1. On the left navigation pane, go to **Inventory** and select **Physical > Nodes**.
2. In the physical nodes page, select the nodes and clusters you want to export.
3. Click **Export > Export Selected**.

Click **Export All** to export the nodes in a table format.

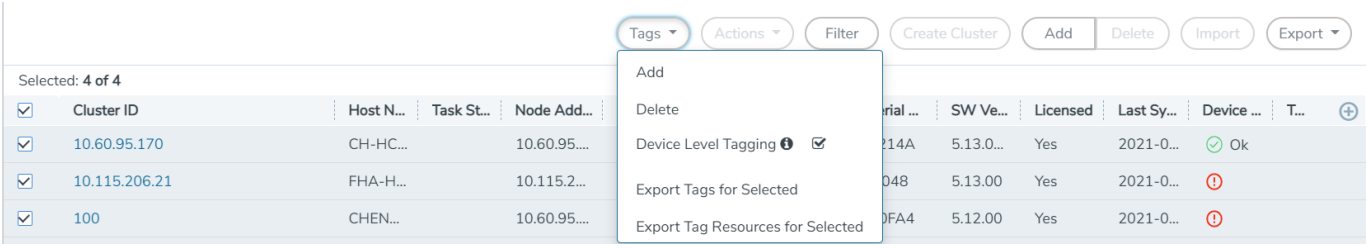
To export selected nodes:

1. Click **Select All.**
2. Click **Export Selected** to export only the selected nodes.



However, the member nodes will not get downloaded as they do not get selected. To select and download the member nodes:

1. Enable **Device Level Tagging** option under the Tags menu.
2. Click **Export Selected** to export the selected nodes, including the member nodes.



Problems with SCP?

After upgrading GigaVUE-FM to a new release, under some circumstances you may find that a previously-managed H Series node no longer accepts SCP commands to backup or restore configuration files. This can happen when the SSH keys in use change, causing a mismatch between the keys stored on the H Series node and those presented by GigaVUE-FM. Use the following steps on the H Series node to remove the GigaVUE-FM server from the H Series node's list of addresses, resolving the issue:

1. Log in to the affected H Series node and switch to Configure mode.
2. Use the **ssh client user admin known-host?** command to discover the IP address for the GigaVUE-FM server. For example:

   (config) # ssh client user admin known-host ? 10.150.100.23 10.150.100.77
3. The question mark (**?**) instructs the H Series node to list the known ssh clients. From the list of IP addresses returned by the CLI, identify the one belonging to GigaVUE-FM and remove it using the **remove** argument. For example, if 10.150.100.77 is the GigaVUE-FM server's IP address:

   (config) # ssh client user admin known-host 10.150.100.77 remove

   Return to GigaVUE-FM and attempt the configuration backup again.

Events

The Events page displays all the events that occur in the physical nodes or clusters. An event is an incident that occur at a specific point in time. Examples of events include:

- Device status change
- Stack image install status
- Fan tray changed

From the left navigation pane, go to **Inventory > Physical > Nodes**.. On the Physical Nodes page click **Events** on the left navigation pane.

For information about the parameters for each event, refer to the *"Event Parameters"* section in the *GigaVUE Administration Guide*. For filtering the events, refer to the *"Filter Events"* section in the *GigaVUE Administration Guide*.

NOTE: The events can be purged or archived only from the Events page. For more information, refer to the *"Archive or Purge Event Records"* in the *GigaVUE Administration Guide*.

Alarms

An Alarm is a response to one or more related events. If an event is considered of high severity, then GigaVUE-FM raises an alarm. Examples of alarms include:

- High or low port utilization
- High or low CPU utilization
- High exhaust temperature

The alarms broadly fall into the following categories: Critical, Major, Minor, or info.

Audit Logs

With Audit Logs, changes and activities that occurred in the physical nodes or clusters due to user actions can be easily tracked for auditing. The logs can also be further filtered to view specific information.

For information about the parameters in the audit log page, refer to the *"Overview of Audit Logs"* section in the *GigaVUE Administration Guide*. Filtering the audit logs allows you to display specific type of logs. For more information, refer to the *"Filtering Audit Logs"* section in the *GigaVUE Administration Guide*.

Search for Specific Nodes Using Keywords

The filter option provides a way for the users to narrow down the display using certain keywords such as Standalone, Clusters, H Series and others. As you click on the Filter button, you will see the quick view window pop-up.

The Filter quick view provides you filter criteria for your search. These options are available in the drop down menu under Criteria. You can further narrow the options using the Model, Software Version #, Cluster ID, Host Name, DNS Name, or Node IP. You are not required to fill in all these options to narrow your search. As you select these options in the quick view, you will see the options narrowing in the main window.

To clear or revert the search, do any of the following:

- To clear a part of the search, use backspace to clear the search entry and re-type a new option.
- To clear all the search criteria, use the Clear button on the top of the quick view window.
- To revert to the main window with the new searches, click on the **X** of the quick view window.

To revert to all nodes visible, you can use the clear filter option on top of the main window or the clear option in the quick view window. When no filters are in place, this option will no longer be visible in the main window.

Search for Ports on a GigaVUE Node

When viewing a node from GigaVUE-FM, you can use Port List Filter feature to display only certain ports that match specified criteria, such as only those ports that are used in a map.

To use the port list filter:

1. Select a physical node from the **Physical Nodes** page and then select the node to view.
2. Go to **Ports > Ports > All Ports**.
3. To filter the ports, click **Filter**. The Filter quick view is displayed.
4. Specify the criteria of the ports you want to filter.

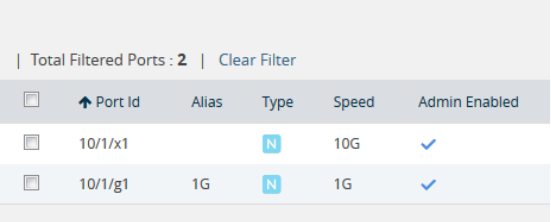
The criteria that you can use to filter the port list is as follows:

Criteria	Description
Box/Slot ID	Display only those ports that match the specified box and slot IDs.
Port Alias	Display port with the specified alias.
Port ID	Display ports with specified number in the port ID. For

Criteria	Description
	example, if you specify 3 the result will also display ports that include the number 3, 13, 23, 30, and so on.
Type	Display ports with the specified port type. Select one of the following: ▪ Network ▪ Tool ▪ Inline Network ▪ Inline Tool ▪ GigaSMART ▪ Hybrid ▪ Stack
Port Used in Map(s)	Display ports based on their usage in maps. The possible selections are: ▪ All — display all ports either unused or in use by maps. This is the default. ▪ In Use — display ports that are in use by any map. ▪ Unused — display ports that are not use by any maps.
Admin Status	Display ports based on their current admin status. The possible selections are: ▪ All — display ports with a status of Enabled or Disabled. This is the default. ▪ Enabled — display ports with admin enabled. ▪ Disabled — display ports with admin disabled.
Speed	Display ports with the selected port speed. The port speeds available depend on the node.
Transceiver Type	Display ports with the selected transceiver type. The transceivers available selection depend on the type of transceivers connected to the ports.

To remove the filter selections, click **Clear**.

After the filter is applied, the Ports page displays only the ports that correspond to the selected filters and shows the total number of ports that meet the criteria. To clear the filters, select **Clear Filter**. The figure below shows the Port pages with two ports that correspond to the current filters: Network Type and Admin Status Enabled.



Overview Page

The **Overview** page displays general information on the specific H Series node, which includes System, Ports, Maps, and Traffic information.

GigaVUE-FM notifies you with a message when a node is down. Click **Proceed** to continue viewing the overview dashboard even when the node is unreachable.

Systems Information

Systems information is displayed on the System widgets.

The Systems widget displays general information about the specific device that you selected from the drop-down list at the top of the widget. If the system is a cluster, you can select a device in the cluster to display on the widget.This widget gives you a quick status if any issues are present in any of the device's components through color indicators; green (running), amber (warning), or red (alert).

NOTE: Ensure that all the nodes and clusters have a Box ID defined. If the Box ID is missing, the Systems widget may not display any information relating to the node.

NOTE: Red alert appears for cards not present.

Field	Description
Host Name	The host name of the box.
Hardware	The hardware type, (for example, GigaVUE-TA1).
Software	The version of the software running on the device.
Memory	Shows the amount of used and free memory.
Load Average	The average load on the system over the last 1 minute, 5 minutes, and 15 minutes.

Field	Description
	NOTE: When a new device is added to GigaVUE-FM, it takes one stats cycle for the average load value to be reflected in the GigaVUE-FM GUI.
Cards	Displays all slots for the specific hardware type including its slot number and the type of card it contains or not. **Note:** When you hover over the card slot, the **temperature** is displayed.
Fan Trays	Indicates that the Fans are On or Off.
Power Supply	Indicates that the power supply is On or Absent. **NOTE:** When one or more power supply units are down, red alert is displayed.

Failure to Authenticate

To view a physical node in the Dashboard System pane, your login credentials must have the appropriate permissions. Otherwise, GigaVUE-FM shows an error message.

There are two possibilities that caused a user authentication error:

- The login user credential for GigaVUE-FM is not "**admin**".
- The password associated with the login user name for GigaVUE-FM is different for the physical device.

NOTE: When a "non-admin" user goes to Physical Nodes page, they are able to view displayed physical nodes with the status as Green, Amber or Red, this is because the physical node information is captured using the default "admin" user role.

Ports Information

The Overview page displays widgets that provide port information for the number of ports down, the number of ports with packet drops, and the number of over utilized receiving (Rx) ports and transmitting (Tx) ports. The ports widgets default to displaying a counter. Clicking on the icon in the upper right-hand corner displays the information as a table. The Ports with Packet drops and Over-Utilized Ports widgets are similar.

Traffic

The Traffic widget the shows most-utilized ports and ordered by traffic count. Each displayed port is labeled with its location, whether it is a transmitting or receiving port, and its percentage of utilization.

NOTE: The Traffic pane is view-only. It reflects traffic activity with port ID at the time of discovery and does not immediately refresh.

Workflows

The Workflows page provides wizards for creating maps. These wizard step you through the work flow to make sure you configure all of the components necessary for configuring out-of-band and inline maps for traffic flow monitoring. The workflows keep track of each step so that you can stop and then return to were you left off in the workflow. However, you can only work on one workflow at a time.

Overview of Workflows

Table 2: Map Types and Map Wizards describes the maps that you can create with the wizards.

Table 2: Map Types and Map Wizards

Map Type	Map Wizard	Description
Out-of-Band maps	Map with rules	Walks you through the steps to select source and destinations ports, then create a Regular By Rule map with those ports.
	Pass-all map	Walks you through the steps to select source and destinations ports, then create a Pass All map with those ports.
	Collector map	Walks you through the steps to select source and destinations ports, then create a Collector map with those ports.
Inline maps	Map with rules	Walks you through the steps to select the destination and source, then use them in an Inline By Rule map. If the Inline Networks or Inline Network Groups have not been created yet, you can create them from the wizard. The wizard selects Symmetric for Traffic Type.
	Pass-all map	Walks you through the steps to select the destination and source port, then use them in an Inline Pass All map. If the Inline Networks or Inline Network Groups have not been created yet, you can create them from the wizard.

Map Type	Map Wizard	Description
		The wizard selects Symmetric for Traffic Type.
	Collector map	Walks you through the steps to select the destination and source ports, then use them in an Inline By Rule map. If the Inline Networks or Inline Network Groups have not been created yet, you can create them from the wizard. The wizard selects Symmetric for Traffic Type.
	Asymmetric Inline map	Walks you through the steps to select the destination and source ports, then use them in an Inline Pass All map. The source port must be an Inline Network port. The destination port must be an Inline Tool, an Inline Tool Group, or an Inline Serial Tool Group. If the Inline Networks or Inline Network Groups have not been created yet, you can create them from the wizard. The wizard selects Asymmetric for Traffic Type.
Basic out-of-band GigaSMART maps:	Map with GigaSMART Apps	Walks you through the steps to select a GigaSMART Group, GigaSMART operation, select the source and destination ports, and then create a Regular By Rule map. The wizard allows you to create the GigaSMART Group and GigaSMART operation if they do not already exist.
	First level map	Walks you through the steps to select or created the source port and select a or create a virtual port for the destination port, and then create a First Level By Rule map.
	Second level map	Walks you through the steps to select or created the virtual port for the source port and select a or create the port for the destination port, and then create a Second Level By Rule map.
Advanced out-of-band GigaSMART maps	Map with NetFlow	Walks you through the steps to select or create an IP interface; a NetFlow exporter, records, and monitor; a GigaSMART Group and Operation, source ports; and then created a Regular by Rule map.
	SSL-based map	Walks you through the steps to create or select a GigaSMART Group, configure SSL, create or select a GigaSMART Operation, and then create a Regular By Rule map.
	Map with ASF	Walks you through the steps to create or select a GigaSMART Group, virtual port, and GigaSMART Operation; configure GigaSMART Application Session Filter; and create the First and Second Level maps needed for implementing an ASF solution.

How to Use Workflows

To start a workflow, click on one of the links, such as **Map with Netflow**. When you click on the link, the Workflow page is displayed. The Workflow panel on the right shows the tasks in the workflow, indicating the current task.

While using workflows, you can only work on one workflow at a time. Also, you cannot roll back changes made to a node after canceling a workflow.

A task in the workflow allows you to select an item or create the item if one does not exist. For example, a GigaSMART Group needs to be selected if one does not exist. In this case, click **Create**. Clicking Create takes you to the GS Group configuration page. Configure the GigaSMART Group and click **Save**.

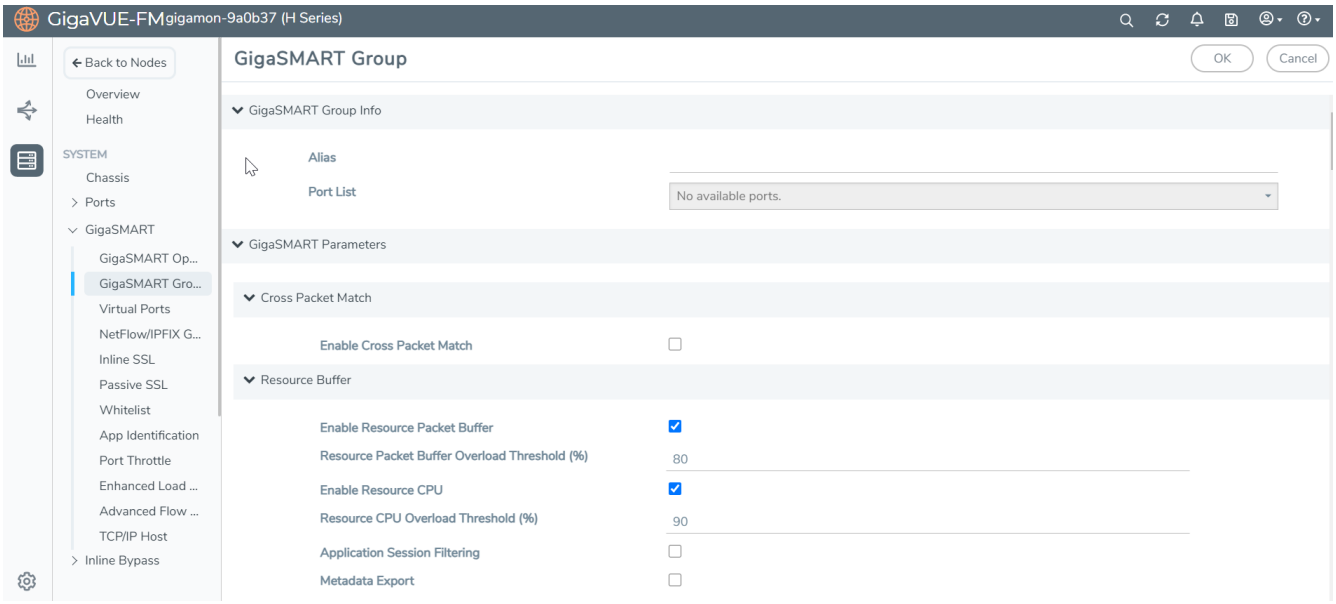


Figure 8 *GS Group Configuration in Workflow*

After saving the configuration, the Workflow moves to the next task and the Workflow panel indicates which tasks have been completed.

You can select a task in a different order than shown in the Workflow panel. For example, you can go to the NetFlow Monitor task. After completing the task, the Workflow returns you to the to the GS Group configuration page with the Monitor field completed.

Workflow allows you to leave the current workflow and return at anytime during a GigaVUE-FM session. The **In Progress** panel indicates the current workflow and the Workflow panel indicates the competed tasks. Figure 9Workflow in Progress shows an example of workflow in progress. You can abandon a workflow by clicking the red x in the **In Progress** panel.

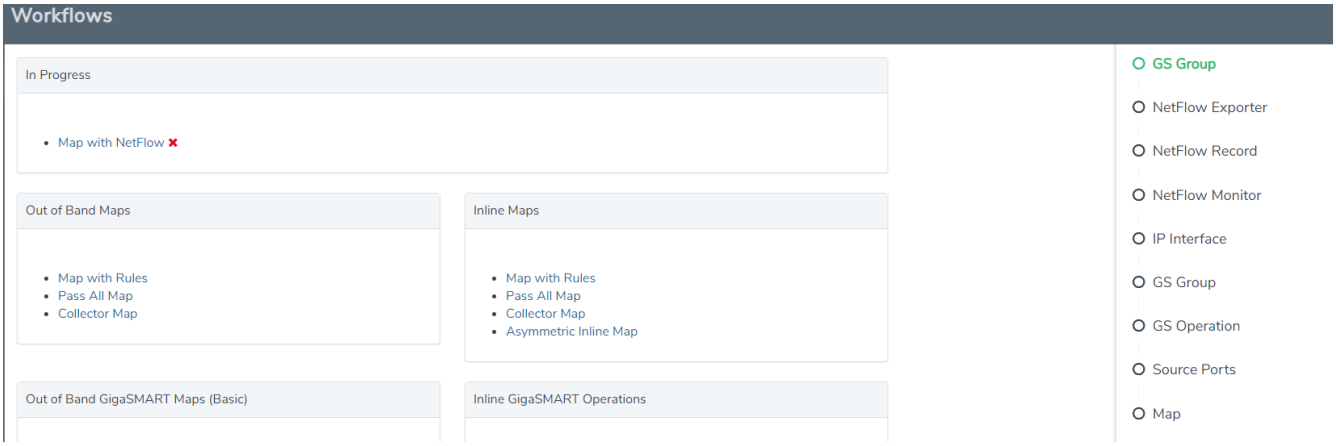


Figure 9 *Workflow in Progress*

When a workflow is completed, a page displays providing you with several choices for the next task. For example, when the Map with Rules workflow is completed, you can go to creating a collect map by clicking the **Create a Collector Map**, return to the Workflow page by clicking **To Work Flows**, or go to the Maps page by clicking **To Maps**.

Chassis Table View

When viewing the Chassis Table View managed from GigaVUE-FM, the Table View includes environment information about the chassis.

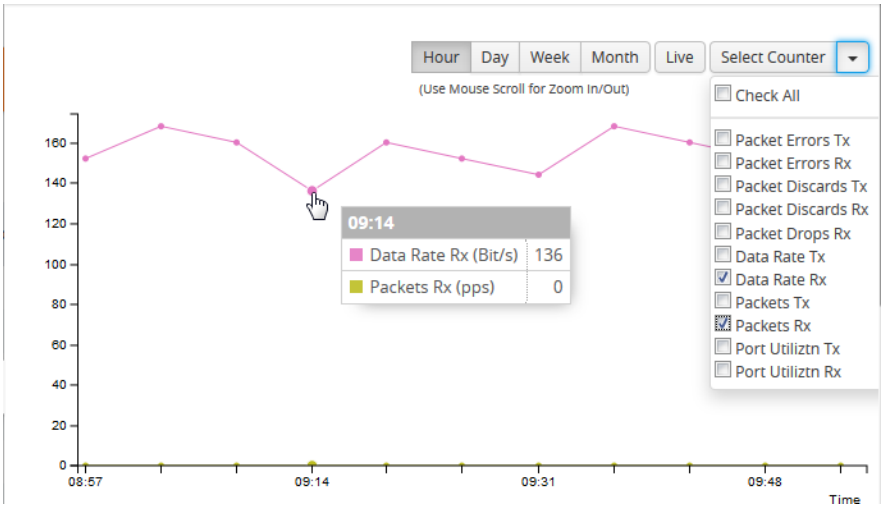
The Table View provides the following information about the chassis and its components:

Chassis Information	Description
Properties	Provides information about the chassis.
Cards	Describes the cards installed in each slot of the chassis, including its current status.
Environment	Provides temperature and voltage information about the chassis.
Power Supplies	Describes the power supplies installed in the chassis, including their current status.
Fan Trays	Describes the fan trays installed in the chassis, including their current status.
Fan RPM	Provides the current RMP of the each fan.

NOTE: When a new device is added to GigaVUE-FM, it takes one stats cycle for the following field values to be reflected in the GigaVUE-FM GUI: Environment, Power Supplies, Fan Trays, Fan RPM.

Live Graphing

When viewing ports on a node running GigaVUE-OS 4.6 or later from GigaVUE-FM, you can select to see the graph data display in real time by clicking **Live**. This changes the updating of the information in the graph from the default to every 10 seconds. You can also select the data to display on the graph by selecting an option from the Select Counter list. Live GraphingThe following figure shows an example, where Live is selected and the Data Rate Rx and Packets Rx counters are selected.



Safe and Limited Modes

Starting in GigaVUE-OS software version 4.7, safe and limited modes are introduced to the cluster environment and standalone nodes.

During clustering operations, in rare scenarios, there can be unrecoverable system errors that can potentially put the cluster or the clustered nodes into unsafe or unstable states. Once in such a state, additional operations or configuration changes can cause the node to crash, the cluster to deform, and the data traffic to be impacted. For example, due to a node attempting to rejoin a cluster, a chassis can end up in a reboot loop. In previous software versions, there was no way to prevent entering the loop.

These modes provide notification, stop further operations from being performed, and give you time to troubleshoot and plan the recovery of the cluster or of any node in the cluster or standalone node.

Two modes are supported. The first is called safe mode and is triggered when the node detects unrecoverable errors, but the existing flow maps are not impacted. The second is called limited mode and is triggered when the node detects continuous system reboots. In this mode, the node will become standalone and only basic configuration will be allowed.

When a cluster is in safe mode, GigaVUE-FM displays a Safe Mode banner and message.

Safe Mode

A node enters safe mode when there are unrecoverable errors. Any node in a cluster can enter this mode.

Examples of unrecoverable errors are when there are inconsistencies between the system and the running configuration or when the cluster configuration did not merge properly with the existing configuration. A node will automatically enter safe mode.

As part of merge error recovery, nodes joining a cluster are automatically restarted so the merge error can be fixed.

When a node is in safe mode:

- The node displays a banner indicating it is in safe mode.
- An SNMP trap is sent to notify the user when the mode changes
- Traffic provisioning is not allowed on the affected node. Any other configuration remains as is.
- Configured traffic continues to be forwarded.
- If the standby node in the cluster is in safe mode, it can still become the leader if the current leader fails or switches over, but the database on the standby node may not be in sync, so it is not recommended to continue in that state. Instead, take immediate action to recover the node.
- In safe mode, the node does not process any incoming traffic configuration from the cluster leader.

When a node is in safe mode and you try do any operations that are not allowed in safe mode, the UI displays the following message:

```
The system has restricted provisioning in safe mode. Contact Gigamon Support on
how to troubleshoot and recover from safe mode.
```

Also, hovering over the status bubble of the nodes on the Physical Nodes page in GigaVUE-FM displays a message that the node is in Safe Mode.

To exit safe mode, reload the node.

Limited Mode

A node automatically enters limited mode when it detects repeated system crashes.

When a node is in limited mode:

- The node displays a banner indicating that it is in limited mode.
- An SNMP trap is sent to notify the user when the mode changes.
- Only basic system provisioning is allowed. Traffic provisioning is not allowed. Only commands that are related to image download, installation, next boot, and reboot are allowed.

Limited mode is triggered when there are three (3) failures/system crashes within 15 minutes. In limited mode, the cluster configuration is ignored. No cluster configuration or GigaVUE-OS configuration is accepted when the node is in limited mode.

When a node is in limited mode, a Limited Mode banner displays in GigaVUE-FM.

Enable SNMP Trap for Safe Mode and Limited Mode

Use the following steps to configure a notification that will be sent to all configured destinations when a node in the cluster changes from operational mode to safe mode or from operational mode to limited mode

The safe mode and limited mode capabilities are enabled through the SNMP trap event Operational Mode Change. To enable the trap on a node, do the following:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**.. On the Physical Nodes page, select the node you want to configure.
2. Select **Settings > Global Settings > SNMP Traps**.
3. Click **Trap Settings**.
4. On the Edit SNMP Traps Settings page, select **Operational Mode Change**.
5. Click **Save**.

When the cluster leader enters safe mode, the SNMP trap will be sent and the leader will be identified as the local node in the trap.

When a node in a cluster enters safe mode, the SNMP trap will be sent and the node will be identified as the local node in the trap. In addition, a notification will be sent to the cluster leader. The node that entered safe mode will be identified by its box ID in the notification to the leader.

Log messages also provide information. The following is a sample log:

Jun 8 13:46:27 GC-TA10-N6 mgmtd[2400]: [mgmtd.INFO]: SAFE mode: Merge error detected !! Triggering SAFE mode ...

Collect Information for Technical Support

Collecting the following information can help Technical Support:

- sysdumps/debug dumps for all nodes in the cluster
- sysdumps for nodes that observed a crash entering safe or limited mode
- debug dumps for nodes that did not observe a crash
- gslogs for application information
- console logs
- CLI histories
- CLU or GigaVUE-FM screen captures
- SNMP captures

Manage Not-reachable Nodes in Cluster

GigaVUE-FM allows you to add and manage standalone nodes and clusters as described in Manage GigaVUE® Nodes and Clusters. Occasionally, the nodes in a cluster may become unreachable due to various reasons.

Until software version 6.0.00, when a node in a cluster managed by GigaVUE-FM is removed from the cluster using the **`no cluster enable`** CLI command, or when any node becomes unreachable to the cluster leader due to other reasons such as power outage, the cluster leader marks the operational status of the node as 'Down' in CLI. After the next config sync cycle, GigaVUE-FM does not manage the node anymore. An event is triggered and corresponding alarms are raised. When a node in a cluster managed by GigaVUE-FM is removed using the GigaVUE-FM GUI (Edit Cluster page), GigaVUE-FM manages the node as a standalone node. The Operational Status of the node is marked as 'Down' in CLI.

When a cluster with one of the nodes with operational status 'Down' is added to GigaVUE-FM, GigaVUE-FM removes the node from the cluster and no longer manages the node. If the node comes up, the node is added to GigaVUE-FM, only if the node is not already managed by GigaVUE-FM. Otherwise, you must manually:

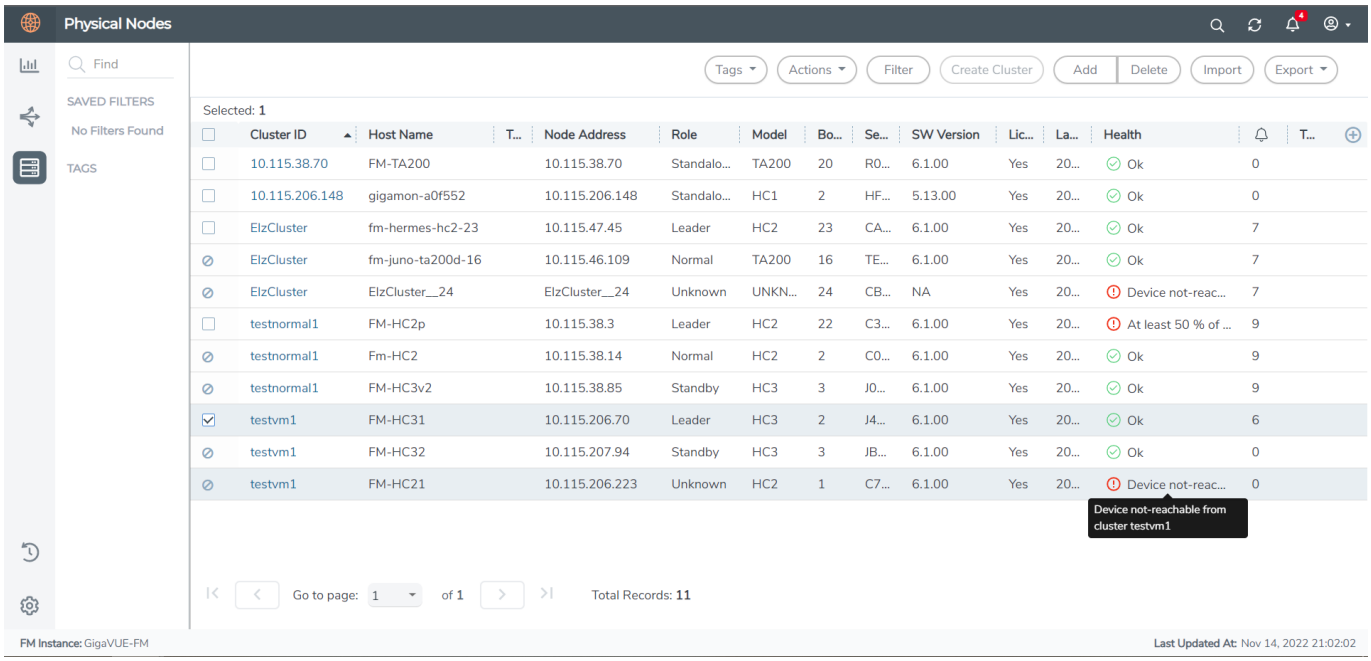
- Add the node as a stand-alone node in GigaVUE-FM
- Add the node to the cluster

Therefore, it is important to keep track of the events and alarms to know about the nodes that have been removed from the cluster, especially in large deployment scenarios.

Starting from software version 6.1.00, when a node become unreachable, the nodes are designated with the following operational statuses (depending on the reason):

Reasons for node becoming unreachable	Operational Status
Nodes in cluster is removed intentionally from the cluster: • Using the **`no cluster enable`** CLI command. • Using GigaVUE-FM GUI (by navigating to the Edit Cluster page)	**Left**
Nodes is offline due to issues such as power outage, management interface being down, or crash.	**Not-reachable**

GigaVUE-FM manages the not-reachable member nodes as part of the cluster and does not remove the nodes from the cluster, that is, GigaVUE-FM manages the not-reachable nodes in faulty/offline state. This allows easy management of the not-reachable nodes. GigaVUE-FM displays clusters with not-reachable nodes in the Physical Nodes page and in the Chassis page.



To view only clusters with not-reachable node from the Physical Nodes page:

1. Navigate to the Physical Nodes page.
2. Click the **Filter** button.
3. Scroll down and enable the following toggle option to view only the clusters with not-reachable nodes:

   Show only the clusters with not-reachable nodes

   Refer to the following notes:

- If the node is already managed by GigaVUE-FM, then the Host Name and IP Address of the not-reachable nodes is displayed in the GUI.
- If you add a cluster with not-reachable node to GigaVUE-FM, then the IP address and Host Name of the not-reachable node will be displayed in the following format: **clusterid_boxid**. The IP address of the not-reachable node will get updated once the node becomes reachable. Role of the not-reachable node will be 'Unknown'. Software Version and other properties will be displayed as NA.

Rules, Notes, and Limitations

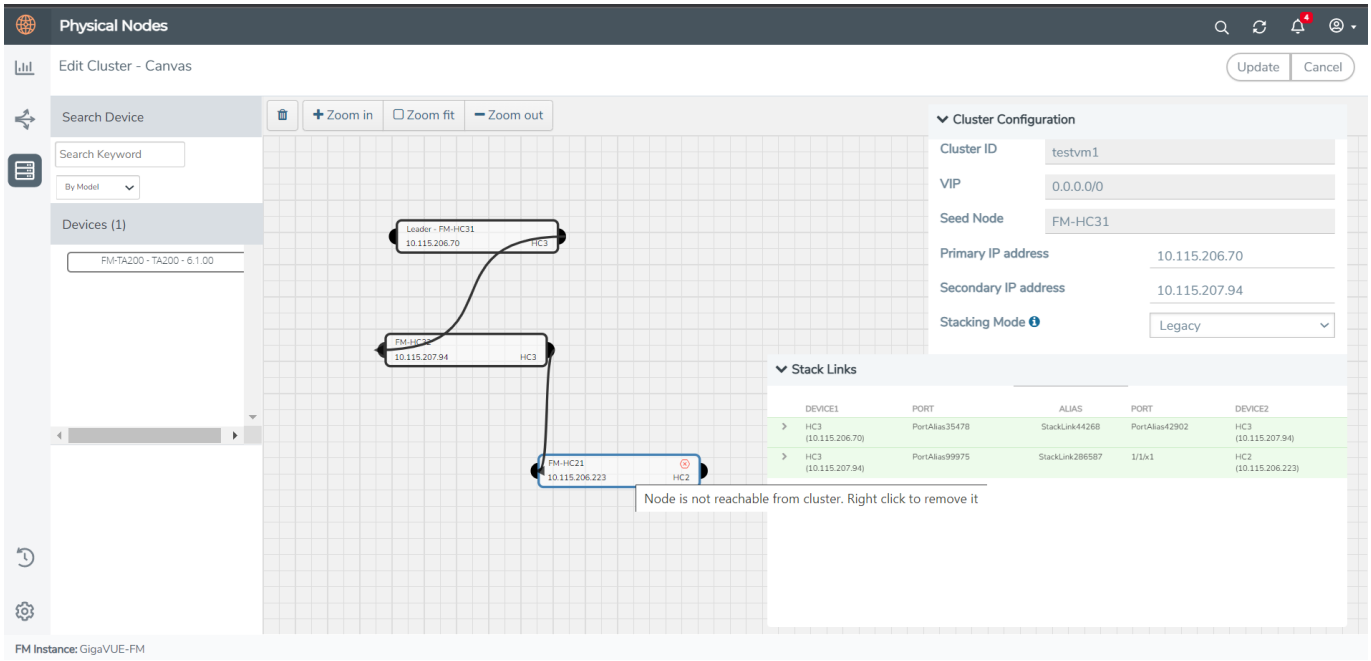
Refer to the following rules, notes and limitations:

- You can only remove the not-reachable member node from the cluster. You cannot perform any other write operation on the not-reachable node.
- You can configure Email notification (instant or digest) for the events triggered for the not-reachable nodes. Refer to the Configure Email Notifications section in the GigaVUE Administration Guide for details.
- The not-reachable member nodes are captured in the physical and Fabric Health Analytics dashboards.
- You can import and export clusters with not-reachable member nodes.
- You must perform a factory reset of the not-reachable nodes when:
  - removing a not-reachable node from a cluster and adding it to a new cluster
  - removing a not-reachable node from a cluster and adding it as a leader node of another cluster
- When rebooting a cluster with not-reachable nodes, GigaVUE-FM prompts you to skip the not-reachable nodes. Reboot will not proceed until you skip the not-reachable nodes.

Remove Offline Chassis

GigaVUE-FM allows you to remove the not-reachable nodes from the cluster. To remove the non-reachable nodes:

1. From the Physical Node page, select the required cluster from which you want to remove the not-reachable nodes.
2. Select **Actions** > **Edit Cluster**. The Edit Cluster - Canvas appears.
3. Right click on the not-reachable node, and click **Delete**.



To perform this action:

- You must be a read-write user with Infrastructure Management category.
- You must remove the associated ports and map configuration of the not-reachable member nodes before removing the nodes from the cluster. Without doing this, GigaVUE-FM does not allow you to remove the nodes.

Alarms and Health Status of Not-reachable Nodes

The following alarms are triggered in the Alarms page depending on the status of the not-reachable nodes:

State	Alarm
Member node in a cluster is not-reachable to the cluster leader. The alarm is cleared when the member node is reachable.	***Cluster Member Not-Reachable***
Member node is reachable by the cluster and added to the same cluster	***Cluster Member Online***

When a not-reachable node becomes part of another cluster:

- The config-sync operation fails.
- The cluster to which the node originally belonged to displays the member node as not-reachable.
- The health status of the new cluster becomes red.

You are responsible for removing the not-reachable member node from the cluster to which it originally belonged to only after which the config sync operation will succeed.

Upgrade Cluster with Not-reachable Nodes

When upgrading a cluster from software version 6.0.00 to 6.1.00:

Member nodes with Operational Status as Down: Cluster upgrade is successful in both CLI and GigaVUE-FM. After the upgrade, the operational status of the member nodes is *left*.

While upgrading a cluster from software version 6.1.00 to higher, if the cluster contains:

- **Member nodes with Operational Status as Not-reachable**: Cluster upgrade will not proceed until the not-reachable nodes are skipped. Refer to the following table for the various scenarios:

Type of Upgrade	Cluster State	Upgrade Status
Immediate Upgrade	All nodes in the cluster are reachable	Upgrade will succeed.
	One or more nodes in the cluster are not-reachable	Enable the "Skip not-reachable nodes to upgrade" check-box for the upgrade to succeed.
Scheduled Upgrade	All nodes in the cluster are reachable.	Upgrade will succeed.
	All nodes in the cluster are reachable. However, it is possible for the nodes to go to not-reachable state at the scheduled upgrade time.	Enable the "Skip not-reachable nodes to upgrade" check-box for the upgrade to succeed.
	One or more nodes in the cluster are not-reachable	Enable the "Skip not-reachable nodes to upgrade" check-box for the upgrade to succeed.

> **NOTE:** After the cluster upgrade operation is completed, you must ensure to upgrade the not-reachable nodes to the required software version. Failure to do so will result in version mismatch conflict.

For GigaSMART Signature image upgrade, you must skip the not-reachable nodes for the upgrade to succeed.

Type of Upgrade	Cluster State	Upgrade Status
Immediate Upgrade	All nodes in the cluster are reachable	Upgrade will succeed.
Scheduled Upgrade	One or more nodes in the cluster are not-reachable	Enable the "Skip not-reachable nodes to upgrade" check-box for the upgrade to succeed.

Backup and Restore of Cluster with Not-reachable Nodes

During GigaVUE-FM backup and restore operation, if not-reachable nodes are part of the cluster:

- The not-reachable member nodes are also backedup. On restoring the backup file, the not-reachable member node will be restored to its original state.
- During the next config sync cycle after restore, GigaVUE-FM updates the status of the not-reachable member nodes and the cluster.

For cluster backup and restore operation, refer to the following table:

Operational State During backup	Operational State During Restore	Final Operational State
All nodes are in 'up' state	Few nodes are in 'left' or 'not-reachable' state	Operational status of the nodes in 'left' or 'not-reachable' state will remain as 'left' and 'not-reachable', respectively.
	Few nodes are removed using **`no chassis box id`** command	Operational status of the removed nodes is configured as 'left'.
Few nodes are in 'left' or 'not-reachable' state	Same state	No change.
	Few nodes are removed using **`no chassis box id`** command	Operational status of the removed nodes is configured as 'left'.
	Some nodes come up	Operational status of the nodes will show as up.

Multi-Path Leaf and Spine

This chapter describes the leaf and spine architecture with multiple paths for achieving high availability in a cluster environment. Refer to the following sections for details:

- Introduction to Multi-Path Leaf and Spine
- Configuration Overview
- Leaf-Spine Cluster Deployment

NOTE: Refer to Regular Cluster Formation Workflow for how to use the Regular Cluster workflow.

Introduction to Multi-Path Leaf and Spine

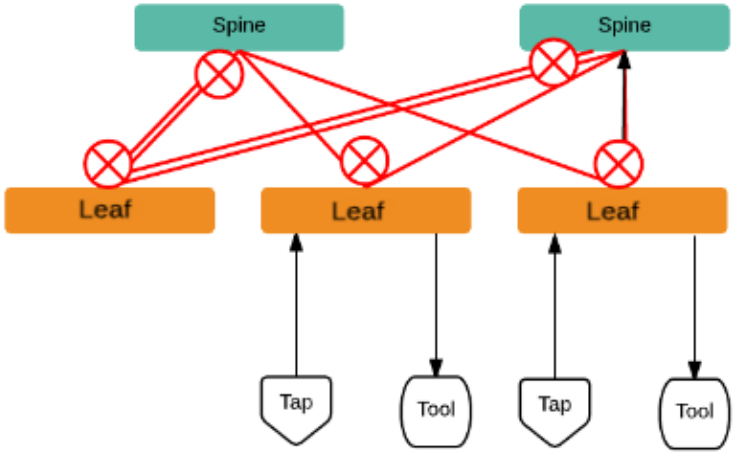
The leaf and spine architecture is a two-layer architecture used for network aggregation. There are two kinds of nodes in this architecture, as follows:

- leaf nodes, which are edge nodes and can also have TAPs or tools attached to them
- spine nodes, which are the nodes to which the leaf nodes attach

NOTE: The connections between leaf node and spine node is through
- stack GigaStreams configured on both the leaf and the spine nodes and
- spine links configured on the leaf side.

With multiple paths between the nodes in a cluster, the leaf and spine architecture protects against failures, such as stack link or spine node failures. In the event of a failure, the traffic fails-over to the other available path.

In this architecture, each leaf node connects to every spine node. This forms a mesh between the leaf and spine nodes. However, no leaf node directly connects to another leaf node and no spine node directly connects to another spine node. An example of a spine and leaf architecture is shown in the following figure.



In a cluster, the number of leaf nodes is typically greater than the number of spine nodes. In the above figure, there are three leaf nodes and two spine nodes. The leaf nodes aggregate to a fewer number of spine nodes.

The spine nodes are generally GigaVUE® TA Series nodes, such as GigaVUE-TA100, while the leaf nodes are GigaVUE® HC Series nodes, such as GigaVUE-HC1 and GigaVUE-HC3, which places the traffic intelligence at the edge.

The figure shows TAPs or tools connecting to the leaf nodes, and the leaf nodes connecting to the spine nodes. Note that TAPs or tools do not connect to the spine nodes.

Instead of one leaf node connecting to one spine node with a single link, in this architecture there are multiple links from the leaf nodes to the spine nodes. The leaf nodes connect to the spine nodes through at least two paths. Some leaf nodes with higher capacity, such as GigaVUE-TA100, can have more paths, as shown by double red lines in the figure.

Traffic between ports on a leaf node will be local to that leaf node, but traffic between leaf nodes will go through the spine nodes.

The traffic from a source leaf node to a destination leaf node flows as follows:

- From a TAP, traffic flows to the source leaf node
- From the source leaf node, traffic is load balanced to all spine nodes
- From a spine node, traffic flows to the destination leaf node
- From the destination leaf node, traffic flows to tool ports

Resiliency is achieved when there are multiple paths from the network to the tools across GigaVUE nodes.

Path Protection

The spine leaf architecture in the cluster environment provides failover for the following:

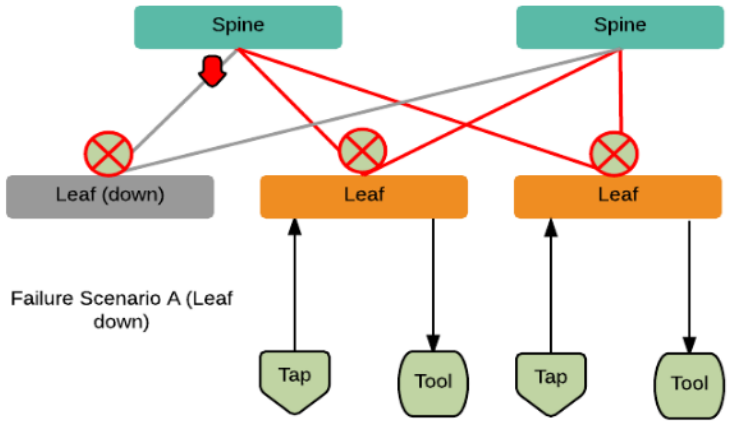
- leaf node failure. Refer to Leaf Node Failure.
- stack link failure on a leaf node (not connected to a tool, but can be connected to a network TAP). Refer to Stack Link Failure on Leaf (TAP Connected).
- spine node failure. Refer to Spine Node Failure.
- stack link failure on a leaf node (connected to a tool). Refer to Stack Link Failure on Leaf (Tool Connected).

Leaf Node Failure

Refer to the figure for a failure in which a leaf node is powered down or rebooted. The leaf node does not have a connected TAP or tool.

NOTE: In the following figures, red arrows indicate traffic direction.

With this type of failure, the stack links connected to the affected leaf node go down, which will be detected by the spine nodes. No action will be required at the spine nodes. Since the links are down, no traffic will be sent to the affected leaf node.



Restoration

Once the leaf node is powered up and booted, it will restore its traffic configuration.

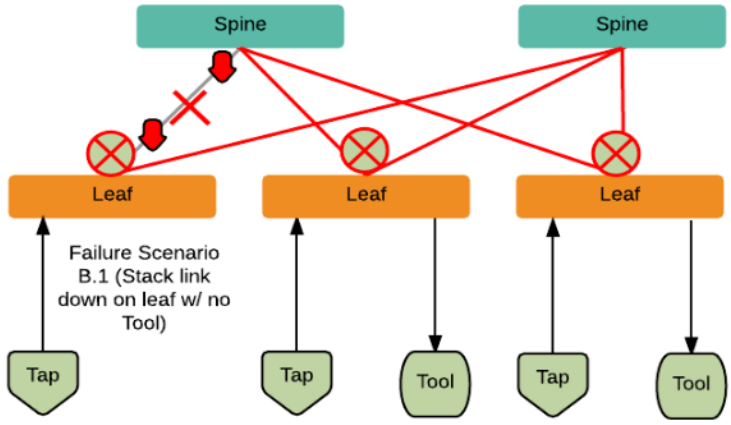
Affected Time

None. Traffic on other leaf nodes will not be affected.

Stack Link Failure on Leaf (TAP Connected)

Refer to the following figure for a failure in which a stack link on a leaf node fails and the leaf node is connected only to a TAP.

With this type of failure, the stack link between the leaf and spine nodes goes down. No action will be required at the spine node. At the leaf node, the affected link will be removed from the stack GigaStream. Traffic will be sent to the other spine node.



Restoration

When the link comes back up, the leaf node will put the link back into the GigaStream.

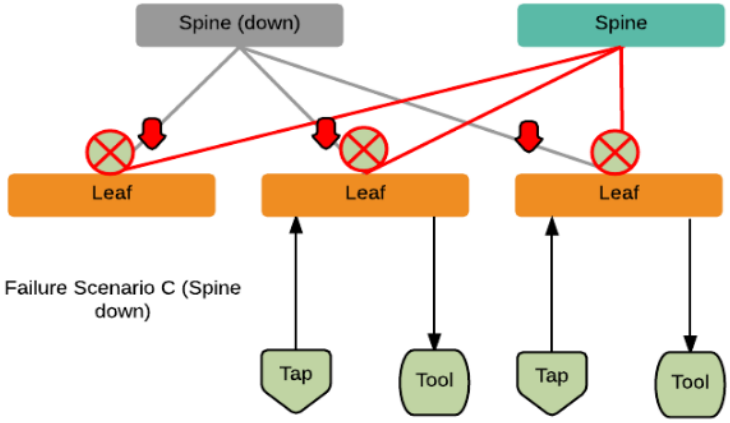
Affected Time

When the link is down, traffic recovers in a similar amount of time as a tool GigaStream.

Spine Node Failure

Refer to the following figure for a failure in which a spine node is powered down or rebooted.

With this type of failure, the stack link between the leaf and spine nodes goes down. The leaf nodes will detect that the stack link is down. The affected link will be removed from the stack GigaStream. Traffic will be load balanced to the other spine node.



Restoration

When the spine node reboots, the cluster will synchronize. When the node converges to the cluster and the configuration synchronizes, traffic will be restored.

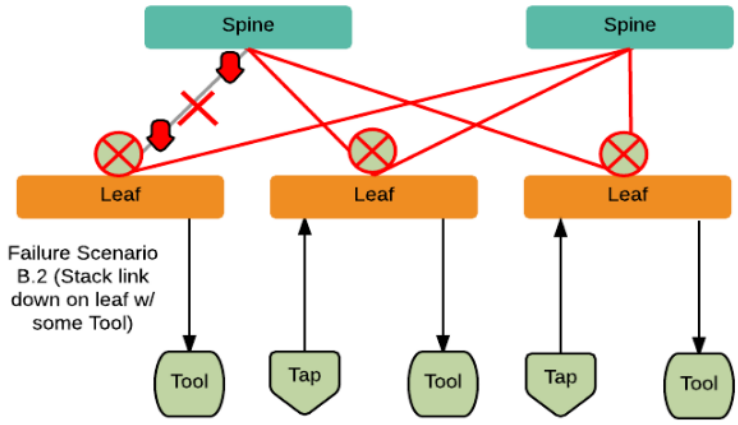
Affected Time

When the spine node is powered down or rebooted, traffic recovers in a similar amount of time as a tool GigaStream.

Stack Link Failure on Leaf (Tool Connected)

Refer to the following figure for a failure in which a stack link between the leaf and spine nodes fails and the leaf node is connected to a tool.

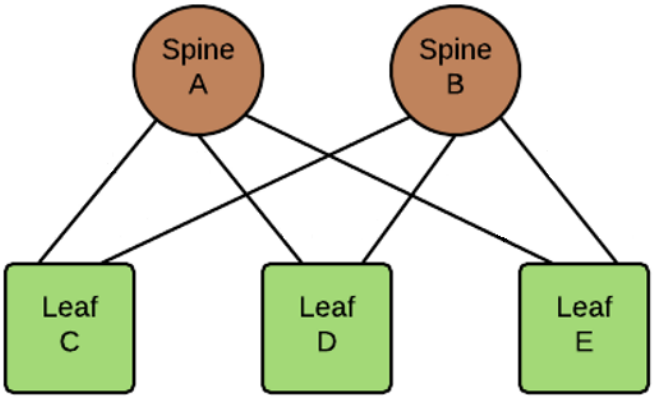
In the current software version, this type of failure is not supported.



Configuration Overview

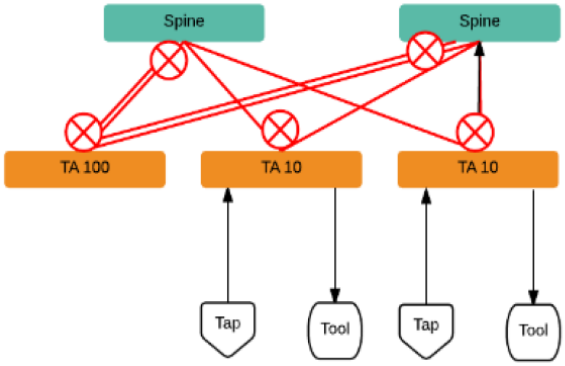
This section provides an overview of the configuration. The configuration is done from the leader in the cluster. Follow this configuration sequence to prevent loops.

This configuration connects nodes using multiple paths. For an example of the configuration, refer to the following figure.

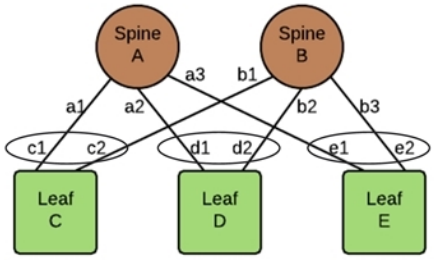


The configuration steps are as follows:

- Configure stack GigaStream. The stack GigaStream connect the spine and leaf nodes. In the following figure, the stack GigaStream are: a1, a2, a3, b1, b2, b3, c1, c2, d1, d2, e1, e2. Even if there is only one port that connects the nodes, you must still configure a stack GigaStream. With a configuration of two spine nodes and three leaf nodes, the number of stack GigaStream is 12.

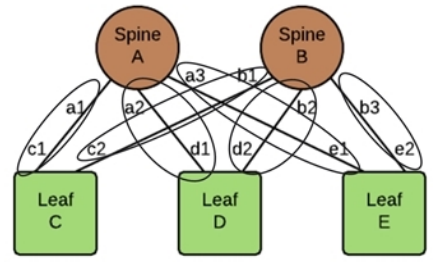


- Configure spine links. On each leaf node, there is one spine link that contains the list of GigaStream connecting the leaf nodes to the spine nodes. The spine links contain multiple stack GigaStream that are bundled together. The spine links are: {c1,c2}, {d1,d2}, and {e1,e2}. The total number of spine links is three for this configuration. The spine links are located at the leaf nodes. Across the spine link members, traffic is load balanced. For this part of the configuration, refer to the circles in the following figure.



NOTE: For the spine links, make sure that all paths are reachable.

- Configure stack links. The stack links are: {a1,c1}, {a2,d1}, {a3,e1}, {b1,c2}, {b2, d2}, and {b3, e2}. The total number of stack links is six for this configuration. For this part of the configuration, refer to the circles in the following figure.



These configuration steps ensure that the spine and leaf nodes are fully meshed.

Notes and Considerations

Refer to the following notes and considerations:

- The multi-path leaf and spine architecture is only supported in an out-of-band cluster.
- The spine link GigaStream must be of type stack. Stack GigaStream carry bi-directional traffic.
- All spine link GigaStream must have the same parameters, such as the same hash value and failover mode.
- Once a spine link is configured, editing of GigaStream parameters is not supported, except for editing the comment.
- Adding a new node to an existing leaf-spine topology is not supported. If you need to add a new node, ensure that you remove the existing configurations and reconfigure them for the new topology.
- GigaStream must be configured before spine links are configured.
- Once a GigaStream is configured in a spine link, it cannot be deleted. To delete a spine link, the stack links must first be deleted.
- A spine link cannot be deleted if a map is using the spine link.

- A spine link cannot be created if a map is using the GigaStream.

The number of spine and leaf nodes is not limited. The ratio of spine and leaf nodes are determined by the cluster traffic needed between the leaf nodes. Larger topologies have the same restrictions as the GigaVUE-OS as follows:

- the total number of nodes in a cluster, for example, 32
- the number of links in a GigaStream (which depends on the GigaVUE node and line card or module, for example, the PRT-HC3-X24 module on GigaVUE-HC3 can have 24 stack GigaStream)

Leaf-Spine Cluster Deployment

This section describes the steps and prerequisites to deploy a leaf-spine cluster.

Refer to Introduction to Multi-Path Leaf and Spine for a conceptual overview of the leaf-spine architecture.

Deployment Checklist

Before forming a Leaf-Spine Cluster, it is strongly recommended that you get familiar with the relevant documentation and review the deployment checklist to prepare for deployment.

Pre-deployment checklist

- Gigamon Fabric Management must be upgraded to GigaVUE-FM 5.3.00 or later
- Gigamon device must be upgraded to GigaVUE-OS 5.2.00 or later
- Advanced Features License must be installed in TA devices
- Physical connection must be established to create stack links
- Devices must have GDP enabled and be physically connected to create links among devices from GigaVUE-FM.

IMPORTANT: Recommendation is to use TA devices as SPINE Nodes and other devices as LEAF Nodes.

Formation Scenario

The Leaf-Spine cluster can be formed with different combinations of devices with four Spine and six Leaf nodes as a 10-node cluster.

The following configuration creates a leaf spine cluster with two spines and three leafs.

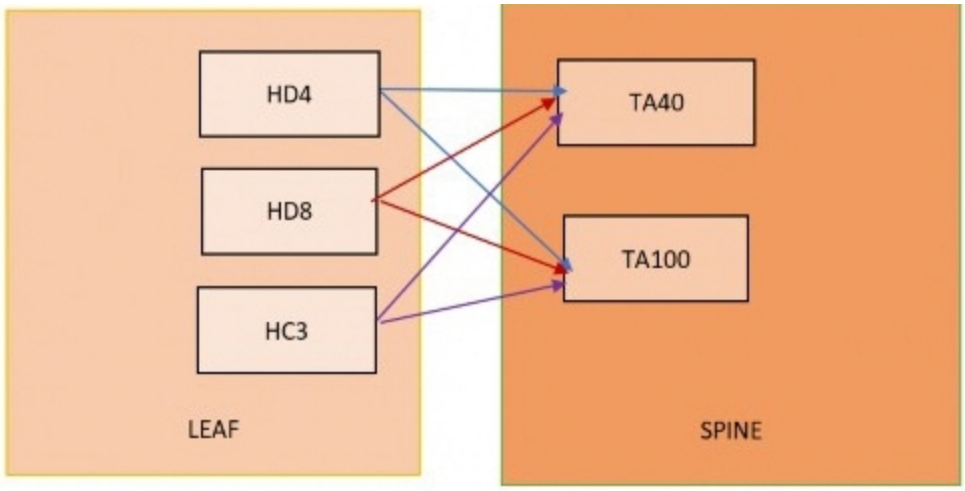


Figure 10 *Leaf spine cluster overview,*

NOTE: GigaStreams support different speeds, as indicated by the different colored connector lines in Figure 10Leaf spine cluster overview, .

Leaf-Spine Cluster Formation Workflow

GigaVUE-FM supports workflow-based configurations for forming a cluster. This workflow walks through the required steps to form a complete Leaf-Spine cluster. Additional procedures for editing and deleting cluster formations are also provided:

- Create a Leaf-Spine Cluster
- Edit a Cluster
- Delete a Node from a Cluster
- How to Change the Leader Preference of a Device

Create a Leaf-Spine Cluster

To create a Leaf-Spine cluster:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**..
2. Click **Create Cluster**.

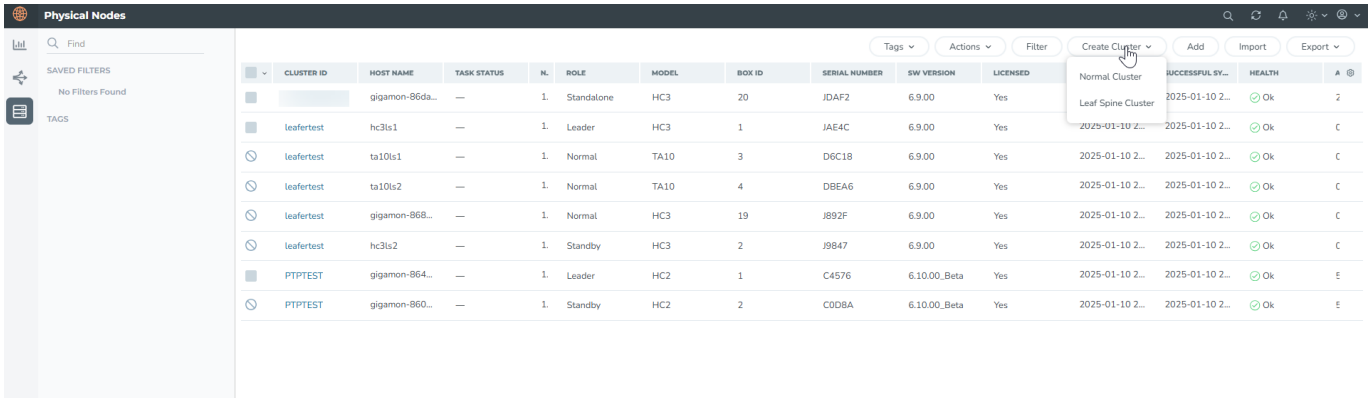
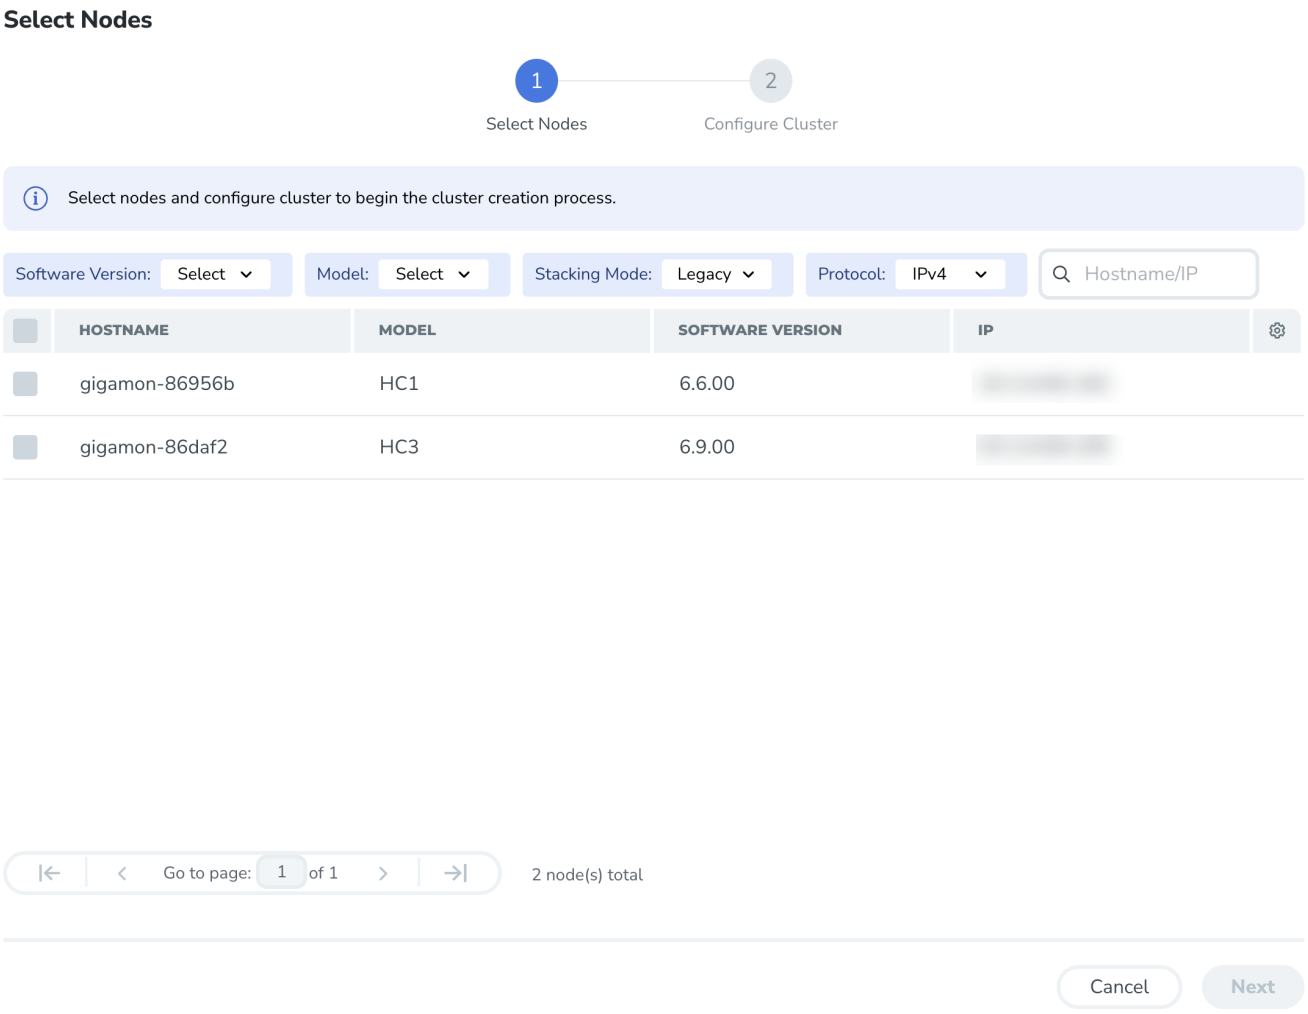


Figure 11 *Create Cluster*

3. From the drop-down list, select **Leaf Spine Cluster**

A **Select Nodes** window appears as shown:



The first step is to select the devices in your cluster.

4. Filter the nodes based on the Software Version, Model, Stacking Mode, Hostname or IP address, and the Protocol through which the nodes communicate with each other.

> **Notes:**
>
> - GigaVUE-TA400 and GigaVUE-HCT devices are not supported in Legacy stacking mode.

5. Select the nodes to include in this cluster.
6. Click **Next**

Cluster Configuration

Use the Cluster Configuration window to configure cluster.

7. Enter a valid **Cluster ID** and Virtual IP (**VIP**) and select the leader in the **Seed Node** list.

NOTE: GigaVUE-TA devices cannot be a leader.

8. After completing the Cluster Configuration details, click **Go to Canvas**.

   NOTE: Use the **Back** button to return to the Select Devices page to revise the selection of devices for this cluster.

Stack Links configuration

Finally, customize the stack links to define how the nodes should be connected.

9. If GDP (Gigamon Discovery Protocol) is enabled at the device chassis level, then the corresponding ports used to create links and ports should be admin enabled. If a physical connection exists in the device, then the links will be shown.
10. If GDP is not enabled in the ports, then the links need to be drawn to connect the devices.

    Connect Leaf devices with Spine devices to create a stack link between them.

    Click the tip of the node and drag your cursor to the second node tip to create a link. After you create the link, a dotted line illustrates the connection.



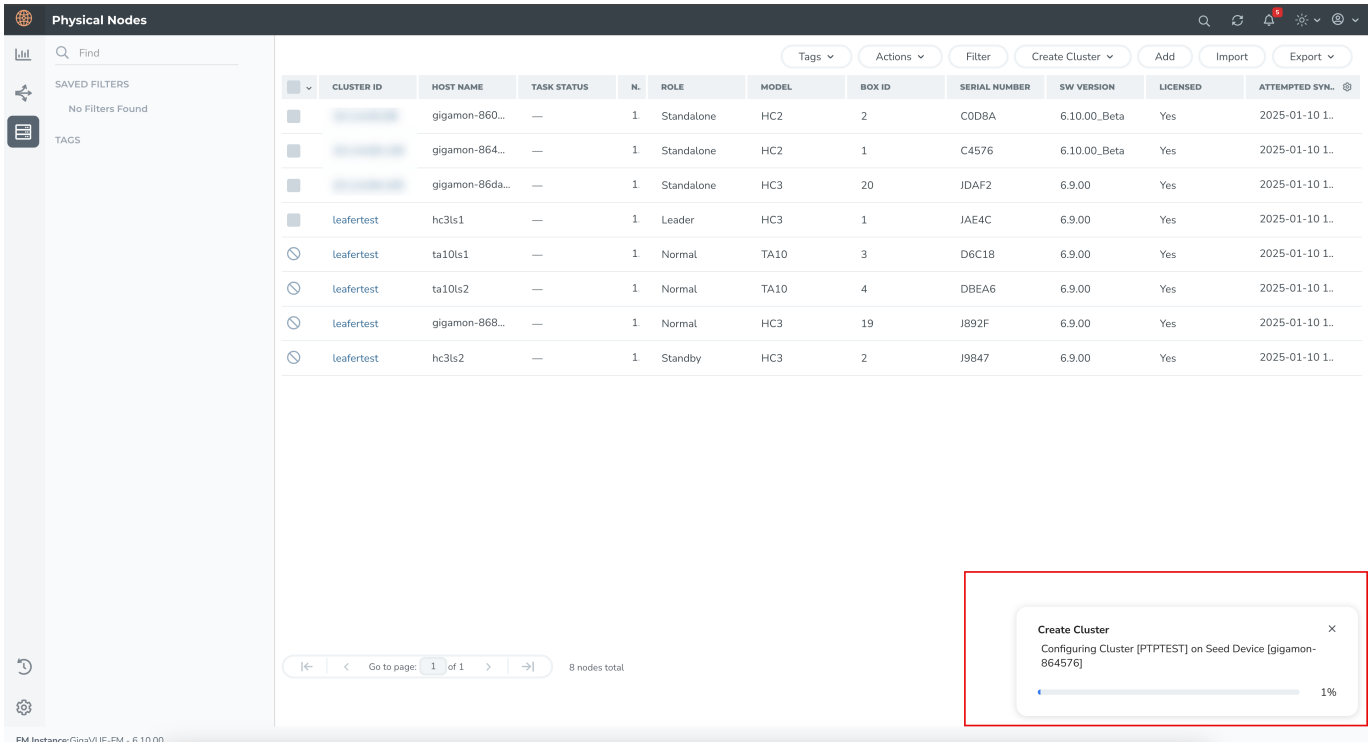
11. Configure the formed links in the Stack Links Configuration pane:
    a. Select a port from each device that are compatible to create a port stack link, for example: x-x ports ,x-q ports, q-c ports, x-c ports.
    b. Select one or more ports from each device to create a stack GigaStream stack link.
    c. Click **Save**.

GigaVUE-FM generates the aliases for each Port stack link and GigaStream stack link as needed.

> **NOTE:** While creating cluster, click **View** stacklinks on the right corner of the canvas to view the details of the stacklinks.. Click **Edit** to edit the details of the devices.

12. Click **Create** to start the cluster creation process.

    The Creating Cluster progress message appears as the cluster is being created.



The Create Cluster progress window in the lower right corner of the page shows the status of every node as it joins the cluster. It takes a few minutes for the cluster to form. The cluster creation process involves the following steps:

- Cluster[clusterName] Creation Successful followed by Seed device
- Verifying Nodes[Will display HostName of all devices]
- Adding Node[HostName] to cluster [clusterName]
- Node[HostName] successfully joined to the cluster.
- Configuring cards for cluster[clusterName]
- Rediscovering cluster[clusterName]
- Configuring ports for cluster[clusterName].
- Configuring ports will display the status of each stack link and GigaStream whether the creation is successful or not.

When the cluster formation process is complete, a notification window confirms the successful creation of the cluster.

13. Click **Go to Cluster** to view the cluster overview.

Edit a Cluster

The Edit cluster option supports only the following operations to the existing cluster:

- Multiple devices can be added to the existing cluster in a single update operation.
- Multiple devices can be added as Leaf, Spine, Leaflet.

- Stack links can be created only from the new device which is added into the cluster wizard.
- Leader preferences can be changed only to device (except) through edit cluster option.
- Stack link alias and GigaStream alias can be edited for newly created links.
- Stacking mode can be changed from Legacy to Default and vice versa.

Notes:

- You cannot change the stacking mode and also add or remove a device simultaneously. Only one task can be performed at a time.

NOTE:

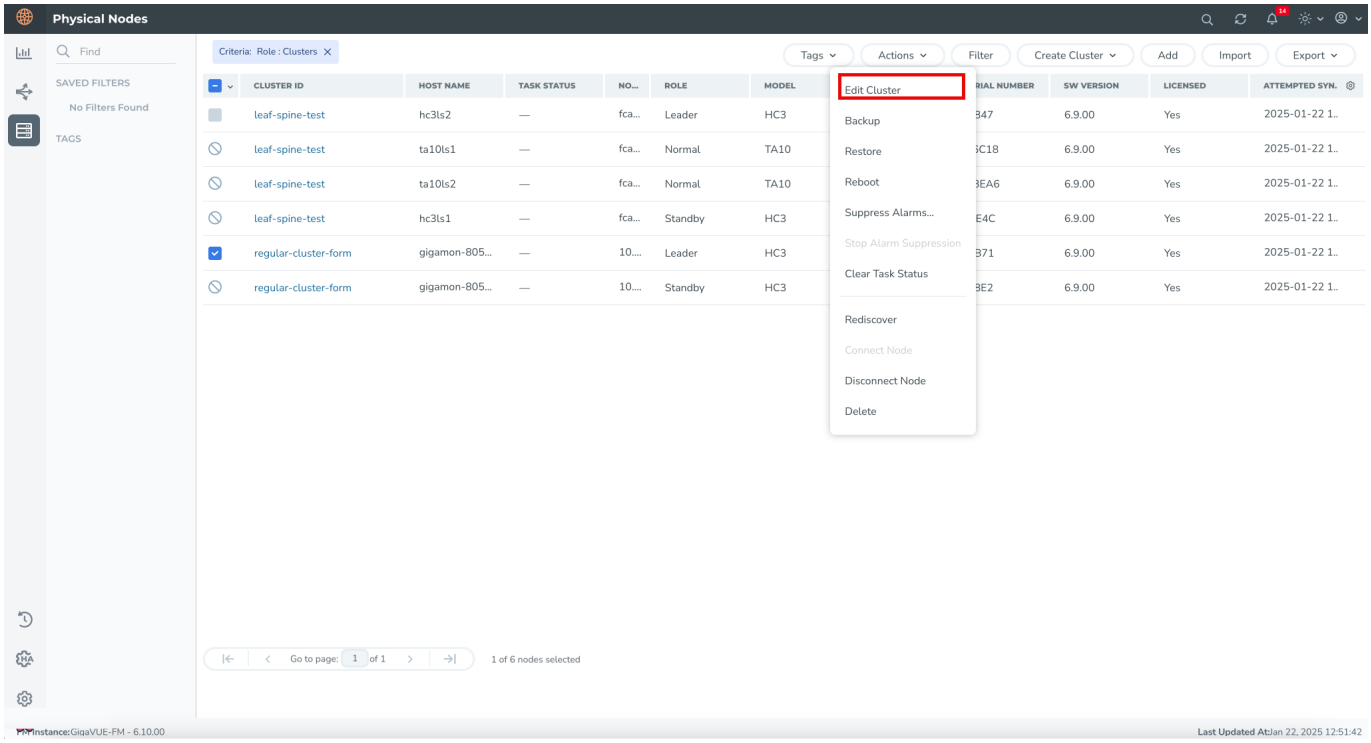
- No option to remove the existing stack links through cluster canvas.
- No option to create links in existing devices.
- Addition and deletion of devices in a single update operation should not be appreciated.
- No option to edit the existing stack link alias and GigaStream alias.

Prerequisites

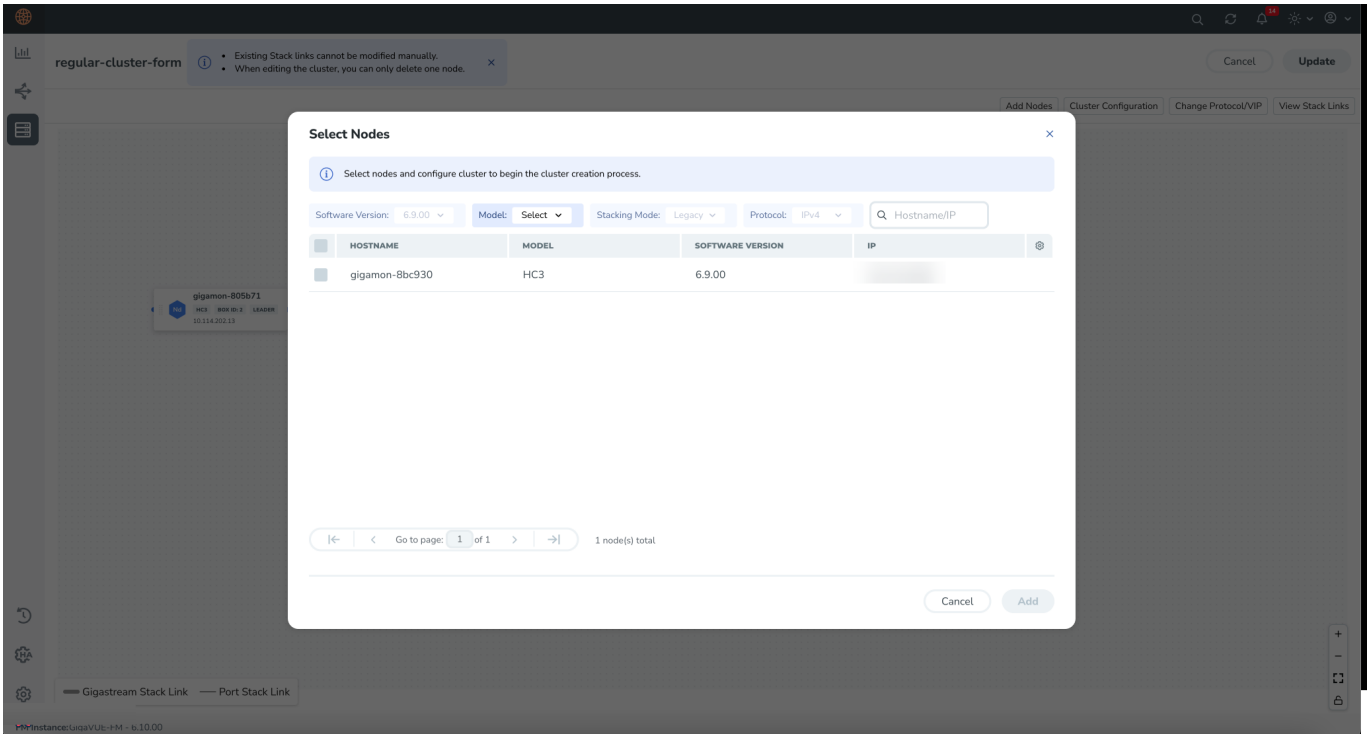
Standalone devices that have maps cannot be added to cluster if ports used in maps overwrites with the selected ports in stack link table.

This workflow describes how to add a node to a existing cluster.

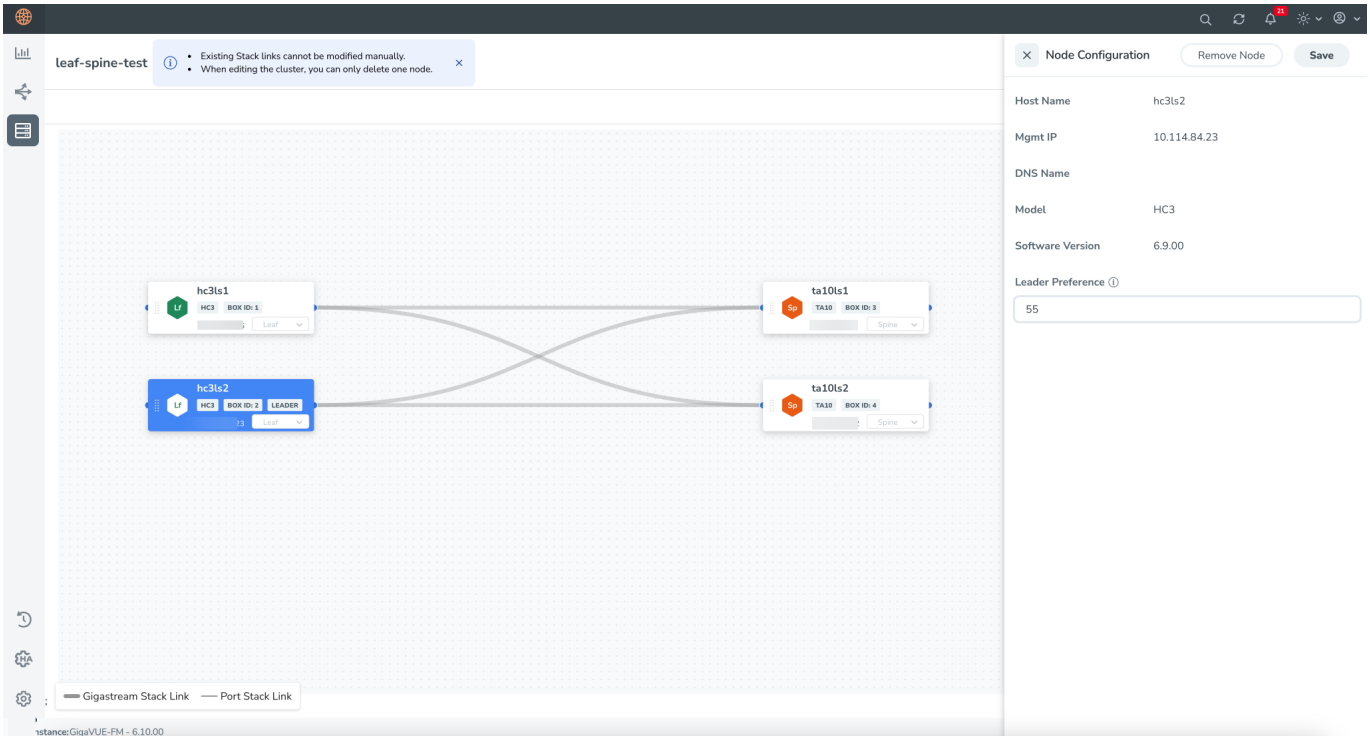
1. Select a cluster and choose **Actions > Edit cluster**.



2. The Edit Leaf Spine Cluster Canvas displays the existing stack link configuration details in the cluster canvas. Click **Add Nodes** to view the standalone devices.
3. Select the required devices from the table in the **Select Nodes** window and click **Add**.



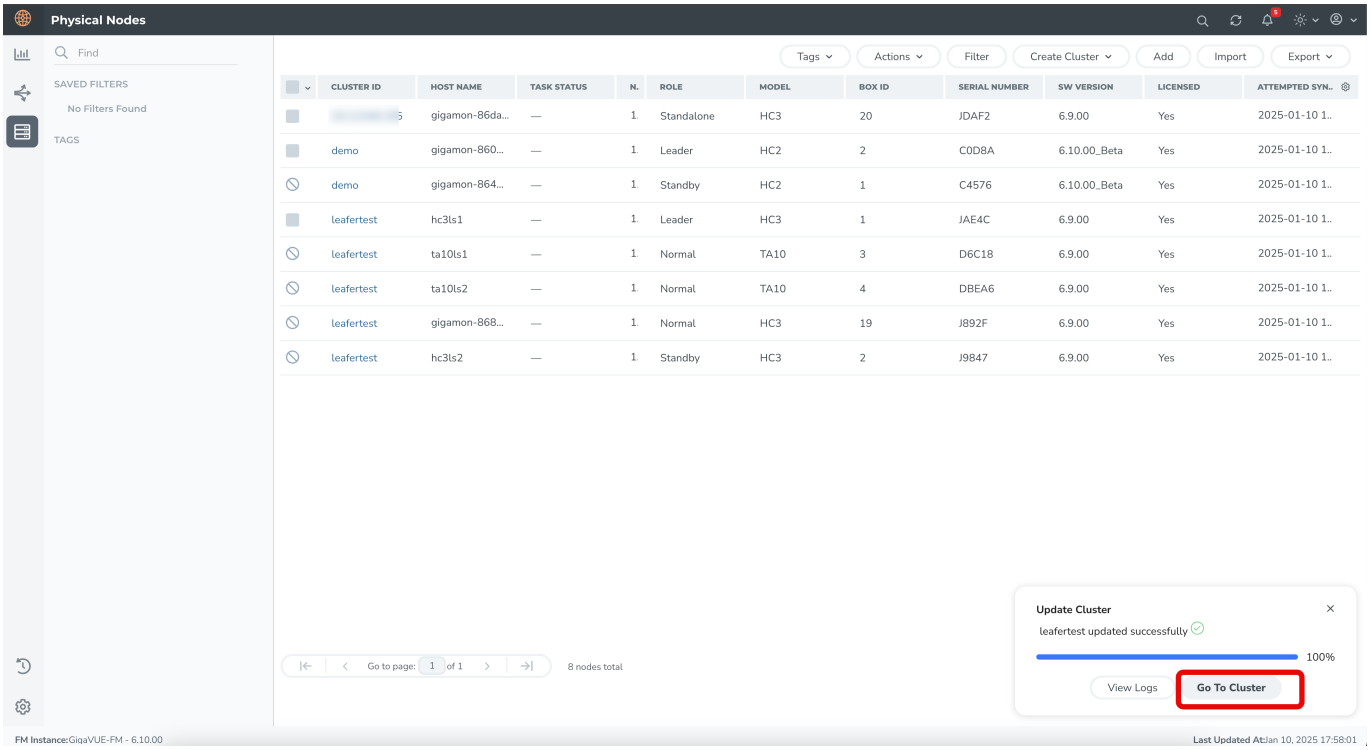
4. In the canvas, draw the links between the newly added device. (**NOTE**: no new link is created for the existing device.)



5. Configure the stack link details in the stack link table.
6. Click **Save**.
7. Change the stacking mode as required.
8. Click **Update** to initiate the update process. A confirmation window appears advising that a backup file is generated with the updated traffic configurations. The configurations saved in the backup file must be applied manually to restore the traffic configurations.
9. Click **OK** to run the cluster update.

   When the cluster update operation starts, a notification window appears at the right corner of the GigaVUE-FM window to show the status progression of each node, card, GigaStream and stack link.

   When the cluster update operation is complete, a notification window confirms the completion of the cluster updates.
10. Click **Go to Cluster** to go to view the cluster overview.



The created GigaStreams appears in the device Port Groups page, and the created stack links appear in the device Stack Links page.

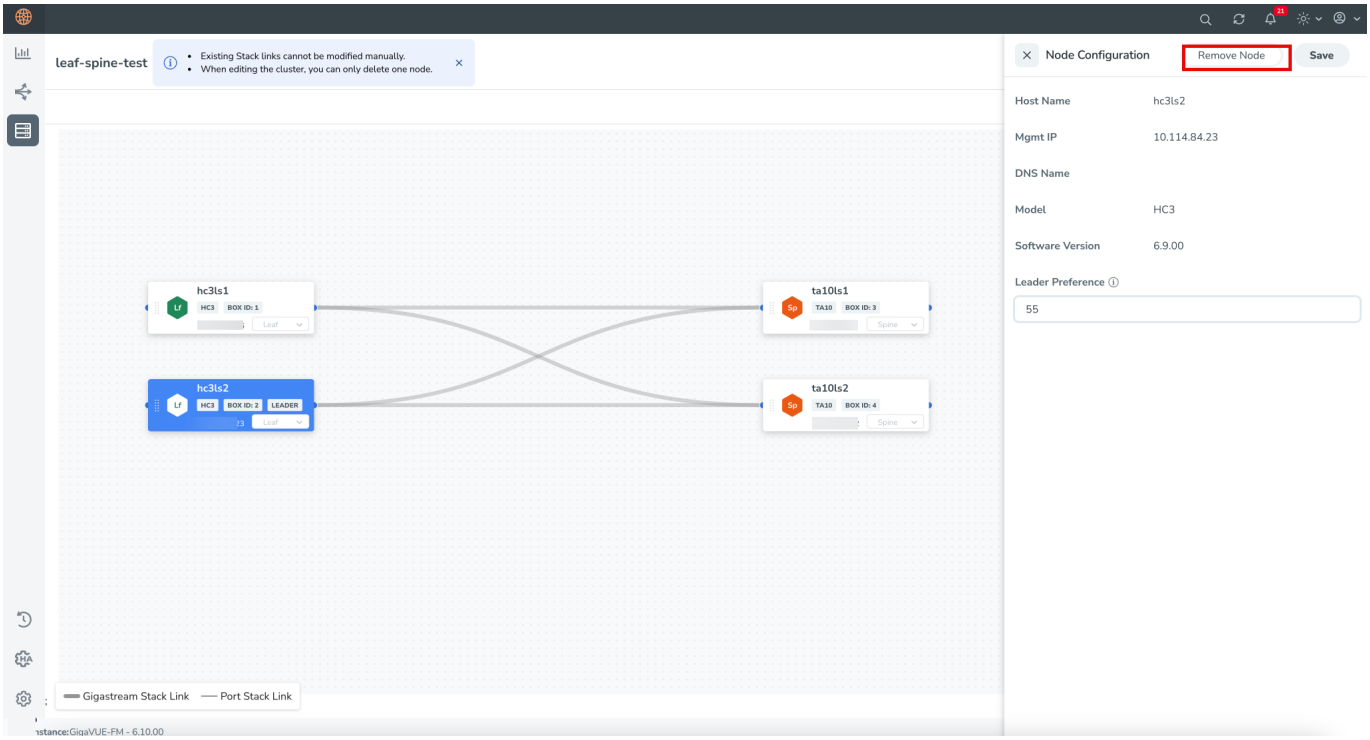
Delete a Node from a Cluster

This workflow describes how to remove a device node from a cluster.

> **NOTE:** Only one device can be removed from the cluster per update operation.

The device should not contain any map configurations in a cluster. Those devices cannot be removed until the maps are present.

1. Select a cluster and choose **Actions > Edit cluster**. Only one device can be deleted from the canvas. It can be either Leaf, Spine or Leaflet.
2. To remove a device, click the device to be removed from the canvas and click **Remove Node**.



The removed device will be deleted from canvas.

3. Click **Update** to initiate the cluster-update operation.

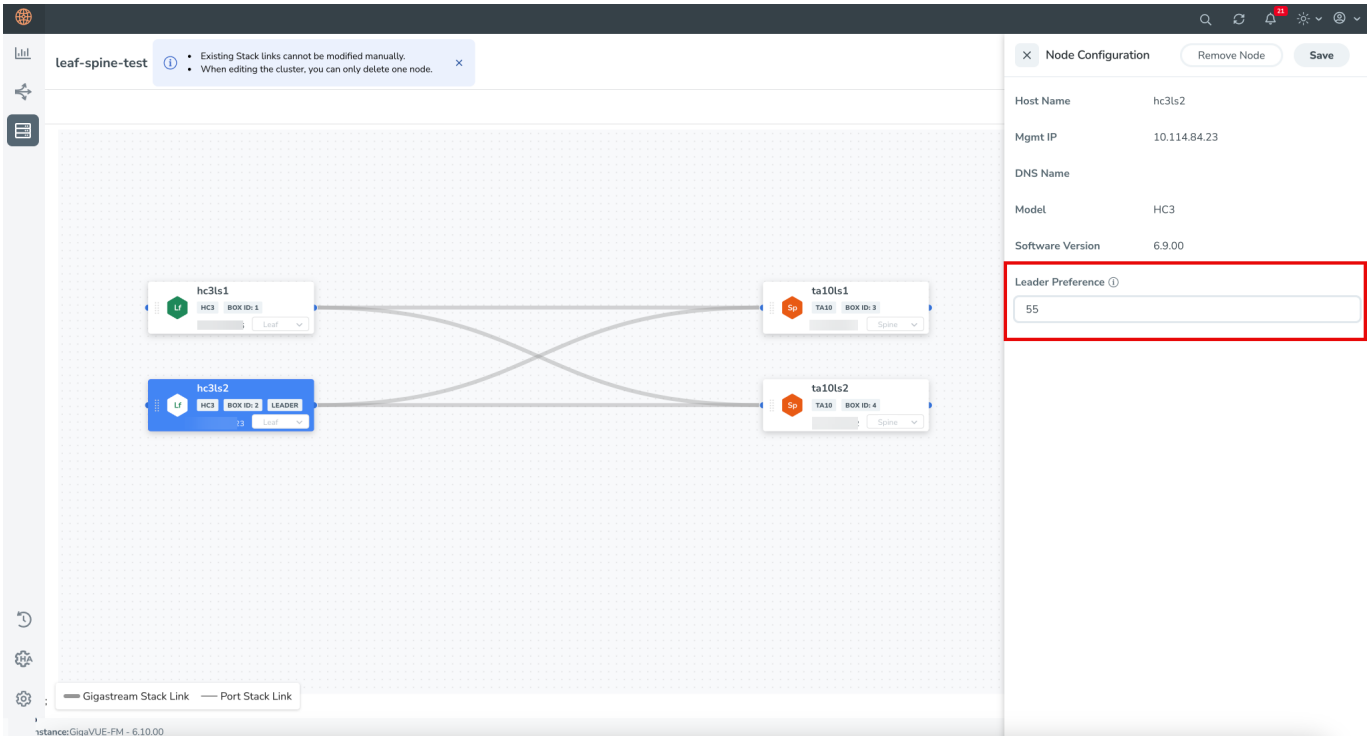
   The Manage Cluster notification window shows the progress of nodes being removed from the cluster. When the device is successfully removed from the cluster, a notification window confirms the successful deletion of the cluster.

4. Click **Go To Cluster** to go into device overview page and see the cluster details.

How to Change the Leader Preference of a Device

This workflow describes how to change the device's leader preference.

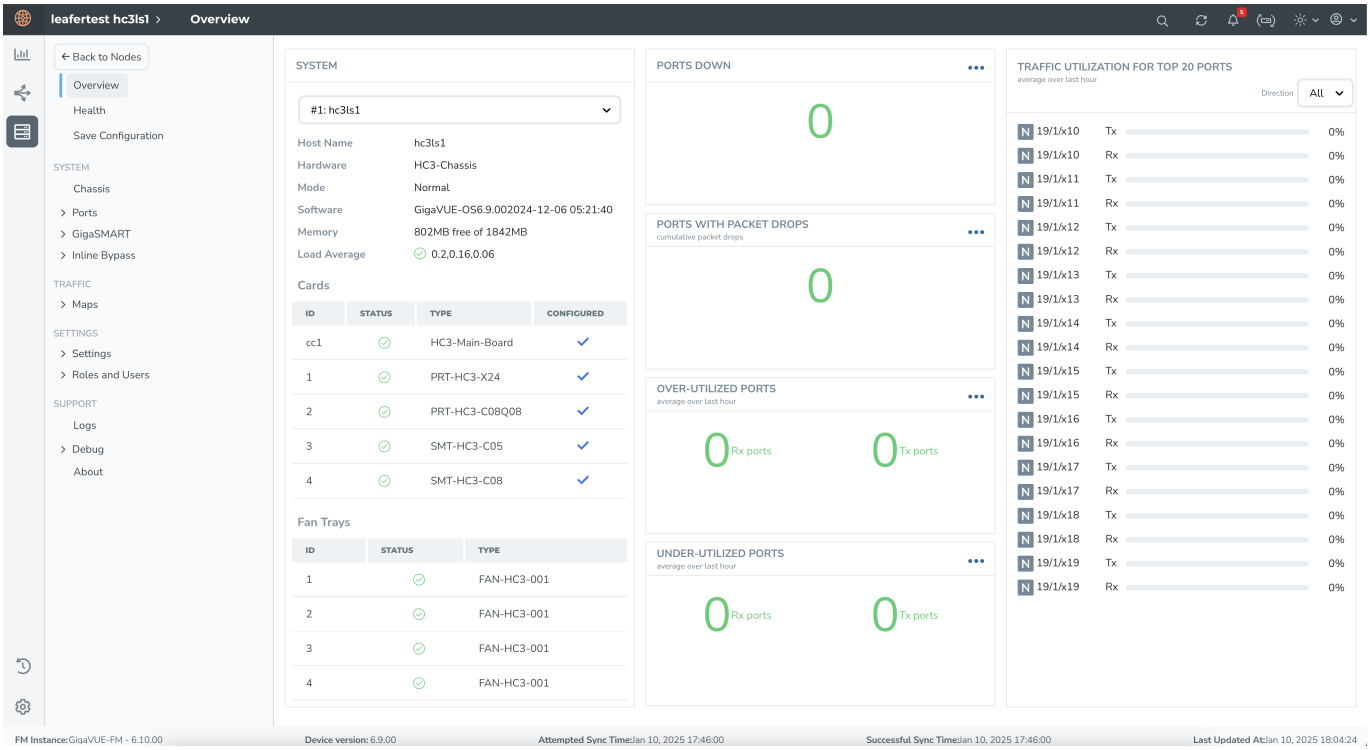
1. Select a cluster and choose **Edit cluster** under Actions.
2. To set the leader preference for a device, click the device and update the **Leader Preference** details in the **Node Configuration** quick view that appears in the side pane.



3. Update the leader preference text box and click **Update** to proceed.

   The Update Cluster notification appears to show the progress of the cluster update. When the process is complete, **Cluster updation completed**, message appears in the notification window.

4. Click **Go To Cluster** to go to the device overview page and see the cluster details.



Spine to Spine and Leaf

This chapter describes the Spine to Spine and Leaf architecture for achieving high availability in a cluster environment. Refer to the following sections for details:

- Introduction to Spine to Spine and Leaf
- Configuration Overview
- Configuration of Spine to Spine and Leaf Architecture
- Leaf-Spine Cluster Deployment

NOTE: Refer to Regular Cluster Formation Workflow for how to use the Regular Cluster workflow.

Introduction to Spine to Spine and Leaf

The Spine to Spine and Leaf architecture is a multi-layer architecture used for network aggregation. This architecture supports leaf nodes and multiple levels of spine nodes. In Spine to Spine and Leaf architecture, connect the leaf and spine as follows:

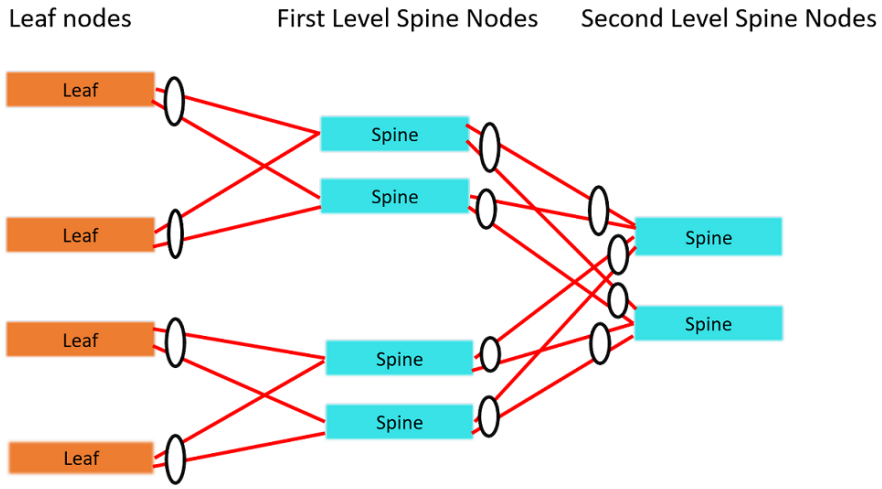
1Leaf nodes to the TAPs or tools.

2First level spine nodes to the leaf nodes and the second level spine nodes.

3Second level spine nodes to all first level spine nodes.

With multiple paths between the nodes in a cluster, the spine to spine and leaf architecture protects against traffic congestion, failures, such as stack link or spine node failures. In the event of a failure, the traffic on one path fails over to the other path. This architecture provides resiliency to the network.

An example of a Spine to Spine and Leaf architecture is shown in the figure.



In a Spine-Leaf cluster, the number of leaf nodes is typically higher than the number of spine nodes. In the figure, there are four leaf nodes, four first level spine nodes, and two second level spine nodes. The leaf nodes aggregate to a fewer number of spine nodes.

For more information on Spine to Leaf architecture, refer to Multi-Path Leaf and Spine.

In the figure, the spine nodes are GigaVUE® TA Series nodes, such as GigaVUE-TA100 or TA200 while the leaf nodes are GigaVUE® HC Series nodes, GigaVUE-HC3 which places the traffic intelligence at the edge.

Traffic between ports on the same leaf node will be local to that leaf node, but traffic between different leaf nodes will go through the spine nodes.

The traffic from a source leaf node to a destination leaf node flows as follows:

- From a TAP, traffic flows to the source leaf node.
- From the source leaf node, traffic is load balanced to the connected spine nodes.
- From the spine node, depending on the configuration, traffic flows to the next level of spines or the destination leaf node.
- From the destination leaf node, traffic flows to the tool ports.

Resiliency is achieved when there are multiple paths from the network to the tools across GigaVUE nodes.

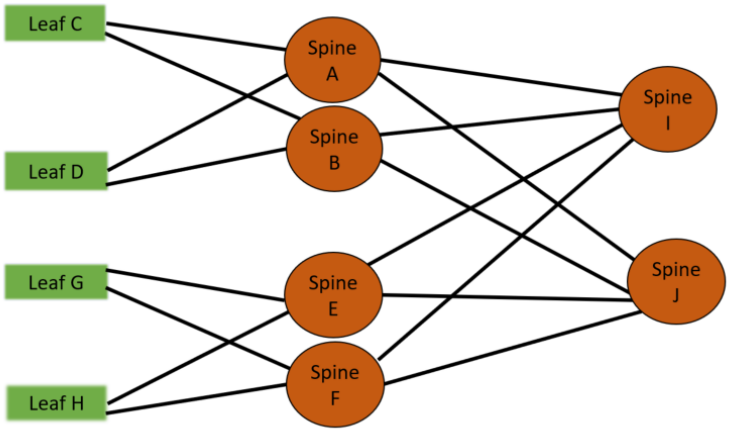
Refer to Path Protectionfor the leaf node failure, stack link failure on a leaf node or spine node.

Configuration Overview

This section provides an overview of the configuration. You must perform the configuration from the leader in the cluster. Follow this configuration sequence to prevent loops:

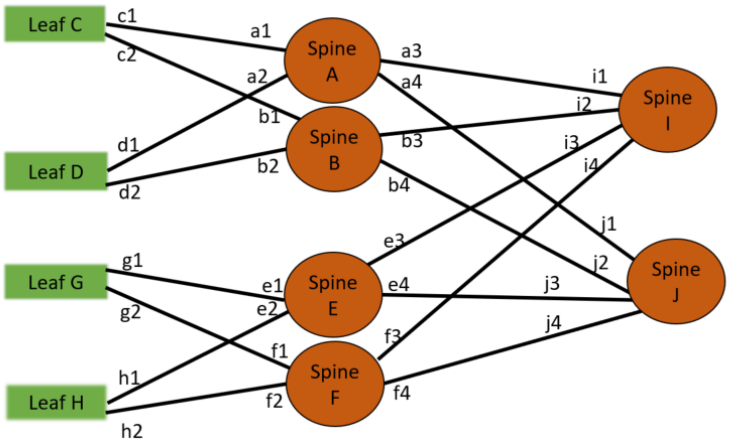
1. Configure Stack GigaStream
2. Configure Spine Links
3. Configure Stack Links

This configuration connects nodes using multiple paths. Refer to the following figure for an example of the configuration.



Configure Stack GigaStream

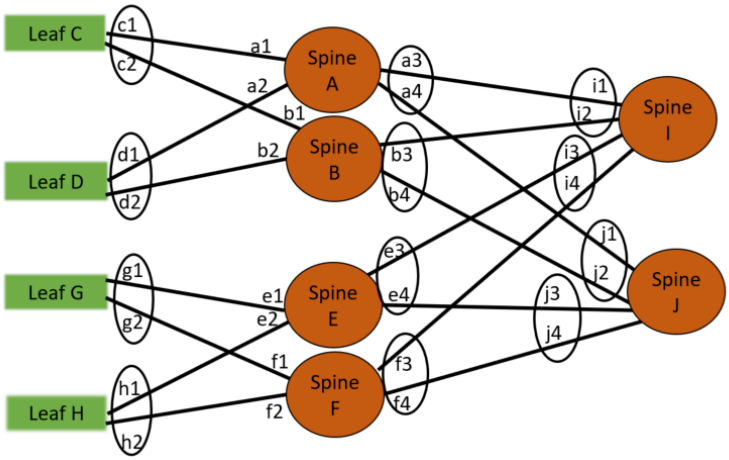
The stack GigaStream connect the spine and leaf nodes. In the following figure, the stack GigaStream are: a1, a2, a3, a4, b1, b2, b3, b4, c1, c2, d1, d2, e1, e2, e3, e4, f1, f2, f3, f4, g1, g2, h1, h2, i1, i2, i3, i4, j1, j2, j3, j4. Even if there is only one port that connects the nodes, you must still configure a stack GigaStream. With a configuration of 6 spine nodes and 4 leaf nodes, the number of stack GigaStream is 32.



Configure Spine Links

On each leaf node, there is one spine link that contains the list of GigaStream connecting the leaf nodes to the spine nodes. The spine links contain multiple stack GigaStream bundled together. The spine links on the leaf nodes are: {c1,c2}, {d1,d2}, {g1, g2} and {h1, h2}.

The spine node also contains the spine links. A spine node connecting to another spine nodes has a spine link. The spine links on the spine nodes are: {a3, a4}, {b3, b4}, {e3, e4}, {f3, f4}, {i1, i2}, {i3, i4}, {j1, j2}, and {j3, j4}. The total number of spine links is twelve for this configuration. Across the spine link members, traffic is load balanced. For this part of the configuration, refer to the circles in the figure.



NOTE: You must not configure any spine links from spine nodes to leaf nodes. For example, in the figure, a1 and a2 should not be configured as spine links.

Configure Stack Links

The stack links are: {a1,c1}, {a2,d1}, {a3,i1}, {a4,j1}, {b1,c2}, {b2, d2}, {b3, i2}, {b4, j2}, {e1, g1}, {e2, h1}, {e3, i3}, {e4, j3}, {f1, g2}, {f2, h2}, {f3, i4}, {f4, j4}. The total number of stack links is sixteen for this configuration. For this part of the configuration, refer to the circles in the above figure.Configuration Overview.

These configuration steps ensure that the spine and leaf nodes are fully meshed.

Notes and Considerations

For the limitations and considerations, refer the Notes and Considerations in Page 8.

Configuration of Spine to Spine and Leaf Architecture

To configure Spine to Spine and Leaf Architecture, follow these steps:

1. Configure Leaf-Spine Cluster formation workflow. To configure, refer Leaf-Spine Cluster Formation Workflow.
2. Execute the Gigamon Automation GigaVUE-FM SDK to establish the connection between the first level of spines and the second level of spines.
3. View the renewed topology in GigaVUE-FM, and reorder them if required.
4. Verify the configurations performed.

You can view the updated topology in GigaVUE-FM only after the execution of the Gigamon Automation GigaVUE-FM SDK to establish the connection between the first level and second level spines.

Limitations

- It is not recommended to have network and tool in one segment and GSOP in other segment. You must follow any one of the below arrangements:
  - Network and GSOP in one segment of the spine, and the tool in other segment of the spine.
  - Network in one segment of the spine, and GSOP and tool in other segment of the spine.

- If the network port is in first level spine of segment 1, then the leaf should be in the leaf of segment 2.
- If the network port is in first level spine of segment 2, then the leaf should be in the leaf of segment 1.

Leaf-Spine Cluster Deployment

This section describes the steps and prerequisites to deploy a leaf-spine cluster.

Refer to Introduction to Spine to Spine and Leaf for a conceptual overview of the leaf-spine architecture.

Deployment Checklist

Before forming a Leaf-Spine Cluster, it is strongly recommended that you get familiar with the relevant documentation and review the deployment checklist to prepare for deployment.

Documentation

- Review the **GigaVUE-FM Release Notes** to familiarize yourself with the functionality around creating and managing clusters.
- Review the “Multi-Path Leaf and Spine” chapter to familiarize yourself with leaf and spine architecture.
- Review the **GigaVUE-FM Release Notes** for any known issue that may impact your use case.

Pre-deployment checklist

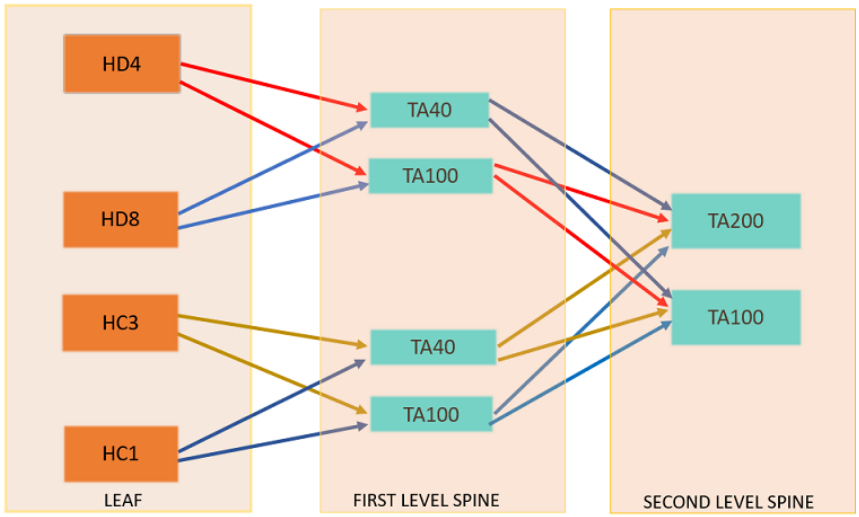
- Gigamon Fabric Management must be upgraded to Gigamon 5.4.01 or later.
- Gigamon device must be upgraded to GigaVUE-OS 5.4.00 or later.
- Advanced Features License must be installed in TA devices.
- Physical connection must be established to create stack links.
- Devices must have GDP enabled and be physically connected to create links among devices from Gigamon.

IMPORTANT: Recommendation is to use TA devices as SPINE Nodes and other devices as LEAF Nodes.

Formation Scenario

The Spine to Spine and Leaf cluster is formed with different combinations of devices with Spine and Leaf nodes as a node cluster.

The following configuration creates a leaf-spine cluster with four leaves, four first level spines, and two second level spines.



NOTE: GigaStreams support different speeds, as indicated by the different colored connector lines in the figure.

Fabric Statistics

This chapter provides the statistics of all types of ports available in GigaVUE H series and TA series.

Refer to the following sections for details:

- About Fabric Statistics
- Display Fabric Statistics for All Ports
- Display Fabric Statistics for a Single Port
- Export Fabric Statistics
- Filter Fabric Statistics
- Clear Fabric Statistics Counters

About Fabric Statistics

GigaVUE-FM provides the ability to view detailed information about the packets dropped, packets discarded, and packets received and transmitted by the following ports:

- Network ports
- Tool ports
- Hybrid ports
- Circuit ports
- Inline-network ports
- Inline-tool ports
- GigaSMART engine ports
- Backplane ports
- XAUII ports
- Stack ports

Network ports, tool ports, hybrid ports, circuit ports, inline-network ports, and inline-tool ports are also called front panel ports. These are the ports that are visible on the front view of the GigaVUE node.

The backplane ports on the control card are connected to the backplane ports on the line card or GigaSMART engine ports. They allow packets to move from card to card, for example, from line card to control card and control card to GigaSMART card. If there are packet errors in the backplane ports, this information can be viewed in the Fabric Statistics page. The format used to represent a backplane port is <box ID>/<slot ID>/<port ID>. For example, 1/1/s1 where s1 refers to port s1.

GigaVUE-FM also collects statistics of XAUII links on the GigaSMART engine ports. Each GigaSMART engine port is made up of 2 to 4 XAUII links, depending upon the platform. They are considered as the child ports of GigaSMART engine ports. The format used to represent a XAUII port is <box ID>/<slot ID>/<port ID>. For example, 1/5/e1x0 where e1x0 refers to the first XAUII port on engine 1.

Display Fabric Statistics for All Ports

Using the Fabric Statistics page, you can view the statistics associated with the port types. If there are packet drops in the ports, click on the port and drill down further to investigate the cause of the packet drops.

To view the Fabric Statistics page:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**. and click on a GigaVUE node.
2. Select **Ports > Ports > Fabric Statistics**.

Table 3: Fabric Statistics Definitions describes the columns in the Fabric Statistics table.

Table 3: Fabric Statistics Definitions

Counter	Definition	Notes
Port ID	Port ID in the format <box id/slot id/port id>	
Alias	Alias of the port	
Unicast Packets Rx	Total unicast packets received	Excludes multicast packets, broadcast packets, packets with FCS/CRC errors, MTU exceeded errors, oversize packets, and pause packets.
Unicast Packets Tx	Total unicast packets transmitted	Excludes multicast packets, broadcast packets, packets with FCS/CRC errors, MTU exceeded errors, oversize packets, and pause packets.
Packets/sec Rx	Total packets received per second Excludes multicast packets, broadcast packets, packets with FCS/CRC errors, MTU exceeded errors, oversize packets, and pause packets.	Excludes packets with FCS/CRC errors.
Packets/sec Tx	Total packets transmitted per second	Excludes multicast packets, broadcast packets, packets with FCS/CRC errors, MTU exceeded errors, oversize packets, and pause packets.
Octets/sec Rx	Total bytes received per second This indicates the rate of incoming bytes in the last second.	
Octets/sec Tx	Total bytes transmitted per second This indicates the rate of incoming bytes in the last second.	
Bits/sec	Total bits transmitted per second.	
Utilization Rx	Percentage of port utilization by packets received	
Utilization Tx	Percentage of port utilization by packets transmitted	
Link Speed	Maximum link speed of the port	
Type	Type of port	
Octets Rx	Total bytes received Includes all valid and error frames with the exceptions noted in the adjacent columns.	Excludes undersize frames.

Counter	Definition	Notes
Tx drop count	Total packets dropped in the Tx packets.	
Tx drop percentage	Total packets dropped in the Tx packets in percentage.	
Error count	The count of error packet received.	
Error percentage	The percentage of error packets received.	
Octets Tx	Total bytes transmitted Includes all valid and error frames with the exceptions noted in the adjacent columns.	Excludes undersize frames.
Non-unicast Packets Rx	Total Broadcast and Multicast packets received	
Non-unicast Packets Tx	Total Broadcast and Multicast packets transmitted	
Rx Drops count	Total received packets dropped	Packets are dropped when a tool port's bandwidth is exceeded due to oversubscription. Packets are dropped when they reach the tool port but before they are sent out.
Tx Drops count	Total transmitted packets dropped	Packets are dropped when a tool port's bandwidth is exceeded due to oversubscription. Packets are dropped when they reach the tool port but before they are sent out.
Discards Rx	Total received packets discarded	This counter increments when a packet is discarded at a tool port due to a tool port map rule.
Discards Tx	Total transmitted packets discarded	This counter increments when a packet is discarded at a tool port due to a tool port map rule.
Error Rx	Total error packets received Error packets include undersize, FCS/CRC, MTU exceeded, fragments, and oversize packets.All packets that list under this counter are discarded and not processed further.	Excludes oversize packets without FCS/CRC. Packets larger than the MTU setting arriving on a network port are counted twice in the counter. So 1000 oversize packets would show up as 2000. This double-counting only happens with Oversize error packets.
Error Tx	Total error packets transmitted Error packets include undersize, FCS/CRC, MTU exceeded, fragments, and oversize packets.All packets that list under this counter are discarded and not processed further.	Excludes oversize packets without FCS/CRC. Packets larger than the MTU setting arriving on a network port are counted twice in the counter. So 1000 oversize packets would show up as 2000. This double-counting only happens with Oversize error packets.

The port types in the Fabric Statistics table are represented as follows:

Table 4: Port Types

Port	Port Type
Network Port	N
Tool Port	T
Hybrid Port	H
Circuit Port	C
Inline Network Port	iN
Inline Tool Port	iT
GigaSMART Engine Port	E (bigger size)
Backplane Port	BP
XAUII Port	E (smaller size)
Stack Port	S

Display Fabric Statistics for a Single Port

To view the Fabric Statistics of a single port:

1. On the Fabric Statistics page, select any port (except for backplane and XAUII).
2. Click on the Port ID field. The port quick view is displayed.

The port quick view provides information about a specific port. Each field is color-coded in the graphical representation.

Port Quick View

The port quick view provides a graphical view of the port statistics. You can choose to view the statistics based on an hour, a day, a week, a month, or live to see the real-time data.The quick view also provides information about port properties, statistics about the traffic received (Rx) and transmitted (Tx), and alarms information. Click Counters drop-down list to select other options.

The following table describes the parameters displayed in the Port quick view:

Field	Description
Port Info	Displays the port information such as the name of the port, if any.
Parameters	Displays information about the current link status of the port, admin status, port type, port speed, port's duplex configuration, auto negotiation configuration, and the force link up configuration.
Transceiver	Displays the transceiver type connected to the port, optical input power, vendor from where the transceiver is purchased, product number, serial number, and temperature.
Filters	Displays the number of filters, pass rules, and drop rules applied to the port.
Port Discovery	Displays the discovery protocol information (GDP, CDP, LLDP) and neighbor device information.
Alarms	Displays the buffer threshold percentage on the ports in the Rx and Tx directions, high and low utilization threshold percentage, and the port utilization in the Rx and Tx directions.
Permissions	Displays the roles and permissions associated with the port, users with permission to lock the port, and users to whom the port is shared.
Related Maps	Displays the related maps associated with the port.
Tags	Displays the port level tag which is configured by GigaVUE-FM user.

Export Fabric Statistics

To export the fabric statistics, click any one of the options as requied:

- **Export All** - Exports statistics of all the maps.
- **Export Selected** -Exports statistics of selected maps.

The Statistics table is downloaded in either CSV or XLSX format.

Filter Fabric Statistics

Using filters, you can search and narrow down the ports displayed in the Fabric Statistics page. To filter the fabric statistics, click **Filter**. A filter quick view appears as shown in Figure 12Fabric Statistics Filter. Specify the criteria to filter the ports.

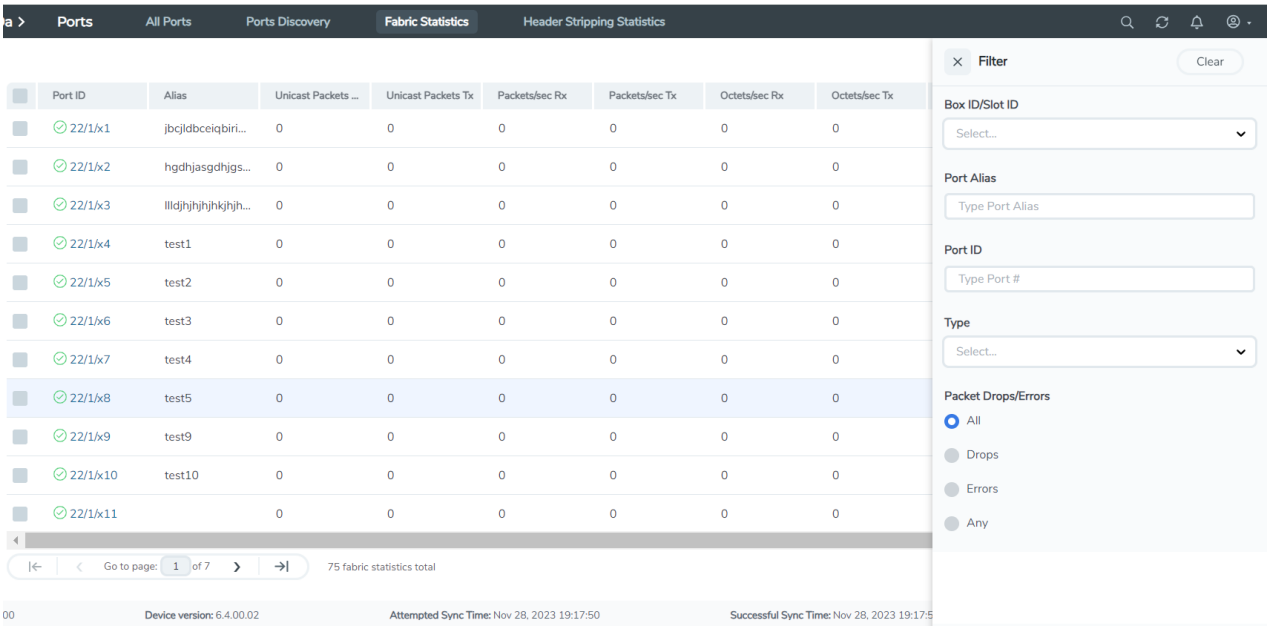


Figure 12 *Fabric Statistics Filter*

The criteria that you can use to filter the fabric statistics is as follows:

Criteria	Description
Box/Slot ID	Displays only those ports that match the specified box and slot IDs.
Port Alias	Displays a port with specified alias.
Port ID	Displays ports with specified port ID. For example, if you specify 3 as the port ID, the result will display ports that include the number 3, 13, 23, 30, and so on.
Type	Displays ports with specified port type. Select one of the following: ▪ Network ▪ Tool ▪ Inline Network ▪ Circuit port ▪ Inline Tool ▪ GigaSMART ▪ Hybrid ▪ Stack ▪ Backplane ▪ XAUII

Criteria	Description
Packet Drops/Errors	Displays only those ports with packet drops or errors or both. Select one of the following: ▪ All — displays ports with both packet drops and errors. This is the default. ▪ Drops — displays ports with only packet drops. ▪ Errors — displays ports with only packet errors. ▪ Any — displays ports with either packet drops or packet errors.

Clear Fabric Statistics Counters

To clear the fabric statistics counters:

1. Select the required port ids for which you need to clear the counters.
2. Click **Clear**. A confirmation message pops-up asking if the statistical counters must be cleared for the selected port ids.
3. Click **Ok**.

Enable Discovery Protocols

You must enable the discovery protocols at the chassis level and at the port level for the nodes and clusters to be discovered.

- Refer to Enable Gigamon Discovery on Chassis for enabling Gigamon Discovery Protocol at the chassis level.
- Refer to Enable LLDP and CDP on Chassis for enabling other discovery protocols at the chassis level.
- Refer to Enable LLDP, CDP, and Gigamon Discovery on Ports for enabling the discovery protocols at the port level.

NOTE: Use the Port Discovery page to bulk configure the discovery protocols (GDP, CDP and LLDP) on all the ports across the devices. Refer to Port Discovery section for more details.

Enable Gigamon Discovery on Chassis

Gigamon discovery is disabled on chassis and ports. The Gigamon discovery packets are transmitted only when Gigamon Discovery is enabled on chassis as well ports.

To enable Gigamon discovery on chassis:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**..
2. Click on a node on which you want to en able Gigamon Discovery. The Overview page of the node is displayed.
3. In the left navigation pane, click **Chassis**.
4. Switch the Chassis view to List View.
5. Under Properties, select the Box ID for which you want to enable Gigamon Discovery.
6. Click **Actions** and select **Enable Gigamon Discovery**.

NOTE: Refer to the Enable LLDP, CDP, and Gigamon Discovery on Ports section for details on how to enable the discovery protocols on ports.
Use the Port Discovery page to bulk configure the discovery protocols (GDP, CDP and LLDP) on all the ports across the devices. Refer to Port Discovery section for more details.

Enable LLDP and CDP on Chassis

You can choose to enable LLDP and CDP on the management interface of a GigaVUE HC Series or GigaVUE TA Series device. This will allow the device to discover neighboring devices in a network.

To enable LLDP and CDP on the management interface of a device:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**..
2. Click on a node on which you want to enable LLDP and CDP. The Overview page of the node appears.
3. In the left navigation pane, select **Settings > Interface > Protocol Configuration**.
4. Select the box ID on which you want to enable LLDP and CDP, and then click **Edit**.
5. Select the **Discovery Protocols** check box, and then select one of the following: **LLDP**,

or **CDP**.

6. Click **OK**. The selected discovery protocol is enabled on the management interface of the device.

NOTE: Refer to the Enable LLDP, CDP, and Gigamon Discovery on Ports section for details on how to enable the discovery protocols on ports.
Use the Port Discovery page to bulk configure the discovery protocols (GDP, CDP and LLDP) on all the ports across the devices. Refer to Port Discovery section for more details.

Enable LLDP, CDP, and Gigamon Discovery on Ports

To enable LLDP, CDP, and Gigamon Discovery on ports:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**..
2. Click on a node on which you want to enable LLDP and CDP. The Overview page of the node appears.
3. In the left navigation pane, select **Ports > All Ports**.
4. On the Ports page, select the Port ID on which you want to enable LLDP, CDP, and Gigamon Discovery. Click **Edit**.
5. Under **Device Discovery**, enable the checkbox for the required discovery protocols. Enable all three checkboxes to enable all protocols.
6. Click **OK**.
7. Click **Save**.
8. Click the **Device Discovery** tab. For each port on which the ports discovery is enabled, the neighbor information is displayed.

Refer to:

- Enable Gigamon Discovery on Chassis for enabling Gigamon Discovery protocol at the chassis level.
- Enable LLDP and CDP on Chassis for enabling other discovery protocols at the chassis level.
- Port Discovery for bulk configuring the discovery protocols on all the ports across the devices in GigaVUE-FM.

Limitations of Discovery

The following limitations apply to Gigamon Discovery:

- Gigamon discovery is not supported on inline-tool and inline-network port type.
- For stack ports, Gigamon discovery will work only if both ends of the link are of type stack.
- If a pass-all map is configured for Gigamon discovery enabled ports, the show map stats command will show an increment in the map rule counter for Gigamon discovery packets although the Gigamon discovery packets are not sent to the tool ports.
- If Gigamon discovery is enabled on a port, for example port 9/3/x1, and the port type is changed to an unsupported port types such as inline-network, Gigamon discovery is immediately disabled, and the following message is displayed:

  (config) # ! GDP is not supported with port type inline-network. Disabling gdp on port 9/3/x1

  Similarly, if the port type is changed to inline-tool, the following message is displayed:

  (config) # ! GDP is not supported with port type inline-tool. Disabling gdp on port 9/3/x1
- The hybrid ports, tool-mirror source ports, and ports configured with force link up are operated in loopback mode. So, when Gigamon discovery is enabled, these port show themselves as neighbors.

The following limitation applies to CDP and LLDP Discovery:

- You cannot change the port type if CDP and LLDP Discovery are enabled. To change the port type, you must first disable CDP and LLDP Discovery on the port.

  (config) # ! Cannot change port type with discovery configured. Disable discovery on port 1/1/x1 before changing type.

Topology Visualization

The Topology page provides enhanced visualization of the nodes, thereby improving the usability of the page. The nodes are grouped based on pre-defined tags. The Topology page provides the following views:

- **Tagged Layout Topology:** The topology layout is based on predefined tags created by the user. Refer to Tagged Layout Topology section. The nodes can be grouped based on:
  - **Nature of deployment:** Location of nodes such as region, state, city, data center.
  - **Placement:** Aggregation nodes, leaf nodes, spine nodes.
- **Smart Layout Topology:** For smaller deployments, GigaVUE-FM creates predefined tags to group the nodes. Refer to the Smart Layout Topology section.

You can use the device level tagging to associate tags to the devices individually, as against the cluster-level tagging option. Refer to Device Level Tagging section for more details.

NOTE: Classic Topology page is deprecated from software version 6.2.00.

Supported Devices

Topology Visualization is supported in the following devices:

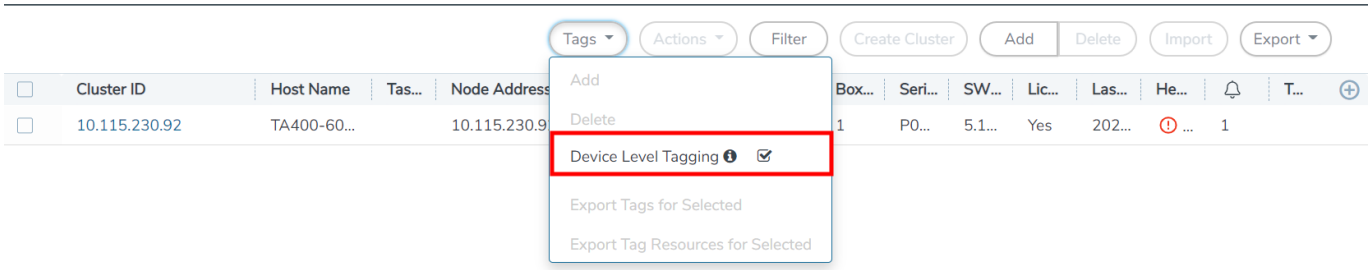
- GigaVUE-HC1 and GigaVUE-HC3
- GigaVUE TA Series devices
- Tool Devices discovered through CDP/LLDP and manual Tool Devices

Device Level Tagging

Each of the individual nodes in a cluster can be associated with a tag key and a tag value. To assign a tag key to a device:

1. On the left navigation pane, go to **Inventory** go to **Physical** > **Nodes**.
2. Click on **Tags** and enable the **Device Level Tagging** option.
3. Select the required node(s). Click **Tags** and select **Add**. The **Add Tags to Resources** page appears.
4. Select the required tag key and the tag value to which the nodes must be tagged.
5. Click **OK**.

NOTE: Only aggregation tags can be configured to the devices.



The nodes in a cluster are now associated to tags which can be used for creating the required topology.

Smart Layout Topology

The **Topology** page displays smart layout topology for smaller deployments, in which the number of devices is less than or equal to fifty. GigaVUE-FM constructs the smart layout topology based on the port configuration. GigaVUE-FM applies the following types of tags to the devices:

- **NETWORK DEVICES**: External network devices are grouped under NETWORK DEVICES. By default, this is disabled. You must manually enable it to view the network nodes.
- **GTAP NODES**: Devices with internal tag value of GTAP_Nodes are grouped under GTAP NODES.

  NOTE: GTAP NODES column is displayed only if GTAP is configured in GigaVUE-FM.
- **TAP NODES:** Devices with maximum number of network ports are grouped under TAP NODES.
- **CORE NODES**: Devices with maximum number of stack and circuit ports are grouped under CORE NODES.
- **TOOL NODES**: Devices with maximum number of tool ports are grouped under TOOL NODES.
- **EXTERNAL TOOLS**: External tools connected to the devices are grouped under EXTERNAL TOOLS.
- **LEAF NODES**: Leaf nodes in a leaf-spine cluster are grouped under LEAF NODES.
- **SPINE NODES**: Spine nodes in a leaf-spine cluster are grouped under SPINE NODES.

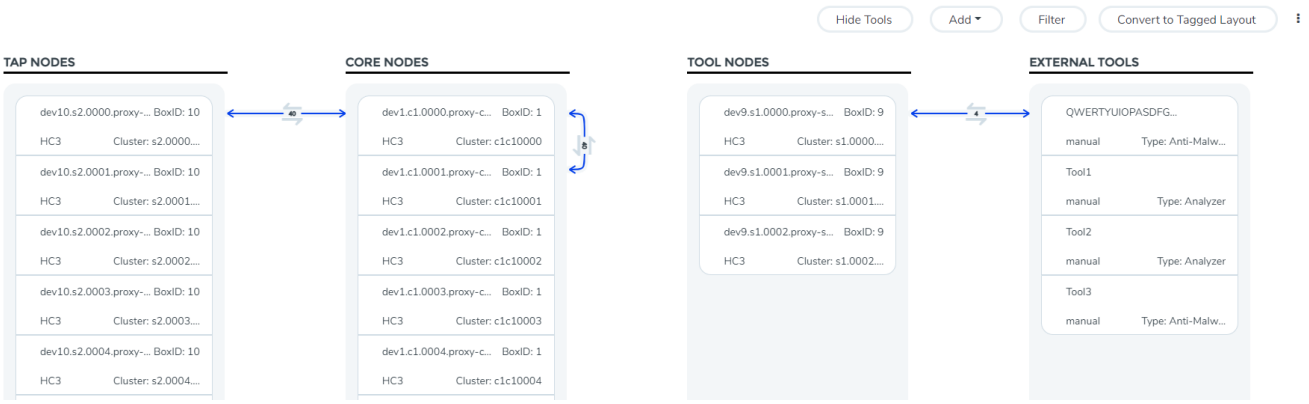
  NOTE: Use the **Show Tools/Hide Tools** option to display the external tools in the Topology page.

The devices that do not belong to either the TAP NODES or the TOOL NODES will be included in the CORE NODES category.

The edit option allows you to change the display names (default names) with the names of your choice for the following categories:

- TAP NODES
- CORE NODES
- TOOL NODES

You cannot edit the display name for the external tools.



- Click on a link or a device. The **Connections Nodes and Analytics** pane is displayed at the bottom. Refer to Smart Layout Topology section for details.

Notes:

- You can drag and drop the devices across the columns to adjust the position of the nodes. The tags will get automatically updated.
- **If the number of nodes exceeds 55**: GigaVUE-FM prompts you with a message indicating that the smart layout topology is applicable only up to 75 nodes. You can either retain the existing layout or click **Convert to Tagged Layout** to create tagged layout topology. This option is available only for admin users with read-write access to the Infrastructure Management resource category.
- **If 76th node is added to GigaVUE-FM**: The newly added node will not be visible in the smart layout topology.
- **If you switch to tagged layout topology**: You cannot switch back to smart layout topology unless the device count becomes less than 50.

Tagged Layout Topology

Tagged layout topology uses tags for grouping the nodes based on the following criteria:

- **Hierarchical tags for grouping based on deployment:** This allows you to drill down from a top-level hierarchical view to a lower-level detailed view.
- **Placement tags for placement of nodes:** This allows the you to see a detailed view of the devices from a left to right, source to tool view. Only single-valued tags can be used for this configuration.

When you first switch to tagged layout topology, if the devices have not been associated with tags, GigaVUE-FM prompts an error message to tag the devices from the physical nodes page.

Note:

- Only users with read-write access to the Infrastructure Management category can configure the required topology using tags. Refer to the 'Tags' section in the *GigaVUE Administration Guide* for detailed steps on how to create tags and how to associate the clusters/devices to tags.
- GigaVUE-FM allows you to configure tags for the individual nodes in a cluster using the Device Level Tagging option. This allows tags to be associated at the device level and allows the devices to be grouped for visualizing the topology. Only aggregation tag type is allowed for tagging of the devices.

Hierarchical View using Hierarchical Tags

Using Hierarchy Tags you can group the nodes based on a specific hierarchical view. For example to create a hierarchy based on the location of the nodes, you must create the following tags ids and tag values:

Hierarchy Level	Tag ID	Tag Values
I	Region	North, East, South, West
II	State	California, New York, Washington
III	City	San Francisco, Los Angels, Albany,
IV	Data Centers	DC1, DC2 in San Francisco DC3 in Los Angels DC4 in Albany DC5 DC6 in Rochester DC7 DC8, DC9 in Seattle

You are responsible for creating the appropriate tag keys and tag values and associate the nodes to the correct tag key and tag value to get the required hierarchical view of the tags. The tags can be either RBAC tags with hierarchical property set to True, or Aggregation tags.

NOTE: GigaVUE-FM does not validate the tag values assigned to the tag keys that are selected for hierarchical tag views. You are responsible for associating the correct tag keys and tag values to the nodes, so that the nodes are grouped appropriately.

Tagged Layout View using Single-valued Placement Tags

Using placement tags, you can place the nodes in distinct columns. For example, to place the nodes based on the type of the node, the tag key must be 'Node Type' and the tag values can be one of the following:

- Tap_processing_node
- Tool_node
- Leaf_node
- Spine_node

NOTE: Only aggregation tags can be configured at the node level tags.

Configure Topologies

To configure tagged layout topologies based on hierarchical and placements view, refer to the following sections:

- Rules and Notes
- Steps

Rules and Notes

Keep in mind the following rules and notes while creating tagged layout topology:

- The tag keys and the tag values required for the hierarchical view and placement view must be created in advance.
- The devices must be associated to the appropriate tag keys and tag values for visualizing the topology.
- Topology configurations related to subgrouping, color customization, alias names are user-specific configurations. That is, configurations done by you will not affect the views of other users.
- Gigamon Discovery Protocol (GDP) must be enabled on the nodes to visualize the links. For bulk configuration of GDP, refer to the Port Discovery page in the *GigaVUE Administration Guide*.
- The 'Device Level Node ID' allows you to associate tags to the devices, as against cluster-level tagging.

- You cannot drag and drop the devices across the hierarchical columns. That is, you cannot move the devices across the regions which is logically not possible from GigaVUE-FM GUI. However, you can drag and drop the devices with in the same row. That is, you can move the placements of the devices within the same location.

Steps

1. On the left navigation pane, go to **Inventory** go to **Physical** > **Topology**.
2. Click **Edit Configured Tags**. This option is available only for admin users with read-write access to the Infrastructure Management resource category. The Topology Tag Configurator page appears.
   - For Hierarchy Tags: Select the required Tag Keys that need to be added to the Selected Hierarchy Tags and click **Add**. Add the Tag Keys in the order based on which you want to drill down.
   - For Placement Tags: Select the required single-valued tags that need to be added to the Placement tags and click **Add.** Add the Tag Values in the order based on which you want to view the nodes in a row.
3. Click **Save**. The configuration is saved and the topology is displayed.



The Physical connections between the nodes across the groups are represented by a single consolidated line. Click on a specific node to view the physical connections of that node with other nodes.

Edit Topologies

To edit the topologies:

1. Click **Edit**.
2. Drag and drop the devices across the hierarchical columns.
3. Click **Save** to save the changes.

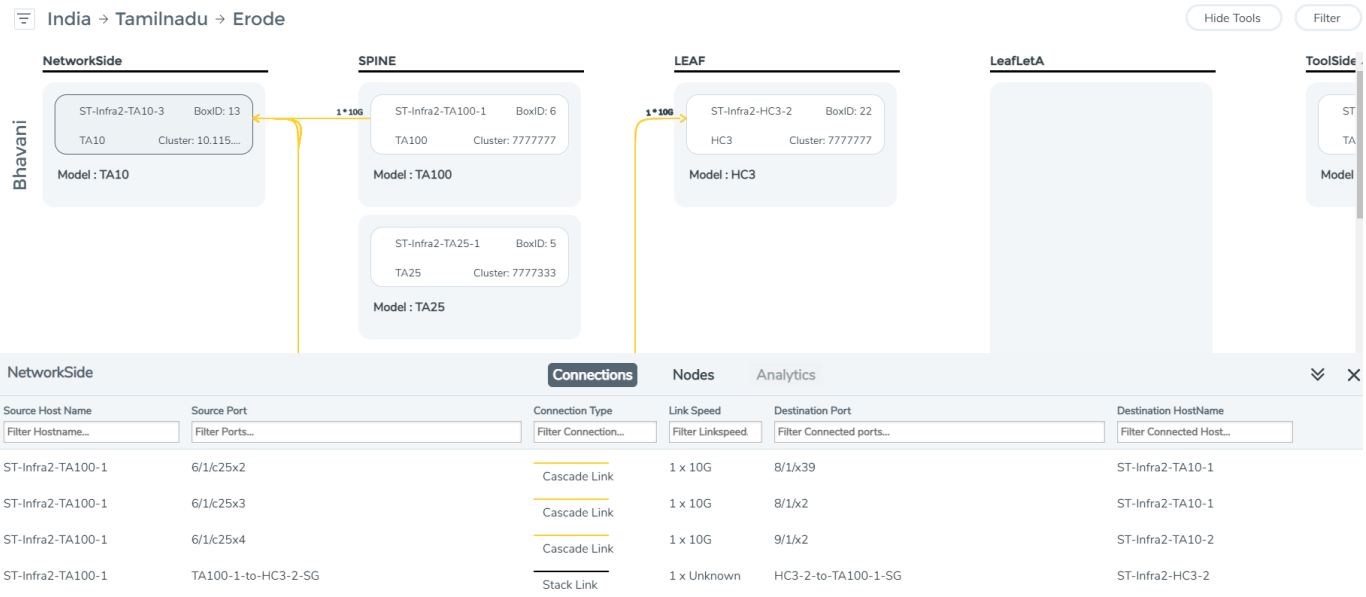
The Edit Configured Tags option is also displayed under the edit option.

View Consolidated Links

From the Topology page you can view the details of the devices and the connections associated with the devices for both smart layout and tagged layout topologies. Click on a link in the Topology page to view the following:

- **Nodes:** Displays details about the connected devices such as Cluster ID, Host name, Node IP, Model, Role, Box ID, and other such details.
- **Connections:** Displays details about the links connecting the Devices such as the host name, port or GigaStream Alias, Type of Connection, Link Speed and other details associated to the connection.
- Hover over the port id to view the health status of the port.
- **Analytics:** Displays the statistical details of the devices and the connections. Click on a specific device or a connection for which you need to view the details. The Analytics tab displays top 5 RX/TX packet drop, Data Rate (Octets) and utilization for a particular column, device or connection based on selection.

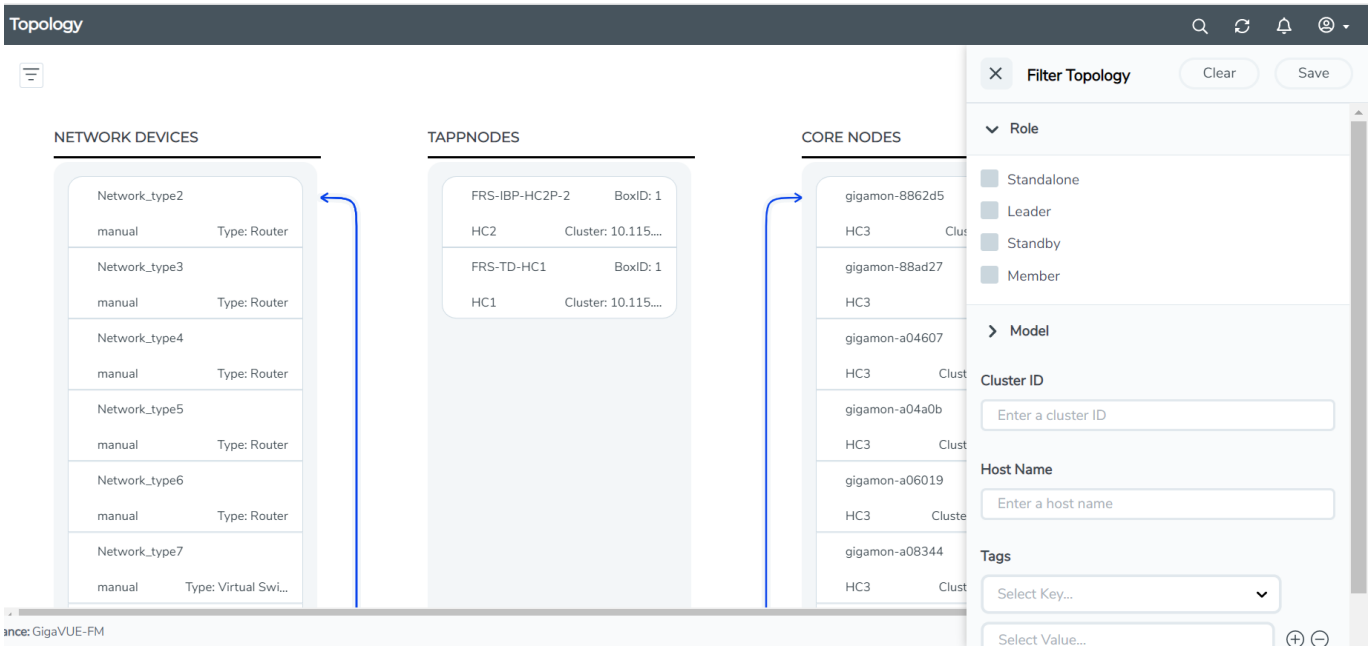
> **NOTE:** For tagged layout topologies, the Analytics tab is enabled only for the tags at the last hierarchical level.



How to Filter

The **Filter** button allows you to filter the devices in the Topology page based on the following criteria:

- **Role:** Select one of the following options:
  - Standalone
  - leader
  - Standby
  - Member
- Model
- Cluster ID
- Host Name
- Tags



The selected filter criteria appear as info panes on the topology page and the filtered node(s) will get highlighted. Click **Save** to save the filter. Click Clear to clear the filter.

Use the **Show Unmatched**/**Hide Unmatched** buttons to show or hide the unmatched nodes and clusters.

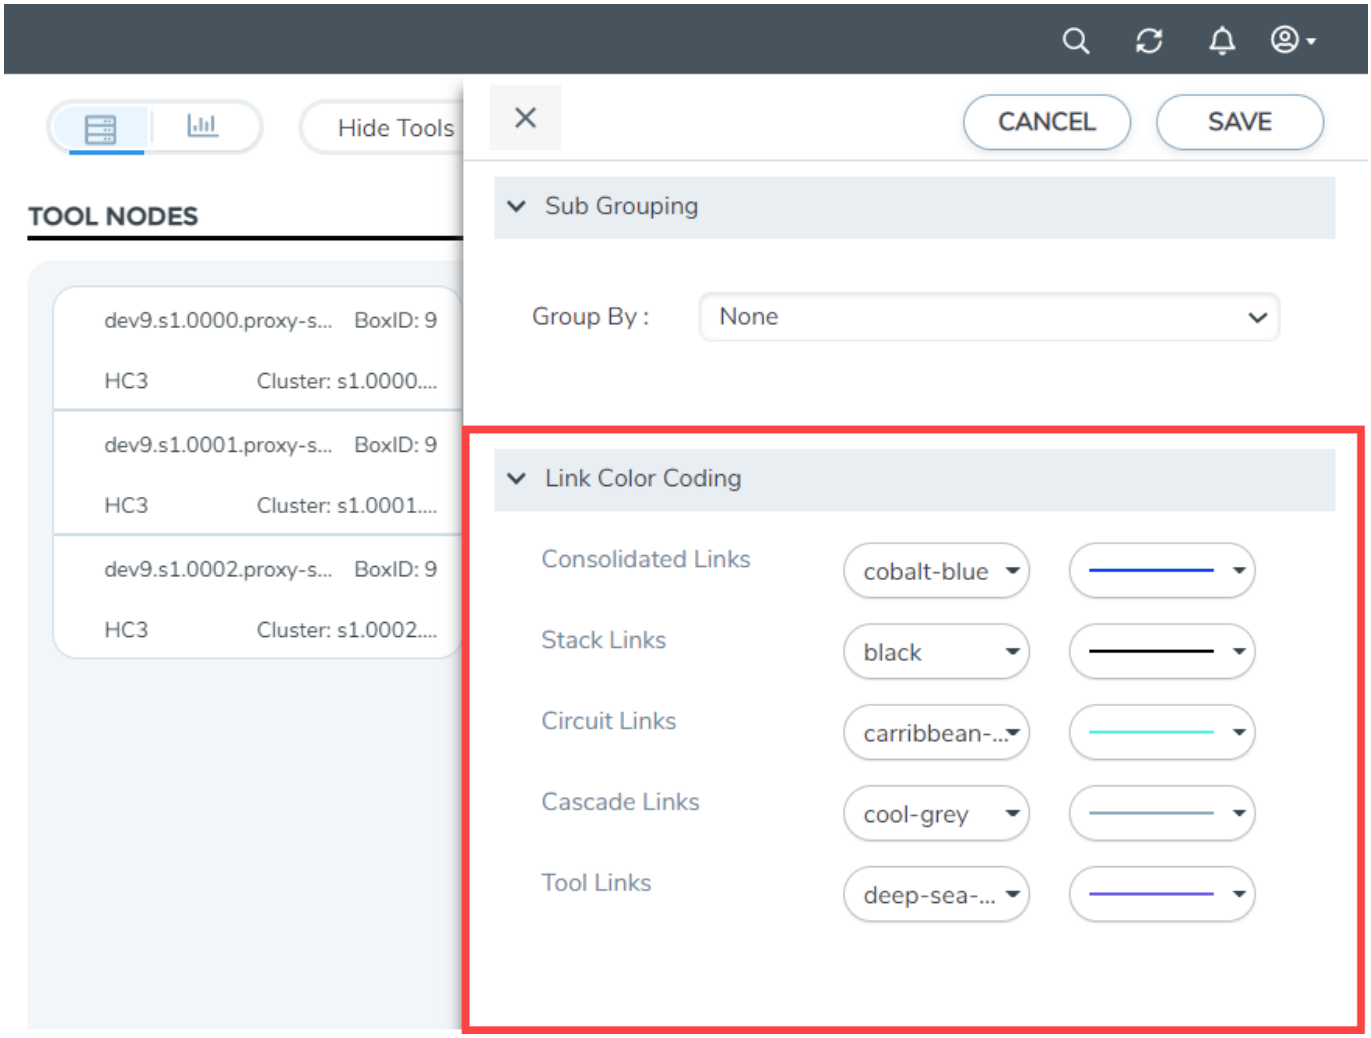
How to Customize the Colors in a Connection

The colors and the format of the connections in the Topology page can be configured as follows:

1. Click the ellipsis on the top right corner of the topology page.
2. Use the **Link Color Coding** option to edit the colors and formats of the following types of connections:
   - Consolidated Links
   - Stack Links
   - Circuit Links
   - Cascade Links
   - Tool Links

The changes are saved and will be applied for the connections when a new topology is created.

NOTE: Some of the colors available in the previous versions have been removed in software version 5.13.00. If you have used those colors in the previous versions, then on upgrading to software version 5.13.00, those colors will be replaced with the supported colors.

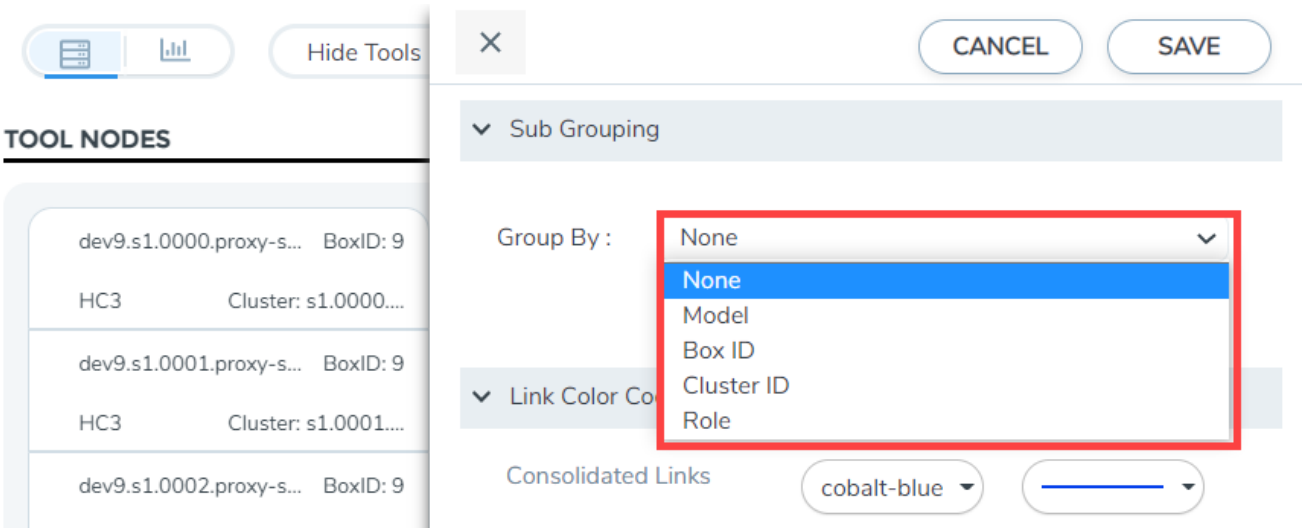


Sub-Grouping of Nodes

The Sub-Grouping option allows you to group the nodes in the Topology page:

1. Click the ellipsis on the top right corner of the topology page.
2. Use the **Sub Grouping** option to select the criteria based on which the nodes must be grouped:
   - Model
   - Cluster Id
   - Box Id
   - Role

The nodes can be grouped in both smart and tagged layouts topology.



Based on the current user who has logged in, the sub-grouping is saved.

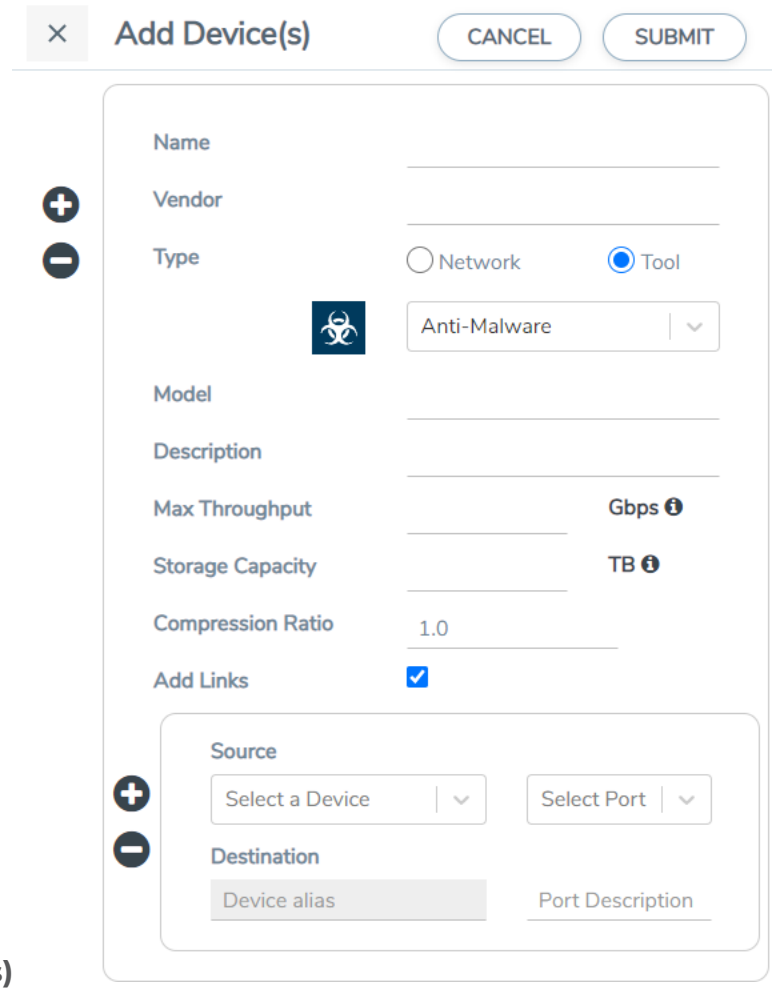
External Tool Links and Manual Device Links

The Topology page displays the links between the nodes and the external tools. The tools can be either manually added or discovered through CDP/LLDP by the tool port:

To add devices and links refer to the following section.

Add Devices and Links

To add devices:



1. Select **Add > Add Device(s)**.
2. You can add additional nodes by clicking the **+** button. To remove a node, click the **-** button.
3. Enter the following information for the node:
    - **Name**—the name of the node.
    - **Vendor**—the node's manufacturer. For Gigamon nodes, you cannot edit this field.
    - **Model**—the model number of the node.
    - **Description**—optional description or additional information about the node.
    - **Max Throughput** - the maximum throughput traffic currently being sent to this tool.
    - **Storage Capacity** - the total processing capacity dedicated to all Gigamon ports physically connect to the tool.

    NOTE: This field is not displayed when selecting a device type **Network**.

    The information entered for the node displays when hovering over the icon on the topology. Also, the type selected determines the icon.

4. Click in the **Type** field to select the type of device, either Network or Tool. Depending on your selection, the following icons are displayed:

Type		Icon
Network	Router	
	Switch	
	Virtual Switch	
	Load Balancer	
	Firewall	
	VOIP	
	Cloud	

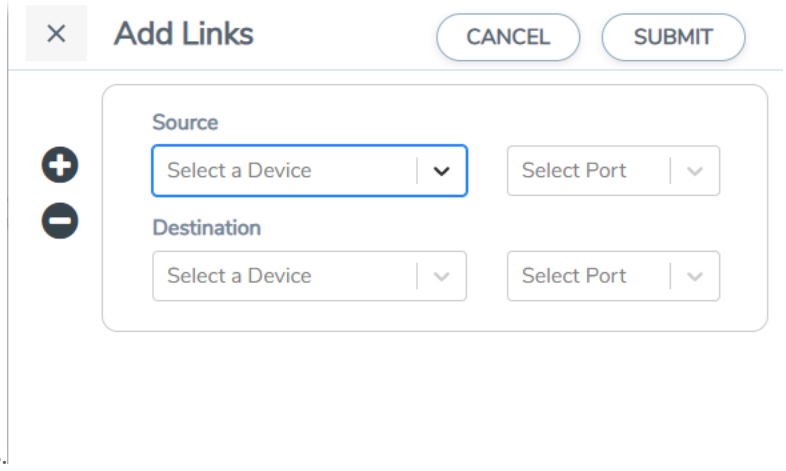
Type		Icon
Tool	Analyzer	
	Anti-Malware	
	Customer Experience Management	
	Application Performance Management	
	Data Loss Prevention (DLP)	
	Forensics	
	Intrusion Detection System (IDS)	
	Intrusion Prevention System (IPS)	
	Next Generation Firewall (NGFW)	
	SIEM	
	Other	

5. Click **Submit**.

You can also add devices from the Network Devices page. Refer to the Configure Network Devices section for more details.

To add links:

To add a link or links between nodes:



1. Click Add Links.
2. Select the **Source** device and the required source port from the drop-down list.
3. Select the **Destination** device and the required destination port from the drop-down list.

   NOTE: You can also select GigaStream as source or destination for the topology and select the required ports.
4. Click **Submit**.

You can also add links from the Connections page. Refer to Configure Connections section for details.

Recently Viewed

The **Recently Viewed** pane allows you to view the list of recently viewed topologies. Use the **Find** in **Recently Viewed** to find the topologies that you last searched for.

Configure Network Devices

The Network Devices page displays a list network devices configured in GigaVUE-FM.

To view the network devices, log into GigaVUE-FM. On the left navigation pane, go to **Inventory** and under **Physical**, select **Network Devices**.

The Network Devices page displays the following information:

Field	Description
Network Name	Name of the network

Field	Description
Vendor	Vendor name
Type	Type of device
Model	Model of the device
Description	Brief description about the device
Discover Type	Type of discovery
System Name	Name of the system (only for CDP/LLDP links from external device)
System Description	Brief description about the system
Platform	System platform
Software version	Software version of the network device

NOTE: The columns in the Network Devices page can be customized based on the type of content you want to view in the table. For customizing the columns, refer to Table View Customization.

Use the following buttons to manage the network devices in GigaVUE-FM:

Button	Description
Filter	Use to filter the network devices. You can filter the devices based on the following fields: • Network Name • Vendor • Type • Model • Discovery Type
Add	Use to add new network devices. Enter the following information: • Name • Vendor • Type • Model • Description You can also add Links when adding the manual devices. Refer to the Configure Connections section for more details.

Button	Description
Actions	Use to Edit/Delete the existing manual devices
Import	Use to import the network devices in to GigaVUE-FM. You can only import .CSV files. **NOTE:** You can only import .CSV files. After the file is imported, a pop-up message appears that shows the number of devices that were successfully imported and the number of devices that failed to import. For failed devices, click on the link in the message, and you will be navigated to the Events page.
Export	Use to export the devices from GigaVUE-FM, either in .CSV or .XLSX format. You can export all the devices or only the selected devices.

Configure Connections

The Connections page displays the links and connections in the topology page.

To view the connections, log into GigaVUE-FM. On the left navigation pane, go to **Inventory** and under **Physical**, select **Connections**.

The Connections page displays the following information:

Field	Description
Source Host Name	Name of the source host
Source Port	Port ID/Alias of the source port. Hover over the port id. A pop-up appears that displays the following details: • Destination Port ID/Alias • Health State • Link State • Destination Port ID/Alias • Health State
Connection Type	Type of connection
Connection Speed	Speed of connection
Connection	Source of connection

Field	Description
Source	
Destination Port	Port ID/Alias of the source port. Hover over the port id. A pop-up appears that displays the following details: • Destination Port ID/Alias • Health State • Link State • Destination Port ID/Alias • Health State
Destination Host Name	Name of the destination host.

Use the following buttons in the Connections page, as required:

Button	Description
Filter	Use to filter the connections listed in the connections page. You can filter the links based on any of the fields.
Add	Use to add new links/connections. Enter the following information: • Source Device • Destination Device • Source Port • Destination Port
Actions	Use to delete or edit the connection.
Import	Use to import the connections in to GigaVUE-FM. **Note**: You can only import .CSV files. The Source Cluster ID and Destination Cluster ID fields are mandatory in the .CSV file. After the file is imported, a pop-up message appears that shows the number of connections that were successfully imported and the number of connections that failed to import. For failed connections, click on the link in the message, you will be navigated to the Events page.
Export	Use to export the connection details from GigaVUE-FM, either in .CSV or .XLSX format. You can export all the connections or only the selected connection.

Configure Tools

The Tools page displays the list tools configured in GigaVUE-FM.

To view the Tools:

1. Log into GigaVUE-FM.
2. On the left navigation pane, go to **Inventory** and under **Physical**, select **Tools**.

The Tools page displays the following information:

Field	Description
Tool Name	Name of the tool.
Vendor	Vendor name.
Type	Type of tool.
Model	Model of the tool.
Description	Brief description about the tool.
Total Ports	Ports associated with tool.
Number of Maps	Number of maps involving the tool port.
Max Throughput	Maximum throughput sent to this tool.
Storage Capacity	Total processing capacity dedicated to all Gigamon ports physically connected to the tool.
Wrap Around Time	Total Storage (User entered data for the Tool)/Current Throughput
Throughput	Traffic throughput.
Compression Ratio	Compression ratio.
Discovery Type	Discovery type.
System Name	System name.
System Description	System description.
Platform	Platform.
Software Version	Software version.

NOTE: The columns in the Tools page can be customized based on the type of content you want to view in the table. For customizing the columns, refer to Table View Customization.

Use the following buttons to manage the tools in GigaVUE-FM:

Button	Description
Filter	Use to filter the tools. You can filter the tools based on the following fields: • Tool Name • Vendor • Type • Model • Description • Ports • Tags
New Device(s)	Use to add new tools. Enter the following information: • Name • Vendor • Type • Model • Description • Max Throughput • Storage Capacity • Compression Ratio • Add Links You can also add Links when adding the tools. Refer to the Configure Connections section for more details.
Actions	Use to perform the following actions: • Edit • Delete • View Details • View Statistics Graph
Import	Use to import the tools in to GigaVUE-FM. **NOTE:** You can only import .XLS files. You cannot import more than 50 tools at a time. On successfully importing the tools, a notification pops up. A notification also pops-up when the import operation fails.

Button	Description
Export	Use to export the tools from GigaVUE-FM, either in .CSV or .XLSX format. You can export all the tools or only the selected tools.

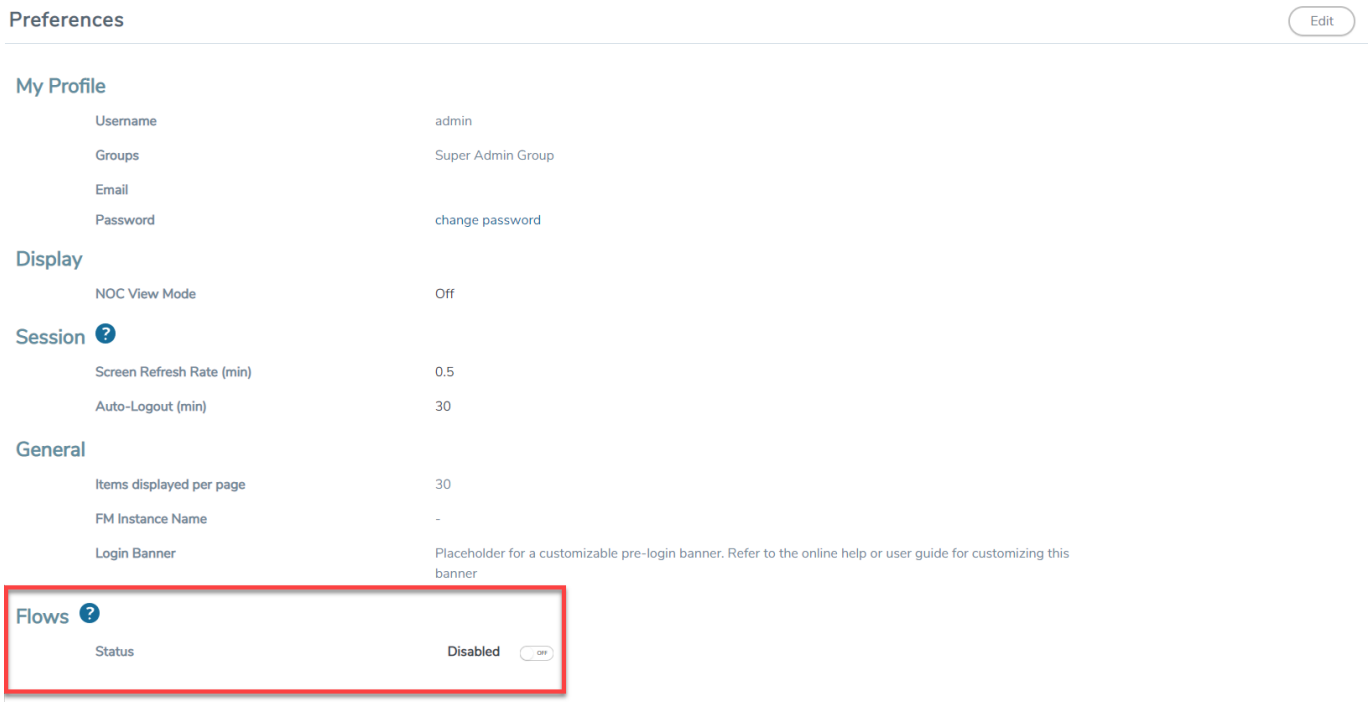
Flows

This chapter describes what is a flow and how to quickly view the packet drops and packet errors that causes the flow to be unhealthy.

Considering the system resources needed for its calculation, flows is disabled by default in GigaVUE-FM. However, you can enable flows if required. For larger deployments, exercise caution before enabling flows.

To enable Flows:

1. Go to **Settings** and select **System** > **Preferences**.
2. Click **Edit**.
3. On the Edit Preferences page, under **Flows**, change the status to Enabled.



NOTE: You cannot disable flows once it is enabled. Contact Customer Support to disable Flows.

Refer to the following sections for details:

- About Flows
- View Flows
- View the Flow Summary and Statistics
- How to Change the Flow Layout
- How to Update Flows
- View Events
- Set Notifications
- Limitations of Flows

About Flows

Flows provide the ability to view the traffic flowing from a network port that receives the traffic from a network TAP to the tool port that sends the traffic to the tools.

A flow is constructed by traversing backward starting from the egress tool port connected to the monitoring tools all the way up to the network ports that receive the traffic from a network TAP. If there are five egress tool ports sending the traffic out to the monitoring tools, there will be five flows displayed in the Flows page.

Figure 13Flow - All Sites shows an example of a flow. It illustrates the network ports that receive the packets from different sources, a set of map rules that filter the packets, and the egress tool port that sends the packets to the monitoring tools. A flow name is determined by the egress tool port ID or alias. In this example, the flow name is SPEGE0070_Te5, which is the name of the egress tool port.

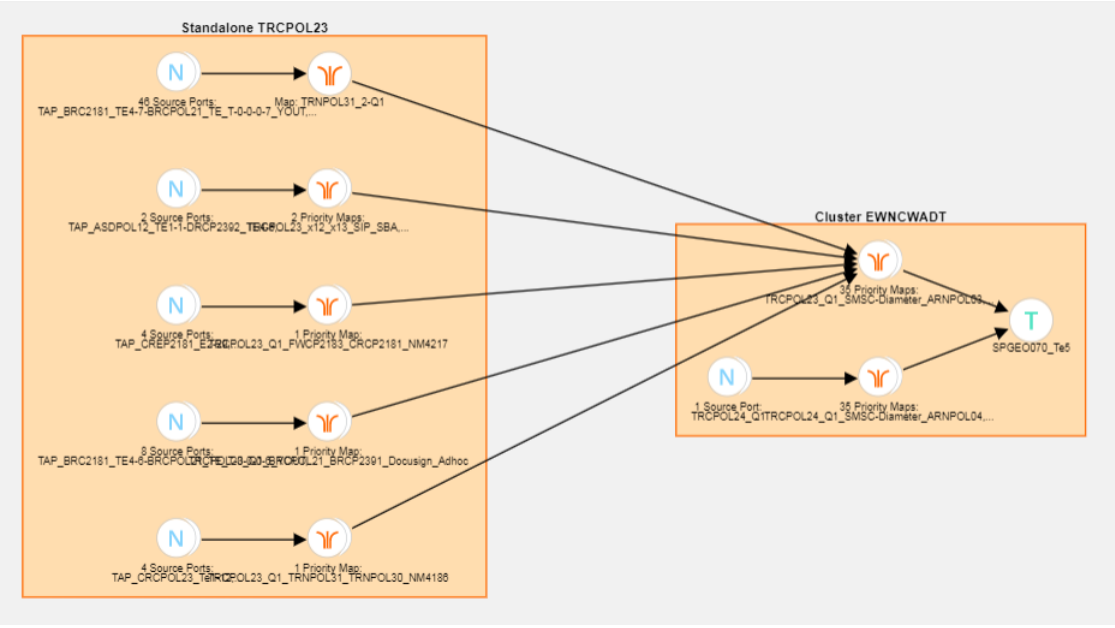


Figure 13 *Flow - All Sites*

Figure 14Standalone GigaVUE Nodes shows a number of GigaVUE nodes in the Topology page that are not connected to each other.

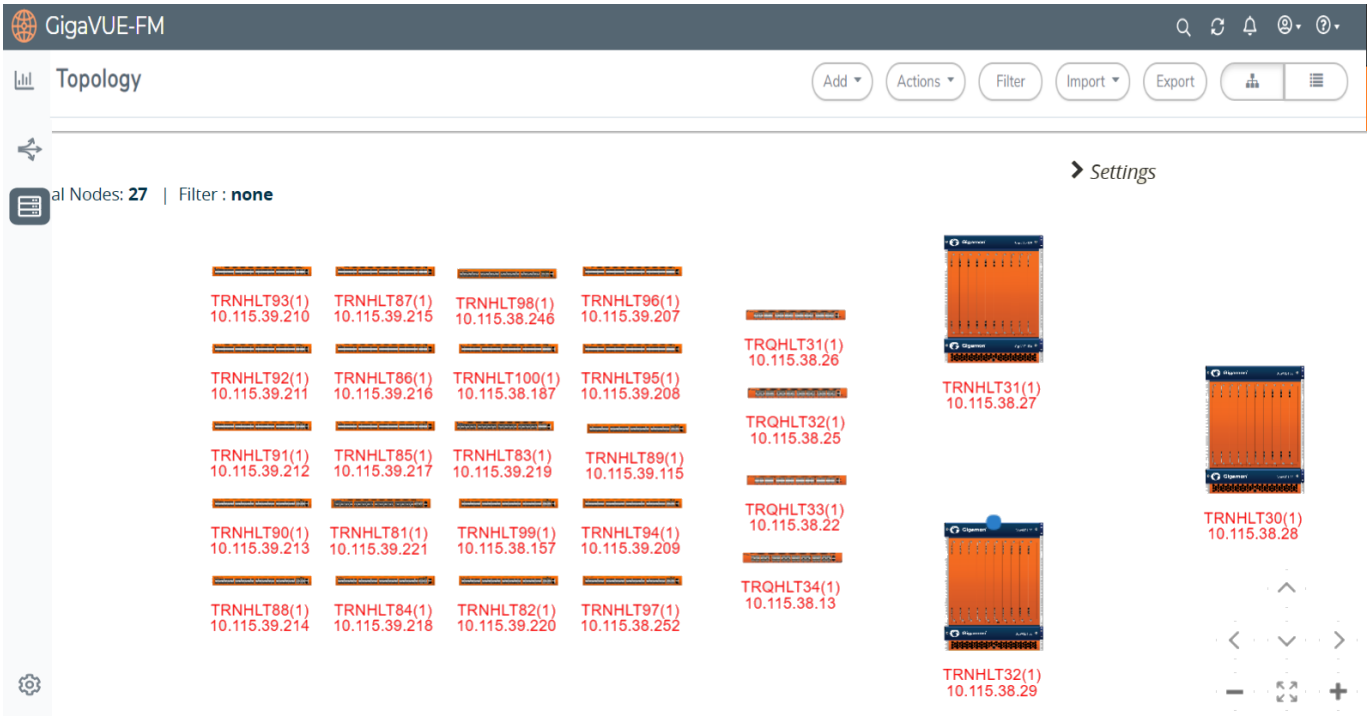


Figure 14 *Standalone GigaVUE Nodes*

Each node seen in Figure 14Standalone GigaVUE Nodes is represented as a separate flow in the Flows page.

Total Flows: **33**

Name	Status	Total Ports	Total Unhealthy Ports	Total Maps	Total Unhealthy Maps
TRQHLT31_TRNFHLT32	Flow is healthy	12	0	10	0
TRNHLT87_TRQHLT31-34	Flow is healthy	36	0	2	0
DAMHLT_Te1&Te12	Flow is healthy	9	0	1	0
TRNHLT94_TRQHLT31-34	Flow is healthy	36	0	2	0
TRNHLT88_TRQHLT31-34	Flow is healthy	36	0	2	0
TRNHLT90_TRQHLT31-34	Flow is healthy	36	0	2	0
TRNHLT85_TRQHLT31-34	Flow is healthy	36	0	2	0
TRNHLT93_TRQHLT31-34	Flow is healthy	36	0	2	0

Figure 15 *Flows Representing the Standalone Nodes*

When those GigaVUE nodes are connected using manual links or Gigamon Discovery links, the number of flows created depends on the number of egress tool ports sending the traffic out to tools. In this example, three flows are created as shown in Figure 14Standalone GigaVUE Nodes.

Total Flows: **3**

Name	Status	Total Ports	Total Unhealthy Ports	Total Maps	Total Unhealthy Maps
DAMHLT_Te1&Tel2	Maps [TRNHLT81_1Q1Q2_TRNHLT31_1Q7Q8_GTP-DCPCF001, TRNHLT81_1Q1Q2_TRNHLT31_1Q7Q8_Permanent] in the flow are unhealthy	843	2	125	2
DAMHLT02_Te1	Maps [TRNHLT81_1Q1Q2_TRNHLT31_1Q7Q8_GTP-DCPCF001, TRNHLT81_1Q1Q2_TRNHLT31_1Q7Q8_Permanent] in the flow are unhealthy	426	2	63	2
DAMHLT01_Te1	Flow is healthy	426	0	63	0

Figure 16 *Healthy Flows - sample illustration of flows after connections are made*

NOTE: Gigamon Discovery is supported only from GigaVUE-FM 5.2 and above.

Figure 22Unhealthy priority Maps shows how unhealthy priority maps are illustrated in a flow view page.

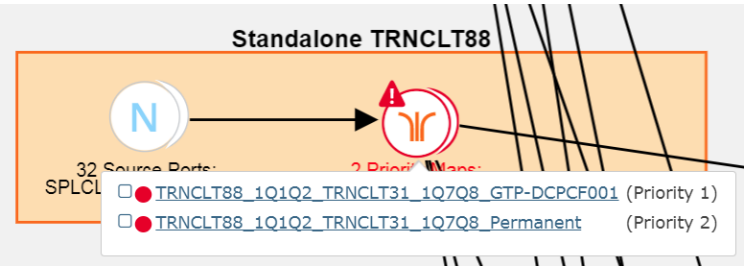


Figure 17 *Unhealthy Maps*

Using flows, you can perform the following:

- Quickly identify the maps and ports that are unhealthy.
- Quickly drill down the unhealthy map to investigate the cause of failure.
- Export a list of unhealthy ports and maps.
- View the data traffic across the flow, starting from the egress tool port that sends the traffic to the monitoring tool to the ingress network ports that receives the traffic from the network taps or span/mirror ports. The arrows in the flow indicate the path of the traffic flow.
- View the pass all maps and the priority maps. A priority map set contains multiple maps configured with the same source ports in the port list.
- Select multiple maps and view the statistics to verify if the traffic is flowing as expected.

- Filter flows by Flow Name, Status or Cluster ID.
- Update flows.

A flow is automatically constructed every 24 hours, and flow health is calculated every 5 minutes. You can also manually trigger flow and graph calculations on demand.

The following components affect the health of a flow:

- Manual links or Gigamon Discovery links that connect the GigaVUE nodes.
- Maps that are participating in the flow.
- Health status of all the underlying components of a map such as GigaStream, port group, GigaSMART group, and vport.

For information on flow health status, refer to Flow Health Status. For instructions on updating flows, see How to Update Flows.

View Flows

On the left navigation pane, click on  select **Physical> Orchestrated Flows > Flow Maps,** and then select the required cluster or node ID to view a Flow. The Flows page provides a summary of the flow name, cluster, host name, status, total ports, total unhealthy ports, total maps, total unhealthy maps and last computed time.

Table 5: Flows Summary Page

Option	Description
Name	Name of the flow.
Cluster	Where a specific flow is ending. Cluster ID
Host Name	Host name of the cluster.
Status	The Heath status of the flow.
Total Ports	The total number of ports involved in the flow are displayed in the Total Ports link.
Unhealthy Ports	Total number of ports that are unhealthy. If there are unhealthy ports, click the Unhealthy Ports link to view the related ports that are unhealthy.

Option	Description
Total Maps	Total number of maps participating in the flow. Click the Related Maps link to view the related maps in the flow.
Unhealthy Maps	Total number of maps that are unhealthy. If there are unhealthy maps, click the Unhealthy Maps link to view the related maps that are unhealthy.
Last Computed Time	

1. Click the numbers under Total Ports, Total Unhealthy Ports, Total Maps, and Total Unhealthy Maps to view detailed information about the maps and ports participating in the flow. Refer to View Maps and Ports.
2. To open a flow, select a flow and click **Actions** > **Open Flow**. Alternatively, click a flow.

View the Flow Summary and Statistics

A flow can display all maps that are created for managing the packet distribution in the GigaVUE nodes. A flow summary provides information only about the total number of maps and ports participating in the flow along with the total number of unhealthy maps and ports in the flow. You can also select multiple maps and view the statistics to check how exactly the packets are flowing.

To view the flow summary:

1. Follow steps 1 to 3 as described in View Flows.
2. In the Flows page, click a flow that you want to view. Alternatively, select a flow and click **Actions** > **Open Flow.** The flow view page is displayed.

   The following table describes the information provided in the Summary tab:

Table 6: Flow Summary

Option	Description
Related Ports	Total number of ports participating in the flow. Click the Related Ports link to view the related ports in the flow. Refer to View Total Ports. **NOTE:** V ports and GigaSMART ports are not considered.
Unhealthy Ports	Total number of related ports that are unhealthy. If there are unhealthy ports, click the Unhealthy Ports link to view the related ports that are unhealthy. Refer to View Total Unhealthy Ports. To know more about how the health of a port is calculated, refer to Port Health Status.
Related Maps	Total number of maps participating in the flow. Click the Related Maps link to view the related maps in the flow. Refer to View Total Maps.
Unhealthy Maps	Total number of related maps that are unhealthy. If there are unhealthy maps, click the Unhealthy Maps link to view the related maps that are unhealthy. Refer to View Unhealthy Maps. To know more about how the health of a map is calculated, refer to Map Health Status.

3. To view the statistics, select the maps. To select the maps, follow one of the following methods:

   Method 1:

   a. Click on a map. A list of priority maps are displayed. Refer to Figure 18Selecting Maps From Maps List.

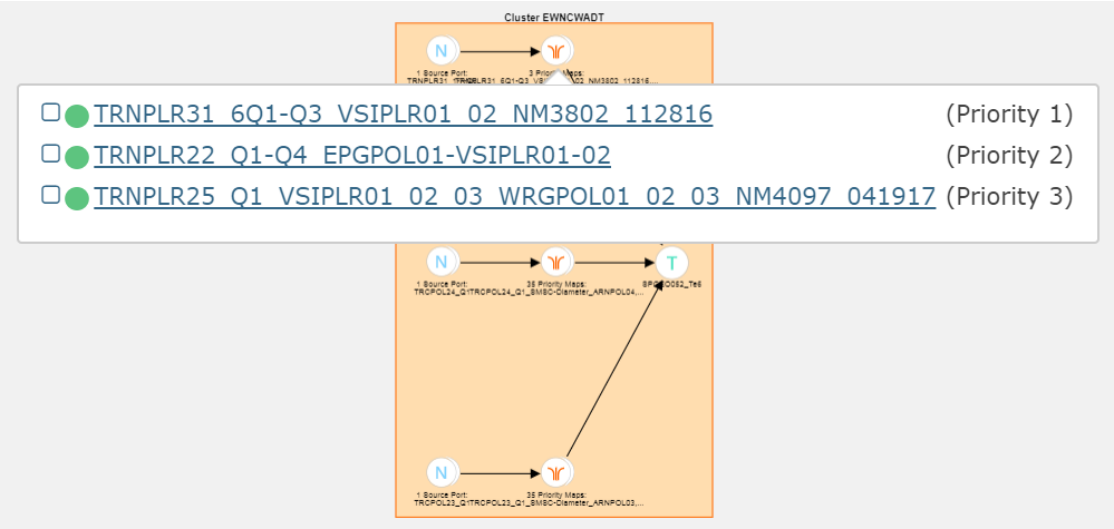


Figure 18 *Selecting Maps From Maps List*

b. Click the check box next to the priority map. The selected map is displayed in the **Items Selected** pop-up. Refer to Figure 19Selected Maps in Items Selected Pop-up.

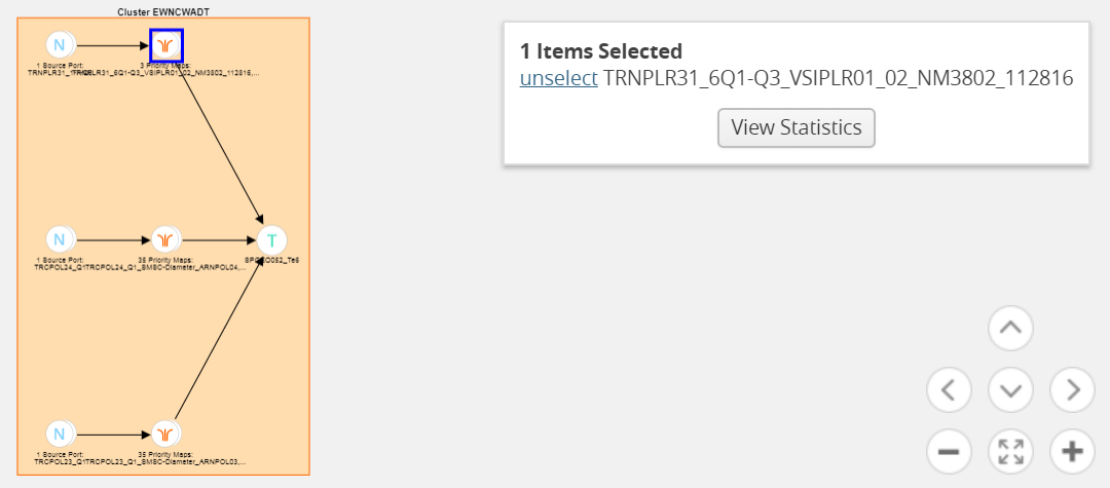


Figure 19 *Selected Maps in Items Selected Pop-up*

c. Repeat step a and b for selecting multiple maps.

Method 2:

a. In the flow view page, double-click on a map. A list of priority maps are displayed. Refer to Figure 20List of Priority Maps.



Figure 20 *List of Priority Maps*

b. Click on a map. The map is selected and is simultaneously displayed in the **Items Selected** pop-up. Refer to Figure 21Maps Selected.

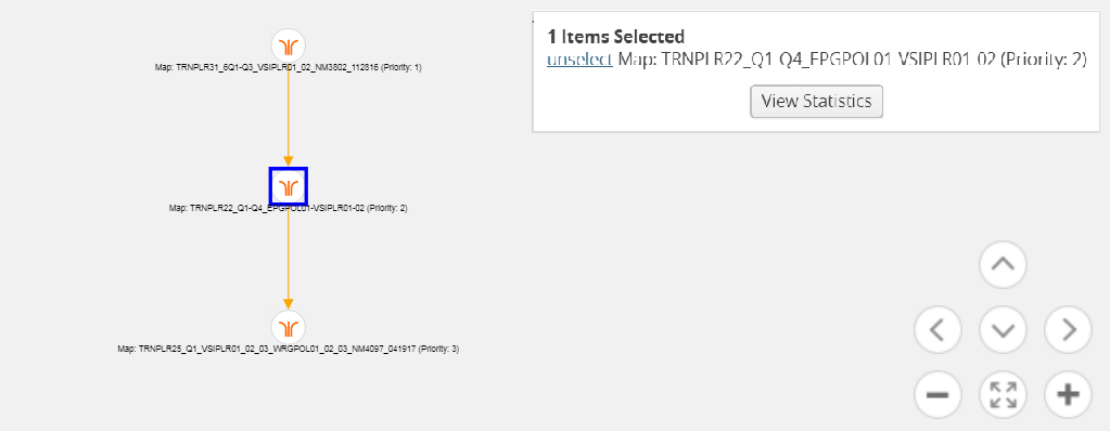


Figure 21 *Maps Selected*

c. Repeat step b to select multiple maps.

4. In the **Items Selected** pop-up, click **View Statistics**. Alternatively, select the **Statistics** tab and click **View Statistics**. The **Statistics** tab displays a graph to show how the packets are flowing from the selected maps.

View Maps and Ports

Figure 22Unhealthy priority Maps shows how unhealthy priority maps are illustrated in a flow view page.

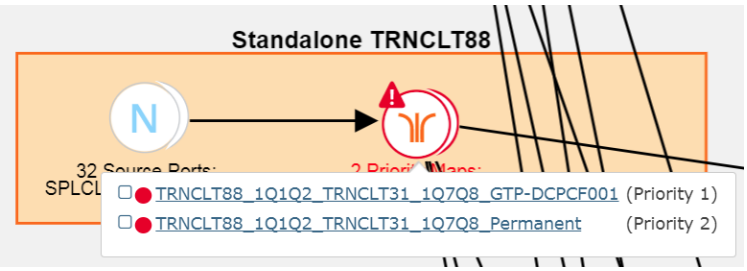


Figure 22 *Unhealthy priority Maps*

Figure 23Unhealthy Pass-all Map shows how a pass-all map is illustrated in a flow view page.

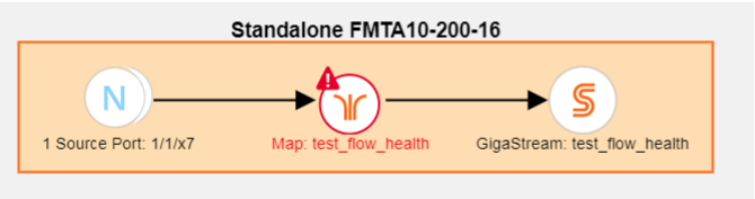


Figure 23 *Unhealthy Pass-all Map*

A priority map set as illustrated in Figure 22Unhealthy priority Maps can contain one or more maps configured with the same source ports. The health of this map is determined by the aggregated health of the priority maps in the map set. Sometimes, this map might be unhealthy because a map that is not participating in the flow is unhealthy. The overall health status of a flow is determined based on the aggregated health of all the maps that are involved in the flow.

There are two ways to view the total number of maps and ports that are healthy and unhealthy.

- In the Flows page, click the total number of ports, unhealthy ports, total maps, and total unhealthy maps to view detailed information about the maps and ports involved in the respective flow.
- Click a flow in the Flows page. In the Summary tab, click the Related Ports, Total Unhealthy Ports, Related Maps, and Total Unhealthy Maps links to view detailed information about the maps and ports involved in the selected flow.

View Total Ports

The total number of ports involved in the flow are displayed in the Total Ports link.

To view the total ports involved in the flow:

1. In the Flows page, click the **Total Ports** link. The All Ports quick view displays a list of all ports that are involved in the flow. It also provides information such as port ID or alias, type, port health status, and the node on which the port is configured.
2. In the All Ports quick view, click a Port ID to open the Port quick view.
3. To return to the All Ports quick view, click **Back**.
4. In the All Ports quick view, click a node under the Node column to open the Overview page.

View Total Unhealthy Ports

The total number of unhealthy ports involved in the flow are displayed in the Total Unhealthy Ports link.

To view the unhealthy ports:

1. In the Flows page, click the **Total Unhealthy Ports** link. The Unhealthy Ports quick view displays a list of unhealthy ports that are involved in the flow. It also provides information such as port ID or alias, type, port health status, and the node on which the port is configured.

2. In the Unhealthy Ports quick view, click a Port ID to open the Port quick view.
3. Click a node under the Node column to open the node's Overview page.

View Total Maps

The total number of maps involved in the flow are displayed in the Total Maps link.

1. In the Flows page, click the **Total Maps** link. The All Maps quick view displays a list of all maps that are involved in the flow. It also provides information such as map ID or alias, type, map health status, and the priority node on which the map is configured.
2. In the All Maps quick view, click a Map alias to open the Map quick view.

View Unhealthy Maps

The total number of unhealthy maps involved in the flow are displayed in the Total Unhealthy Maps link.

In the Flows page, click the **Total Unhealthy Maps** link.

The Unhealthy Maps quick view displays a list of all unhealthy maps that are involved in the flow. It also provides information such as map alias, type, map health status, and the priority node on which the map is configured.

In the Unhealthy Maps quick view you can:

- Click **Export** to export an Excel report listing the unhealthy maps.
- Click a a map alias link in the Alias column to open the Map quick view. Refer to View Total Maps.
- Click a node under the Node column to open the Overview page.

Filter Flows

To filter flows by name, do the following:

1. On the left navigation pane, click on  select **Physical> Orchestrated Flows > Flow Maps,** and then select the required cluster or node ID.
2. Click the **Filter** button at the top of the pane. The Filter panel displays.

3. Start typing the first few characters of the name of a flow you wish to filter on. The flow list is sorted and all the flows containing the characters you entered displays in order at the top on the page.

To filter flows based on their status:

4. Click the **Status** drop down and select one of the Status colors to display only those flows that correspond to the color selected.

To filter flows based on the Cluster ID:

5. Click the **Cluster ID** drop down and select one of the Cluster IDs to display only those flows that correspond to the Cluster ID you selected.

How to Change the Flow Layout

By default, a flow layout is arranged in a directional flow. The ingress network port receiving the traffic from the network TAP is aligned to the left and the tool port sending the traffic to the monitoring tool is aligned to the right. The maps configured on the standalone nodes and clusters are displayed within the respective containers. The arrows in the layout shows the direction in which the packets are moving from network ports to tool ports. You can drag and drop the cluster or node containers to change the alignment. You can also drag and drop the ports and maps within the container to change the layout.

To reset, save, or restore the layout:

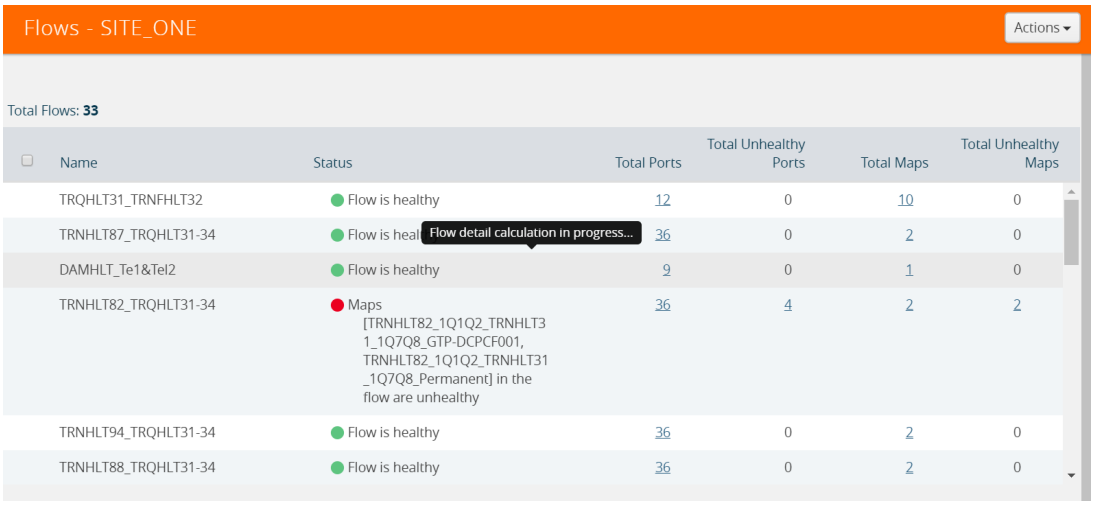
1. In the Flow view page, click and drag a port or a map to change the layout as desired.
2. To save the new flow layout, click **Actions > Save current layout**. The current alignment of the flow is saved.
3. To align the layout to default, click **Actions > Automatically align layout**. The layout is automatically set to the default.
4. To restore the saved layout, click **Actions > Restore saved layout**.

How to Update Flows

Flows are computed based on the Gigamon discovery or manual links connecting the GigaVUE nodes and the maps participating in the flow. The changes made to the flow are not seen unless they are automatically updated after every 24 hours or manual updated from the Flows page.

To manually update a flow:

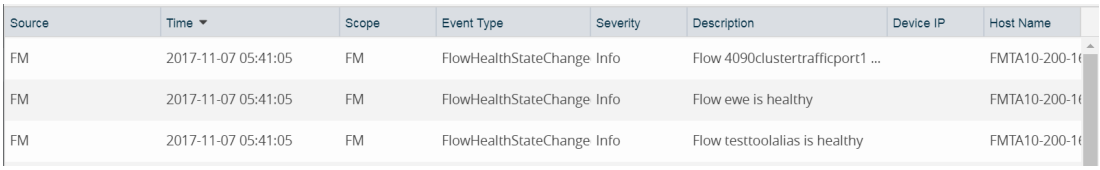
1. On the left navigation pane, click on  select **Physical> Orchestrated Flows > Flow Maps,** and then select the required cluster or node ID.
2. Click **Actions** > **Update**. A message is displayed to indicate that the flow calculation is in progress. The check boxes to select a flow disappear when the flow is being updated. They slowly start appearing when the update is completed for that particular flow.



View Events

On the Events page, the FlowHealthStateChange event type indicates that there is a change in the health status of a flow. For more information about Events, refer to the *"Events"* section in the *GigaVUE Administration Guide*.

Refer to the figure for the flow health change event on the **Events** page.



Set Notifications

On the Notifications page, the email notification for the **Flow Health State Changed** event type can be configured to automatically send emails to the specified addresses when there

is a change in the health of any of the flows managed by GigaVUE-FM. For more information about configuring notifications, refer to the *"Notifications"* section in the *GigaVUE Administration Guide*.

Refer to Figure 24Flow Health State Changed Notification for the flow health state changed notification on the **Notifications** page.

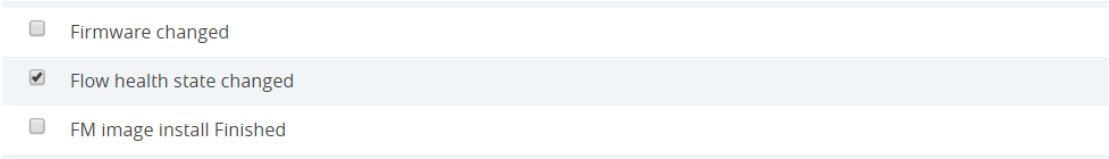


Figure 24 *Flow Health State Changed Notification*

Limitations of Flows

If Gigamon Discovery is enabled on GigaVUE nodes and a link is created between a source and a destination node that are of hybrid port types, the flow ends at the source node hybrid port. As hybrid port acts as an indirect traffic source port and as a tool port, the flow construction fails to identify the direction of the traffic flow. The workaround is to add a manual link to indicate the direction of the traffic flow.

Device Logs and Event Notifications

This chapter describes how to view the logs for the selected node.

By default, GigaVUE-FM acts as a syslog server. However, you can use your own syslog collector instead of using GigaVUE-FM. This optimizes the performance of GigaVUE-FM.

Refer to the following sections:

- Stream Device Logs to GigaVUE-FM
- View Device Logs
- Device Log Host Servers
- Storage Management for Device Logs
- Manage Device Log Output

Stream Device Logs to GigaVUE-FM

GigaVUE HC Series device/cluster nodes provide comprehensive logging capabilities to keep track of system activity. Logging is particularly useful for troubleshooting system issues, as well as maintaining an audit trail. You can specify what types of events are logged, view log records by priority, date, or name, and upload log files to a remote host for external troubleshooting.

When HC Series devices are added to GigaVUE-FM, the default settings ensure that log records stream to the GigaVUE-FM server and are written on the device's file system in a messages file. Log records can be used to analyze the system behavior directly from the GigaVUE-FM interface instead of needing to sign in to each device.

External third-party servers can also be added to host the streamed logs. Refer to Add an External Logging Host Server to a Node.

In this section:

- Cluster Behavior
- Standardized Logs
- Device Log Categories
- Device Log Message Types
- Device Logging Levels
- Device Logging Processes

Cluster Behavior

Logging configuration is local to the node. Each node has its own logging configuration.

Important: This is a different behavior than in software versions prior to 5.3 in which the logging configuration on the leader was synchronized to the other nodes in the cluster.

Standardized Logs

Standardized log messages can be streamed to GigaVUE-FM. The log messages follow the industry standard described in RFC5424. The format includes structured data, timestamp, version, and message ID. The timestamp includes milliseconds for increased accuracy.

In GigaVUE-FM, the log information is displayed in a table with configurable, sortable columns and extensive filter options.

Refer to View Device Logs for information about viewing and filtering logs in GigaVUE-FM.

Refer to the **GigaVUE-OS-CLI User's Guide** for information about the standardized log message format in the CLI view.

Device Log Categories

Table 7: Device Log Categories describes the categories in device/cluster log messages as per the standardized log format. Refer to Standardized Logs for standardization information.

Table 7: Device Log Categories

Message ID	Description
GENERAL-ERR	Errors that are common across applications
HIGH-TEMPERATURE	High Temperature
PACKETDROP	Packet Drop
LINK	Link
INIT	Initialization
RESOURCE-UTIL	Resource Utilization
REQUEST	Query system information

Device Log Message Types

Table 8: Device Log Message Types describes the message types in device/cluster log messages as per the standardized log format. Refer to Standardized Logs for standardization information.

Table 8: Device Log Message Types

Message ID	Description
memoryAllocError	Memory allocation error
memoryAccessError	Unexpected NULL access
fileAccessError	File access error

Message ID	Description
switchCpuHighTemperature	High Switch CPU Temperature
opticsHighTemperature	High Optics Temperature
gigasmartCpuHighTemperature	High GigaSMART CPU Temperature
ambientCpuHighTemperature	High Ambient CPU Temperature
exhaustCpuHighTemperature	High Exhaust CPU Temperature
egressPacketDrop	Packet Drop at Egress
ingressPacketDrop	Packet Drop at Ingress
linkChangeNotify	Link Change Notification
mgmtModule	Failed to load mgmt module
systemCpuUtil	High System CPU Utilization
systemMemoryUtil	High System Memory utilization
processCpuUtil	High CPU Consumption by a process
processMemoryUtil	High Memory Consumption by a process
processAccessError	Process access errors
systemAccessError	Process access errors
queryFail	Failed to fetch system data
moduleInitFail	Failed to load a cli module

Device Logging Levels

Table 9: Logging Levels shows the standard logging levels that are used to rank logged events by degree of severity. When configuring the device/cluster logs to stream, the highest severity level will capture fewer logs than the lowest severity. For example, if the lowest level (info or debug) is selected in your configuration, all levels of logs, from the selected level up to the highest level, will be streamed. If the highest level (emergency) is selected, only emergency logs will be streamed.

NOTE: If you change the logging severity of the nodes through GigaVUE-OS CLI, then the same will not be reflected in GigaVUE-FM. Therefore, it is recommended to change the logging severity of the nodes through GigaVUE-FM.

Table 9: Logging Levels

Log-Level	Description
emergency	Emergency – the system is unusable. The severity level with the least logging – only emergency level events/commands are logged.
alert	Action must be taken immediately.
critical	Critical conditions.
error	Error conditions.
warning	Warning conditions.
notice	Normal but significant condition.
info	Informational messages.
debug	Debug-level messages. Authorized for factory use only.

Device Logging Processes

Logs that are specific to a process are logged with the process ID. Table 10: Logging Processes provides a list of logging processes and their descriptions.

Table 10: Logging Processes

Logging Process	Description
acctd	AAA Accounting daemon
avd	Active Visibility daemon
cli	Command Line Interface
clusterd	Cluster daemon
debuggabilityd	Debuggability daemon
frm	Foreign Resource Manager
gsd	GigaSMART daemon
gprof	Profiler
httpd	HTTP daemon
licd	License daemon
mgmtd	Management daemon

Logging Process	Description
ndiscd	Network Discovery daemon
netdevd	Netdev daemon
notf_mgr	Notification Manager
ntpd	Network Time Protocol daemon
peripd	Peripheral daemon
persistd	Persistence daemon
pm	Process Manager
ptpd	PTP Protocol daemon
restapid	REST API daemon
sched	Scheduler daemon
snmpd	SNMP daemon
statsd	Statistics daemon
syncd	Sync daemon
syshth	System Health
ugwd	Unified Gateway daemon
wizard	Wizard
wsmd	Web session Manager daemon
xd	XML Gateway
xinetd	Extended Internet Service daemon

View Device Logs

GigaVUE-FM provides the ability to view logs for each device/cluster node. These logs list all the user events and enable Gigamon Technical Support to troubleshoot in case of any problems.

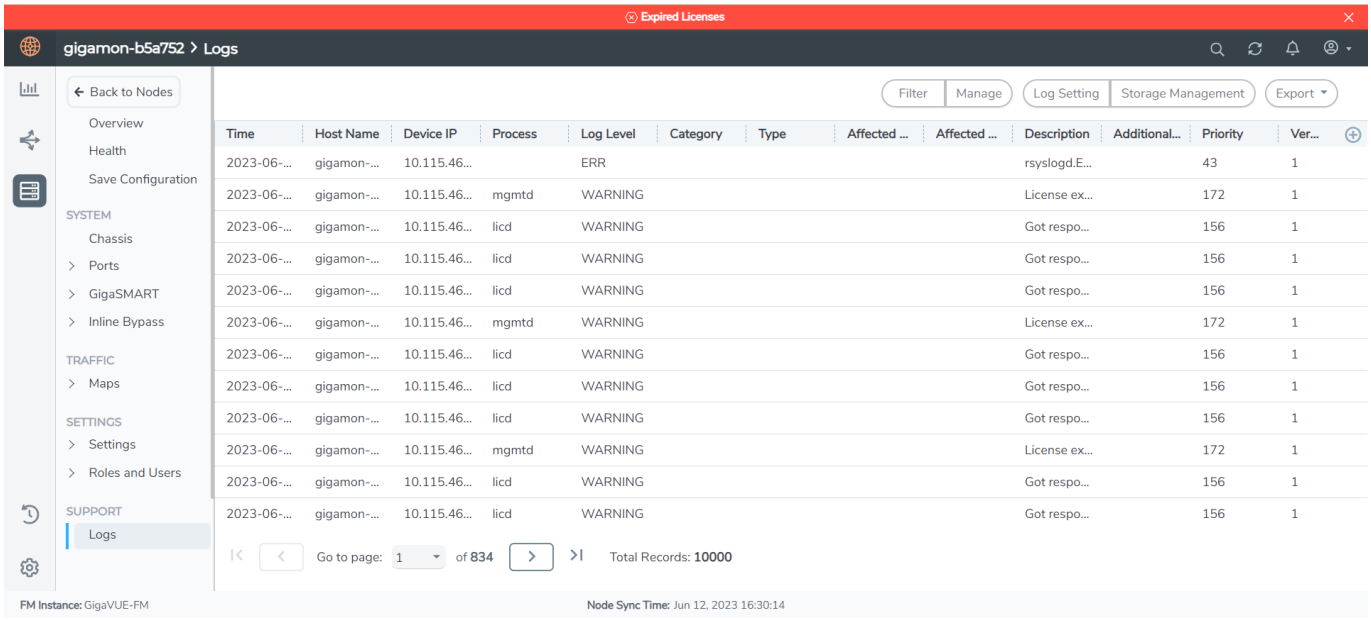
NOTE: Log streaming is supported on GigaVUE HC Series devices only.

In addition to describing how to view logs, this section describes the following supported functionality related to understanding and configuring the Logs view.

- Arrange Columns in the Logs View

To view the logs of a single node:

1. On the left navigation pane, go to **Inventory** and select **Nodes** which displays the list of physical nodes managed by GigaVUE-FM.
2. Select a node. The Single Node view displays overview statistics about the selected node.
3. In the left navigation pane, click **Logs**.



The Logs view displays the latest log information for the selected GigaVUE-OS node in table format.

1. (Optional) Click the heading cell for any column to sort the list by that attribute. Refer to Arrange Columns in the Logs View for details.
2. (Optional) Click **Filter** to refine the list of logs by any of the available attributes. The Filter quick view displays the attributes of the logs.
   a. Enter a value in any of the available attribute fields to narrow the list of logs according to that value.

Filterable Attributes	Description
Host Name	Enter a hostname or partial host name.
Device IP	Enter a device IP address or partial address.
Time	Specify a start date and time and an end date and time to define a time period.

Filterable Attributes	Description
	IMPORTANT: Timestamps are shown in the time zone of the client browser's computer and not the timezone of the node reporting the event. The timestamp is based on the correctly configured clock on the GigaVUE-FM server and converted from UTC to the client computer's configured timezone.
Process	Specify a process to target, such as clusterd, httpd, ntpd, restapid, snmpd, sshd, ugwd, or wsmd. Partial text entries are acceptable. Refer to Device Logging Processes
Log Level	Select one from the list of log levels, which are ranked by degree of severity. Refer to Device Logging Levels.
Category	Select a category of issue from the list of values, such as: High Temperature, Link, Packet Drop, or Resource Utilization. Refer to Device Log Categories.
Priority	Enter a priority number as defined by RFC 5424.
Version	Enter a version number.
Type	Select an event type from the list of values, such as CPU utilization high, memory utilization high, egress or ingress packet drop, and so on. Refer to Device Log Message Types.
Affected Entity Type	Select an affected entity type from the list of values. Options are: CPU, Memory, Port, or Optics
Affected Entity	Enter a specific affected entity

b. Click **Filter** to apply the filter on the current list of logs.

   The list of logs is now limited to the set matching your filter criteria.

c. To clear the filter, click **Clear** in the Filter dialog.

Arrange Columns in the Logs View

The Logs view has sortable, configurable columns. The options below describe how to use these options to personalize your view of the logs:

- Click the heading cell for any column to sort the list by that attribute.

- To move the order of a column, click the heading cell for that column and drag it to the location you prefer, then release the mouse click to drop it in the new location.
- Configure the columns to display in this view:
  a. Click **+** icon on the right above the table.
  b. The Column option appears.
  c. Amend your column settings:

     To add or remove a column, click the column label for the column to get selected and unselected accordingly.

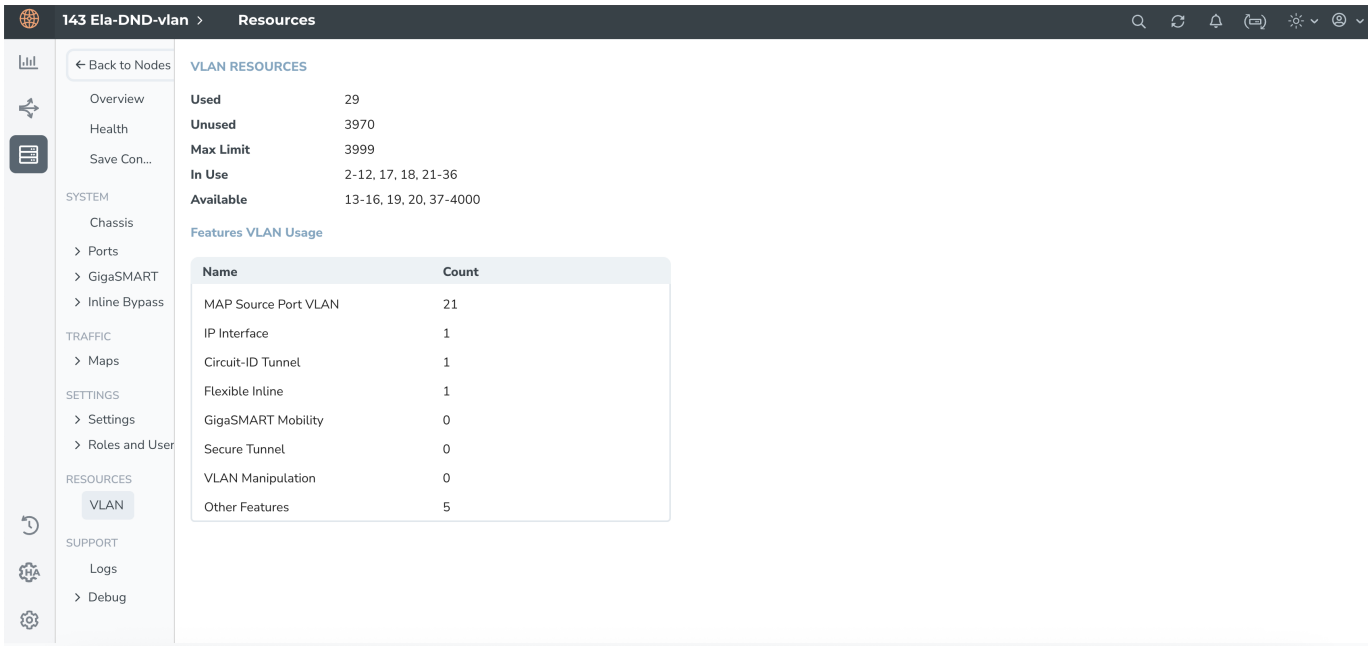
     To revert to the GigaVUE-FM default setting, click **Reset Columns to Defaults**.

VLAN Resources

GigaVUE-FM displays the VLAN resources that are used, and available. It also displays feature-wise VLAN usage.

To view the vlan resources:

1. On the left navigation pane, go to **Inventory** and select **Nodes** which displays the list of physical nodes managed by GigaVUE-FM.
2. Select a node.
3. In the left navigation pane, click **Resources > VLAN**.



The VLAN view presents the most recent information on VLAN resources for the selected GigaVUE-OS node, as illustrated above.
The following are the details of the different features and functionalities for which internal VLANs are allocated:
MAP Source Port VLAN: VLANs allocated to each network port of a flow map, the ports belonging to inline-network, inline-tool, and port-pair.
IP Interface VLAN: VLANs allocated for IP Interfaces.
Circuit-ID Tunnel VLAN: VLANs allocated for circuit-ID tunnels, which are used in fabric maps.
Flexible Inline VLANs: VLANs allocated for flexible inline solution deployments.
GigaSMART Mobility: VLANs allocated for GigaSMART mobility feature.
Secure Tunnel VLANs: VLANs allocated to the secure tunnels.
VLAN Manipulation VLANs: VLANs allocated for advanced VLAN manipulation feature.
Other Features: VLANs are allocated for GigaSMART operations on GigaVUE-HC1-Plus, GigaVUE-HCT, GigaVUE-TA25, and GigaVUE-TA25E platforms. This includes those assigned to network ports that are currently not part of any MAPs.
Refer to show vlan-resource.

Device Log Host Servers

All device/cluster log messages for all H Series nodes are streamed to the GigaVUE-FM server by default and are written on the device's file system in a messages file.

NOTE: Starting in software version 5.11.00, you can choose to have GigaVUE-FM as the syslog server or you can configure external syslog servers.

External third-party servers, like Splunk, can be added to host the streamed logs. You may want to have certain types of log messages streamed to different servers for different purposes. The Log Settings enable you to add additional host servers for the log streams and to configure the type of logs each server will host.

When an external server is specified, the GigaVUE® HC Series node will send logged events through UDP, TCP, or SSH to the specified destination.

In this section:

- Add an External Logging Host Server to a Node
- Edit Host Server Settings

NOTE: For a description of log levels, refer to Logging Levels. For a description of Syslog Server configuration options, refer to Host Server Options.

Add an External Logging Host Server to a Node

This topic describes how to configure a system log server as a destination for logging in GigaVUE-FM.

To add a host server:

1. On the left navigation pane, go to **Inventory** navigate to **Physical > Nodes** and select a node.
2. Click **Logs** from the Single Node view.
3. Click **Log Setting**.

   The Log Settings view appears.
4. Click **Add Server**.

   The Add Log Server quick view appears.
5. Specify the server attributes of the log server you wish to add. The parameters vary depending on the protocol you select.

   Select the logging protocol: **UDP**, **TCP (default)**, or **SSH**.

   For **UDP**, do the following:

   a. Enter the external server's IP address in the **Server Address** field.

   b. Enter the port number in the **Port Number** field.

   c. Select the required host from the **Hosts** drop-down.

   For **TCP**, do the following:

   a. Enter the external server's IP address in the **Server Address** field.

   b. Enter the port number in the **Port** field.

   c. Select the required host from the **Hosts** drop-down.

   For **SSH**, do the following:

   a. Enter the external server's IP address in the Server Address field**.**

   b. Enter the port number in the **Port Number** field.

   c. Enter the user name in the Username field.

   d. Select the required host from the **Hosts** drop-down.
6. Click **OK** to save the settings and add the server.

   The additional server appears in the server list.

7. Specify the optional settings for the server. Refer to Host Server Options for a description of the options.
8. Click **OK** to save the settings.

Host Server Options

The following table describes the optional host server settings for device logs.

Table 11: Syslog Server Settings

Setting	Description
Logging	This option enables or disable the server. ▪ Click **Enable** to enable the server. (Default) ▪ Click **Disabled** to disable the server. ▪ Click **Delete** to remove the server from this node.
Server Log Level	Select the level of alert you wish to tack on this server: ▪ Info ▪ Notice ▪ Warning (default) ▪ Error ▪ Alert ▪ Critical ▪ Emergency **Note**: Refer to Device Logging Levels.
Protocol	Select the protocol for communicating with this server: ▪ TCP ▪ UDP (default) **Note**: GigaVUE-FM can receive logs in TCP as well as UDP on port 5672. However by default it is UDP.
Port	Enter the Port Number.
Ssh Enabled	Check the **Ssh Enabled** check box to enable Ssh on this server.

Edit Host Server Settings

To edit log settings:

1. On the left navigation pane, go to **Inventory** and navigate to **Nodes** and select a node.
2. Click **Logs** from the Single Node view.

3. Click **Settings**.

   The Log Settings view appears.

   The top portion of the Logs Setting page shows the global Storage Management settings. The table in the lower portion lists the servers

4. Edit the settings for the server directly in the server list table. Refer to Host Server Options for a description of the options.
5. Click **OK** to save the settings.

Storage Management for Device Logs

Device/cluster log messages for HC Series nodes in GigaVUE-FM are continuously being recorded. As a result, the logs can take up a lot of storage space over time. You may want to delete old records on a regular basis to clear-up storage space. You may want to export log records as a back-up before performing a delete operation or to preserve for external analysis.

GigaVUE-FM Storage Management allows you to define how the stored logs are managed. You can specify a schedule for purging old device logs. You can also specify an SFTP server to export the log records prior to purging.

NOTE: GigaVUE-FM Storage Management is used for all storage settings, including device logs, alarm/event notifications, and statistics. This topic, however, describes the storage management for device logs, in particular.

Access Storage Management

There are two ways to access the Storage Management settings for device logs. These paths are described in Accessing Storage Management .

Table 12: Accessing Storage Management

Path	Steps
From  **:**	**1.** Click  **2.** Select **System** > **Storage Management**. Refer to the Storage Management section in the GigaVUE Administration Guide for

Path	Steps
	more details.
From the device logs page:	**1.** From the left navigation pane, go to **Inventory > Physical > Nodes**.. **2.** Select a node. **3.** Click **Logs** from the left navigation pane in the node view. **4.** Click **Settings**. **5.** The Storage Management page under GigaVUE-FM settings appears. **Note**: To perform a one-time manual clean-up of old logs on a specific node, refer to Manage Device Log Output.

The configurations performed from this page is applied globally across all the nodes managed by GigaVUE-FM.

Manage Device Log Output

Log management allows you to export and delete logs that are stored in GigaVUE-FM. Use this process to manually clean-up old logs on a specific node before a specified cut-off date and time. You can also specify an SFTP server and file path to export the records prior to purging.

NOTE: Refer to Storage Management for Device Logs for global Storage Management settings that control purging of old logs on a scheduled basis across all devices.

To edit the log management settings:

1. On the left navigation pane, go to **Inventory** and navigate to **Physical Nodes** and select a node.
2. Click **Logs** from the node view.
3. Click **Manage**.

   The Manage Logs view appears.
4. Complete the fields on this page to specify how to manage the logs:

Setting		Description
Time Range		
	Select records older than	Specify a cut-off date and time for deleting log records. Log records that were created prior to the specified date and time will be deleted. The time is based on the

Setting		Description
		current timezone in GigaVUE-FM.
Export Records To		
	SFTP Server Address	The records are exported to a CSV file. Specify the ftp/sftp location to send the CSV files. For example: **`sftp://username@121.0.0.1/path/directory`**
	Username and Password	If this is a secure server, which is recommended, specify the username and password for accessing the server.
	File Path	Specify the path on the server to store the file. For example: **/root/dir/archive.zip**
Purge Selected Records		
	Check this check box to enable purging the selected records. **Important**: If this option is selected, records will be deleted immediately when you click **OK**.	

5. Click **OK** to save the settings.

Traffic Filtering

This section describes the core intelligence provided by GigaVUE-OS and managed through GigaVUE-FM that can be used to optimize your network traffic flow.

Core intelligence is...

- Flow Mapping® that defines policies to extract flows of interest
- GigaStream load balancing that distributes flows across tools
- Terabit-scale configurations with Clustering and Fabric Maps
- Inline Bypass configurations that optimize threat prevention tools, enforcement point
- Visibility across physical, virtual and cloud infrastructure

You can use the GigaVUE-OS traffic management capabilities within GigaVUE-FM by accessing the device that has been added to the GigaVUE-FM fabric manager from the GigaVUE-FM interface. The Traffic option appears in the navigation pane of the device view on supported devices.

You can access traffic operations from the GigaVUE-FM interface.

Go to **Traffic** and select **Physical**. The following options are available:

- Orchestrated Flows
- Flow Maps
- Active Visibility
- Tunnel Monitoring

Topics:

- Ports and GigaStreams
- Flow Mapping®
- Flexible Inline Arrangements
- Inline Bypass Solutions
- Timestamps
- Fabric Maps
- Orchestrated Configurations

See also:

- GigaSMART®
- About Cluster

Ports and GigaStreams

This chapter provides the following information:

- About Ports
- Managing Ports
- Port Discovery
- Ingress and Egress VLAN
- How to Use GigaStream
- Port Statistics and Counters
- Monitor Port Utilization
- Packet Capture (PCAP)

About Ports

This section provides an overview of the various port types, describes the steps involved in configuring ports, and provides details about port filters and port status. This section includes the following major topics:

- About Network and Tool Ports
- Port Aliases
- Work with Hybrid Ports
- Port Filters
- Status of Line Cards/Nodes and Ports

About Network and Tool Ports

Packets arrive at the Gigamon Deep Observability Pipeline at **network ports** and are directed to monitoring and analysis tools connected to **tool ports** by flow maps. Figure 1GigaVUE-OS Packet Distribution illustrates the concept of data flows between network and tool ports. Data arrives from different sources at the network ports on the left and is forwarded to different tools connected to the tool ports on the right.

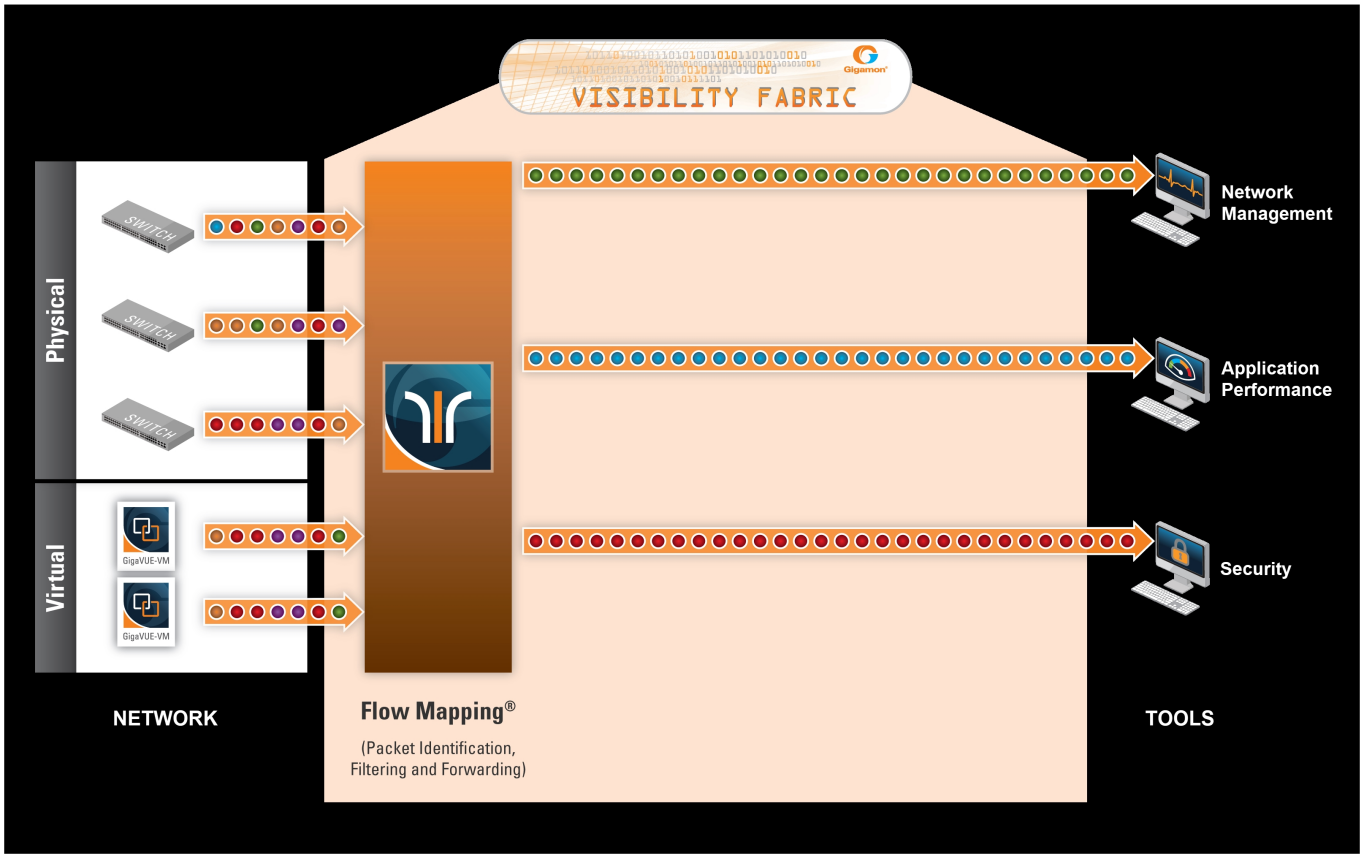


Figure 1 *GigaVUE-OS Packet Distribution*

Network (Ingress) Ports Defined

Network ports are where you connect data sources to GigaVUE nodes. For example, you could connect a switch's SPAN port, connect an external TAP, or simply connect an open port on a hub to an open port on a line card. Regardless, the idea is the same – network ports are where data arrives at the GigaVUE node.

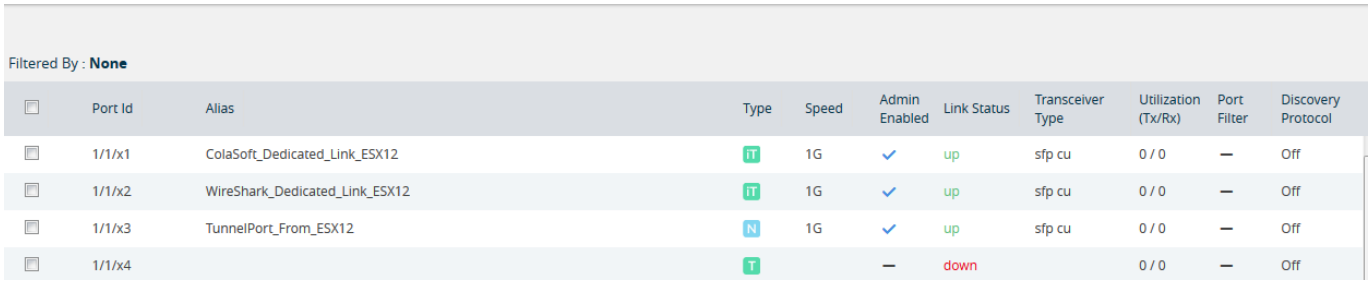
> **NOTE:** In their standard configuration, network ports only accept data input – no data output is allowed.

Tool (Egress) Ports Defined

Tool ports are where you connect destinations for the data arriving on network ports on GigaVUE nodes. For example, an IT organization could assign one set of tool ports to its Security Team for an intrusion detection system, a forensic data recorder, and a traditional protocol analyzer while a separate set of tool ports assigned to the Application Performance

Management team is used for a flow recorder and a long-term packet capture device. Regardless of the specific tool connected, the idea is the same – tool ports are where users select different portions of the data arriving on network ports.

NOTE: Tool ports only allow data output to a connected tool. Any data arriving at the tool port from an external source will be discarded. In addition, a tool port's **Link Status** must be **up** for packets to be sent out of the port. You can check a port's link status on the Ports page by selecting **Ports > Ports > All Ports** and looking at the Link Status field. Figure 2Port Link Status shows an example where the link status is up for ports 1/1/x1, 1/1/x2, and 1/1/x3 bu down for port 1/1/x4.



Filtered By : **None**

	Port Id	Alias	Type	Speed	Admin Enabled	Link Status	Transceiver Type	Utilization (Tx/Rx)	Port Filter	Discovery Protocol
☐	1/1/x1	ColaSoft_Dedicated_Link_ESX12	iT	1G	✓	up	sfp cu	0 / 0	–	Off
☐	1/1/x2	WireShark_Dedicated_Link_ESX12	iT	1G	✓	up	sfp cu	0 / 0	–	Off
☐	1/1/x3	TunnelPort_From_ESX12	N	1G	✓	up	sfp cu	0 / 0	–	Off
☐	1/1/x4		T		–	down		0 / 0	–	Off

Figure 2 *Port Link Status*

Ports on GigaVUE® TA Series Traffic Aggregator Nodes

In earlier software versions, GigaVUE® TA Series Traffic Aggregator nodes did not support tool ports. Instead, they supported gateway ports as displayed in Figure 3GigaVUE-TA Packet Distribution and described in Concepts Illustrated in Figure 3GigaVUE-TA Packet Distribution.

All gateway ports on GigaVUE® TA Series nodes are tool ports. For details, refer to Notes and Consideration on GigaVUE® TA Series Nodes.

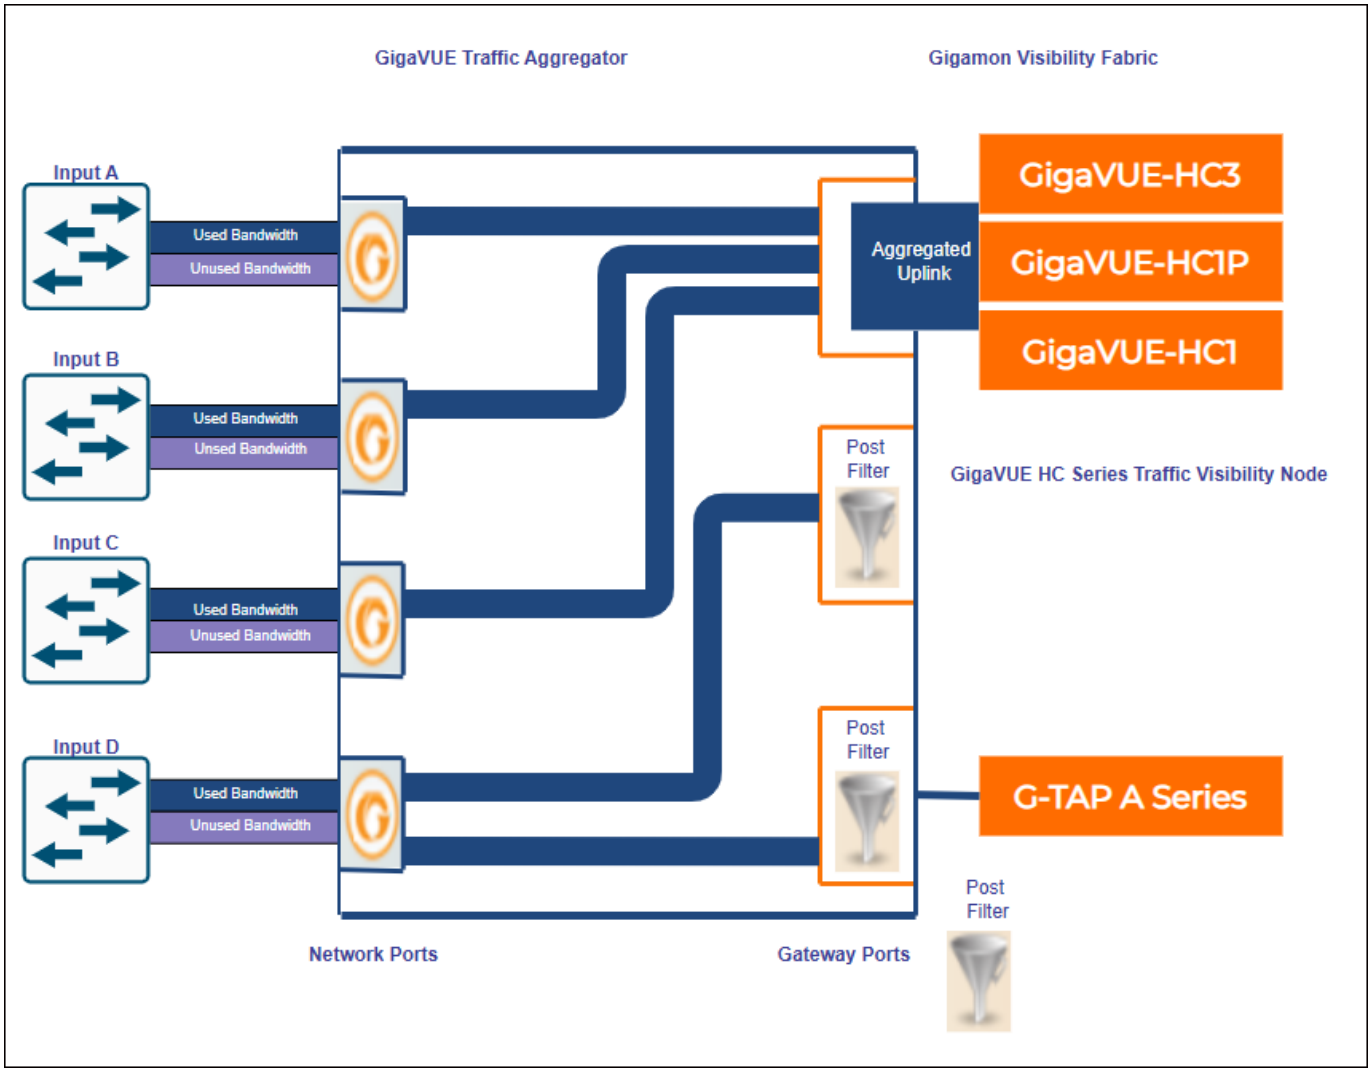


Figure 3 *GigaVUE-TA Packet Distribution*

Concepts Illustrated in Figure 3GigaVUE-TA Packet Distribution

Figure 3GigaVUE-TA Packet Distribution illustrates the concept of data flows. Data arrives from different sources at the network ports on the left and is forwarded to different Gigamon nodes connected to the tool ports (formerly gateway ports) on the right.

The following are important points about setting up packet distribution on GigaVUE® TA Series nodes:

- Traffic from multiple ingress ports can be sent to the same tool port for aggregated uplink to the Gigamon Platform fabric.

In this example, the traffic from Inputs A, B, and C is all sent to the same tool port. In turn, this tool port is connected to a GigaVUE® HC Series node so that the combined traffic from these inputs is available to the full suite of Flow Mapping® tools provided by the GigamonDeep Observability Pipeline.

- Traffic arriving at a single network port can be sent to multiple destination tool ports.

  Note that in Figure 3GigaVUE-TA Packet Distribution, the traffic arriving on **Input D** is sent to two different tool ports.

- Filters can be applied to tool ports:

  Filters applied to tool ports, inline network hybrid and circuit ports are called **egress-filters**. Egress-filters are useful if you want to send the same traffic to multiple tool ports and have each one allow or deny different packets based on specified criteria. You can use up to **20 egress-filters** at a time on GigaVUE® TA Series nodes.

NOTE: In Figure 3GigaVUE-TA Packet Distribution, egress-filters are set to focus on different parts of the data stream arriving at **Input D** – traffic on a VLAN range, a subnet range, and so on.

Notes and Consideration on GigaVUE® TA Series Nodes

GigaVUE® TA Series nodes support network, tool, stack, and hybrid port types. Refer to the following notes and considerations for GigaVUE® TA Series nodes (including GigaVUE-TA100, GigaVUE-TA200 and GigaVUE-TA400):

- Gateway ports on GigaVUE TA Series nodes are removed and converted to tool ports. In addition, gateway mirrors are removed and converted to tool mirrors.
- Tool ports on GigaVUE® TA Series nodes can continue to be used to aggregate traffic (as displayed in Figure 3GigaVUE-TA Packet Distribution and described in Concepts Illustrated in Figure 3GigaVUE-TA Packet Distribution).
- Tool ports on GigaVUE® TA Series nodes can also be used to directly connect to tools, such as firewalls, Intrusion Prevention Systems, or Application Performance Monitors.
- Hybrid ports are fully supported in both standalone and cluster mode on GigaVUE® TA Series nodes.
- GigaVUE® TA Series nodes can continue to be clustered with GigaVUE® HC Series nodes.
- When GigaVUE® TA Series nodes are in a cluster, bidirectional traffic flow is enabled on the stack links of GigaVUE® TA Series nodes.
- Map rules using GigaVUE® TA Series tool ports in the egress direction are supported.

Hybrid Ports

Hybrid ports are created by creating a dual function tool port. A physical tool port is set as a virtual network port which can then send traffic to other tool ports using secondary maps. A hybrid port is operated in loopback mode. This is only available if the GigaVUE H series node

is upgraded to minimum of 4.2 release. For more details on how to setup hybrid ports and the caveats, refer to the *GigaVUE-OS CLI Reference Guide*.

Stack Ports

Stack ports are used to carry traffic arriving at a network port on one GigaVUE node to a tool port on another GigaVUE node in a cluster.

Inline Network Ports

Inline networks, inline tools, and inline maps work together to form an Inline Bypass solution. The Inline Bypass solution has an overall state, which can change in response to hardware conditions and user configuration. Inline network ports are ports to which end-point devices are attached in an Inline Bypass solution.

NOTE: Inline network ports are supported only on GigaVUE-HC1, GigaVUE-HC1P, GigaVUE-HCT, GigaVUE-HC3, GigaVUE-TA25, GigaVUE-TA200, GigaVUE-TA400, and GigaVUE-TA400E.

Inline Tool Ports

Inline tool ports are ports to which inline tools are attached in an Inline Bypass solution.

NOTE: Inline tool ports are only supported on GigaVUE-HC1, GigaVUE-HC1P, GigaVUE-HCT, GigaVUE-HC3, GigaVUE-TA25, and GigaVUE-TA200, GigaVUE-TA400, and GigaVUE-TA400E.

Circuit Ports

Required License: Advanced Feature License for GigaVUE® TA Series Nodes

Circuit ports are used to send or receive traffic between two clusters. The circuit ports are configured at the sending and receiving ends of two clusters and the clusters are connected through a circuit tunnel. Circuit ports send or receive only the traffic that is tagged with a circuit-ID. In a map, if a circuit port is used as a source port, it acts as a network port, and decapsulates the traffic that contains a circuit-ID. If a circuit port is used as a destination port, it acts as a tool port, encapsulates the traffic, and strips the circuit-ID.

Circuit ports are supported on the following:

- All GigaVUE HC Series and TA Series nodes.
- As a source port in a regular map and as a destination port in a regular collector map.

- GigaStreams, port filter, and port groups.

GigaSMART Engine Ports

GigaSMART Engine ports are used when configuring GigaSMART groups. These ports cannot be edited. On the Ports page, the GigaSMART engine ports populates only the Port ID, Type, and Link Status fields.

Port Lists

Many map commands require a port-list (for example, rule and shared-collector arguments all require them). You can define the port lists using any combination of port IDs and port aliases. In GigaVUE-FM, port lists are created in the **Source** and **Destination** fields when editing or creating a new map. The following are considerations when creating a port list:

- When creating a **Pass All** map, you can specify a network port list, hybrid port list, or an inline network alias in the **Source** field. In the **Destination** field for a **Pass All** map, you can specify a tool port list, hybrid port list, an inline tool alias, an inline tool group alias, or an Inline Bypass.
- Circuit ports are supported as source ports on Regular maps and as destination ports on Regular Collector maps.
- The **Source** and **Destination** fields lets you select multiple non-contiguous ports. To enter port IDs in a list, simply select the port from the drop-down list after clicking in the field. If the port has a alias, it is shown in the list along with the ID.
- GigaSMART load balancing port groups can have ports with different rates.

Port Aliases

GigaVUE-FM lets you configure textual aliases for all port types. Aliases can be used in place of the numerical **bid/sid/pid** identifier required in many packet distribution.

To set up Port Aliases in GigaVUE-FM:

1. Select **Ports > Ports > All Ports**
2. Select the **Port ID** that needs an alias
3. Select **Edit**.
4. In the **Alias** field, enter an alias for the port, and the click **Save**.

Work with Hybrid Ports

A hybrid port is a physical port that has a dual function as an indirect traffic source port and a tool port. Hybrid means that a network port (ingress) can become a tool port (egress) to which map rules can be applied. Hybrid ports are introduced in software version 4.2. Hybrid ports are supported on GigaVUE HC Series nodes and GigaVUE TA Series devices.

A hybrid port is operated in loopback mode. The network data coming from the internal loopback is available to be used in maps.

Hybrid ports help alleviate the number of ports needed. For example, without hybrid ports, if you had traffic coming in with an MPLS header, but wanted to filter on a particular subnet, you would create a map to remove the MPLS header, physically loop the traffic back from the tool port to a new network port, and create another map to filter on the subnet. This same functionality can now be accomplished with hybrid ports.

If you have been using IP/UDP tunneling to encapsulate whole Ethernet frames and want to filter packets to destination tool ports after being decapsulated by GigaSMART, you can now use hybrid ports.

Hybrid ports can also be used to duplicate traffic from a network source. Using hybrid ports, you can create maps in parallel. For example, all HTTP traffic can be sent to one tool port unmodified and the same HTTP traffic can be sent to another tool port sliced at 100 bytes.

Using hybrid ports, you can create maps in a GigaSMART chaining.

As soon as a hybrid port is configured, it is internally changed to loopback mode. This means that the link is *Up* with or without SFPs inserted. (If SFPs are not inserted, the traffic runs at the maximum speed supported.) Traffic flows out of a hybrid port (Tx direction) and the duplicated flow loops back to it (Rx direction). This is similar to tool mirrors.

WARNING: Do not connect cables to hybrid ports coming from network ports. All cabling attached to hybrid ports must be attached to tools.

When a port is configured as a hybrid port type, it can be used as follows:

- as a map source and destination (for regular maps, as well as map-passall, and map-scollector)
- in a GigaStream
- in a port group
- with an egress port filter

Maps using hybrid ports, regardless of source or destination, can be applied to a GigaSMART operation.

When using hybrid ports, consider the following:

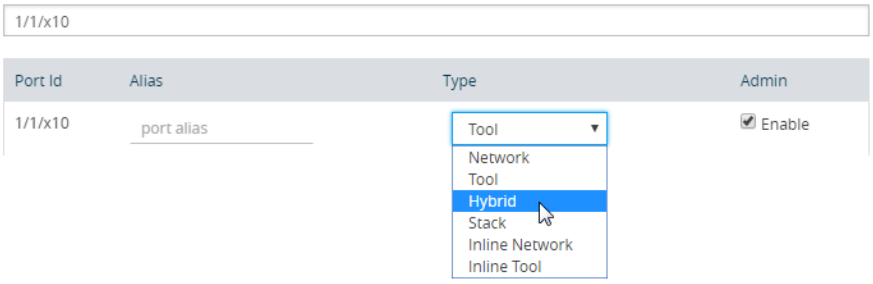
- Be aware not to configure traffic loops, such as such as H1 → H2, H2 → H3, H3 → H1.Do not use the same hybrid port as ingress as well as egress on all maps, such as such as H1 → H1.
- Once a hybrid port is used in a map or other traffic object, the port type cannot be changed.
- Hybrid ports can be used in inline OOB.
- Hybrid ports cannot be used in any inline objects, such as inline-network or inline-tool.
- Hybrid ports support ingress port VLAN tagging.
- Hybrid ports cannot be used in a tool mirror because that is for tool ports only.
- Hybrid ports cannot be used in a port pair because that is for network ports only.
- Hybrid ports are not supported on 100Gb ports with CFP2 transceivers.

In a cluster environment, hybrid ports can be configured across nodes.

To configure a hybrid port, do the following

1. Select **Ports > Ports > All Ports**.
2. Click **Quick Port Editor**.
3. Enter the port ID, the Quick Search field. For example, 1/1/x10.

- Click in the **Type** field and select **Hybrid** as shown in the following figure.



4. Click **Save**.

Hybrid ports can be used in the following:

- Regular map
- Regular map with GigaSMART operation
- First level and second level maps with vports

When configuring a map, use a hybrid port as follows:

- In the **Source** field when it is used as an indirect traffic source port
- In the **Destination** argument when it is used as a tool port

NOTE: You cannot use the same hybrid port in one map as both **Source** and **Destination**, or create a loop from multiple maps.

There is no limitation to the number of maps that can be used as second level maps to which packets can be forwarded.

Port Filters

Flow Mapping® provides the ability to apply filters to egress ports (tool, hybrid, circuit, and inline network), passing or dropping traffic after it has been forwarded from a network port.

Port-filters provide a convenient way to narrow down the traffic seen by egress ports without having to change an entire map. However, they are less efficient and scalable than flow maps – focus on using flow maps as your first packet distribution technique.

Port Filter—Rules and Notes

Keep in mind the following notes when managing port-filters:

- The **filter** is only supported for egress ports (tool, hybrid, circuit, and inline network) – network ports use maps to direct traffic.
- You can only configure egress port filters on a single port at a time. The **filter** argument is blocked when used the with multiple tool ports or port groups.
- In cases of inline network LAG and inline network groups, the port filters must be applied on each of the inline network ports that are part of the inline network LAG or inline network group.
- Port filters for inline network ports are supported on GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA400, and GigaVUE-TA400E.
- The GigaVUE-TA25 and GigaVUE-TA25E ports cannot be part of destination ports of first level maps if the source port is on another node (i.e combination of VPort and GigaVUE-TA25 and GigaVUE-TA25E destination port in the “to” ports list) in legacy cluster.
- The outer VLAN tool port filter cannot be used to match the ingress VLAN tag configured on the source port in GigaVUE-TA400 in 5.14 release.
- IP fragmentation tool port filter is not supported on GigaVUE-TA400 in 5.14 release.

- The following limitation is applicable only for double tag mode (software version 5.14.00 onwards).

  Egress port filters are supported on GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-HCT, and GigaVUE-HC1-Plus, except that

  a. VLAN rules are not supported with port filters, and

  b. Either IPv4 or IPv6 type port filter rules are supported only if L2 circuit encapsulation tunnels or GigaSMART maps are used else both IPv4 and IPv6 rules are supported.

Port-Filter Maximums

Table 1: Port-Filter Maximums per GigaVUE Node provides the maximum port-filters for the different GigaVUE nodes:

Table 1: Port-Filter Maximums per GigaVUE Node

GigaVUE Node	Maximum Number of Port-Filters
GigaVUE-HC1	• 448 for IPv4 rules • 255 for IPv6 rules • 448 for IPv4+IPv6 Pass rules. **NOTE:** For an IPv4 and IPv6 combination the maximum filters allowed is 448. In such combination the maximum limit is 254 for IPv4 filters and 255 for IPv6 filters. While configuring an IPv4 + IPv6 combination ensure that the individual filter limits are not crossed.
GigaVUE-HC3 (CCv1 and CCv2)	
GigaVUE-HC1-Plus	• 448 per Pseudo slot
GigaVUE-HCT	• 448 per chassis
GigaVUE-TA25	**Without Advanced Feature License:** • 20 per chassis **With Advanced Feature License** • 448 per chassis
GigaVUE-TA25E	**Without Advanced Feature License:** • 20 per chassis **With Advanced Feature License** • 448 per pseudo-slot
GigaVUE-TA100	**Without Advanced Feature License:** • 20 per chassis **With Advanced Feature License:** • 448 for IPv4 rules per pseudo-slot • 255 for IPv6 rules per pseudo-slot • 448 for IPv4+IPv6 Pass rules per pseudo-slot.

GigaVUE Node	Maximum Number of Port-Filters
GigaVUE-TA200	**NOTE:** For an IPv4 and IPv6 combination the maximum filters allowed is 448. In such combination the maximum limit is 254 for IPv4 filters and 255 for IPv6 filters. While configuring an IPv4 + IPv6 combination ensure that the individual filter limits are not crossed.
GigaVUE-TA200E	**Without Advanced Feature License:** • 20 per chassis **With Advanced Feature License:** • 448 per pseudo slot
GigaVUE-TA400	**Without Advanced Feature License:** • 20 per pseudo slots **With Advanced Feature License:** • 475 for IPv4 rules per pseudo slot • 475 for IPv6 rules per pseudo slot.
GigaVUE-TA400E	**Without Advanced Feature License:** • 20 per pseudo slots **With Advanced Feature License:** • 475 for IPv4 rules per pseudo slot • 475 for IPv6 rules per pseudo slot.

NOTE: A single filter applied to multiple tool ports counts multiple times against the 100-filter limit.

How to Apply Port Filters

To apply a port filter, do the following:

1. From the device view, go to **Ports > Ports > All Ports.**
2. Select the egress port (tool, hybrid, circuit, and inline network) to which you want to apply a filter, and then click **Edit**.
3. Under the Filters section on the Ports page, click **Add Rule**.
4. Select and configure any of the following required rule:

Table 2: Port-Filter Rule

Rule	Action
circuit-id	Configure circuit id
Description	Add a description to the Map Rule
dscp	Configure DiffServ Code Point bits

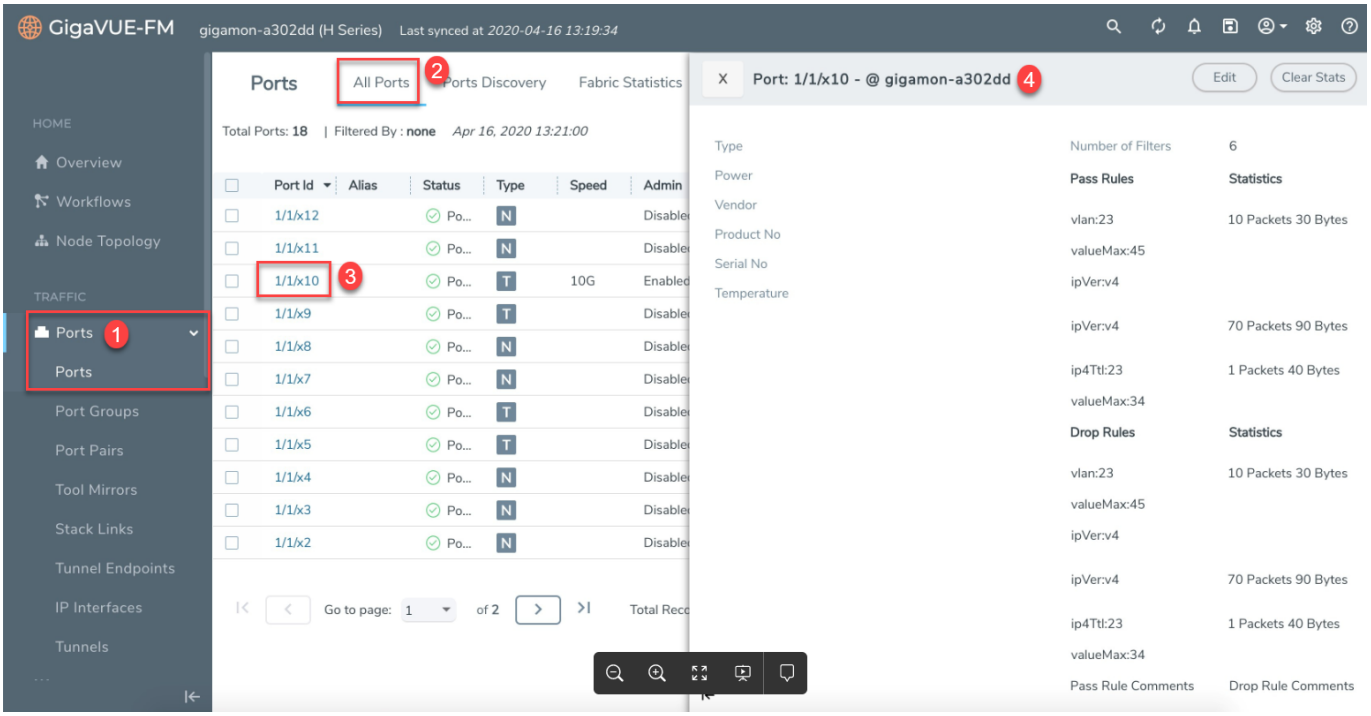
Rule	Action
ethertype	Configure Layer 2 ethernet type
ip6dst	Configure destination IPv6 address
ip6src	Configure source IPv6 address
ipdst	Configure destination IPv4 address
ipfrag	Configure IP fragmentation bits
ipsrc	Configure source IPv4 address
ipver	Configure IP version number
l2gre-id	Configure l2gre id
macdst	Configure destination MAC address
macsrc	Configure source MAC address
portdst	Configure destination port number or port range
portsrc	Configure source port number or port range
protocol	Configure internet protcol number
tcpctl	Configure TCP control bits
tosval	Configure type of service bits
ttl	Configure time to live value or range
vlan	Configure vlan id or id range
vxlan-id	Configure vxlan id

5. Add a new port-filter using the specified criteria as follows:
   - Use a **drop** rule to deny packets matching the specified criteria.
   - Use a **pass** rule to allow packets matching the specified criteria. All other packets are denied.
6. Click **Save**.

View Port Filter Statistics

You can view the port filter counters based on the filter rules configured for the port. To view the port filter statistics:

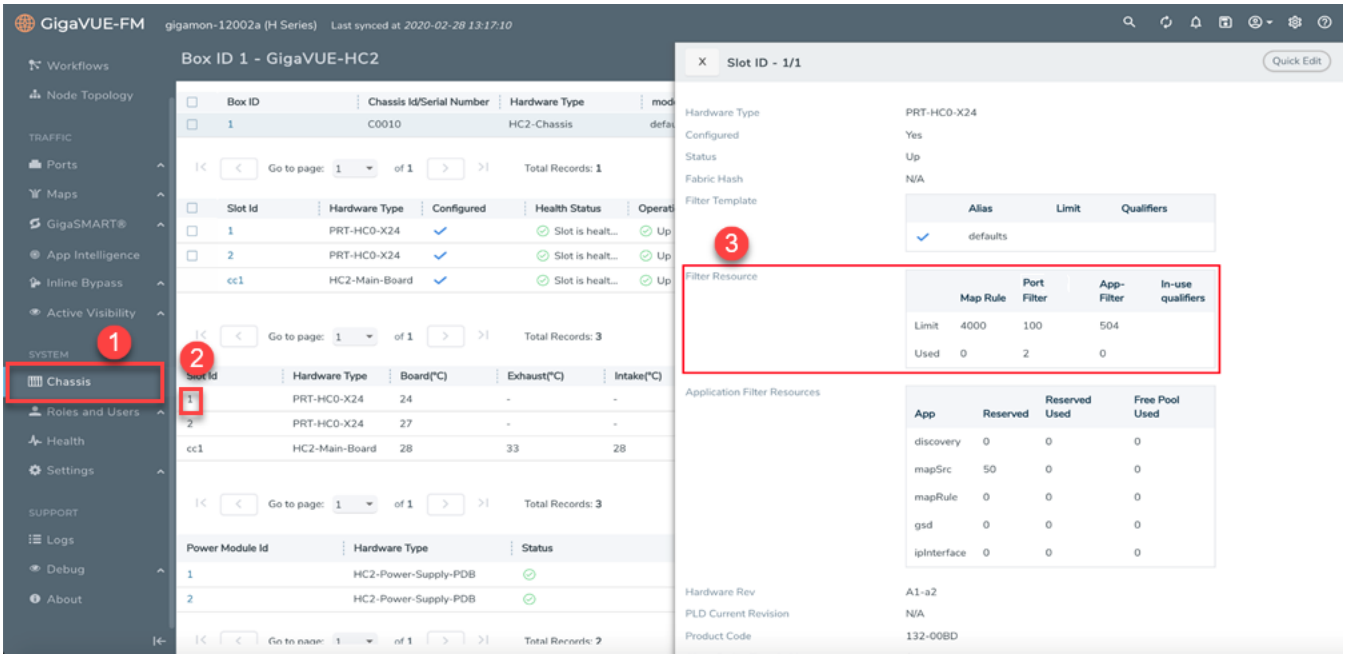
1. From the device view, go to **Ports > Ports > All Ports**.
2. Click the port ID for which you want to view the filter counters. The Port ID quick view appears. Refer to the following figure:



View Filter Resources for a Slot

You can view the maximum filter resources available and the filter resources used for a slot in the Slot ID quick view. To access the Slot ID quick view:

1. From the device view, go to **Chassis**. The Box ID page appears.
2. Click the required slot ID. The Slot ID quick view appears.
3. Go to the Filter Resource section to view the filter resources limit and the filter resources used. Refer to the following figure:



Status of Line Cards/Nodes and Ports

You can review the current status of the node, line cards, modules, and ports through either the CLI or the GigaVUE-FM. This will ensure that all units have been properly configured and that the node is ready for further configuration.

To check the line cards/modules and ports with GigaVUE-FM, using the Ports page or the Chassis page (select **Chassis** in the navigation pane).

How to Check Port Status with Ports Page

To check the port status with the Ports page, do the following:

1. Select **Ports > Ports > All Ports** to open the Ports page.
2. Locate the port to check by entering the port ID or port alias in the search field.
3. If you need to change the port type or enable the port:
   a. Click **Quick Port Editor**.
   b. In the Quick Port Editor, enter the port ID or port alias in the Quick search field to find the port.

c. Set the port type by selecting type from the drop-down list in the **Type** field. enable the by selecting **Enable** as needed.

How to Check Port Status with Chassis Page

To check the status of the ports and cards with the Chassis page, do the following:

1. Select **Chassis** from the navigation pane to open the Chassis view shown in Figure 4Chassis Page.
2. Use the view buttons on the Chassis page to check the status of the cards as well as the ports. When viewing a node in cluster, there is a drop down option to select a specific node in a cluster configuration.

For details about the Chassis page, refer to the *"Chassis"* section in the *GigaVUE Administration Guide*.

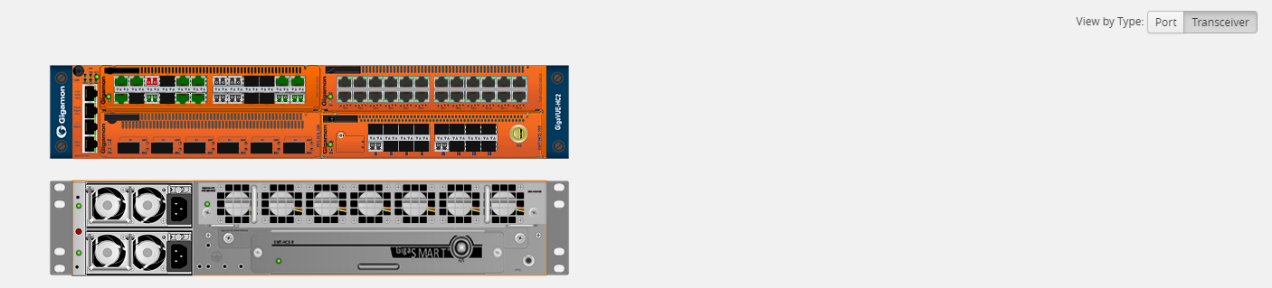


Figure 4 *Chassis Page*

Managing Ports

The Ports pages allows you to manage and configure ports for various functions. In case of inline supported devices, the pages for managing Inline Bypass ports are available by selecting Inline Bypass from the main navigation pane.

Before beginning with managing and configuring ports, make sure that the user role, which is assigned to a User Group, has the permission for the specific ports on the system. For details on the port-based access levels, refer to the *GigaVUE-OS CLI Reference Guide and "Managing Roles and Users"* in the *GigaVUE Administration Guide*.

This section provides a description of the Ports pages in the GigaVUE-FM UI. It covers the following topics:

- Ports
- Port Groups
- Port Pairs
- Tool Mirrors

- Stack Links
- IP Interfaces
- Circuit Tunnels

NOTE: Starting in software version 5.5.01, any change in the port health status is indicated immediately in the following port pages:

- Ports > All Ports
- Port Groups > All Port Groups
- Port Groups > GigaStream
- Port Pairs
- Tool Mirrors
- Stack Links
- IP Interfaces

The link in the notification opens the port quick view.

Ports

The Ports tab lets you select the All Ports and Ports Discovery pages. You can also control which ports display.

All Ports

The All Ports page displays when you select All Ports. The Ports page shows a table with detailed information about each port ID on a specific device. Only the GigaVUE HC Series and GigaVUE TA Series devices are presented in the Port Page view as shown in Table 3: Descriptions of Ports Page Columns. You can control which ports display on the page by selecting a set of filters or configure the ports through the Quick Port Editor or selecting Edit for a selected port. For details about filtering ports, refer to Port List Filter. For details about the Quick Port Editor, refer to Quick Port Editor.

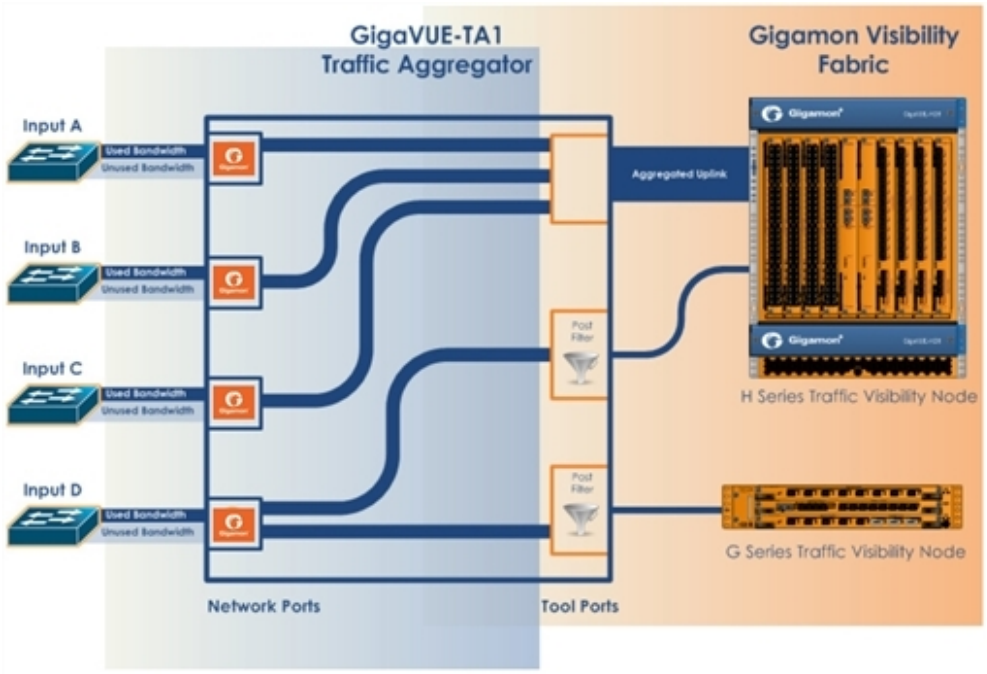


Figure 5 Ports Page

The port type determines which columns are populated with data in the table. The columns are populated as follows:

- **Engine** ports populate the Port ID, Type, and Link columns.
- **Network** ports populate the Port ID, Alias, Type, Speed, Admin Enabled, Link Status, Transceiver Type, Utilization, Port Filter, and Discovery Protocol.
- **Tool** port populate the Port ID, Alias, Type, Speed, Admin Enabled, Link Status, Transceiver Type, Utilization, Port Filter, and Discovery Protocol.
- **Stack** port populates Port ID, Alias, Type, Admin Enable, Link Status, Utilization, Port Filter, and Discovery Protocol.
- **Hybrid** port populates Port ID, Alias, Type, Admin Enable, Link Status, Utilization, Port Filter, and Discovery Protocol.
- **Circuit** port populates Port ID, Alias, Type, Admin Enable, Link Status, Utilization, Port Filter, and Discovery Protocol.
- **Inline Network** port Port ID, Alias, Type, Admin Enable, Link Status, Utilization, Port Filter, and Discovery Protocol.
- **Inline Tool** port Port ID, Alias, Type, Admin Enable, Link Status, Utilization, Port Filter, and Discovery Protocol.

NOTE: Not all port types are supported on all platforms. Inline network and inline tool ports are supported on GigaVUE HC Series nodes, GigaVUE-TA25, GigaVUE-TA25E and GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA400, and GigaVUE-TA400E.

Table 3: Descriptions of Ports Page Columns provides descriptions of the columns on the ports page.

Table 3: Descriptions of Ports Page Columns

Column	Description
Port ID	The port number is in <box ID>/<slot ID>/<port_ D> format in the CLI and GigaVUE-FM. For the GigaVUE TA Series slot ID is always 1 as they are not modular. For the GigaVUE-HC1, and GigaVUE-HC3, the line cards or modules are identified by the slot number. In a standalone (default) configuration, box ID is always designated as 1 but can be changed through the CLI (it cannot be changed through GigaVUE-FM). In a cluster configuration, the box ID can vary.
Alias	Alias name of the port, if any.
Status	Health status of the port. **Note**: When a ports administrative status changes from 'enabled' to 'disabled', the health status of the administratively disabled ports remains in Green. However, the health status of administratively disabled port is indicated as NA if **Exclusion Rules** are enabled for the ports that are admin disabled in the Alarms Page. Refer to the 'Manage Alarms' section in the GigaVUE Administration Guide for more details.
Type	List the type of port such as network, tool, stack, inline network, inline tool, circuit, or hybrid. You can set the port type through the Quick Port Editor or selecting the port on the Port page and clicking **Edit**. The port type is set by selecting the type in the **Type** field.
Speed	Current setting for the port's speed.
Admin	Indicates whether the port is administratively enabled or disabled.
Force Link Up	Indicates the 'force link up' setting for the port. When enabled, this option forces connection on the optical port.
Ude	Indicates whether the port is enabled for unidirectional (Ude) or bidirectional traffic. Enabled means unidirectional; Disabled means bidirectional. **Note**: Do not disable UDE on Out Ports when using BiDi RX/TX optics connecting towards passive TAPs. For example, if BiDi RX/TX optics are connected to both the Tap-facing and Network-facing ports of the TAP-M506/506T, disabling UDE on the TAP-facing ports will cause a production link failure.
FEC	Configures forward error correction (FEC) on the port to ensure error-free traffic over long distance. The values are: • **CL91**—Enables FEC on the port and supports 25Gb, 100Gb, and 4x25Gb speeds. • **CL74**— Enables FEC on the port and supports 25Gb, and 4x25Gb speeds. • **CL108** — Enables FEC on the port and supports 25Gb speed only on GigaVUE-TA25 and GigaVUE-TA25E. • **OFF**—Disables FEC on the port and supports 25Gb, 100Gb, and 4x25Gb speeds.
Link Status	The current status of the link connected to the port, either port link up or port link down.
Transceiver Type	The type of transceiver installed in this port.
SFP Power	SFP power for copper transceivers

Column	Description
	Note: When a new device is added to GigaVUE-FM, it takes one stats cycle for the SFP power value to be reflected in the GigaVUE-FM GUI.
Port Filter	Indicates if the egress port filter (tool, hybrid, circuit, and inline network), is associated with this port.
Discovery Protocol	Protocol used to discover neighboring nodes using CDP or LLDP. This feature is available for network, tool and circuit ports.
Box Hostname	Host name of the device
Gigamon Discovery	Indicates if Gigamon Discovery protocol is enabled
Tags	Tag associated with the port.

Port Quick View

The Quick View for ports displays when you click on a row in the Ports page to quickly get more information about a specific port. The quick view shows the port properties, statistics information on receiving (Rx) and/or Transmitting (Tx) ports and alarms information. The quick view also shows a graphical representation of port statistics. Refer to Figure 6Ports Quick View for an example.

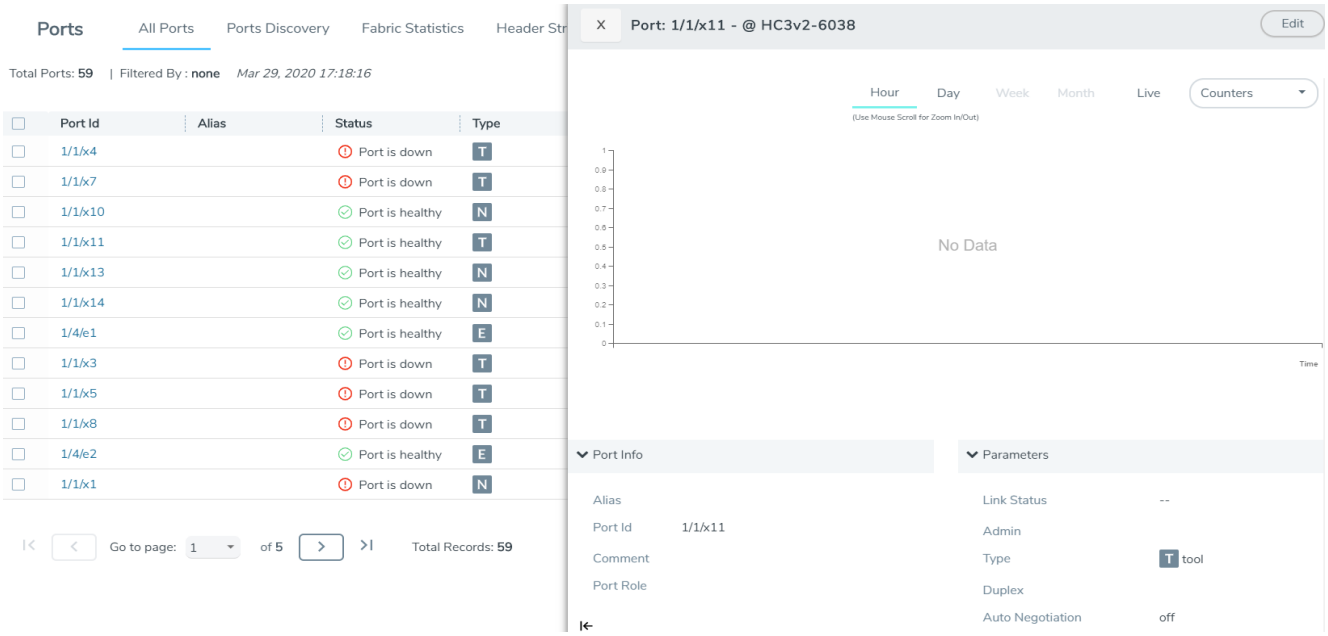


Figure 6 *Ports Quick View*

Port List Filter

The ports that display on the Ports page can be filtered so that only ports that meet certain criteria display on the page, such as port type and admin status. To filter the ports, select **Filter**. This opens the Filter view shown in Figure 7Port List Filter where you can specify how to filter ports displayed on the Ports page.

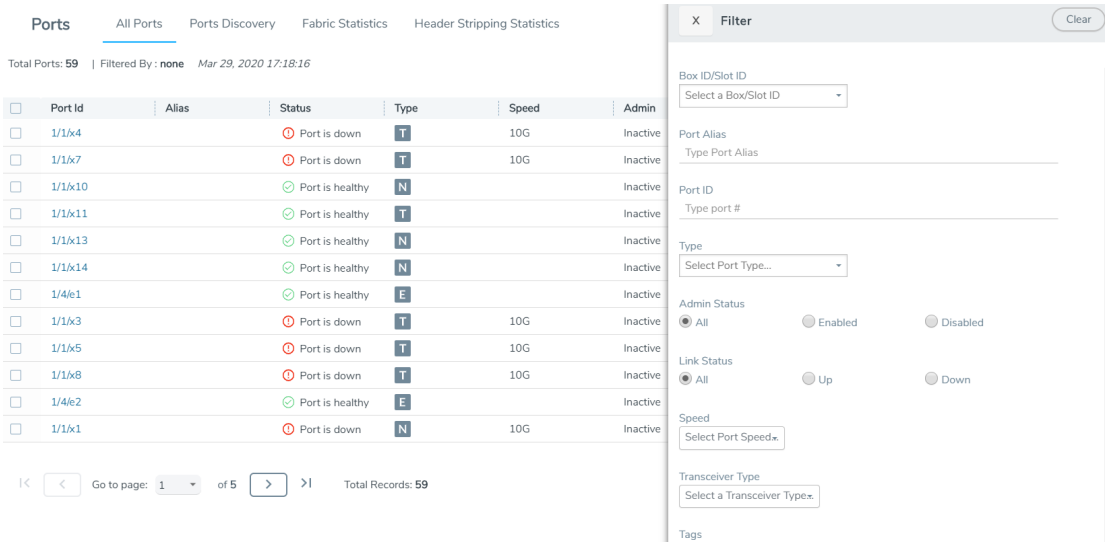


Figure 7 *Port List Filter*

The criteria that you can use to filter the port list is as follows:

Criteria	Description
Box/Slot ID	Display only those ports that match the specified box and slot IDs.
Port Alias	Display port with the specified alias.
Port ID	Display ports with specified number in the port ID. For example, if you specify 3 the result will also display ports that include the number 3, 13, 23, 30, and so on.
Type	Display ports with the specified port type. Select one of the following: • Network • Tool • Inline Network • Inline Tool • GigaSMART • Hybrid • Circuit • Stack
Admin Status	Display ports based on their current admin status. The possible selections are: **All** — display ports with a status of Enabled or Disabled. This is the default. **Enabled** — display ports with admin enabled **Disabled** — display ports with admin disabled
Link Status	Display ports based on their current link status: The possible selections are: **All** — display ports with a status of Up or Down. This is the default.

Criteria	Description
	Enabled — display ports with a link status of up. **Disabled** — display ports with a link status of down.
Speed	Display ports with the selected port speed. The port speeds available depend on the node.
Transceiver Type	Display ports with the selected transceiver type. The transceivers available selection depend on the type of transceivers connected to the ports.
Tags	Display ports associated with the selected tag key and tag value.

To filter the ports, enter the information to use for filtering the ports and select the radio buttons. For example, in Figure 8Filtering by Network Port Type and Admin Status Enabled, the filters selected are Network Type and Admin Status Enabled. Click the **Clear** button to remove the filter selections.

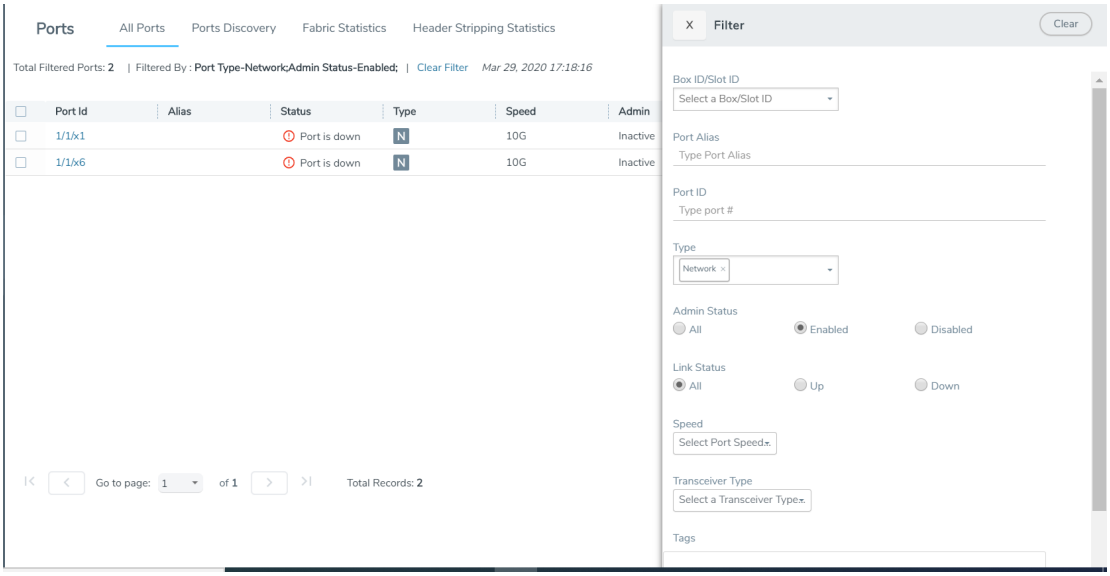


Figure 8 *Filtering by Network Port Type and Admin Status Enabled*

After the filter is applied, the Ports page displays only the ports that correspond to the selected filters and shows the total number of ports that meet the criteria. To clear the filters, select **Clear Filter**. Figure 9Filtered Ports List shows the Port pages with two ports that correspond to the current filters: Network Type and Admin Status Enabled.

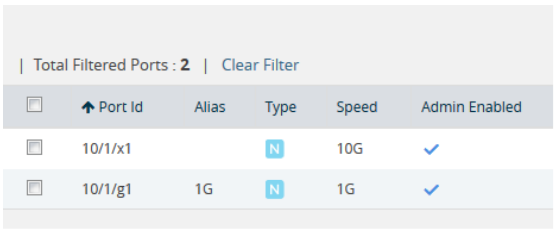


Figure 9 *Filtered Ports List*

Quick Port Editor

From the Ports page, you can open the Port Type Editor to quickly change the port types in a chassis. To set the port type for ports in a chassis, do the following:

1. Click **Quick Port Editor**.
2. For each port on which you want to set the port type, select the type from the drop-down list. In Figure 10 Port Type Selection, port 1/1/x9 is being changed from a network port to at tool port.

   To find a specific port, you can use the Quick Search to find a specific port by entering the port ID or alias in the Quick search field.

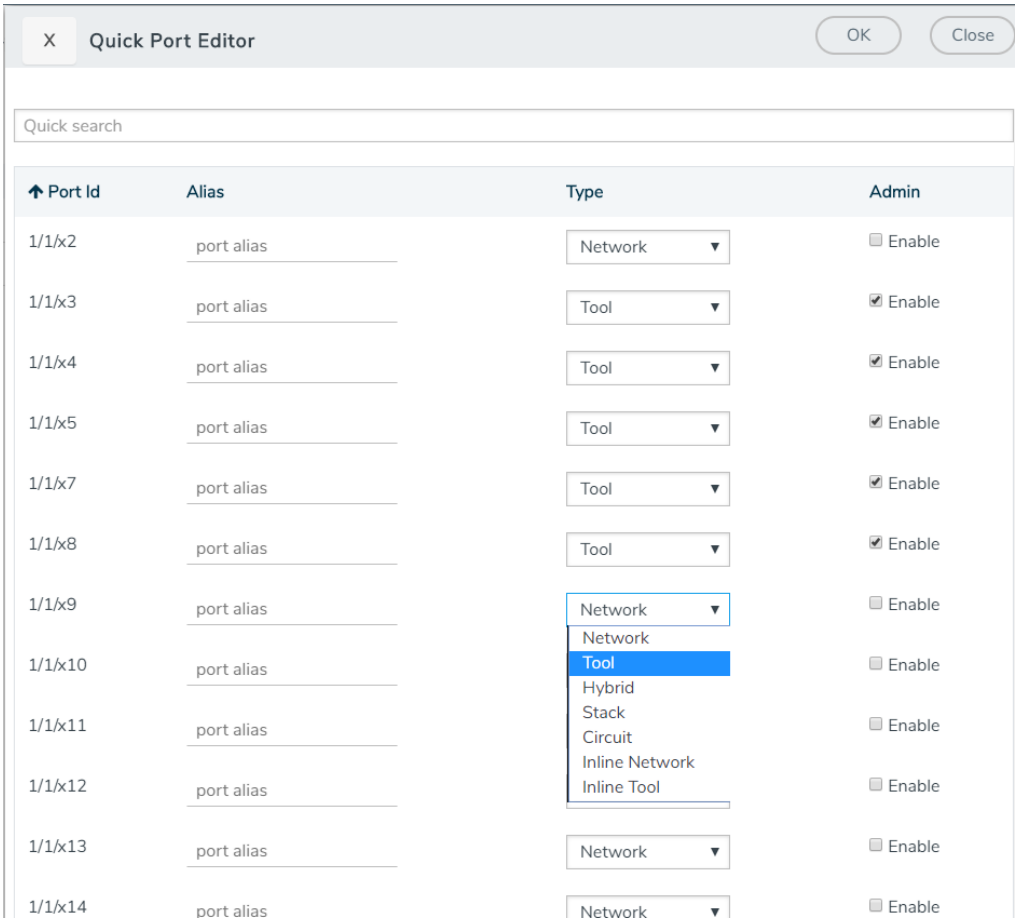


Figure 10 *Port Type Selection*

3. To enable the port, select **Enable**.
4. Click **OK**.

Each port can also be assigned an alias. Any port types set in the CLI or through the GigaVUE-FM APIs are reflected on this page. For more information, refer to Port Aliases for port aliases and to the *GigaVUE-FM REST API Reference in GigaVUE-FM User's Guide* and the *GigaVUE-FM Reference* for APIs.

Configure Ports

From the Ports page, you can either configure or edit a specific port by selecting a port and clicking the **Edit** button on the Ports page or on the Quick Port Editor. Table 4: Port Configuration Options describes the options on the configuration page.

Table 4: Port Configuration Options

Field	Description
Alias	The alias configured for this port, if any. Aliases can be used in place of the numerical bid/sid/pid identifier required in many packet distribution commands in the CLI. For example, instead of configuring a connection between, say, 1/1/x1 and 1/2/x4, you could connect Gb_In to Stream-to-Disk. Note that aliases can only be applied to single ports. They cannot be applied to groups of ports. Port alias can be up to 128 characters long including special characters. Aliases are case sensitive.
Admin	Check to enable the port.
Type	Specifies whether the port is configured as an Inline Network port, Inline Tool port, Network port, Tool port, Stack port, Circuit port, or Hybrid port.
Speed	Specifies the speed for the selected port. For copper ports, you can click to change the speed as long as Auto Negotiation is disabled.
Duplex	Specifies the port's duplex configuration. Only full duplex is supported. Starting in software version 5.2, half duplex support is removed from all GigaVUE nodes. If half duplex was configured in a previous software version, it will remain intact following the upgrade to 5.2 or higher release. Update to full duplex, if required.
Auto Negotiation	Select to enable auto-negotiation for the selected port. When auto-negotiation is enabled, duplex and speed settings are ignored. They are set through auto-negotiation. For 1Gb fiber ports, auto-negotiation is not supported on Gigamon Platforms..
Force Link Up	When enabled, this option forces connection on an optical port. Use this option when an optical GigaPORT tool port is connected to a legacy optical tool that does not transmit light. This option is not available for 10Gb capable ports with a 1Gb SFP installed.
Ude	When selected, this option indicates the port is unidirectional (UDE). When deselected (disabled), the port is bidirectional. UDE is enabled by default. Note: This option is available for GigaVUE-HC3, GigaVUE-TA100, GigaVUE-TA200, GigaVUE-TA25, GigaVUE-TA400, and GigaVUE-TA400E platforms with 100Gb BiDi (QSB-512) transceiver. If used with passive taps, ports used for monitoring should be set to Network port with UDE enabled. **Important:**If you clear the UDE check box, the laser will start to transmit, which may affect the remote connectivity.
FEC	Configures forward error correction (FEC) on the port to ensure error-free traffic over long distance. The values are: **CL91**—Enables FEC on the port. **CL74—** Enables FEC on the port .

Field	Description
	CL108 — Enables FEC on the port for **GigaVUE-TA25.** **OFF**—Disables FEC on the port. This option is available only on 25Gb and 100Gb transceivers.
Timestamp	Use the timestamp options when a GigaPORT-X12-TS line card is installed. For deta ils about the GigaPORT-X12-TS line card, refer to the *GigaVUE-OS CLI Reference Guide*. The timestamp options are as following: **Append Ingress**—Use this option to add a timestamp to ingress packets a GigaPORT-X12-TS ports. This applies to ports x1..x12 when configured as network ports. **Strip Egress**—Use this option to strip timestamps from egress packets. **Source ID Egress**—Use this option to specify a custom source ID to be included in the timestamp appended by the GigaPORT-X12-TS. The source ID identifies the ingress port on the GigaVUE HC Series node where the timestamped packet arrived. The timestamp always includes a source ID. If you do not specify a custom value, the GigaPORT-X12-TS generates one automatically using the following formula: **(Box ID * 2048) + (Slot ID * 256) + Port Number** **Important:** Only apply the Strip Egress option to packets with time stamps appended. The strip egress feature strips the last 14 bytes of each packet regardless of whether a timestamp has been added.
VLAN Tag	Use VLAN tags to identify, differentiate, or track incoming sources of traffic. When the traffic reaches the tools or the maps, you can filter on the VLAN tags for the corresponding ports you want to measure. The port must be a network or inline-network or hybrid type of port. Ingress port VLAN tagging is supported for IPv4 and IPv6 packet types, including non-tagged packets, tagged packets, and Q-in-Q packets. Ingress port VLAN tagging is not supported on inline network ports, hybrid ports, or on network ports that are connected via port-pairs. The same VLAN tag can be assigned to multiple network ports. However, each port can only have one VLAN tag. VLAN tagging is supported in a cluster. • To add, VLAN IDs for a Port, enter the VLAN ID in this field. • To modify, update the VLAN ID in this field and Save. It will take effect. • To delete, remove any values for the VLAN ID in this field and Save. It will remove the VLAN Tag from this Port. VLAN tags are only available on network ports. **Note**: On GigaVUE-TA25 and GigaVUE-TA25E, ingress VLAN tagging is not supported on network ports associated with GSOP maps, and GSOP maps are not supported on network ports configured with ingress VLAN tag. This limitation is not applicable to the E-tag mode.
TPID	Select the **TPID** for the VLAN Tag. The default value of TPID is 0x8100. You can select the other supported values 0x9100 and 0x88a8 from the drop-down list.
Port Discovery	Select to enable discovery of neighbors associated with the port. Neighbor discovery is only available on network ports
Discovery Protocols	When port discovery is enabled, use the Discovery Protocol options to set up CDP or LLDP or both (All) on the port. The results are shown on the Ports Discovery page.
Buffer Threshold	Specifies the alarm buffer threshold on a port. You can specify the alarm buffer

Field	Description
	threshold in the Rx and Tx directions on network and stack type ports and in the Tx direction on tool type ports. By default, the threshold is set to **0**, which disables the threshold
Utilization Threshold	Sets the utilization percentage for this port at which the GigaVUE HC Series node will generate high or low utilization alarms for the port. For more information about port utilization, refer to Monitor Port Utilization. By default, the threshold is set to **0**, which disables the threshold.
Lock Port	Restricts use of the port for only your user account as follows: • Users with the admin role can lock any port in the system. Users with the Default/Operator role assigned can only lock ports to which their account has been granted access. • Administrators can lock a port for another user by including the optional user. You can optionally share a locked port by specifying users in the **Lock shared with Users** field or selecting users to share the lock with through their assigned roles. For more information about who to set lock sharing, refer to Managing Ports.
Tags	Select the required tag key and tag value to which the port must be associated to. The tag key and the tag value will be displayed depending on the role and the corresponding access rights of the user.

Ports Discovery

The Ports Discovery page displays the port neighbor information for each port that has discovery enabled. For each network port, tool port or circuit port on which discovery is enabled, neighbor information is collected. Information for up to five of the most recent neighbors is retained for each port.

The following are limits on the amount of discovery information that is retained:

- For each port, discovery information for a minimum of two neighbors and a maximum of 20 neighbors is retained.
- For a chassis, discovery information for a maximum of 2K neighbors is retained.

Neighbor information is removed or replaced as follows:

- When the neighbor information expires due to the TTL.
- When the number of neighbors for the chassis reaches the 2K maximum and a new neighbor is discovered. In this case, the following can occur:
  - If there are currently two or more discovered neighbors for a port, the newly discovered neighbor replaces the neighbor information for the least recently updated neighbor.
  - If there are currently less than two discovered neighbors for a port, the newly discovered neighbor is added (actually exceeding the 2K limit to guarantee a minimum of two neighbors per port).

NOTE: Aging (the discovery protocol time-to-live) determines how long neighbor information is valid.

For information about the discovery protocols and enabling port discovery, refer to Port Discovery

Statistics

The Statistics page displays the statics for all the ports on the node, providing the following information about a packets transmitted or received on a port:

Column	Definition	Notes
Octets (Rx/Tx)	The count of packets/bytes received and transmitted by this port.	Error packets are not transmitted, therefore they are not counted. Excludes undersize frames.
Octets/sec (Rx/Tx)	The count of packets/bytes received and transmitted by this port per sec.	Error packets are not transmitted, therefore they are not counted.
Unicast Packets (Rx,/Tx)	The count of packets/bytes received and transmitted by this port.	Excludes multicast packets, broadcast packets, packets with FCS/CRC errors, MTU exceeded errors, oversize packets, and pause packets.
Non-Unicast Packets (Rx/Tx)	Total Non-unicast packets received or transmitted.	
Packets/sec (Rx/Tx)	The rate which packets are received or transmitted.	
Packet Drops (Rx)	Total Dropped Packets	Packets are dropped when a network port's bandwidth is exceeded due to oversubscription. Packets are dropped when they reach the port but before they are sent out.

Column	Definition	Notes
Discards (Rx/Tx)	Total received and transmitted packets discarded. This counter increments when a packet is discarded at the tool port due to egress port filter.	Discards are counted in the following cases: Traffic arriving at a network port that is not logically connected using a map or map passall. Map rules applied on a network port. In packets on a tool port. Pause frames.
Errors (RX/TX)	Total Error Packets Received or Transmitted. Error packets include undersize, FCS/CRC, MTU exceeded, fragments, and oversize packets.	Excludes oversize packets without FCS/CRC. Packets larger than the MTU setting arriving on a network port are counted twice in the counter. So 1000 oversize packets would show up as 2000. This double-counting only happens with Oversize error packets.
Utilization (Rx/Tx)	Percentage of port utilization by packets received or transmitted	

Port Groups

The Port Groups selection in the top navigation bar provides access to the All Port Groups and GigaStream pages, for creating port groups and GigaStreams, respectively.

All Port Groups

Selecting **Port Groups** under **Port Groups** opens the **Ports Group** page by default. This page is used to create a port group, which can simplify administration of GigaVUE ports. Administrators can create groups of ports that can then be quickly assigned to different user groups. With clustered ports potentially numbering into the hundreds, port groups provide a useful shorthand when assigning multiple ports to different user groups. (To create user groups, select **Roles and User** from the navigation pane, and refer to *Managing Roles and Users*" in the *GigaVUE Administration Guide* for more details.) The following are the different types of port groups:

- **Network Port Group**—contains only network ports.
- **Tool Port Group**—contains only tool ports or tool GigaStream, which is a combination of multiple tool ports.
- **Hybrid Port Group**—contains only hybrid ports or hybrid GigaStream, which is a combination of multiple hybrid ports.
- **Circuit Port Group**—contains only circuit ports or circuit GigaStream, which is a combination of multiple circuit ports.

- **Load Balancing Port Group**—contains tool ports for load balancing. The maximum number ports allowed for port load balancing is 16.

However, port groups that include GigaStream can only be used with GTP Overlap Flow Sampling maps. For more information about GTP Flow Sampling, GTP Whitelisting and GTP Overlap Flow Sampling maps, refer to GigaSMART GTP Whitelisting and GTP Flow Sampling.

Create Port Groups

To create a Port Group, do the following:

1. Select **Ports > Port Groups > Port Groups**.
2. Click **New**.

   The Port Group configuration page appears.
3. Configure the Port Group:
   a. Enter an alias in the **Alias** field.
   b. (Optional) Enter a description in the **Description** field.
   c. Select the Port type from Network , Tool , Hybrid or Circuit.
   d. Click in the **Ports** field and select the ports for this port group .

   NOTE:
   - Use only ports from the same map when you create a Network Port Group. The port group won't work as expected if you select ports from different maps. Even when you use ports from the same map, ensure you use the exact set of ports defined in the map. Do not modify the port set when creating the port group.
   - The port should not be part of the IP interface.
   - Before upgrading to the latest version, you must remove either the port group or the IP interface if the ports are identical.

   e. (Optional) Click in the **Tags** field and select the tags for this port group.
4. Click **Apply**.

The Port Group created is added to the list view.

Edit Port Group

You can edit the optional fields in the Port Group such as Description, Ports, Tags, Tunnel endpoints etc. To edit, do the following:

1. Select **Ports > Port Group**. The Port Group page displays .

2. Select a Port Group and then click **Actions > Edit.**

   NOTE: To edit a port group that's part of a map, first remove it from the map. You can't make changes to the port group while it's associated with a map.

3. After editing your configurations ,click **Apply.**

Clone Port Group

In some cases, you may want to create a Port Group that is similar to an existing one. To do this use the Clone feature.

1. Select **Ports > Port Group**.
2. Select the Port Group that you want to copy, and the click **Actions > Clone**.
3. Enter a new alias in the **Alias** field.
4. (Optional) Add or update Description in the **Description** field.
5. (Optional) Click in the **Ports** field and select the ports for this port group .
6. (Optional) Click in the **Tags** field and select the tags for this port group.
7. Click **Apply**.

Use the following buttons to manage the Port Groups.

Button	Description
Tags	Use to associate tags to Port Groups and to remove the tags from Port Groups when not required. To add tags: **1.** Select the required Port Groups to which the tags must be applied. **2.** Click the **Tags** drop-down menu button and select **Add**. **3.** In the Add Tags to Resources page, select the required Tag Keys and Tag Values. **4.** Click **Ok** to save the configuration. The Port Group is associated with the selected tags. To remove tags: **1.** Select the required Port Groups that are already associated with tags. **2.** Click the **Tags** drop-down menu button and select **Delete**. **3.** In the Delete Tags from Resources page, select the required Tag Keys and Tag Values that must be removed from Port Groups. **4.** Click **Ok** to save the configuration.
New	Use to create a new Port Group. Refer to Create Port Groups for

Button	Description
	detailed information.
Actions	Use the **Actions** drop-down button to perform the following options: • **Clone**: Use to clone the selected Port Group. Refer to Clone Port Group for details. • **Edit**: Use to edit the selected Port Group. Refer to Edit Port Group for details. • **Delete**: Use to delete the selected Port Group.
Export	Use to export the Port Groups in CSV or XLSX format. The following options are available: • **Export All**: All Port Groups in the list view are exported. • **Export Selected**: Only the selected Port Groups are exported.

Port Group Statistics

To view the statistical details of the port group members:

1. Select **Ports > Port Groups > Statistics**. The Port Group Statistics page is displayed. The following fields are displayed:

- Port Group Members
- Total Bytes
- Total Packets
- Total Sessions
- Active Sessions

2. Use the **Actions** drop-down button to perform the following:

- **Clear**: Select a Port Group Member and click **Clear** to clear the port group members statistics counters.
- **Clear All**: Click **Clear All** to clear the port group member statistics counters of all the port group members.

3. Use the **Export** button to export the statistical details.

Port Pairs

A port-pair is a bidirectional connection in which traffic arriving on one port in the pair is transmitted out the other (and vice-versa) as a passthrough TAP. Keep in mind the following rules and notes for port-pairs:

- You can configure whether a port-pair uses link status propagation. Link port propagation does the following:
  - Enabled—when one port in the pair goes down, the other port goes down.
  - Disabled—when one port in the pair goes down, the other port is unaffected.
- Port-pairs can be established between ports using different speeds. For example, from a 100Mb port to a 1Gb port. However, the system will warn you when creating such port-pairs. Depending on traffic volume, port-pairs between ports using different speeds can cause packet loss when going from a faster port to a slower port. For example, going from 1Gb to 100Mb, from 10Gb to 1Gb, and so on.

Create Port Pair

To configure a port pair, do the following:

1. Select **Ports > Port Pairs**.
2. Click **New**.
3. On the Port Pair page, do the following:
   a. (Optional) Type an alias in the **Alias** field to help identify this port pair.
   b. (Optional) Type a description in the **Description** field.
   c. Click in the **First Port** field and select a network port.
   d. Click in the **Second Port** field and select another network port.
   e. (Optional) Enable **Link Failure Propagation**.

   Port pairs can operate with or without line failure propagation (LFP) as follows:
   - With LFP enabled, link failure on one of the ports in the port pair automatically brings down the opposite side of the port pair.
   - With LFP disabled, the opposite port is not brought down automatically.

   Note: A port pair created on a copper TAP has LFP enabled by default.
4. Click **Apply**.

Use the following buttons to manage the Port Pairs.

Button	Description
Tags	Use to associate tags to Port Pairs and to remove the tags from Port Pairs when not required. To add tags: **1.** Select the required Port Pairs to which the tags must be applied. **2.** Click the **Tags** drop-down menu button and select **Add**.

Button	Description
	3. In the Add Tags to Resources page, select the required **Tag Keys** and **Tag Values**. 4. Click **Ok** to save the configuration. The Port Pair is associated with the selected tags. To remove tags: 1. Select the required Port Pairs that are already associated with tags. 2. Click the **Tags** drop-down menu button and select **Delete**. 3. In the Delete Tags from Resources page, select the required Tag Keys and Tag Values that must be removed from Port Pairs. 4. Click **Ok** to save the configuration.
New	Use to create a new Port Pair. Refer to Create Port Pair for detailed information.
Actions	Use the **Actions** drop-down button to perform the following options: • **Edit**: Use to edit the selected Port Pair. • **Delete**: Use to delete the selected Port Pair.
Export	Use to export the Port Pairs in CSV or XLSX format. The following options are available: • **Export All**: All Port Pairs in the list view are exported. • **Export Selected**: Only the selected Port Pairs are exported.

Tool Mirrors

In addition to maps, the GigaVUE-OS also includes a special Tool Mirror packet distribution feature. A Tool Mirror can be used to send all packets on one tool port to another tool port (or multiple tool ports) or GigaStream on the same box. Tool Mirrors can still be applied to network ports even if they are already in use with an existing connection or map. Use tool-mirror connections between tool ports/GigaStreams on the same node, cross-box tool-mirror connections are not supported.

Tool-mirror can be created from:

- Tool port to tool port or ports on the same node.
- Tool port to GigaStream or GigaStreams on the same node.

  The destination for a tool-mirror must always be either a tool port or a GigaStream.
- Tool Mirrors can cross line cards/modules – they can start on one line card and end on another in the same node. However, they cannot cross nodes in a cluster.
- Tool Mirrors on GigaVUE-HC1 can be created on tool ports or GigaStream ports.
- Tool Mirrors are not allowed from Tool GigaStream to tool port.

- Tool Mirrors are not supported on tool ports with copper SFPs installed.

Create Tool Mirror

To create a Tool Mirror, do the following:

1. Select **Ports > Tool Mirrors**.

   The Tool Mirrors page displays a list of the currently configured Tool Mirrors.
2. Click **New**.

   The Tool Mirror configuration page appears.
3. Configure the Tool Mirror:
   a. Enter an alias in the **Alias** field.
   b. (Optional) Enter a description in the **Description** field.
   c. Click in the **Source Tool Ports** field and select the source tool ports for this tool mirror.
   d. Click in the **Mirror Destination Ports** field and select the destination tool ports for this tool mirror.
4. Click **Apply**.

Edit Tool Mirror Description

Description in the Tool Mirror configuration is optional. However, you can add description at any time, or edit the existing description. To add or edit description, do the following:

1. Select **Ports > Tool Mirrors**. The Tool Mirrors page displays a list of the currently configured Tool Mirrors.
2. Select a Tool Mirror in the list of Tool Ports, and then click **Actions>Edit.**
3. Enter or change a description in the **Description** field. (You cannot make any other changes to the Tool Mirror.)
4. Click **Apply.**

Clone Tool Mirror

In some cases, you may want to create a Tool Mirror that is similar to an existing one. To do this use the Clone feature.

1. Select **Ports > Tool Mirrors**.
2. Select the Tool Mirror that you want to copy, and the click **Actions>Clone**.
3. Enter a new alias in the **Alias** field.

4. (Optional) Add or update Description in the **Description** field.
5. Make change to the **Source Tool Ports** and **Mirrored Destination Ports** as needed.
6. Click **Apply**.

Use the following buttons to manage the Tool Mirrors.

Button	Description
New	Use to create a new Tool Mirror. Refer to Create Tool Mirror for detailed information.
Actions	Use the **Actions** drop-down button to perform the following options: • **Clone**: Use to clone the selected Tool Mirror. Refer to Clone Tool Mirror. • **Edit**: Use to edit the selected Tool Mirror. Refer to Edit Tool Mirror Description. • **Delete**: Use to delete the selected Tool Mirror.
Export	Use to export the Tool Mirrors in CSV or XLSX format. The following options are available: • **Export All**: All Tool Mirrors in the list view are exported. • **Export Selected**: Only the selected Tool Mirrors are exported.

Stack Links

Use stack-links to connect multiple GigaVUE nodes in a unified cluster. The stack-links carry traffic entering one system and bound for another via a map. Stack management traffic uses its own dedicated network connections through the Stacking ports on the Control Cards.

You can construct stack-links either out of single stack ports or a stack GigaStream. However, because of the incredible 10Gb port density offered by the GigaVUE HC Series, using only one 10Gb port for a stack connection could cause a serious bottleneck.

NOTE: Packet loss may be seen on stack links when traffic exceeds 95% of the line rate. This is because each packet has a 16 bytes higig header added to it, which reduces the throughput.

NOTE: In a stack port configuration, with if a Gen 2 and Gen 3 card from the same chassis, use different VLAN IDs. You cannot use the default VLAN IDs.

A stack GigaStream dramatically increases the bandwidth available for stack connections, letting you connect GigaVUE nodes in a cluster and still take advantage of the 10Gb port density. Alternatively, nodes with 40Gb or 100Gb ports can take advantage of their high bandwidth for stack-links. (For more details about clustering, refer to the *GigaVUE-OS CLI Reference Guide*.)

Stack links are supported at speeds of 10Gb, 40Gb, and 100Gb. Refer to the *Hardware Installation Guide* for each GigaVUE node for information on stack link support.

Stacking is not supported on GigaVUE-TA400 in 5.14 release.

When using stack GigaStream for stack links, you must create a stack GigaStream on each side of the stack link and each must consist of the same number of ports running at the same speed.

To create a stack link, do the following:

1. Select **Ports > Stack Links**.
2. Click **New**.
3. Enter an alias for the stack link in the **Alias** field.
4. Select the **Type** for this stack link.
    - **Stack Ports** specifies that the stack link is between two ports.
    - **Stack GigaStream** specifies that the stack link is between GigaStream.
5. In the First Member and Second Member fields, select the ports or GigaStream for the stack link, depending on the type selected in Step 4
6. Click **Save**.

IP Interfaces

You can configure IP interface in the control card. All the control operations such as the gateway resolution, tunnel health check, and NetFlow exporter SNMP requests are handled in the control card. Similarly, the ARP/NDP timer configuration is also moved to the control card.

You can configure the IP Interfaces for the following tunnels:

- GigaSMART tunnel - Refer to the Configure IP Interfaces for GigaSMART Tunnels.
- Layer 2 Generic Routing Encapsulation (L2GRE) tunnel - Refer to theAbout Layer 2 Generic Routing Encapsulation (L2GRE) Tunnels

- Virtual Extensible LAN (VXLAN) tunnel - Refer to the About Virtual Extensible LAN (VXLAN) Tunnels.

Configure IP Interface

Prerequisites

- Before you configure an IP interface, you must configure a network and a tool port.
- To configure IP interfaces for GigaSMART tunnels, you must create a GigaSMART group and a NetFlow exporter. Associate GigaSMART engines to theGigaSMART group.

To configure an IP interface:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**. Select the node for which you want to configure the IP interface.
2. From the left navigation pane, go to **Ports > IP Interfaces**.
3. In the IP Interfaces page, click **New**. The New IP Interface page appears.
4. In the **Alias** and **Description** fields, enter a name and description for the IP interface.
5. Click **Port Editor**. A quick port editor appears. You can enable the port for which you want to configure the IP interface. Click **Save**.
6. From the **Ports** drop-down list, select the tool or the network port that you had enabled.

   NOTE:
   - The port should not be part of the port group.
   - Before upgrading to the latest version, you must remove either the port group or the IP interface if the ports are identical.

7. Select the type as either **IPv4** or **IPv6**.
8. Enter the **IP Address**, **IP Mask**, and **Gateway** for the IP interface.

   You can specify the subnet mask using either of the following formats:
   - netmask – For example, 255.255.255.248
   - mask length – For example, /29

NOTE: For IPv6 interface configured on Generation 3 GigaSMART card (SMT-HC1-S), prefix length of 128 is not supported.

9. Enter the **MTU** for the IP interface (100 - 9600 bytes). The MTU is fixed at 9600 for all network/tool ports on the following platforms:

- Certified Traffic Aggregation White Box
- GigaVUE-HC1
- GigaVUE-HC3

- GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA100, GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA400, and GigaVUE-TA400E.

RECOMMENDATION: Set the MTU to 9400 on all platforms.

10. If you are configuring IP Interface for GigaSMART tunnels:
    a. From the **GS Groups** drop-down list, select theGigaSMART groups that you have configured.

> **NOTE:** You can associate multiple GigaSMART groups to the IP interface.

    b. From the **Exporters** drop-down list, select the NetFlow exports that you have created.
11. Click **Apply** to configure the IP interface associated with a network or tool port and add it to the list of currently configured IP interfaces.

To edit the existing IP Interfaces, select the IP Interface and click **Actions > Edit**.

To delete the IP interfaces, select the IP Interface and click **Actions > Delete**.

IP Interface Statistics

To view the statistics such as the packets and bytes decapsulated on the IP interface:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**. Select the node for which you want to view the IP interface statistics.
2. From the left navigation pane, go to **Ports > IP Interfaces > Statistics**. The IP Interface Statistics page appears.

To clear the statistics for a specific IP Interface, select the IP Interface and click **Actions > Clear**.

To clear all the statistics, select all the IP interfaces and click **Actions > Clear All**.

To export the IP Interface statistics, click **Export**.

Circuit Tunnels

Circuit tunnels are used to route traffic between two clusters. The traffic is tapped and sent through network ports on the TAP landing nodes in a cluster. Based on the flow map configuration, traffic is filtered at the TAP landing nodes and sent to the circuit ports. The circuit ports encapsulate the traffic with a Circuit ID and routes the encapsulated traffic

through a circuit tunnel. At the receiving end, the traffic is decapsulated and sent to the tool ports. The circuit tunnels are bidirectional. For more information about circuit tunnels, refer to About Circuit-ID Tunnels.

Port Discovery

This section describes port discovery for the GigaVUE HC Series, providing information about discovery protocols and how to enable discovery through GigaVUE-FM. For details refer to the following:

- Port Discovery with LLDP and CDP
- Enable Port Discovery
- Port Discovery Support

Port Discovery with LLDP and CDP

The GigaVUE HC Series and GigaVUE TA Series devices are capable of snooping Link Layer Discovery Protocol (LLDP) packets and Cisco Discovery Protocol (CDP) packets. If the devices in your network use either of these protocols, a GigaVUE HC Series or a GigaVUE TA Series node can identify its immediate neighbors and their capabilities. Snooped LLDP and CDP information includes the remote port and chassis IDs, as well as other selected information, if it is included by the sender. This information can be used to determine the origin of traffic flows.

All GigaVUE HC Series and GigaVUE TA Series nodes support LLDP and CDP port discovery,

LLDP and CDP are physical topology discovery protocols (Layer 2). The protocols are unidirectional. Devices send their identity and capabilities in a packet. The GigaVUE HC Series or GigaVUE TA Series node receives the packet and extracts information from it, such as the chassis ID and port ID of a neighbor. The information from the neighbors varies depending on what is sent in the packet.

An LLDP packet supports the following capabilities in a type-length-value (TLV) structure. The first four capabilities are mandatory.

- Chassis ID
- Port ID
- Time-to-Live (TTL)
- End of TLVs
- Port description

- System name
- System description
- System capabilities available
- System capabilities enabled
- VLAN name
- Management address
- Port VLAN ID
- Management VLAN ID
- Link Aggregation port ID
- Link Aggregation status
- Maximum Transmission Unit (MTU)

A CDP packet supports the following capabilities in a TLV structure:

- Device ID
- Port ID
- Platform
- Software version
- Native VLAN ID
- Capabilities
- Network prefix address
- Network prefix mask
- Interface address
- Management address

The LLDP/CDP discovery packets are copied and parsed by the node, and the neighbor information is cached. Discovery packets are not terminated on the GigaVUE HC Series or the GigaVUE TA Series node, nor are they removed from the ingress data stream.

Notes:

- Port discovery can be enabled on network, tool, and circuit type ports.
- Use port discovery on ports fed by SPAN ports or aggregators with caution. LLDP/CDP information received from a SPAN port may be misleading, depending on how it is configured. When a large range of ports are SPANed, different and conflicting LLDP/CDP information may be received. LLDP/CDP is best used on TAPed network interfaces.

Enable Port Discovery

Port discovery is disabled by default. It can be enabled on network, tool, and circuit type ports.

NOTE: The ports do not have to be included in a map.

1. Select **Ports > All Ports**.
2. On the Ports page, click on the Port ID of the port on which you want to enable port discovery.

   The Quick View window displays for the port ID.
3. Select **Edit** from the top right corner of the Quick View Window.
4. To enable ports discovery do the following under **Device Discovery**:

   a. Select **Enable**

   b. For Discovery Protocols, select one of the following: **All**, **LLDP**, or **CDP**.

   Enable Port DiscoveryFigure shows ports discovery enabled using the LLDP protocol for the port 1/2/x1.
5. Click **OK**.

Limits of Discovery Information

The following are limits on the amount of discovery information that is retained:

- For each port, discovery information for a minimum of two neighbors and a maximum of 20 neighbors is retained.
- For a chassis, discovery information for a maximum of 2K neighbors is retained.

Neighbor information will be removed or replaced as follows:

- when the neighbor information expires due to the TTL
- when the number of neighbors for the chassis reaches the 2K maximum and a new neighbor is discovered. In this case, the following can occur:
  - if there are currently two or more discovered neighbors for a port, the newly discovered neighbor will replace the neighbor information for the least recently updated neighbor
  - if there are currently less than two discovered neighbors for a port, the newly discovered neighbor will be added (actually exceeding the 2K limit in order to guarantee a minimum of two neighbors per port)

NOTE: Aging (the discovery protocol time-to-live) determines how long neighbor information is valid.

Port Discovery Support

This section describes port discovery for a cluster and port discovery for SNMP.

Port Discovery for a Cluster

LLDP and CDP discovery can be enabled on any network, tool, or circuit ports in a cluster. The discovery information will be aggregated and available on the cluster leader.

Port Discovery Supported for SNMP

The information from LLDP discovery is supported in the standard MIB and can be retrieved with SNMP **Get**.

The name of the MIB file that needs to be loaded in order to poll the LLDP information with SNMP is as follows:

- LLDP-MIB

The information from CDP discovery is supported in Cisco private MIBs and can be retrieved with SNMP **Get**.

The names of the Cisco MIB files that need to be loaded in order to poll the CDP information with SNMP are as follows:

- CISCO-CDP-MIB
- CISCO-SMI
- CISCO-SMI-MIB
- CISCO-TC
- CISCO-TC-MIB
- CISCO-VTP-MIB

Ingress and Egress VLAN

This section describes ingress port VLAN tagging and egress port VLAN stripping. Refer to the following sections for details:

- About Ingress Port VLAN Tagging

    - Ingress Port VLAN Tagging
    - Adding VLAN Tags
    - Deleting VLAN Tags
- Using VLAN Tags in Maps
- Ingress Port VLAN Tag Limitations
    - Second Level Maps
    - Double-Tagged Packets
    - IP Interfaces
- Configure Egress Port VLAN Stripping
    - Enable Egress Port VLAN Stripping
    - Disable Egress Port VLAN Stripping
    - Display Egress Port VLAN Stripping
- Egress Port VLAN Stripping Limitations
- How to Use Both Ingress Tagging and Egress Stripping

About Ingress Port VLAN Tagging

You can add VLAN tags to ingress packets on a per-port basis. You manually associate VLAN IDs with specific ports of type network, inline-network, or hybrid or hybrid GigaStream.

Use VLAN tags to identify, differentiate, or track incoming sources of traffic. When the traffic reaches the tools or the maps, you can filter on the VLAN tags for the corresponding ports you want to measure.

Ingress port VLAN tagging is supported for IPv4 and IPv6 packet types, including non-tagged packets, tagged packets, and Q-in-Q packets. Ingress port VLAN tagging is not supported on network ports that are connected via port-pairs.

Each port can only have one VLAN tag. The same VLAN tag can be assigned to multiple network ports or hybrid ports or to both ports in an inline network port pair.

VLAN tagging is supported in a cluster.

Refer to Figure 11Using Ingress Port VLAN Tagging for an example. In the example, traffic from San Jose is tagged with VLAN 1001 and traffic from San Francisco is tagged with VLAN 1002.

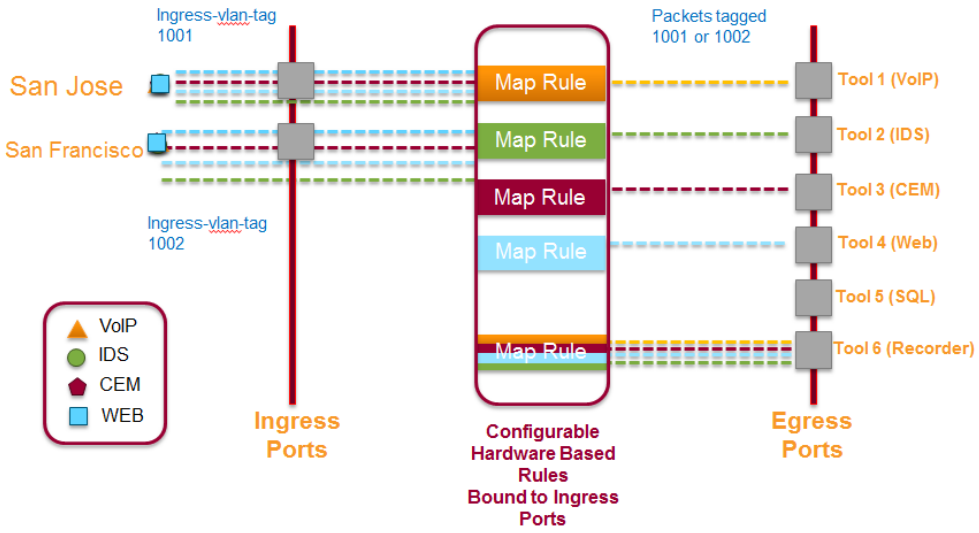


Figure 11 *Using Ingress Port VLAN Tagging*

Ingress Port VLAN Tagging

The port type must be a network or hybrid or inline network type of port. Each network port can only have one VLAN tag. Once a VLAN tag is configured, it can be modified by overriding the existing one with a new VLAN ID.

The VLAN ID is specified in the **VLANTag** field of the port configuration page for an network or inline network port. The value of the VLAN ID is specified as a number between 2 and 4000.

Adding VLAN Tags

To add/modify VLAN tags, follow these steps:

1. Go to **System > Ports > All Ports**.
2. Click on the **Port ID**. Ensure that this Port ID is set as a network port. The Quick View window for the Port ID displays.
3. Select **Edit** from the top right corner of the Quick View Window.
4. Add the **VLAN ID** to the parameter field **VLAN Tag** and click **Save**.

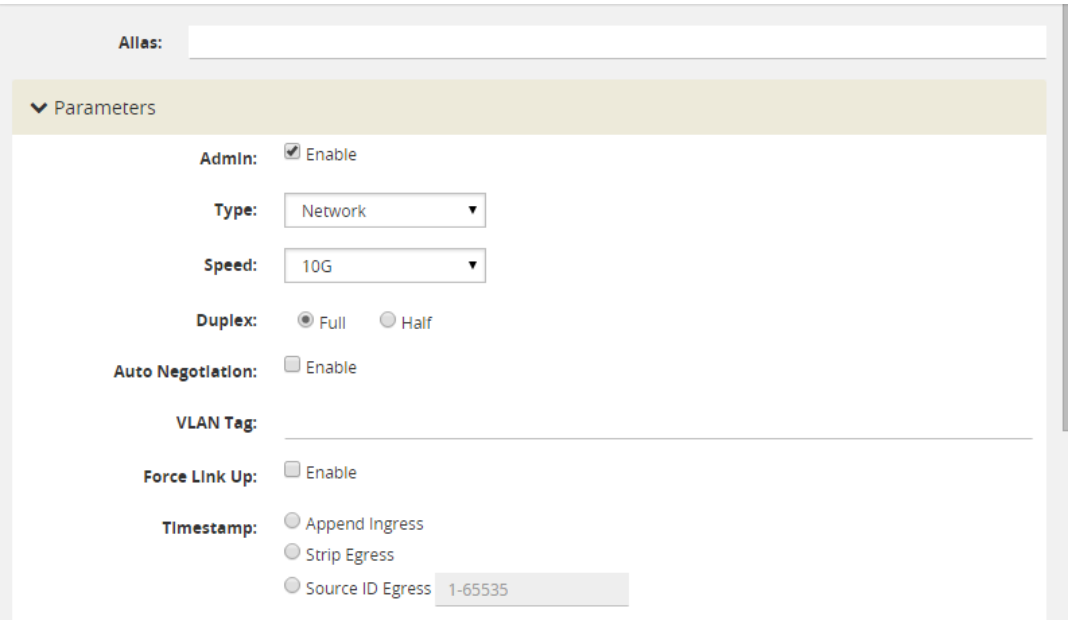


Figure 12 *Port-ID Configuration*

Deleting VLAN Tags

Once a VLAN tag is configured, it can be deleted by removing the value from the **VLAN Tag** field and saving the port configuration.

Using VLAN Tags in Maps

Ingress port VLAN tags are supported in first level maps, including the following:

- map
- map-passall
- map-scollector
- GigaSMART operation (gsop-enabled) maps

For example, if the traffic from network port 2/1/q3, (which has VLAN tag 1001 configured), is forwarded to tool port 2/1/q4. The traffic at tool port 2/1/q4 will have the added VLAN tag 1001. (Even though the VLAN tag is configured on the network port, it is added when the traffic exits the tool port.)

NOTE: Traffic from a network port will not match a map rule that filters on a VLAN tag configured on the network port.

Ingress Port VLAN Tag Limitations

The following sections describe limitations of ingress port VLAN tagging:

- Second Level Maps
- Double-Tagged Packets
- IP Interfaces
- Local Tool Port Ingress VLAN Tag

Second Level Maps

VLAN tagging is not supported for second level maps, which are maps from a virtual port (vport).

For tagged network ports, if the ingress traffic is going to a second level map, the packets will not be tagged at the egress ports of the second level map. This is a limitation of GigaSMART operations using maps with vports.

Double-Tagged Packets

If incoming packets already have two VLAN tags, such as with Q-in-Q, the addition of a third VLAN tag can cause problems with the following:

- Layer 3/Layer 4 filtering
- GigaStream hashing (all packets may be sent to only one tool port)

IP Interfaces

For IP interfaces, a VLAN tag added at the network port of the encapsulation path (n1 in Figure 13 IP Interfaces) will become part of the payload going to the decapsulation path. But a VLAN tag added at the network port of the decapsulation path (n2 in Figure 13 IP Interfaces) will be available at the end tool port for filtering (t2 in Figure 13 IP Interfaces).

Refer to Figure 13 IP Interfaces. VLAN tag (vlan1) added at the encap network port (n1) is encapsulated in the tunnel payload and cannot be used for filtering at the decap side. VLAN tag (vlan2) added at the decap network port (n2) can be used in a filter rule to send packets to tool port (t2).

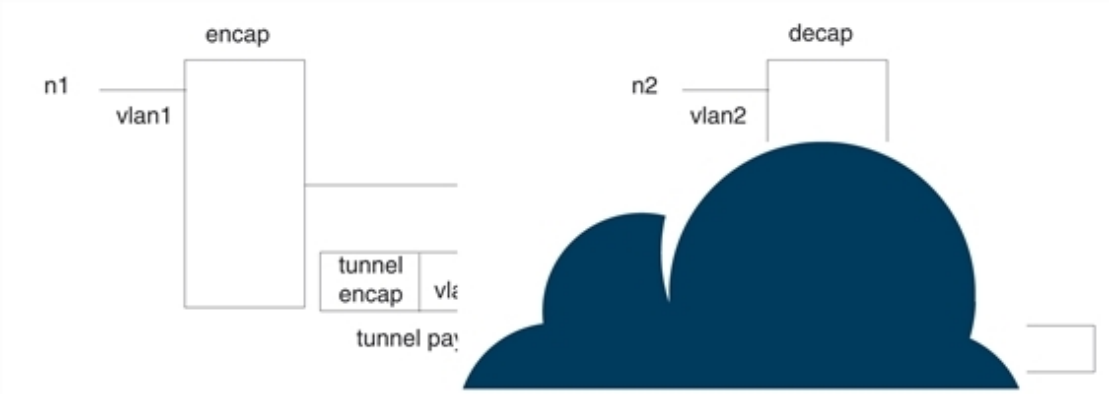


Figure 13 *IP Interfaces*

Local Tool Port Ingress VLAN Tag

Ingress VLAN tag is not working on the local tool port if both MPLS and Ingress VLAN tag are enabled on the same network port of GigaVUE-HC1-Plus, GigaVUE-HCT, GigaVUE-TA25, and GigaVUE-TA25E.

Note: Circuit ports to make Circuit Mapping functionality to work is not supported.

Configure Egress Port VLAN Stripping

You can enable or disable outer VLAN stripping on specified egress ports. The port type must be tool or hybrid.

Use egress port VLAN stripping to strip an outer VLAN tag without using a GigaSMART stripping operation.

Enable Egress Port VLAN Stripping

To enable egress port VLAN stripping:

1. Select a tool or hybrid port on the Ports page.
2. Click **Edit**.
3. Under Parameters, for Egress Vlan Tag, select Strip. Refer to Figure 14Select Strip Egress Vlan Tag.

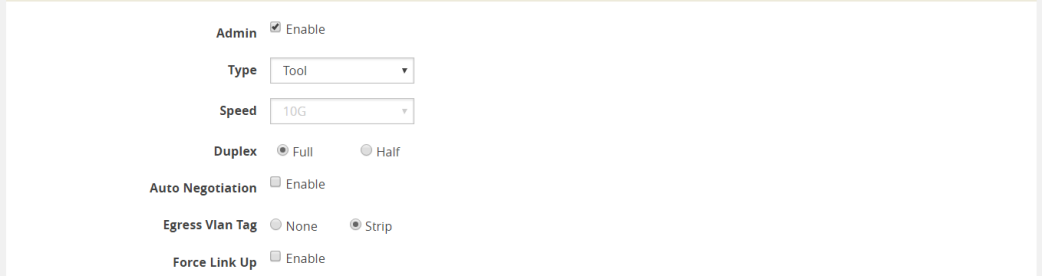


Figure 14 *Select Strip Egress Vlan Tag*

4. Click **OK**.

Disable Egress Port VLAN Stripping

Once egress port VLAN stripping is enabled, it can be disabled. In Figure 14Select Strip Egress Vlan Tag, for Egress Vlan Tag, select None.

Display Egress Port VLAN Stripping

To display egress port VLAN stripping configuration:

1. Double-click a tool or hybrid port on the Ports page.
2. View the configuration under Port Info.

Egress Port VLAN Stripping Limitations

The following are limitations of egress port VLAN stripping:

- Enabling both ingress port VLAN tagging and egress port VLAN stripping on the same port is not supported.
- Egress port VLAN stripping does not support inline tool ports or stack ports.
- If a port is configured for egress port VLAN stripping, configuring a port filter with either pass or drop VLAN rules is not recommended.
- In Clustering, Egress port VLAN stripping with outer tag ethertype of 0x88A8 or 0x9100 is supported if the tool port and the network port are configured in the same device.
- An outer tag of ethertype 0x8100 works without any limitation when egress port VLAN stripping is configured along with ingress port VLAN tagging on the upstream GigaVUE node.
- In GigaVUE-TA25, GigaVUE-HC1-Plus, GigaVUE-HCT, and GigaVUE-TA25E engines, Port level egress VLAN stripping is only supported if the network and tool ports are not configured on the same node in double tag mode.

How to Use Both Ingress Tagging and Egress Stripping

When ingress port VLAN tagging is enabled on a network port and egress port VLAN stripping is enabled on a tool port on the same GigaVUE node, refer to the Table 5: VLAN Stripping Table:

Table 5: VLAN Stripping Table

Tool →	Stripping Enabled			Stripping Disabled		
Network	**Untagged**	**Single Tag**	**Double Tag**	**Untagged**	**Single Tag**	**Double Tag**
Ingress VLAN tag enabled	None	None	Customer VLAN tag	Ingress VLAN tag	Ingress VLAN tag + Customer VLAN tag	Ingress VLAN tag + Customer VLAN tag + Service VLAN tag
Ingress VLAN tag disabled	None	None	Customer VLAN tag	None	Customer VLAN tag	Service VLAN tag

NOTES:

- The inner VLAN tag is classified as the Customer VLAN tag (ethertype 0x8100)
- The outer VLAN tag is classified as the Service VLAN tag (ethertype 0x8100, 0x88A8, or 0x9100)

How to Use GigaStream

This section describes how to create and manage GigaStream. A GigaStream groups multiple ports into a logical bundle. Refer to the following sections for details:

- About GigaStream
- Regular GigaStream
- Controlled GigaStream
- Advanced Hashing
- Weighted GigaStream
- GigaStream Rules and Maximums

GigaStreams

A GigaStream is a bundle of multiple ports on a GigaVUE-OS node. You can create the following types of GigaStreams:

- **Tool GigaStream**—It is a bundle of multiple tool ports used as a single logical group. This type of GigaStream as a single addressable destination, allowing you to overcome tool port oversubscription issues.
- **Hybrid GigaStream**—It is a bundle of multiple hybrid ports that are combined into a single logical group in all H Series nodes.
- **Stack GigaStream**—It is a bundle of multiple stack ports used as a single logical group. Stack-links can use GigaStream to distribute data between multiple GigaVUE® HC Series nodes operating in a cluster. With the number of 10Gb/40Gb/100Gb ports possible in a GigaVUE HC Series chassis, using only one 10Gb port for a stack-link could cause a serious bottleneck. A GigaStream dramatically increases the bandwidth available for stack-link connections, letting you connect H Series nodes in a cluster and still take advantage of the 10Gb port density.
- **Circuit GigaStream**—It is a bundle of multiple circuit ports that are combined into a single logical group. The circuit ports send or receive traffic that is tagged with the circuit ID.
- **Controlled GigaStream**—Controlled GigaStream provides more control of the traffic stream by specifying the size of a hash table and allowing the assignment of hash IDs to the ports in a GigaStream. This makes it possible to keep the hashing algorithm from reapplying the algorithm to the ports if one of the ports in the GigaStream goes down.

Accessing the GigaStream page

To access the GigaStream page:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**..
2. In the Physical Nodes page, click the required standalone node or the cluster ID for which you want to configure a GigaStream.
3. On the left-navigation pane, go to **Ports > Port Groups > GigaStream™**.

The GigaStream page displays the list of configured GigaStreams.

The GigaStream page has the following buttons, which can be used to perform specific options.

Button	Description
Tags	Use to associate tags to GigaStream and to remove the tags from

Button	Description
	GigaStream when not required. To add tags: 1. Select the required GigaStream to which the tags must be applied. 2. Click the **Tags** drop-down menu button and select **Add**. 3. Click **Ok**. To remove tags: 1. Select the required GigaStream that are already associated with tags. 2. Click the **Tags** drop-down menu button and select **Delete**. 3. Click **Ok**.
New	Use to create a new GigaStream. Refer to Configure Regular GigaStream for detailed information.
Actions	Use the **Actions** drop-down button to perform the following options: • **Clone**: Use to clone the selected GigaStream. • **Edit**: Use to edit the selected GigaStream • **Delete**: Use to delete the selected GigaStream
Advanced Hash Settings	Use to configure the advanced hash settings. Refer to Advanced Hashing section for detailed information.
Export	Use to export the GigaStream in CSV or XLSX format. The following options are available: • **Export All**: All GigaStreams in the list view are exported. • **Export Selected**: Only the selected GigaStreams are exported.

About GigaStream

There are two types of GigaStream:

- Regular GigaStream
- Controlled GigaStream.

Both types of GigaStream bundle multiple ports to provide logical bandwidth. Packets arriving through network ports are processed with various map rules and then directed to ports. All traffic streams destined to a GigaStream are hashed among the bundled ports.

Regular GigaStream

Regular GigaStream groups multiple ports running at the same speed into a single logical bundle called a GigaStream. Regular GigaStream can be used as either a packet egress destination (tool GigaStream) or as a stack-link between two GigaVUE-OS nodes operating in a cluster (stack GigaStream).

NOTE: The existing tool and stack GigaStream are now referred to as regular GigaStream. The term GigaStream is used when something applies to both types

For details on regular GigaStream, refer to Regular GigaStream.

Controlled GigaStream

Controlled GigaStream provides GigaStream controlled traffic distribution. Controlled GigaStream samples traffic based on hash settings and helps to ensure that traffic sent to each tool is within the capacity of the tool.

For details on controlled GigaStream, refer to Controlled GigaStream.

Controlled GigaStream provides greater flexibility in allocating the bandwidth assigned to tools within the GigaStream. Regular GigaStream assumes that each tool in the GigaStream is sent an equal fraction of the traffic. Controlled GigaStream allows different tools to be sent different fractions of the traffic.

Note: Regular GigaStream and controlled GigaStream differ in the following ways:

- traffic distribution based on hashing
- traffic fail over when a port goes down
- editing a configuration, such as adding ports on the fly

Regular GigaStream

Regular GigaStream can be used as either a packet egress destination (tool GigaStream) or as a stack-link between two GigaVUE-OS nodes operating in a cluster (stack GigaStream).

All ports in a GigaStream must be running the same speed, such as 10Gb or 40Gb, and must use the same port type, either tool or stack. All ports in a GigaStream can be on different modules of the same GigaVUE-HC3 node.

With regular GigaStream, the hashing is computed based on traffic. Incoming packets arriving through network ports are processed with various map rules and then directed to ports. The result of the hash distributes traffic equally across the GigaStream members. Refer to Figure 15Regular GigaStream Overview.

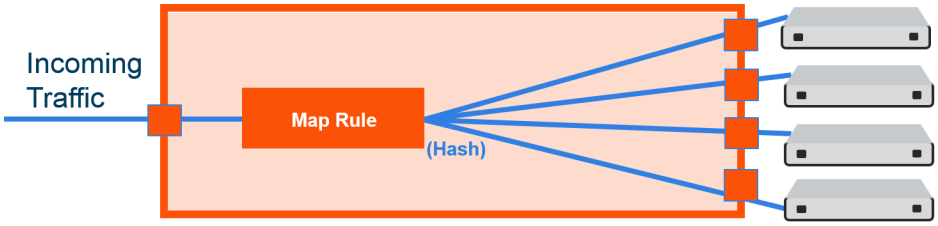


Figure 15 *Regular GigaStream Overview*

With the advanced hashing mode, the hash is a result of multiple parameters such as source MAC address, destination MAC address, source IP address, destination IP address, protocol, or other criteria. The hashing algorithm determines the destination tool port for a particular packet. All packets matching a particular set of hashing criteria will be sent to the same port. Sessions are maintained within a stream.

For example, a regular tool GigaStream is configured with ports x1 to x4. The hash table of size 4 is evenly divided among the 4 ports, and the traffic is distributed accordingly.

Regular Tool GigaStream

A regular tool GigaStream can be used as a single addressable destination, allowing you to overcome tool port over subscription issues.

NOTE: A regular tool GigaStream can consist of tool ports or hybrid ports.

Refer to Figure 16Regular Tool GigaStream Illustrated.

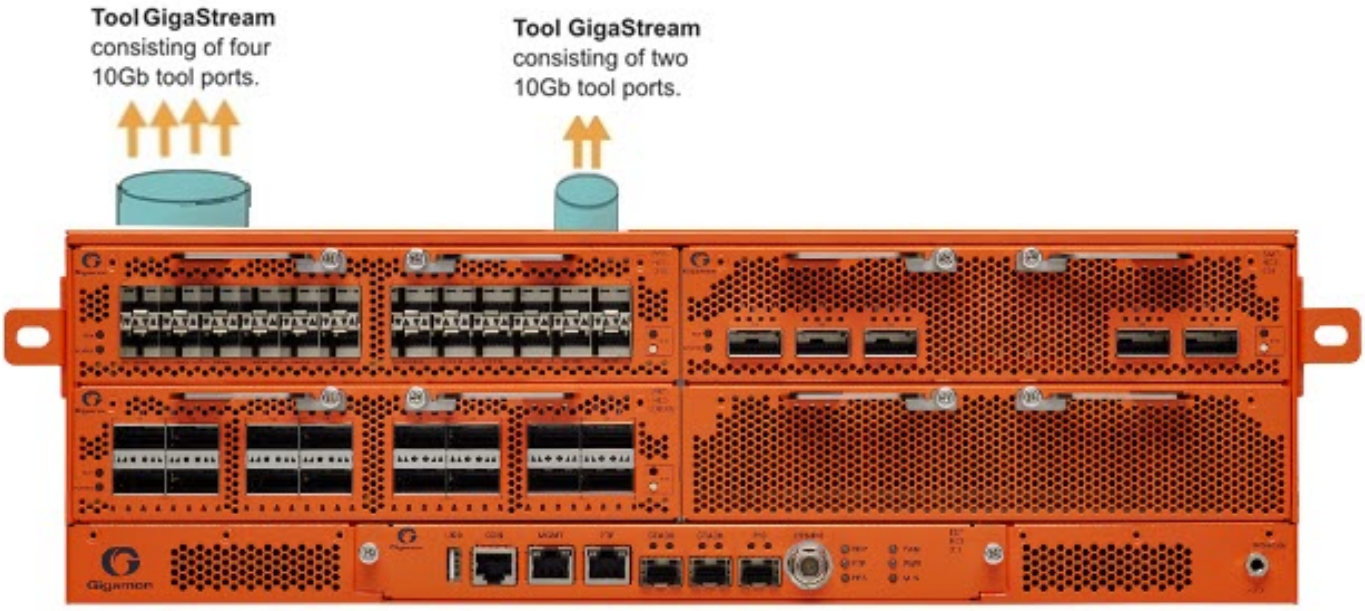


Figure 16 *Regular Tool GigaStream Illustrated*

Regular Stack GigaStream

A regular stack GigaStream can use stack-links to distribute data between GigaVUE-OS nodes operating in a cluster. With the terabits of throughput possible in a GigaVUE HC Series node, using only one 10Gb port for a stack-link could cause a bottleneck. A regular stack GigaStream dramatically increases the bandwidth available for stack-link connections, providing greater flexibility and throughput within a cluster.

Refer to Figure 17Regular Stack GigaStream Illustrated.

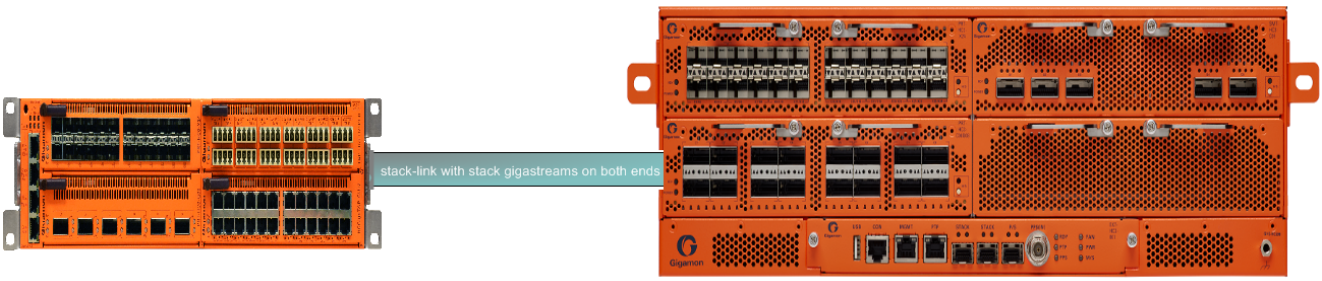


Figure 17 *Regular Stack GigaStream Illustrated*

Configure Regular GigaStream

All ports combined in a GigaStream must be running at the same speed, using the same port types. Port speeds less than 1000Mb are not supported.

Also, refer to Advanced Hashing for optional advanced hashing settings and Weighted GigaStream for optional weighting settings.

Before you configure a regular GigaStream, ensure that you have configured the required ports—tool, hybrid, stack, or Circuit. For information about configuring a port, refer to Configure Ports.

To configure a regular GigaStream:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**..
2. In the Physical Nodes page, click the required cluster ID for which you want to configure a regular GigaStream.
3. On the left-navigation pane, go to **Ports > Port Groups > GigaStream™**.
4. Click **New**. The **GigaStream™** page appears.
5. In the **Alias** and **Description** fields, enter the name and description of the regular GigaStream that you want to configure.
6. Select the type of GigaStream that you want to configure. For example, if you want to configure a regular tool GigaStream, select **Tool GigaStream**.
7. From the **Ports** drop-down list, select the ports that you have configured. For example, select the required hybrid ports to configure a regular hybrid GigaStream.
8. From the **Weighting** drop-down list, select one of the following options:
    - **Equal**—Traffic is distributed equally to all the ports in the regular GigaStream.



You can enter a value for the Variance Threshold in % only if:

- Equal Weighting is selected
- GigaStream is a regular GigaStream
- Device Version is 6.0 and above

If the threshold is above or below the Variance threshold, and if the traffic distribution is not uniform across the ports in the GigaStream, traps will be sent from the devices to GigaVUE-FM

    - **Relative**—Traffic is distributed to the ports in the regular GigaStream based on the relative weight or ratio assigned to the respective ports. The valid range is 1–256.
    - **Percentage**—Traffic is distributed to the ports in the regular GigaStream based on the percentage assigned to the respective ports. The valid range is 1–100.

If you select **Relative** or **Percentage** as the weighting option, enter the hash weights for the ports that appear in the table below the **Weighting** drop-down list.

9. In the **Drop Weight** field, enter the relative weight to drop the traffic. For example, if you enter 2 in this field, 2% of the total traffic entering the regular GigaStream will be dropped.

   NOTE: The **Weighting** and the **Drop Weight** fields are not available when you configure a regular stack GigaStream.

   For example, you want to send only 25% of the traffic to a tool group with four tool ports, 1/1/c1..c4. Depending on the Weighting option you choose, enter the hash weights as follows:
   - **Relative**—Enter the hash weights for the ports as 1 for 1/1/c1..c4 and **Drop Weight** as 12. This means that the traffic, 1/16 goes to tool port 1/1/c1, 1/16 goes to tool port 1/1/c2, and so on. 12/16 traffic gets dropped.
   - **Percentage**—Enter the hash weights for the ports in percentage. 6% for 1/1/c1..c4 and **Drop Weight** as 76%. This means that 6% of the traffic goes to tool port 1/1/c1, 6% goes to tool port 1/1/c2, and so on. 76% of the traffic gets dropped.

10. Select the required **Tag** and **Tag Value** that needs to be assigned to the GigaStream. For more information about Tags, refer to the Tags section in the GigaVUE Administration Guide.
11. Click **Apply** to save the configuration.

The configured regular GigaStream appears in the table in the GigaStream™ page.

Edit Regular GigaStream

GigaVUE-FM provides support for editing a regular GigaStream. You can add and delete tool ports from GigaStreams without the need to recreate the GigaStream with a new map.

Support is available to edit Tool and Hybrid GigaStream's attached to the following maps types:

1. First level GigaSMART map
2. Second level GigaSMART map (includes GTP overlap sampling maps)
3. Regular map (including port mirroring, collector, etc.)

NOTE: Regular GigaStream and controlled GigaStream are interchangeable. You can change the type of GigaStream from regular to controlled in real-time.

To edit a regular tool GigaStream attached to a map:

1. Select **Ports > Port Groups > GigaStreams**.

2. Select a GigaStream and click **Actions** > **Edit** to open the GigaStream configuration page.
3. Click in the **Ports field** and **Change**, **Add** or **Delete** tool ports associated with the GigaStream as needed.
4. Click **Apply**.

Edit Regular Stack GigaStream

You can edit regular stack GigaStreams that are configured on either sides of a stack link. When a stack GigaStream is attached to a map, you can directly add or delete stack ports from the stack GigaStream.

To edit a regular stack GigaStream:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**..In the Physical Nodes page, click the required cluster ID.
2. On the left-navigation pane, go to **System>Ports > Stack Links**. The Stack Links page appears.
3. Select the alias of the stack link that is grouped in the stack GigaStream that you want to modify, and then click **Edit**. The Stack Link page appears.
4. From the **Ports** drop-down list, add or delete the required stack ports, and then click **OK**.

Traffic Distribution Across Regular GigaStream

All the traffic streams destined to the GigaStream are distributed among the bundled ports based on hashing, as defined in the advanced hash settings or Weighted GigaStream, as defined in the **Weighting** field.

The hash is performed across multiple fields, such as IP address, port number, protocol, MAC address, and other criteria. The best practice is to include both the source and destination fields, such as source IP address and destination IP address, within the advanced hash settings. Because the hash calculation is symmetrical with respect to source and destination addresses, all packets belonging to the same session will be sent to the same tool.

For more information on hashing, refer to Advanced Hashing.

The GigaVUE-OS nodes distribute traffic between the ports in a regular GigaStream using one of the following criteria:

- The criteria configured using the Advanced Hash Setting page for the selected line card or chassis. (Click **Advanced Hash Setting** on the GigaStream page to open the Advance Hash Setting.) Because traffic is hashed across member ports rather than

divided evenly, the bandwidth available for a regular GigaStream is not a straight multiple of the number of ports in the bundle – some flows will use more bandwidth than others.

> **NOTE:** The GigaVUE HC Series node tries to distribute incoming traffic evenly across all tool ports in the GigaStream. However, live network traffic is often unpredictable, including bursty periods for certain sessions. Because of this, the distribution patterns described are not ironclad – variations in traffic will result in variations in distribution.

The distribution described in this section applies to GigaVUE-HC3 nodes, GigaVUE-HC1 nodes, and GigaVUE TA Series nodes for regular tool GigaStream and regular stack GigaStream.

- Weighting mode and hash weights assigned to the different ports in the regular GigaStream. For more information about Weighted GigaStream, refer to Weighted GigaStream.

Regular GigaStream Failover Protection—Resiliency

Regular GigaStream has built-in failover protection or resiliency. When there is a failover of a port that is part of a regular GigaStream, the traffic is redistributed to the other tool ports without disturbing the session continuity.

When a tool port goes down, the sessions allocated to the failed port are redistributed among the remaining available ports. With the resiliency functionality, the redistribution of traffic occurs without disturbing the session continuity of the active ports.

Recovery of a regular GigaStream is automatic. When a down link returns, the traffic will be reassigned to their original ports automatically.

The resiliency functionality is supported on all GigaVUE HC Series, and TA Series nodes and on Hybrid GigaStream, Tool GigaStream, Circuit GigaStream and Cluster GigaStream.

Regular Circuit GigaStream

A regular circuit GigaStream uses circuit ports to pass traffic between two clusters. For more information refer to the following topics:

- About Circuit-ID Tunnels
- Fabric Maps Prerequisites

Controlled GigaStream

Controlled GigaStream provides controlled traffic distribution, which gives more granular control over hashing to the tool ports.

- All GigaVUE® HC Series and TA Series nodes support controlled GigaStream.
- GigaVUE nodes with controlled GigaStream are supported in a cluster environment.
- Controlled GigaStream can only be used as a packet egress destination (tool GigaStream). All port speeds are supported.

Controlled GigaStream delivers traffic flows based on hash settings and hash bucket assignments. Assigning varying numbers of hash buckets helps to ensure that traffic sent to each tool is within the capacity of the tool. Leaving some hash buckets unassigned creates a simple traffic sampling configuration, as traffic flows for unassigned hash buckets are discarded.

With controlled GigaStream, hashing is computed based on traffic. There is a configurable number of hash buckets, from 1 to 256. The hash size of a controlled GigaStream specifies the number of logical tools the traffic will be distributed across. Refer to Figure 18Controlled GigaStream Overview.

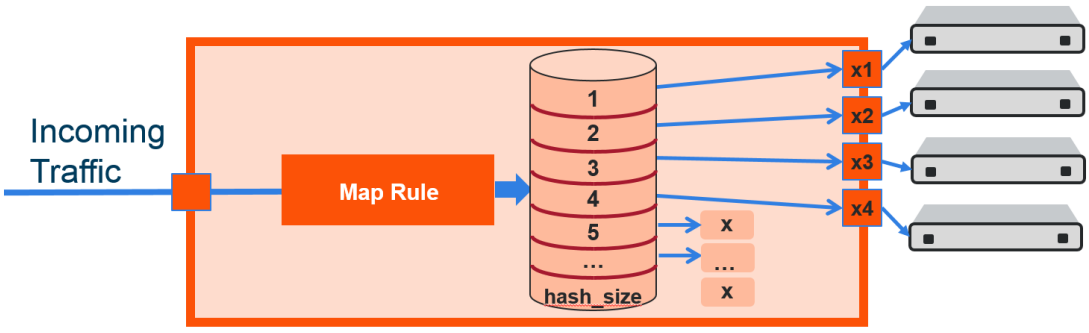


Figure 18 *Controlled GigaStream Overview*

Controlled GigaStream can manage network port bandwidth hashed to GigaStream tool ports. For example, if there is 10Gb of distributed traffic coming in on network ports directed to a GigaStream, and the tools connected to each tool port of the GigaStream can handle only 2Gb of bandwidth, the GigaStream can distribute the streams to 5 tool GigaStream ports. The ingress bandwidth divided by the number of tools determines the number of hash buckets.

Not all hash buckets need to be mapped to ports. In Figure 18Controlled GigaStream Overview, four buckets are mapped to ports, while the remaining buckets are black holed. This provides a form of sampling, that is, only a sample of traffic is sent to the tools.

To determine the best hash size to use for your monitoring needs, divide the maximum bandwidth being sent to the tools by the bandwidth that can actually be consumed by the tools. For example, if you have 150Gb of traffic, but the tools can only process 3Gb, the recommended hash size is 150/3 = 50. To have completely even distribution across the logical tools, round up to the nearest power of 2. In this example, round up a hash size to 64.

When you have more bandwidth than the tools can process, you can use controlled GigaStream to restrict the amount of traffic sent to each tool. The hash size is determined by:

- the amount of traffic to be monitored, for example 300Gb
- the maximum bandwidth of the monitoring tools, for example 2.5Gb

Then divide (300/2.5=120), and round up to a power of 2 (for example, 16, 32, 64, 128). In this case, the hash size would be 128.

Another use for a controlled GigaStream is to increase the reliability of tool ports. For example, a trunk size of 5 is configured on 4 ports with 1 hash bucket each, port x1 is allocated or mapped to hash bucket ID 1, port x2 is mapped to hash bucket ID 2, port x3 is mapped to hash bucket ID 3, and port x4 is mapped to hash bucket ID 4. Hash bucket ID 5 is not mapped to a port. It can be reserved to be mapped to a port later. Until then, any traffic hashed to bucket 5 will be black holed. Refer to Figure 19Controlled GigaStream Example.

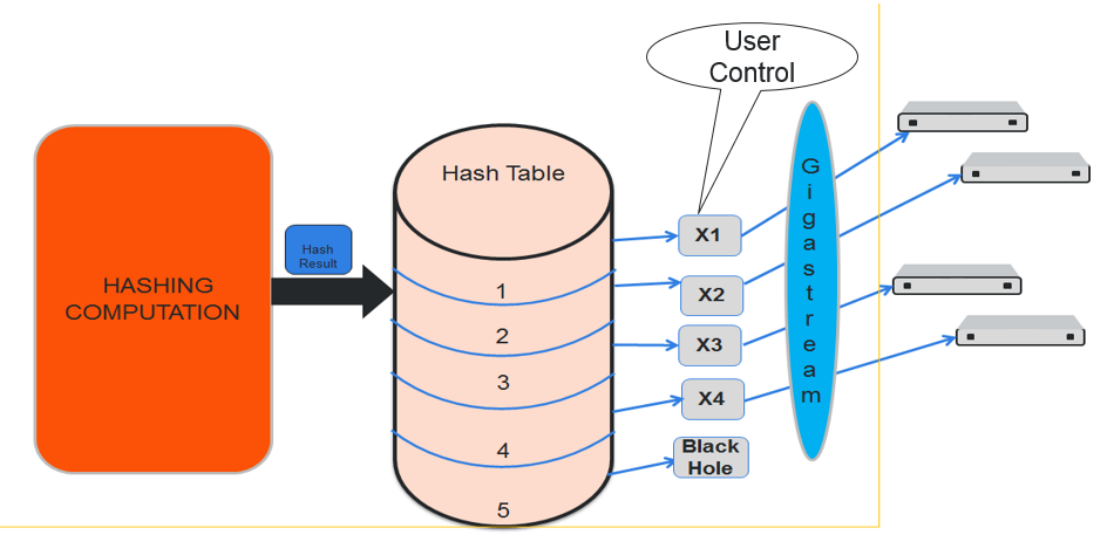


Figure 19 *Controlled GigaStream Example*

Generally, if there are only four tools available, with controlled GigaStream, a GigaStream trunk size of 5 can be configured and allocated to the available four tool ports of the GigaStream. The fifth tool port can be reserved and you can attach that port to the GigaStream whenever it is needed. The existing traffic streams are not impacted.

GigaStream controlled traffic distribution provides enhanced control of traffic hashed across the trunk ports as compared to regular GigaStream. The GigaStream trunk size is configurable, and ports can be dynamically added and deleted.

To configure controlled GigaStream, there are two parameters needed as follows:

- the hashing trunk size, which defines the number of hash buckets to be configured. It defines the maximum number of hash bucket IDs, from 1 to 256.
- the hash bucket ID, which specifies the mapping to ports. Each member of the trunk is mapped to a hash bucket ID. Mapping a port to a hash bucket ID makes it part of a GigaStream. The mapping to ports is static. When a port goes down, traffic is not re-hashed to the remaining ports.

In general, a controlled GigaStream is defined with a hash size equal to the number of trunk ports expected. Mapping a hash bucket ID to each trunk port will evenly distribute the traffic among the ports.

A particular trunk member can be mapped to multiple hash bucket IDs. If one tool can handle 4Gb, 2 hash bucket IDs can be mapped to that tool port. A trunk member that is configured to two hash bucket IDs will be two times more likely to receive hashed traffic as compared to a trunk member with one hash bucket ID. Thus, more traffic can be sent to the higher capacity tools in the GigaStream.

NOTE: A hash bucket ID cannot be mapped to multiple ports.

For more information, refer to Traffic Distribution Across Controlled GigaStream.

Notes and Considerations for Controlled GigaStream

Refer to the following notes and considerations for controlled GigaStream:

- Controlled GigaStream can be used with a regular map, map-passall, or maps collector.
- Controlled GigaStream supports tool ports, but not inline tools or inline tool groups.
- Controlled GigaStream does not support stack or hybrid port types.
- Controlled GigaStream does not support a GigaSMART operation (gsop), or the first and second level maps associated with it.
- The maximum hash size is 256 per trunk.
- All the tool ports participating in the GigaStream must be on the same node. GigaStream can be created across GigaVUE-HC3 modules.
- All participating ports in the GigaStream must be running the same speed and must use the same port type.

- A **Controlled GigaStream** checkbox is used on the GigaStream configuration page for controlled GigaStream, enabling the prefix mode for specifying hash size and hash bucket IDs.
- Before attaching a controlled GigaStream to a map, it should be configured with at least one port.
- Controlled GigaStream can be modified on the fly, even after it is attached to a map. Refer to Edit Regular GigaStream.
- If the GigaStream is already attached to a map, the last mapped hash bucket ID cannot be deleted. That is, do not delete all the ports from a controlled GigaStream.

Controlled GigaStream Configuration

To configure a controlled tool GigaStream, specify hash size and hash bucket ID.

Following are the steps to configure Controlled GigaStream:

1. Use the Quick Port Editor to configure ports as type tool for the controlled GigaStream.
2. Select **Ports > Port Groups > GigaStreams** and click **New** to open the GigaStreams configuration page.
3. Enter an name for the GigaStream in the **Alias** field. For example stream 2.
4. (Optional) Enter a comment in the **Description** field. For example, controlled GigaStream.
5. For **Type**, select **Tool GigaStream**.
6. Select **Controlled GigaStream**. The **Port** field changes to **Hash Size**.
7. In the **Hash Size** field, specify a hash size value. The range is 1 through 256.

   The hash size value determines the number of hash bucked IDs and ports available for assigning to the GigaStream. For example, in Figure 20Controlled GigaStream with Five Hash IDs and Port IDs, the Hash Size is set to 5 so the GigaStream page displays five Hash IDs and five Port ID fields.

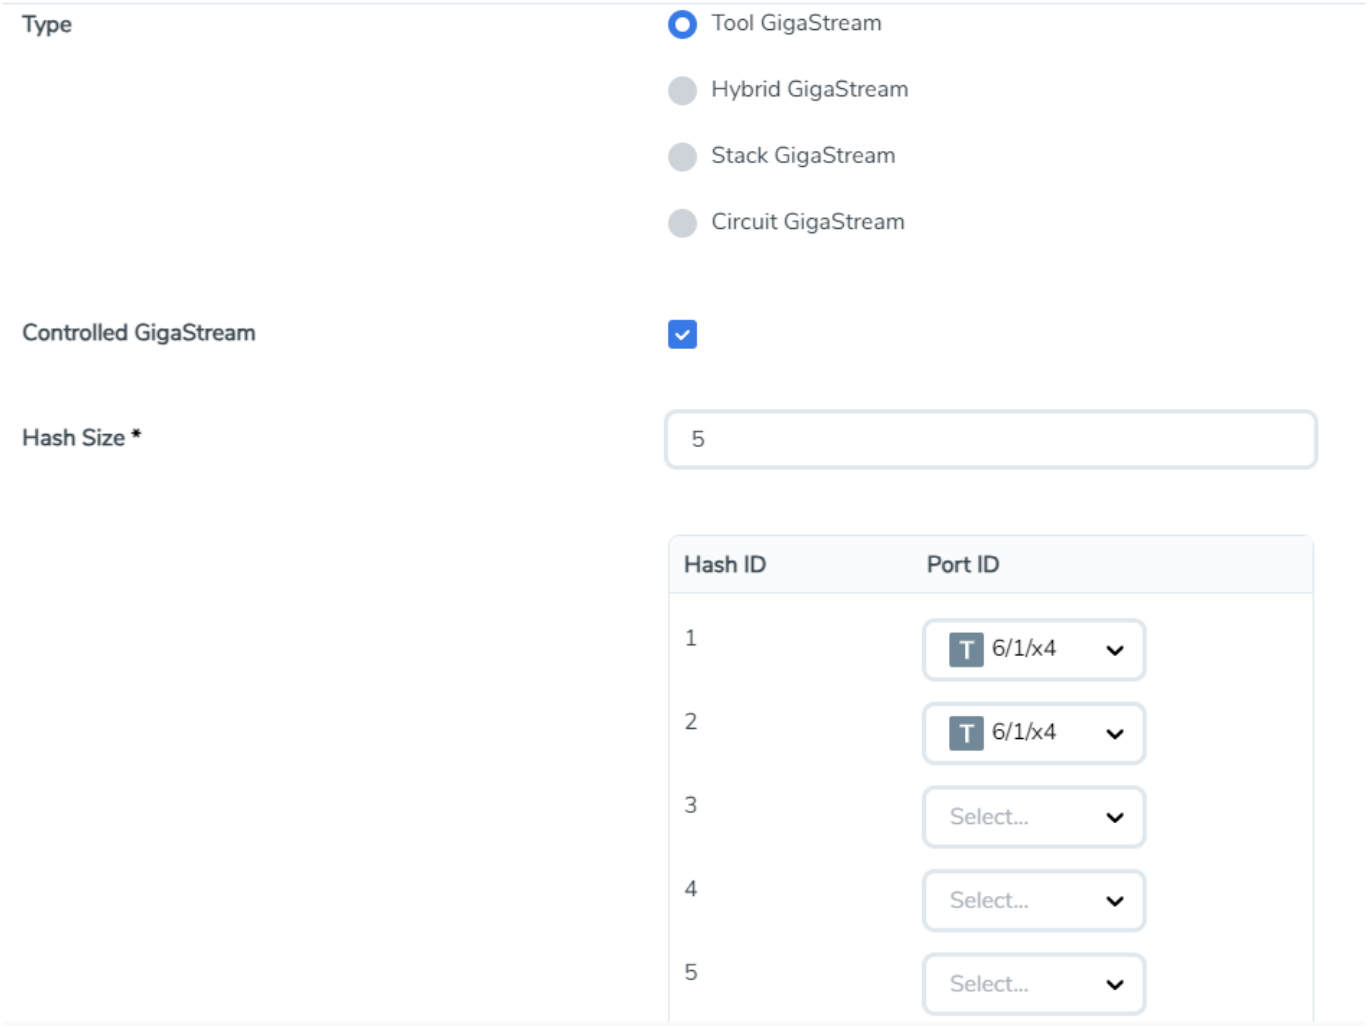


Figure 20 *Controlled GigaStream with Five Hash IDs and Port IDs*

8. Assign ports to the hash bucket IDs by clicking in each **Port ID** field and selecting a tool port.

   The example in Figure 21Ports Assigned to Hash Bucket IDs assign tool ports to hash buckets 1 though 4. Hash bucket 5 has no port assigned to it.

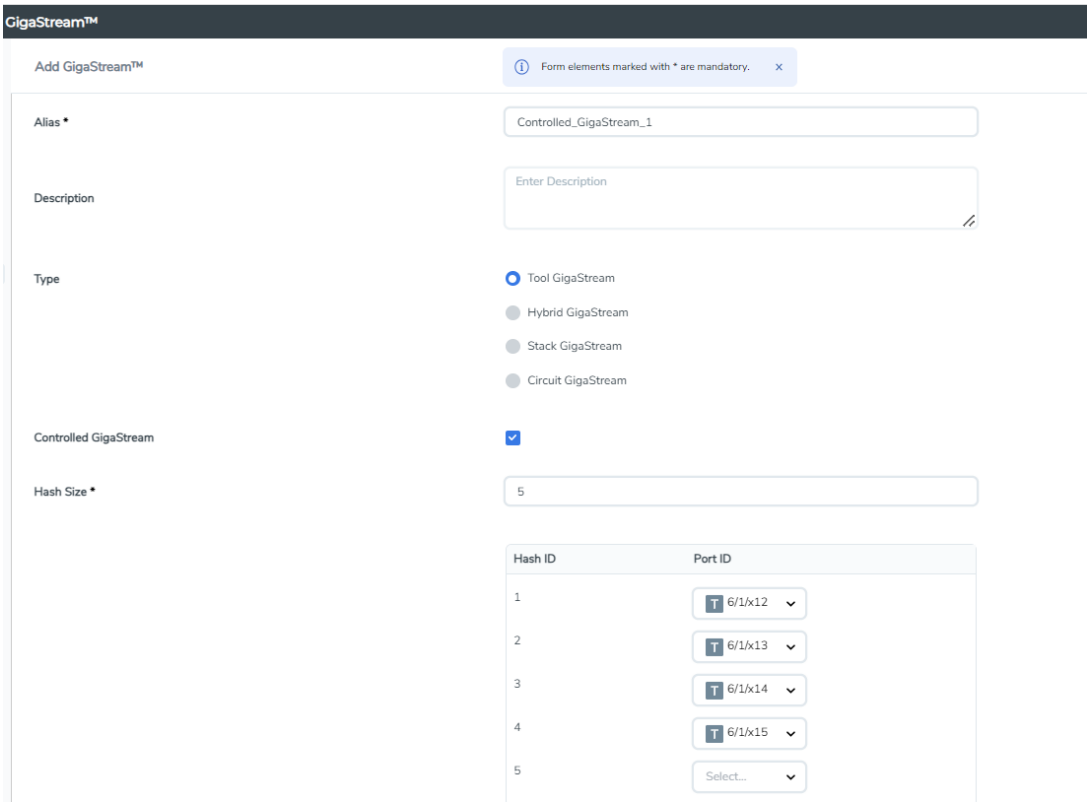


Figure 21 *Ports Assigned to Hash Bucket IDs*

9. Click **Ok** to save the configuration.

After saving the controlled GigaStream, it appears on the GigaStreams page.

NOTE: For controlled GigaStream, the GigaStream page shows a Failover Status of disabled. When a port goes down, traffic is not re-hashed. Refer to Failover and Controlled GigaStream.

Edit Controlled GigaStream

A controlled GigaStream can be edited, even when the GigaStream is attached to a map. Unlike regular GigaStream, you can make changes without deleting the map or the GigaStream.

You have the control to map unused hash bucket IDs to any tool port dynamically, without deleting the trunk. This modification of a tool port mapping to a hash bucket ID will not affect the streams flowing on the hash bucket IDs that are mapped to other ports. In addition, you can replace the mapping of any hash bucket ID to a port, dynamically.

If one of the GigaStream ports goes down, all the hash bucket IDs mapped to that port will be black holed until they are re-mapped to a new port, or until the port comes back up. This means that the packets sent to the remaining tools are unaffected.

If one port is receiving a lesser amount of bandwidth, the traffic can be reallocated to it. For example, if port x4 is underutilized, you have the control to reconfigure hash bucket ID 5 to also map to port x4. Then port x4 receives all the traffic that is hashed to hash bucket IDs 4 and 5.

You also have the flexibility to change the size of the trunk anytime, but this will require reprogramming of the whole hash table, so that might impact the existing streams.

Increasing the size of the trunk creates new hash buckets, which can be mapped to new or existing GigaStream tool ports. You can increase the bucket size per GigaStream. For example, if the bucket size is 4, you can increase it to 5.

If the size of the trunk has to be decreased, you have to take extra caution when releasing the hash bucket IDs gracefully, since they are mapped to GigaStream tool ports.

If you decrease the bucket size, empty out the bucket by unmapping buckets to ports. Also, do not reduce the hash size to less than the last occupied hash bucket ID.

NOTE: There is some packet drop associated with the following type of controlled GigaStream editing:

- adding a new port
- deleting an existing port
- changing the hash size

The Port Statistics page may display Discards in these cases, but not when additional buckets are added to the same port, or when a port goes down.

Traffic Distribution Across Controlled GigaStream

Controlled GigaStream has N buckets (where N is from 1 to 256) distributed across one or more ports, logical or physical.

Controlled GigaStream uses advanced hashing with 1 to 256 buckets. With controlled GigaStream, you define the number of buckets first, unlike with regular GigaStream.

Controlled GigaStream can manage network port bandwidth hashed to GigaStream tool ports. For example, if you are monitoring 500Gb traffic and have 10 tools, 50Gb per tool would be required. But if the tools cannot handle 50Gb, packets will be lost randomly.

With controlled GigaStream, first determine how much traffic the tools can process. For example, perhaps each tool can process 5Gb of traffic.

The formula is ingress bandwidth divided by tool capability. For example, 500Gb/5Gb = 100 tools. But if you only have 10 tools, you create a controlled GigaStream of 100 logical tools or 100 logical hash buckets, and then map only 10 of them.

Taking the number of buckets and dividing it by 100 tools (256/100 = 2.56) results in some buckets of 3, some buckets of 2.

The recommendation is to round to an even divisor of 256 (2, 4, 8, 16, 32, 64, 128, or 256). In this example, instead of using 100, use 128, so each bucket will be 2 (256/128 = 2).

Hash buckets IDs are mapped to ports as follows:

Buckets	Ports
1	x1
2	x2
3	x3
4	x4
5	x5
6	x6
7	x7
8	x8
9	x9
10	x10
11	unmapped
...	unmapped
128	unmapped

The hashing to the 10 connected tools captures the traffic associated with those sessions.

There are no tools associated with the remaining buckets, so that traffic is black holed. Unlike regular GigaStream, you do not have to allocate ports to the remaining buckets.

Multiple buckets can be mapped to one physical port as follows:

Buckets	Ports
1	x1

Buckets	Ports
2	x1
3	x2
4	x3
...	...

In this example, port x1 receives 5Gb of traffic, while ports x2 and x3 receive 2.5Gb each.

Note that the mapping does not have to be consecutive as follows:

Buckets	Ports
1	x1
2	x2
3	x3
4	x1
...	...

In this example, port x1 also receives 5Gb of traffic, while ports x2 and x3 receive 2.5Gb each.

Through the mapping of buckets to ports, you can control the overall distribution of traffic to a given port.

Failover and Controlled GigaStream

Unlike the regular GigaStream, failover will not be triggered during a port down event. With controlled GigaStream, there is no rehashing or redistribution of traffic. In other words, the sessions flowing to other tool ports will not be disturbed and do not risk becoming oversubscribed.

Controlled GigaStream maintains hashing. When a port goes down, traffic is not re-hashed, but is black holed. Unlike with regular GigaStream, you do not need to enable **Force Link Up** on ports in order to counteract the default failover protection.

Advanced Hashing

Both regular GigaStream and controlled GigaStream use advanced hashing, which lets you select the criteria on which the hash is based, such as source and destination IP address, source and destination MAC address, source and destination port, and protocol.

GigaVUE-OS nodes distribute traffic between the ports in a GigaStream based on the hashing criteria configured using the **Advanced Hash Settings** page for the selected line card or chassis. GigaStream hashing is applicable for the following port types:

- Tool port
- Hybrid port
- Circuit port

To open the Advanced Hash Settings page, select **Ports > Port Groups > GigaStream** and click **Advanced Hash Settings**. (For more details, refer to Advanced Hash Settings.) On the GigaVUE nodes, GigaStream hashing is per chassis, not per line card.

The **Advanced Hash Settings** let you select the different packet criteria used to send matching flows to the same destination port within a GigaStream.

By default, the GigaVUE HC Series node hashes traffic based on source and destination IP addresses, IP protocol, and source and destination ports.

How to Change Advanced Hash Criteria

You can select the criteria for the advanced hash algorithm by using **Advanced Hash Settings**. The advanced hash method you specify is used for all GigaStream in place on the specified line card or chassis.

Advanced Hash Settings

The Advanced Hash Settings page is where criteria for the advanced hash algorithm is set. To open the page, select **Ports > Port Groups > GigaStream** and click **Advanced Hash Settings**.

The following table describes the fields in the Advanced Hashing Settings page.

Field	Description
Box	Identifies chassis to which the advanced algorithm will be applied.

Field	Description
Slot	Identifies the line card to which the advanced hash algorithm will apply. Each line card in certain GigaVUE HC Series nodes has its own individual advanced hash algorithm. On GigaVUE-HC1, and GigaVUE-HC3, GigaStream hashing is per chassis, not per line card. For example, the slot field will only show cc1 when configuring an GigaVUE-HC3.
Type	Type can be one of the following: • **Default**: Sets the advanced hash algorithm to its default settings. By default, the advanced hash algorithm includes source/destination IPv4/IPv6 addresses and ports. • **Custom**: Clears the settings from the advanced hash and allows you to select your own criteria. • **All**: Selects all criteria. • **None**: Clears all fields from the advanced hash.
IPv4	This area of the page lets you select the following criteria. Use the toggle button to select your options: • IPv4 Source Address—Adds IPv4 source IP • IPv4 Destination Address—Adds IPv4 destination IP • IPv4 Protocol—Adds IPv4 protocol • IPv4 Source Port—Adds IPv4 source port • IPv4 Destination Port—Adds IPv4 destination port
IPv6	This area of the page lets you select the following criteria. Use the toggle button to select your options: • IPv6 Source Address—Adds IPv6 source IP • IPv6 Destination Address—Adds IPv6 destination IP • IPv6 Next Header—Adds IPv6 next header field. • IPv6 Source Port—Adds IPv6 source port • IPv6 Destination Port—Adds IPv6 destination port
Layer2	This area of the page lets you select the following criteria: • Source MAC Address—Adds L2 source MAC • Destination MAC Address—Adds L2 destination MAC • EtherType—Adds L2 ethertype field.
MPLS	This area of the page lets you select the following criteria: • MPLS Hash—Adds MPLS labels (up to three)
GTP TEID	This area of the page lets you select the following criteria: • GTP TEID—Adds GTP tunnel endpoint identifier
Ingress Port	This area of the page lets you select the following criteria:

Field	Description
	• Ingress Ports—Adds ingress port. For advanced hash fields, the "ingressport" refers to the ingress port specified in the map's from port-list. In a multi-path leaf and spine cluster, when a spine node uses the "ingressport" field for advanced hashing, it references from the port-list defined in the map(s) responsible for forwarding traffic across the spine.

Advanced Hash Examples

The following are some different advanced hash examples. Note that the advanced hash method usually combines multiple criteria.

The example in Figure 22Advanced Hash with IPv4 Source and Destination Addresses sets a **Custom** advanced hash method for slot cc1 in box ID 6 that distributes traffic based on matching IPv4 source and destination addresses.

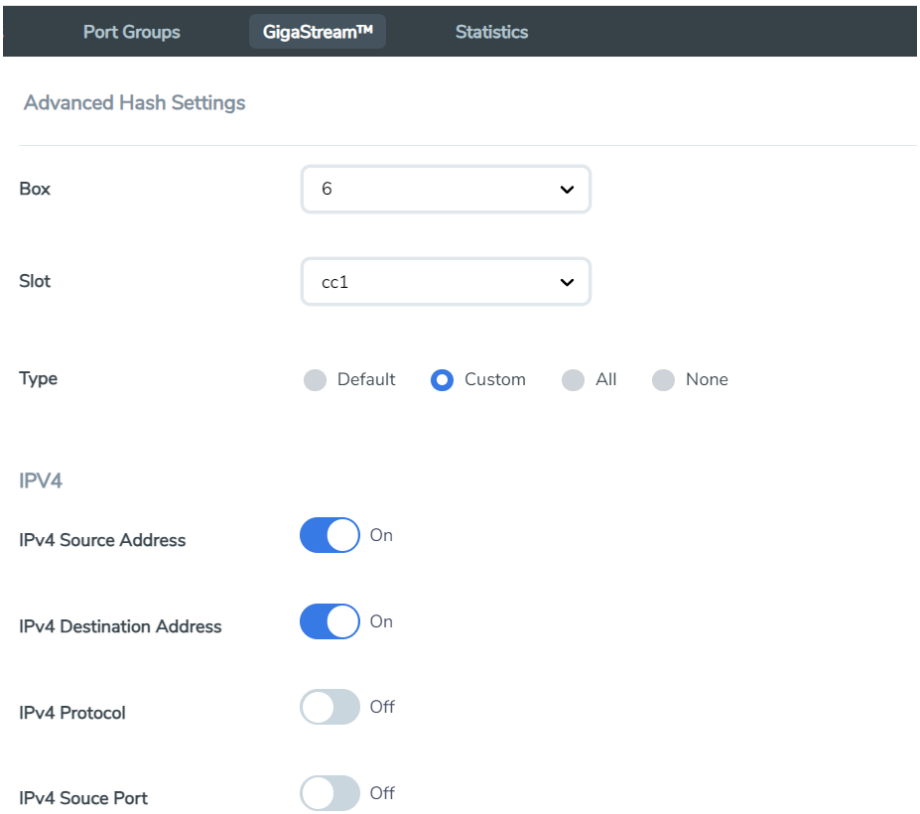


Figure 22 *Advanced Hash with IPv4 Source and Destination Addresses*

The example in Figure 23Advanced Hash Default Criteria sets the advanced hash for slot cc1 in box ID 6 to the **Default** criteria.

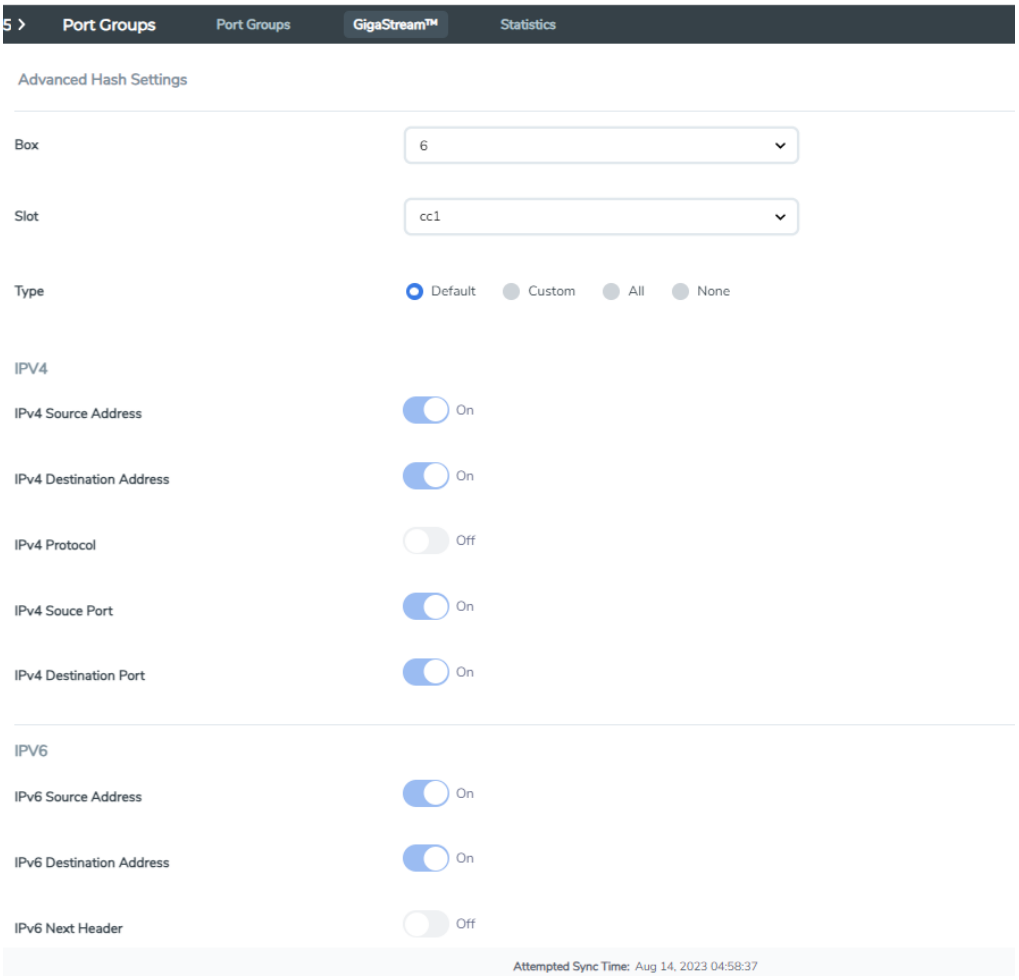


Figure 23 *Advanced Hash Default Criteria*

The example in Figure 24Advanced Hash with Source and Destination MAC Address sets a **Custom** advanced hash for box ID 3 that distributes traffic based on matching source and destination MAC addresses.

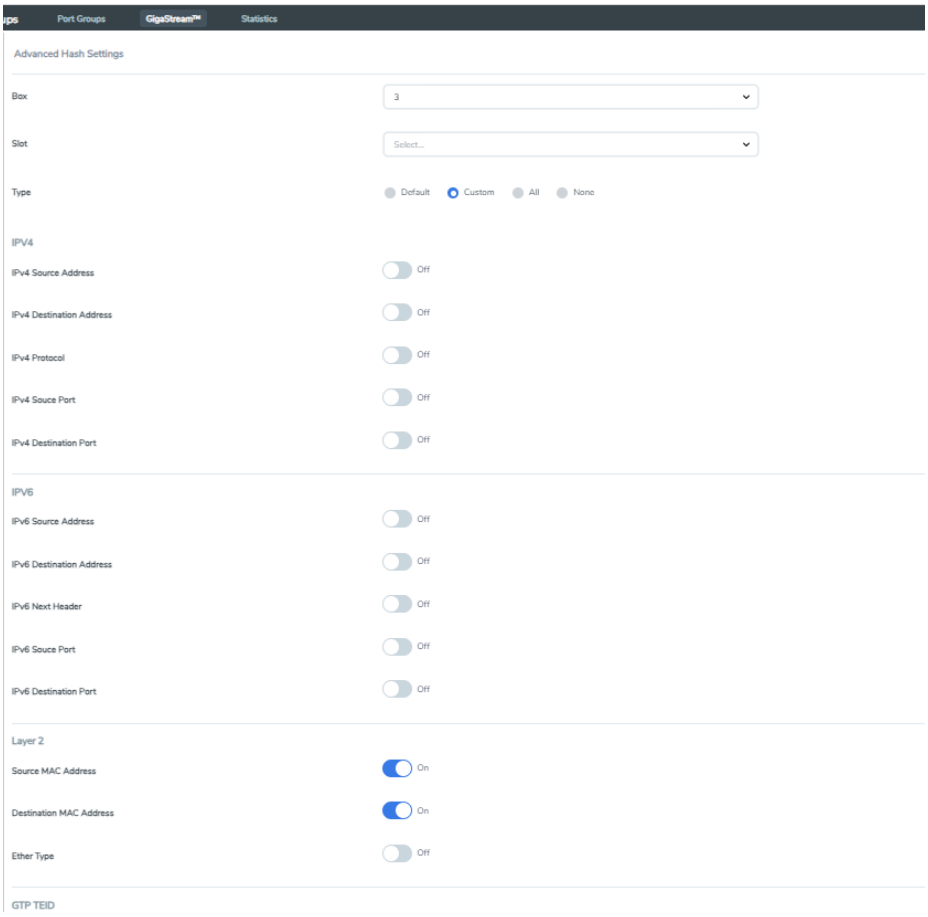


Figure 24 *Advanced Hash with Source and Destination MAC Address*

Hashing Behavior

Table 6: Hashing Behavior Based on Hash Criteria Field Combinations shows the possible hash criteria field combinations and the corresponding hashing behavior based on packet type for advanced hashing for non-MPLS packets. (Refer to Table 7: Hashing Behavior Based on Hash Criteria Field Combinations for MPLS packets.)

Table 6: Hashing Behavior Based on Hash Criteria Field Combinations

Hash Criteria Fields	Packet Type	Hashing Behavior
Source MAC Address, Destination MAC Address	MAC	Hash on Source MAC Address, Destination MAC Address
Source MAC Address, Destination MAC Address	MAC + IPv4	Hash on Source MAC Address, Destination MAC Address

Hash Criteria Fields	Packet Type	Hashing Behavior
Source MAC Address, Destination MAC Address	MAC + IPv6	Hash on Source MAC Address, Destination MAC Address
IPv4 Source Address, IPv4 Destination Address	MAC	No hash
IPv4 Source Address, IPv4 Destination Address	MAC + IPv4	Hash on IPv4 Source Address, IPv4 Destination Address
IPv4 Source Address, IPv4 Destination Address	MAC + IPv6	No hash
IPv6 Source Address, IPv6 Destination Address	MAC	No hash
IPv6 Source Address, IPv6 Destination Address	MAC + IPv4	No hash
IPv6 Source Address, IPv6 Destination Address	MAC + IPv6	Hash on IPv6 Source Address, IPv6 Destination Address
Source MAC Address, Destination MAC Address, IPv4 Source Address, IPv4 Destination Address	MAC	Hash on Source MAC Address, Destination MAC Address
Source MAC Address, Destination MAC Address, IPv4 Source Address, IPv4 Destination Address	MAC + IPv4	Hash on IPv4 Source Address, IPv4 Destination Address
Source MAC Address, Destination MAC Address, IPv4 Source Address, IPv4 Destination Address	MAC + IPv6	Hash on Source MAC Address, Destination MAC Address
Source MAC Address, Destination MAC Address, IPv6 Source Address, IPv6 Destination Address	MAC	Hash on Source MAC Address, Destination MAC Address
Source MAC Address, Destination MAC Address, IPv6 Source Address, IPv6 Destination Address	MAC + IPv4	Hash on Source MAC Address, Destination MAC Address
Source MAC Address, Destination MAC Address, IPv6 Source Address, IPv6 Destination Address	MAC + IPv6	Hash on IPv6 Source Address, IPv6 Destination Address

Hash Criteria Fields	Packet Type	Hashing Behavior
Source MAC Address, Destination MAC Address, IPv4 Source Address, IPv4 Destination Address, IPv6 Source Address, IPv6 Destination Address	MAC	Hash on Source MAC Address, Destination MAC Address
Source MAC Address, Destination MAC Address, IPv4 Source Address, IPv4 Destination Address, IPv6 Source Address, IPv6 Destination Address	MAC + IPv4	Hash on IPv4 Source Address, IPv4 Destination Address
Source MAC Address, Destination MAC Address, IPv4 Source Address, IPv4 Destination Address, IPv6 Source Address, IPv6 Destination Address	MAC + IPv6	Hash on IPv6 Source Address, IPv6 Destination Address

NOTE: No hash means that the packets will be sent to the first port in the GigaStream.

Notes and Considerations for Advanced Hashing

Refer to the following notes and considerations for advanced hashing:

- With symmetric hashing, packets with their source and destination IP addresses and Layer 4 (L4) ports interchanged will go to the same GigaStream port. It is recommended to enable source and destination IPv4 or IPv6 pairs and L4 source and destination ports.
- Symmetric hashing is enabled for all GigaStream on all GigaVUE® HC Series and TA Series nodes.
  - For non-MPLS IPv4 and IPv6 packets, the hashing is fixed to the following 3-tuple: ipsrc, ipdst, and protocol or ip6src, ip6dst, and protocol. All other traffic follows the advanced hash settings.
- Gigamon devices do not support hashing of the IP header fields when there is a PPPOE header. Only Layer 2 (L2) fields can be used for such packets.

Advanced Hashing with MPLS

Starting in software version 5.1, GigaStream MPLS hashing adds the ability to hash on MPLS labels as well as the following IP address fields inside an MPLS tunnel: **ipsrc**, **ipdst**, **ip6src**, and **ip6dst**.

Advanced hashing with MPLS is supported on all GigaVUE® HC Series and TA Series nodes.

Use the **Advanced Hash Settings** to specify MPLS, which can detect up to three MPLS labels. Packets with one to three MPLS labels can be hashed, along with IP address fields, if present. If a packet has more than three MPLS labels, IP address fields after the third MPLS label cannot be hashed. Refer to Hashing Behavior Based on Hash Criteria Field Combinations for details of the hashing behavior.

MPLS labels will be used as part of the GigaStream hash criteria if the MPLS field is configured and the packet has EtherType 0x8847.

MPLS hashing applies to the following:

- regular GigaStream
- controlled GigaStream
- stack GigaStream
- inline tool groups

Table 7: Hashing Behavior Based on Hash Criteria Field Combinations shows the possible hash criteria field combinations and the corresponding hashing behavior based on packet type for advanced hashing with MPLS packets. (Refer to Table 6: Hashing Behavior Based on Hash Criteria Field Combinations for non-MPLS packets.)

Table 7: Hashing Behavior Based on Hash Criteria Field Combinations

Hash Criteria Fields	Packet Type	Hashing Behavior
IPv4 Source Address,IPv4 Destination Address	MAC + Label1 + Label2 + Label3 + IP + L4 + Payload	Hash on IP
Source MAC Address,Destination MAC Address	outer MAC + Label1 + Label2 + Label3 + inner MAC + IP + L4 + Payload	Hash on outer MAC
Source MAC Address,Destination MAC Address	MAC + Label1 + Label2 + Label3 + IP + L4 + Payload	Hash on MAC
MPLS Hash	MAC + Label1 + Label2 + Label3 + IP + L4 + Payload	Hash on MPLS Label1, Label2, Label3
MPLS Hash	MAC + Label1 + Label2 + IP + L4 + Payload	Hash on MPLS Label1, Label 2
MPLS Hash	MAC + Label1 + IP + L4 + Payload	Hash on MPLS Label1
MPLS Hash, EtherType, IPv4 Protocol	MAC + Label1 + Label2 + Label3 + Label4 + IP + L4 + Payload	Hash on MPLS Label1, Label2, Label3
MPLS Hash,IPv4 Source Address,IPv4 Destination Address	MAC + MPLS Label1 + IP + L4 + Payload	Hash on MPLS Label1, IPv4 Source Address, IPv4 Destination Address
MPLS Hash,IPv4 Source Address,IPv4 Destination Address	MAC + Label1 + Label2 + IP + L4 + Payload	Hash on MPLS Label1, Label2, IPv4 Source Address, IPv4 Destination Address

Hash Criteria Fields	Packet Type	Hashing Behavior
MPLS Hash,IPv4 Source Address,IPv4 Destination Address	MAC + Label1 + Label2 + Label3 + IP + L4 + Payload	Hash on MPLS Label1, Label2, Label3, IPv4 Source Address, IPv4 Destination Address
MPLS Hash,IPv4 Source Address,IPv4 Destination Address	MAC + Label1 + Label2 + Label3 + Label4 + IP + L4 + Payload	Hash on MPLS Label1, Label2, Label3
MPLS Hash,IPv4 Source Address,IPv4 Destination Address	MAC + MPLS Label1 + MAC + IP + L4 + Payload	Hash on MPLS Label1
MPLS Hash,IPv4 Source Address,IPv4 Destination Address	MAC + Label1 + Label2 + MAC + IP + L4 + Payload	Hash on MPLS Label1, Label2
MPLS Hash,IPv4 Source Address,IPv4 Destination Address	MAC + Label1 + Label2 + Label3 + MAC + IP + L4 + Payload	Hash on MPLS Label1, Label2, Label3
MPLS Hash,Source MAC Address,Destination MAC Address	MAC + Label1 + Label2 + Label3 + IP + L4 + Payload	Hash on MPLS Label1, Label2, Label3
MPLS Hash,IPv4 Source Port,IPv4 Destination Port	MAC + Label1 + Label2 + Label3 + Label4 + IP + L4 + Payload	Hash on MPLS Label1, Label2, Label3
IPv4 Source Port,IPv4 Destination Port	MAC + Label1 + Label2 + Label3 + IP + L4 + Payload	No Hash

NOTES:

- No hash means that the packets will be sent to the first port in the GigaStream.
- If an MPLS packet has a router alert label as one of its labels, the router alert label is skipped and the other available labels are used for hashing. For example, if a packet has four labels and the second label is the router alert, the first, third, and fourth labels are used for hashing.
- Starting with 6.1 release, all platforms (excluding TA25, TA25E, HC1P) support symmetric hashing of MPLS and non-MPLS mixed IP traffic (both IPv4, IPv6) using Source IP address, Destination IP address (2-tuple), Source Port, Destination Port (4-tuple), Protocol/Next-header (5-tuple) hash field combinations.

  TA400 supports this with or without MPLS Header Stripping, other platforms require MPLS Header Stripping. TA400 supports this with MPLS layer2 as well as MPLS layer3 IP traffic, other platforms support this with MPLS layer 3 traffic only.

Advanced Hashing with GTP TEID

Starting in software version 5.2, GigaStream GTP TEID hashing adds the ability to hash on GTP tunnel endpoint identifiers (TEIDs). GPRS Tunneling Protocol (GTP) is an IP/UDP-based protocol for mobile data.

The TEID field in a GTP header is a unique identifier for mobile subscribers and is used to multiplex different connections on the same GTP tunnel. Use GTP TEID advanced hashing to load balance GTP packets across all GigaStream ports.

Advanced hashing with GTP TEID is supported on GigaVUE® HC Series and TA Series nodes, with the following distinctions:

- The following nodes are supported: GigaVUE-HC1, GigaVUE-HC1-Plus, GigaVUE-HCT, GigaVUE-HC3 (with control card version 1 or 2), GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA100, GigaVUE-TA200, GigaVUE-TA200E, and GigaVUE-TA400.

Use the **Advanced Hash Settings** to specify the **GTP TEID** field. It must be specified with one of the port source and destination pairs: either **IPv4 Source Port** and **IPv4 Destination Port** for IPv4 or **IPv6 Source Port** and **IPv6 Destination Port** for IPv6.

When the **GTP TEID** field is configured, GTP packets will use it for hashing along with configured IP fields, instead of the Layer 4 (L4) source and destination for GTP packets. The hashing for all non-GTP packets will be based on L4 source and destination port.

GTP TEID hashing is supported only for GTP-User packets, version 1 (V1), with L4 source and destination port 2152. The hashing functionality works only for non-fragmented packets.

Refer to Hashing Behavior Based on Hash Criteria Fields: GTP TEID. for details of the hashing behavior.

GTP TEID hashing applies to the following:

- regular GigaStream
- controlled GigaStream
- stack GigaStream (except GigaVUE-HC3)
- inline tool groups

NOTE: GTP TEID hashing is not supported for stack GigaStream on GigaVUE-HC3 in this software version due to the 3-tuple limitation listed in Notes and Considerations for Advanced Hashing.

Table 8: Hashing Behavior Based on Hash Criteria Fields: GTP TEID. shows the possible hash criteria field combinations and the corresponding hashing behavior based on packet type for advanced hashing with GTP packets.

Table 8: Hashing Behavior Based on Hash Criteria Fields: GTP TEID.

Hash Criteria Fields	Packet Type	Hashing Behavior
IPv4 Source Address,IPv4 Destination Address,IPv4 Source Port,IPv4 Destination Port	MAC + IP + L4 + Payload	Hash on IPv4 Source Address,IPv4 Destination Address,IPv4 Source Port,IPv4 Destination Port
IPv4 Source Address,IPv4 Destination Address,IPv4 Source Port,IPv4 Destination Port,GTP TEID	MAC + IP + L4 + GTP + Payload	Hash on IPv4 Source Address,IPv4 Destination Address, GTP TEID
IPv4 Source Address,IPv4 Destination Address,IPv4 Source Port,IPv4 Destination Port,GTP TEID	MAC + IP + L4 + Payload	Hash on IPv4 Source Address,IPv4 Destination Address,IPv4 Source Port,IPv4 Destination Port
IPv6 Source Address,IPv6 Destination Address,IPv6 Source Port,IPv6 Destination Port	MAC + IP + L4 + Payload	Hash on IPv6 Source Address,IPv6 Destination Address,IPv6 Source Port,IPv6 Destination Port
IPv6 Source Address,IPv6 Destination Address,IPv6 Source Port,IPv6 Destination Port,GTP TEID	MAC + IP + L4 + GTP + Payload	Hash on IPv6 Source Address,IPv6 Destination Address, GTP TEID
IPv6 Source Address,IPv6 Destination Address,IPv6 Source Port,IPv6 Destination Port,GTP TEID	MAC + IP + L4 + Payload	Hash on IPv6 Source Address,IPv6 Destination Address,IPv6 Source Port,IPv6 Destination Port

Packet Distribution and the Advanced Hash Algorithm

- When an **IPv4 Fragmentation** map rule is used to send traffic to an advanced hash tool GigaStream, all fragments are consolidated to a single port within the GigaStream.
- Packets with multiple VLAN tags (such as Q-in-Q) will experience uneven traffic distribution. For this traffic, GigaSMART load balancing is recommended.

Weighted GigaStream

Weighted GigaStream provides you the ability to distribute traffic to the ports by assigning either an equal weight or a custom weight to the ports. You can assign custom weight in percentage or ratio. If a port in a weighted GigaStream goes down, the traffic from the port will be redistributed to other healthy ports in the weighted GigaStream. The port assigned with maximum weight receives more traffic than the ports assigned with lesser weight.

Weighted GigaStream is supported on the following:

- All GigaVUE HC Series and GigaVUE TA Series nodes.
- Regular tool GigaStream, regular hybrid GigaStream, and regular circuit GigaStream.

Use the **Weighting** and **Drop Weight** fields in the GigaStream configuration page to configure a weighted GigaStream. For instructions, refer to Configure Regular GigaStream.

You can also choose to rehash the traffic when you find that the traffic distribution is not ideal. When you rehash the traffic, GigaVUE-FM reassigns the hash buckets to the ports. For example, the following table shows that the ports are assigned with sequential hash buckets:

Port	Hash Buckets
x1	1, 2, 3, 4
x3	5, 6, 7, 8
x4	9, 10, 11, 12

In such cases, the traffic distribution may not be ideal. You can choose to rehash the traffic. The following table shows how the hash buckets are reassigned when you rehash the traffic:

Port	Hash Buckets
x1	1, 4, 7, 10
x3	2, 5, 8, 11
x4	3, 6, 9, 12

GigaStream Rules and Maximums

The following rules apply to regular GigaStream and controlled GigaStream:

GigaStream Rule	Description
Port Location	All participating ports must be on the same GigaVUE node. On the GigaVUE-HC3, GigaStream can be across modules.
Speed Requirements	▪ All participating ports must be running the same speed (1Gb, 10Gb, 40Gb. or 100Gb) and must use the same port types (for example, all g, x, q, or c). ▪ A stack GigaStream must consist of ports with 10Gb speed or higher. ▪ The system will not let you change the speed of any port participating in a

GigaStream Rule	Description
	tool GigaStream. Keep in mind that the only ports that allow speed changes through are g***x*** ports. ▪ You can use g***x*** ports in a regular tool GigaStream, but only with ports running at the same speed (and no slower than 1000Mb).
Addressing	Once a port belongs to a GigaStream, it must be addressed by its GigaStream alias. It can no longer be addressed as an individual port. For example, if tool port 1/1/x4 is part of a tool GigaStream, the GigaVUE® HC Series node prevents you from using it as the destination for a map rule.
Stack Ports	Stacking ports must be 10Gb or higher. Therefore, the Maximum Stack Ports per GigaStream for any 1Gb port is N/A in : Maximum Ports per GigaStream to.GigaStream Rules and Maximums
SFP+	For SFP+ ports that can operate at 10Gb or 1Gb, refer to the values in the Maximum Tool Ports per GigaStream column for the 10Gb Ports rows in : Maximum Ports per GigaStream to GigaStream Rules and Maximums.

For GigaStream maximums, refer to Maximum Ports per GigaStream, which have been updated in software version 5.1.

Maximum Ports per GigaStream

Table 9: : Maximum Ports per GigaStream lists the maximum ports per GigaStream for GigaVUE nodes.

Table 9: : Maximum Ports per GigaStream

Platform	Maximum Stack Ports per GigaStream	Maximum Tool Ports per GigaStream
GigaVUE-TA 10	32	256
GigaVUE-TA 40	64	256
GigaVUE-TA 100	64	256
GigaVUE-TA200	64	256
GigaVUE-TA200E	128	128
GigaVUE-TA25	56	56
GigaVUE-TA25E	64	64
GigaVUE-TA400	128	128

Platform	Maximum Stack Ports per GigaStream	Maximum Tool Ports per GigaStream
GigaVUE-TA400E	128	128
GigaVUE-HC1-PLUS	64	256
GigaVUE-HCT	64	256
GigaVUE-HC1	64	256
GigaVUE-HC 3	64	256

Port Statistics and Counters

This section describes the counters displayed for the Port Statistics information. This page provides information similar to the output from the **show port stats** command from the CLI.

The major sections in This section include:

- Display Port Statistics
- How to Clear Traffic Counters

Display Port Statistics

From the **Ports** page, you can view the port statistics for either of the following:

- A single port. Refer to Display Port Statistics for a Single Port
- All the ports. Refer to Display Statistics for All Ports

Display Port Statistics for a Single Port

From the **Ports > Ports > All Ports**, select any port by clicking on the row. The quick view window that appears provides port statistics for that particular port. Each field is color-coded in the graphical representation.

By hovering over the graph, numerical value for each of the data points is visible as shown in Figure 25Numerical Values for Data Points on the Statistics Graph.

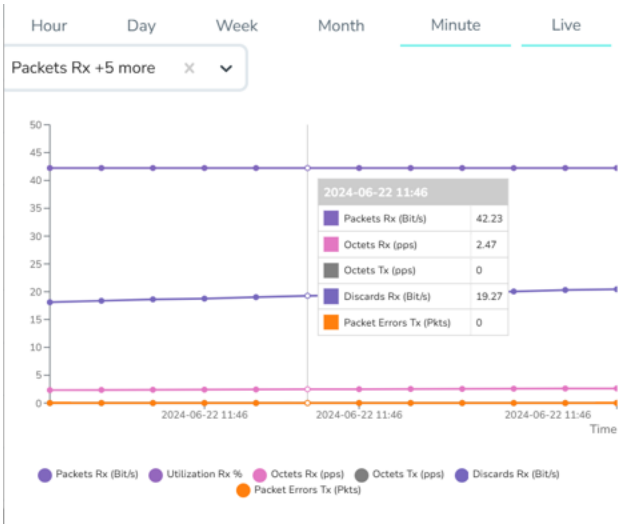


Figure 25 *Numerical Values for Data Points on the Statistics Graph*

NOTE: GigaVUE-FM will display 9-11 data points in the graph.

You can modify the time lapse for measuring various data points by selecting the Minutes, Hour, Day, or Week button.

Table 10: Port Statistics Definitions describes the port statistics available.

Table 10: Port Statistics Definitions

Counter	Definition	Notes
Packet Errors	Total Error Packets Received or Transmitted This indicates hardware detected errors. Error packets include undersize, FCS/CRC, MTU exceeded, fragments, and oversize packets. All packets that list under this counter are discarded and not processed further.	Excludes oversize packets without FCS/CRC. Packets larger than the MTU setting arriving on a network port are counted twice in the counter. For example, 1000 oversize packets would show up as 2000. This double-counting only happens with Oversize error packets.
Discards	Discards Received or Transmitted	Discards are counted in the following cases: • Traffic arriving at a network port that is not logically connected using a map or map passall. • Map rules. • Packets on a tool port enabled by force link up. • Drop rules in an egress filter. • Pause frames.
Packet Drops	Total Dropped Packets Received or Transmitted	Packets are dropped when a port's bandwidth is exceeded due to oversubscription. Packets are dropped when they reach the port but

Counter	Definition	Notes
		before they are sent out.
Octets	Total Bytes Received or Transmitted Includes all valid and error frames with the exceptions noted in the adjacent columns.	Excludes undersize frames.
Buffer Usage	Percentage of buffer space used by packets transmitted or received	The buffer is used when the port reaches 100 percent utilization such as during a microburst. If the buffer reaches 100 percent utilization, packets will be dropped.
Utilization	Percentage of port utilization by packets received or transmitted	
Packets	Total Packets Received or Transmitted Excludes multicast packets, broadcast packets, packets with FCS/CRC errors, MTU exceeded errors, oversize packets, and pause packets.	Excludes packets with FCS/CRC errors.

Display Statistics for All Ports

To view the statistics for all ports select, **Ports > Ports > Statistics**. The **Ports Statistics** page displays, which shows a table with the statistics for each port as shown in Figure 26Port Statistics for All Ports. For the definitions of the statistics shown in the table, refer to Table 10: Port Statistics Definitions.

Port ID	Octets		Octets /sec		Unicast Packets		Non-Unicast Packets		Packets /sec		Packet Drops	Discards		Error		Utilization	
	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Rx	Tx	Rx	Tx	Rx	Tx
1/1/x20	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
vTunnelEn...	2.38 G	83.03 K	36.42 K	0	2.26 M	1.18 K	310.9 K	0	49	0	0	0	0	0	0	0.03	0
toRSASecu...	290.98 K	0	0	0	0	0	1.42 K	0	0	0	0	1.42 K	0	0	0	0	0
vTunnelEn...	2.38 G	0	36.42 K	0	2.26 M	0	310.89 K	0	49	0	0	2.57 M	0	0	0	0.03	0
Demo_To...	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1/2/q1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	-	-
1/2/q2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	-	-
1/2/q3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	-	-
1/2/q4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	-	-

Showing **1** - **30** of **72** 1 2 3

Figure 26 *Port Statistics for All Ports*

The information on the statistics page can be filtered as well as downloaded to an Excel spreadsheet. To export the statistics, click **Export**. The statics table is downloaded with a filename in the format Port_Stats_<yyyymmddhhmmss>; for example, Port_Stats_ 20161003172336.

To filter the port statistics, click **Filter**. A Filter Quick View opens, where you can specify how to filter ports displayed on the Statistics page.

The criteria that you can use to filter the port statistics is as follows:

Criteria	Description
Box/Slot ID	Display only those ports that match the specified box and slot IDs.
Port Alias	Display port with the specified alias.
Port ID	Display ports with specified number in the port ID. For example, if you specify 3 the result will also display ports that include the number 3, 13, 23, 30, and so on.
Type	Display ports with the specified port type. Select one of the following: • Network - ports that receive incoming traffic from various sources. • Tool - egress ports where data is sent out to monitoring and analysis tools. • Inline Network - ports to which end-point devices are attached in an Inline Bypass configuration. • Circuit port- ports to send or receive traffic between clusters. • Inline Tool - ports that are used to connect inline security or monitoring tools in an Inline Bypass solution. • GigaSMART - ports are internal ports on GigaSMART line cards or modules used to power GigaSMART features. • Hybrid - physical port that serves a dual function as both an indirect traffic source port and a tool port. • Stack - ports to connect multiple GigaVUE nodes in a unified cluster. • Backplane - transfer traffic from front ports to the GigaSMART engines through XAUII interfaces. • XAUII - facilitate communication between GigaSMART engines, particularly when a second-level map is implemented.
Admin Status	Display ports based on their current admin status. The possible selections are: • All — display ports with a status of Enabled or Disabled. This

Criteria	Description
	is the default. • Enabled — display ports with admin enabled • Disabled — display ports with admin disabled
Link Status	Display ports based on their current link status: The possible selections are: • All — display ports with a status of Up or Down. This is the default. • Up— display ports with a link status of up. • Down— display ports with a link status of down.

How to Clear Traffic Counters

To clear the traffic counters, do the following:

1. On the left navigation pane, go to  **> Ports> Ports.**
2. Go to **Fabric Statistics** page and select a port.
3. Click **Clear** to clear the traffic counters in the Hardware.

The traffic counters will now be cleared.

Header Stripping

With the Header Stripping functionality enabled on the network and hybrid ports, the GigaVUE devices can identify and remove headers from tagged or tunneled (encapsulated) packets. Refer to the following sections for details:

- About VXLAN Header Stripping
- About MPLS Header Stripping

About VXLAN Header Stripping

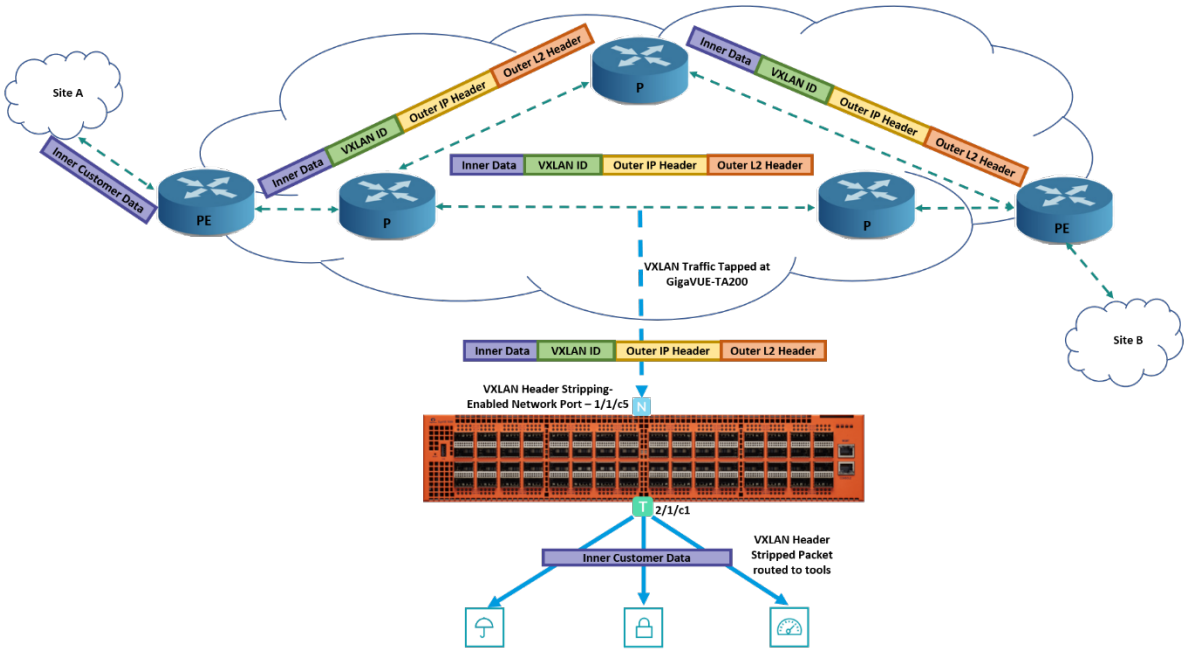
The VXLAN Header Stripping functionality can identify and remove headers from the VXLAN tagged packets that are tapped from the VXLAN-based enterprise networks and are routed to the respective tools for analysis. This functionality is useful when working with tools that either cannot recognize the VXLAN headers or must engage in additional processing to

analyze the VXLAN traffic. The GigaVUE-OS device is configured with the required traffic intelligence capability to strip the VXLAN header from the incoming packets. It then sends the inner payload to the tools based on the map rules configured.

The incoming VXLAN packets are discarded until the GigaVUE-OS device acquires the capability to remove the VXLAN header.

Once the required intelligence is acquired, Header Stripping can happen at line rate.

The following figure illustrates the VXLAN Header Stripping functionality:



In this diagram, the VXLAN encapsulated traffic from the network is tapped on the network port, 1/1/c5 in the GigaVUE-TA200 device. The VXLAN Header Stripping functionality is enabled in this network port. Based on the traffic intelligence capability, the GigaVUE-TA200 device strips the VXLAN header from the incoming packets that has the L4 destination port as 4789. You can choose to configure the L4 destination port at the node level. You can also define non-standard port (47889) instead of standard port (4789). The non-standard port is also known as iVXLAN. The GigaVUE-TA200 device routes the inner payload to the respective tools based on the map rules configured.

Any other traffic that enters the VXLAN Header Stripping-enabled network port will also be processed similar to a normal by-rule map.

The following table provides the capabilities available for VXLAN tunnel decapsulation as against the capabilities available for VXLAN Header Stripping:

Capabilities	With GigaSMART	Without GigaSMART
IPv4 support	Yes	Yes
IPv6 support	Yes	No
Header stripping on GigaVUE-TA Series	No	Yes
Header stripping at line rate	No	Yes

GigaVUE-OS has a scan interval between 300 to 1000000 seconds to optimize the VXLAN ID processing capability. You must configure 0 to disable the scan interval.

NOTE: GigaVUE-OS restarts the VXLAN traffic intelligence capability on every reload.

VXLAN Tunnel Decapsulation Versus VXLAN Header Stripping

The difference between the VXLAN tunnel decapsulation and VXLAN Header Stripping is that in the case of VXLAN tunnel decapsulation, the traffic originates from and terminates at the GigaVUE-OS devices. So, the GigaVUE-OS devices are aware of the VXLAN IDs based on which the traffic is decapsulated. In the case of VXLAN Header Stripping, the GigaVUE-OS devices are configured with traffic intelligence capability to strip the VXLAN header from the incoming packets with any VXLAN IDs.

VXLAN Header Stripping – Rules, Notes, and Limitations

Keep in mind the following rules and notes when working with VXLAN Header Stripping:

- VXLAN Header Stripping is supported on GigaVUE-HC1-Plus, GigaVUE-HCT, GigaVUE-HC1, GigaVUE-HC3, GigaVUE-TA100, GigaVUE-TA200, GigaVUE-TA200E,GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA400, and GigaVUE-TA400E devices.
- The destination IP based statistics is not supported on GigaVUE-TA25E, GigaVUE-HC1-Plus, GigaVUE-TA25, and GigaVUE-HCT.
- The VXLAN header stripped packet does not match the VLAN qualifier in maps due to chip set limitation in GigaVUE-TA25E, GigaVUE-HC1-Plus, GigaVUE-TA25, and GigaVUE-HCT.
- Gigamon®'s traffic intelligence processing may lead to initial packet drops. This is not applicable for GigaVUE-TA400.
- On a network or hybrid port that taps the network traffic, you can enable either VXLAN Header Stripping or MPLS Header Stripping, but you cannot enable both the functionalities.

- Network ports configured with VXLAN Header Stripping and MPLS Header Stripping functionalities cannot be part of the same map. You must create separate maps for these ports.
- After a header is removed from the packet, FCS is recomputed by the hardware and not by the GigaVUE-OS.
- VXLAN Header Stripping is not supported on network ports that are part of pass-all maps except on GigaVUE-TA400.
- VXLAN Header Stripping is not supported with IPv6 addresses except on GigaVUE-TA400.
- You cannot associate an IP interface with a network or hybrid port that is enabled with VXLAN Header Stripping.
- Filter rule is not supported on hybrid port that is enabled with VXLAN Header Stripping.
- Ingress VLAN tagging is not supported.
- VXLAN Header Stripping is not supported for Q-in-Q traffic except on GigaVUE-TA400.
- You cannot enable VXLAN Header Stripping on a port that is part of a port-pair.
- Reassembly of fragmented packets after VXLAN Header Stripping is not supported.
- A maximum of up to 4096 dynamic VXLAN IDs are supported for VXLAN Header Stripping. On GigaVUE-TA400, and GigaVUE-TA400E, all VXLAN IDs are supported.
- If a map has both Header Stripping-enabled ports and other network ports, the VXLAN traffic that enters the other network ports will not be sent to the shared collector except on GigaVUE-TA400.
- VXLAN Header Stripping does not work if you configure a map with any rule that includes the qualifying attributes of the VXLAN header because such rules override the traffic intelligence capability. For example, if you configure a map with pass rule as IPv4 Destination, IPv4 Source, MAC Destination, or MAC Source and the source port of the map is overlapped/matched with VXLAN headers, the Header Stripping functionality does not work. This is an exception for GigaVUE-TA400.
- VXLAN Header Stripping does not work if you configure a map with only drop rules and choose the **Pass Traffic** option so that the traffic is passed through the port when there are no matching rules. For more information, refer to Map Types. This is an exception for GigaVUE-TA400.
- Following table provides the maximum number of static IP addresses that can be configured for each platform for the VXLAN Header Stripping functionality:

Platform	Maximum number of static IP addresses supported
Platform	Maximum number of static IP addresses supported
GigaVUE-HC1	3966
GigaVUE-HC3, GigaVUE-TA100, and GigaVUE-TA200	1918

Configure VXLAN Header Stripping

Before you configure VXLAN Header Stripping, refer to *VXLAN Header Stripping – Rules, Notes, and Limitations*.

The following table summarizes the required tasks to configure VXLAN Header stripping to strip the incoming VXLAN packets:

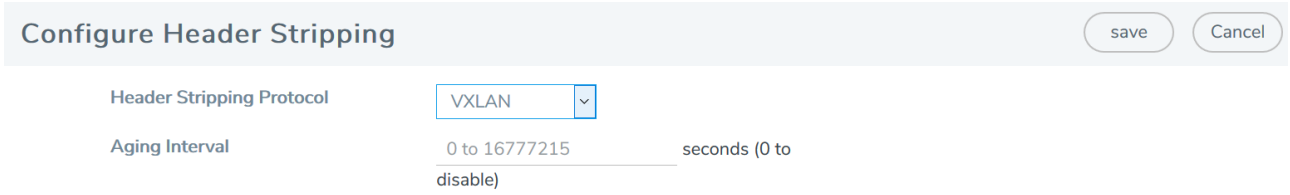
S.No	Task	Refer to..
1.	Configure the required network or hybrid port.	Configure Ports
2.	Enable VXLAN Header Stripping on the network or hybrid port that you configured in task 1.	Enable Header Stripping Protocol on Ports
3.	It is recommended that you configure the aging interval for the VXLAN Header Stripping to refresh the traffic intelligence capability for the device.	*Configure Ageing Interval for VXLAN Header Stripping*
4.	Configure a new map. Keep in mind the following details when you configure a map: • Ensure that you select the network or hybrid port on which you enabled VXLAN Header Stripping as the source port of the map. • Ensure that you do not select any rule that includes the qualifying attributes of the VXLAN header because such rules override the traffic intelligence capability. For more details, refer to *VXLAN Header Stripping – Rules, Notes, and Limitations*.	Create a new map
5.	View the VXLAN and destination IP statistics to know the number of packets transmitted for a VXLAN ID and a destination IP address, and the transmitted bytes per packet.	*View Header Stripping Statistics*

Configure Ageing Interval for VXLAN Header Stripping

It is recommended that you configure the ageing interval to refresh the traffic intelligence capability based on which GigaVUE-OS devices strips the VXLAN headers from the incoming VXLAN traffic.

To configure the ageing interval:

1. From the left navigation pane, go to **System > Chassis**. Toggle to view the list view.
2. Select the Box ID on which you want to configure the Header Stripping protocol.
3. From the **Actions** drop-down list, select **Configure Header Stripping**. The Configure Header Stripping page appears.
4. From the **Header Stripping Protocol** drop-down list, select **VXLAN**.
5. In the **Aging Interval** field, enter the interval in seconds to refresh the traffic intelligence capability for the chassis.
6. Click **Save**.



View Header Stripping Statistics

To view the Header Stripping statistics for a VXLAN ID and destination IP address:

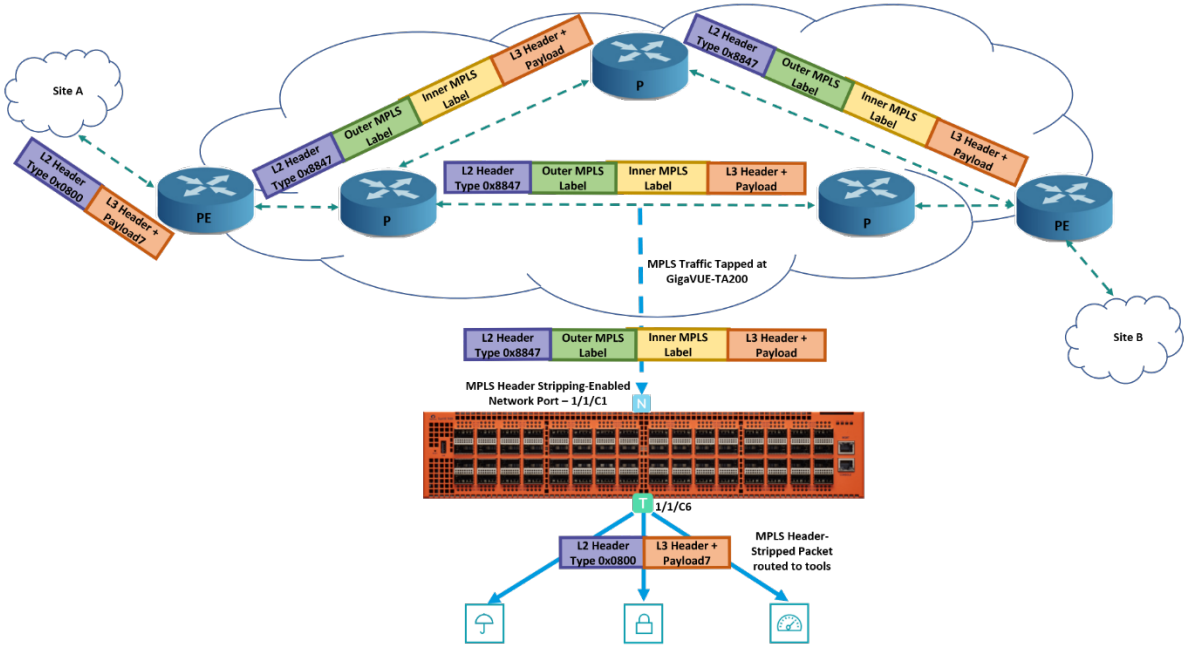
1. From the device view, go to **Ports > Ports > Header Stripping Statistics**.
2. From the **Statistics** drop-down list, select one of the following options:
   - **IP Statistics** – View the number of packets transmitted and the bytes per packet for a destination IP address.
   - **VXLAN Statistics** – View the number of packets transmitted and the bytes per packet for a VXLAN ID.

NOTE: The destination IP based statistics is not supported on GigaVUE-TA25E, GigaVUE-HC1-Plus, GigaVUE-TA25, and GigaVUE-HCT.

About MPLS Header Stripping

The MPLS Header Stripping functionality can identify and remove headers from the MPLS traffic tapped at the network and routed to the respective tools for analysis. The GigaVUE-OS device strips the MPLS header based on the MPLS label configured for the device. The outer and the inner MPLS labels must be configured for the device. The MPLS traffic is then routed to the respective tools based on the map rules configured to match the MPLS payload.

The following figure illustrates the MPLS Header Stripping functionality.



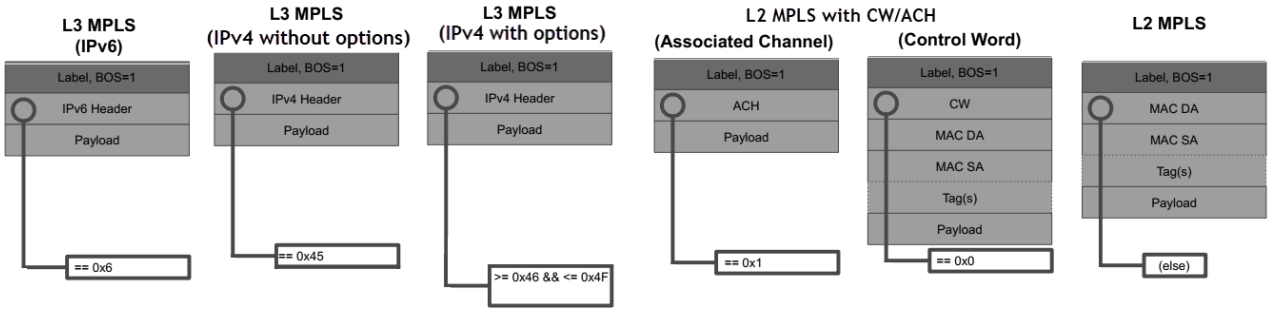
In this diagram, the L3-MPLS encapsulated traffic from the network is tapped on the network port, 1/1/c1 in the GigaVUE-TA200 device. The L3-MPLS Header Stripping functionality is enabled on this network port. Based on the MPLS label configured for the GigaVUE-TA200 device, the device strips the MPLS header from the incoming packets and routes the MPLS payload to the respective tools based on the map rules configured.

Any other traffic that enters the L3-MPLS Header Stripping-enabled network port will also be processed similar to a normal by-rule map.

Advanced MPLS Header Stripping for GigaVUE-TA400

The GigaVUE-TA400 introduces advanced MPLS Header Stripping capabilities, enabling support for all Layer 2 and Layer 3 MPLS traffic types. Header stripping applies only to the outer MPLS header in nested MPLS scenarios, and up to 7 MPLS labels can be stripped. You can configure the Advanced MPLS Header Stripping options as required. If multiple options are configured, stripping logic will be applied per the diagram below.

NOTE: When L3 MPLS options are configured, then any traffic that is not L3 MPLS will be as is without stripping.



Below are the processing rules for MPLS packets based on specific header values after the Bottom of Stack (BOS) label:

L3 MPLS IPv6 - If the first nibble after the BOS label is 0x6, it strips MPLS headers.

L3 MPLS IPv4 without Options - If the first byte after the BOS label is 0x45, it strips MPLS headers.

L3 MPLS IPv4 with Options - If the first byte after the BOS label is in the range 0x46 to 0x4F, it strips MPLS headers.

L2 MPLS with CW/ACH

- If the first nibble after the BOS label is 0x1, it strips the outer ethernet and MPLS headers along with Associated Channel Header(ACH).
- If the first nibble after the BOS label is 0x0, it strips the outer ethernet and MPLS headers along with Control Word (CW).

L2 MPLS without CW/ACH - If the first nibble after the BOS label does not have a value matching in the list [0x0, 0x1, 0x45, 0x46 to 0x4F, and 0x6], it strips the outer ethernet and MPLS headers.

The table below describes the Header strip configuration and its potential misclassification.

Header strip configuration	MPLS packet type classification
L2 MPLS without CW/ACH, L3 MPLS IPv4 without options	L2 MPLS packets without CW/ACH having Packet DMAC starting with 0x45 will be misidentified as MPLS IPv4 without options.
L2 MPLS without CW/ACH, L3 MPLS IPv4 with options	L2 MPLS packets without CW/ACH having Packet DMAC starting with ranges 0x46-0x4F will be misidentified as MPLS IPv4 with options.
L2 MPLS without CW/ACH, L3 MPLS IPv6	L2 MPLS packets without CW/ACH having Packet DMAC starting with 0x6 will be misidentified as MPLS IPv6.
L2 MPLS with CW/ACH, L2 MPLS without CW/ACH	L2 MPLS packets without CW/ACH having Packet DMAC starting with 0x0 or 0x1 will be misidentified as L2 MPLS with CW/ACH.
L2 MPLS without CW/ACH	L2 MPLS packets with CW/ACH will be misidentified as L2 MPLS without CW/ACH and only the Outer Ethernet, MPLS headers will be stripped without the PW CW from the packet.

To configure Advanced MPLS Header Stripping options, refer to Enable Header Stripping Protocol on Ports.

MPLS Header Stripping – Rules, Notes, and Limitations

Keep in mind the following rules and notes when working with MPLS Header Stripping:

- MPLS Header Stripping is supported on GigaVUE-HC1-Plus, GigaVUE-HCT, GigaVUE-HC1, GigaVUE-HC3, GigaVUE-TA100, GigaVUE-TA200, GigaVUE-TA200E,GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA400, and GigaVUE-TA400E devices.
- On GigaVUE-TA25, the outer tag is also stripped along with MPLS labels of incoming tagged MPLS traffic.
- On GigaVUE-HC1-Plus, GigaVUE-TA25, GigaVUE-TA25E, and GigaVUE-HCT devices, ingress VLAN tag functionality will not be supported on MPLS header strip enabled ports.
- In legacy stacking mode, the MPLS header stripped packets cannot traverse a cluster stack link and reach the destination ports in remote node on GigaVUE-TA25E, GigaVUE-HC1-Plus, GigaVUE-TA25, and GigaVUE-HCT. It is expected that the source and destination ports reside in the same device. This is not a limitation in case of ETAG

mode. The customer VLAN gets stripped off in ETAG mode when the source port is enabled for MPLS Header Stripping in GigaVUE-TA25E, GigaVUE-HC1-Plus, GigaVUE-TA25, and GigaVUE-HCT.

- In GigaVUE-TA25E, the MPLS traffic received on the MPLS header strip enabled port for which the corresponding MPLS labels are not programmed in the device already, hits the shared collector and gets forwarded to tools with an additional VLAN tag. This is not applicable in Etag
- MPLS Header Stripping is supported only on traffic with a maximum of up to two MPLS labels (inner and outer MPLS labels) configured for a device. On GigaVUE-TA400, MPLS Header Stripping is supported with a maximum of seven MPLS labels per packet.
- MPLS Header Stripping is not supported with IPv6 addresses except on GigaVUE-TA400.
- A maximum of up to 4096 MPLS labels can be configured for a device. If the label configuration exceeds the maximum limit specified and if the new labels overlap with the already configured labels, you must delete the existing labels and reconfigure the complete set of labels without exceeding the maximum limit. This is not applicable for GigaVUE-TA400.
- All the MPLS labels that are used to strip the MPLS header must be configured at the chassis-level so that all the network or hybrid ports enabled with the MPLS Header Stripping functionality is aware of the labels. ForGigaVUE-TA400, the labels do not need chassis-level configurations to perform MPLS Header Stripping.
- MPLS Header Stripping is supported only for L3 MPLS and L3 MPLS VPN traffic.
- On a network or hybrid port that taps the network traffic, you can enable either VXLAN Header Stripping or MPLS Header Stripping, but you cannot enable both the functionalities.
- If you want to switch between the Header Stripping protocol configuration for a network or hybrid port, you must disable the existing configuration on the port, and then enable the required Header Stripping protocol.
- Ports enabled with VXLAN Header Stripping functionality and ports enabled with MPLS Header Stripping functionality cannot be part of the same map. You must create separate maps for these ports.
- MPLS Header Stripping is not supported on network or hybrid ports that are part of pass-all maps except for L2 MPLS Header Stripping on GigaVUE-TA400.
- In legacy cluster when MPLS Header Stripping enabled on GigaVUE-TA25/GigaVUE-TA25E source/network ports to remote tool ports, packets are discarded at stack ports.
- MPLS Header Stripping is not supported for Q-in-Q traffic except on GigaVUE-TA400.
- In an E-tag cluster MPLS Header Stripping enabled on any of GigaVUE-HC3, GigaVUE-TA100, GigaVUE-TA200 source/network ports, incoming MPLS Q-in-Q traffic will get discarded at local tool/stack port.
- You cannot enable MPLS Header Stripping on a port that is part of a port-pair.
- If you configure a Regular By-Rule map with TTL as one of the map rule, the TTL value will match the MPLS packets' TTL value and not the IP packets' TTL value.

- On GigaVUE-TA400, and GigaVUE-TA400E, both L3 MPLS outer Header Stripping and L2 MPLS outer Header Stripping of traffic carrying with and without Pseudowire MPLS control word (PWMCW) or Pseudowire associated channel header (PWACH) are supported.

Configure MPLS Header Stripping

Before you configure MPLS Header Stripping, refer to *MPLS Header Stripping – Rules, Notes, and Limitations*.

The following table summarizes the required tasks to configure MPLS Header stripping to strip the incoming MPLS packets:

S.No	Task	Refer to..
1.	Configure the required MPLS labels at the chassis level.	Configure MPLS Labels on a Chassis
2.	Configure the required network or hybrid port.	Configure Ports
3.	Enable MPLS Header Stripping on the network or hybrid port that you configured in task 2.	Enable Header Stripping Protocol on Ports
4.	Configure a map. Keep in mind the following details when you configure a map: • Ensure that you select the network or hybrid port on which you enabled MPLS Header Stripping as the source port of the map. • Configure the map rules with the following parameters: o L2 parameters such as EtherType, MAC Destination, MAC Source, VLAN, or Inner VLAN. o MPLS payload parameters such as IP Version, IPv4 Destination, IPv4 Source, IPv6 Destination, IPv6 Source, Protocol, UDA1, or UDA2. For more details, refer to *MPLS Header Stripping – Rules, Notes, and Limitations*.	Create a New Map

Create a New Map

The following are the steps for creating a map:

1. Check the status of the nodes and ports that you plan to use with the map.

   For information about how to check the status of the nodes and ports, refer to Status of Line Cards/Nodes and Ports.
2. From the device view, go to **Maps > Maps** to open the Maps page.
3. Click **New**.

4. Enter the Map Information:

   a. Enter an alias for the map.

      Use an alias that helps identify the task and destination. For example, netflix_traffic_to_wireshark.

   b. (Optional) Enter a description of the map. When adding a description, consider the following:

- Use up to 128 characters, including spaces.
- Enclose the description in quotation marks if the description is longer than one word.
- To include double quotation marks (") inside the quotation marks, precede it with a backslash (\).

      See also Configure MPLS Header Stripping.

   c. Select the **Type**.

      The map type can be Regular, First Level, Second Level, or Inline.

      For detailed information about the types of maps, refer to Map Types

   d. Select the map **Subtype**.

      The map subtype can be **By Rule**, **Pass All**, or **Collector**.

      For detailed information about Pass All, refer to About Map-passall Maps. For detailed information about Collector, refer to About Shared Collectors.

   e. Enable the Pass Traffic checkbox if no rules match.

   f. Enable the Control Traffic checkbox to pass the GTP control traffic (GTP-c) to all GigaSMART engines in a GTP engine group. A GTP engine group has multiple GigaSMART engine port members.

   NOTE: The Control Traffic checkbox is applicable only for GTP and is displayed only if the map type is configured as First level, and the map sub type is configured as By Rule.

5. Specify the **Map Source and Destination**.

   a. From the **Source** and **Destination** drop-down list, select the required source and destination ports for the map. To create a port list, click **Port Editor**.

   NOTE: You can add a maximum of 324 ports in the **Source** drop-down list, if the ports are not attached to a GigaStream.

   NOTE: For details about port types that are supported for the different types of maps, refer to Port Lists.

b. If you have selected a circuit port in the **Destination** drop-down list, select the required circuit tunnel from the **Encapsulation Tunnel** drop-down list to encapsulate the traffic.

> **NOTE:** For details about circuit tunnels, refer to About Circuit-ID Tunnels.

c. If the map is used to redirect the decapsulated traffic to the required tool ports, ensure that you select the IP interface in the **Source** drop-down list. You must have attached the IP interface to the VXLAN or L2GRE tunnel. For details about VXLAN or L2GRE tunnels, refer to About Virtual Extensible LAN (VXLAN) Tunnels and About Layer 2 Generic Routing Encapsulation (L2GRE) Tunnels. From the **Destination** drop-down list, select the required tool ports.

d. If the map will use a GigaSMART operation, select the Header Stripping (MPLS) operation that you have created earlier from the **GigaSMART Operations (GSOP)** drop-down list.

e. If you select the tool GigaStream for the destination, you can view the utilization value for the GigaStream.

6. Add rules to the map.

   To add rules to the map, do any of the following:
   - Use the **Quick Editor**. For details, refer to the Configure MPLS Header Stripping.
   - Import a map template by clicking **Import**.
   - Create a rule by clicking **Add a Rule**.

   For detailed information about map rules, refer to Map Rules.

7. Set the **Map Order** by selecting the priority from the Priority list.

   For details about map priority, refer to Map Priority.

8. Set the **Map Permissions**.

   For details about map permissions, refer to Port Access and Map Sharing.

9. Set the **Map Tag**.

   Select the required tag key and tag value to which the map must be associated. The tag key and the associated tag values must be created in advance in GigaVUE- FM. Refer to the "Tags" and "Role Based Access Control" sections in the GigaVUE Administration Guide for more details.

   > **NOTE:** When you associate a map to a tag value, then the ports, port groups, port pairs, GigaStreams that belong to the map are also associated to the tags.

Configure MPLS Labels on a Chassis

You must configure MPLS labels based on which the device strips the MPLS header from the incoming traffic. This configuration is not applicable for GigaVUE-TA400.

To configure MPLS labels on a chassis:

1. From the left navigation pane, go to **System > Chassis**. Toggle to view the list view.
2. Select the Box ID on which you want to configure the Header Stripping protocol.
3. From the **Actions** drop-down list, select **Configure Header Stripping**. The Configure Header Stripping page appears.
4. From the **Header Stripping Protocol** drop-down list, select **MPLS-L3**.
5. Select the **Apply to All Box** check box to configure the MPLS labels on all the chassis.
6. In the **Labels** field, enter the MPLS labels that you want to configure on the required box. You can also choose to add a range of labels.
7. Click **Save**.



Enable Header Stripping Protocol on Ports

You can enable Header Stripping protocol on any network or hybrid ports. If you enable the Header Stripping protocol on a port, you cannot configure the VXLAN ID or L2GRE ID for that port.

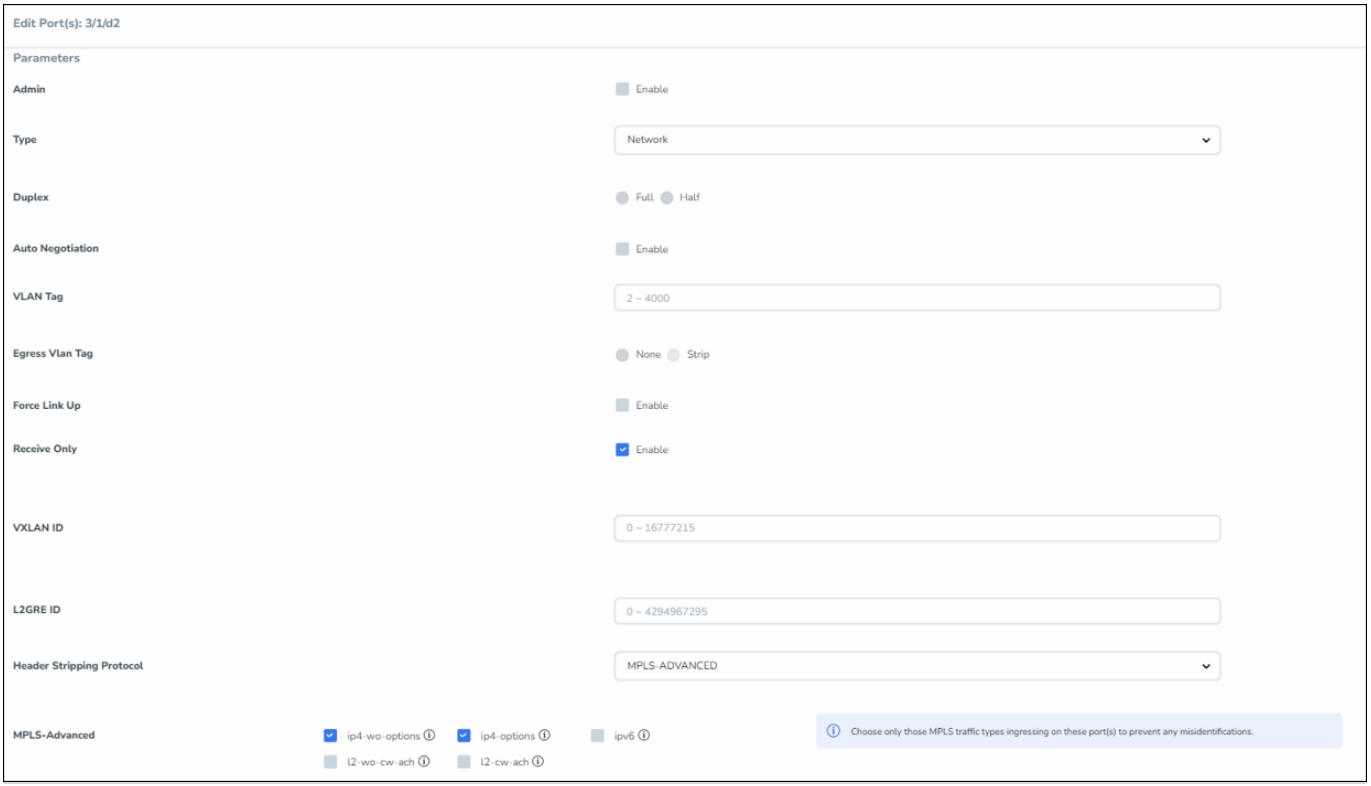
To enable a Header Stripping protocol on a network or hybrid port:

1. From the device view, go to **Ports > Ports > All Ports**.
2. Select the required network or hybrid port on which you want to enable the Header Stripping protocol, and then click **Edit**.

3. Under the Parameters section, from the **Header Stripping Protocol** drop-down list, select one of the following options:

   - **None** – Header stripping protocol is not enabled on the port.
   - **MPLS-L3** – Port is enabled with MPLS Header Stripping protocol to strip the MPLS header from the MPLS layer 3 packets tapped on this port. Select this option if L3 MPLS traffic type only is present in network..
   - **VXLAN** – Port is enabled with VXLAN Header Stripping protocol to strip the VXLAN header from the VXLAN packets tapped on this port.
   - **MPLS-Advanced** – Enable advanced MPLS Header Stripping on the port. This functionality applies only to the GigaVUE-TA400 devices. Refer to Advanced MPLS Header Stripping for GigaVUE-TA400. This option supports the stripping of various MPLS packet formats and provides additional sub-options:
     - **IPv4 Without Options** – Strips the MPLS Header from the L3 MPLS IPv4 (without options) packets. Select this option only if L3 MPLS IPv4 (without options) traffic type is present in network. IPv4 without options has a standard 20-byte header for efficient routing.
     - **IPv4 With Options** – Strips the MPLS Header from the L3 MPLS IPv4 (with options) packets. Select this option only if L3 MPLS IPv4 (with options) traffic type is present in network. IPv4 with options includes additional fields (up to 60 bytes) for advanced functions like timestamping.
     - **IPv6** – Strips the MPLS header from the L3 MPLS IPv6 packets. Select this option only if L3 MPLS IPv6 traffic type is present in network. The standard IPv6 header (40 bytes) follows the MPLS label stack and contains source and destination IPv6 addresses, flow labels, and other fields.
     - **L2 Without CW/ACH** – Strips the outer Ethernet and MPLS headers from the incoming L2 MPLS without Control Word(CW) or Associated Channel Header(ACH) packets. Select this option only if L2 MPLS (without CW/ACH) traffic type is present in network. L2 MPLS without CW or ACH provides a basic MPLS label-switched path for transporting Layer 2 packets without the extra control, sequencing, or signaling functionalities provided by the CW and ACH headers.
     - **L2 With CW/ACH** – Strips the outer Ethernet and MPLS headers along with Control Word (CW) or Associated Channel Header(ACH) from the incoming L2 MPLS with CW or ACH packets. Select this option only if L2 MPLS (with CW/ACH) traffic type is present in network. L2 MPLS with CW ensures packet sequencing and integrity for Layer 2 frames, while ACH enables Operations, Administration, and Maintenance (OAM) transport for network monitoring and management, especially in MPLS traffic.

4. Click **OK** to save the configuration.



Tunnels

Tunneling is a communication protocol that is used to transmit data from one network to another by encapsulating the data. A tunnel is a virtual interface. You can create tunnels for both encapsulation and decapsulation.

Required License: Advanced Feature License on GigaVUE TA Series Nodes

This chapter describes about the different types of native tunnels, which are independent of GigaSMART operations. It also describes how to configure these tunnels for encapsulating and decapsulating traffic. Refer to the following sections for details:

- About Circuit-ID Tunnels
- About Layer 2 Generic Routing Encapsulation (L2GRE) Tunnels
- About Virtual Extensible LAN (VXLAN) Tunnels
- Create Tunnel
- Create VXLAN / L2GRE Group
- Configure L2GRE / VXLAN Identifier
- View VXLAN / L2GRE ID Statistics

About Circuit-ID Tunnels

Circuit-ID tunnels are used to route traffic between two clusters. The traffic is tapped and sent through network ports that are configured on the cluster in the encapsulation side. Based on the flow map configuration, traffic is filtered and then sent through the circuit ports that are configured as the destination port in the map. These circuit ports encapsulate the traffic with a circuit-ID and transmit the encapsulated traffic through the circuit tunnel that connects two clusters. At the receiving end of another cluster in the decapsulation side, the circuit port that is configured as the source port decapsulates the traffic and sends the traffic to the appropriate tools through the tool ports.

The following figure illustrates the circuit Flow Mapping® between two clusters using circuit-ID.

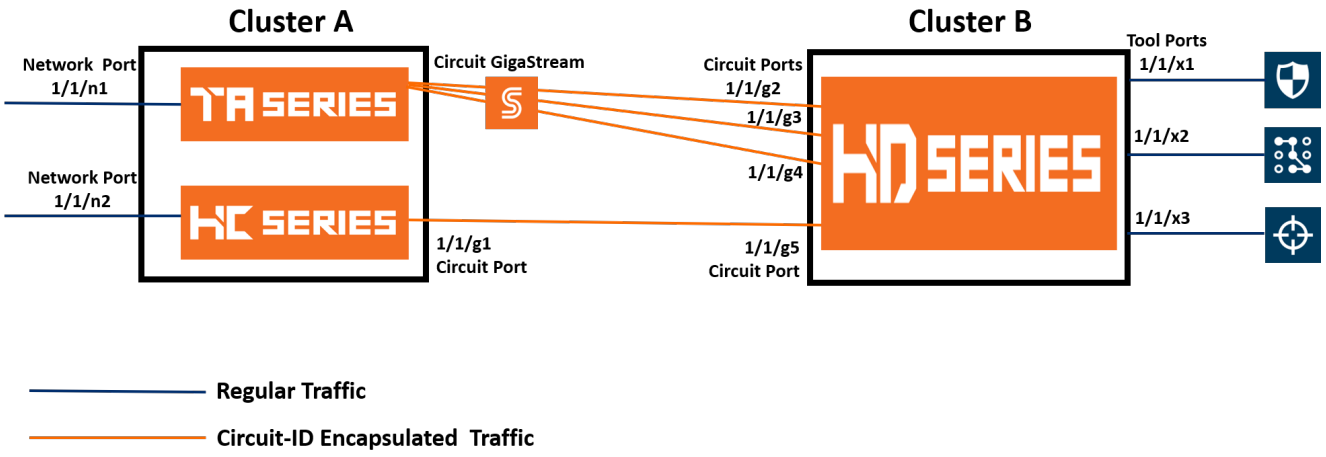


Figure 27 *Circuit Flow Mapping®*

In this example, the GigaVUE TA Series and GigaVUE HC Series nodes reside in cluster A. The tapped traffic is sent through network ports, 1/1/n1 and 1/1/n2. Based on the rules configured in the map, the traffic is filtered at the nodes in cluster A. The filtered traffic is then sent through the circuit port, 1/1/g1 and circuit GigaStream that are configured as the destination port in the map. These circuit ports encapsulate the traffic with circuit-ID and transmit the encapsulated traffic through the circuit tunnel that connects cluster A and cluster B. At the receiving end of cluster B, the circuit ports, 1/1/g2, 1/1/g3, 1/1/g4, and 1/1/g5 that are configured as the source ports decapsulate the traffic, strip the circuit-ID, and send the traffic to the appropriate tools through the tool ports, 1/1/x1, 1/1/x2, and 1/1/x3.

Refer to the following sections for details about the Circuit-ID tunnel encapsulation and decapsulation:

- Circuit-ID Tunnels—Rules and Notes
- Circuit-ID Tunnel Encapsulation

- Circuit-ID Tunnel Decapsulation

Circuit-ID Tunnels—Rules and Notes

Keep in mind the following rules and notes when working with Circuit-ID tunnel encapsulation and decapsulation:

- A maximum of 1024 fabric maps can be created between any two given clusters or nodes in non-shared mode.

NOTE: Prior to GigaVUE-OS release 6.12, a maximum of 512 unique circuit IDs are supported within each cluster or standalone device for encapsulation and decapsulation.

- If a network port receives a double-tagged packet that is encapsulated with a circuit-ID, the five tuple hashing will not work only in the second cluster, that is the cluster in the decapsulation side over stack GigaStream or tool gigastream. Hence, traffic cannot be filtered using the IP/L4 parameters. After decapsulation, Flow Mapping® filters the traffic based on circuit-ID.
- It is not supported for inline scenarios.

Keep in mind the following rules and notes when working with Circuit-ID tunnel encapsulation:

- Circuit-ID tunnel encapsulation is not supported on Pass All maps.
- Port filter configured on circuit port for VLAN pass/drop will try to match the encap circuit-id instead of packet outer VLAN.

Keep in mind the following rules and notes when working with Circuit-ID tunnel decapsulation:

- Circuit-ID tunnel decapsulation is not supported on Pass All and Shared Collector maps.
- A circuit-ID must be paired with a circuit port, only in one circuit-ID tunnel.

Keep in mind the following rules and notes when working with Fabric map circuit ID allocation:

- Circuit ID is used internally to pass traffic from one cluster to another. You can configure Circuit ID ranges from 2 to 4000, or set your own custom range, with each topology able to reuse ID ranges (e.g., 2-513) across multiple topologies.
- Circuit ID allocation is managed globally, beginning from the lower limit of the defined range. This ensures that IDs are allocated efficiently across different topologies. It is allocated automatically.

Circuit-ID Tunnel Encapsulation

Before creating a Circuit-ID tunnel for encapsulation, refer to the Circuit-ID Tunnels—Rules and Notes.

The following table summarizes the required tasks to configure a circuit-ID tunnel for encapsulating the traffic:

S.No	Task	Refer to...
1.	Configure the required circuit ports and circuit GigaStream.	• Configure Ports • Configure Regular GigaStream
2.	Configure a circuit-ID tunnel for encapsulation. Ensure that you select the mode as **Encap**.	Create Tunnel
3.	Configure a map to encapsulate the traffic and attach the circuit-ID tunnel to the map.	Create a new map

Circuit-ID Tunnel Decapsulation

Before creating a Circuit-ID tunnel for decapsulation, refer to the Circuit-ID Tunnels—Rules and Notes.

The following table summarizes the required tasks to configure a circuit-ID tunnel for decapsulating the traffic:

S.No	Task	Refer to...
1.	Configure the required circuit ports and circuit GigaStream.	Configure Ports
2.	Configure a circuit-ID tunnel for decapsulation. Ensure that you select the mode as **Decap** and attach the circuit ports or circuit GigaStream that you configured in step 1 to the circuit-ID tunnel.	Create Tunnel
3.	Configure a map to decapsulate the traffic and ensure that you specify the circuit-ID as a pass/drop rule in the map.	Create a new map

About Layer 2 Generic Routing Encapsulation (L2GRE) Tunnels

L2GRE tunnels are used to route traffic from any remote device to a GigaVUE HC Series or GigaVUE-TA Series device over the internet. The device at the remote site encapsulates the filtered packets, adds a L2GRE encapsulation header, and forwards it to the corresponding circuit port that is used for GRE encapsulation. The encapsulation header consists of Ethernet + IP + GRE headers. The parameters of the encapsulated header are user-configurable, such as the IPv4 address of the IP interface on the destination GigaVUE device and the GRE key that identifies the source of the tunnel.

The encapsulated packet is sent out of the circuit port, which is connected to the public network (the Internet). This packet is routed in the public network to reach the main office site. The packet is ingressed at the circuit port of the GigaVUE device at the main office. The destination IP address of the received packet is checked against the IP configured for the circuit port. If they match, decapsulation is applied. The Ethernet + IP + GRE header is stripped and the remaining packet is sent to the tool port.

If the destination IP address of the received packet does not match with the IP address configured for the IP interface, the packet is dropped.

The following figure illustrates the L2GRE tunnel encapsulation and decapsulation.

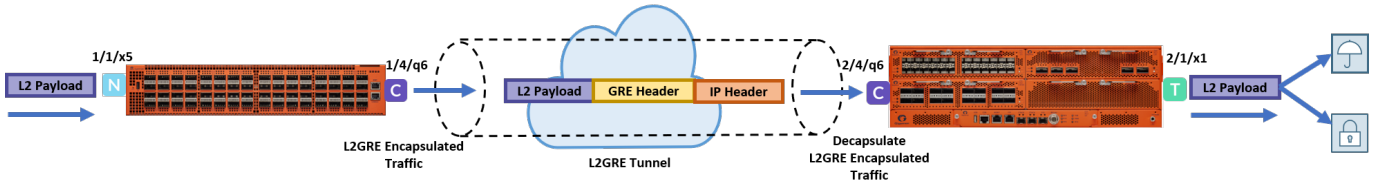


Figure 28 *L2GRE Tunnel Encapsulation and decapsulation*

In this diagram, traffic is tapped on a GigaVUE-TA200 device at a remote site, and then it is tunneled through L2GRE encapsulation across the network before it reaches the GigaVUE-HC3 device at the main office site, which is connected to the actual tools. The tunnel decapsulation is executed on an ingress circuit port (IP interface). After tunnel decapsulation the packet is presented to the Flow Mapping® module to filter based on map rule parameters.

Refer to the following sections for details about the L2GRE tunnel configuration:

- L2GRE Tunnel Configuration—Rules and Notes
- Limitation
- Configure L2GRE Tunnel to Encapsulate Traffic
- Configure L2GRE Tunnel to Decapsulate Traffic

L2GRE Tunnel Configuration—Rules and Notes

Keep in mind the following rules and notes when working with L2GRE tunnels:

- L2GRE tunnels are supported only on GigaVUE-HC1-Plus, GigaVUE-HCT, GigaVUE-HC1, GigaVUE-HC3CCv1 and CCv2, GigaVUE-TA100, GigaVUE-TA200, GigaVUE-TA200E GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA400, GigaVUE-TA400E, and DELL S4112F-ON devices.
- A maximum of 1500 L2GRE IDs are supported.
- IPv6 L2GRE tunnels are supported only when using the GigaSMART engine and require an Advanced Tunneling license.
- Ingress VLAN tagging and Tool Mirror features are not supported with L2GRE tunnels.
- Filtering of Q-in-Q packets is not supported with L2GRE tunnels except on GigaVUE-TA400.
- Map-passall is not supported for the circuit port that encapsulates or decapsulates the L2GRE packet.

- When configuring a map for L2GRE encapsulation, you cannot configure a combination of a regular tool port and L2GRE encapsulation tunnel as part of the "To" ports.

- Any encapsulated packet that exceeds the MTU value configured for the IP interface will be discarded because IP fragmentation and reassembly of packets are not supported.
- L2GRE tunnel encapsulation is not supported on circuit GigaStreams.
- Flow mapping that is configured on the circuit port used for L2GRE decapsulation will filter only the inner packet attributes along with L2GRE-ID. Any other non-tunneled packets that ingress on this circuit port will not be filtered or redirected to tool ports, even if it matches the rules configured on the map.
- GigaSMART operations cannot be combined with L2GRE decapsulation in the same map.
- L2GRE tunnel decapsulation is supported only on encapsulated packets that are not tagged. On GigaVUE-TA400, GigaVUE-TA400E L2GRE tunnel decapsulation is supported on encapsulated packets that are both tagged and untagged.
- Inner VLAN qualifier is not supported on the port in which the L2GRE tunnel decapsulation is enabled except on GigaVUE-TA400.
- L2GRE ID qualifier is available as part of existing static templates. Following table provides details about the platforms for which the static templates are available:

Table 11: L2GRE ID Qualifier Availability

Template	Platform	
	GigaVUE-HC1	GigaVUE-HC3/ GigaVUE-TA100/ GigaVUE-TA25/ GigaVUE-TA25E/ GigaVUE-TA200/ GigaVUE-TA200E/ GigaVUE-TA400/ GigaVUE-TA400E
IPv4	Not available	Available
IPv6	Available	Available
IPv4+UDA	Not available	Available
IPv4+MAC	Available	Available
UDA	Available	Available

Limitation

When the encapsulation device fragments your traffic, the L2GRE Tunnels used to decapsulate the traffic does not support re-assembly. To avoid this, you can use GigaSMART L2GRE decapsulation, which reassembles the fragmented packets. Refer to GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation for more detailed information on how to configure GigaSMART L2GRE Tunnel Decapsulation. You can also configure the highest possible MTU value before tapping the traffic to the virtual machine so that packets are not fragmented.

Configure L2GRE Tunnel to Encapsulate Traffic

Before creating a L2GRE tunnel to encapsulate traffic, refer to the L2GRE Tunnel Configuration—Rules and Notes.

The following table summarizes the required tasks to configure a L2GRE tunnel for encapsulating the traffic:

S.No	Task	Refer to...
1.	Configure the required circuit port.	Configure Ports
2.	Configure an IP interface and attach the circuit port that you created in task 1.	Configure IP Interface
3.	Configure a L2GRE tunnel for encapsulation and attach the IP interface that you created in task 2. Ensure that you select the Type as **L2GRE** and the Mode as **Encapsulation**.	Create Tunnel
4	Create a L2GRE group for a device and add all the L2GRE IDs that are specific to the device.	Create VXLAN / L2GRE Group
5.	Configure the L2GRE ID either at the chassis-level or at the port-level to help identify the encapsulation tunnel. You must use one of the L2GRE IDs that you have already added in the L2GRE group for the device in task 4.	Configure L2GRE / VXLAN Identifier
6.	Configure a map to encapsulate the L2GRE traffic. Ensure that you add the following details when configuring a map: • If you have configured the L2GRE ID for the chassis, you can configure any network port on the chassis as the source port of the map. • If you have configured the L2GRE ID for a network port, you must configure that network port as the source port of the map. • Configure the encapsulation tunnel that you created in task 3 as the destination port of the map. • Specify the required pass/drop rule in the map to filter the traffic based on inner packet attributes or L2GRE ID for the template configured.	Create a new map

Configure L2GRE Tunnel to Decapsulate Traffic

Before creating a L2GRE tunnel to decapsulate traffic, refer to the L2GRE Tunnel Configuration—Rules and Notes.

The following table summarizes the required tasks to configure a L2GRE tunnel for decapsulating the traffic:

S.No	Task	Refer to...
1.	Configure the required circuit port.	Configure Ports
2.	Configure an IP interface and attach the circuit port that you created in task 1.	Configure IP Interface
3.	Configure a L2GRE tunnel for decapsulation and attach the IP interface that you created in task 2. Ensure that you select the Type as **L2GRE** and Mode as **Decapsulation**.	Create Tunnel
4	Create a L2GRE group for a device and add all the L2GRE IDs that are specific to the device.	Create VXLAN / L2GRE Group
5.	Configure a map to decapsulate the traffic. Ensure that you add the IP interface that you created in task 2 as the source of the map and specify the required pass/drop rule in the map to filter the traffic based on inner packet attributes or L2GRE ID for the template configured.	Create a new map

Orchestrated Workflow Configuration of L2GRE Tunnels

The L2GRE tunnels can now be configured through the Orchestrated Configuration page. To configure L2GRE Tunnels, follow the below steps:

- In GigaVUE-FM go to **Traffic>Physical> Orchestrated Flows> Tunnels**.
- Click on **New**.
- Select **Embedded L2GRE**. The New Tunnel configuration page appears. Enter the required information as described below:

Field	Description
Tunnel Name	The name of the tunnel. **Note**: Alias must not have spaces or any of these characters: *"?;,/%@.
Tunnel Description	The description of the tunnel endpoint.
Traffic Direction	Select the traffic direction of the tunnel. Choose DECAP for decapsulation or choose ENCAP for encapsulation.

Field	Description
	• If you choose Decap, select the IP Interface from the list. You can create one using the Create an IP interface option if you do not have an IP interface. Refer to IP Interfaces to learn more about creating an IP Interface. • If you choose Encap enter the Source Port. You can select from the drop-down list. To edit the Port type or to enable the port for Admin privileges use the Port Editor window. o You can also add a new L2GRE ID by clicking on the **ADD L2GRE ID**. **Note**: The L2GRE ID configured at source port value will take precedence over the L2GRE ID that is configured at chassis level.
Nodes	Select the Nodes on which you intend to create the tunnel. For L2GRE Encapsulation tunnel you will have to configure the L2GRE ID. • Click on **Config L2GRE ID> ADD L2GRE ID.** • Fill in the require parameters. • Click on **ADD**. • The L2GRE IDs will list in the Configuration page.
ID	For L2GRE Decapsulation tunnel you will have to select the L2GRE ID. To configure a new L2GRE ID do the following: • Click on **ADD L2GRE ID.** • Fill in the require parameters. • Click on **ADD**. **Note**: Configuring a L2GRE ID is not mandatory if the map rules are configured. Either configure the rules or a L2GRE ID.
Rules	Configure the map rules that needs to be adhered to while decapsulating or encapsulating the traffic. Click on **Rules Editor** to configure the rules. **Note**: Configuring rules are mandatory for Encapsulation tunnel, For a Decapsulation tunnel rules are not mandatory if a L2GRE ID is already configured.
Encapsulation IP interface	Enter the interface IP address of the node (Destination IP) for an encapsulation tunnel.
Remote Tunnel IP	Enter the Remote Tunnel IP values for and encapsulation tunnel.
Destination Port	Select the destination port for the decapsulation tunnel. You can edit the port details from the Port editor screen.

The configured tunnels provide you a **Details View** and **Troubleshoot View**. Click on the tunnel profile to view the below:

- **Details View**: Displays the configured parameters of the configured tunnel.
- **Troubleshoot View**: Displays the tunnel's statistics. Use the **Clear Stats** button to reset the statistics.

About Virtual Extensible LAN (VXLAN) Tunnels

VXLAN is a simple tunneling mechanism that allows overlaying a Layer 2 (L2) network over a Layer 3 (L3) underlay with the use of any IP routing protocol. It uses MAC Address-in-User Datagram Protocol (MAC-in-UDP) encapsulation. A remote device, such as the Gigamon cloud, GigaVUE TA Series, GigaVUE HC Series, or a customer-specific device, encapsulates the filtered traffic, adds an encapsulation header that consists of Layer 2 + IP + UDP + VXLAN headers. The encapsulated packet is sent out of the circuit port, which is connected to the public network (the Internet). This packet is routed in the public network to reach the main office site. The packet is ingressed at the circuit port configured in the GigaVUE-H Series or GigaVUE-TA Series device at the main office. After validation of the source port, destination port, and VXLAN Network Identifier (VNI) of the packet, the VXLAN tunnel header will be removed and the inner payload will be sent to the tools based on the map rules configured.

The following figure illustrates the VXLAN tunnel encapsulation and decapsulation.

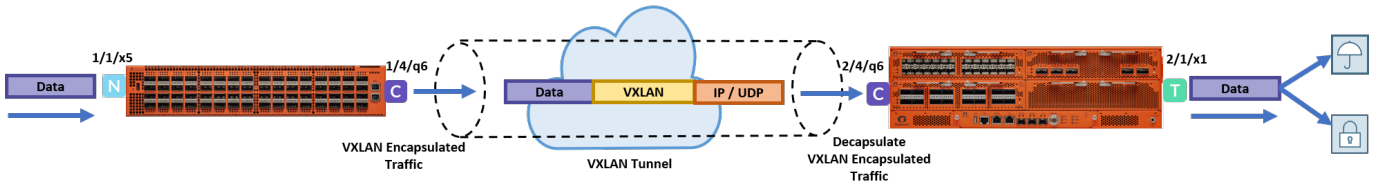


Figure 29 *VXLAN Tunnel Encapsulation and Decapsulation*

In this diagram, traffic is tapped on a GigaVUE-TA200 device at a remote site, and then it is tunneled through VXLAN encapsulation across the network before it reaches the GigaVUE-HC3 device at the main office site, which is connected to the actual tools. The tunnel decapsulation is executed on an ingress circuit port (IP interface). After tunnel decapsulation, the packet is presented to the flow mapping module to filter based on map rule parameters.

Refer to the following sections for details about VXLAN tunnels:

- VXLAN Tunnel Configuration—Rules and Notes
- Configure VXLAN Tunnel to Encapsulate Traffic
- Configure VXLAN Tunnel to Decapsulate Traffic

VXLAN Tunnel Configuration—Rules and Notes

Keep in mind the following rules and notes when working with VXLAN tunnels:

- VXLAN tunnels are supported only on GigaVUE-HC1, GigaVUE-HC3, GigaVUE-HCT, GigaVUE-HC1-Plus, GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA200, GigaVUE-TA200E,GigaVUE-TA400, and GigaVUE-TA400E devices.
- A maximum of 1500 VXLAN IDs are supported.
- Flow mapping that is configured on the circuit port used for VXLAN decapsulation will filter only the inner packet attributes along with VXLAN-ID. Any other non-tunneled packets that ingress on this circuit port will not be filtered or redirected to tool ports, even if it matches the rules configured on the map.
- IPv6 protocol is not supported with VXLAN tunnels. However, IPv6 protocol is supported with GigaSMART VXLAN tunnels.
- Ingress VLAN tagging and Tool Mirror features are not supported with VXLAN tunnels.
- Any encapsulated packet that exceeds the MTU value configured for the IP interface will be discarded because IP fragmentation and reassembly of packets are not supported.
- VXLAN tunnel encapsulation is not supported on circuit GigaStreams.
- VXLAN tunnel decapsulation is supported only on encapsulated packets that are not tagged. OnGigaVUE-TA400, and GigaVUE-TA400E VXLAN tunnel decapsulation is supported on encapsulated packets that are both tagged and untagged.
- GigaSMART operations cannot be combined with VXLAN decapsulation in the same map.
- Map-passall is not supported for the circuit port that encapsulates or decapsulates the VXLAN packet.
- When a circuit port is configured for VXLAN tunnel decapsulation, you cannot use the port in any other regular map in which a network port is configured as the source port.
- Inner VLAN qualifier is not supported on the port in which the VXLAN tunnel decapsulation is enabled except on GigaVUE-TA400, and GigaVUE-TA400E.
- VXLAN ID qualifier is available as part of existing static templates. Following table provides details about the platforms for which the static templates are available:

Table 12: VXLAN ID Qualifier Availability

Template	Platform	
	GigaVUE-HC1	GigaVUE-HC3/GigaVUE-TA100/GigaVUE-TA200/ GigaVUE-TA200E/ GigaVUE-TA25/GigaVUE-TA25E/ GigaVUE-TA400/GigaVUE-TA400E
IPv4	Not available	Available
IPv6	Available	Available

Template	Platform	
	GigaVUE-HC1	GigaVUE-HC3/GigaVUE-TA100/GigaVUE-TA200/ GigaVUE-TA200E/ GigaVUE-TA25/GigaVUE-TA25E/ GigaVUE-TA400/GigaVUE-TA400E
IPv4+UDA	Not available	Available
IPv4+MAC	Available	Available
UDA	Available	Available

Configure VXLAN Tunnel to Encapsulate Traffic

Before creating a VXLAN tunnel to encapsulate traffic, refer to the VXLAN Tunnel Configuration—Rules and Notes.

The following table summarizes the required tasks to configure a VXLAN tunnel for encapsulating the traffic:

S.No	Task	Refer to...
1.	Configure the required circuit port.	Configure Ports
2.	Configure an IP interface and attach the circuit port that you created in task 1.	Configure IP Interface
3.	Configure a VXLAN tunnel for encapsulation and attach the IP interface that you created in task 2. Ensure that you select the Type as **VXLAN** and the Mode as **Encapsulation**.	Create Tunnel

S.No	Task	Refer to...
4	Create a VXLAN group for a device and add all the VXLAN IDs that are specific to the device.	Create VXLAN / L2GRE Group
5.	Configure the VXLAN ID either at the chassis-level or at the port-level to help identify the encapsulation tunnel. You must use one of the VXLAN IDs that you have already added in the VXLAN group for the device in task 4.	Configure L2GRE / VXLAN Identifier
6.	Configure a map to encapsulate the VXLAN traffic. Ensure that you add the following details when configuring a map: • If you have configured the VXLAN ID for the chassis, you can configure any network port on the chassis as the source port of the map. • If you have configured the VXLAN ID for a network port, you must configure that network port as the source port of the map. • Configure the encapsulation tunnel that you created in task 3 as the destination port of the map. • Specify the required pass/drop rule in the map to filter the traffic based on inner packet attributes or VXLAN ID for the template configured.	Create a new map

Configure VXLAN Tunnel to Decapsulate Traffic

Before creating a VXLAN tunnel termination, refer to the .

The following table summarizes the required tasks to configure a VXLAN tunnel for decapsulating the traffic:

S.No	Task	Refer to...
1.	Configure the required circuit port.	Configure Ports
2.	Configure an IP interface and attach the circuit port that you created in task 1.	Configure IP Interface
3.	Configure a VXLAN tunnel for decapsulation and attach the IP interface that you created in task 2. Ensure that you select the Type as **VXLAN** and Mode as **Decapsulation**.	Create Tunnel

S.No	Task	Refer to...
4	Create a VXLAN group for a device and add all the VXLAN IDs that are specific to the device.	Create VXLAN / L2GRE Group
5.	Configure a map to decapsulate the traffic. Ensure that you add the IP interface that you created in task 2 as the source of the map and specify the required pass/drop rule in the map to filter the traffic based on inner packet attributes or VXLAN ID for the template configured.	Create a new map

Orchestrated Workflow Configuration of VXLAN Tunnels

The VXLAN tunnels can now be configured through the Orchestrated Configuration page. To configure VXLAN Tunnels follow the below steps:

- In GigaVUE-FM go to **Traffic>Physical> Orchestrated Flows> Tunnels**.
- Click on **New**.
- Select **Embedded VXLAN**. The New Tunnel configuration page appears. Enter the required information as described below:

Field	Description
Tunnel Name	The name of the tunnel. **NOTE:** Alias must not have spaces or any of these characters: *'"?;:,/%@.
Tunnel Description	The description of the tunnel endpoint.
Traffic Direction	Select the traffic direction of the tunnel. Choose DECAP for decapsulation or choose ENCAP for encapsulation. • If you choose Decap, select the IP Interface from the list. You can create one using the Create an IP interface option if you do not have an IP interface. Refer to IP Interfaces to learn more about creating an IP Interface. • If you choose Encap enter the Source Port. You can select from the drop-down list. To edit the Port type or to enable the port for Admin privileges use the Port Editor window. o You can also add a new VXLAN ID by clicking on the **ADD VXLAN ID**. **NOTE:** The VXLAN ID configured at source port will take precedence over the VXLAN ID configured at chassis level.
Nodes	Select the Nodes on which you intend to create the tunnel. For VXLAN Encapsulation tunnel you will have to configure the VXLAN ID.

Field	Description
	• Click on **Config VXLAN ID > ADD VXLAN ID.** • Fill in the require parameters. • Click on **ADD**. • The VXLAN IDs will list in the Configuration page.
ID	For VXLAN Decapsulation tunnel you will have to select the VXLAN ID. To configure a new VXLAN ID do the following: • Click on **ADD VXLAN ID.** • Fill in the require parameters. • Click on **ADD**. **NOTE:** Configuring a VXLAN ID is not mandatory if the map rules are configured. Either configure the rules or a VXLAN ID.
Rules	Configure the map rules that needs to be adhered to while decapsulating or encapsulating the traffic. Click on **Rules Editor** to configure the rules. **NOTE:** Configuring rules are mandatory for Encapsulation tunnel, For a Decapsulation tunnel rules are not mandatory if a VXLAN ID is already configured.
Encapsulation IP interface	Enter the interface IP address of the node (Destination IP) for an encapsulation tunnel.
Remote Tunnel IP, L4 Port	Enter the IP address , L4 Port values for an encapsulation tunnel.
Destination Port	Select the destination port for the decapsulation tunnel. You can edit the port details from the Port editor screen.

The configured tunnels provide you a **Details View** and **Troubleshoot View**. Click on the tunnel profile to view the below:

- **Details View**: Displays the configured parameters of the configured tunnel.
- **Troubleshoot View**: Displays the tunnel's statistics. Use the **Clear Stats** button to reset the statistics.

Create Tunnel

You can create tunnels on the encapsulation side and/or decapsulation side. To create a tunnel:

1. From the device view, go to **Ports > Tunnels**.
2. Click **New**. The Tunnel configuration page appears.
3. In the **Alias** and **Description** fields, enter the name and description of the tunnel.
4. From the **Type** drop-down list, select one of the following options:
   - **Circuit-ID—**to create Circuit-ID tunnels.

- **VXLAN—**to create VXLAN tunnels.
- **L2GRE—**to create L2GRE tunnels.

5. Select one of the following modes:
   - **Encapsulation**—Select this mode to send the encapsulated traffic to the destination node that resides in another cluster.
   - **Decapsulation**—Select this mode to decapsulate the traffic received from the source node that resides in another cluster.
6. In the **Circuit-ID** field, enter the circuit ID used to encapsulate or decapsulate the traffic. You can enter multiple circuit-IDs when you create a circuit-ID tunnel for decapsulation.

   NOTE: This field is available only when you select **Circuit-ID** in the **Type** drop-down list.
7. In the **Destination IP Address** field, enter the IPv4 address of the destination node.

   NOTE: This field is available only when you select **VXLAN** or **L2GRE** in the **Type** field and the **Mode** as **Encapsulation**.
8. In the **L4 Source Port** field, enter the layer 4 source port number.

   NOTE: This field is available only when you select **VXLAN** in the **Type** field and the **Mode** as **Encapsulation**.
9. From the **Attached entity** drop-down list, select the required entity based on the tunnel type you are creating:
   - If you are creating a Circuit-ID tunnel, this field is available only when you select the **Mode** as **Decapsulation**. Ensure that you attach the required circuit ports or circuit GigaStreams.
   - If you are creating a VXLAN tunnel or a L2GRE tunnel, ensure that you select the IP interface to which you have attached the circuit port.
10. Click **OK**.

The newly added tunnel appears in the Tunnels listing page.

Create VXLAN / L2GRE Group

You can create a VXLAN or L2GRE group for a GigaVUE HC Series or a GigaVUE TA Series device and add all the VXLAN IDs or the L2GRE IDs that are specific to the device. You can add a maximum of 1500 IDs per system. To create a VXLAN or L2GRE group:

1. From the device view, go to **Ports > Tunnels > VXLAN / L2GRE Groups**.
2. Click **New**. The VXLAN / L2GRE Group page appears.
3. In the **Alias** and **Description** fields, enter the name and description of the group.
4. Select either **VXLAN** or **L2GRE** based on the group that you want to create.
5. From the **Box ID** drop-down list, select the required box ID of the device for which you want to create the group.
6. Click **Add ID**, and then enter the VXLAN ID or L2GRE ID for the device.
7. Repeat step 6 to add multiple VXLAN IDs or L2GRE IDs for the device, and then click **OK**.

NOTE: Ensure that you create a VXLAN or L2GRE group and add the required IDs to the group before you assign the VXLAN or L2GRE IDs to the map you create for tunnel termination.

The ingress packets to the VXLAN tunnels must have a matching L4 destination port configured without which the packets will be discarded. To configure the **L4 Destination Port** :

1. Click **Actions** on the VXLAN / L2GRE Group page and select **Configure VXLAN Global**.
2. Select the required **Box ID**.
3. Enter the **L4 Destination Port** value. Default value is 4789.
4. Click **OK.**

NOTE: The L4 Destination port is configured globally on the chassis and thus affects all the VXLAN tunnels on the box.

Configure L2GRE / VXLAN Identifier

The L2GRE or VXLAN IDs that you have added in the L2GRE or VXLAN group created for a specific device must be configured either at the chassis-level or at the network port-level. The ID that you configure for the network port overrides the ID configured for the chassis. The L2GRE or VXLAN IDs help to identify the respective tunnels.

Configure L2GRE / VXLAN Identifier for a Chassis

To configure L2GRE or VXLAN ID for a chassis:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**..In the Physical Nodes page, select the node for which you want to configure the L2GRE or VXLAN ID.

2. From the left navigation pane, go to **System > Chassis**.
3. Go to List View, and then select the Box ID for which you want to configure the L2GRE or VXLAN ID.
4. From the **Actions** drop-down list, select **Configure L2GRE/VXLAN ID**.
5. Enter either the L2GRE ID or the VXLAN ID in the respective field.

   NOTE: You cannot specify both L2GRE ID and VXLAN ID for a chassis.

6. Click **OK**.

Configure L2GRE / VXLAN Identifier for a Network Port

To configure L2GRE or VXLAN ID for a network port:

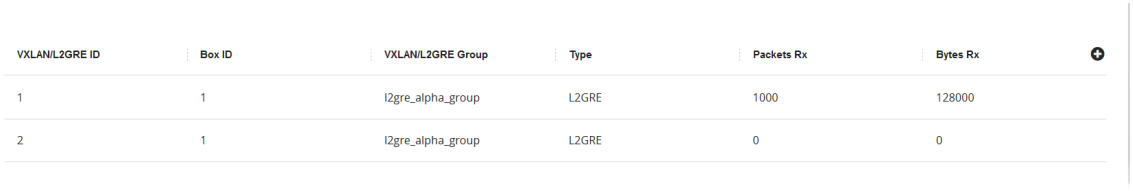
1. From the left navigation pane, go to **Inventory > Physical > Nodes**.. In the Physical Nodes page, select the node for which you want to configure the L2GRE or VXLAN ID.
2. Go to **System > Ports > Ports > All Ports**.
3. Select the network port for which you want to configure the L2GRE or VXLAN ID, and then click **Edit**. The Ports page appears with the details of the port selected.
4. Under the Parameters area of the page, enter either the L2GRE ID or the VXLAN ID in the respective field.

   NOTE: You cannot specify both L2GRE ID and VXLAN ID for a network port.

5. Click **OK**.

View VXLAN / L2GRE ID Statistics

You can view the statistics such as the number of packets and bytes that were decapsulated using a VXLAN ID or a L2GRE ID for a device. To view the VXLAN /L2GRE ID Statistics page, go to **Ports > Tunnels > Statistics**. Figure 30VXLAN / L2GRE ID Statistics illustrates the statistics.



VXLAN/L2GRE ID	Box ID	VXLAN/L2GRE Group	Type	Packets Rx	Bytes Rx
1	1	l2gre_alpha_group	L2GRE	1000	128000
2	1	l2gre_alpha_group	L2GRE	0	0

Figure 30 *VXLAN / L2GRE ID Statistics*

The following table describes the VXLAN or L2GRE ID statistics:

Statistic	Description
VXLAN/L2GRE ID	The VXLAN or L2GRE identifier that is specific to the device.
Box ID	The box identifier of the device.
VXLAN/L2GRE Group	The name of the group created for the device.
Type	The type of group—VXLAN or L2GRE.
Packets Rx	The number of packets that were decapsulated using the VXLAN or L2GRE ID.
Bytes Rx	The total bytes that were decapsulated using the VXLAN or L2GRE ID.

To export the statistics to a CSV file, click the '+' icon and then select **Download all data as csv**.

Note: The bulk CSV download window does not generate the correct CSV file.

To clear the VXLAN statistics, click **Clear Stats > Clear All VXLAN Stats**. To clear the L2GRE statistics, click **Clear Stats > Clear All L2GRE Stats**.

Tunnel Monitoring

A tunnel is a logical entity configured to establish secure connections across networks. Traffic is encapsulated at the sending end of the tunnel and decapsulated at the receiving end of the tunnel. GigaVUE-FM supports configuring various tunnel types. Refer to the Tunnels section for details.

With the Tunnel Monitoring feature, GigaVUE-FM uses a tunnel discovery mechanism to identify the tunnels created across the managed nodes and clusters. The following tunnel types are supported in this release:

- Embedded tunnels.
  - Circuit-id tunnels
  - L2GRE tunnels
  - VXLAN tunnels
- GigaSMART tunnels.
  - GigaSMART L2GRE tunnel
  - GigaSMART VXLAN tunnel
  - GigaSMART TLS_PCAPNG Tunnel
- Mixed Tunnels: Tunnel type on encapsulation end and decapsulation end can be either Embedded tunnel or GigaSMART tunnel.
- Hybrid Tunnels: Egress Tunnels (L2GRE/VxLAN/TLS_PCAPNG) originating from GigaVUE V Series Nodes and terminating at GigaVUE HC Series or GigaVUE TA Series devices.
- Secure Tunnels.

The identified tunnels are logically grouped based on the following grouping criteria and listed in the Tunnel Monitoring page. The Tunnel Monitoring page provides a unified view of both ends of the tunnels created across the devices managed by GigaVUE-FM.

Tunnel Type	Grouping Criteria
Embedded Tunnels	
- Circuit-ID Tunnels	• Circuit-ID • Physical Connectivity discovered by GDP. **Note**: Physical connectivity is an optional parameter for grouping the circuit-id tunnels
- L2GRE/VXLAN Tunnels	• Destination IP • L2GRE-ID/VXLAN-ID

GigaSMART Tunnels	
- **L2GRE/VXLAN Tunnels**	• Destination IP • L2GRE-ID/VXLAN-ID • Application ports (VXLAN Tunnel)
-**TLS_PCAPNG Tunnels**	• Destination IP • Application Ports
Hybrid Tunnels	
- L2GRE/VXLAN/ from GigaVUE V Series Nodes and terminating at GigaVUE HC Series and GigaVUE TA Series devices	• Destination IP • L2GRE-ID/VXLAN-ID
-TLS_PCAPNG Tunnel form GigaVUE V series Nodes	• Destination IP • Application Ports

During addition of nodes to GigaVUE-FM, the tunnels configured in the devices are automatically detected, logically grouped and added to the Tunnel Monitoring page. The tunnel logical group gets deleted during node deletion. GigaVUE-FM identifies the tunnel related configuration changes during the config sync cycle and accordingly updates the tunnel logical groups in the Tunnel Monitoring page.

NOTE: For the tunnels listed in the tunnel monitoring page, you can identify the traffic drop at the decapsulation end of the tunnel by creating Alert Policies. Refer to the Traffic Drop Identification section for details.

Rules, Notes, and Limitations

Refer to the following Rules, notes, and limitations:

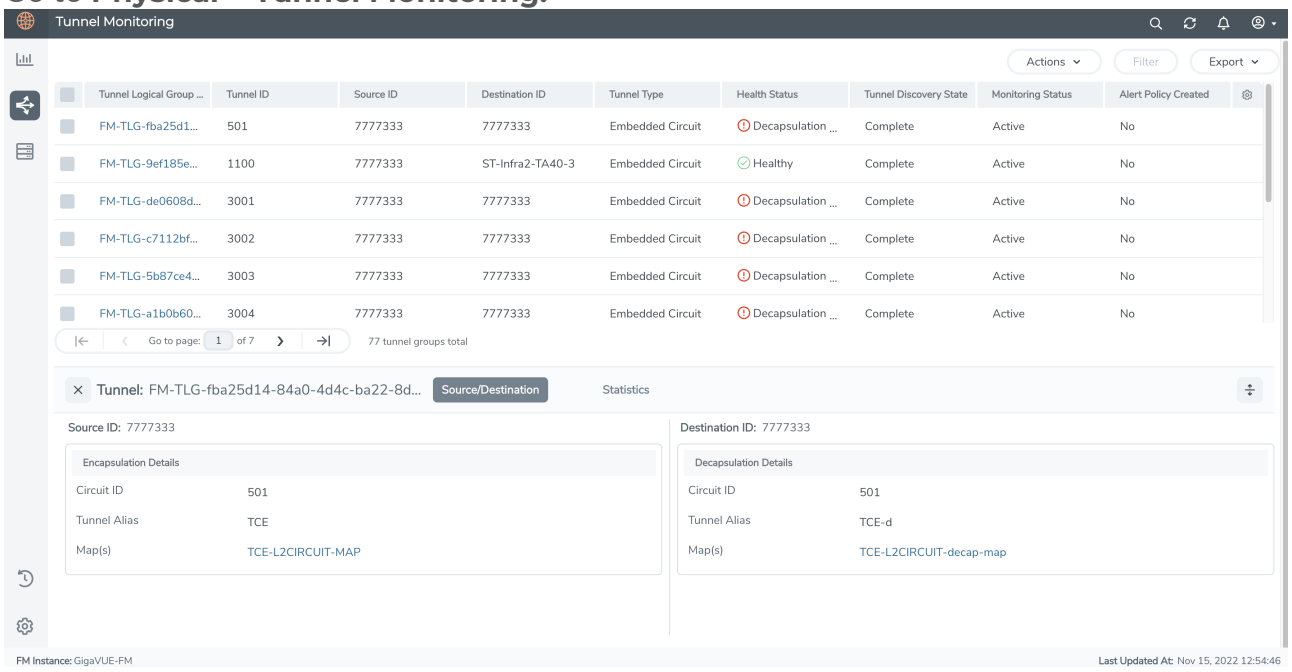
- The tunnel discovery mechanism identifies tunnels and the tunnel applications that were configured at the devices that are being managed by GigaVUE-FM. The tunnels are identified irrespective of the device software version.
- GigaVUE-FM identifies tunnels configured between the physical devices. GigaVUE-FM also identifies tunnels configured between GigaVUE V Series Nodes and physical devices. Tunnels deployed in VMware (ESXi and NSX-T) are only supported.
- The following tunnel types are not supported in this release:
  - Tunnels configured between the GigaVUE V Series Nodes.
  - Tunnels configured between GigaVUE V Series Nodes and UCT-Vs
- GigaVUE-FM detects tunnels that use 'Regular' maps for encapsulation and decapsulation.
- Circuit ID tunnels that are created manually will be discovered and monitored. Circuit ID tunnels created as part of Fabric Map configurations are not within the scope of the Tunnel Monitoring page.

- For embedded tunnels (L2GRE/VXLAN), GigaVUE-FM detects the decapsulation configuration only when pass rule is configured with L2GRE/VXLAN Id on the maps at the decapsulation end.

View Tunnel Monitoring

To access the tunnel monitoring page:

1. Go to **Traffic**
2. Go to **Physical > Tunnel Monitoring.**



The list of tunnels are displayed.

Field	Description
Tunnel Logical Group Name	The Logical group name of the tunnel is automatically generated by GigaVUE-FM. Use the edit option to edit the alias name of the tunnel.
Tunnel ID	Tunnel Identifier
Source ID	Host name or cluster Id (Cluster type) of nodes participating in tunnel encapsulation.
Destination ID	Host name or cluster Id (Cluster type) of nodes participating in tunnel decapsulation.
Tunnel Type	Type of tunnel. Can be: • Embedded (Circuit-Id, VXLAN, L2GRE) • GigaSMART (TLS_PCAPNG,VXLAN, L2GRE, ERSPAN - Decapsulation end) • Mixed • Hybrid
Health Status	Health status of the tunnel group. Displayed only if monitoring status is active.

Tunnel Discovery State	Tunnel discovery state. Can be: • **Complete:** Tunnel has both encapsulating and decapsulating end • **Incomplete**: Tunnel has only one end.
Monitoring Status	Monitoring status of the tunnel, based on which the Health Status is displayed
Alert Policy Created	Displayed as 'Yes' or 'No' based on Alert Policy created.

Click on a Tunnel Logical Group name to view the following details:

- **Source/Destination**: Displays the source and destination details of a tunnel logical group.
- **Statistics**: Displays the following statistical details about the tunnel logical group based on statistics collected by GigaVUE-FM during the last statistical poll cycle:
  - Aggregated packets/octets Tx of encapsulation end.
  - Aggregated packets/octets Rx of decapsulation end.
  - Processed Traffic % displays the amount of packets received at the decapsulation end.

Use the following buttons to manage the displayed tunnels:

Button	Description
Actions	Use the **Actions** drop-down button to activate and deactivate the selected tunnels. To create Alert Policy for tunnels: 1. Select the required Tunnel Logical Group. 2. Click the Actions drop-down. 3. Click Create Alert Policy. 4. Enter the required details. **NOTE:** You can create Alert Policies only if the Tunnel Discovery State is 'Complete'. Refer Traffic Drop Identification to the section for details.
Filter	Use to filter the Tunnels based on the following criteria: • Tunnel ID • Tunnel Type • Tunnel Discovery State • Monitoring Status • Source ID • Destination ID
Export	Use to export the tunnels (either all or the selected tunnels) in CSV or XLSX format.

Packet Capture (PCAP)

Starting from software version 5.16.00, you can configure Packet Capture (PCAP) from GigaVUE-FM. Both GigaVUE-FM and the devices must be running software version 5.16.00 and greater. Use the PCAP feature to analyze the network traffic and to troubleshoot any performance issues.

GigaVUE-FM allows you to configure packet capture at the ingress port or egress port or both. The port must be a physical port. The port type used for packet capture can be network, tool, hybrid, inline tool, or inline network port. Packet capture is not supported on GigaSMART ports or back plane ports.

NOTE: PCAP feature is enabled by default. To disable or re-enable PCAP, contact Gigamon customer support. Once disabled, the corresponding PCAP configurations will not work.

Supported Devices

Packet capture functionality is supported on the following devices:

- GigaVUE-HC1
- GigaVUE-HC1-Plus
- GigaVUE-HCT
- GigaVUE-HC3
- GigaVUE-TA25
- GigaVUE-TA25E
- GigaVUE-TA100
- GigaVUE-TA200
- GigaVUE-TA200E
- GigaVUE-TA400
- GigaVUE-TA400E

You can configure PCAP on both standalone nodes as well as on nodes that belong to a cluster. For non-leader ports, PCAP can be configured only from the leader node.

To configure packet capture, you must define filters to capture specific traffic based on rules. You can specify the following criteria in the rules:

Criteria	Description
Source MAC address	The source and destination MAC address.
Destination MAC address	
VLAN ID	VLAN ID value.
Inner-VLAN	Inner VLAN ID value.
Layer 2 ethernet type	Layer 2 Ethernet type value.
Source IPv4 address	The source and destination IPv4 address. You can also specify a wild card with an IP mask.
Destination IPv4 address	
Internet protocol	Valid Internet protocol.
IP version number	IP version for traffic, either IPv4 or IPv6.
IP fragmentation bits	Match IP fragments.
Time to Live (TTL) value	Time to Live (TTL—IPv4) or Hop Limit (IPv6) value in an IP packet.
DiffServ Code Point (DSCP) bits	Decimal DSCP value.
Layer 4 destination port number	Layer 4 destination port number.
Layer 4 source port number	Layer 4 source port number.
TCP flags	TCP flags to indicate the state of connection.

You can specify the criteria in any combination. Packets matching the defined criteria are captured and saved as pcap files.

Refer to the following sections for details:

- Rules, Notes, and Limitations
- Packet Capture (PCAP)
- Configure PCAP Profile
- View PCAP Files
- Configure PCAP over IP Interface
- View PCAP over IP Interface
- Delete PCAP over IP Interface

Rules, Notes, and Limitations

Refer to the following rules and notes:

- The PCAP feature supports up to 16 capture files per device. Each capture file can have up to 40000 packets. A capture file is maintained per PCAP session. Each session can have up to 64 filter rules per direction. Each capture file can be viewed, deleted, or uploaded out of the device for offline use.
- The PCAP feature supports up to 16 active capture sessions at a time per port on all GigaVUE platforms except GigaVUE-TA400, and GigaVUE-TA400E. The GigaVUE-TA400, and GigaVUE-TA400E platform currently supports one active PCAP session per port at a time.
- The PCAP configuration doesn't persist across node reboots or upgrades.
- The PCAP feature is not supported on stack ports in legacy stacking mode.
- The PCAP feature does not support 'vlan' and 'inner-vlan' filter rules on a tool, hybrid or circuit port in the 'tx' direction.
- The PCAP feature on tool ports does not capture the vlan tag specified with the ingress-vlan-tag feature. To overcome this, redirect the traffic to another hybrid port along with other tool ports and capture the packets on the hybrid port ingress.
- The PCAP feature on GigaVUE-HC1-Plus, GigaVUE-HC1, GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA400, and GigaVUE-TA400E platforms contain extra VLAN header added in capture files. Untagged packet captures contain vlan-tag 1 header added and tagged packet captures contain an outer tag duplicated.
- The PCAP feature on port discovery protocols (LLDP/CDP/GDP) enabled ports will not capture the discovery protocol control packet in the PCAP file.
- The PCAP feature may miss some packets in the capture file depending on the rate of traffic being captured.
- When the PCAP feature is enabled for both directions to capture TX and RX, bi-directional traffic must be sent.
- The PCAP feature supports a maximum of four active capture sessions when the direction is configured as TX or Both.
- GigaVUE-OS cannot capture ARP/ICMP control packets on an IP Interface.
- When PCAP is configured to capture TX direction traffic on an IP interface for Embedded Tunnels, only the actual packets are visible, not the L2GRE/VxLAN header. This is because the packet is mirrored to the PCAP interface before encapsulation occurs. This does not affect packet capture for GigaSMART tunnels, where the actual packets with headers can be seen.
- Q-in-Q packets cannot be captured on the egress port when any filter rule is applied.

Configure PCAP Profile

To configure PCAP through GigaVUE-FM:

1. From the device view, go to **Ports > Ports > All Ports.**
2. Select the required port/ports for which you need to configure PCAP.

   NOTE: You can configure PCAP only for a maximum of four ports at a time.

3. Click **Action** and select **Configure PCAP**.

   The **Action** button is disabled:

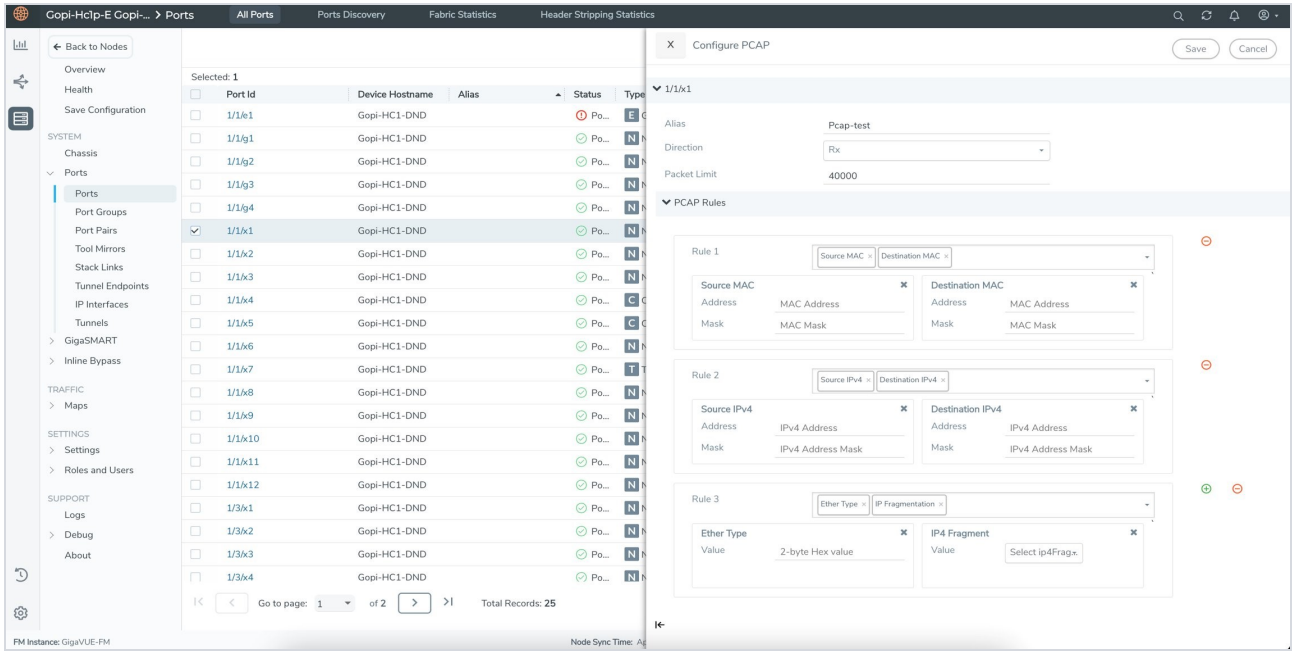
   - If you select more than four ports.
   - If you do not select any port.
   - If you select GigaSMART Engine ports or other unsupported port types.

   If the PCAP feature is disabled by the customer support team, a banner notification is displayed.

   The **Action** button is hidden:

   - For G-TAP devices.
   - For devices running software version less than 5.16.00 and managed by GigaVUE-FM.

4. Select or enter the following details:



Field	Description
Alias	Name of the packet capture filter
Direction	The direction of traffic. Can be: • Rx • Tx • Both
Channel Port	The channel port identifier for the packet capture filter. The channel port is any unused port that does not have any map configuration. The channel port must be on the same node as the capture port. The channel port must be administratively enabled and must remain enabled while a packet capture filter is configured. You must specify one channel port for each transmitted or both direction. channel port is not needed for received direction. **Note**: If a PCAP configuration is deleted, the channel ports configured in the PCAP will go down.
Packet Limit	The number of packets to capture. The valid range is 1 to 40000 for all the platforms. Use the packet limit to stop packet capture after a specified number of packets have been captured. The default value is 40000 for all the platforms.
PCAP Rules	The rules are based on which the traffic will be filtered. You can add multiple filters to the same PCAP. Select the required rule: • **Source MAC**: The source MAC address and MAC netmask. • **Destination MAC**: The destination MAC address and MAC netmask. • **VLAN**: The VLAN ID value as a number between 1 and 4094. • **Inner VLAN**: The inner VLAN ID value as a number between 1 and 4094. • **Ether type**: The layer 2 ethernet type value. • **Source IPv4**: The source IPv4 address and IP mask or a wildcard with an IP mask. • **Destination IPv4**: The destination IPv4 address and IP mask or a wildcard with an IP mask.

Field	Description

	• **Protocol**: The valid protocols and their hex values are as follows:
	  • ipv6-hop (0x0
	  • icmp-ipv4 (0x1)
	  • igmp (0x2)
	  • ipv4ov4 (0x4)
	  • tcp (0x6)
	  • udp (0x11)
	  • ipv6 (0x29)
	  • rsvp (0x2E)
	  • gre (0x2F)
	  • icmp-ipv6 (0x3A)
	  • A custom-defined value can also be defined in 1 byte hex.
	• **IP version**: The IP version for traffic, either IPv4 or IPv6.
	• **IP4 Fragment**: IP fragments, such as no-frag, all-frag, all-frag-no-first, first-frag, and first-or-no-frag.
	• **TTL**: The Time to Live (TTL—IPv4) or Hop Limit (IPv6) value in an IP packet, as a number between 0 and 255.
	• **DSCP**: The decimal DSCP value. Any value within the four Assured Forwarding (af) class ranges or (ef) for Expedited Forwarding. The valid DSCP values by Assured Forwarding Class are as follows:
	  • Class 1—11, 12, 13
	  • Class 2—21, 22, 23
	  • Class 3—31, 32, 33
	  • Class 4—41, 42, 43
	  • Expedited Forwarding—ef
	• **Port Source**: The Layer 4 source port number, from 0 to 65535. A range of ports is not supported.
	• **Port Destination**: The Layer 4 destination port number, from 0 to 65535. A range of ports is not supported.
	• **TCP Control**: TCP control bits, such as SYN, FIN, ACK, URG, as 1 byte hex values.

5. Click **Save** to save the configuration.

The captured packets are stored as pcap files. When multiple filters are configured, the traffic matching each filter is stored in a pcap file for each session under */var/log/tmp* directory in the device. Refer to View PCAP Files for details on viewing the PCAP files.

To configure PCAP from device CLI, refer to the GigaVUE-OS CLI Reference Guide.

View PCAP

To view the configured PCAPs:

1. Click **Action** and select **View PCAP**.
2. The configured PCAPs can be viewed.

Delete PCAP

To delete the configured PCAPs:

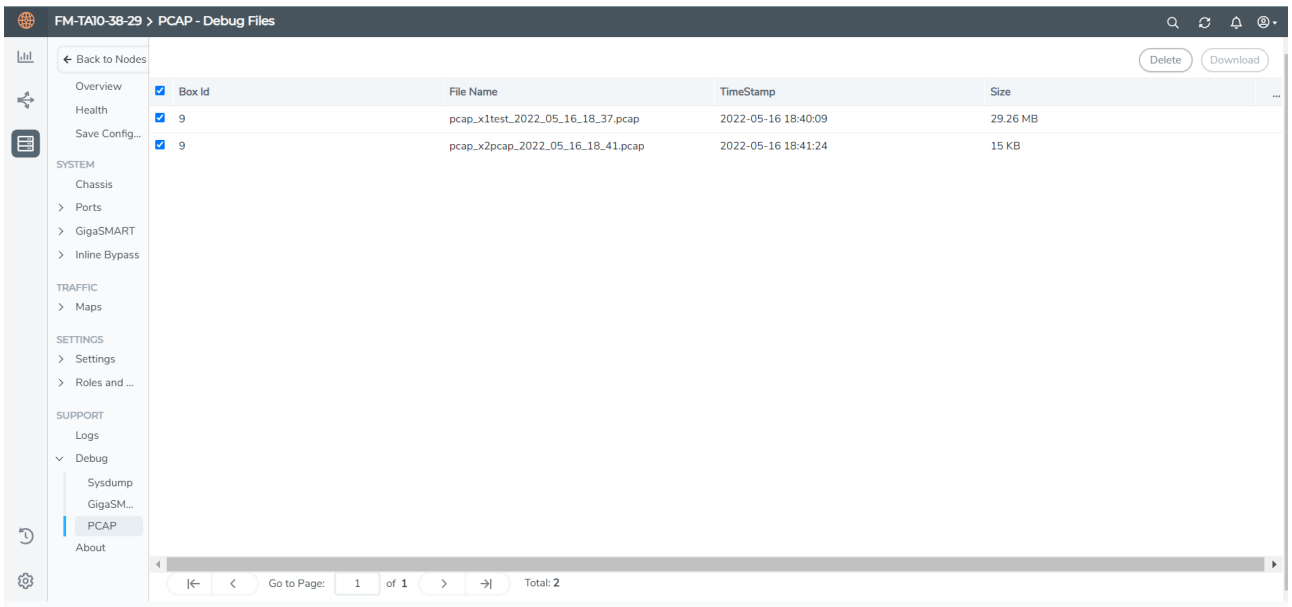
1. Click **Action** and select **Delete PCAP**.
2. Select the required PCAP configurations that you want to delete.

Refer to the GigaVUE-OS CLI Reference Guide for details on configuring PCAP from CLI.

View PCAP Files

You can view and download the PCAP files from GigaVUE-FM. To view the PCAP files:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**.
2. Select a cluster ID, and then from the left navigation pane, go to **Support > Debug > PCAP**.



3. Select the required PCAP file(s):
   - Click **Download** to download the file. You can download only one file at a time.
   - Click **Delete** to delete the PCAP files.

Configure PCAP on IP Interface

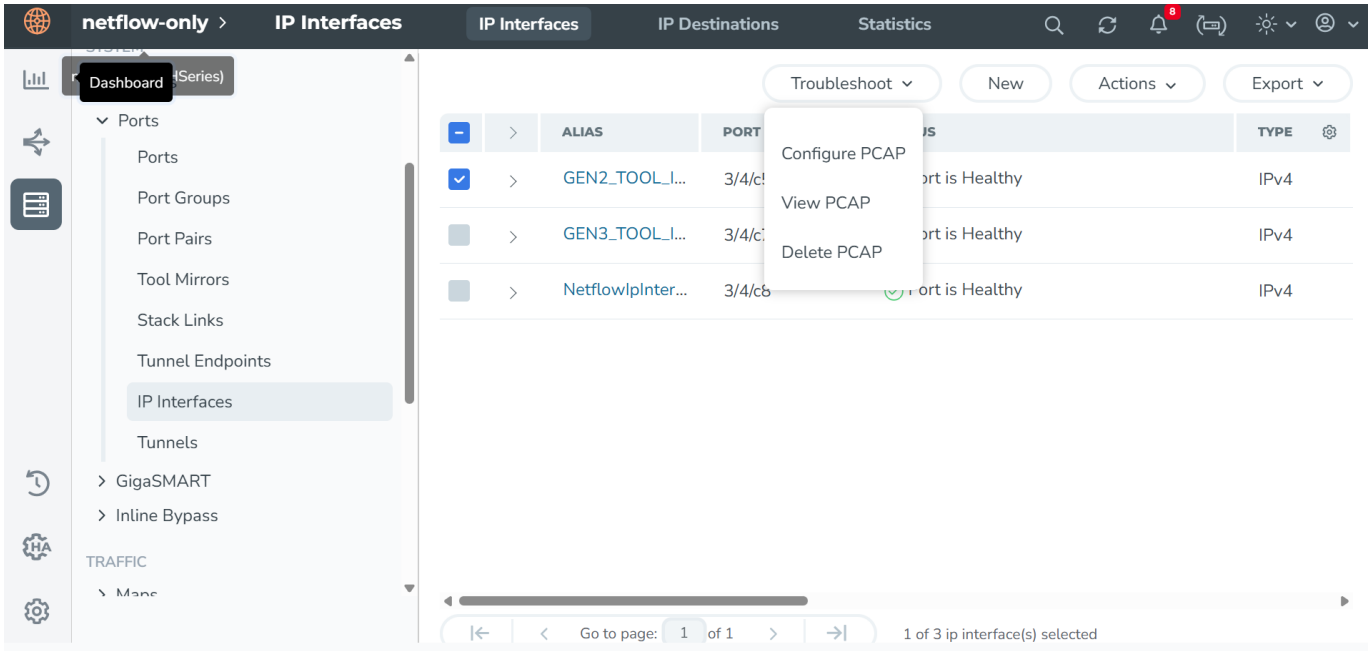
The PCAP over IP Interface feature extends GigaVUE-OS packet capture capabilities to ports associated with IP interfaces, enabling capture of traffic received from and transmitted to those interfaces for deep analysis and troubleshooting. It preserves existing PCAP behavior and performance characteristics while ensuring that critical control traffic such as ARP and ICMP remains unaffected, and is supported across all GigaVUE HC Series and GigaVUE TA Series devices.

To configure PCAP IP Interface, do the following:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link where you want to create the packet capture. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > Ports > IP Interfaces**. The IP Interfaces page appears.
4. Select the alias checkbox, that you want to configure the PCAP IP Interface.

NOTE: A maximum of four alias checkbox can be selected to configure the PCAP IP Interface.

5. Click **Troubleshoot**.



6. Select **Configure PCAP** from the list.
   - Alias - Enter the name for the alias.
   - Direction - (Optional) Select either Rx, Tx, or Both.

NOTE: If you select Tx or Both directions, then you need to select the channel port from the list.

- Channel Port - Select the channel port from the list.
- Packet Limit - (Optional) Enter the value from 1 to 40000. If no value is provided, by default it will take as 40000.
- Rules - (Optional) You can add a rule. Click Add a Rule, and select Condition from the list, and provide the values.

7. Click **Save**.

View PCAP on IP Interface

To view PCAP IP Interface, do the following:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link where you want to view the packet capture. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > Ports > IP Interfaces**. The IP Interfaces page appears.
4. Select the alias checkbox, that you want to view the PCAP IP Interface. The view PCAP configuration page appears.
5. Click **Close**.

Delete PCAP on IP Interface

To delete PCAP IP Interface, do the following:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link where you want to delete the packet capture. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > Ports > IP Interfaces**. The IP Interfaces page appears.
4. Select the alias checkbox, that you want to delete the PCAP IP Interface. The delete PCAP configuration page appears.
5. Select the checkbox on the right-side to delete the PCAP over IP Interface.
6. Click **Delete**.

Flow Mapping®

This chapter provides the following information about flow mapping:

- About Flow Mapping®
- Manage Maps
- Flow Mapping® FAQ
- Configure Active Visibility

About Flow Mapping®

This section describes what is Flow Mapping® and how to apply it to GigaVUE® HC Series and TA Series nodes. Refer to the following sections for details:

- Flow Mapping® Overview
- Get Started with Flow Mapping®
- Flow Map Syntax and Construction
- Work with Map-Passalls and Port Mirroring in GigaVUE-FM
- Port Access and Map Sharing
- Map Examples

Flow Mapping® Overview

Flow Mapping® is the technology found in GigaVUE nodes that takes line-rate traffic at 1Gb, 10Gb, 40Gb, or 100Gb from a network TAP or a SPAN/mirror port (physical or virtual) and sends it through a set of user-defined map rules to the tools and applications that secure, monitor, and analyze IT infrastructure. Flow Mapping® provides **superior granularity and scalability** above and beyond the capabilities of connection and ACL filter based technologies by addressing the problems inherent when going beyond small numbers of connections or when more than one traffic distribution rule is required.

Flow Mapping® can granularly filter and forward traffic to specific monitoring tools through thousands of map rules with criteria based on over 30 predefined Layer 2, Layer 3, and Layer 4 parameters including IPv4/IPv6 addresses, application port numbers, VLAN IDs, MAC addresses and more. Users can also define custom rules that match specific bit sequences in the traffic streams, applying Flow Mapping® to tunneled traffic, specialized applications and even higher-layer protocols.

In addition, IT operations can deploy Visibility as a Service (**VaaS**), taking advantage of the Flow Mapping® **role based access control (RBAC)** features on GigaVUE nodes. Users can be given access to traffic based on their needs without interfering with the monitoring operations of the other teams. This helps protect compliance and privacy protocols and allows teams to dynamically react when needed for increased efficiency.

Flow Mapping® can be combined with **GigaSMART** technology to provide packet modification and intelligent capabilities like de-duplication and packet slicing, making tools more efficient by reducing the number and size of packets they have to store and process. Header stripping and de-tunneling functions provide tools access to protocols and data they would otherwise by blind to.

When multiple GigaVUE nodes are in a **stacked or clustered configuration**, Flow Mapping® enables traffic to be sent from any network port to any tool port, expanding visibility beyond a single rack, row, or data center. GigaSMART can be leveraged on all traffic flow, accepting traffic from any network port and regardless of where the GigaSMART hardware is located within the stack or cluster.

To access Flow Maps

1. Go to **Traffic** and select **Physical** > **Flow Maps**. This displays the list of Devices/Cluster Nodes managed by this instance of GigaVUE-FM.
2. Click the Cluster ID of any node to open the node. The following options are displayed on left navigation pane:
   - Maps
   - Workflows
   - Map Templates
   - Filter Templates

Get Started with Flow Mapping®

You can manage packet distribution in both the GigaVUE-OS command-line interface (CLI) and GUI (GigaVUE-FM). Both interfaces allow you to perform all packet distribution tasks:

- designating ports as network or tool ports
- setting up map rules
- mapping network ports to tool ports
- many other such functions

NOTE: These tasks can also be performed with the GigaVUE-FM APIs. For more information about the APIs, refer to *GigaVUE-FM REST API Reference in GigaVUE-FM User's Guide*.

For the setup of a few simple maps, refer to the following sections for examples:

- Example: How to Create a Simple Map
- Example: How to Handle Overlaps when Sending VLANs and Subnets to Different Tools

Check Status of Nodes and Ports

Before configuring maps, check the status of line cards, modules, and ports. To view the status of line cards and modules, select **Chassis** from navigation pane and select Table view. For more information about the Chassis page, refer to the *"Chassis"* section in the *GigaVUE-OS and GigaVUE-FM Administration Guide*. Go to **System > Ports > Ports > All Ports** to see a table with details about each port. For more information about ports, refer to About Ports and Managing Ports.

Designated Port Types

Ports on GigaVUE-OS nodes can be one of the following types:

- network
- tool
- stack
- inline-network
- inline-tool
- hybrid
- circuit

NOTE: Not all port types are supported on all platforms. Inline-network and inline-tool port types are only supported on GigaVUE® HC Series nodes, GigaVUE-TA25, GigaVUE-TA200, GigaVUE-TA400, and GigaVUE-TA400E.

About Shared Collectors

GigaVUE nodes let you create map rules that direct traffic on any network port or ports to any tool ports. Traffic can be dropped intentionally using the drop rule or any packets that do not match any other rule in the map can be sent to the collector. Shared collectors are set up to capture any packets that do not fulfill the map criteria but may be required by other tools.

NOTE: If a shared-collector destination for a set of network ports is not defined, non-matching traffic is silently discarded.

When assigning the priorities to map rules on GigaVUE® HC Series nodes, GigaVUE-TA25 and GigaVUE-TA200, the first rule setup will also have the highest priority unless specified by the user. The shared collector rule is the only exception because it will always have the

lowest priority even if configured first. This means that an incoming packet will be matched against all the rules in the same map and when not matched with any rules, it be forwarded to the designated tool port for the collector.

A GigaStream or multiple sets of GigaStream can also be set as destination for a collector port by using the GigaStream alias.

In cases, where multiple network ports are sharing the multiple maps, packets that do no fit any of the maps can be sent to the shared collector.

No Map Statistics for Shared-Collector Only

A shared collector is intended to be used with other maps. For example, use a shared collector with a regular map containing at least one rule. If there is a shared collector but no other map, there will be no map statistics.

Shared Collector Configuration

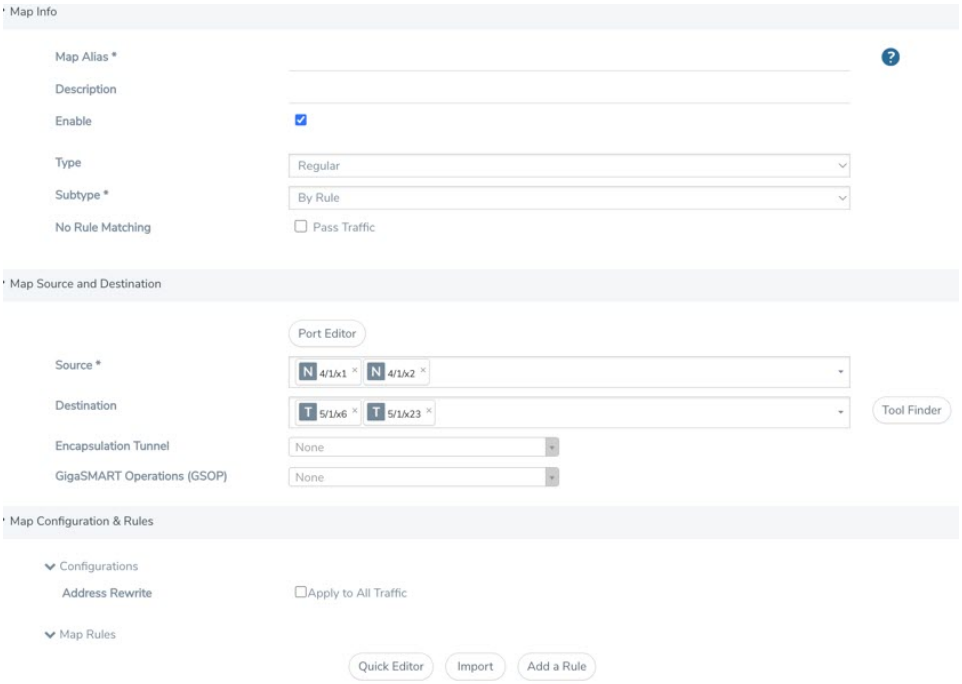
A shared-collector is a special type of map configured with only a set of **Source** ports shared-collector ports or GigaStream. Rules, priority settings, GigaSMART operations and destination ports are not allowed in shared-collector maps. In GigaVUE-FM, the collector ports can be selected from a list of tool or hybrid ports.

To create a shared-collector map, do the following:

1. Select **Maps > Maps > Maps**.
2. Click **New**.
3. Type an alias in the Alias field to identify this map. For example, shared_collector.
4. For **Type**, select **Regular**.
5. For **Subtype**, select **Collector**.
6. Click in the **Source** field and select a network port .

NOTE: Ports already used as source ports in the orchestrated configuration will not be listed in the drop-down.

7. Click in the **Destination** field and select the tool or hybrid ports that will be the shared-collector ports for the map. The map should look similar to the map shown in the following figure.



NOTE: Shared-collector maps do not include any rules, priority settings, or GigaSMART operations. These are grayed out in the UI when Collector is select for the map's subtype.

8. Click **Save**.

In Figure 1Shared Collector Map, the Maps page shows the shared-collector map shared_ collector with the standard components. The **Source** ports match those used by a set of normal flow maps. The **Destination** ports are the collector ports are where you want to send any packets not matching the normal flow maps.

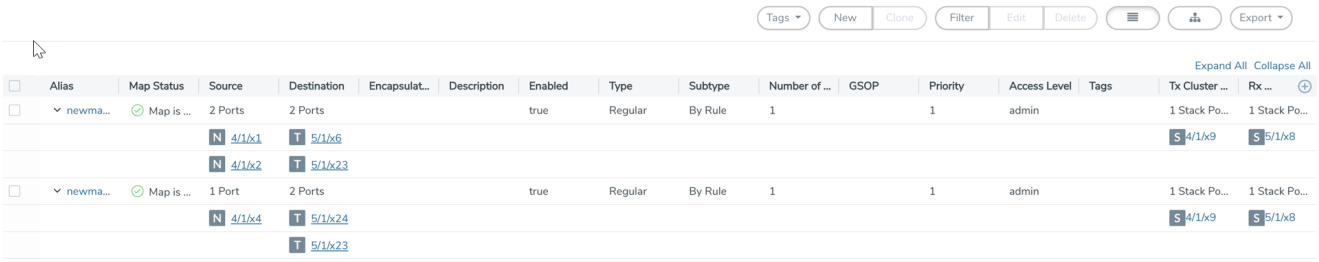


Figure 1 *Shared Collector Map*

About Map-passall Maps

Map-passall maps provide a way to specify a destination for packets without any filtering on a set of network ports. As indicated by the name, all traffic is passed through. A map-passall may share the network ports with a map which filters using map rules.

The same logic as set for Shared Collector can be set for map-passall. That is, that Map-passall can be set to a GigaStream alias or multiple GigaStream aliases or a single tool port or multiple tool ports. To set the **Destination,** use the same map range configuration.

Map-Passall and Regular Byrule Map

When you configure a Regular Byrule Map with more than one Network port, these individual Network ports are part of a Separate Map-Passall with different destinations. Therefore, all traffic received on each Network port is expected to reach all Map-Passall destinations.

Map-Passall and Shared-Collector Only

If a map-passall and a shared-collector both use the same network source port and there are no other maps, such as a regular map containing at least one rule, all traffic will be passed to the shared-collector.

Map-passalls Configuration

A map-passall map is a special type of map configured with only a set of **Source** ports and **Destination ports** or **GigaStream**. Map rules and GigaSMART operations are not allowed in passall maps.

Map-passalls

The web-based GigaVUE-FM interface for GigaVUE-OS nodes provides an all traffic **Pass All** subtype selection for regular maps that performs the same function as a **map-passall** in the CLI. Making this selection inGigaVUE-FM turns the map into a map-passall, delivering all traffic from the selected network ports to another tool port or GigaStream on any line card in the same node, irrespective of the other packet distribution. Although the names are different inGigaVUE-FM and the CLI, the two features are identical.

Define Map Source Port Lists

You can configure more than one map with the same source ports in the **Source** field in a map.

The source port list of one map must be exactly the same as the source port list of another map (have the same ports as well as the same number of ports) and it must not overlap with the source port list of any other map.

For example if map1 has **Source** ports 1/1/x4, 1/1/x5, and 1/1/x6 already configured:

- map2 **Source** ports 1/1/x5, and 1/1/x6 cannot be configured because ports x5 and x6 are configured already in map1 as source ports.

- map3 **Source** ports 1/1/x3, 1/1/x4, 1/1/x5 cannot be configured because ports x4 and x5 are configured already in map1 as source ports.

Share Network Ports Between Maps

Network ports can be shared between a regular map and a map passall, as follows:

- the regular map has network ports 1/1/x3..x4
- the passall map has network port 1/1/x4

When there are overlapping network ports and shared tool ports between a regular map and passall map, the map passall tool ports or GigaStream will receive traffic from the network ports configured on the regular map, in addition to the traffic from its own network port or ports.

In the configuration above, the map passall will also receive traffic from 1/1/x3.

NOTE: A shared collector map and a regular map should have exactly the same set of network ports. (This is the correct use case.) Overlapping network ports should not be configured for collector maps. (This is an incorrect use case whether the overlapping ports are a subset or a superset of the network ports.)Network ports cannot be shared between a regular map and a shared collector map.

Share Tool Ports Between Maps

When a map passall and shared collector share the same tool ports, removing the shared tool ports from the passall map may affect the shared collector traffic. The workaround is to not share the same tool ports between a map passall and a shared collector.

Map Priority

Packets matching multiple maps in a configuration are sent to the map with the highest priority when the network ports are shared among multiple maps with pass-by map rules. By default, the first map configured has the highest priority; however, you can adjust this.

In GigaVUE-FM, the UI displays the maps from highest priority to lowest as top to bottom.

Maps sharing the same source port list are grouped together for the purpose of prioritizing their rules. Traffic is subjected to the rules of the highest priority map first and then the rules of the next highest priority map and so on. Within a map, drop rules are applied first and then pass rules, in other words, drop rules always have higher priority than pass rules. Currently when a map's source port list is defined the map is grouped/prioritized with other maps sharing the same source port list. Newly configured maps are added as the lowest priority map within the group when initially configured unless changed by the user.

NOTE: Shared collector will always go to the lowest priority when setting up maps.

Adjust Map Priority in GigaVUE-FM

Before you get started adjusting map priority, start by reviewing the current map priorities in place by opening the Maps page and viewing the priority of the maps in the Priority field. For example, Figure 2Map Priorities shows three maps MyMap1, MyMap2, and MyMap3 with the same source port 1/1/x8. The Priority column in the table shows the current priority of each map.

Alias	Comments	Type	Sub Type	Source	No of Rules	GSOP	Priority	Access Level	Destination
GTP-Sampling-2		secondLevel	flowSample	vport_elias_test	1	gtp_flow_sampling_elias_test	1	admin	T 1/1/x4
MyMap1		regular	byRule	N 1/1/x8	0		1	admin	T 1/1/x7
MyMap2		regular	byRule	N 1/1/x8	0		2	admin	T 1/1/x10
MyMap3		regular	byRule	N 1/1/x8	0		3	admin	T 1/1/x15
OpenStack_vTraffic_toWireshark		regular	byRule	N vTunnelEndpointForOpenStack	1	GigavueVM_Tunnel	1	admin	T toRSASecurityAnalytics

Figure 2 *Map Priorities*

Then, once you have reviewed the existing hierarchy of map priorities, you can fine-tune the priority of maps by using the **Priority** field in the map to select one of the following:

- Highest (top)—set the map to the highest priority
- After map - <map-alias> — set the map priority after the map with the specified alias.
- Lowest (bottom)—set the map priority to the lowest priority

Flow Map Syntax and Construction

This section provides information about map types, map rules, and working with map passalls.

Map Types

Map configuration consists of several parameters, including the following:

- **Source**—Specifies the source ports for the map.
- **Destination**—Specifies the destination ports for the map.
- **Type**—Specifies the type of map.
- **Rules**—Specifies the rules for the map.
- **Subtype**—Specifies map subtype.
- **GSOP**—Specifies a GigaSMART operation.

There are four types of maps, with the map type being determined by the **Source** and **Destination** as well as the map rules as follows:

- **Regular**—Specifies a regular map type, with the **Source** parameter specifying network or hybrid ports, or single inline-network or single inline-tool ports (for out-of-band maps) and the **Destination** parameter specifying tool, hybrid, or GigaStream ports.

  When the subtype for a **Regular** map type is **By Rule**, a **Pass Traffic** option for **No Rule Matching** is available. Selecting Pass Traffic specifies what to do with traffic that does not match any rule in a map that only has drop rules. This argument changes the default behavior of drop to pass in a drop-only map. If you do not use this argument and there are only drop rules in a map, the default behavior is that all traffic not matching the rules will be dropped, or, if a shared collector is configured, traffic will be sent to the shared collector. However, if you use the Blacklisting option and there are only drop rules in a map, traffic will be passed rather than dropped.
- **inline**—Specifies an inline map type, with the **Source** parameter specifying inline-network pairs or inline-network-groups and the **Destination** parameter specifying inline-tool pairs, inline-tool-group, inline-serial, or bypass.
- **First Level**—Specifies a first level map type, with the **Source** parameter specifying network or hybrid ports and the **Destination** parameter specifying virtual ports, used with GigaSMART operations.

  When a First Level map type is selected, a **Control Traffic** option is available. This option is for GTP applications to pass GTP control traffic (GTP-c) to all GigaSMART engines in a GTP engine group. To enable GTP-c, select **Control**.
- **flexinline**—Specifies a flexible inline map, with the **from** parameter specifying inline-network and the **a-to-b** or **b-to-a** parameters specifying an ordered list of inline-tool or inline-tool-group.

Define Map Source Port Lists for First Level Maps

You can configure more than one map with the same source ports in the **Source** field in a map.

The source port list of one map must be exactly the same as the source port list of another map (must have the same ports as well as the same number of ports) and it must not overlap with the source port list of any other map.

For example:

- map1 has Source ports 1/1x1, 1/1x2 already configured.
- map 2 with Source ports 1/1x1, 1/1x2 can also be configured.
- map3 with Source port 1/1/x3 can be configured. If you try to add source ports 1/1x1, 1/1x2 as part of map3 then it is not possible because 1/1x1, 1/1x2 are already part of map 1.

NOTE: This is also applicable for regular maps.

- **Second Level**—Specifies a second level map type, with the **Source** parameter specifying virtual ports, used with GigaSMART operations, and the **Destination** parameter specifying tool, hybrid, or GigaStream ports.
- **Transit Level** — Specifies the transit level map type, with the source port specifying a virtual port (VP1) and destination port specifying an another virtual port (VP2) with GSRULE and GSOP associated to perform the GigaSMART operations. In the transit-level map you can interconnect different v-ports to chain the operations as needed. For unsupported GigaSMART operations, you can:
  - create a transit level for a set of operation
  - create a second level map for another set of operation
  - combine the two maps

NOTE: If a transit-level or second-level map is used in the configuration, APF must be enabled in the GSOP when used as part of the transit-level or second-level map.

NOTE: There is also a **Template** map type for creating map templates.

The following table provides information about the GigaSMART operations supported by transit-level maps in SMT-HC0-Q02X08 of HC2P, HC1-X12G4 of HC1 with native, SMT-HC3-C08, SMT-HC1-S, and SMT-HC1A-R cards:

GS Apps Supported in Transit Map	Transit Map in HC1-X12G4 of GigaVUE-HC1 with native	Transit Map in SMT-HC3-C08	Transit Map in SMT-HC1-S	Transit Map in SMT-HC1A-R
De-duplication	√	√	√	√
Slicing	√	√	√	√
Masking	√	√	√	√
Header Stripping	√	X	X	X
Trailer	√	X	X	X
Adaptive Packet Filtering (APF)	√	√	√	√
Application Session Fitlering (ASF)	√	√	√	√
Application Filtering Intelligence (AFI)	√	√	√	√
Application Metadata Intelligence (AMI)	√	√	√	√

NOTE: V-port chaining across GigaSMART groups is not supported in SMT-HC3-C08 and SMT-HC1-S cards. V-port chaining across GigaSMART groups in different cards is also not supported in SMT-HC3-C08 and SMT-HC1-S cards. Application Filtering Intelligence (AFI) that is configured in Application Intelligence (AI) session uses Hybrid port.

Map types are described in Table 1: Matrix of Map Types.

Table 1: Matrix of Map Types

Type	Subtype	Map Rule	GSOP	Source	Destination
Regular	By Rule, or Pass All, or Collector	Rule	yes	network ports, or hybrid ports, or single inline-network ports or single inline-tool ports (for out-of-band maps)	tool ports, or hybrid ports, or GigaStream
Inline	By Rule, or Pass All, or Collector	Rule	no	inline-network pairs, or inline-network-groups	inline-tool pairs, or inline-tool-groups, or bypass, or inline-serial tools
First Level **Note**: The Gen3 GigaSMART cards support a maximum of 230 first-level maps per gsgroup and vport, when gsop is associated with flow sample or flow whitelist applications.	By Rule	Rule	no	network ports, or hybrid ports	virtual ports, tool ports, or hybrid ports, or GigaStream, (collector is not allowed)
Transit Level	By Rule	gsrule	yes	virtual ports	virtual ports
Second Level	By Rule, or Collector	gsrule	yes	virtual ports	tool ports, or hybrid ports, or GigaStream, or port group
	Flow Filter	flowrule	yes		
	Flow Sample	flowsample	yes		
	Flow Sample Overlap	flowsample	yes		
	flowSample-sip	flowsample	yes		
	Flow Whitelist	whitelist	yes		
	Flow Whitelist Overlap		yes		
flexinline	byRule collector	ordered list of inline-tool or inline- tool-group in a-to-b or b-to-a direction or both	rule -	inline-network alias	inline-tool pairs, or inline-tool-groups, or bypass

Null Port in Maps

A map can be configured with only the source port without specifying any destination port (null-port in command-line reference). A map is still created, but the traffic is dropped.

If a null-port is specified in map, the GigaSMART operation is performed on the traffic, and the traffic is finally dropped. After dropping the traffic, the GigaSMART packet drop counter is incremented. If GigaSMART operations are chained, then after applying all the GigaSMART operations on traffic, the traffic is dropped.

NOTE: This feature is supported on GigaVUE HC series devices on both stand-alone nodes and clusters. Null port is applicable for regular maps and second-level maps involving GigaSMART operations.

Map Subtypes

The map subtype describe in Matrix of Map Types is optional. It specifies the following:

- **By Rule**—Specifies a rule-based map subtype, which is supported on the following map types:
  - **First Level**, **inline**, and **Regular** map types.
  - **Transit Level** map type.
  - **Second Level** map type.
- **Pass All**—Specifies a passall map subtype, which applies to **regular** and **inline** map types. The **Pass All** subtype is not supported on **First Level** and **Second Level** map types. With this subtype, map priority cannot be configured or modified.
- **Collector**—Specifies a collector map subtype, which applies to **Regular**, **inline**, and **Second Level** map types. The **Collector** subtype is not supported on the **First Level** map type. With this subtype, map priority cannot be configured or modified.
- **Flow Filter**—Specifies a flow filtering map subtype, which applies to **Second Level** map types. Specify the **Flow Filter** map subtype when using a **Flow Filter** parameter.
- **Flow Sample**—Specifies a flow sampling map subtype, which applies to **Second Level** map types. Specify the **Flow Sample** map subtype when using a **Flow Sample** parameter.
- **Flow Sample Overlap**—Specifies a flow sampling overlap map subtype, which applies to **secondLevel** map types. Specify the **flowSample-ol** map subtype when using a **flowsample** rule.
- **Flow Sample sip**—Specifies a SIP flow sampling map subtype, which applies to **secondLevel** map types.
- **Flow Sample Overlap**—Specifies a flow sampling overlap map subtype, which applies to **Second Level** map types. Specify the **Flow Sample Overlap** map subtype when using a flowsample rule.
- **Flow Whitelist**—Specifies a whitelist map subtype, which applies to **Second Level** map types. Specify the **Flow Whitelist** map subtype when using a whitelist rule.

- **Flow Whitelist Overlap**—Specifies a whitlelist overlap map subtype, which applies to Second Level map types, Specify **Flow Whitelist Overlap** when using a whitelist rule.
- **Flow Whitelist-sip**—Specifies a SIP flow whitelist map subtype, which applies to **secondLevel** map types.

The default map subtype is **By Rule**.

NOTE: Maps with subtype **Flow Sample Overlap** or **Flow Whitelist Overlap** do not support map editing.

If a By Rule map which has overlapping network and tool ports with pass-all map is disabled, traffic will still be forwarded to tool ports. Re-configure the maps to avoid traffic loss.

Map Type and Subtype Modification

Once a map is created, the map **Type** and **Subtype** cannot be modified. However, you can delete the map and recreate it with a different **Type** and **Subtype**.

Backwards Compatibility

For backwards compatibility, the map **type** parameter does not have to be configured. The **type** and **subtype** will be determined by the system based on the remainder of the map configuration parameters. If not enough information is available, the default values of **regular** and **byRule** will be assumed for the type and subtype.

Minimum Requirements for Map Creation

A map must be configured with at least a **from** parameter. Even if other parameters such as **to**, **rule**, or **use** are configured, without **from**, the map will not be created.

Map Rules

This section provides information about the different types of map rules that you can specify when creating maps in GigaVUE-FM. The following topics are covered:

- Other Types of Map Rules for GigaSMART Operations
- IPv4/IPv6 and Map Rules
- Set Map Rules for TCP Control Bits
- How to Use Bit Count Netmasks
- How to Combine Rules and Rule Logic
- How to Mix Pass and Drop Rules
- Configuring Port Criteria and Bi-Directional Rules in By Rule Maps

- Work with User-Defined Pattern Match Rules
- Inner Header and MPLS Header Filtering

Other Types of Map Rules for GigaSMART Operations

There are other types of rules for GigaSMART operations as follows:

- Adaptive Packet Filtering (APF) and Adaptive Session Filtering (ASF). For details, refer to GigaSMART Adaptive Packet Filtering (APF) and GigaSMART Application Session Filtering (ASF) and Buffer ASF.
- GTP Correlation. For details, refer to GigaSMART GTP and CUPS Correlation.
- GTP whitelisting and GTP flow sampling. For details, refer to GigaSMART GTP Whitelisting and GTP Flow Sampling.

IPv4/IPv6 and Map Rules

GigaVUE-OS provides a variety of criteria for pass/drop rules specific to IPv6 traffic, including:

IPv6 Entity	Rule Condition
IPv6 Source/Destination Addresses	IPv6 Source/IPv6Destination
IPv6 Flow Labels	IPv6 Flow Label
IPv6 Traffic	IP Version and Version it set to v6

In addition to the explicit IPv6 filters listed in the table, you can use the **IP Version** condition to change how some of the other attributes are interpreted.

When **IP Version** is used by itself in a map rule, it returns all traffic matching the specified IP version, **4** or **6**. However, when **IP Version** is set to **6**, several of the other arguments are interpreted differently when used in the same rule, as follows:

Condition	IP Version set to 4 (or not specified)	IP Version set to 6
Port Destination/Port Source	Matches all IPv4 traffic on the specified port number.	Matches all IPv6 traffic on the specified port number.
IPv4 Protocol	When used with the **<1-byte-hex>** argument, matches against the protocol field in the standard IPv4 header.	When used with the **<1-byte-hex>** argument, matches against the Next Header field in the standard IPv6 header.

Condition	IP Version set to 4 (or not specified)	IP Version set to 6
	Note: These fields perform essentially the same service in both versions, specifying what the next layer of protocol is. However, they have different names and are found at different locations in the header. Refer to Protocol Map Rules and IPv6 for a list of the useful values for the <1-byte-hex> field.	
IPv4 TTL	Matches against the standard TTL (time-to-live) field in the IPv4 header.	Matches against the standard Hop Limit field in the IPv6 header.
	Note: These fields perform essentially the same service in both versions, specifying how long a datagram can exist.	

NOTE: The **IP Version** argument is implicitly set to **4** – if you configure a map rule without **IP Version** specified, the GigaVUE® HC Series node assumes that the IP version is 4.

Protocol Map Rules and IPv6

The predefined protocol map-rules available for IPv4 (GRE, RSVP, and so on) are not allowed when **IP Version** is set to **6**. This is because with the next header approach used by IPv6, the next layer of protocol data is not always at a fixed offset as it is in IPv4.

To address this, the GigaVUE HC Series node provides the <1-byte-hex> option to match against the standard hex values for these protocols in the Next Header field. The standard 1-byte-hex values for both IPv4 and IPv6 are set by selecting the option from the **Value** list or specifying a decimal value when **IPv4 Protocol** is selected. The following are options or values that you can select from the **Value** list.

- IPv6Hop 0
- ICMP_IPv4 1
- IGMP 2
- IPv4 4
- TCP 6
- UDP 17
- IPv6 41
- RSVP 46
- GRE 47
- ICMP_IPv6 58

You can also specify a custom value (0-255) by selecting **Custom**. The following are some additional values and their meanings:

- 43: Routing Option (v6 only)
- 44: Fragment (v6 only)
- 50: Encapsulation Security Payload (ESP) Header

- 51: Authentication (v6 only)
- 59: No Next Header (v6 only)
- 60: Destination Option (v6 only)

Set Map Rules for TCP Control Bits

Select the **TCP Control** to set map rules matching one-byte patterns for the standard TCP control bits. The following table summarizes the bit positions of each of the flags, along with their corresponding hexadecimal patterns.

NOTE: Rules using the **TCP Control** must also include **IPv4 Protocol** and the Value set to **TCP 6**.

Flag	Bit Position	Pattern	TCP Control Mask
Congestion Window Reduced	X...	0x80	0x3f
ECN Echo	.X..	0x40	0x3f
Urgent Pointer	**..X. ....**	0x20	0x3f
Acknowledgment	**...X ....**	0x10	0x3f
Push	**.... X...**	0x08	0x3f
Reset	**.... .X..**	0x04	0x3f
SYN	**.... ..X.**	0x02	0x3f
FIN	**.... ...X**	0x01	0x3f

Examples

The map rule shown in Figure 3Map Rule with SYN Bit Set matches packets with only the SYN bit set:

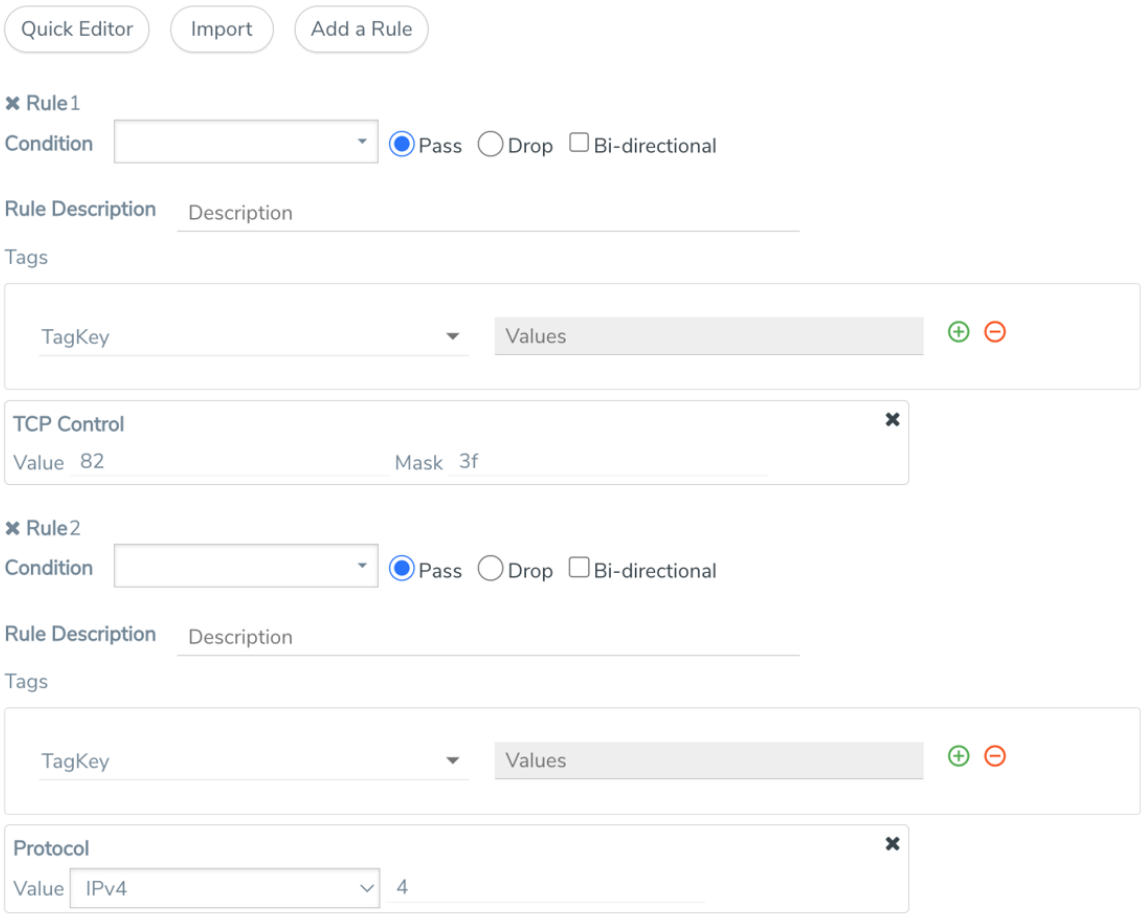


Figure 3 *Map Rule with SYN Bit Set*

Many packets will have some combination of these bits set rather than just one. So, for example, the map rule in Figure 4Map Rule Matching All Packets with Both ACK and SYN Bits set matches all packets with both the ACK and SYN bits set.

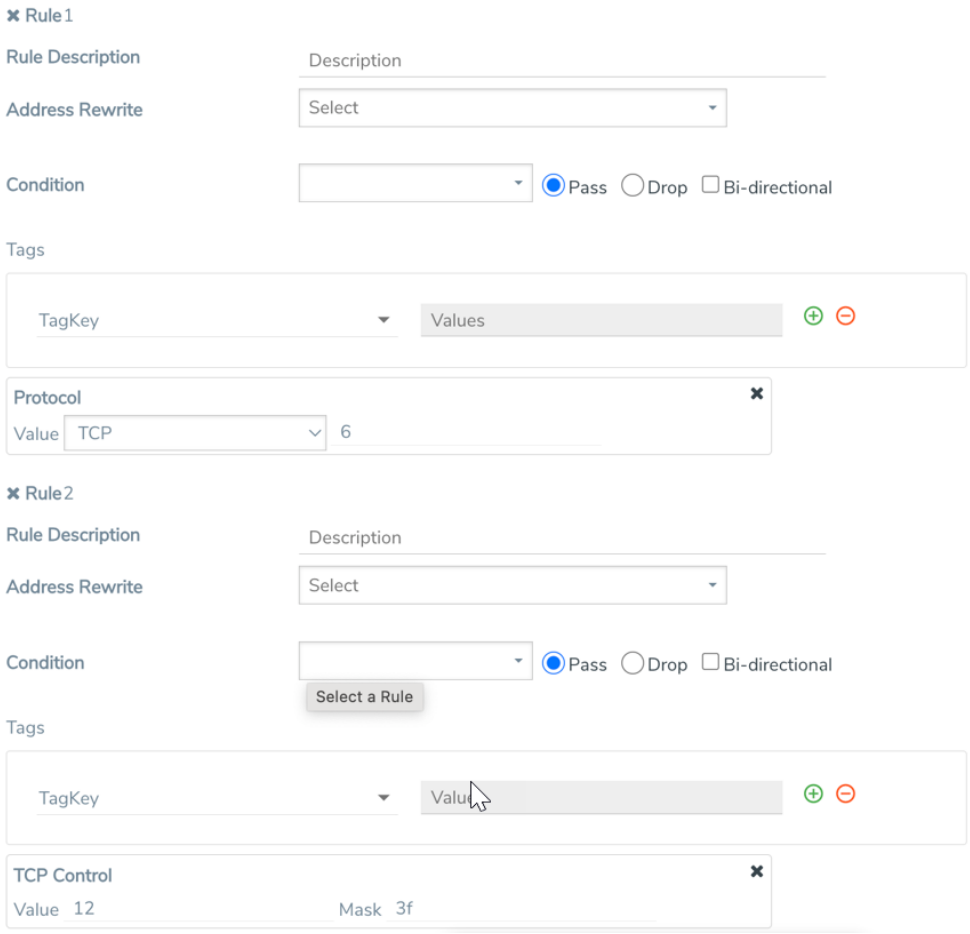


Figure 4 *Map Rule Matching All Packets with Both ACK and SYN Bits set*

How to Use Bit Count Netmasks

The following table summarizes the bit count netmask value for standard dotted-quad IPv4 netmasks. You can enter IP netmasks in the bit count format by using the **/nn** argument.

Bit count netmasks are easier to visualize for IPv6 addresses, specifying which portion of the total 128 bits in the address correspond to the network address. So, for example, a netmask of /64 indicates that the first 64 bits of the address are the network address and that the remaining 64 bits are the host address. This corresponds to the following hexadecimal netmask:

ffff:ffff:ffff:ffff:0000:0000:0000
ffff:ffff:ffff:ffff:0000:0000:0000

Standard Netmask	Bit Count Netmask
255.255.255.255	/32
255.255.255.254	/31
255.255.255.252	/30
255.255.255.248	/29
255.255.255.240	/28
255.255.255.224	/27
255.255.255.192	/26
255.255.255.128	/25
255.255.255.0	/24
255.255.254.0	/23
255.255.252.0	/22
255.255.248.0	/21
255.255.240.0	/20
255.255.226.0	/19
255.255.192.0	/18
255.255.128.0	/17
255.255.0.0	/16
255.254.0.0	/15
255.252.0.0	/14
255.248.0.0	/13
255.240.0.0	/12
255.226.0.0	/11
255.192.0.0	/10
255.128.0.0	/9
256.0.0.0	/8
254.0.0.0	/7
252.0.0.0	/6
248.0.0.0	/5
240.0.0.0	/4
226.0.0.0	/3
192.0.0.0	/2
128.0.0.0	/1
0.0.0.0	/0

How to Combine Rules and Rule Logic

When working with maps, you can easily combine multiple criteria into a single rule. GigaVUE-OS processes rules as follows:

- Within a single rule, criteria are joined with a logical **AND**. A packet must match each of the specified criteria to satisfy the rule.
- Within a map, rules are joined with a logical **OR**. A packet must match at least ONE of the rules to be passed or dropped.

NOTE: When used in a map rule with multiple criteria, the **ipver** argument changes the interpretation of some map rule arguments. Refer to IPv4/IPv6 and Map Rules for details.

Examples of Map Rule Logic

For example, the rules shown in the following table are both set up with criteria for **vlan 100** and **portsrc 23**.

- The first example combines the two criteria into a single rule. This joins the criteria with a logical **AND**.
- The second example creates two separate rules – one for each of the criteria. This joins the criteria with a logical **OR**.

How to Mix Pass and Drop Rules

GigaVUE-OS lets you mix pass and drop rules on a single port. Mixing pass and drop rules can be useful in a variety of situations. Figure 5Pass and Drop Rules in a Map shows a pass rule set up to include all traffic matching a particular source port range combined with a drop rule configured to exclude ICMP traffic.

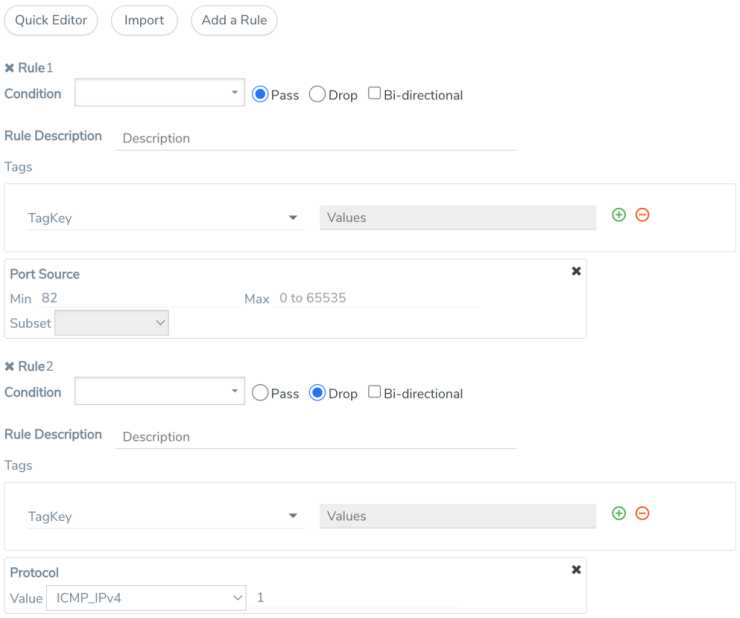


Figure 5 *Pass and Drop Rules in a Map*

Drop Rules Have Precedence!

Keep in mind that within a map, drop rules have precedence over pass rules. So, if a packet matches both a pass and a drop rule in the same map, the packet is dropped rather than passed.

Configuring Port Criteria and Bi-Directional Rules in By Rule Maps

When configuring a By Rule map in the device, it's essential to understand how to accurately set port criteria and leverage the Bi-Directional option to ensure comprehensive traffic matching.

Defining Port Ranges and Single Ports

Some rule conditions, such as Port Destination or Port Source, require you to input Minimum and Maximum values to define port ranges. Here's how to use these fields effectively:

Single Port Matching:

To match a specific port (e.g., port 80), enter the same value in both the Min and Max fields.

Example:

Min: 80

Max: 80

Port Range Matching:

To match a range of ports (e.g., 1000 through 2000), enter the lower value in Min field and the higher value in Max field.

Example:

Min: 1000

Max: 2000

This method ensures precision in traffic filtering and is applicable to all rule types using Min/Max configurations.

Using the Bi-Directional Option

The Bi-Directional checkbox significantly simplifies rule setup by enabling mirrored traffic matching. Typically, you must choose between Source or Destination for conditions like ports or IPs. Enabling Bi-Directional creates a reverse match automatically.

Example:

If a rule matches traffic with destination port 80, enabling Bi-Directional will also match traffic where port 80 is the source. This eliminates the need to manually define both directions.

Work with User-Defined Pattern Match Rules

GigaVUE-OS lets you create pass and drop map rules with *pattern matches* to search for a particular sequence of bits at a specific offset in a packet. You can configure up to two user-defined, 16-byte **pattern matches** in a map rule. A **pattern** is a particular sequence of bits at a specific location in a frame.

NOTE: Refer to User-Defined Pattern Match Examples for step-by-step instructions on creating a real-world pattern-match map rule.

User-defined pattern matches consist of the following components:

Step	Description
Pattern	Use the **UDA1 Value** and **UDA2 Value** fields for map rule to set up the actual bit patterns you want to search for. Refer to User-Defined Pattern Match Examples for details.
Mask	Use the **UDA1 Mask** and **UDA2 Mask** fields for map rules to specify which bits in the pattern must match to satisfy the map rule.
Offset	Use the **UDA1 Offset** and **UDA2 Offset** fields for map rules to specify where in the packet bits

Step	Description
	must match. **Note**: The GigaVUE® HC Series node accepts a maximum of two offsets per device. When both of the available offsets for the device are in use with existing map rules, you will not be able to add a new rule with a different value for **UDAx Offset** until at least one of the **UDAx Offset** is freed up from all existing map rules.

User-Defined Pattern Match Syntax

The user-defined pattern match syntax is as follows:

- Both the **UDA*x* Value** and **UDA*x* Mask** arguments are specified as 16-byte hexadecimal sequences. Specify the pattern in four 4-byte segments separated by hyphens. For example:

  0x01234567-89abcdef-01234567-89abcdef
- Masks specify which bits in the pattern must match. The mask lets you set certain bits in the pattern as wild cards – any values in the masked bit positions will be accepted.
  - Bits masked with binary 1s must match the specified pattern.
  - Bits masked with binary 0s are ignored.
- You can set up the two global offsets allowed per device at 4-byte boundaries beginning at frame offset 2 and ending at offset 110. The resulting data range for pattern matches is from byte 3 through byte 126.
  - Multiple offsets must be set either equal to one another, or set beyond the boundaries of each other. For example, if **UDA1 Offset** starts at byte 2, the **UDA2 Offset** can only start either at byte 2 or at any point beginning with byte 18 (which would be the next 4-byte boundary after the 16-byte pattern used at **UDA1 Offset**).
  - Offsets are always frame-relative, not data-relative.
  - In many cases, you will be looking for patterns that do not start exactly on a four-byte boundary. To search in these position, you would set an offset at the nearest four-byte boundary and adjust the pattern and mask accordingly.

User-Defined Pattern Match Rules

Keep in mind the following rules when creating user-defined pattern matches:

- Offsets are specified in decimal; patterns and masks are specified in hexadecimal.
- All hexadecimal values must be fully defined, including leading zeroes. For example, to specify 0xff as a 16-byte value, you must enter 00000000-00000000-00000000-000000ff.
- User-defined pattern-match criteria are not allowed in egress port-filters (tool, hybrid, circuit, and inline network).
- You can use user-defined pattern matches as either standalone map rules or in tandem with the other available predefined criteria for map rules (for example, port numbers, IP addresses, VLAN IDs, and so on).

- You can use up to two separate user-defined pattern matches in a single map rule. When two user-defined pattern matches appear in the same map rule, they are joined with a logical AND. However, the two patterns cannot use the same offset.
- User-defined pattern matches are combined in map rules using the same logic described in How to Combine Rules and Rule Logic.
- Avoid using user-defined pattern matches to set map rules for elements that are available as predefined criteria (for example, IP addresses, MAC addresses, and so on).
- GigaVUE HC Series nodes accept a maximum of two offsets per line card. When both of the available offsets for a line card are in use with existing map rules, you will not be able to add a new rule with a different value for **UDAx Offset** until at least one of the **UDAx Offsets** is freed up from all existing map rules.
- On GigaVUE-HC1-Plus, GigaVUE-HCT, GigaVUE-TA25, and GigaVUE-TA25E, UDA1 supports a 12-byte data match of tagged traffic with an offset starting from 16 to 116. UDA1 does not support untagged traffic. UDA2 supports a 4-byte data match with an offset starting from 0 to 60.
- Due to limitations in the platform, map rules to match the data pattern on Inner-L3 and Inner-L4 qualifiers with UDA base does not match the Inner headers for the following types of encapsulation:
    - ERSPAN
    - LISP
    - L2GRE

However, the inner headers of ERSPAN, LISP, L2GRE can be made to match with outer L3 and Outer L4 UDA base by appropriately adjusting the offsets.

- When configuring a map rule with common UDA offset, you cannot combine source ports from GigaVUE-HC1-Plus or GigaVUE-HCT or GigaVUE-TA25 devices together with source ports from other devices in a single map. You must create separate maps for GigaVUE-HC1-Plus or GigaVUE-TA25 devices and other devices.

User-Defined Pattern Match Examples

In this example, a 3G carrier is monitoring the Gn interface between the SGSN and the GGSN in the mobile core network and wants to split traffic from different subscriber IP address ranges to different tool ports. However, because the subscriber IP addresses are tunneled using the GPRS Tunneling Protocol (GTP), standard IP address map rules will not work. The addresses are always at the same offsets, though, so we can construct UDA pattern match rules to match and distribute the traffic correctly.

For example, suppose we want to apply the following rules to all traffic seen on network port 1/5/x1:

- Send all traffic to and from the 10.218.0.0 IP address range inside the GTP tunnel to tool port 1/5/x4.

- Send all traffic to and from the 10.228.0.0 IP address range inside the GTP tunnel to tool port 1/5/x9.

Keep in mind that we also know the following about tunneled GTP traffic:

- The offset for source IP addresses inside the GTP tunnel is 62.
- The offset for destination IP addresses inside the GTP tunnel is 66.

The following example explains how to construct two maps that will distribute traffic using UDA pattern match rules.

Description	UI Step
Map #1 – GTP_Map218 Our first map will send traffic to and from the 10.218.0.0 IP address range inside the GTP tunnel to tool port 1/5/x4.	
Create a map with the alias **GTP_Map218**.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Enter GTP-MAP218 in the Alias field.
Specifies the map type and subtype.	**4.** Select Regular for Type. **5.** Select By Rule for Subtype.
Specify that this map will match packets arriving on network port 1/5/x1.	**6.** Select 1/5/x1 for Source.
Specify that packets matching this map will be sent to tool port 1/5/x4.	**7.** Select 1/5/x4 for Destination.
Next, add the map rules for our first address range – 10.218.0.0. This IP address translates to 0ada in hex. The first rule matches the 10.218.0.0 address at the source address offset of 62 in the GTP tunnel.	**8.** Click Add a Rule to add the first rule. **9.** Select Pass. **10.** Select UDA1 and set the values: Value: 0ada0000-00000000-00000000-00000000 Mask: ffff0000-00000000-00000000-00000000 Offset: 62
The second rule matches the same address range (10.218.0.0) but at the destination address offset of 66 in the GTP tunnel. Notice that we have still specified the offset as 62 and have simply masked out to the correct location of the destination address. This way, we have still only used one of the two possible offsets in place for the GigaVUE HC Series node at any one time.	**11.** Click Add a Rule to add the second rule. **12.** Select UDA1 and set the values: o Value: 00000000-0ada0000-00000000-00000000 o Mask: 00000000-ffff0000-00000000-00000000 o Offset: 62
Save the map.	**13.** Click Save.
Map #2 – GTP_Map228	

Description	UI Step
Our second map will send traffic to and from the 10.228.0.0 IP address range inside the GTP tunnel to tool port 1/5/x9.	
Create a map with the alias **GTP_Map228**.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Enter GTP-MAP228 in the Alias field.
Specifies the map type and subtype.	**4.** Select Regular for Type. **5.** Select By Rule for Subtype.
Specify that this map will match packets arriving on network port 1/5/x1.	**6.** Select 1/5/x1 for Source.
Specify that packets matching this map will be sent to tool port 1/5/x9.	**7.** Select 1/5/x9 for Destination.
Now, create rules for the second address range – 10.228.0.0 (0ae4 in hex). As with the first range, create separate rules for the source and destination offsets inside the GTP tunnel. This address range is being sent to 1/1/x4.	**8.** Click Add a Rule to add the first rule. **9.** Select Pass. **10.** Select UDA1 and set the values: Value: 0ae40000-00000000-00000000-00000000 Mask: ffff0000-00000000-00000000-00000000 Offset: 62
Here is the companion rule for the destination address offset of 66.	**11.** Click Add a Rule to add the second rule. **12.** Select UDA1 and set the values: o Value: 00000000-0ae40000-00000000-00000000 o Mask: 00000000-ffff0000-00000000-00000000 o Offset: 62
Save the map.	**13.** Click Save.

User-Defined Pattern Match on GigaVUE-TA400, and GigaVUE-TA400E

The following table outlines the attributes, base offsets, offset ranges, and data widths for user-defined pattern matching (UDA) on the GigaVUE-TA400, and GigaVUE-TA400E chassis:

Attribute	Base Offset	Offset Range	Data Width
uda1	Starting of Outer Ethertype header	0-160	8 bytes
uda2	Starting of Outer L3 header	0-160	8 bytes
inner-uda1	Starting of inner L3 header	0-120	8 bytes
inner-uda2	Starting of inner L3 header	0-120	8 bytes

The data to be matched must be specified with trailing 8-byte zeros. For example to specify 0x54206874 as a match value, enter 00000000-54206874-00000000-00000000.

Flow Map Rule Combinations

The following table shows the unsupported UDA combinations:

Map1	Map2	Support Status
uda2	inner-uda2	Not Supported
inner-uda1	uda1	Not Supported
uda1	inner-uda1	Not Supported
inner-uda2	uda2	Not Supported

NOTE: When working with UDA configurations, the combination of (uda1 + inner-uda1) and (uda2 + inner-uda2) is not supported at the chassis level. You may upgrade devices with these unsupported combinations, but further editing of the rules will be restricted due to potential traffic impacts, as inner-uda takes precedence. You are allowed to delete the mapping and rules for these unsupported configurations.

Inner Header and MPLS Header Filtering

GigaVUE-OS now supports the following capabilities for GigaVUE-TA400, and GigaVUE-TA400E devices:

Inner Header Attributes

GigaVUE-OS can filter based on inner payload packet fields for tunneled traffic. This enables the extraction of flows based on inner header packet attributes without stripping the encapsulated header. GigaVUE-OS can filter traffic based on the following inner header packet attributes:

- **Inner Ethertype**: Filter based on the Ethertype of the inner packet.
- **Inner IP Version**: Filter based on the IP version (IPv4 or IPv6) of the inner packet.
- **Inner IP Fragmentation Bits**: Filter based on the fragmentation bits in the inner IP header.
- **Inner Source IP**: Filter based on the source IPv4 address of the inner packet.
- **Inner Destination IP**: Filter based on the destination IPv4 address of the inner packet.
- **Inner IP Protocol**: Filter based on the protocol (e.g., TCP, UDP, ICMP) in the inner IP header.
- **Inner Source IPv6**: Filter based on the source IPv6 address of the inner packet.
- **Inner Destination IPv6**: Filter based on the destination IPv6 address of the inner packet.

- **Inner L4 Source Port**: Filter based on the Layer 4 (TCP/UDP) source port of the inner packet.
- **Inner L4 Destination Port**: Filter based on the Layer 4 (TCP/UDP) destination port of the inner packet.
- **Inner UDA1**: Filter based on the first user-defined attribute (UDA) in the inner packet.
- **Inner UDA2**: Filter based on the second user-defined attribute (UDA) in the inner packet.

For filtering based on the DSCP value in the payload IPv4 or IPv6 headers, the inner UDA attributes can be used. These attributes take the offset from the start of the inner L3 header, ensuring accurate filtering.

Example:

The map rule shown below matches the Inner Header packets:

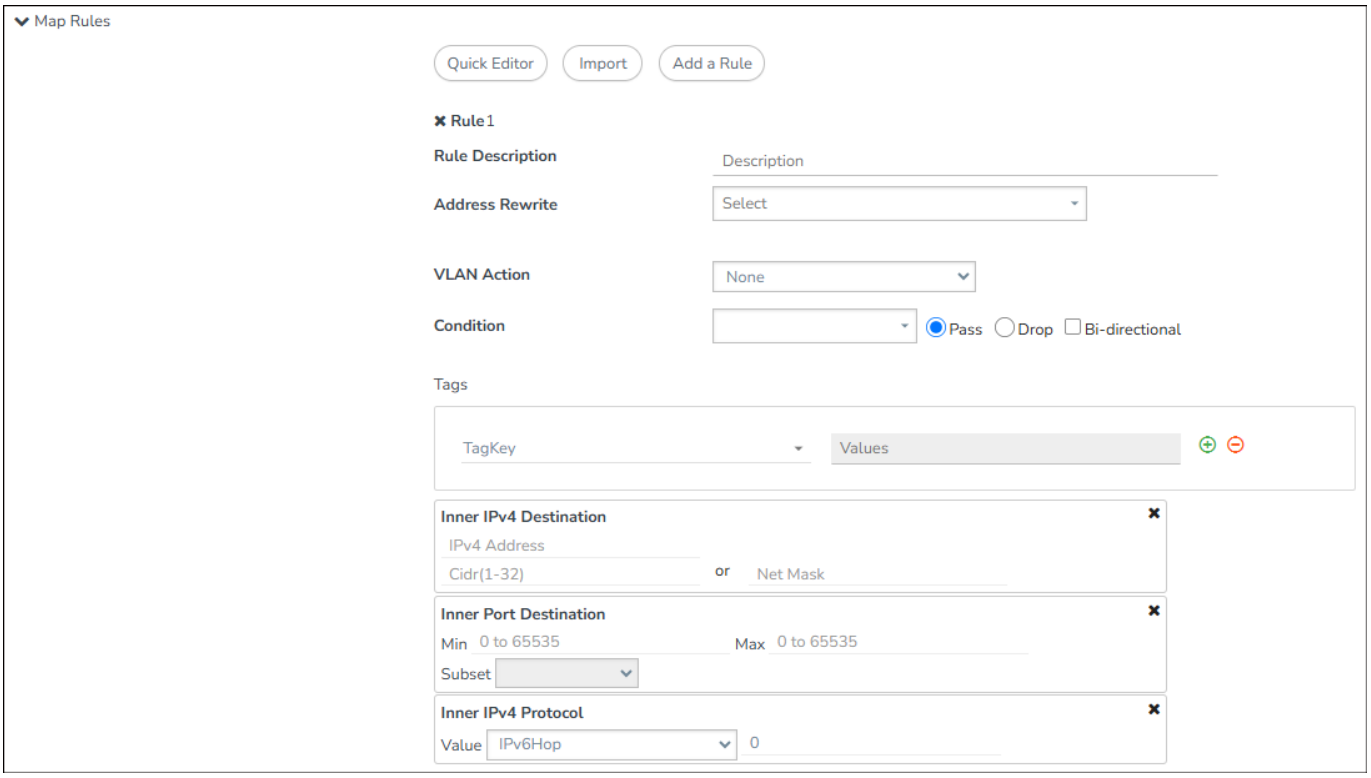


Figure 6 *Map Rule with Inner IPv4 Destination, Inner Port Destination, and Inner Ipv4 Protocol*

MPLS Header Attributes

GigaVUE-OS also supports filtering based on MPLS header attributes, allowing traffic to be forwarded without stripping the MPLS header. The following attributes can be filtered:

- **MPLS Label ID**: Filter based on the MPLS Label ID.
  - Position any: Specifies all levels (1-7) of the MPLS header in the packet.

- Position 1-7: Specifies the depth of the MPLS header in the packet (1 = outermost label, 7 = innermost label).

- **MPLS Label EXP (Experimental)**: Filter based on the EXP bits in the MPLS header.
  - Position any: Specifies all levels (1-7) of the MPLS header in the packet.
  - Position 1-7: Specifies the depth of the MPLS header in the packet (1 = outermost label, 7 = innermost label).
- **MPLS Label BOS (Bottom of Stack)**: Filter based on the BOS bit in the MPLS header.
  - Position any: Specifies all levels (1-7) of the MPLS header in the packet.
  - Position 1-7: Specifies the depth of the MPLS header in the packet (1 = outermost label, 7 = innermost label).
- **MPLS Label TTL (Time to Live)**: Filter based on the TTL value in the MPLS header.
  - Position any: Specifies all levels (1-7) of the MPLS header in the packet.
  - Position 1-7: Specifies the depth of the MPLS header in the packet (1 = outermost label, 7 = innermost label).

Traffic can be filtered based on these attributes and sent to the desired tools without modifying or stripping the MPLS header.

Example:

The map rule shown below matches the MPLS Header packets:

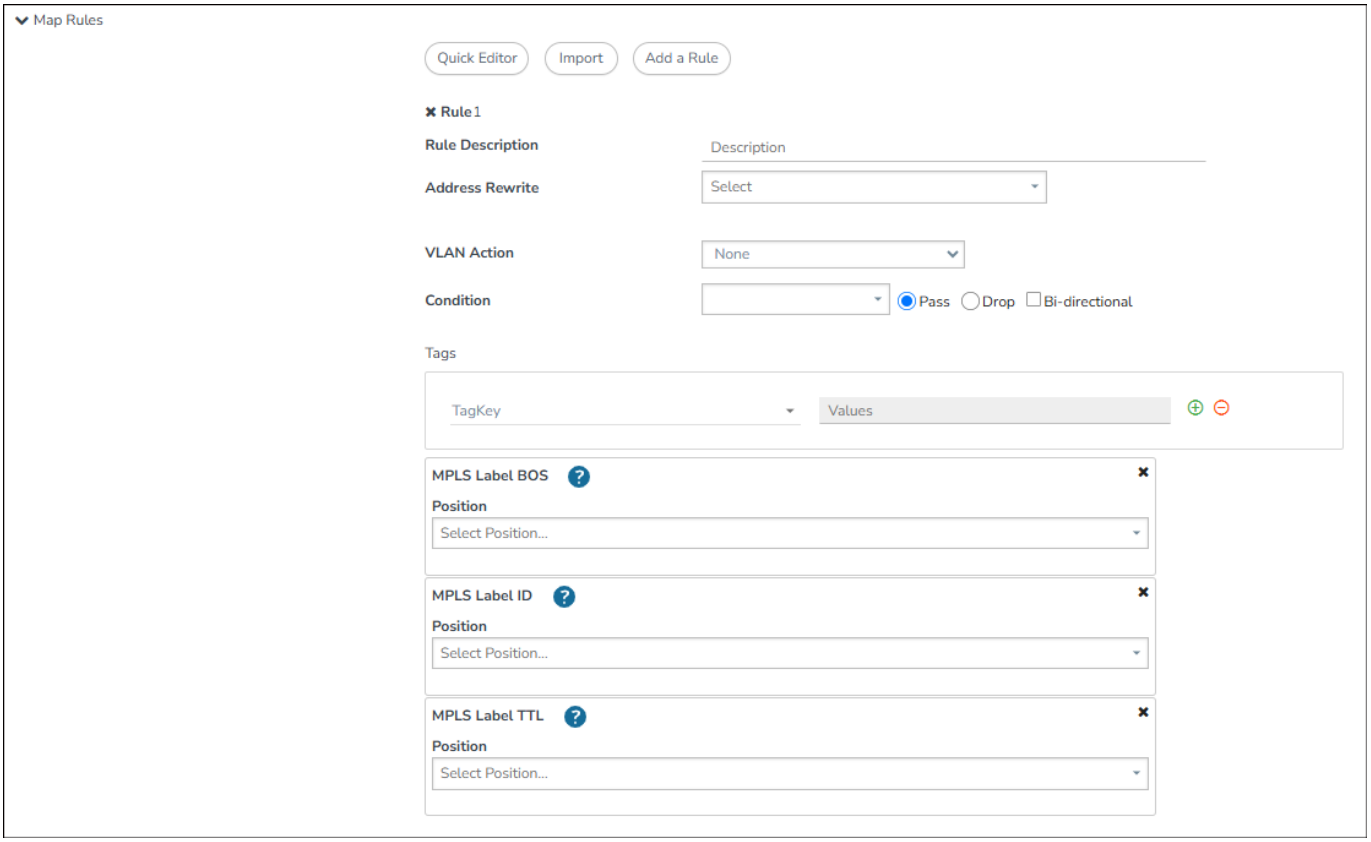


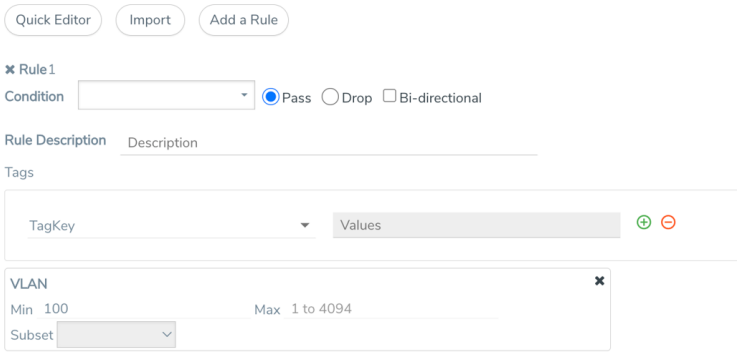
Figure 7 *Map Rule with MPLS Label BOS, MPLS Label ID and MPLS Label TTL*

To configure an Inner Header qualifier and MPLS Header qualifier using GigaVUE-OS CLI, refer to the Configure Inner Header Qualifier and MPLS Header Qualifier chapter in the GigaVUE-OS CLI Reference Guide.

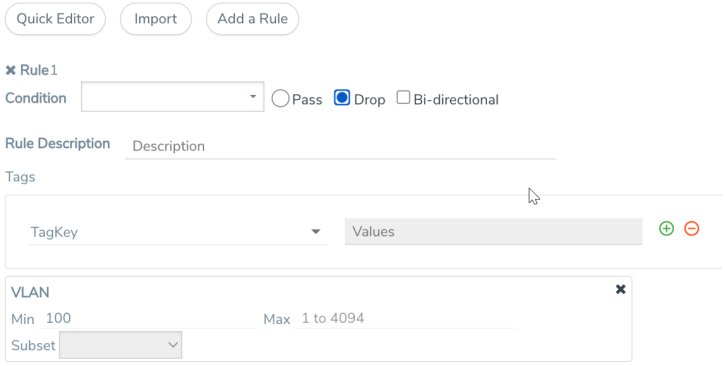
How to Handle Q-in-Q Packets in Maps

The Q-in-Q packets in maps are handled as follows:

- For traffic that matches the map pass rule shown in the following figure, Q-in-Q packets of TPID ethertype 0x8100, 0x88A8, and 0x9100 are passed.



- For traffic that matches the map drop rule shown in the following figure, Q-in-Q packets of TPID ethertype 0x8100, 0x88A8, and 0x9100 are dropped.



You do not specify TPID EtherTypes 0x8100, 0x88A8, and 0x9100 explicitly in a rule. If you specify these values in the **Value** field an **EtherType** rule, the map is blocked and one of the following error messages is displayed:

Invalid ethertype : '0x8100'. Please use attribute 'vlan' instead.

Invalid ethertype : '0x88A8'. Please use attribute 'vlan' instead.

Invalid ethertype : '0x9100'. Please use attribute 'vlan' instead.

NOTE: The **Value** field accepts values with out the leading 0x only.

In summary, for single-tagged (0x8100) or double-tagged (0x88A8 and 0x9100) VLAN packets, you only configure the VLAN as the matching criteria, not the ethertype.

For handling of priority tagged packets, refer to Priority Tagged Packets.

For filtering of Q-in-Q packets on inner VLAN tag, refer to Flow Mapping® on Inner VLAN Tags.

Comparison of Q-in-Q Tagging

The following table details the various combinations and corresponding behaviors depending on the packet content:

Packet Content	Rule: pass vlan 100	Rule: pass ethertype 0x0800	Rule: pass vlan 100 ethertype 0x0800
No tags, ethertype 0800	drop	pass	drop
One tag: TPID 8100, VID 100, ethertype 0800	pass	pass	pass
One tag: TPID 9100, VID 100, ethertype 0800	pass	pass	pass
One tag: TPID 88a8, VID 100, ethertype 0800	pass	pass	pass
Two tags: outer TPID 8100 VID 100, inner TPID 8100 VID 200, ethertype 0800	pass	pass	pass
Two tags: outer TPID 9100 VID 100, inner TPID 8100 VID 200, ethertype 0800	pass	pass	pass
Two tags: outer TPID 88a8 VID 100, inner TPID 8100 VID 200, ethertype 0800	pass	pass	pass
Two tags: outer TPID 8100 VID 200, inner TPID 8100 VID 100, ethertype 0800	drop	pass	drop
Two tags: outer TPID 8100 VID 200, inner TPID 88a8 VID 100, ethertype 0800	drop	drop	drop
Two tags: outer TPID 88a8 VID 200, inner TPID 8100 VID 100, ethertype 0800	drop	pass	drop

Packet Content		Rule: pass vlan 100	Rule: pass ethertype 0x0800	Rule: pass vlan 100 ethertype 0x0800
Two tags: outer TPID 8100 VID 100, 88a8 VID 100, ethertype 0800	inner TPID	pass	drop	drop
Two tags: outer TPID 8100 VID 100, 9100 VID 100, ethertype 0800	inner TPID	pass	drop	drop

Priority Tagged Packets

Priority tagged packets are handled by the GigaVUE node. These packets have a user priority of 0 to 7 in the packet. Single tagged packets or double tagged packets with a VLAN ID of zero or a non-zero value will be sent accordingly to the tool ports.

Flow Mapping® on Inner VLAN Tags

Flow mapping on inner VLAN tags is supported for filtering on Q-in-Q traffic.

- For packets that have both an inner and an outer VLAN tag, the outer tag is detected when the ethertype is 0x8100, 0x88A8, or 0x9100. The inner tag is detected only when the ethertype is 0x8100.
- If the inner VLAN tag ethertype is not 0x8100, then further encapsulations are not detected.

To specify an inner VLAN tag, add a new map rule (pass or drop) of type Inner VLAN.

1. Add a new map rule (pass or drop) of type Inner VLAN.
2. Select a VLAN (Min) or a range of VLANs (Min and Max). Subset, even or odd, is optional.

The inner VLAN range is supported with any other qualifier with a range, such as VLAN or portsrc.

NOTE: There is no filtering after the two VLAN tags (inner and outer).

Filtering on inner VLAN uses application filter resources. To track resource usage, go to **Chassis > Quick Port Editor** for a particular box ID, card and slot.

Each map rule uses a number of entries. A single inner VLAN uses one entry per map rule. A range of inner VLANs uses two or more entries per map rule. For the same map source, identical inner VLAN or inner VLAN range spread across different rules will consume the same map rule resources.

A maximum of 454 application filter resource entries is available if no other application filters are using resources. The number of entries in the output of **Application Filter Resources** might be impacted by the other applications listed, such as GSD or Discovery.

The application filter resources are as follows:

- GSD—for GigaSMART tunnels
- Map Src—for network port source local to the node or slot (one entry per unique network port source). Note that 50 is always reserved per node or slot.
- Map Rule—for each inner VLAN rule
- Discovery—for LLDP/CDP

The following GigaVUE nodes have a maximum limit of 454 entries (the limit of 504 minus the 50 reserved):

- GigaVUE TA Series—per node
- GigaVUE-HC1—per node
- GigaVUE-HC3—per slot

Inner VLAN Limitation

Overlapped inner VLAN range is not supported within a map or set of maps that has the same network source. An identical VLAN range (and values) is supported.

For example, the following two rules are not supported because the inner VLAN range overlaps:

- Rule1: rule add pass inner-vlan 100 portsrc 1000
- Rule2: rule add pass inner-vlan 100..110 portsrc 1100

To overcome this, specify the rules as follows:

- Rule1: rule add pass inner-vlan 100 portsrc 1000
- Rule2: rule add pass inner-vlan 100 portsrc 1100
- Rule3: rule add pass inner-vlan 101..110 portsrc 1100

NOTE: You cannot use map rule editing to change an existing inner VLAN range to a range that overlaps with the original range. To edit an inner VLAN range, delete the rule and create a new rule with the new range.

Work with Map-Passalls and Port Mirroring in GigaVUE-FM

In addition to regular maps, GigaVUE-FM also makes it possible to create map-passall maps and configure tool-mirror ports for packet distribution.

A map-passall map lets you send all packets on a network port one or more tool ports or tool GigaStream on the same node or between the nodes in a cluster, irrespective of the packet distribution already in place for the ports.

Tool-mirror ports let you configure all a pass-all between two tool port or tool port and a too GigaStream on the same node, irrespective of the packet distribution already in place for the ports.

These map-passall maps and tool mirror ports are particularly useful in the following situations:

- Redirecting all traffic to IDS monitors regardless of any map rules applied to network ports.
- Temporary troubleshooting situations where you want to see all traffic on a port without disturbing any of the maps already in place for the port.

This section includes the following topics:

- Syntax for Maps-passalls and Port Mirroring
- Rules for Map-Passalls and Port Mirroring

Syntax for Maps-passalls and Port Mirroring

Refer to the following sections for details on map-passalls and port mirrors:

- About Map-passall Maps
- Managing Ports

Rules for Map-Passalls and Port Mirroring

Keep in mind the following rules for the map-passalls and tool mirrors:

- You can set up a map passall from:
  - Network port(s) to tool port(s) on the same node.
  - Network port(s) to one or more GigaStream.
- You can set up a tool mirror from:
  - Tool port to tool port(s) on the same node.

- Tool port to GigaStream(s) configured with the advanced-hash algorithm on the same node.

NOTE: The destination for a map-passall or tool-mirror can be a tool port, a hybrid port, a circuit port, a tunnel, a tool GigaStream, or a hybrid GigaStream.

- You cannot set up a map-passall or tool mirror from network port to network port. To be able to create such functionality, refer to Port Access and Map Sharing.
- A map-passall can cross line cards or modules – they can start on one line card/module and end on another in the same node. Also, they can cross nodes in a cluster.
- A tool mirror can cross line cards or modules – they can start on one line card or module and end on another in the same node. They cannot, however, cross nodes in a cluster.
- Tool mirrors are not allowed from tool GigaStream to tool port.
- Tool mirrors are not supported on tool ports with copper SFPs installed or on 100Gb ports with CFP2 transceivers.
- A map-passall cannot be used with a port that is part of a port-pair.

View and Delete Map-passalls

Map-passalls are created by selecting **Subtype Pass All** for a **Type Regular** map. You can view the map by clicking on the map alias on the Maps page to open the Quick View for the map. To delete a pass-all map, select the map on the Map page and click **Delete**.

Port Access and Map Sharing

There are two ways to define a user's access to ports and maps:

- Port-based access levels
- Map sharing

Both methods assign permissions to user roles, as defined by the user groups, rather than specific user accounts.

Port-based Access Levels

Users are assigned roles based on their user group. Each user group is given permission to specific ports on the node. There are four port-based permission levels:

- Level 1—Can view the port but cannot make any changes to port settings or maps. When applied to a network port, can view maps attached to the network port. This level is used for users who only need to monitor the activities of the port.

- Level 2—Can use the port for maps, create tool-mirror to/from port, and change egress port filters. Can configure port-lock, lock-share, and all traffic objects except port-pair. Also includes all Level 1 permissions.
- Level 3—Can configure port parameters (such as administrative status of the port, speed, duplex, and autonegotiation, as well as create port pairs. Also includes all Level 2 and Level 1 permissions.
- Level 4—Can change the port type. Also includes all Level 3, 2, and 1 permissions.

Table 2: Port-based Permission Levels summarized the permissions for each of the levels.

Table 2: Port-based Permission Levels

Permissions	Level 1	Level 2	Level 3	Level 4	Admin
View port	✓	✓	✓	✓	✓
View maps attached to network port	✓	✓	✓	✓	✓
Create/edit map attached to port	×	✓	✓	✓	✓
Create tool-mirror to/from port	×	✓	✓	✓	✓
Change egress filters	×	✓	✓	✓	✓
Edit port parameters	×	×	✓	✓	✓
Create port pairs	×	×	✓	✓	✓
Change port type	×	×	×	ü	✓

How to share Maps

Maps can be shared with one or more user groups. When sharing a map, the map owner or Admin designates which user groups have which permissions. There are four map-sharing permission levels:

- **Read Only** – Can view the map but cannot make any changes.
- **Listen** – Can add or remove tool ports they own*. This is equivalent to “subscribing” to a map.
- **Read/Write** – Can delete and edit the map, can remove any network ports, can add network ports they own*, and can add or remove tool ports they own*.
- **Read/Write/Owne**r – Can perform all the Read/Write functions and assign map sharing permission levels.

Requires Level 2 or Level 3 access, based on User Group membership.

Table 3: Permission Levels for Map Sharing. summarizes the permission levels for map sharing.

Table 3: Permission Levels for Map Sharing.

Permissions	Read Only	Listen	Read/Write	Read/Write/Owner
View map	✓	✓	✓	✓
Add tool port*	×	✓	✓	✓
Remove tool port	×	✓*	✓	✓
Remove network port	×	×	✓	✓
Add network port*	×	×	✓	✓
Delete/edit map	×	×	✓	✓
Share map	×	×	×	✓

Only applies to ports to which the user has Level 2 or Level 3 access.

NOTE: In Table 3: Permission Levels for Map Sharing., tool port includes ports of type tool and inline-tool. Network port includes ports of type network and inline-network.

The admin user can also assign map sharing permissions.

Users with Level 1 (or greater) access to a given network port can also view, but not edit, maps associated with that network port. This is independent of the map sharing permissions.

Map sharing permissions override and supersede role based access controls. Thus, a user group can be assigned Read/Write access to map even if they do not have any access rights to any of the associated network or tool ports. However, adding tool ports to a map or removing network or tool ports from a map requires Level 2 or Level 3 permissions, as defined by the user group, for the ports to be added or removed.

Map Examples

This section provides the following map examples:

- Example: How to Create a Simple Map
- Example: How to Handle Overlaps when Sending VLANs and Subnets to Different Tools

Example: How to Create a Simple Map

In this example, a few simple maps are illustrated to show how to create and display the packet distribution in place on the node.

When you set up flow maps from the perspective of your tools, start by asking yourself which traffic you would like a particular tool to see. Then, select the necessary traffic from network ports.

For example, the scenario in this example is as follows:

- An application performance management tool is connected to **tool port 2/4/x6** that focuses on traffic from **VLANs 100..199** on **network ports 2/2/x10** and **2/2/x12**.
- An application performance Management tool connected to **tool port 2/4/x18** that focuses on traffic from **VLANs 200..299** on **network ports 2/2/x10** and **2/2/x12.**

Consider a GigaVUE-HC3 device with the following configurations:

- BoxID of the GigaVUE-HC3 device is set as 2.
- The network ports are set on the second blade in the node.
- The tool ports are set on the fourth blade in the node.

Therefore, the ports in this scenario are represented as follows:

- The network port IDs are set as 2/2/x2, 2/2/x10, and 2/2/x12.
- The tool port IDs are set as 2/4/x6, 2/4/x8, 2/2/x20, and 2/4/x24

The maps for this scenario the map types are set as follows:

- The maps with the VLAN rules specified have the subtype set as **By Rule**.
- The shared collector maps are set as subtype **Collector** and no rules added.
- The passall maps are defined as subtype **Pass All** and no rules added.

For details about the different map types, refer to Matrix of Map Types and Map Subtypes.

Except for the following map types, the map type defaults to **Regular**

- Only maps with Inline Bypass solutions can be set as type **Inline**.
- Only maps that have GigaSMART operations defined for second level maps can be set as **First Level** or **Second Level** maps.

The following are the steps to create the simple maps as shown in Map Examples:

1. Check the port types for the each of the ports that the maps will use and set them if necessary.

   To set the port types, select Port > Ports > All Ports and use the Port Type Editor to set the port types. For more information about port types, refer to

2. Create a **Regular** map with a **By Rule** subtype to pass the VLAN100 traffic as shown in Figure 8Map for VLAN 100..199.

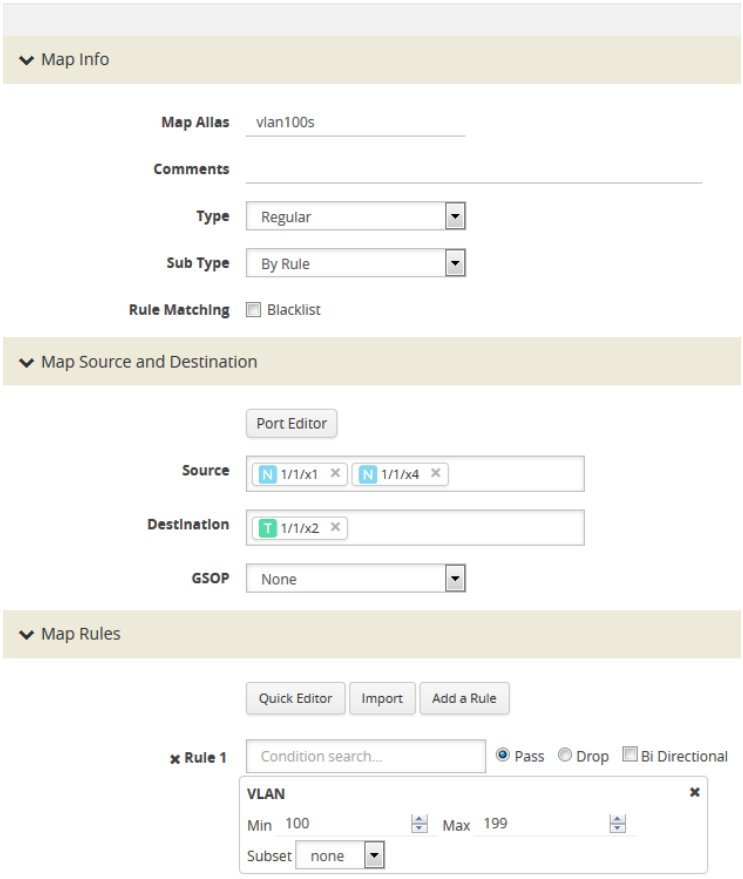


Figure 8 *Map for VLAN 100..199*

3. Create a **Regular** map with a **By Rule** subtype to pass the VLAN200 traffic.
4. Create a map with type **Regular** and subtype **Collector** to create a shared collector map. For more information about shared collector maps, refer to About Shared Collectors.
5. Create a map with type **Regular** and subtype **Pass All** to create the passall map.

Example: How to Handle Overlaps when Sending VLANs and Subnets to Different Tools

Figure 9Sending Subnets and VLANs to Different Ports shows how to use map priority when handling packets matching criteria in multiple maps. In this example, we want to achieve the following results:

- Send packets on the 172.16.0.0 subnet to 1/2/x1
- Send packets on the 172.17.0.0 subnet to 1/2/x2
- Send packets on VLAN 100 to 1/2/x3

The trick is in how to handle packets on either 172.16.0.0 or 172.17.0.0 **and** VLAN 100. In this example, we use map priority to ensure that packets such as this are sent to both of their desired destinations.

Notice that the first two maps configured in Figure 9Sending Subnets and VLANs to Different Ports are set up to handle this situation. For example, **map1** has a pass rule that accepts packets on 172.16.0.0 and VLAN 100. It sends matching packets to both 1/2/x1 (the destination we wanted for the 172.16 subnet) and 1/2/x3 (the destination we wanted for VLAN 100). Because this map was entered before **map3**, it has higher priority, ensuring the packet goes to both 1/2/x1 and 1/2/x3 and not just the 1/2/x3 destination specified by **map3**.

The same principle is applied in **map2** for packets on 172.17.0.0 and VLAN 100.

NOTE: If we did not observe the order of map entry shown in Figure 9Sending Subnets and VLANs to Different Ports, we could always adjust the priority as needed using the instructions in Example: How to Create a Simple Map.

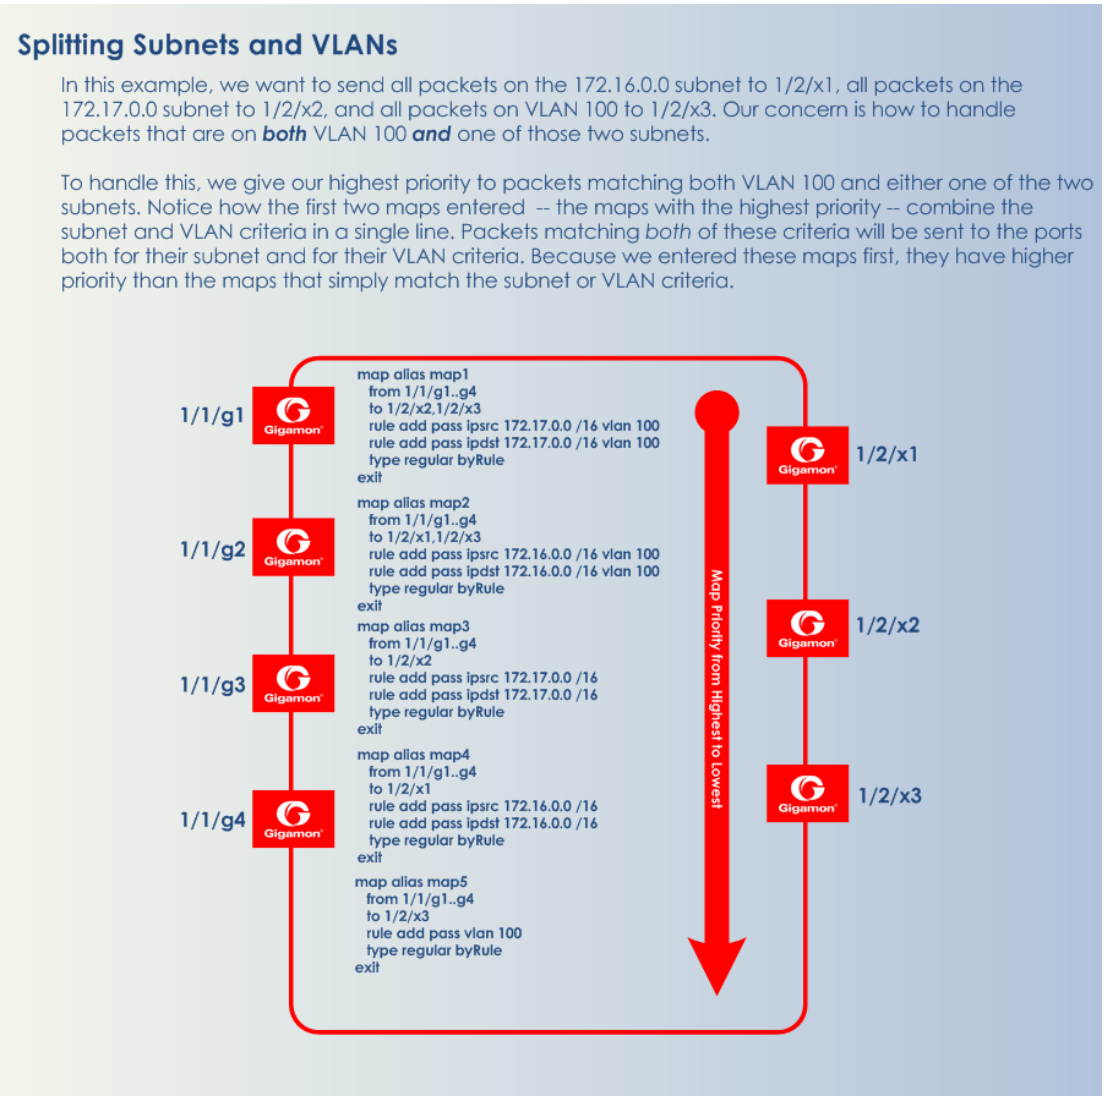


Figure 9 *Sending Subnets and VLANs to Different Ports*

Manage Maps

This section provides a description of the Maps pages in the GigaVUE-FM UI. It covers the following topics:

- Map Views
- Manage Maps
- Map Templates
- Manage Map Rule Resources
- Filter Templates
- Review Map Statistics with Map Rule Counters

Map Views

The Maps page displays the maps created using the CLI, or GigaVUE-FM.

To view the map, do the following:

1. Go to **Inventory > Physical > Nodes**.
2. Click the Cluster ID link, go to **Traffic > Maps > Maps**.

NOTE: Starting in software version 5.5.01, any change in the map health status is indicated immediately in the Maps page.

The maps can be displayed in List or Map Topology views.

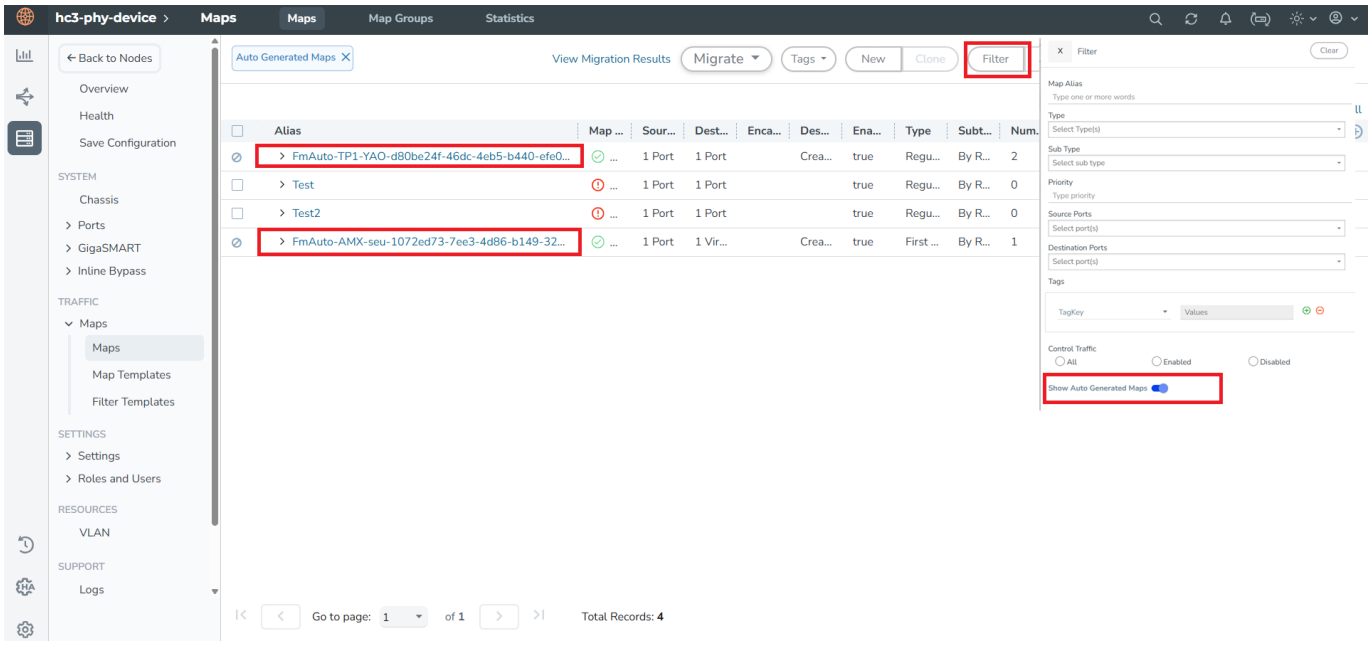
List View

The List View is the default view of the maps when the Maps page is opened. This view shows the basic information about each map:

- Alias
- Type
- Subtype
- Source and destination ports
- Control Traffic: Displays if the control traffic is enabled or disabled for a particular map. This is applicable only for First Level, By Rule maps.

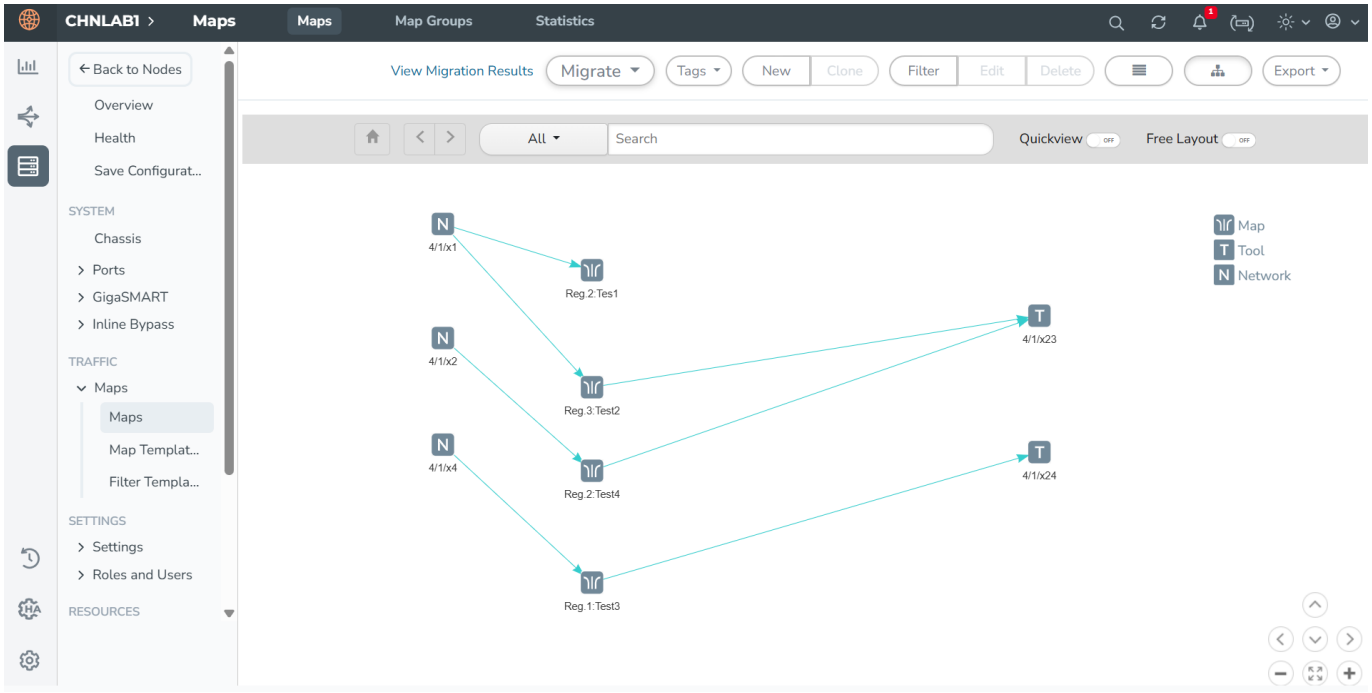
Map Topology View

Click the icon. The Topology View of the maps appears. Each map in the Maps page gets displayed as Topologies. The Map Topology presents a graphical representation of all maps available from the device, including both system-defined and user-created maps. The system-defined maps are displayed when you click the **filter**. A pop-up appears, and then you enable **Show Auto-Generated Maps**.



You can interact with the topology to efficiently search for ports, navigate through the layout, and inspect individual elements in detail.

If a particular map is selected, only that map is displayed with its details, instead of showing all maps. You can use the back and forward arrows to navigate to the previous and next maps, respectively. The **Home** option allows you to return to the main maps page.



GigaVUE-FM provides multiple options for viewing map topology, including QuickView and Free Layout.

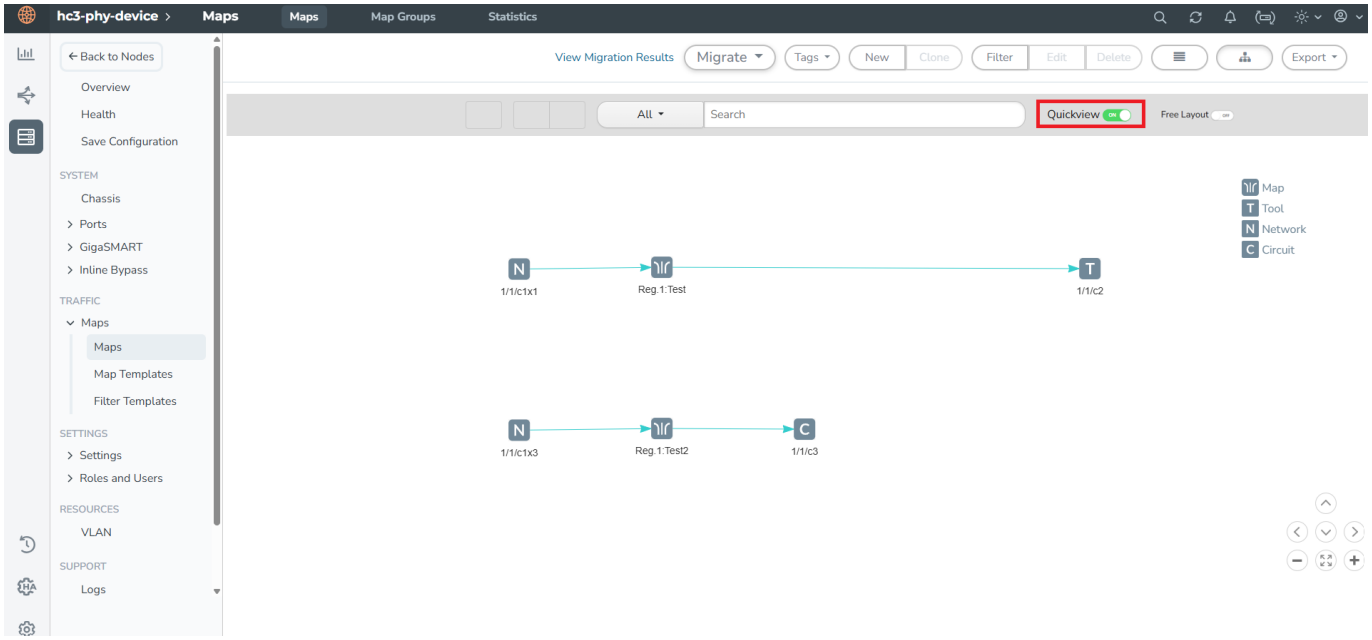
Quickview

If QuickView is enabled and you select a port, its details are displayed in a pane. In QuickView, you can move ports up and down only.

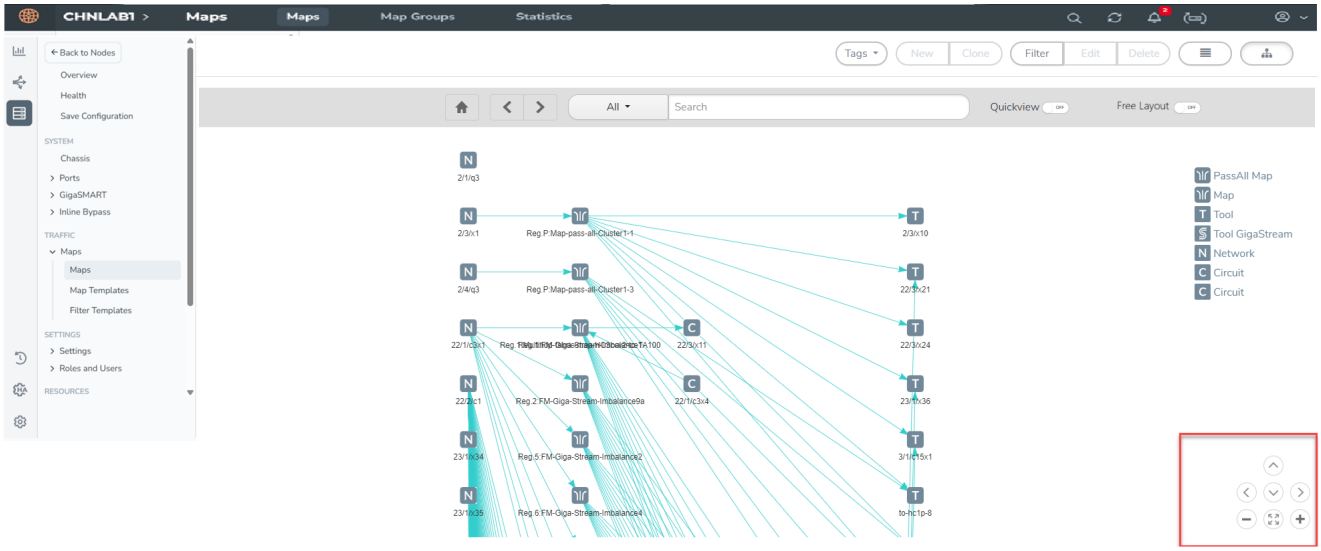
- When you click on any node in the Topology View, a Quick View panel opens, displaying details for that element.
- The Quick View panel opens, showing details such as port details, map configuration, and associated connections.
- Use the zoom and scroll buttons at the bottom of the page to zoom in and out.
- Click on a map to view the connections.
- Use the Search field to enter the name or type of map you want to locate. The topology view filters and displays only matching maps.

Without enabling QuickView, selecting a map display only that map in the topology view.

The topology section provides access to all maps created on the device by the user. You can filter for a specific type of map to view its details. Only the searched map will be displayed. You can search using port alias, port ID, etc. For example, if you search for the x23 port, only the map containing that port will be shown.



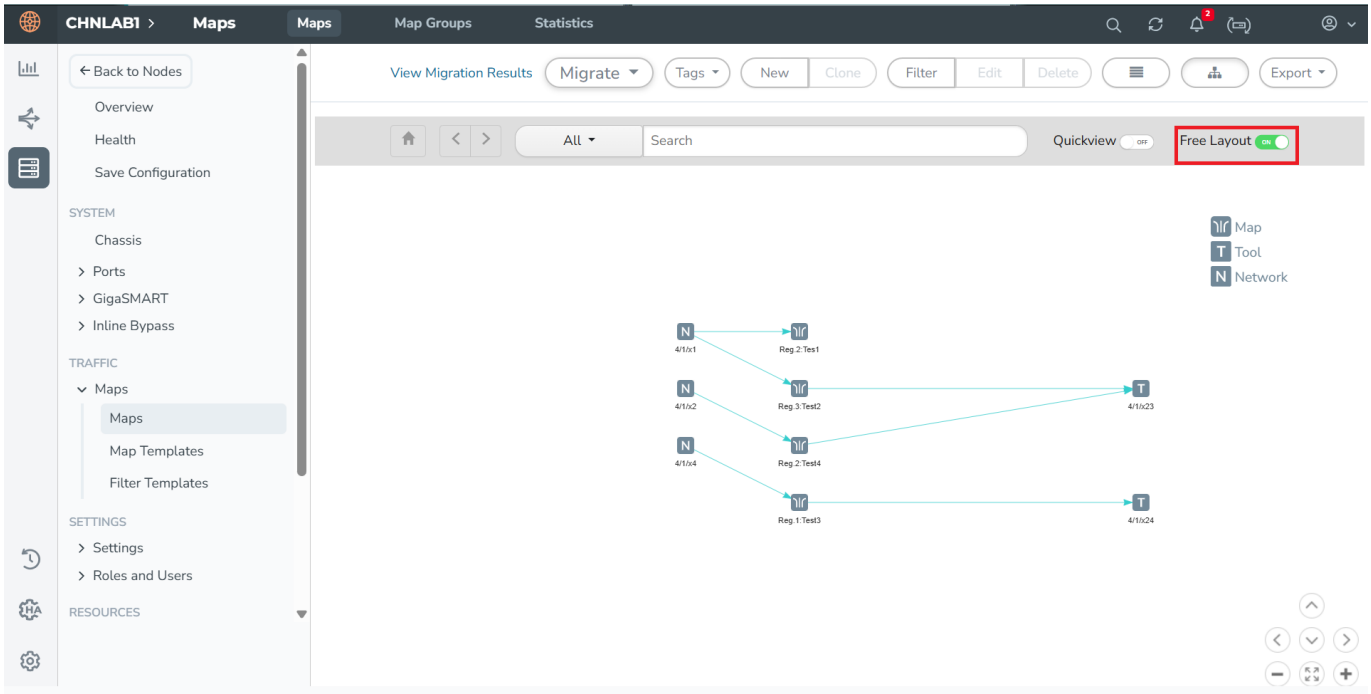
NOTE: Searching narrows the display to matching items only, making it easier to locate specific maps in large deployments.



Free Layout

NOTE: Starting from 6.14, a new **Free Layout** option in the topology view lets you freely rearrange these canvas elements to match your preferred visual layout. Your layout is stored as a user preference and is automatically restored when you return to the Maps topology view.

1. Click on the icon. The Topology View of the maps appear. Each map in the Maps page gets displayed as Topologies.
2. In the topology view, turn on the Free Layout option (toggle). When Free Layout is enabled, the elements on the topology canvas become freely movable.
3. Click and drag map elements (such as network, map layers, tool ports, virtual ports, and hybrid ports) to the desired positions.
4. Arrange the elements in any logical pattern you prefer. For example, vertically from network (source) to tool port (destination) instead of the default horizontal layout.
5. Use the zoom and scroll buttons at the bottom of the page to zoom in and out.
6. Click on a map to view the connections.



Reset to Default Layer

You can return to the default topology layout, do the following:

1. In the Map Topology view, disable the **Free Layout**.
2. The topology resets to the default layout, clearing any applied filters or navigation state.

Manage Maps

This section provides the basic steps for doing the following tasks:

- Create a new map
- Clone Map
- Edit Maps
- Delete Maps
- Create Map Groups

It also includes information about the following:

- Description to Map Rules
- How to Use the Quick Editor for Pass and Drop Rules

Create a new map

Prerequisites:

1. Check the status of the nodes and ports that you plan to use with the map. For information about how to check the status of the nodes and ports, refer to Status of Line Cards/Nodes and Ports.
2. From the device view, go to **Maps > Maps** to open the Maps page.
3. Click **New**.

Create a new map – Field References

The following table lists and describes the fields you must complete to configure new map.

Field	Description
Map Info	
Map Alias	A map alias specifies the name of the map. The alias must be unique and can contain up to 128 alphanumeric characters. Aliases are case-sensitive. Most special characters are supported in map aliases. However, map aliases that are only one period (.) or two periods (..) should not be created. These aliases cannot be accessed for editing. Use an alias that helps identify the task and destination. For example netflix_traffic_to_wireshark.
Description	(Optional) The description for the map. When adding description, consider the following: • Use up to 128 characters, including spaces. • Enclose the description in quotation marks, if the description is longer than one word.

Field	Description
	• To include double quotation marks (") inside the quotation marks, precede it with a backslash (\). See also Description to Map Rules.
Enable	Checkbox to enable or disable the map.
Type	Select the map **type**. The map type can be **Regular, First Level, Second Level, Inline, Inline First Level, Inline Second Level, or Transit Level**. For detailed information about the types of maps, refer to Map Types
Subtype	Select the **subtype**. The map subtype can be **By Rule, Pass All, or Collector**. For detailed information about Pass All, refer to About Map-passall Maps. For detailed information about Collector, refer to About Shared Collectors.
No Rule Matching	Checkbox for excluding rule-based matching. Enable the Pass Traffic checkbox if no rules are matching.
Control Traffic	Enable the Control Traffic checkbox to pass the GTP control traffic (GTP-c) to all GigaSMART engines in a GTP engine group. A GTP engine group has multiple GigaSMART engine port members. **Note**: The Control Traffic checkbox is applicable only for GTP and is displayed only if the map type is configured as First level, and the map sub type is configured as By Rule.
Map Source and Destination Map	
Source	From the **Source** drop-down list, select the required source ports for the map. To create a port list, click **Port Editor**. **Note**: You can add a maximum of 324 ports in the Source drop-down list, if the ports are not attached to a GigaStream. **Note**: For details about port types that are supported for the different types of maps, refer to Port Lists.
Destination	From the **Destination** drop-down list, select the required destination ports for the map. To create a port list, click **Port Editor**.
Encapsulation Tunnel	If you have selected a circuit port in the **Destination** drop-down list, select the required circuit tunnel from the **Encapsulation Tunnel** drop-down list to encapsulate the traffic. If the map is used to redirect the decapsulated traffic to the required tool ports, ensure that you select the IP interface in the **Source** drop-down list.

Field	Description
	You must have attached the IP interface to the VXLAN or L2GRE tunnel. For details about VXLAN or L2GRE tunnels, refer to About Virtual Extensible LAN (VXLAN) Tunnels and About Layer 2 Generic Routing Encapsulation (L2GRE) Tunnels. From the **Destination** drop-down list, select the required tool ports.
GigaSMART Operations (GSOP)	If the map will use a GigaSMART operation, select the operation from the **GigaSMART Operations (GSOP)** drop-down list.
Map Configuration and Rules	
Address Rewrite	(Optional) Enable the address rewrite checkbox and select either MAC Address or IPv4 Address. For detailed information about MAC Address Rewrite, Refer to MAC Address Rewrite. For detailed information about IP Address Rewrite, Refer to IP Address Rewrite.
VLAN Action	Select **Add** option from the VLAN Action field to add a new VLAN tag to the outgoing traffic. For detailed information about VLAN Manipulation, Refer to VLAN Manipulation.
Map Rules	To add rules to the map, do any of the following: • Use the **Quick Editor**. For details, refer to the How to Use the Quick Editor for Pass and Drop Rules. • Import a map template by clicking **Import**. • Create a rule by clicking **Add a Rule**. For detailed information about map rules, refer to Map Rules.
Map Order	
Priority	Set the **Map Order** by selecting the priority from the Priority list. For details about map priority, refer to Map Priority.
Map Permissions	
Owner	Select user group(s) who own the map.
Edit	Select user group(s) who can edit the map.
Listen	Select user group(s) who can listen to the traffic.

Field	Description
View	Select user group(s) who can view the map.
Map Tag	
Tags	Select the required tag key and tag value to which the map must be associated. The tag key and the associated tag values must be created in advance in GigaVUE- FM. Refer to the "Tags" and "Role Based Access Control" sections in the GigaVUE Administration Guide for more details. **Note**: When you associate a map to a tag value, then the ports, port groups, port pairs, GigaStreams that belong to the map are also associated to the tags.

Clone Map

You can create a copy of an existing map by doing the following:

1. Select **Maps > Maps** to open the Maps page.
2. Select the check box of the map that you to clone.
3. Click **Clone**.
4. Make changes to the map as necessary, such as specifying an alias.
5. Click **Save**.

Edit Maps

To edit an existing map:

1. Click the List view button.
2. Select the map on the Maps page by either selecting the check box and then clicking **Edit**, or click on the row in the table and clicking **Edit** in the Map Quick View.
3. Make the changes to the map.

   When editing a map, you can only modify the following:
   - Map alias
   - Description
   - Change the source and destination.
   - Select a different GS Operation.
   - Modify rules.
   - Add rules.
   - Delete rules.
   - Permissions as allowed.
4. Click **Save**.

Delete Maps

Keep in mind the following rules and notes before you delete a map:

- If you have configured an Active Visibility policy action for a map, ensure that you delete the Active Visibility policy before you delete the map. If you delete the map before deleting the Active Visibility policy, it would result in errors in the policy action.
- If you delete a map, which has a network port that is shared with another map, the map statistics will be reset for the other map.

To delete a map or maps, do the following:

1. Go to **Maps**, select **Maps > Maps** to open the Maps page.
2. Select the check box for the map or maps to delete, and then Click **Delete**.
3. When the message appears, asking if you want to delete the selected maps, click **OK**.

NOTE: In the GRIP configuration, when you delete a map on the primary node, irrespective of the inline-network traffic-path, the traffic is switched to the secondary node. The port utilization must be 0% on the primary node and active on the secondary node.

Create Map Groups

Map Groups are a collection of maps that are used with GTP Flow Sampling Overlap and GTP Whitelisting Overlap GigaSMART solutions.

Use the Map Groups page to create a group of maps for GTP forward listing and GTP flow sampling. All the maps in a map group receive traffic according to map rules, rather than map priority. Thus, multiple copies of a GTP packet can be sent to more than one tool. This functionality is referred to as overlapping maps.

The virtual port for specified as the source for GTP forward listing and GTP flow sampling must have **GTP Overlap** enabled.

When creating Map Groups keep the following in mind:

- A map group can be associated with only one GigaSMART group (gsgroup).
- All maps within a map group must be connected to the same vport.
- A map group can consist of only one GTP forward listing map or only one GTP flow sampling map but it cannot contains two maps of the same type.
- Once you have created a map group, you cannot edit it to change the type or subtype of the map. However, you can add or edit the map rules for a map, which is configured in a map group. Similarly, you cannot edit the map group alias but you can edit the aliases of the maps that are configured in the map group.

- If multiple map groups are configured, the maps within each map group must point to the same port groups as the other map groups.

To create a Map Group, do the following:

1. Go to **Maps**, select **Maps > Map Groups** to open the Maps page.
2. Click **New**.
3. Enter an alias for the map group.

   Use an alias that helps identify the map group.
4. (Optional) Enter a description about the map group. Refer to Description to Map Rules for the considerations regarding description.
5. Click in the **Maps** field and select the maps to add to the map group.
6. Click **Apply**.

Description to Map Rules

Use description to label the purpose of a rule or the type of traffic covered by a rule. To add a map rule description to a map select **Maps > Maps > Edit**.

Consider the following when adding map rule description:

- Use up to 128 characters, including spaces.
- Enclose the description in quotation marks, if the description is longer than one word.
- To include double quotation marks (") inside the quotation marks, precede the quote mark with a backslash (\).

Error Messages

Error messages are displayed when a description is invalid, for example:

- if the description is longer than one word and does not include double quotation marks
- if the description is longer than 128 characters
- if the rule with which the description is included is not valid.

Edit Map Rule Description

To edit a map rule description, do the following:

1. Select **Maps > Maps**
2. Select the map to edit.
3. Click **Edit**.
4. Change the description in the **Description** field

5. Click **Save** to recreate with a different description.

How to Use the Quick Editor for Pass and Drop Rules

When creating a map, you can use the Quick Rule Editor to quickly select custom port numbers for a map rule or add a range of IP address.

When manually configuring an Inline Bypass rule-based map in the device, it's important to understand how the Quick Rule Editor behaves based on the map's configuration mode:

a. If the map type is set to '**Inline**' mode before using the Quick Rule Editor. Any rules added through the Quick Rule Editor will be applied with default, non-editable values, resulting in "**read-only**" rules.

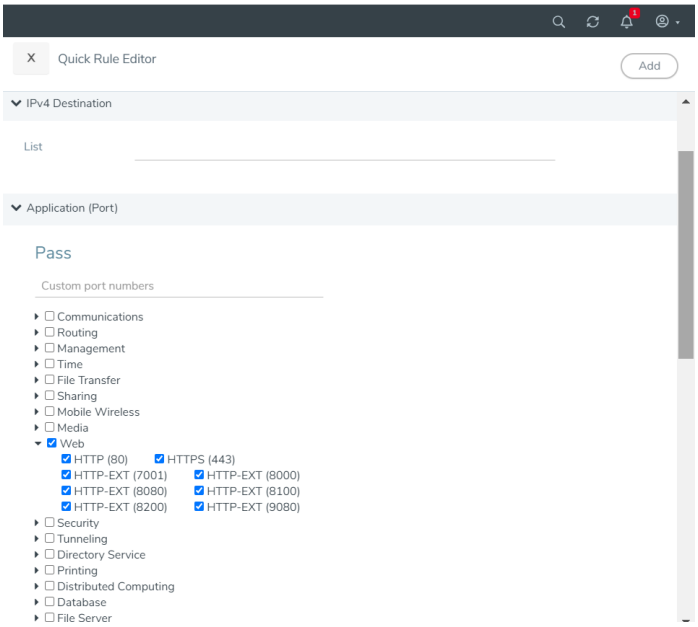
b. If the map remains in its default mode (not set to 'Inline') when using the Quick Rule Editor. Rules can be added and subsequently edited as needed.

How to Use the Quick Editor to Add Port Numbers

To use the Quick Rule Editor, do the following:

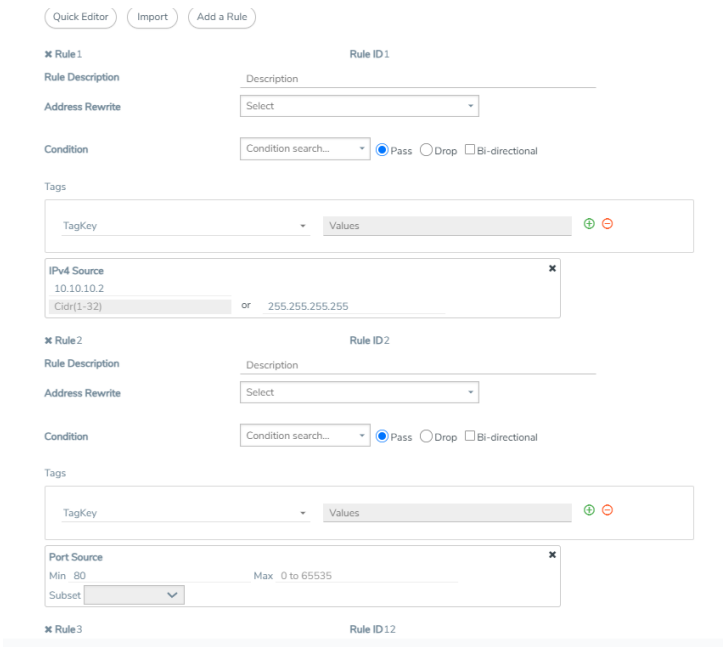
1. While on an Edit Map or New Map page, click **Quick Editor** under Map Rules. This opens the **Quick Rule Editor**.
2. On the Quick Rule Editor view, select the port number or numbers to add for a pass or drop rule or both.

   The Quick Rule Editor has two columns of custom port numbers, one for pass rules and one for drop rules. In each column, the ports are categorized by type. For example, **Web** provides a list of HTTP ports as shown in the following figure, where HTTPS port 443 is selected for pass and HTTP port 80 is selected for drop.



3. Click **Add**.

   A rule with a port source is added for each custom port number selected in the Quick Rule Editor. If the port was selected from the custom port numbers under Pass, the rule is a pass rule. If the port was selected from the port numbers under Drop, the rule is a drop rule. The following figure shows two rules added by the example shown in the previous step.



How to Use the Quick Port Editor to Add IP Address

Rather than enter IP address for map rules one at a time, the Quick Rule Editor makes it possible to enter a range to quickly add the IP addresses, saving time.

To enter a range of IP address with the Quick Rule Editor, do the following:

1. While on an Edit Map or New Map page, click **Quick Editor** under Map Rules. This opens the **Quick Rule Editor**.
2. Enter an IP address range in the **List** field under IPv4 Source or IPv4 Destination or both. For example, 10.10.10.9..11/32 in the **List** field under IPv4 Source .
3. Click **Add**.

   The Quick Rule Editor adds the IPv4 Source rules with the IP addresses. For example, if you entered 10.10.9..11 for the IPv4 Source, the editor adds three Ipv4 Source rules with the IP addresses 10.10.10.9/32, 10.10.10.10/32, 10.1010.11/32.

Map Templates

Map templates can be created by admin users. Once created, any user creating a map can use any template.

Admin users can define standardized traffic flows, applications, and rules that will be convenient for users when creating their maps. To do this, the admin creates map templates; later, users can use one or more templates as the basis for their maps.

Templates are created using the same rules and parameters as regular maps, but they do not have any network or tool ports. GigaSMART operations are also not included in templates.

The rules defined in the template become the starting point for the map. The rules can be edited or removed and new rules can be added to the map.

No changes made to the map will be reflected back in the original template. Once the map is created, it has no association with the original template from which it was created. Any changes to a template will not be reflected in any maps created with the previous version of the template.

Create Map Templates

To create a map template, do the following:

1. Select **Traffic > Maps > Map Templates**.
2. Click **New**. The New Map Template opens.
3. Enter an alias in the **Map Template Alias** field.
4. (Optional) Enter description about the template.
5. Add map rules by clicking **Add a Rule** for each rule that you want to add.
6. After you are done creating rules, click **Save**.

   The new map is included on the Map Templates page.

Edit Map Templates

To edit a map template, do the following:

1. Select **Traffic > Maps > Map Templates** to open the Map Templates page.
2. On the Map Template page, select the template to edit.
3. Click **Edit**.
4. Modify the map template by adding or deleting rules or description. (You cannot change the alias.)
5. Click **Save**.

Map Template Quick View

When you click on the alias of a map template on the Map Templates page, a Quick View displays that shows the comments (if any) and rules.

Manage Map Rule Resources

The resources available on a GigaVUE HC Series line card or node changes depending on the combination of criteria in place on the line card or node as a whole. In general, adding or removing MAC address, UDA pattern match, or IPv6 criteria in the map rules bound to a line card or node changes the type of filter template used on the line card or node. This can result in a brief interruption of traffic as the new template is applied.

Template Groups

The template groups are listed sequentially from least resource-intensive to most resource-intensive:

- **IPv4 Only** – This is the default filter template, including all IPv4 arguments without any MAC addresses, UDA data patterns, or IPv6 arguments. This template can support the IPv4 and related filter criteria, including VLAN tags, source/destination ports, protocol criteria, and so on.
- **IPv4 and MAC Addresses** – This template combines MAC address criteria with the standard IPv4-related criteria. When MAC address criteria are in use, map rule resources are decreased.
- **IPv4 and Single UDA Data Pattern** – This template combines one of the two available UDA data patterns with the standard IPv4-related criteria. Using a single UDA criteria does not affect the total number of drop map rules available, but it does decrease the number of pass map rules available.
- **Both UDA Data Patterns** – This template uses both UDA data patterns but removes the ipv4 argument. Drop map rule availability is not affected by adding a second UDA data pattern, but pass map rules are decreased again from what was available when only a single UDA was used.
- **IPv6 Arguments** – This template adds the use of the IPv6 argument. IPv6 criteria are resource-intensive, significantly decreasing both drop and pass map rule capacity, as shown in the following table. Note also the changes in available criteria and available resources.

The map rule criteria available in each filter template (or "group") is shown in the following table.

Table 4: Map Rule Criteria for Default Templates

Rules - Pass and Drop	IPv4	IPv4 + MAC	IPv4 + UDA	UDA	IPv6
IPv4	✓	✓	✓		✓
IPv6					✓
MAC		✓			
UDA1			✓	✓	
UDA2				✓	
VLAN	✓	✓	✓	✓	✓
Inner VLAN	✓	✓	✓	✓	✓
L4 Port dst and src	✓	✓	✓	✓	✓
Ethertype		✓	✓		
IP Ver	✓	✓	✓	✓	✓
Protocol	✓	✓	✓	✓	✓
DSCP	✓	✓	✓	✓	✓
ToS	✓	✓	✓	✓	✓

Rules - Pass and Drop	IPv4	IPv4 + MAC	IPv4 + UDA	UDA	IPv6
TCP Ctl	✓	✓	✓	✓	✓
IP Frag	✓	✓	✓	✓	✓
TTL	✓	✓	✓	✓	✓
IPv6 Flow Label		✓	✓		

The number of rule entries in a cluster is shown in the following table.

Table 5: Rule Entries

# Rule Entries	IPv4	IPv4 + MAC	IPv4 + UDA	UDA	IPv6
GigaVUE-HCT	3071	3071	3071	3071	3071
GigaVUE-HC1-Plus	3071	3071	3071	3071	3071
GigaVUE-HC3	4096(1024 per slot)	4096(1024 per slot)	4096(1024 per slot)	4096(1024 per slot)	4096(1024 per slot)
GigaVUE-HC1 Node	16384	8192	8192	8192	8192
GigaVUE-TA100 Node	1024	1024	1024	1024	1024
GigaVUE-TA200 Node	1024	1024	1024	1024	1024
GigaVUE-TA200E	1024	1024	1024	1024	1024
GigaVUE-TA25	3071	3071	3071	3071	3071
GigaVUE-TA25E	3071	3071	3071	3071	3071
GigaVUE-TA400	10229	6133	10229	10229	6133
GigaVUE-TA400E	10229	6133	10229	10229	6133

NOTE: The above table lists the values for the Default template with license.

Refer to the following notes:

- You can use flexible filter templates on GigaVUE-HC3 to support up to 6K rules per slot.
- GigaVUE-TA100, GigaVUE-TA200 and GigaVUE-TA200E support up to 6K rules per pseudo-slot or 24K rules per node.
- GigaVUE-TA25 supports up to 18K rules per node and GigaVUE-TA25E supports up to 36k rules per node.

Refer to Filter Templates for details.

Add Tags to Map Rules

You can associate tags to the map rules. Each rule within a map can be associated with different tag keys and tag values. Tagging at rule level allows you to view the statistics of the flow of traffic based on the rules associated to the traffic.

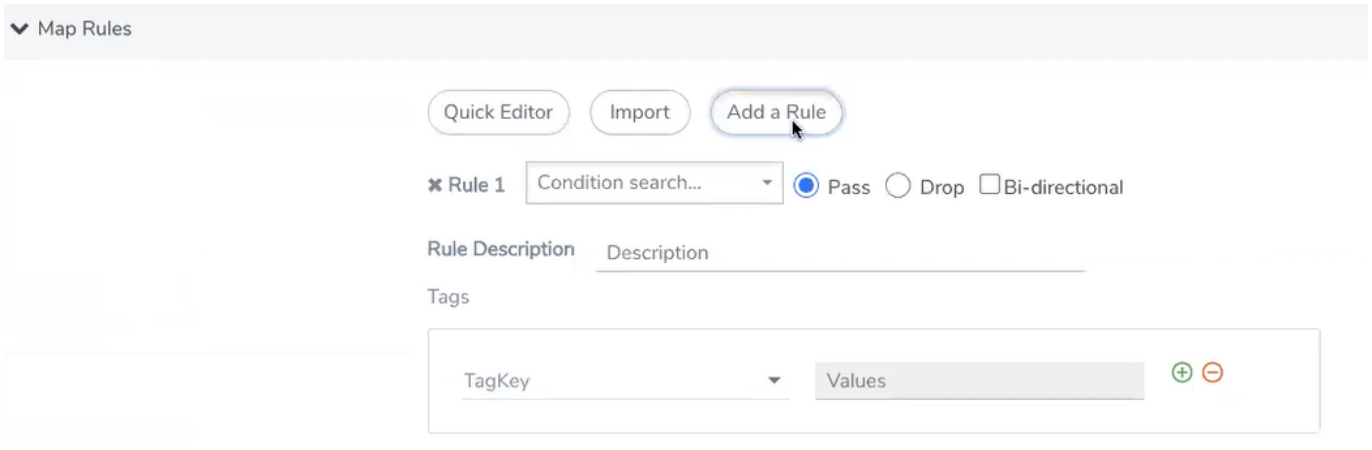
Consider the following maps and map rules associated to tags:

Maps	Tag Key	Tag Value
Map 1 Src: 1/1/x1 Dst: 1/1/x2	Maptag	M1
- Rule 1 Vlan 100	Flow	Flow1
- Rule 2 Vlan200		Flow2
Map 2 Src: 1/1/x3 Dst: 1/1/x4	Maptag	M2
- Rule 1 Vlan 100	Flow	Flow1
- Rule 2 Vlan200		Flow2

With rule-level tagging, you can view the traffic flow associated with Rule 1 of both the Maps based on the tag value Flow1, even though the rules are defined as part of different maps. Refer to the Map Rule Statistics Dashboard section for viewing statistics related to the flow of traffic based on the rules defined.

To associate tags to map rules:

1. When creating a map, click on **Add a Rule.**
2. Add the required tags to the rule.



Notes

Refer to the following notes:

- Rule level tagging is applicable for:
  - Flow Maps
  - Fabric Maps
- Rule level tagging is applicable for the following map types:
  - Regular Maps
  - First Level Maps
- You can associate only aggregation tags with map rules. However, traffic statistics displayed in the Map Rule Statistics dashboard is based on the following
  - RBAC tags associated at the map level.
  - Aggregation tags associated at the rule level.

Map Rule Statistics Dashboard

The Map Rule Statistics dashboard displays statistical data related to the map rules and the associated traffic. To view the Map Rule Statistics dashboard:

1. On the left navigation pane, select  -> **Analytics** -> **Dashboards**.
2. Click **Dashboards** from the top menu. Select **Map Rule Statistics** from the list of dashboards.

The **Map Rule Statistics** dashboard is displayed. The visualizations are displayed in the following two tabs:

- Metric
- Trend

The Metric tab displays the following visualizations

- **Total Traffic**: Total traffic in packets per second
- **Average Traffic**: Average traffic in packets per second
- **Associated Map Rules**: Map rules associated with the traffic

The Trend tab displays the following visualizations:

Visualizations	Displays...
Maximum vs. Average Traffic in pps	Maximum traffic versus average traffic in packets per second on an hourly interval.
Maximum vs. Average Traffic in bps	Maximum traffic versus average traffic in bytes per second on an hourly interval.
Traffic Distribution at Rule Level in pps	Traffic distribution at rule level in packets per second.
Top-N Rules by Maximum Traffic in pps	Top 5 rules contributing to the maximum traffic in packets per second
Top-N Rules by Average Traffic in pps	Top 5 rules contributing to the average traffic in packets per second

Use the following 'Control Visualizations' to filter the traffic:

- Flow Sample
- Cluster ID
- Host Name
- Map Rule Sample

Customize the **Flow Sample and Map Rule Sample** options to filter the data based on the tags configured in GigaVUE-FM. Refer to the Filter Data Using Tags in Control Filters section for more details.

NOTE: Refer to the Analytics section for information about filtering and searching in the dashboards page.

Filter Templates

Filter templates maximize the number of map rules, optimize filter resources, and enhance the scalability and flexibility of flow mapping. Flexible filter template is supported in GigaVUE-HC1-Plus, GigaVUE-HCT, GigaVUE-HC1, GigaVUE-HC3, GigaVUE-TA100, GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA400, and GigaVUE-TA400E.

Refer to Manage Map Rule Resources for template groups on other GigaVUE nodes.

Flexible filter templates increase the number of map rules and also eliminate current restrictions on map rule combinations, such as ipv6+MAC or ipv6+UDA.

Refer to the section Flow Mapping® FAQ for the number of map rules supported.

Flow Mapping® uses filter templates to determine the traffic to filter based on qualifiers specified in the template. A filter template has a specific set of qualifiers used to apply to map rules. You can control the template that you apply to a specific slot on GigaVUE-HC3 or a specific pseudo-slot on GigaVUE-TA100, GigaVUE-TA200, GigaVUE-TA400, and GigaVUE-TA400E. For GigaVUE-HC1, you can apply the filter template only at the control card level which will be applied across all the line cards.

Flexible filter templates offer five default templates. Custom templates can also be created that have a qualifier set selected from the list of available qualifiers.

For the GigaVUE-TA400, and GigaVUE-TA400E platform, custom templates can be created with the inner qualifiers and MPLS headers. Inner qualifiers provide extended support for packet processing in the Flexible filter template. MPLS header attributes now support up to 7 levels, allowing for detailed and specific MPLS header field mapping and processing.

Refer to the following sections for details:

- Filter Template Qualifiers and Defaults
- Filter Template Configuration
- Filter Template Limits
- Filter Template Rules and Recommendations
- Filter Template Best Practices
- Filter Templates in a Cluster
- Filter Templates Formulas
- Troubleshoot Custom Filter Template Issues

Filter Template Qualifiers and Defaults

Refer to the rows in Map Rule Criteria for Default Templates for the list of qualifiers for filter templates. Refer to the columns in Table 4: Map Rule Criteria for Default Templates for the default templates and the qualifiers that are predefined for the defaults.

NOTES:

- The default templates cannot be deleted.
- The **ipver** qualifier is implicitly included in all default and custom templates.

Filter Template Configuration

To configure filter templates:

1. On the left navigation pane, click  and select **Physical**.
2. Select the required GigaVUE node.
3. Select **Maps > Filter Templates**.
4. To add a custom template, click **New**.
5. Specify an alias, an optional description, then select qualifiers. Click **OK**.
6. To apply a custom template to a slot or pseudo-slot, select it and click **Apply**.

   > **NOTE:** For GigaVUE-HC1, you can apply a filter template only at the control card level which will be applied across all the line cards

7. Select the slot or pseudo-slot and click **OK**.

   The Filter Templates page displays the applied slot or pseudo-slot. You can edit an existing custom filter or delete it. A template can be deleted if it is not currently in use, meaning that it has not been applied.
8. To display filter templates, click on a row in the Filter Templates page.

Filter Template Limits

The number of qualifiers in a template limits the total number of rules that can be defined. The maximum rule limit on the GigaVUE-HC3, GigaVUE-TA100, or GigaVUE-TA200 is 1K (1024) per slot or pseudo-slot when using the default templates.

Custom templates allow the creation of templates with only those qualifiers needed for the rules that you plan to use in flow maps. The qualifiers specified in a flexible template can increase or decrease the maximum rule limit, depending on the qualifiers selected. With flexible filter templates, it is possible to reach a maximum limit of 6K rules per slot on the GigaVUE-HC3 node and 6K rules per pseudo-slot , or 24K total rules on the GigaVUE-TA100 ,GigaVUE-TA200 and GigaVUE-TA200E node, 18k rules per node on GigaVUE-TA25 node, 36k rules per node on GigaVUE-TA25E node, on GigaVUE-TA400, and GigaVUE-TA400E node, 22K rules per pseudo-slots.

How to Understand Map Filter Resources

Starting in software version 5.0, when a filter template is applied, filter resources display the total number of map rules used in a map as well as the limit. If the limit is 1024, 1023 is displayed, even though the actual limit is 1022, or two less than the limit. This discrepancy is

due to extra resources needed for internal usage.

Filter Template Rules and Recommendations

When creating flexible filter templates, keep the following rules and recommendations in mind:

- Filters are applied to a specific slot or pseudo-slot, not to the node.
- By default, all slots will be in the pre-defined **ipv4** template.
- There is a limit of 512 custom templates.
- Custom templates can have duplicate sets of qualifiers.
- The filter limit is calculated when the template is created. In most cases, a higher-cost qualifier set (for example, IPv6, UDA, or MAC are higher cost) consumes more resources and leads to a lower filter limit.
- Flexible filter templates have no effect on existing flow mapping behavior, including pass versus drop map rules, map priority, network port sharing, GigaSMART operations, or first level and second level maps.
- When deploying a Resilient Inline Arrangements (RIA) map in IPv6, specifying MAC source address, MAC destination address and EtherType as qualifiers will not be accepted.
- When configuring filter templates, certain combinations of qualifiers are not supported on some of the platforms even though the total bits consumed by the qualifiers is less than (480-54) bits. This is due to the limitations in the hardware. For example, a flexible filter template configured with qualifiers 'ipsrc ipdst portsrc portdst uda1 uda2' is supported on GigaVUE-HC3.
- To use Flexible Filter Template along with inline-netlag, both VLAN and inner VLAN must be added in addition to other needed qualifiers.

  NOTE: To verify if a flexible filter template is supported on a specific platform and the number of rules supported, you must create the template with the desired qualifiers and execute the `show filter-template limit` command. If the number of rules is N/A, then it indicates that this combination of qualifiers is not supported on the corresponding platform.

Filter Template Best Practices

The following are best practices for optimizing filter resources using filter templates.

First determine all the needed qualifiers, then create a template, apply the template, and configure the map rules.

Example 1 - Applicable for the all the platforms

- Connect network ports of a slot to flows of the same application. For example, if you have two flows:
  - one is filtered on **macsrc** and **macdst**

    - the second one is filtered on **ipdst** and **ipsrc**
- In case both flows connect to ports on the same slot, that slot will have to have a template of **macsrc**, **macdst**, **ipsrc**, and **ipdst**, with a limit of 1024 rules.
- However, filter resources can be optimized by connecting these two flows to ports on different slots with one template for **macsrc** and **macdst** and the other template for **ipsrc** and **ipdst**. Both templates will have a limit of 3072 rules.

Example 2 - Applicable for the GigaVUE-TA400 platform

- Connect network ports of a slot to flows of the same application. For example, if you have two flows:
    - one is filtered on **mpls-label-id** and **mpls-label-exp**
    - the second one is filtered on **inner ipsrc** and **inner ipdst**
- In case both flows connect to ports on the same slot, that slot will have to have a template of **mpls-label-id**, **mpls-label-exp**, **inner ipsrc**, and **inner ipdst**.
- However, filter resources can be optimized by connecting these two flows to ports on different slots with one template for **mpls-label-id** and **mpls-label-exp** and the other template for **inner ipsrc** and **inner ipdst**.

Example 3 - Applicable for the GigaVUE-TA400 platform

- Connect network ports of a slot to flows of the same application. For example, if you have two flows:
    - one is filtered on **macsrc** and **macdst**
    - the second one is filtered on **inner ipdst** and **inner ipsrc**
- In case both flows connect to ports on the same slot, that slot will have to have a template of **macsrc**, **macdst**, **inner ipsrc**, and **inner ipdst**.
- However, filter resources can be optimized by connecting these two flows to ports on different slots with one template for **macsrc** and **macdst** and the other template for **inner ipsrc** and **inner ipdst**.

The following are best practices for adding more rules if a limit has been reached:

- Create a new template with all the qualifiers that are in use on a specified slot.
- Issue the **show filter-resource** command to obtain the list of qualifiers in use.
- Issue the **filter-template alias <alias> qualifiers add** command with that list of qualifiers.
- Issue the **show filter-template limit** command to check if the new template allows a higher limit. If it does, apply the filter using the **card slot <slot ID> filter-template** command.

Filter Templates in a Cluster

Filter template configuration is synchronized across the cluster. However, a cluster can have different GigaVUE nodes, so one set of qualifiers may or may not be valid on all nodes.

Filter Templates Formulas

The formulas in this section can help you determine the number of map rules that are supported, based on the qualifiers specified in the filter template. Use the formulas as guidelines.

The number of map rules depends on the number of qualifiers a template can support.

Table 6: Maximum Cost Supported per Map Rule for Different Qualifier Combinations

Qualifier Combination	Maximum cost supported per map rule
Outer Header attribute with IPv6 qualification	10
Outer Header attribute without IPv6 qualification	12
Inner Header and MPLS Header attribute	18

NOTE: Combining outer header attributes with inner or MPLS header attributes in a map rule is supported, and the maximum cost per map rule will increase accordingly.

The cost of each qualifier depends on the number of bits it consumes. The following table lists the number of bits each qualifier consumes and the cost for each qualifier.

Table 7: Bits Consumed and Cost per Qualifier

Qualifier	Bits	Cost
dscp	6	1
ethertype	16	1
ip6src	128	4
ip6dst	128	4
ip4src	32	1
ip4dst	32	1
macdst	48	2
macsrc	48	2
macsrc and macdst	96	3

Qualifier	Bits	Cost
portsrc	16	1
portdst	16	1
portsrc and portdst	32	1
protocol	8	1
tos	8	1
ttl	8	1
vlan	16	1
vxlan	48	1
l2gre	24	1
uda1	128	4
uda2	128	4
inner-vlan	16	1
qset1*	58*	*

* qset1 is made up of the following: tos: 8, ipfrag: 2, tcpctl: 8, ttl: 8, ip6fl: 32. The cost depends on the combination of the qualifiers used.

Table 8: Bits Consumed and Cost per Inner Qualifier and MPLS header for the GigaVUE-TA400

Qualifier	Bits	Cost
inner ethertype	16	1
inner ip6src	128	4
inner ip6dst	128	4
inner ipfrag	32	2
inner ip4dst	32	1
inner ip4src	32	1
inner portdst	16	1
inner portsrc	16	1
inner portdst and inner portsrc	32	1

Qualifier	Bits	Cost
inner protocol	8	1
inner uda1	64	4
inner uda2	64	4
mpls-label-id	20	1 per level
mpls-label-exp	3	1 per level
mpls-label-bos	1	1 per level
mpls-label-ttl	8	1 per level

Refer to Filter Template Limits section to know the maximum rules supported for the selected filter template.

Troubleshoot Custom Filter Template Issues

Starting with release 6.10, certain combinations of custom filter-template configurations assigned to slots on GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-HCT, and GigaVUE-HC1-Plus devices are no longer supported. This change is due to the introduction of new feature sets, which increase the overall size of the filter-template. Specifically, unsupported configurations include custom filter-templates that use the maximum combination of qualifier sets exceeding the supported TCAM width in BCM.

Examples for combinations which have become unsupported: [Any one of existing qualifier can be removed to support this combination]

{SIP6, MACSRC, MACDST, PROTOCOL, L4 PORTDST, UDA1, UDA2}

{SIP6, DIP6, PROTOCOL, L4 PORTSRC, UDA1, UDA2}

{SIP6, DIP6, MACSRC, ETHERTYPE, UDA1, UDA2}

Whenever you configure an unsupported custom filter-template and assign it to a slot, the system will restrict the configuration and display an error message.

Sample CLI Output:

```
HC1P (config) # filter-template alias ft1 qualifiers add ip6src macsrc macdst ipdst
portdst protocol uda1 uda2
HC1P (config) # card slot 1/1PS1 filter-template ft1
% Filter-template 'ft1' is not valid for the specified slot.
```

A CLI command is available to help determine whether a custom filter-template is supported on a GigaVUE chassis without applying it to a slot. A result of “N/A” indicates that the template is unsupported.

```
HC1P (config) # show filter-template limit
Total Filter-Template-Limit Count: 6
=================================================================================================================================================================
Template Alias HC1      HCT     HC2     HC2P    ---------------------- HC3 ----------------------- ---------------------- HD ----------------------
               CC1      CC1     CC1     CC1     C16     C08     C05     C08FGS  X24     X16     X16Q4LR X16C4LR 12X04   X04G44  Q02X32  C02X08  C06X24  Q08     TS
-----------------------------------------------------------------------------------------------------------------------------------------------------------------
ipv4           16384    3071    4096    16384   1024    1024    1024    1024    1024    1024    1024    1024    2048    2048    2048    2048    16384   2048    2048
ipv6            8192    3071    2048    8192    1024    1024    1024    1024    1024    1024    1024    1024    512     512     512     512     4096    512     512
ft1              N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A
uda             8192    3071    2048    8192    1024    1024    1024    1024    1024    1024    1024    1024    512     512     512     512     4096    512     512
ipv4+mac        8192    3071    2048    8192    1024    1024    1024    1024    1024    1024    1024    1024    1024    1024    1024    1024    8192    1024    1024
ipv4+uda        8192    3071    2048    8192    1024    1024    1024    1024    1024    1024    1024    1024    1024    1024    1024    1024    8192    1024    1024
-----------------------------------------------------------------------------------------------------------------------------------------------------------------
=================================================================================================================================================================
Template Alias -----HC1P----- LY2      TA1     TA10    TA40    TAC     TACX    TA200   TA25    DELL_S4112F TA400   TA200E  TA25E   TA400E
               PS1     PS2     CC1     CC1     CC1     CC1     CC1     CC1     CC1     CC1     CC1     CC1     CC1     CC1     CC1
-----------------------------------------------------------------------------------------------------------------------------------------------------------------
ipv4           3071    3071    2048    2048    2048    4096    1024    1024    1024    3071    8192    10229   1024    3071    10229
ipv6           3071    3071    512     512     512     2048    1024    512     1024    3071    4096    6133    1024    3071    6133
ft1             N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A     6133    N/A     N/A     6133
uda            3071    3071    512     512     512     2048    1024    512     1024    3071    4096    10229   1024    3071    10229
ipv4+mac       3071    3071    1024    1024    1024    2048    1024    1024    1024    3071    4096    6133    1024    3071    6133
ipv4+uda       3071    3071    1024    1024    1024    2048    1024    1024    1024    3071    4096    10229   1024    3071    10229
-----------------------------------------------------------------------------------------------------------------------------------------------------------------
```

Upgrade process (as applicable)

If you were using any unsupported combination of custom filter-templates in your environment prior to release 6.13 and then upgraded to 6.13, the restoration process will complete successfully, but internal errors may be observed. After the upgrade, you must manually verify whether any unsupported custom filter-templates were applied to a slot by running the CLI command below. Check the “Map Rules Limit” value in the output—if it is displayed as “0”, the template is unsupported.

```
HC1P (config) # show filter-resource
=============================================================================================
Box/  Map Rules    Filter      App-filter   Port Filter        In-Use
Slot  Used Limit   Template    Used Limit   Used Limit         Qualifiers
---------------------------------------------------------------------
 1/1PS1 9     0    ft1         2    504     0    448           ipdst, ip6src, uda1 and uda2
 1/1PS2 6     3071 ft2         2    504     0    448           ip6src and macsrc
HC1P (config) # show filter-template alias ft1 brief
Total Filter-Template Count: 1
=============================================================================================
Template Alias       Qualifier List
---------------------------------------------------------------------------------------------
ft1                  circuit-id,ip6src,ipdst,macdst,macsrc,portdst,protocol,uda1,uda2
---------------------------------------------------------------------------------------------
```

Recovery Procedure

You must remove all flow map configurations created on top of the unsupported filter-template, configure the filter-template as “defaults” for the respective slot, and then follow one of the options below:

- Replace the existing custom filter-template qualifier set with a reduced combination, or split it into two different custom filter-templates that are supported on the chassis. Then associate the supported filter-template to the slot and manually reconfigure all flow map configurations.
- Migrate the unsupported custom filter-template to a different chassis such as GigaVUE-HC3, GigaVUE-TA200, and GigaVUE-TA400, which may support the combination (verify using the show filter-template limit CLI command). After migration, reconfigure the flow map configurations manually. For example,

```
HC1P (config) # filter-template alias ft1 qualifiers add ip6src macsrc macdst
ipdst portdst protocol uda1 uda2
HC1P (config) # card slot 1/1PS1 filter-template ft1
% Filter-template 'ft1' is not valid for the specified slot.
```

As shown in the example above, filter-template ft1 is unsupported on the GigaVUE-HC1-Plus chassis but is supported on GigaVUE-TA400/GigaVUE-TA400E.

Review Map Statistics with Map Rule Counters

Map Statistics can be viewed in the following ways:

- The Statistics page.

  For details, refer to Viewing Map Statistics with the Statistics Page
- The Map Quick View.

  For details, refer to Viewing Map Statistics with Quick View

A single packet may match multiple rules in the map and will not cause multiple rule counters to increment. Only the last rule which is the highest priority in the order will increment. The flow map rule priority is based on the order it was created. Thus, the sum of the map rule counters across all the map rules may be higher than the total number of packets received and transmitted by the map.

NOTE: Drop rules have a higher priority than pass rules.

For example, consider the following three map rules:

- **Rule 1** – ipsrc 10.10.0.0 /24 bidir
- **Rule 2** – ipsrc 10.10.0.100 /32 bidir
- **Rule 3** – portsrc 80

A packet with ipsrc 10.10.0.100 and portsrc 20 will match Rule 1 and Rule 2, which will be forwarded to the tool port or ports. The counters for Rule 2 will only be incremented.

There are several reasons a map rule counter may not increment:

- Traffic matching the rule is not currently present in the production network.
- The network port is not monitoring the network at the proper location to see the traffic specified by the map.
- A higher-priority map is forwarding the packet before it can be inspected by this particular map.
- The map rule itself may be specified incorrectly.

Viewing Map Statistics with the Statistics Page

To review map statistics indicating the total packets and total octets handled by maps and the number of rules in the map, select **Maps > Maps > Statistics**. The Statistics page displays the map statistics in a table format listing the maps by their alias. Clicking on a map alias opens a Quick View for that map. To clear map counters, click the **Clear** button.

Viewing Map Statistics with Quick View

To review map rule counters indicating the number of rule matches for a map as packets are inspected and forwarded, select the map and view the information in the Quick View window.

Flow Mapping® FAQ

This section answers frequently asked questions by users migrating to the Flow Mapping® model.

How Many Map Rules are Supported?

The maximum number of map rules supported per line card or standalone node are shown in Maximum Number of Nodes and Map Rules Supported in a Cluster.

Table 9: Maximum Map Rules

Node Type	Maximum Combined Rules for Default Template Without a License	Maximum with Flexible Filter Templates Without a License	Maximum with Advanced Features License for Flexible Filter Templates
GigaVUE-HCT Node	6k (6142)	6k (6142)	N/A
GigaVUE-HCT Module	3k (3071)	3k (3071)	N/A
GigaVUE-HC1-Plus Node	6k (6142)	36k (36862)	N/A
GigaVUE-HC3 Node	4K (4096)	24K (24576)	N/A
GigaVUE-HC3 Module	1K (1024)	6K (6144)	N/A
GigaVUE-HC1	16K (16383)	16K (16383)	N/A
GigaVUE TA Series	256	The Flexible Filter Templates are available for the following TA platforms: GigaVUE-TA100 GigaVUE-TA200 GigaVUE-TA25 GigaVUE-TA400 GigaVUE-TA200E GigaVUE-TA25E	The Advanced Features License is available only for the GigaVUE-TA Series. The extended max rule limit varies based on the GigaVUE-TA Series node.
GigaVUE-TA25	256	256	18k (18431) rules per node
GigaVUE-TA25E	256	256	36k (36862) rules per node
GigaVUE-TA100	256	256	6K (6144) per pseudo-slot with Flexible Filter Template and Advanced Feature License. 1K (1024) per pseudo-slot with only Advanced Feature License.
GigaVUE-TA200	256	256	6K (6144) per pseudo-slot with Flexible Filter Template and Advanced Feature License. 1K (1024) per pseudo-slot with only Advanced Feature License.
GigaVUE-TA200E	256	256	6K (6144) per pseudo-slot with Flexible Filter Template and Advanced Feature License. 1K (1024) per pseudo-slot

Node Type	Maximum Combined Rules for Default Template Without a License	Maximum with Flexible Filter Templates Without a License	Maximum with Advanced Features License for Flexible Filter Templates
			with only Advanced Feature License.
GigaVUE-TA400	256	256	22k map rules (22517) per pseudo-slot (or) 88k map rules in total for 4 available pseudo slots of TA400 with Flexible Filter Template (depends on the qualifier list. Reduces when more qualifiers included) and Advanced Feature License. 10K map rules (10299) per pseudo-slot (or) 40k map rules in total for 4 available pseudo slots of TA400 with only Advanced Feature License.
GigaVUE-TA400E	256	256	22k map rules (22517) per pseudo-slot (or) 88k map rules in total for 4 available pseudo slots of GigaVUE-TA400E with Flexible Filter Template (depends on the qualifier list. Reduces when more qualifiers included) and Advanced Feature License. 10K map rules (10299) per pseudo slot (or) 40k map rules in total for 4 available pseudo slots of GigaVUE-TA400E with only Advanced Feature License.

GigaVUE-TA400: The reference logs as below:

With Pseudo slot:

ST-32N-HC3v2-5-22 [32N-LSCLUSTER999999999: leader] (config) # show filter-resource box-id 38

Slot : 38/1PS1

Filter-Template : portdst (portdst)

In-Use Qualifiers : NONE

Map Rules Usage (used/limit) : 0/22517

Port Filter Usage (used/limit) : 0/950

App-filter usage (used/limit) : 0/1016

Slot : 38/1PS2

Filter-Template : portdst (portdst)

In-Use Qualifiers : NONE

Map Rules Usage (used/limit) : 0/22517

Port Filter Usage (used/limit) : 0/950

App-filter usage (used/limit) : 0/1016

Slot : 38/1PS3

Filter-Template : portdst (portdst)

In-Use Qualifiers : NONE

Map Rules Usage (used/limit) : 0/22517

Port Filter Usage (used/limit) : 0/950

App-filter usage (used/limit) : 0/1016

Slot : 38/1PS4

Filter-Template : portdst (portdst)

In-Use Qualifiers : NONE

Map Rules Usage (used/limit) : 0/22517

Port Filter Usage (used/limit) : 0/950

App-filter usage (used/limit) : 0/1016

ST-32N-HC3v2-5-22 [32N-LSCLUSTER999999999: leader] (config) #

Without Pseudo Slots:

ST-32N-HC3v2-5-22 [32N-LSCLUSTER999999999: leader] (config) # show filter-resource box-id 38

Slot : 38/1PS1

Filter-Template : NONE (NONE)

In-Use Qualifiers : NONE

Map Rules Usage (used/limit) : 0/10229

Port Filter Usage (used/limit) : 0/950

App-filter usage (used/limit) : 0/1016

Slot : 38/1PS2

Filter-Template : NONE (NONE)

In-Use Qualifiers : NONE

Map Rules Usage (used/limit) : 0/10229

Port Filter Usage (used/limit) : 0/950

App-filter usage (used/limit) : 0/1016

Slot : 38/1PS3

Filter-Template : NONE (NONE)

In-Use Qualifiers : NONE

Map Rules Usage (used/limit) : 0/10229

Port Filter Usage (used/limit) : 0/950

App-filter usage (used/limit) : 0/1016

Slot : 38/1PS4

Filter-Template : NONE (NONE)

In-Use Qualifiers : NONE

Map Rules Usage (used/limit) : 0/10229

Port Filter Usage (used/limit) : 0/950

App-filter usage (used/limit) : 0/1016

ST-32N-HC3v2-5-22 [32N-LSCLUSTER999999999: leader] (config) #

The limit for GigaVUE TA Series standalone nodes is 256 combined pass/drop rules but is up to 2048 with the Advanced Features License installed.

Refer to Manage Map Rule Resources for managing map rules.

The maximum number of nodes and map rules supported when in a cluster is as follows:

Table 10: Maximum Number of Nodes and Map Rules Supported in a Cluster

When a Cluster is Configured with:	Number of Nodes	Maximum Map Rules
Out-of-Band Cluster Management	32	38000
Inband Cluster Management	16	38000

The maximum number of map rules supported in a cluster apply to all nodes in the cluster including GigaVUE HC Series nodes: GigaVUE-HC3, GigaVUE-HC1, GigaVUE-HC1-Plus, GigaVUE-HCT and GigaVUE TA Series nodes: GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA100, GigaVUE-TA200, GigaVUE-TA200E, and GigaVUE-TA400 including Certified Traffic Aggregation White Box (white box).

How Many Rules Can Each Map Have?

The maximum number of rules per map is 4K (4096), except on products that only support a total of 2K map rules. GigaVUE-HC3 supports 6K rules per map. Refer to Rules Per Map.

Table 11: Rules Per Map

Node Type	Rules per map
GigaVUE-HCT Node	4K
GigaVUE-HCT Module	4K
GigaVUE-HC1-Plus Node	4K
GigaVUE-HC1-Plus Module	4K
GigaVUE-HC3 Node	6K
GigaVUE-HC3 Module	6K
GigaVUE-HC1	4K
GigaVUE TA Series	4K
GigaVUE-TA25	4K
GigaVUE-TA25E	4K
GigaVUE-TA100	4K

Node Type	Rules per map
GigaVUE-TA200	4K
GigaVUE-TA200E	4K
GigaVUE-TA400	4K

How Many Maps Can Run at Once?

A maximum of 16000 maps can be configured on a cluster or standalone device.

What Criteria can be Filtered in Q-in-Q Packets?

Maps on GigaVUE nodes can match Layer 3/Layer 4 criteria in packets using Q-in-Q with up to two tags. For more information refer to How to Handle Q-in-Q Packets in Maps.

How Many Maps Can Share a Network Port?

A maximum of 512 maps can share the same source port.

How Many Network Ports and Tool Ports Can Be in a Map?

If the ports are not in a GigaStream, the number of individual map ports in the **Source** or **Destination** field of a map is limited to 64 on all GigaVUE HC Series nodes. The individual ports can be any of the following port types: network or tool.

On GigaVUE-HC3, if the ports are in a GigaStream, the limit is 95.

Are Port-Filters Supported?

Yes.

Egress port-filters (tool, hybrid, circuit, and inline network), are less efficient and scalable than flow maps, but they do provide a convenient way to narrow down the traffic seen by the tools/GigaVUE® HC Series nodes without having to change an entire map. Refer to Port Filters for details.

Each GigaVUE-HC3 module, or GigaVUE-HB1 node supports 100 combined egress port-filters. A single filter applied to multiple tool ports counts multiple times against the 100-filter limit.

In the the GigaVUE-HC1 or GigaVUE-HC3 node, the limit is 400 filters.

The GigaVUE TA Series can only support 20 egress port-filters. When the GigaVUE-TA100 or GigaVUE-TA200 , GigaVUE-TA200E are in a cluster, they can support 400 filters. GigaVUE-TA25, GigaVUE-TA25E supports 20 port filters by default and up to 100 port filters with advanced feature license.

Does Flow Mapping® Support Passalls?

Yes.

Flow Mapping® supports passalls with the following:

- The map **Subtype Pass All** option for network to tool port passalls.
- A tool-mirror connection between to tool ports. The Tool Mirrors page replaces the tool-to-tool port passalls. To create a Tool Mirror, select **Ports > Tool Mirrors** to go to the Tool Mirror page shown in Figure 10Tool Mirror Page.

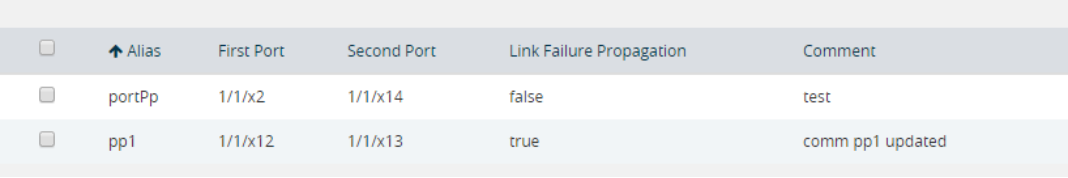


Figure 10 *Tool Mirror Page*

Does Flow Mapping® Support port-pairs?

Yes.

Select **Ports > Port Pairs** to go to the Port Pairs page shown in Figure 11Creating Port Pairs. For more information about port pairs and details on configuring two ports as an inline TAP, refer to Port Pairs.



Alias	First Port	Second Port	Link Failure Propagation	Comment
portPp	1/1/x2	1/1/x14	false	test
pp1	1/1/x12	1/1/x13	true	comm pp1 updated

Figure 11 *Creating Port Pairs*

Does Flow Mapping® Support UDA Pattern Matches?

Yes.

Pattern match rules are still supported in map rules. They can be used in both pass rules and drop rules. Refer to Work with User-Defined Pattern Match Rules for details.

Are Maps Supported Across Nodes in a Cluster?

Yes.

Clusters of GigaVUE nodes operate as a unified fabric. Use the standard box ID/slot ID/port ID syntax to create packet distribution, just as on a standalone node. Maps can have network ports and tool ports on different physical nodes within the cluster.

Similarly, a map does not need to keep its network and tool ports on the same physical node in order to take advantage of GigaSMART operations – a GigaVUE TA Series node in a cluster can take advantage of the GigaSMART processing available on a GigaVUE-HC3 module, or GigaVUE-HC1 node elsewhere in the same cluster.

Can a GigaStream Act as a Shared Collector?

Yes.

Multiple individual ports for a Shared Collector can be setup or a GigaStream tool group. Refer to About Shared Collectors for details.

What Are the GigaStream Maximums?

The number of GigaStream per line card, module, or node varies by product. Refer to GigaStream Rules and Maximums for the details.

DoesGigaVUE-FM Provide the Same Features as the GigaVUE-OS CLI?

The GigaVUE-FM provides all of the mapping features of the CLI in an intuitive and highly-usable setting. The map creation wizard speeds map creation.

What order are the map rules displayed in "show running config"?

You can view the rules within a map, displayed in the order created.

Active Visibility

Active Visibility is a framework that allows your visibility network to adapt to dynamic events. The framework is designed to react to evenfts and take actions in response to events in your visibility network.

An Active Visibility policy defines conditions and actions. When conditions are met, actions are executed. The policy specifies both the conditions and the actions and ties them together.

The conditions and actions are pre-defined. Conditions are events that are used to trigger changes to configuration. Actions can notify users of certain events and/or modify the configuration in response to certain events. Conditions can be port-based or time-based.

For example, if a tool port is overloaded, you might want to reduce the traffic sent to the tool. You can configure two maps, one targeted for your high priority traffic and the other targeted for your low priority traffic. Then you can specify conditions to monitor tool port utilization. When traffic is below a threshold, both maps can be enabled, thus all traffic will be sent to the tools. But when traffic is above the threshold, you can specify an action so that only the map targeted for the high priority traffic is enabled and thus only that traffic will be sent to the tool.

Another example is that you might want to use a different set of maps and map rules to provide visibility during different times, such as during working hours, or on weekends. If there is going to be a backup on the weekend, you can specify a policy to disable a map at a specific time or day.

NOTE: GigaVUE-FM GUI for Active Visibility is available only until software version 5.4.00. Refer to the *"Configuring Active Visibility"* section in the *GigaVUE-OS CLI Reference Guide* for details on further software enhancements of this feature.

Active Visibility—Rules and Notes

Keep in mind the following rules and notes when working with Active Visibility:

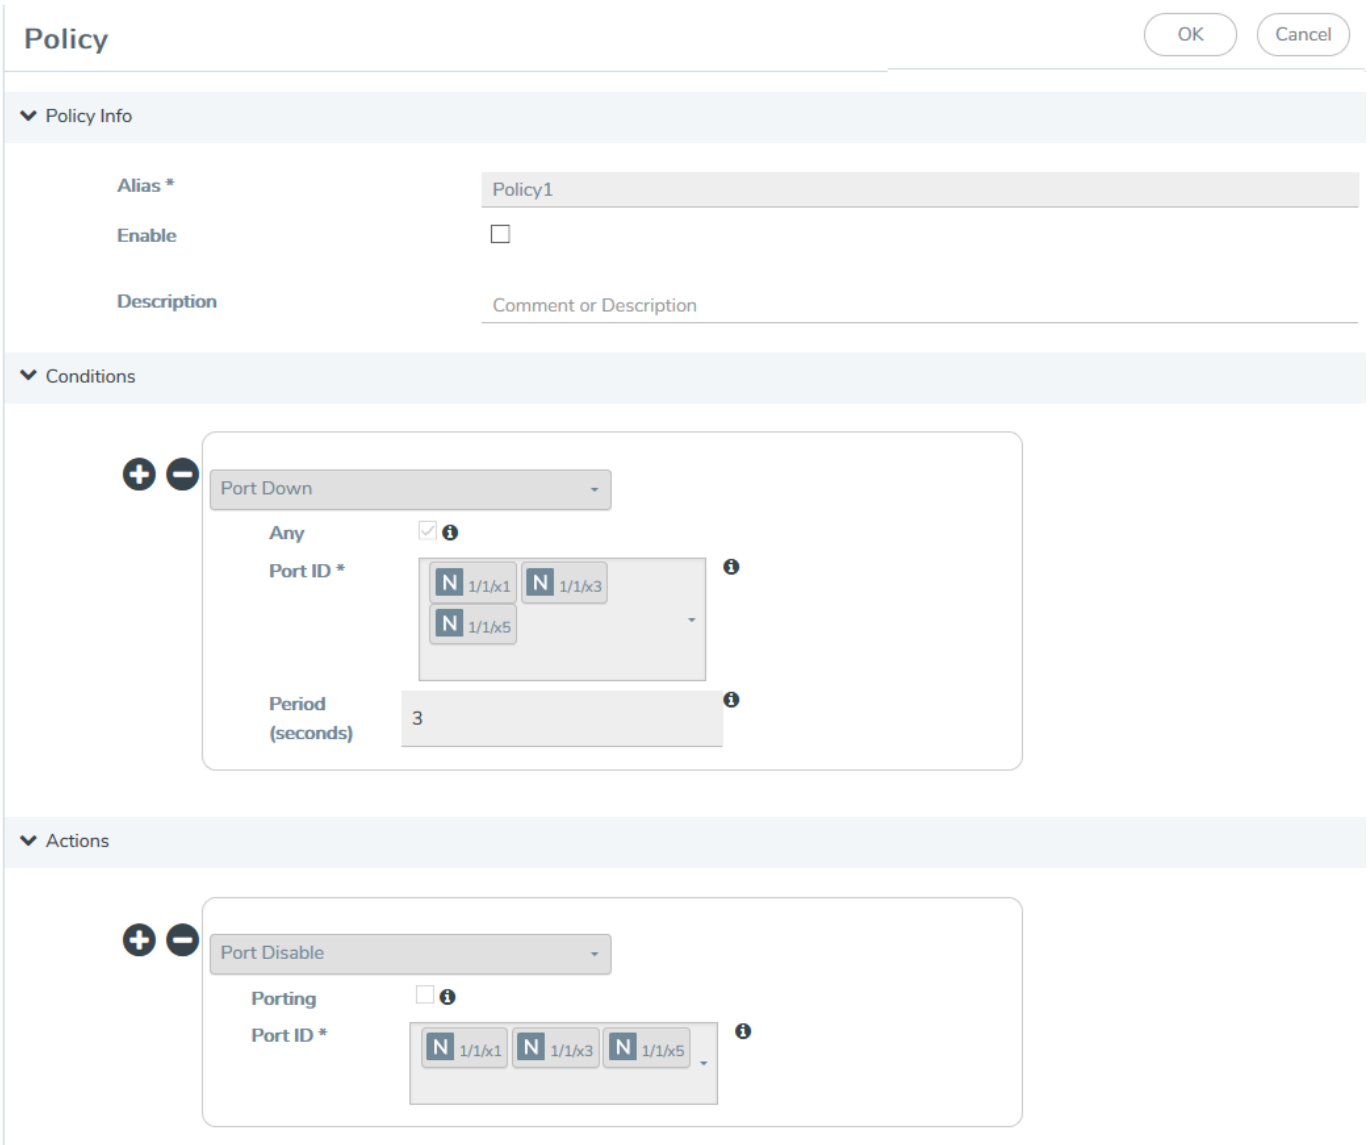
- A policy must have at least one condition. An action is not required.
- Up to five (5) conditions can be specified in a policy. The parameters and values that are specified in the condition depends on the condition. The policy is executed only when all conditions are met.
- Up to five (5) actions can be specified in a policy. The parameters and values that are specified in the action depends on the action.
- Within a policy, there is one unique condition. For example, there can only be one **PortUp** condition, not multiple **PortUp** conditions within a policy.
- Within a policy, there can be multiple actions, including the same action. For example, there can be multiple **PortEnable** actions within a policy.
- A policy is triggered if all the conditions are met. Then all the actions are executed. If there are five conditions in the policy, they all have to be met before the action or actions are executed. If there are multiple actions, they are executed in sequence, as specified in the policy. The policy specifies the priority of the actions.
- When there are multiple actions, they will continue to be executed even if there is an error. For example, if there are three actions and there is an error with the second action, the first and third actions will be executed. When the policy executes, each action will have their status reported (success or failure).
- Multiple policies can be in effect at the same time. The policies can monitor different conditions and take different actions. Up to 100 policies can be defined for a cluster.
- A policy must be enabled for it to become active. However, when you first create a policy, you might want to disable it so that it does not become active right away.
- When a policy is triggered, an SNMP event and email notification (policytrigger) can optionally be generated.
- When you want to delete a map that is attached to a policy, you must first delete the policy and then delete the map. If you delete the map without deleting the policy, there will be configuration synchronization issues in GigaVUE-FM. In such a case, delete the policy using GigaVUE-OS CLI. Refer to the *"policy"* command in the *GigaVUE-OS CLI Reference Guide*.

Configure Policies

To configure a policy:

1. Go to **Traffic** 
2. On the left navigation pane, go to **Physical > Active Visibility,** and then select the required cluster or node ID.
3. In the Policies page, click **New**.
4. Enter an alias and description for the policy.

5. When you first define a policy, you will probably want the policy to be disabled, so clear the **Enable** check box.
6. Select the required condition(s) from the drop-down list. The parameters and values are populated based on the condition(s) you select. For example, if you select the condition, PortTxUtilLow, there is an **Any** checkbox and fields for **Port ID**, **Threshold (%)**, and **Period (seconds)**. To view the list of pre-defined conditions along with their descriptions and comments, refer to View List of Conditions and Actions.
7. Select the required action(s) from the drop-down list. The parameters and values are populated based on the action(s) you select. For example, if you select the action, Port Disable, there is a **Porting** checkbox and a field for **Port ID**. To view the list of pre-defined actions along with their descriptions and comments, refer to View List of Conditions and Actions.
8. When you have added the required conditions and actions, click **OK**.



The newly configured policy is displayed in the Policies page. When a policy has not been executed, it will show a status of NOT EXECUTED. When a policy has been executed, it will show a status of SUCCESS or FAILURE.

Following are the additional actions that you can perform in the Policies page:

- Enable a policy—When you are ready to enable a policy, select the policy, and then go to **Actions > Enable**.
- Edit a policy—Select the policy and then click **Edit**. Make the required changes, and click **OK**.
- Clone a policy—Select the policy and click **Clone**. Change the policy alias, make any changes to conditions or actions, and click **OK**.
- Delete a policy—Select the policy and click **Delete**. A confirmation displays. To confirm, click **OK**.
- View policy details—Click the policy alias to view the details.
- View policy report—To view information on a policy that has been executed, click **Report History**. The policy report history is displayed.

View List of Conditions and Actions

To view the list of pre-defined conditions and actions along with their descriptions and comments:

1. Go to **Traffic** 
2. On the left navigation pane, go to **Physical > Active Visibility,** and then select the required cluster or node ID. You can view the policies details.
3. On the top, select **Help** > **Conditions** to view the list of pre-defined conditions. Select **Help** > **Actions** to view the actions that need to be performed.

MAC Address Rewrite

Media Access Control (MAC) address rewrite converts the incoming traffic's MAC address (source , destination, or both) with a user configured MAC address. The modified packets are then delivered as per flow mapping configurations. This allows the user to maintain confidentiality of the outgoing MAC address.

MAC address rewrite can be enabled in two ways:

- **Rule based**- The MAC address rewrite functionality is enabled for traffic that qualifies a specific rule in a map. This can be enabled only for pass rules. Rule based MAC address re-write allows modifying the rule, source, and destination MAC address.

- **Map Based**- The MAC address rewrite functionality is enabled for traffic that qualifies any of the rules configured in regular by-Rule maps and shared collectors. The configuration applies to all the rules that are part of the map except for drop rules. Map based MAC address re-write allows modifying the source and destination MAC address and can also be applied to a deployed map. Refer to Map MAC Address Source and Destination Compatibility Matrix for more information.

Table 12: Map MAC Address Source and Destination Compatibility Matrix

Source	Destination	Supported
Network	Tool/Hybrid, Tool GigaStream/Hybrid GigaStream, Tool with Egress VLAN strip/Tool with Egress Port filters.	Yes
Network	L2 Circuit Encapsulation Tunneling	No
Hybrid	Tool/Hybrid, Tool GigaStream/Hybrid GigaStream, Tool with Egress VLAN strip/Tool with Egress Port filters.	Yes
Network /Hybrid	Port-group(without smart-lb enabled).	Yes
IP interface (decapsulation Tunnel)	Tool/Hybrid	Yes
L2GRE/VXLAN	L2GRE/VXLAN encapsulation tunnel.	No
VXLAN Header/MPLS Header stripping	Tool/Hybrid	No
Network Port with Ingress VLAN tag	Tool/Hybrid	Yes
L2-Circuit Tunnel	Tool/Hybrid/GigaStream	Yes
VXLAN/L2GRE Tunnel decapsulation with IP interface	Tool/Hybrid/GigaStream	No
Port-Group	Tool/Hybrid/GigaStream	Yes

NOTE: If you have configured both map level and rule level MAC address rewrite functionality in the same map, then rule-based configuration takes priority.

Configuring MAC Address Re-write

Media Access Control address is a six byte hardware identification field with 12 hexadecimal digits that uniquely identifies a device in the network. You can rewrite the MAC source and destination fields to configurable MAC address as follows:

1. To enable MAC address rewrite functionality through GigaVUE-FM:
   a. **Map based Configuration**- To configure MAC address rewrite based on maps follow the below steps:
      - Navigate to  **> Physical > Nodes.**
      - Select the required cluster or device. Navigate to **Maps** and click create **New Map**. Scroll down to **Map Configuration & Rules**.

- Under **Configuration**, enable the **'Address Rewrite'** checkbox.
- Select either MAC Source, Mac Destination, or both.
- Specify the MAC Source and Destination.
- Click on **OK** to complete the configuration.

b. **Rule based Configuration**- To configure MAC address rewrite based on map rules follow the below steps

- Navigate to  **> Physical > Nodes.**
- Select the required cluster or device. Navigate to **Maps** and click create **New Map**. Scroll down to **Map Configuration & Rules**.
- Under **Map Rules**, click **Add a Rule**.
- Enable the **'Address Rewrite'** checkbox.
- From Map Rules section select either MAC Source , Mac Destination, or both.
- Specify the MAC Source and Destination.
- Click on **OK** to complete the configuration.

2. To enable MAC address rewrite through GigaVUE-OS -CLI enter the map prefix mode with the command **config map alias<map>** and then enter any one of the following commands such as:

```
rewrite-dstmac xx:xx:xx:xx:xx:xx
rewrite-srcmac xx:xx:xx:xx:xx:xx
no rewrite-dstmac
no rewrite-srcmac
```

Refer to GigaVUE-OS CLI Reference Guide for more information.

License

You do not need a license to enable this feature for GigaVUE HC Series. To enable this feature for GigaVUE TA Series ensure you have Advanced Features License.

Limitations

The following are the limitations of MAC Address rewrite.

- Pass-all maps are not supported.
- GSOP enabled maps are not supported.
- VXLAN/L2GRE decapsulation and encapsulation tunnels are not supported
- Inline, Flex Inline maps and OOB copy maps are not supported.
- First level, second level and transit maps are not supported.

- This feature is not supported with Fabric Maps, L2 Circuit Tunnel Encapsulation, MPLS and VXLAN Header Stripping enabled-port configurations.
- A paired port receives rewritten mac address when creating a port-pair with a network port in map/rule-based mac-rewrite byrule map.

IP Address Rewrite

Internet Protocol (IP) address rewrite converts the incoming traffic's IP address (source, destination, or both) with a user configured IP address. The modified packets are then delivered as per flow mapping configurations. This allows the user to maintain the confidentiality of the outgoing IP address.

IP address rewrite can be enabled in two ways:

- **Rule based**- The IP address rewrite functionality is enabled for traffic that qualifies a specific rule in a map. This can be enabled only for pass rules. Rule based IP address re-write allows modifying the rule, source, and destination IP address.
- **Map Based**- The IP address rewrite functionality is enabled for traffic that qualifies any of the rules configured in regular by-Rule maps and shared collectors. The configuration applies to all the rules that are part of the map except for drop rules. Map based IP address re-write allows modifying the source and destination IP address and can also be applied to a deployed map. Refer to Map IP Address Source and Destination Compatibility Matrix for more information.

Table 13: Map IP Address Source and Destination Compatibility Matrix

Source	Destination	Supported
Network	Tool/Hybrid, Tool GigaStream/Hybrid GigaStream, Tool with Egress VLAN strip/Tool with Egress Port filters.	Yes
Hybrid	Tool/Hybrid, Tool GigaStream/Hybrid GigaStream, Tool with Egress VLAN strip/Tool with Egress Port filters.	Yes
Network with L2GRE/VXLAN enabled	L2GRE/VXLAN Encapsulation Tunnel	No
Network with VXLAN Header Strip enabled Port/MPLS Header Strip	Tool/Hybrid	No
Network Port with Ingress VLAN tag	Tool/Hybrid	Yes
Network	L2 Circuit Encapsulation Tunneling	No
L2-Circuit Tunnel	Tool/Hybrid/GigaStream	Yes
VXLAN/L2GRE Tunnel decapsulation with IP interface	Tool/Hybrid/GigaStream	No
Port-Group	Tool/Hybrid/GigaStream	Yes
Network /Hybrid	Port-group (without smart-lb enabled).	Yes

NOTE: If you have configured both map level and rule level IP address rewrite functionality in the same map, then rule-based configuration takes priority.

Configuring IP Address Re-write

Internet Protocol address is a four-byte hardware identification field with 8 hexadecimal digits that uniquely identifies a device in the network. You can rewrite the IP source and destination fields to configurable IP address as follows:

1. To enable IP address rewrite functionality through GigaVUE-FM:
   a. **Map based Configuration**- To configure IP address rewrite based on maps follow the below steps:
      - Navigate to  **> Physical > Nodes.**
      - Select the required cluster or device. Navigate to **Maps** and click create **New Map**. Scroll down to **Map Configuration & Rules**.
      - Under **Configuration**, enable the **'Address Rewrite'** checkbox.
      - Select either IPv4 Source, IPv4 Destination, or both.
      - Specify the IPv4 Source and Destination.
      - Click on **OK** to complete the configuration.
   b. **Rule based Configuration**- To configure IP address rewrite based on map rules follow the below steps
      - Navigate to  **> Physical > Nodes.**
      - Select the required cluster or device. Navigate to **Maps** and click create **New Map**. Scroll down to **Map Configuration & Rules**.
      - Under **Map Rules**, click **Add a Rule**.
      - Select the IPV4 address from the Address Rewrite drop down list.
      - From Map Rules section select either IPv4 Source , IPv4 Destination, or both.
      - Specify the IPv4 Source and Destination.
      - Click on **OK** to complete the configuration.
2. To enable IP address rewrite through GigaVUE-OS -CLI enter the map prefix mode with the command **config map alias<map>** and then enter any one of the following commands such as:

```
rewrite-dstip x.x.x.x
rewrite-srcip x.x.x.x
no rewrite-dstip
no rewrite-srcip
```

Refer to GigaVUE-OS CLI Reference Guide for more information.

License

You do not need a license to enable this feature for GigaVUE® HC Series. To enable this feature for GigaVUE® TA Series ensure you have Advanced Features License.

Limitations

The following are the limitations of IP Address rewrite.

- Pass-all maps are not supported.
- GSOP enabled maps are not supported.
- VXLAN/L2GRE Encapsulation and decapsulation tunnels are not supported
- Inline, Flex Inline maps and OOB copy maps are not supported.
- First-level, second-level, and transit maps are not supported.
- This feature is not supported with Fabric Maps, L2 Circuit Tunnel Encapsulation, MPLS, and VXLAN Header Stripping enabled-port configurations.
- A paired port receives rewritten IP address when creating a port-pair with a network port in map/rule-based ip-rewrite byrule map.
- IP rewrite is not supported in GigaVUE-TA400 and GigaVUE-HCT devices.
- IP rewrite with single tagged traffic and ingress VLAN tag is not supported in GigaVUE-HC1-Plus,GigaVUE-TA25, and GigaVUE-TA25E.
- IP rewrite with IPv6 and L3-mpls traffic is not supported.
- Cluster and double-tagged traffic are not supported.
- When a pass-all map shares the same source ports with a by-rule map with IP rewrite enabled, the IP rewrite effect will be applied to the destination port of both the pass-all map and the by-rule map.

VLAN Manipulation

Required License:

- GigaVUE HC Series - Base License
- GigaVUE TA Series - Advanced Feature License

Overview

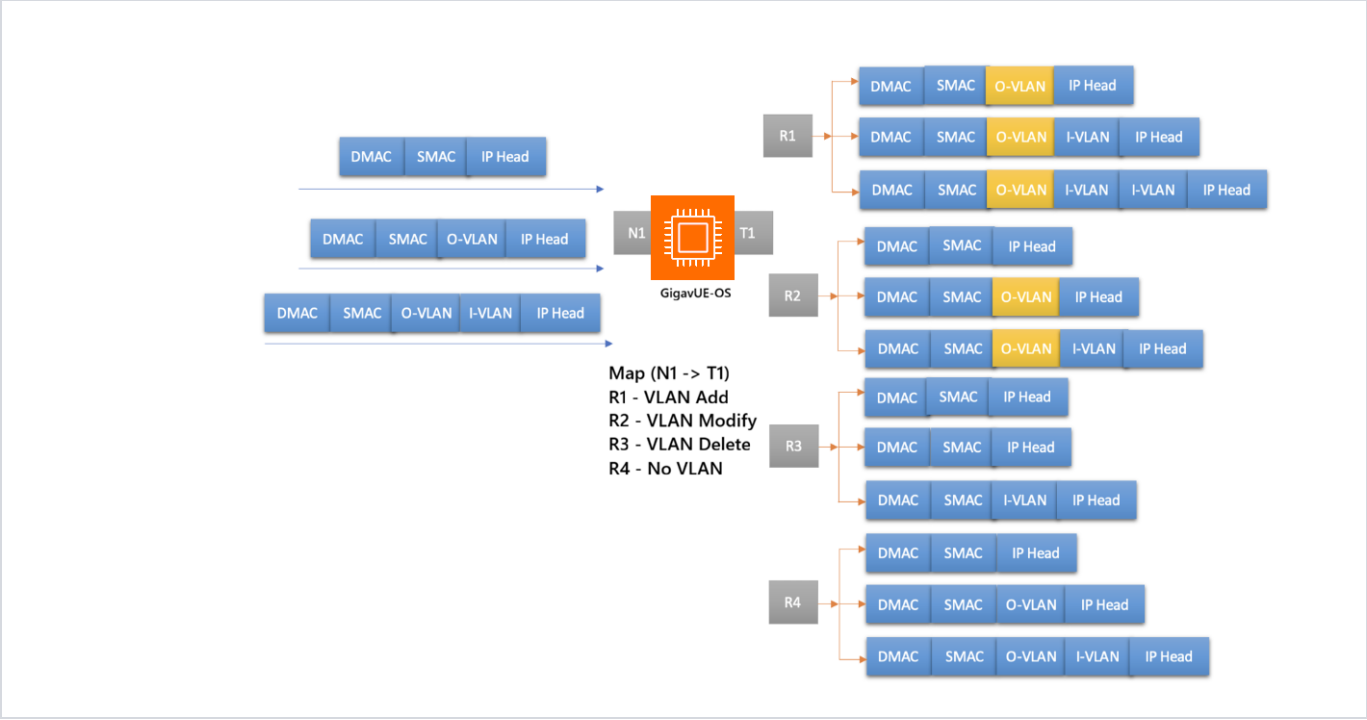
The VLAN manipulation adds a new VLAN tag, modifies/strips the existing outer VLAN tag based on the user configured VLAN manipulation operation. The modified packets are then delivered according to the Flow Mapping® configurations. This process ensures that the confidentiality of outgoing traffic is maintained while allowing efficient traffic segmentation.

NOTE: From 6.8 Release, VLAN manipulation add is supported in standalone mode. From 6.10 Release, VLAN manipulation add is supported in Cluster mode.

NOTE: Starting from 6.13 Release, VLAN manipulation Modify, VLAN manipulation Strip is supported only in standalone mode. It is not supported in cluster mode.

VLAN manipulation can be configured based on map rules and maps:

- **Rule-Based**- In this method, the VLAN manipulation is configured for traffic that matches a specific rule on a map. This can be configured only for pass rules. Rule-based VLAN manipulation adds a new VLAN tag, Modifies/Strips the existing outer VLAN tag based on the user configured VLAN manipulation operation. The figure below illustrates how the Outer VLAN (O-VLAN) can be added to incoming traffic using the VLAN Add operation.



Rule-Level Add VLAN Manipulation

To add VLAN manipulation on rule-level,

1. On the left navigation pane, go to > **Physical > Nodes.**
2. Select the required cluster or device.
3. Go to **Maps** and click **New Map** to create the new map.

NOTE: By default, a Regular ByRule Map will be created.

4. In the **Map Info** dialog, do the following:
   - **Map Alias** - Enter the alias for the new map.
   - (Optional) **Description** - Enter the description for the new map.
   - **Subtype** - Select the **By Rule** from the list.
5. In the **Map Source and Destination** section, do the following:
   - Select the **Source** port from the list.
   - (Optional) Select the **destination** port from the list.
6. In the **Map Configuration & Rules** section, go to the Map Rules section do the following:
   - Click **Add a Rule**.
   - (Optional) **Rule Description** - Enter a rule description.
   - From the **VLAN Action** field, select **Add**.
   - In the **VLAN ID** field, enter a **VLAN ID** value between 1 and 4095.
   - From the **Tag Protocol Id** drop-down list, select the **TPID** value for the VLAN Tag. The default value is 0x8100, but you can also select the other supported values 0x9100 and 0x88a8 from the drop-down list.
   - **Condition** - Select the **condition** from list, and choose **Pass**. Without a condition, the rule will throw an error.
7. Click **OK** to complete the configuration.

To configure VLAN manipulation using GigaVUE-OS CLI, refer to the *"Configure VLAN Manipulation"* section in the GigaVUE-OS CLI Reference Guide.

Rule-Level Modify VLAN Manipulation

To modify VLAN manipulation on rule-level,

1. On the left navigation pane, go to > **Physical > Nodes.**
2. Select the required cluster or device.
3. Go to **Maps** and click **New Map** to create the new map.
4. In the **Map Info** dialog, do the following:

- **Map Alias** - Enter the alias for the new map.
- (Optional) **Description** - Enter the description for the new map.
- **Subtype** - Select the **By Rule** from the list.

5. In the **Map Source and Destination** section, do the following:
   - Select the **Source** port from the list.
   - (Optional) Select the **destination** port from the list.
6. In the **Map Configuration & Rules** section, go to the Map Rules section do the following:
   - Click **Add a Rule**.
   - (Optional) **Rule Description** - Enter a rule description.
   - From the **VLAN Action** field, select **Modify**.
   - In the **VLAN ID** field, enter a **VLAN ID** value between 1 and 4095.
   - From the **Tag Protocol Id** drop-down list, select the **TPID** value for the VLAN Tag. The default value is 0x8100, but you can also select the other supported values 0x9100 and 0x88a8 from the drop-down list.
   - **Condition** - Select the **condition** from list, and choose **Pass**. Without a condition, the rule will throw an error.
7. Click **OK**.

Rule-Level Strip VLAN Manipulation

To delete VLAN manipulation on rule-level,

1. On the left navigation pane, go to  **> Physical > Nodes.**
2. Select the required cluster or device.
3. Go to **Maps** and click **New Map** to create the new map.
4. In the **Map Info** dialog, do the following:
   - **Map Alias** - Enter the alias for the new map.
   - (Optional) **Description** - Enter the description for the new map.
   - **Subtype** - Select the **By Rule** from the list.
5. In the **Map Source and Destination** section, do the following:
   - Select the **Source** port from the list.
   - (Optional) Select the **destination** port from the list.
6. In the **Map Configuration & Rules** section, go to the **Map Rules** section do the following:

- Click **Add a Rule**.
- (Optional) **Rule Description** - Enter a rule description.
- From the **VLAN Action** field, select **Strip**.
- **Condition** - Select the **condition** from list, and choose **Pass**. Without a condition, the rule will throw an error.

7. Click **OK** to complete the configuration.

- **Map-Based**- In this method, VLAN manipulation is configured for traffic that qualifies under rules defined in regular by-rule maps and shared collectors. This configuration applies to all rules within the map except for drop rules. Map-based VLAN manipulation adds a new VLAN tag, Modifies/Strips the existing outer VLAN tag based on the user configured VLAN manipulation operation. For more information, refer to the Map VLAN manipulation Source and Destination Compatibility Matrix .

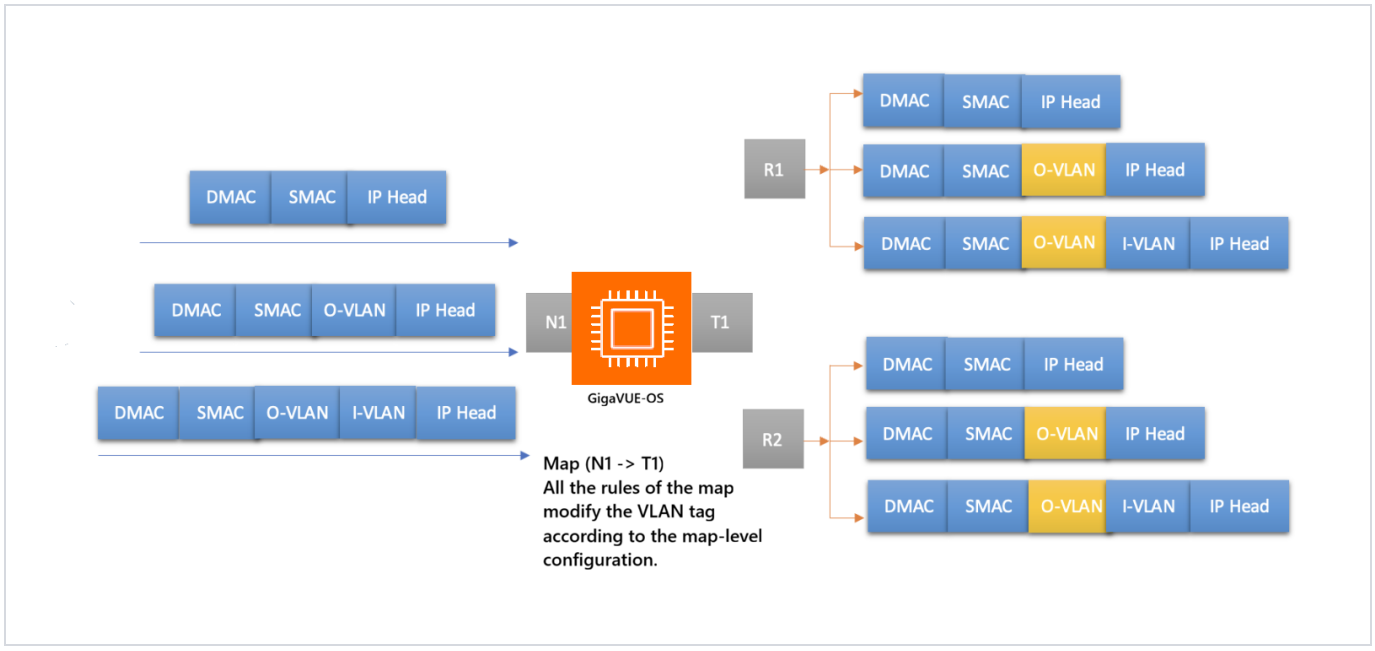
NOTE: If you have configured both map level and rule level VLAN manipulation functionality in the same map, then rule-based configuration takes priority.

Map-Level VLAN Manipulation

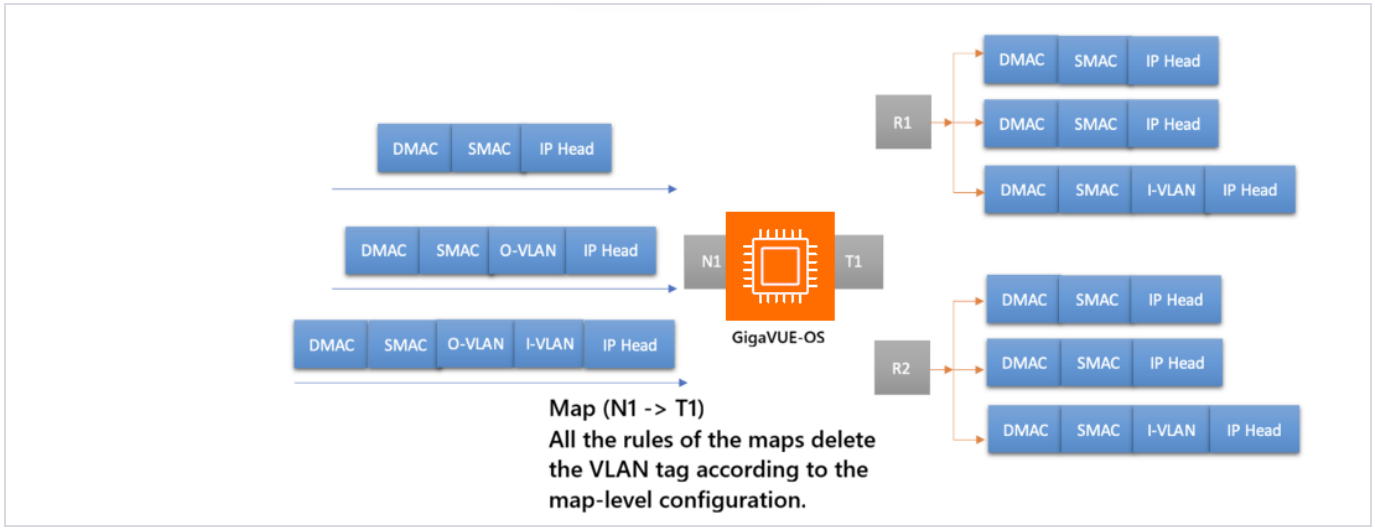
- **Map-Level Add VLAN Manipulation**

The figure below illustrates how the Outer VLAN (O-VLAN) can be added to the rules of the incoming traffic using the VLAN Add operation within the map-based configuration.

- **Map-Level Modify VLAN Manipulation**



- **Map-Level Strip VLAN Manipulation**



To configure VLAN manipulation on map-level,

1. On the left navigation pane, go to > **Physical** > **Nodes**.
2. Select the required cluster or device.
3. Go to **Maps** and click **New Map** to create a new map.

NOTE: By default, a Regular ByRule Map will be created.

4. In the **Map Info** dialog, do the following:
   - **Map Alias** - Enter the alias for the new map.

- (Optional) **Description** - Enter the description for the new map.
- **Subtype** - Select the **By Rule** from the list. By default, the subtype will be By Rule.

5. In the **Map Source and Destination** section, do the following:
   - Select the **Source** port from the list.
   - (Optional) Select the **Destination** port from the list.
6. In the **Map Configuration & Rules** section, go to the **Configurations** section do the following:

   Add:
   - From the **VLAN Action** field, select **Add**.
   - In the **VLAN ID** field, enter a **VLAN ID** value between 1 and 4095.
   - From the **Tag Protocol Id** drop-down list, select the **TPID** value for the VLAN Tag. The default value is 0x8100, but you can also select the other supported values 0x9100 and 0x88a8 from the drop-down list.

   Modify:
   - From the **VLAN Action** field, select **Modify**.
   - In the **VLAN ID** field, enter a **VLAN ID** value between 1 and 4095.
   - From the **Tag Protocol Id** drop-down list, select the **TPID** value for the VLAN Tag. The default value is 0x8100, but you can also select the other supported values 0x9100 and 0x88a8 from the drop-down list.

   Strip:
   - Select **Strip**, from the **VLAN Action** list.

   NOTE: For Strip, there is no need for a VLAN ID, and TPID is not required.
7. Once you add the map rule configuration, do the following:
   - Click **Add a Rule**.
   - In the **Condition**, select **IPv4 Source** from the list.
   - Enter the **IPv4 Address**, and **Cidr**, or **Net Mask** values.
8. Click **OK** to complete the map configuration.

Table 14: Map VLAN manipulation Source and Destination Compatibility Matrix

Source	Destination	Supported
Network	Tool/Hybrid, Tool GigaStream/Hybrid GigaStream, Tool with egress VLAN strip/Tool with egress Port filters.	Yes
Hybrid	Tool/Hybrid, Tool GigaStream/Hybrid GigaStream, Tool with egress VLAN strip/Tool with egress Port filters.	Yes

Source	Destination	Supported
Network/Hybrid Port with ingress VLAN tag	Tool/Hybrid	Yes
Port-Group	Tool/Hybrid/GigaStream	Yes
Network /Hybrid	Port-group (without smart-lb enabled).	Yes

Limitations

The following are the limitations of VLAN manipulation.

- Pass-all maps are not supported.
- GSOP-enabled maps are not supported.
- VXLAN/L2GRE Encapsulation and Decapsulation Tunnels are not supported.
- Inline, Flex Inline maps, and OOB copy maps are not supported.
- First-level, second-level, and transit maps are not supported.
- Fabric Maps, L2 Circuit Encapsulation and Decapsulation Tunnels are not supported.
- MPLS, and VXLAN Header Stripping enabled-port configurations do not support this feature.
- VLAN manipulation is not supported in GigaVUE-HC3 ccv1 device.
- VLAN manipulation with IP rewrite is not supported.
- Port filter with VLAN Qualifier is not supported.
- When VLAN manipulation with ingress VLAN Tag is configured, VLAN Manipulation will take higher precedence.
- When VLAN manipulation with egress VLAN Strip is configured, VLAN Manipulation will take higher precedence.
- If advanced VLAN manipulation is configured on either regular by-rule maps or shared collector maps, then both Passall maps or Port Pair should not use the same network ports as those deployed in regular and collector maps. Similarly, if a collector map or regular map is configured on the same network ports as a Passall map or Port Pair, VLAN manipulation should not be configured. Configuring VLAN manipulation in the above two scenarios may result in traffic discards on the destination ports of the Passall map.

Monitor Port Utilization

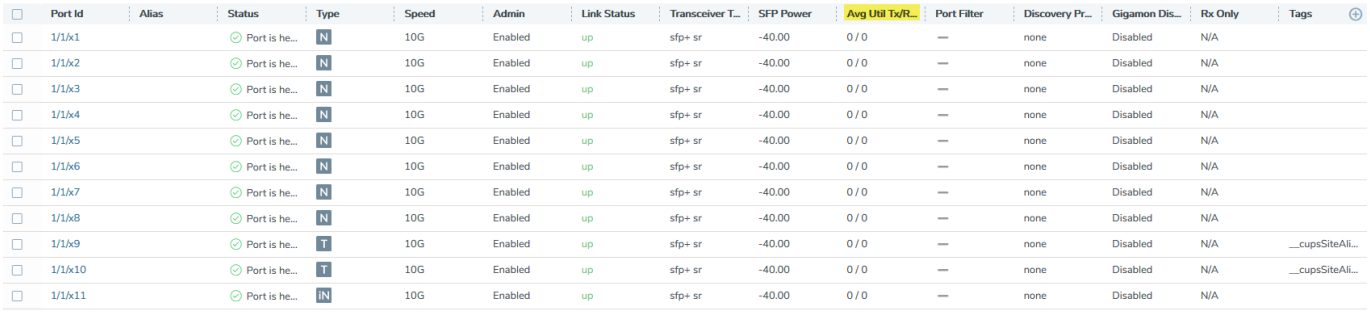
This section describes how to monitor port utilization and buffer thresholds on the GigaVUE HC Series and GigaVUE TA Series nodes. Refer to the following sections for details:

- Port Utilization Availability by Port Type
- Set Port Utilization Thresholds
- Configure Alarm Buffer Thresholds

- Set Alarm Buffer Thresholds
- Microburst

Port Utilization Availability by Port Type

You can view port utilization for all network, tool, hybrid, and stack link ports on the GigaVUE HC Series or GigaVUE TA Series nodes. From the device view, go to **Ports > All Ports** and view the Utilization (Tx/Rx) column in the table.



Port Id	Alias	Status	Type	Speed	Admin	Link Status	Transceiver T...	SFP Power	Avg Util Tx/R...	Port Filter	Discovery Pr...	Gigamon Dis...	Rx Only	Tags
1/1/x1		Port is he...	N	10G	Enabled	up	sfp+ sr	-40.00	0 / 0	—	none	Disabled	N/A	
1/1/x2		Port is he...	N	10G	Enabled	up	sfp+ sr	-40.00	0 / 0	—	none	Disabled	N/A	
1/1/x3		Port is he...	N	10G	Enabled	up	sfp+ sr	-40.00	0 / 0	—	none	Disabled	N/A	
1/1/x4		Port is he...	N	10G	Enabled	up	sfp+ sr	-40.00	0 / 0	—	none	Disabled	N/A	
1/1/x5		Port is he...	N	10G	Enabled	up	sfp+ sr	-40.00	0 / 0	—	none	Disabled	N/A	
1/1/x6		Port is he...	N	10G	Enabled	up	sfp+ sr	-40.00	0 / 0	—	none	Disabled	N/A	
1/1/x7		Port is he...	N	10G	Enabled	up	sfp+ sr	-40.00	0 / 0	—	none	Disabled	N/A	
1/1/x8		Port is he...	N	10G	Enabled	up	sfp+ sr	-40.00	0 / 0	—	none	Disabled	N/A	
1/1/x9		Port is he...	T	10G	Enabled	up	sfp+ sr	-40.00	0 / 0	—	none	Disabled	N/A	__cupsSiteAli...
1/1/x10		Port is he...	T	10G	Enabled	up	sfp+ sr	-40.00	0 / 0	—	none	Disabled	N/A	__cupsSiteAli...
1/1/x11		Port is he...	N	10G	Enabled	up	sfp+ sr	-40.00	0 / 0	—	none	Disabled	N/A	

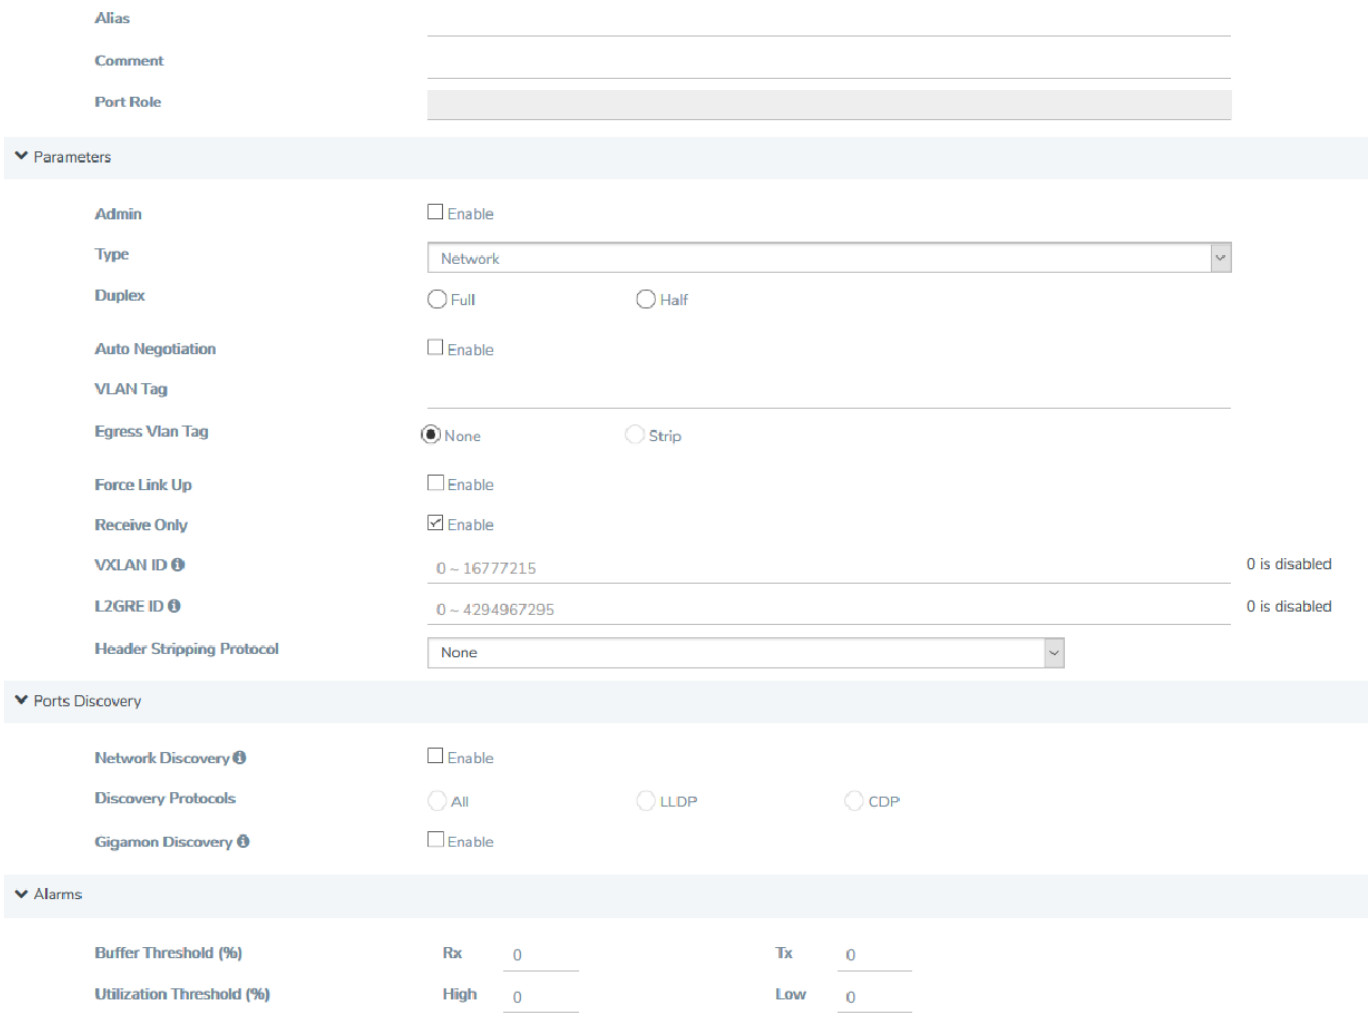
It is the utilization for all requested ports with the port number, port type, port speed, receive (Rx) utilization percentage (network and stack ports), transmit (Tx) utilization percentage (tool, hybrid, and stack ports), alarm threshold, and the last time the threshold was exceeded on either the transmit or receive channel.

Set Port Utilization Thresholds

To set the Alarms for port utilization, do the following:

1. Select **Ports > Ports > All Ports**. Select Port ID of the port on which you want to set the utilization threshold.
2. Click **Edit** to open the port editor.
3. Under **Alarms**, in the **Utilization Threshold** field, enter the percentage at which the GigaVUE HC Series node logs an alarm for the port. By default, the thresholds are 0, which means disabled.

   NOTE: Network ports always use an Rx threshold. Tool ports always use Tx. Stack ports and hybrid ports use both Rx and Tx, and the same threshold is used for each. GigaSMART engine ports use high and low utilization threshold percentage.



Utilization Alarm/SNMP Trap Generation

The GigaVUE HC Series or GigaVUE TA Series node generates a utilization alarm each time the configured threshold is exceeded for more than six consecutive seconds. Once the percentage utilization falls below the configured threshold for at least six consecutive seconds, a second alarm is generated to indicate that utilization has returned to normal. Once utilization has returned to normal (six consecutive seconds below the threshold), a new utilization alarm can be generated once the measured rate again remains above the threshold for six consecutive seconds.

Utilization alarms are written to syslog and forwarded to all SNMP management stations configured as notification destinations. For SNMP traps to be generated, forwarded, and

displayed correctly in your SNMP management station, you must configure SNMP server and notification destinations, enable SNMP notifications and events. Refer to the *"Use SNMP"* section in the *GigaVUE Administration Guide*.

To generate an SNMP trap, you must first enable the required event for SNMP notifications either on the device or on the GigaVUE-FM instance that manages the device. For instructions about how to enable SNMP notification on a device, refer to the *"Enable or Disable Events for SNMP Notifications"* section in the *GigaVUE Administration Guide*. For instructions about how to enable SNMP notifications on the GigaVUE-FM that manages the device, refer to the *"SNMP Traps"* section in the *GigaVUE Administration Guide*.

Configure Alarm Buffer Thresholds

Often network ports are utilized at rates below 50%. If several network ports are aggregated, there is a risk of oversubscribing the tool ports. Alarm buffer thresholds are used to monitor the congestion within the GigaVUE node caused by microbursts or by oversubscription of tool ports.

The buffer usage on any port remains at zero until the maximum line rate of the port is reached. When the usage crosses 100% either instantaneously, in the microburst case, or prolonged, in the oversubscription case, there is congestion.

The internal buffer on the GigaVUE node can absorb a certain number of packet bursts. During congestion, packets are buffered in the chassis and the buffer usage is reported on the corresponding ports and in the corresponding direction: rx (ingress) and tx (egress).

Reporting the buffer usage provides a trend of how the microbursts are causing congestion, so more tool ports can be added before packets are dropped. Buffer usage is measured in intervals of 5 seconds. The peak buffer usage within a 5-second interval is reported.

When buffer usage is less than or equal to zero, there is no congestion, so no packets are dropped due to buffer unavailability.

When buffer usage is greater than zero, there is congestion. When buffer usage is greater than zero on any port in any direction, there is a chance that the packets (that caused the buffer usage to increase) are dropped due to unavailable buffers. However, it is unlikely to see packet drops due to buffer unavailability when the buffer usage on a port is less than 5%.

The buffer usage feature is supported on all ports and module types on the GigaVUE-HC3.

To configure buffer thresholds, refer to Set Alarm Buffer Thresholds.

Use the SNMP throttling functionality in GigaVUE-FM to reduce the flooding of SNMP traps. For details, refer to the *"SNMP Throttling"* section in the *GigaVUE Administration Guide*.

Set Alarm Buffer Thresholds

Use the Alarms section of the Ports configuration page to set rx (ingress) and tx (egress) alarm buffer threshold on a port and utilization threshold. You can specify the alarm buffer threshold in the rx and tx directions on network and stack type ports and in the tx direction on tool type ports. By default, the threshold is set to 0, which disables the threshold.

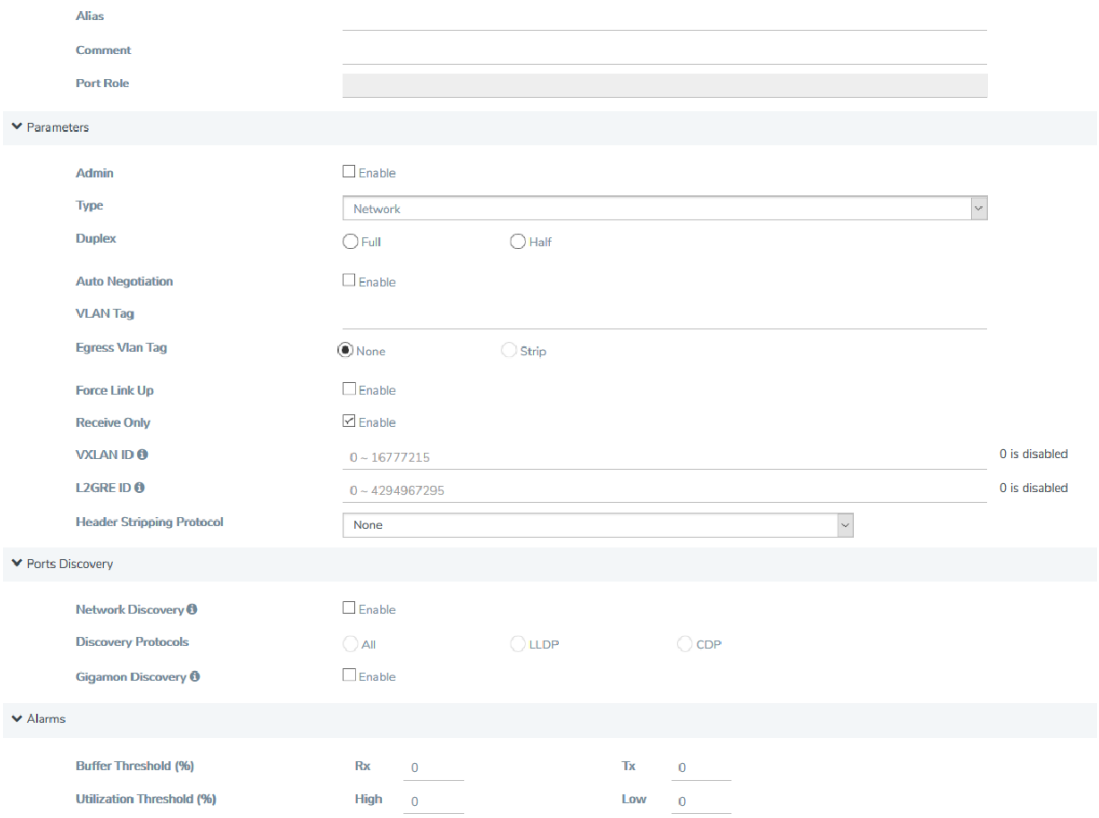
When a buffer usage threshold has exceeded its configured percentage, a message is logged, and optionally, an SNMP trap is sent to all configured destinations.

The SNMP trap will be sent when a threshold is exceeded in any 5-second interval. Once the trap is sent, there is a 30 second hold-off time before the trap is sent again.

For information about how to set SNMP traps, refer to the *"Use SNMP"* section in the *GigaVUE Administration Guide*.

To set the alarm buffer threshold and the usage thresholds on a port do the following:

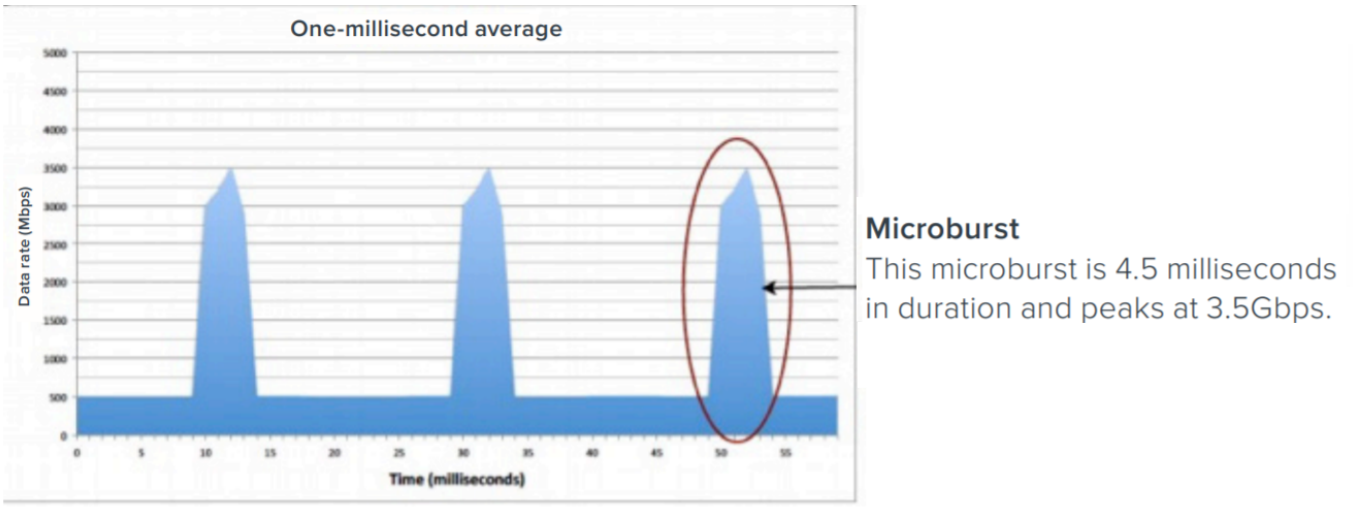
1. Select **Ports > Ports > All Ports**.Select the Port on the Ports page.
2. In the Alarms section of the Port configuration page, enter the **Rx** and **Tx** values for the **Buffer Threshold** and the **High** and **Low Values** for the **Utilization Threshold**.
3. Click **Save**.



Microburst

A microburst is a very intense traffic transmitted in a short period of time. It is a situation in which a large amount of burst data is received in milliseconds, where the burst data rate is much higher than the egress port's bandwidth.

For example, two 100Gb ingress ports transmit traffic to a 10Gb egress port at an average rate of 1Gbps each. The aggregate bandwidth utilized will be 2Gbps (1Gbps x 2 network ports), which is well within the bandwidth of the 10Gb egress port. However, if there is a sudden spike in traffic from one of the ingress ports, that is, traffic is transmitted at a rate of 3.5Gbps in 4.5 milliseconds (as shown in the figure below), it is referred to as a microburst.



When the traffic is aggregated, packets are received at the egress port, usually from multiple ingress ports, in parallel rather than one after another. This results in queuing of packets. GigaVUE-OS devices have some amount of buffer to queue these packets. Each port accesses this buffer dynamically, sharing this resource. This shared buffer is allocated to individual ports based on threshold values. If a port's shared buffer usage reaches its threshold value, the subsequent packets that the port receives will be dropped by the device due to unavailability of buffer resources.

During microbursts, the port utilization will appear to be low, but packets will be dropped and the `IfOutPktDrops` counter will continue to increment.

Following factors impact the packets drop rate:

- Intensity of the microburst traffic.
- Number of egress ports having microbursts at a given point in time.
- Buffer absorption capability of the GigaVUE-OS device.

Best Practices to Improve Burst Tolerance

Keep in mind the following best practices when you design your topology to improve burst tolerance:

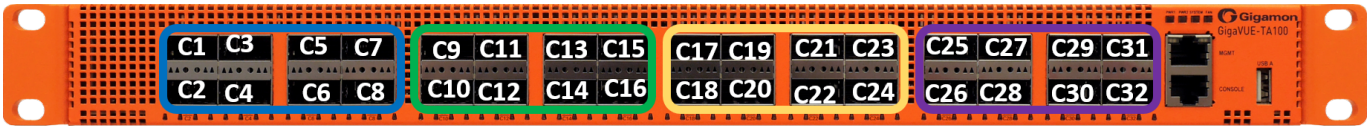
- If you have traffic flowing from a port with higher Gigabit to ports with lower Gigabit, that is, from a 40Gb port to four 10Gb ports, and bursty traffic is expected, one of the port can be internally modified into 4 x 25Gb and made as a hybrid port so the traffic flows from a 40Gb port to a 25Gb port, and then to a 10Gb port. However, you must create more than one map to implement this type of configuration. For more information about Hybrid ports, refer to Work with Hybrid Ports.
- Increase the tool ports bandwidth by adding multiple tool ports to tool GigaStream so that traffic is distributed across multiple ports, and thereby minimizes the risk of microburst. Use the Advanced Hashing feature to select the hashing criteria. For more information, refer to Advanced Hashing.
- Ensure that the egress ports are spread across different logical ports so that the load balancing and buffer utilization is effective. For example, you have a 40Gb egress port that is broken out in to four 10Gb ports, ensure that two of the 10Gb ports are in one logical group and the remaining two 10Gb ports are in another logical group. For details about the logical grouping of ports, refer to Logical Grouping of Ports in GigaVUE-HC3, GigaVUE-TA100, and GigaVUE-TA200.
- Ensure that both ingress and egress ports through which traffic is flowing into and out of the network must be in the same logical group. Following table provides an example of the ingress and egress ports mapping in GigaVUE-TA100. For details about the logical grouping of ports, refer to Logical Grouping of Ports in GigaVUE-HC3, GigaVUE-TA100, and GigaVUE-TA200

Egress Port	Ingress Port	Logical Group
C1	C5	Group 0
C9	C14	Group 1
C18	C23	Group 2
C25	C32	Group 4

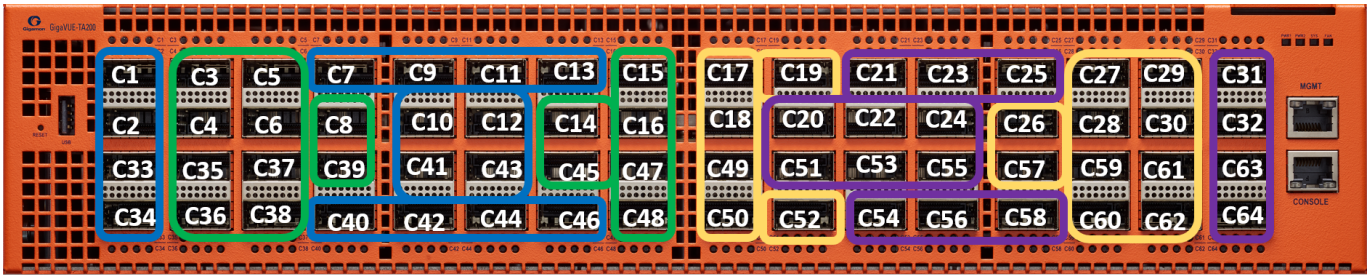
Logical Grouping of Ports in GigaVUE-HC3, GigaVUE-TA100, and GigaVUE-TA200

The ports in GigaVUE-HC3, GigaVUE-TA100, and GigaVUE-TA200 are grouped in to four logical groups—Group 0, Group 1, Group 2, and Group 3. You must ensure that the ingress and egress ports are spread across these logical groups to improve burst tolerance.

- GigaVUE-HC3—The ports in slot 1 are logically grouped into Group 0, ports in slot 2 into Group 1, ports in slot 3 into Group 2, and ports in slot 4 into Group 3.
- GigaVUE-TA100—Following diagram illustrates the logical grouping of ports in GigaVUE-TA100:



- GigaVUE-TA200—Following diagram illustrates the logical grouping in GigaVUE-TA200:



Legend

Logical Group 0

Logical Group 1

Logical Group 2

Logical Group 3

Flexible Inline Arrangements

This chapter provides an overview about the flexible inline arrangements, the supported platforms and software versions, the supported and unsupported functionalities, and limitations. It also provides details about how to configure the flexible inline maps and how to visualize the flexible inline arrangements canvas.

IMPORTANT: It is recommended to define inline configurations through GigaVUE-FM.

- If you configure a flexible inline arrangement solution using the GigaVUE-OS CLI, you cannot view or manage it using GigaVUE-FM.
- If you modify a flexible inline arrangement solution using the GigaVUE-OS CLI, you cannot view the changes in GigaVUE-FM.

NOTE: If your nodes and GigaVUE-FM instance are running software version 5.6.xx or earlier, ensure that you add all the nodes to GigaVUE-FM, upgrade the GigaVUE-FM to 5.8.xx, and then upgrade the nodes to 5.8.xx. This is to ensure that the flexible inline maps added to the nodes before the upgrade are visible in the canvas in GigaVUE-FM.Refer to the GigaVUE-OS Upgrade guide for details.

Refer to the following sections for details:

- Flexible Inline Arrangements
- Benefits of Flexible Inline Arrangements
- About Flexible Inline Maps
- Flexible Inline Arrangements—Rules and Notes
- Visualize the Flexible Inline Arrangements Canvas
- Configure Flexible Inline Flows
- Troubleshoot Flexible Inline Flows
- Backup and Restore Flexible Inline Flows

NOTE: Flexible inline arrangement is an advanced approach than the (Classic) Inline Bypass Solutions . Classic Inline Bypass functionality remains intact for backwards compatibility. For information on configuring classic Inline Bypass solutions , refer to Inline Bypass Solutions .

Flexible Inline Arrangements

Flexible inline arrangements are an approach to guide multiple inline traffic flows through a user-defined sequence of inline tools and inline tool groups. It uses the same software constructs as the existing Inline Bypass solution, such as inline network, inline tool, and inline tool group. Flexible inline arrangements support physical protection based on the specialized hardware on BPS modules. It also supports both protected and unprotected inline network links.

Flexible inline arrangements offer an alternative to classic Inline Bypass. Classic Inline Bypass functionality remains intact for backwards compatibility. For information on configuring Inline Bypass solutions (classic), refer to Inline Bypass Solutions.

Using flexible inline maps, traffic from the same inline network can traverse different sequences of inline tools and share tools across traffic flows or with other inline networks.

You can identify specific flows of traffic using Layer 2 to Layer 4 rules, then designate the tools that will inspect the traffic, and specify the order of the tools. For example, you can send Web traffic (defined by L4 port, 80 and/or 443) through a Web Application Firewall (WAF) and an Intrusion Prevention System (IPS), have backup traffic that might bypass inspection, and send all other traffic through the same IPS.

Figure 12Flexible Inline Arrangements Scenario illustrates a complex inspection scenario that is enabled by flexible inline arrangements.

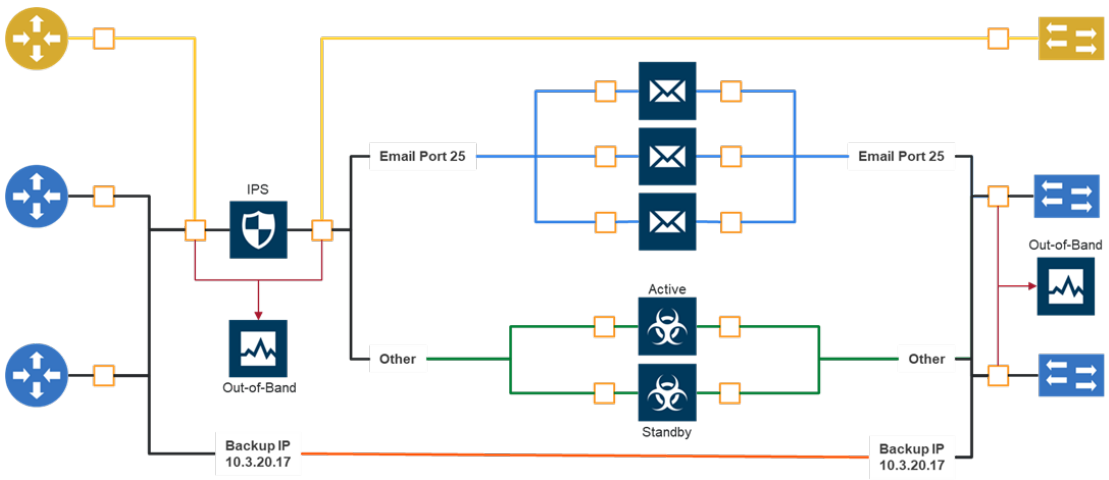


Figure 12 *Flexible Inline Arrangements Scenario*

In this example, on the left, there are three network links, all of which share the IPS. The details are as follows:

- At the top of the figure, the yellow line represent a flow of network traffic that only needs IPS inspection.
- In the middle of the figure, the Email and Other traffic represents network traffic flows that go through IPS inspection first, and then the Email traffic goes to dedicated inspection tools, here shown as three tools in an inline tool group, and the Other traffic goes to a threat protection active/standby pair.
- At the bottom of the figure, daily Backups of already-inspected traffic goes to bypass.

Although not shown graphically in Figure 12Flexible Inline Arrangements Scenario, the traffic in the reverse direction can have a different order of tools than the west-east traffic.

Refer to the Gigamon Validated Design for more information.

- Deploying GigaSECURE Inline SSL Solution using Flexible Inline
- Guiding Relevant Inline Traffic to Tools

Flexible Inline Arrangement vs Inline Bypass Solution

A brief difference between Flexible and Classic Inline Arrangements.

Flexible Inline Arrangement	Inline Bypass Solution (Classic)
Flexible inline arrangements offer flexibility in how traffic is guided through inline tools and the order in which they can be defined to move.	Inline Bypass solutions involve bidirectional traffic between two networks, intercepted by a Gigamon node, and guided through one or more inline tools. The classic Inline Bypass solution allows users to define what traffic is inspected by which inline tools and which should be bypassed and not inspected.
Guides traffic through any arbitrary sequence of inspection tools.	The Gigamon node sends uninspected traffic from one side of the network to an inline tool (such as an IPS), and then sends the inspected traffic to the other side of the network.
Shares inline tools across multiple inline network links and across multiple inline maps.	A sequence of inline tools / inline tool groups used for inspecting traffic on a given inline network / inline network group cannot be reused for inspecting traffic on another inline network / inline network group.
Distributes traffic across multiple tools to meet bandwidth and throughput demands.	The decrypted traffic needs to be received at all the attached inline tools and cannot provide flexibility to guide the traffic to individual tools.

Supported Platforms

Flexible inline arrangements are supported on the following nodes:

- GigaVUE-HC1-Plus
- GigaVUE-HCT
- GigaVUE-HC1
- GigaVUE-HC3(CCv1 and CCv2)
- GigaVUE-TA200
- GigaVUE-TA25
- GigaVUE-TA200E
- GigaVUE-TA25E
- GigaVUE-TA400
- GigaVUE-TA400E

The inline networks, inline tools, and inline tool groups involved in the flexible inline maps must be on the same node.

Flexible Inline Arrangement License

Flexible Inline Arrangements on the GigaVUE TA Series requires an Inline Bypass License. Refer the table below:

GigaVUE TA Series Node	Monthly Subscription (Term License)	Perpetual Subscription
GigaVUE-TA25	IBP-TAX20-SW-TM	IBP-TAX20
GigaVUE-TA25E	IBP-TAx20-SW-TM	IBP-TAX20
GigaVUE-TA200	IBP-TAC20-SW-TM	IBP-TAC20
GigaVUE-TA200E	IBP-TAC20E-SW-TM	IBP-TAC20E
GigaVUE-TA400	IBP-TAC40-SW-TM	IBP-TAC40
GigaVUE-TA400E	IBP-TAC40E-SW-TM	IBP-TAC40E

Licensing Rules and Notes

- Without the license you will not be allowed to create Inline-net and inline-tool ports.
- If your license has expired or uninstalled then the Inline-net and inline-tool ports will move to inactive state port params edit such as speed change and admin enable option will be blocked except for port type modification.

Software Version

GigaVUE-FM and GigaVUE-OS running on software version 5.3.xx or higher support the flexible inline arrangement functionality.

Flexible Inline Solution Supported in Clustered Nodes

The GigaVUE® HC Series, GigaVUE-TA25, GigaVUE-TA200, GigaVUE-TA400, and GigaVUE-TA400E nodes can now be clustered with Flexible Inline solution for the configuration path. The traffic path is limited to a single node such that all inline ports should reside in the same node. This feature is applicable in Out-of-Band and Leaf-Spine cluster. The Flexible Inline features that can be configured are as follows:

Supported Feature	Reference
Flexible Inline Solution	Flexible Inline Arrangements
Single VLAN Tag	Configure Inline Single Tag.
Non Shared Tool	Configure Inline Tool Ports and Inline Tools.
Resilient Weighted Hashing	Refer to 'Resilient Weighted Hashing' in Configure Inline Tool Group.
Out-of-Band Copy The oob-copy target ports must reside on the same node as the respective flexible inline type maps. The traffic originated from oob-copy can be exposed to out-of-band tool residing on other nodes by using hybrid ports. This feature also allows the flexibility to have different oob-copies on each direction (a-to-b and b-to-a) from GigaVUE-FM for bidirectional flex maps. Starting in software version 6.4, the tool/hybrid port which is part of a regular byRule or passall map with ingress VLAN tag enabled on the network port can also be used as an OOB-copy port of flexible inline maps.	Refer to Example 7—Protected Flexible Inline, Out-of-Band Copy in GigaVUE-OS CLI Reference Guide.
Resilient Inline Arrangement. In a Resilient Inline Arrangement, both GigaVUE® HC Series nodes should be in different clusters or either one of the node can be a standalone node.	Refer to 'Resilient Inline Arrangement' in Configure Resilient Inline Arrangement.
Network Link Aggregation Group (LAG)	Refer to Configure Inline Network Link Aggregation Group (LAG).
Hashing	Refer to Asymmetrical Hashing options in Configure Inline Tool Group.

For 5.13.00 version, only GigaVUE® HC Series devices can be configured for Flexible Inline solution in a cluster. Flexible Inline solution in a cluster is not applicable for GigaVUE TA Series device but, it can be part of a GigaVUE® HC Series cluster that is configured for this functionality.
For 5.14.00 version, only GigaVUE® HC Series and GigaVUE-TA200 devices can be configured for Flexible Inline solution in a cluster. Flexible Inline solution in a cluster is not applicable for other GigaVUE TA Series device but, it can be part of a GigaVUE® HC Series cluster that is configured for this functionality.
For 5.16.00 version, only GigaVUE® HC Series, GigaVUE-TA200 and GigaVUE-TA25 devices can be configured for Flexible Inline solution in a cluster. Flexible Inline solution in a cluster is not applicable for other GigaVUE TA Series device but, it can be part of a GigaVUE® HC Series cluster that is configured for this functionality

Limitations

The inline-network, inline-tool and destination tool ports in oob-copy must reside on the same node and cannot be chosen from across devices in a cluster.

Functionalities Not Supported by Flexible Inline Arrangement

Certain functionality is not supported by flexible inline arrangements. In the following cases, use classic inline bypass instead:

- for inline tools that cannot tolerate the addition of VLAN tags to the traffic

Also not supported:

- **Combining Classic and Flex**: Using both Classic Inline and Flexible Inline arrangements on the same device is not supported even in a standalone environment.

Benefits of Flexible Inline Arrangements

Flexible inline arrangements offer flexibility in how traffic is guided through inline tools. It has the following benefits compared to classic Inline Bypass:

- Guides traffic through any arbitrary sequence of inspection tools.
- Shares inline tools across multiple inline network links and across multiple inline maps.
- Distributes traffic across multiple tools to meet bandwidth and throughput demands.

Flexible inline arrangements use the same software constructs as the classic Inline Bypass solution, such as inline network, inline tool, and inline tool group. However, inline network group and inline serial constructs are not needed.

Inline network groups have changed with flexible inline arrangements. Now every inline network is independent and can share any combination of tools in any order. The concept of inline network group is supported by creating multiple flexible inline maps. Also, multiple inline networks can be grouped into an inline network bundle. You can configure one inline map for the network bundle with the inline network bundle as the source.

Inline serial is not needed because the flexibility offered with flexible inline arrangements allows for the same configuration without the inline serial construct.

Figure 12Flexible Inline Arrangements Scenario illustrates the benefits of flexible inline arrangements by showing the kinds of deployment scenarios that can be enabled with this approach.

About Flexible Inline Maps

Traffic flows are the building blocks of flexible inline arrangements. Flows can be based on any flow mapping criteria, such as TCP port, IP subnet, or VLAN. There is a one-to-one correspondence between a traffic flow and a flexible inline map.

A flexible inline map is a new map type. Flexible inline arrangements allow inline maps from inline networks to arbitrary sequences of shared (overlapping) sequences of inline tools and inline tool groups.

Using flexible inline maps, you can identify specific flows of traffic using Layer 2 (L2) to Layer 4 (L4) rules, then designate the tools that will inspect the traffic, and specify the order of traffic to the tools.

You can configure a flexible inline map with a specific inline tool that is part of an inline tool group, which is associated with another flexible inline map. For example, you have created an inline tool group, ITG1 in which inline tools, IT1, IT2, and IT3 are grouped together. You can configure a flexible inline map, Map1 with inline network, IN1 as the source and ITG1 as the destination. You can configure a second flexible inline map, Map2 with IN1 as source and IT1 as destination. Such configuration is useful to guide specific traffic to a particular inline tool and the rest of the traffic to the inline tool group in which the inline tool is associated.

To properly guide traffic through the inline tools, each flow of traffic is assigned a VLAN tag. VLAN tags can be automatically assigned or can be user-defined. You can use flexible inline single tags to map incoming VLANs on the network side to the outgoing VLANs on the tool side.

With flexible inline arrangements, VLAN tags are associated with each inline map, not with each inline network port as in the case of classic Inline Bypass. A single inline network port can have multiple inline maps, each with a separate VLAN tag.

A VLAN tag automatically assigned to an inline map can be manually added to another inline map. For example, if an inline map is given the VLAN tag 2000, the same VLAN tag can be manually added to another inline map. Here, the system would prioritize the manually added VLAN tag and modify the auto-assigned VLAN Tag. This is not applicable for VLAN’s assigned to internal inline maps such as import/export maps and inner non-proxy maps configured by GigaVUE-FM for Resilient Inline Arrangement (RIA) and RIA SSL solutions.

For example, traffic flows can be defined with the following VLAN tags:

- Unspecified traffic—VLAN 101
- Web traffic—VLAN 102
- Email traffic—VLAN 103
- Database traffic—VLAN 104

NOTE: The VLAN tags are added to the traffic before it is sent to the tools and are removed before it is sent back to the network.

Types of Flexible Inline Maps

To define a traffic flow, you must configure a flexible inline map. Following are the two types of flexible inline maps:

- byRule—Use the byRule map type to define a flow using map rules. Any standard L2-L4 mapping rules can be specified in the map, such as, IPv4, IPv6, L4 port, or UDA.
- collector—Use the collector map type for all other traffic. A collector is the lowest priority of map and does not have a map rule definition. Use a collector to catch any traffic that does not go to any other map. You can define a flexible inline collector map without any other maps in place. This provides a map passall, provided there are no rule-based maps. If you want all the traffic to go to the same tools, you only need to configure a collector.

Flexible inline arrangements guide rule-based or collector-based inline traffic flows through unidirectional or bidirectional sequences of inline tools or inline tool groups. The traffic path can be set up independently for side A to side B and side B to side A directions, meaning that the traffic flow can be either symmetrical or asymmetrical.

You specify the ordered list of inline tools or inline tool groups that will inspect a particular flow of traffic. Additionally, you can specify if the A-to-B and B-to-A directions have the same order or the reverse order. Reverse order is the order of inline tools as they are wired in a physical network if a Gigamon network packet broker was not present.

For example, in the A-to-B direction, if the tools are specified in the following order: T1, T2, T3, the same order in the B-to-A direction will be: T1, T2, T3, while the reverse order in the B-to-A direction will be: T3, T2, T1. Or, you can specify the order of the tools explicitly, for example, the B-to-A direction can be: T2, T1, T3.

You can create separate flexible inline maps for each flow of traffic to be inspected by a sequence of inline tools. Create maps until you have accounted for all the flows of traffic. Any unspecified traffic will go to the collector. You can also specify map priorities for the flexible inline maps.

Configure Flexible Inline Maps

Before you configure a flexible inline map, ensure that you create the required inline network. For information, refer to Configure Inline Network Ports and Inline Network .

To configure a flexible inline map:

1. Go to **Physical > Orchestrate > Inline Flows**, and then click **Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the Flexible Inline Canvas that is displayed, select the required device for which you want to configure the flexible inline map.
3. Drag and drop the required inline network into the Flexible Inline Canvas.
4. Drag and drop the Flex Map into the Flexible Inline Canvas.
5. In the **Properties** pane that appears, enter a name and description for the map in the **Alias** and **Description** fields.
6. Enter the required **Tool Side VLAN Tag**.
7. Select the **TPID** for the Tool Side VLAN Tag. The default value of TPID is 0x8100. You can select the other supported values 0x9100 and 0x88a8 from the drop-down list.
8. Add the required rules for the inline map. You can also choose to import an existing map template. For instructions on how to create a map template, refer to Create Map Templates.
9. Click **OK** to save the configuration.

Flexible Inline Arrangements—Rules and Notes

The table below lists the main rules, limits, and unsupported features for flexible inline arrangements, organized for quick reference.

Rule / Feature	Details / Limitations
Exclusive Use	Inline tools in a flexible inline map cannot be used in classic inline or inline decryption maps. All inline networks and tools must belong to only one type of map.
Collector Maps	Only one unidirectional collector map is allowed for the same inline network. To use different VLANs in each direction, create separate unidirectional maps with unique VLAN tags. Tags can be set manually or assigned automatically by GigaVUE-FM.
Unsupported Features (	- Physical Bypass (no BPS card)

Rule / Feature	Details / Limitations
GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA25E, GigaVUE-TA25, GigaVUE-TA400, GigaVUE-TA400E)	- Flexible and Resilient Inline SSL (no GigaSMART card) - GRIP (no BPS card) - ICAP (no GigaSMART card) - Classic Inline Bypass
VLAN Tagging and OOB Copy Limits (GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-HC1-Plus)	Flexible Inline Single VLAN Tag with monitoring mode may send incorrect VLAN tags. OOB copy packets may also have wrong tags. You cannot use BYPASS WITH MONITORING with MONITORING mode on the tool. OOB copy from inline network is not allowed in this mode.
Inline Map Limits — Bidirectional	GigaVUE-TA25,GigaVUE-TA25E, GigaVUE-HC1-Plus → 126 maps GigaVUE-HC1,GigaVUE-HC3 (CCv1 & CCv2), GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA400, GigaVUE-TA400E → 256 or 512 maps depending on setup.
Inline Map Limits — Unidirectional	GigaVUE-TA25,GigaVUE-TA25E, GigaVUE-HC1-Plus→ 252 mapsGigaVUE-HC1,GigaVUE-HC3 (CCv1 & CCv2),), GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA400, GigaVUE-TA400E → 512 or 1024 maps depending on setup.
Flexible Inline SSL Limits	Not supported with Inline Network LAG. Setting inline tools to “Drop” in the chain does not block Inline SSL traffic.
Filtering Limits(GigaVUE-TA400, GigaVUE-TA400E)	VLAN-based filtering in the Egress Port Filter for OOB copies is not supported. If one tool in the map is in monitoring mode, all tools must use the same mode. Asymmetric hashing (a-srcip-b-dstip and b-srcip-a-dstip) is not supported.
Protocol Pass-Through(GigaVUE-TA400, GigaVUE-TA400E)	CDP pass-through is not supported when the source is an Inline Network LAG. Bypass for LACP, CDP, and LLDP is supported.
Scaling Limits — GigaVUE-TA400, GigaVUE-TA400E	Max Inline Networks and Tools: 48 Max Inline Network LAG list: 24 Max Inline tools or tool groups per direction: 16 Max OOB copy entries per direction: 17 Max OOB copy ports per entry: 128

Visualize the Flexible Inline Arrangements Canvas

The GigaVUE-FM user interface provides clear visualization of inline maps. The user interface makes it easy to visualize and configure inline maps and tools. The drag-and-drop capability lets you define and add tools to maps, in any order.

Configure Flexible Inline Flows

The different flexible inline flows are configured and managed using GigaVUE-FM. For end-to-end workflows, examples, and best practices, refer to the GigaVUE Inline Solutions Guide.

Troubleshoot Flexible Inline Flows

The flexible inline canvas provides you with the ability to view the status of the flexible inline flow deployments, port statistics, and the details of the cluster-level maps used in the deployments. You can use these details to troubleshoot any issues or failures in your flexible inline flows.

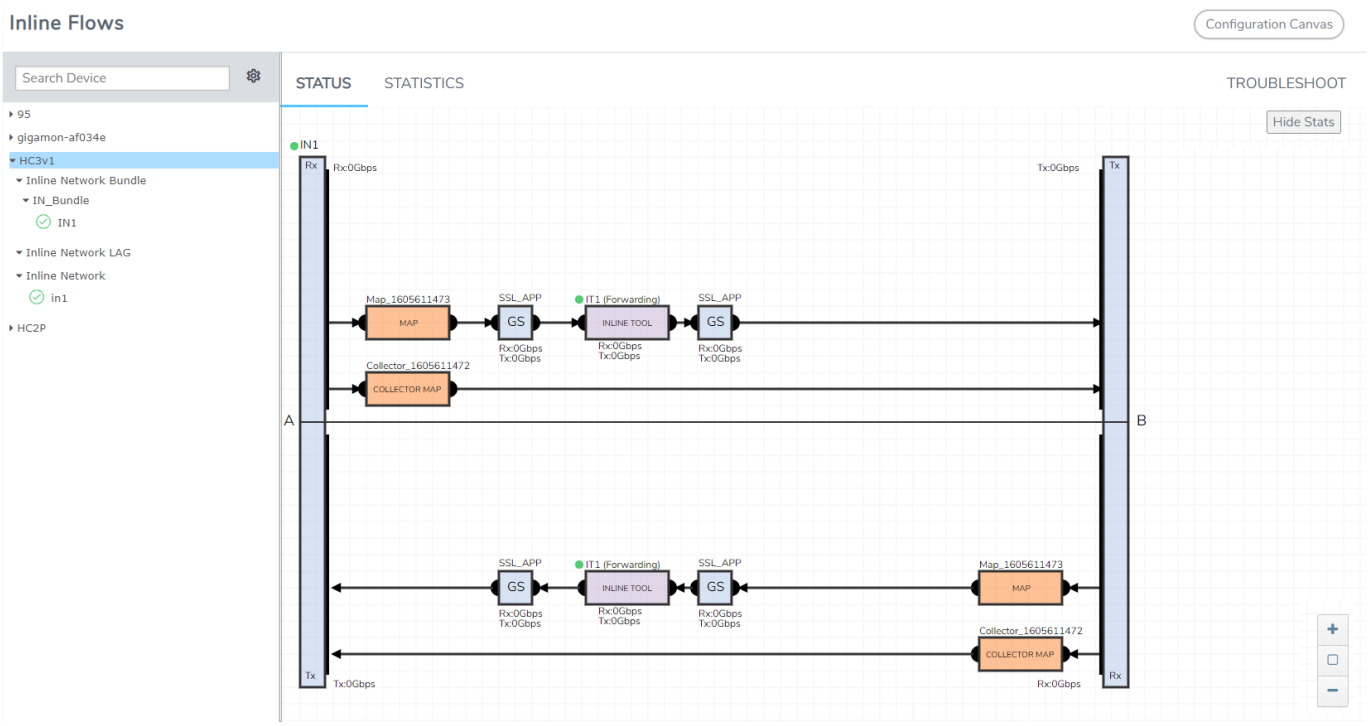
The flexible inline canvas has the following tabs:

Status

The Status tab provides details of the forwarding states of inline network. Click the required component that is part of the selected inline network to view the component's properties. Refer to View the Forwarding States of Inline Networks.

It also provides the status of the components that are part of the selected inline network. Hover over the status to view the description.

Click the **Show Stats/Hide Stats** toggle button to view the Rx/Tx rate for the components that are part of the flexible inline flow deployment. Refer to the following figure for details.



Rules and Notes

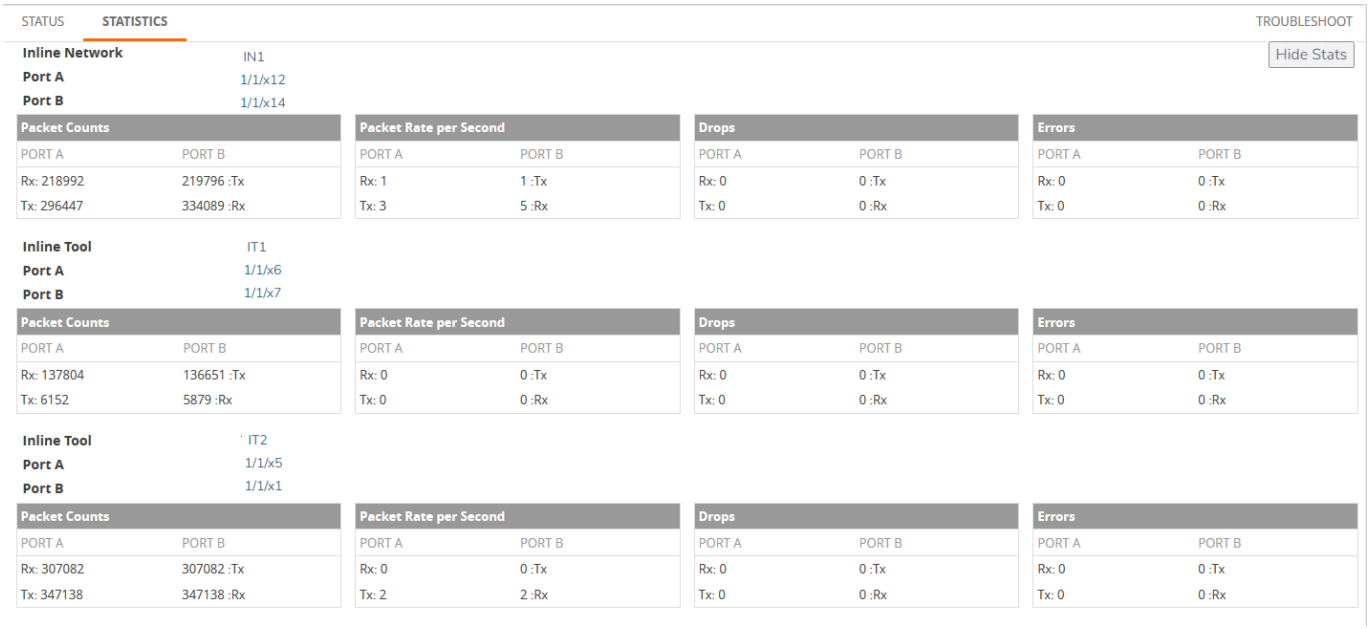
Keep in mind the following points when you are viewing the details in the Status tab:

- The ability to view the Rx/Tx rate for the components is supported only for GigaVUE-HC1, and GigaVUE-HC3 devices.

- The Rx/Tx rate for inline components are displayed in GigaVUE-FM only for flexible inline deployments. In addition, the historical trends are not displayed in GigaVUE-FM. However, you can view the Rx/Tx rate for classic inline components using the GigaVUE-OS CLI commands. For details, refer to the following topics in the *GigaVUE-OS CLI Reference Guide*:
    - show
    - ib pathway
    - inline tool group
    - inline tool
    - inline serial
    - inline network group
    - inline network
    - icap
- The Rx/Tx rate that is displayed is not for a map or a flexible inline flow, but it is a cumulative value that is shared across multiple maps and flexible inline flows.
- The Rx/Tx rate for a GigaSMART group is a cumulative value for the traffic flowing between both a-b and b-a directions.
- The Rx/Tx rate for OOB copy is supported only for single tool.
- You cannot view the health status of GigaSMART.
- The near real time data is not displayed for LED's health status and failover state representation of the inline component. The data gets updated during the GigaVUE-FM configuration synchronization period. Alternatively, when the Rx/Tx rate for an inline component drops to zero, you can choose to refresh the Status tab to view the updated data.

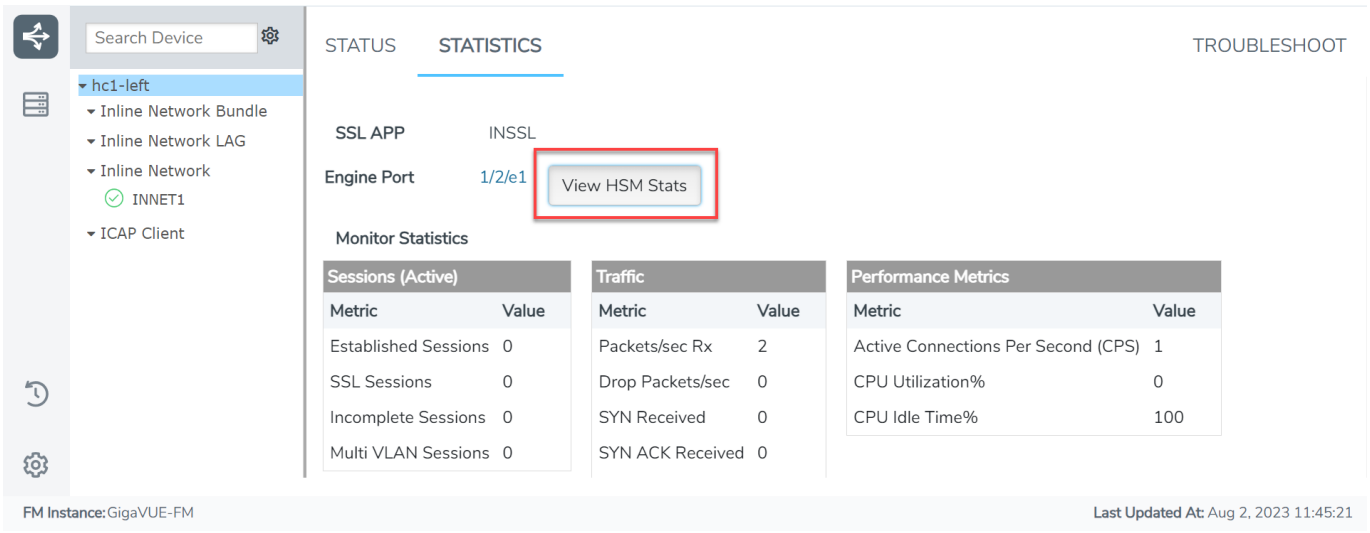
Statistics

The Statistics tab provides statistical information of the inline network ports, inline tool ports, and the virtual ports used in the selected inline network. It also provides the inline decryption session statistics for the inline network. The inline network, inline tools, ICAP Client, and the ports aliases are displayed as clickable links. Use these links to access the quick view of the respective component. Refer to the following figure for details.



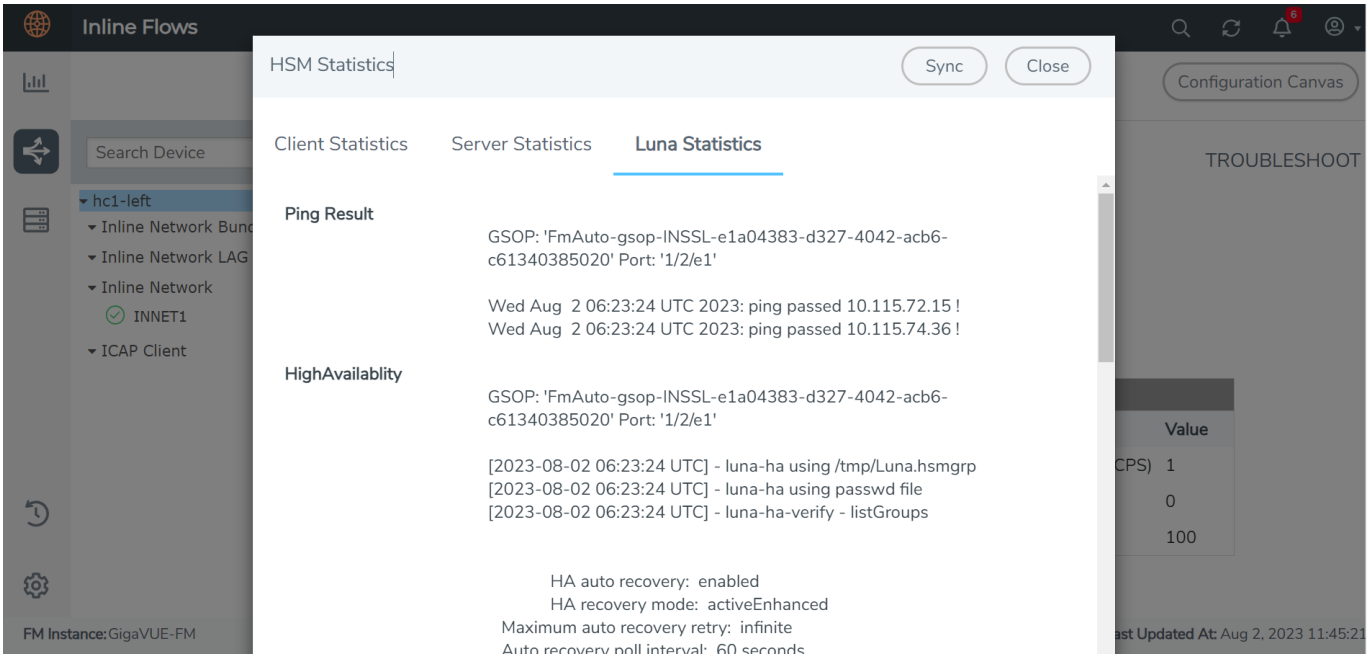
HSM Statistics

To view HSM Statistics, click the **View HSM Stats** button on the Statistics tab.



The HSM Statistics window includes the following tabs:

- **Client Statistics** tab, which includes information on number of requests received and responses sent, metrics, and error details.
- **Server Statistics** tab, which includes details on the number of requests sent and responses received, delay, and average time drop details.
- Starting in software version 6.4, **Luna Statistics** is included in the HSM Statistics window. The **Luna Statistics** tab provides statistical information on the Ping Result, High Availability, and the verification details.



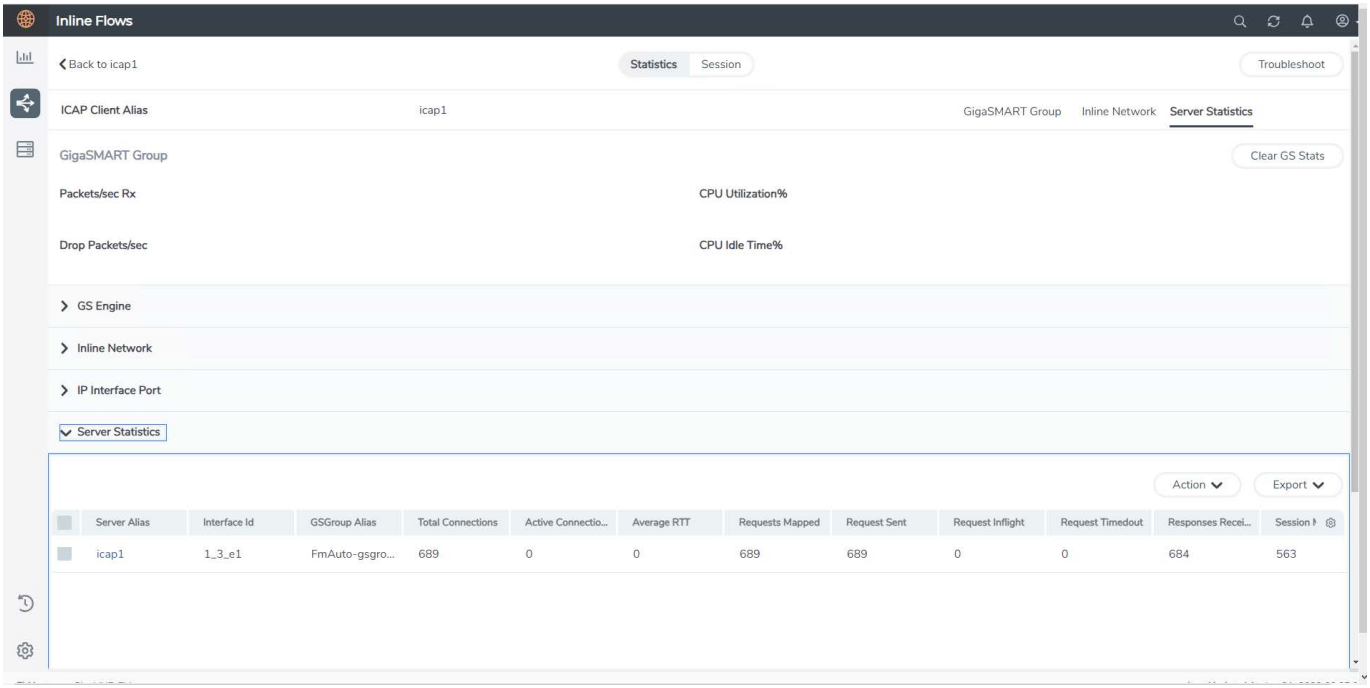
ICAP Show Stats

The Show Stats option will be displayed after the ICAP Client app is configured and deployed. The Show Stats option has the following tabs:

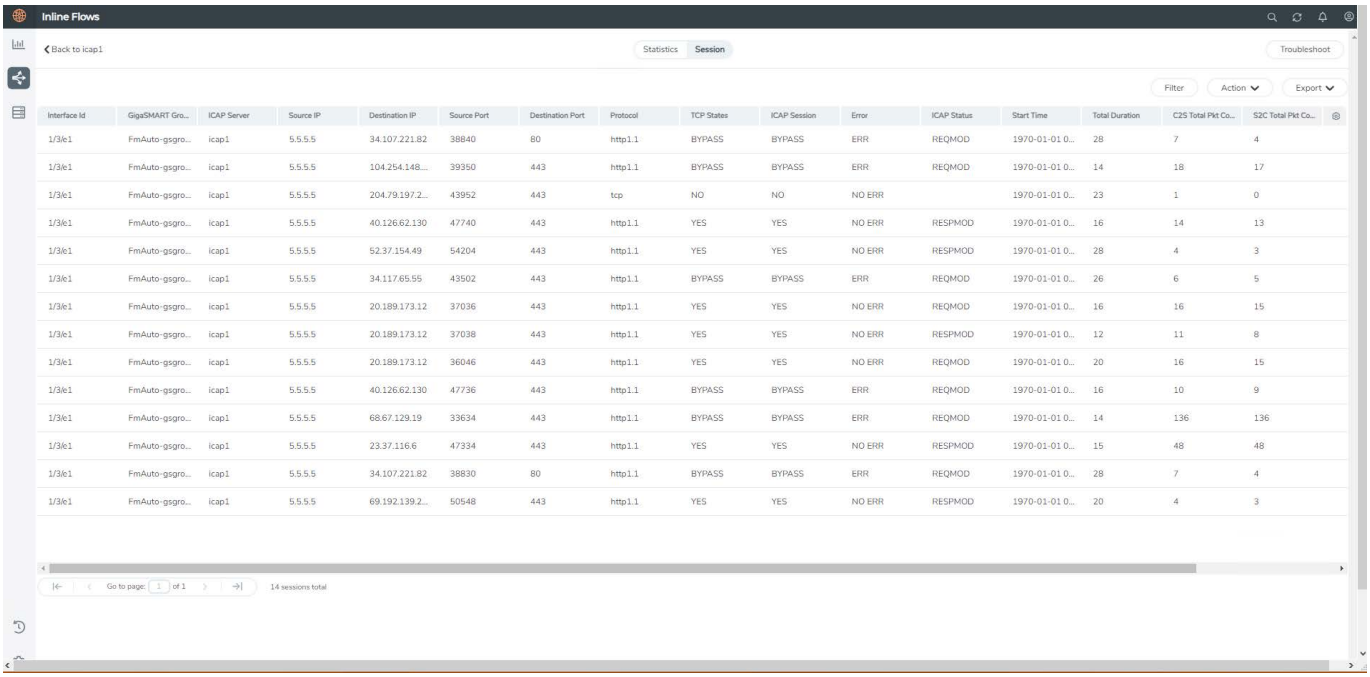
- Statistics
- Session

The **Statistics** tab provides statistical information on the GigaSMART group, GS engines, inline network, IP interface port, and server statistics. Click the drop-down menu of each component listed on the statistics page to know more about the configuration details. You can also view the statistics from the sidebar on the Inline Flows page.

The GS Engine drop-down menu provides information on the ICAP session statistics. Click **Clear GS Stats** to clear the statistics of the GS group.



To view ICAP session statistics details, click the **Sessions** tab.



Refer to the below table for more details on the sessions tab.

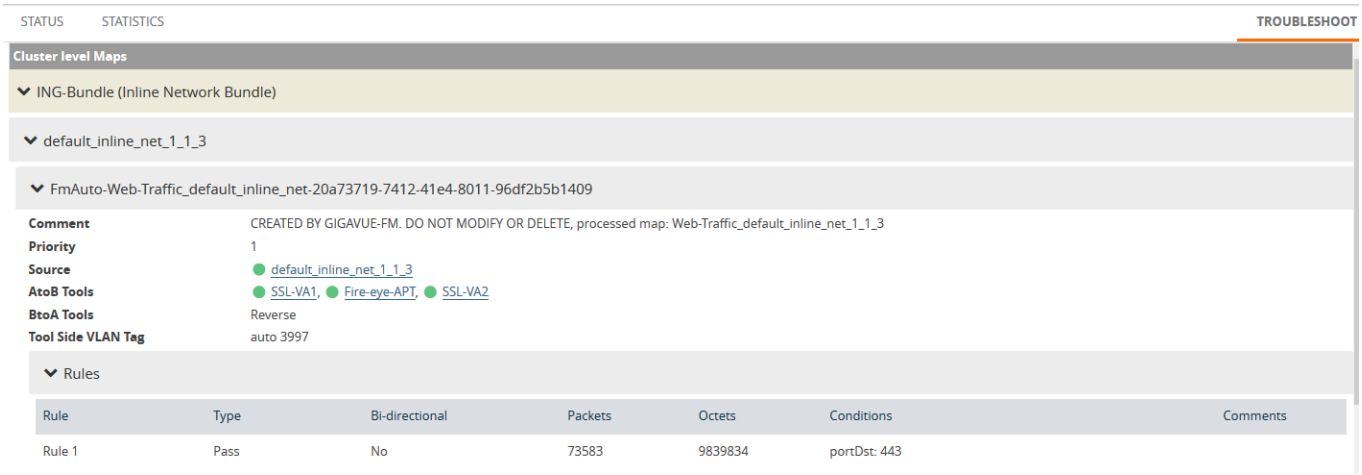
Field	Description
Filter	To view the statistics of a particular IP with selected parameters.
Action	**Upload to server**: To export the log details to an external server.
Export	To export the current session details to local.

Click the **Troubleshoot** button to view more on the below map configuration details for the ICAP Client app:

- Map Alias
- From
- To
- GSOP

Troubleshoot

The Troubleshoot tab provides details about the device-level maps that are created for the selected inline network. The details include the source of the map, inline tools or inline tool groups used in the A to B and B to A directions, tool side and network side VLAN tags, and OOB copies. Click the required component to view its configuration details. It also provides details about the map statistics. In case of inline tool group, you can view the list of inline tools that are associated with the inline tool group. Click the required inline tool from the list to view its configurations. Refer to the following figure for details.



To troubleshoot any issue or failure in your flexible inline flow:

1. From the flexible inline canvas, go to the **Status** tab to check the forwarding states of the required inline network.
2. Go to the **Statistics** tab to check the port statistics to ensure that there are no drops, errors, or discards.

   NOTE: Click the **Statistics** tab again to refresh the data.

3. Go to the **Troubleshoot** tab to check the required map configurations and isolate the issue.

For instructions on how to troubleshoot a specific issue, refer to Example: Troubleshoot Traffic Issues Between Side A and Side B.

Example: Troubleshoot Traffic Issues Between Side A and Side B

This section provides you an example of how to troubleshoot a specific issue using the flexible inline canvas.

Consider that you have an inline network IN1 with one by-rule map and one collector map. The traffic is not flowing from Side A to Side B. To troubleshoot this issue:

1. Go to **Physical > Inline Flows**.
2. Select the required device, and then drill down to the inline network that has the traffic issue.
3. Click **Status** to view the forwarding states of the inline network. Ensure that the forwarding state of the inline network is Normal. For details of the forwarding states, refer to View the Forwarding States of Inline Networks.
4. Click **Statistics** to view the total packet count of the inline network. The Rx count of Port A of the inline network must match the Tx count of Port B. If the count does not match, the traffic is blocked between Port A and Port B.

   NOTE: To refresh the statistical data, click **Statistics** again.

5. Click **Troubleshoot** to view the configurations of the by-rule and collector maps configured for the inline network.
   a. Check that the by-rule map has packet count.
   b. If the by-rule map has packet count, check the inline tools from A to B and B to A directions of the by-rule map to ensure that the inline tools are in up state and the **Flex Traffic Path** is set to **To inline tool**.

6. Click **Statistics** to view the port statistics of the required inline tool. If the Rx count of Port A of the inline tool does not match the Tx count of Port B of the inline tool, the traffic is dropped at the inline tool. Check the inline tool and take the required action.
7. If the Rx count of Port A of the inline tool matches the Tx count of Port B of the inline tool, repeat steps 5 and 6 for the collector map configured for the inline network.
8. If the Rx count of Port A of the inline tool configured for the collector map matches the Tx count of Port B of the inline tool, contact Gigamon Technical Support.

View the Forwarding States of Inline Networks

To view the forwarding states of inline networks in the flexible inline canvas, choose the required inline network, and then click **Status**. Refer to the following figure.

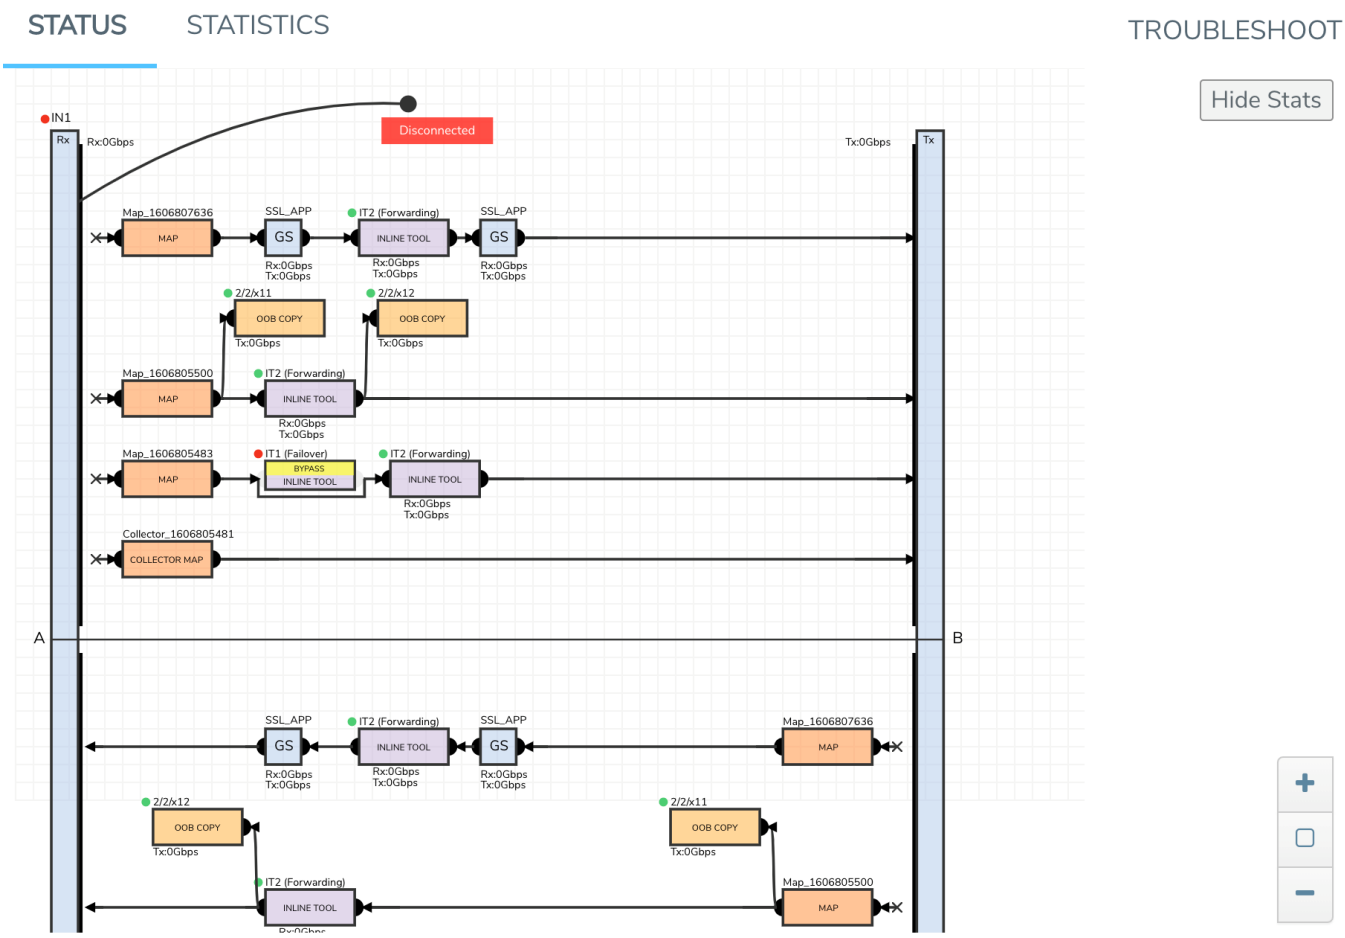


Figure 13 *Inline Network Forwarding States*

Following inline network states are not explicitly shown in the flexible inline canvas:

- Normal—If the state of all inline tools are up, the inline network is in Normal state.
- Abnormal—If any inline tool involved in flexible inline maps (directly or indirectly as a member of an inline tool group) is operationally down and there is no network-level failover action in effect, the inline network is in an Abnormal state.

Following table provides the list of forwarding states of inline network and their description.

Table 15: Forwarding States of Inline Networks

Inline Network Physical Bypass	Inline Network Traffic Path	Far-End Status of Links Connected to Inline Network Ports	Operational State of *Not* Forced Inline Tools and Inline Tool Groups Involved in Maps from the Inline Network	Forwarding State	Description
enable	any inline network traffic path configuration	any combination of far-end ports status	any combination of operational state of the inline tools or inline tool groups involved in the maps originating from the inline network	PHYSICAL BYPASS	all traffic exchanged directly between the end nodes without being noticed by the switching fabric (GigaVUE node acting as a wire or fiber)
disable	traffic path set to drop	any combination of far-end ports status	any combination of operational state of the inline tools or inline tool groups involved in the maps originating from the inline network	DISABLED	all traffic arriving at the inline network ports is dropped
disable	traffic path set to bypass, monitoring, or to-inline-tool	at least one far-end port is down	any combination of operational state of the inline tools or inline tool groups involved in the maps originating from	DISCONNECTED	No traffic is exchanged between the nodes

Inline Network Physical Bypass	Inline Network Traffic Path	Far-End Status of Links Connected to Inline Network Ports	Operational State of *Not* Forced Inline Tools and Inline Tool Groups Involved in Maps from the Inline Network	Forwarding State	Description
			the inline network		
disable	traffic path set to bypass	both far-end ports are up	any combination of operational state of the inline tools or inline tool groups involved in the maps originating from the inline network	FORCED BYPASS	All traffic that matches any of the maps originating from the inline network is redirected through a logical bypass
disable	traffic path set to monitoring	both far-end ports are up	any combination of operational state of the inline tools or inline tool groups involved in the maps originating from the inline network	FORCED BYPASS WITH MONITORING	A copy of the traffic originating from the inline network bypasses the sequence of inline tools and inline tool groups and is re-directed to the opposite-side inline network port. Another copy of the traffic is directed to the sequence of inline tools and inline tool groups, except that no traffic of the second copy is sent to the exit port.

Inline Network Physical Bypass	Inline Network Traffic Path	Far-End Status of Links Connected to Inline Network Ports	Operational State of *Not* Forced Inline Tools and Inline Tool Groups Involved in Maps from the Inline Network	Forwarding State	Description
disable	traffic path set to to-inline-tool	both far-end ports are up	all inline tools involved (directly or indirectly as members of inline tool groups) in the maps originating from the inline network are in the *up* operational state	NORMAL	The traffic is guided between the source inline network port and the destination inline network port according to the status of the inline tools and inline tool groups. **NOTE:** The state of all inline tools must be *up*, including inline tools configured as spare in an inline tool group, inline tools or inline tool group members in the **a-to-b** and **b-to-a** lists configured with any traffic path other than to-inline-tool.
disable	traffic path set to to-inline-tool	both far-end ports are up	at least one of the inline tools or	NETWORK PORTS FORCED DOWN	No traffic is exchanged between the

Inline Network Physical Bypass	Inline Network Traffic Path	Far-End Status of Links Connected to Inline Network Ports	Operational State of *Not* Forced Inline Tools and Inline Tool Groups Involved in Maps from the Inline Network	Forwarding State	Description
			inline tool groups involved in the maps originating from the inline network configured with the traffic path parameter to-inline-tool and failover action of network-port-forced-down is in the *down* operational state		inline network ports, and the inline network ports are brought down
disable	traffic path set to to-inline-tool	both far-end ports are up	**a.** none of the inline tools or inline tool groups involved in the maps originating from the inline network configured with to-inline-tool and failover action network-port-forced-down is in the *down* operational state **b.** at least one of the inline tools or inline tool groups involved in the maps originating	FAILURE INTRODUCED DROP	All traffic arriving at the inline network ports is dropped

Inline Network Physical Bypass	Inline Network Traffic Path	Far-End Status of Links Connected to Inline Network Ports	Operational State of *Not* Forced Inline Tools and Inline Tool Groups Involved in Maps from the Inline Network	Forwarding State	Description
			from the inline network configured with to-inline-tool and failover-action of network-drop is in the *down* operational state		
disable	traffic path set to to-inline-tool	both far-end ports are up	**a.** none of the inline tools or inline tool groups involved in the maps originating from the inline network configured with to-inline-tool and failover action of network-port-forced- down or network-drop is in the *down* operational state **b.** at least one of the inline tools or inline tool groups involved in the maps originating from the inline network configured with	FAILURE INTRODUCED BYPASS	All traffic that matches any of the maps originating from the inline network is redirected through a logical bypass

Inline Network Physical Bypass	Inline Network Traffic Path	Far-End Status of Links Connected to Inline Network Ports	Operational State of *Not* Forced Inline Tools and Inline Tool Groups Involved in Maps from the Inline Network	Forwarding State	Description
			to-inline-tool and failover action of network-bypass is in the *down* operational state		
disable	traffic path set to to-inline-tool	both far-end ports are up	any combination of conditions not listed for the forwarding state definitions of PHYSICAL BYPASS, DISABLED, DISCONNECTED, FORCED BYPASS, FORCED BYPASS WITH MONITORING, NORMAL, NETWORK PORTS FORCED DOWN, FAILURE-INTRODUCED DROP, or FAILURE-INTRODUCED BYPASS	ABNORMAL	The traffic is guided between the source inline network port according to the status of the inline tools and inline tool groups **NOTE:** If any inline tool involved in flexible inline maps (directly or indirectly as a member of an inline tool group) is in the *down* operational state and there is no network-level failover action in effect, the inline network is in the ABNORMAL state.

NOTE: When the Inline Network traffic path is set to monitoring and the Inline Tool Failover action is configured as 'Network Port Force Down,' the Inline Tool will continue to receive traffic, even if it is disabled or in an operational down state. This behavior is expected.

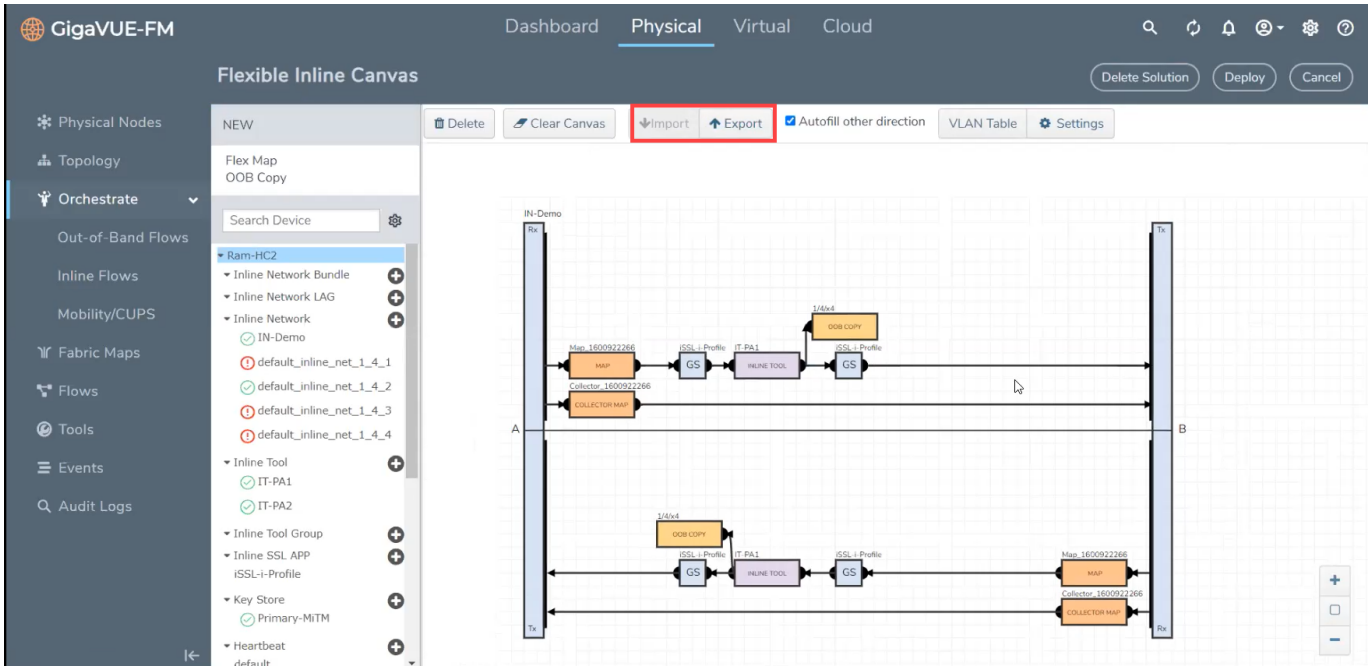
Import and Export Flexible Inline Solution

The flexible inline canvas allows you to import and export a flexible inline solution. The exported solution gets downloaded in to your local folder as an YAML file, which can be imported and deployed again in the following scenarios:

- Retrieve a solution that was deleted unintentionally
- Deploy the solution in another device
    - Re-deploy a solution in the device after GigaVUE-FM is upgraded to a new version (in case of issues in the existing solution)

The downloaded YAML file contains the following information:

- Software version of the GigaVUE-FM instance and the device on which the solution was created
- Member, port and configuration information of the various inline configurations such as the Inline Network Alias, Member Ports, Inline Tool Group Alias, and other such details.



Rules and Notes

- When deploying a solution after GigaVUE-FM is upgraded to a new version, the source port in the target device must not have been used in any solution.
- When importing an YAML file, the following configurations specified in the YAML file must match the configurations in the target device:
  - Inline network alias
  - Inline network bundle alias
  - Inline tool alias
  - Inline network LAG alias
  - Inline SSL App alias

Import and Export a Flexible Inline Solution

To import a file:

1. Click the **Import** button on the canvas.
2. Select the .YAML file, containing the required flexible inline solution, saved in your local folder. The solution appears in the canvas.
3. Click **Deploy** to deploy the solution, again.

You can also edit an imported solution on the canvas, and re-export the same as a new solution.

To export a solution:

1. Create a flexible inline solution and **Deploy** it.
2. Click the **Export** button. The solution is downloaded as an YAML file.
3. Save the file to the required location.

Backup and Restore Flexible Inline Flows

You must backup the devices and GigaVUE-FM at the same time. Ensure that the devices are not undergoing configuration edits during backup or restore. For more information, refer to the *"Restore Devices and GigaVUE-FM for Traffic Management Solutions"* section in the *"GigaVUE Administration Guide"*.

Inline Bypass Solutions

This chapter provides the following information about inline bypass solutions:

- About Inline Bypass Solutions
- Configure Inline Bypass Solutions

About Inline Bypass Solutions

Security tools such as firewalls and intrusion protection systems (IPSs) are often connected inline on production networks, with traffic flowing from the network segment through the tool and then back onto the production network.

Inline bypass solutions involve bidirectional traffic between two networks, intercepted by a Gigamon node, and guided through one or more inline tools.

Inline bypass is a pillar of the GigaSECURE Security Delivery Platform. For an overview of GigaSECURE, refer to About Inline Bypass Solutions.

Inline bypass is supported on GigaVUE® HC Series nodes: GigaVUE-HC3, GigaVUE-HC1 and GigaVUE-HC1-Plus.

NOTE: Classic Inline Bypass Solution is not supported in GigaVUE-HCT devices.

This section describes inline bypass solutions. Refer to the following sections for details:

- Introduction to Inline Bypass Solutions
- Logical Bypass and Physical Bypass
- Inline Networks
- Simple and Complex Inline Bypass Solutions

NOTE: The configuration of inline bypass solutions can be complex. Follow the order of configuration outlined in Configuration Steps and demonstrated in Inline Bypass Solution Examples.

Introduction to Inline Bypass Solutions

Inline bypass solutions place the Gigamon node inline between two sides of a network. The Gigamon node sends uninspected traffic from one side of the network to an inline tool (such as an IPS), and then sends the inspected traffic to the other side of the network.

The reasons for deploying Gigamon's inline bypass solutions are as follows:

- to protect against inline tool failures (including loss of link to inline tools and any inline tool problems detected through heartbeat monitoring)
- to bypass the inline tools for traffic that does not need to be examined by the tools or that cannot be processed by the tools
- to distribute the traffic load among multiple tools
- to send specific traffic to specialized inline tools
- to guide traffic serially through the inline tools, with the traffic from one tool flowing to the next, so that all tools see the same traffic
- to share inline tools among multiple inline network links
- to implement tiered network security by combining inline and out-of-band, allowing the traffic to be examined by both types of tools

For improved reliability and ease of maintenance, the inline bypass solutions also provide the following:

- protection against power loss through physical bypass
- protection against power loss through the high availability solution, Gigamon Resiliency for Inline Protection (GRIP™)

Capabilities of Inline Bypass Solutions

The inline bypass solution offers the following capabilities:

- Guides traffic through one or more inline tools according to user-defined bidirectional connectivity arrangements among respective inline network ports and inline tool ports.
- Reacts to failure conditions, such as the failure of inline tools, or the failure of links leading to the end-point devices between which the inline tool is inserted.
- Configures maps to associate inline networks with inline tools and inline tool groups.
- Supports bidirectional heartbeat and negative heartbeat to monitor inline tool health.
- Supports protected and unprotected inline networks.

- Supports physical bypass with specialized hardware on GigaVUE-HC3, , GigaVUE-HC1 and GigaVUE-HC1-Plus(bypass combo modules equipped with optical protection switches)
- Supports physical bypass on and GigaVUE-HC1 (TAP modules equipped with electrical relays).
- Distributes traffic across multiple inline tools, providing load sharing and traffic redistribution should a tool fail.
- Supports out-of-band maps for selective forwarding of inline traffic to monitoring tools.
- Supports sharing of inline tools by multiple inline networks.
- Supports N+1 and 1+1 inline tool redundancy for inline tool groups.
- Supports guiding inline traffic through inline tools in a serial fashion.
- Supports inline flow mapping through rule-based and shared collector inline maps.
- Provides a high availability solution, Gigamon Resiliency for Inline Protections (GRIP).

Logical Bypass and Physical Bypass

Logical bypass lets traffic bypass the inline tool should it experience a failure. A failure is declared if the GigaVUE node either loses connectivity to the inline tool or fails to receive a heartbeat from the tool.

Logical inline bypass does not require any specialized hardware and can be facilitated on regular ports on GigaVUE-HC3, GigaVUE-HC1 and GigaVUE-HC1-Plus modules. On GigaVUE-HC1, logical inline bypass can be configured on the base module.

Physical bypass provides protection against failure of the GigaVUE node, such as if power is lost. On GigaVUE® HC Series nodes, it is implemented with specialized hardware called bypass combo modules.

The specialized hardware triggers a bypass when power is lost to the node. When the physical bypass is activated, traffic flows from one side of the network to the other, but without monitoring.

The GigaVUE-HC3 bypass combo module is a bypass switch (BPS) module, as follows:

- Bypass Combo Module with two 100Gb/40Gb SR4 MPO inline network port pairs and sixteen regular SFP+ (10Gb) port cages (BPS-HC3-C25F2G)
- Bypass Combo Module with two 40Gb or 100Gb LR4 inline network port pairs and sixteen 10Gb/25Gb cages (BPS-HC3-Q35C2G, and BPS-HC3-C35C2G).

The GigaVUE-HC1 bypass combo module is a bypass switch (BPS) module, as follows:

- Bypass Combo Module with two SX/SR (50/125µm multimode) inline network port pairs and four regular SFP+ (1Gb/10Gb) port cages (BPS-HC1-D25A24)
- Bypass Combo Module with six single-mode LR 1Gb/10Gb inline network port pairs BPS-HC1-D35C60 (LX/LR) and BPS-HC1-D25A60 (SX/SR) .

The following bypass modules are available for the GigaVUE-HC1-Plus:

- Bypass Combo Module with two SX/SR (50/125µm multimode) inline network port pairs and four regular SFP+ port cages (BPS-HC1-D25A24). It requires GigaVUE-HC1-Plus to be running software version 6.3 or higher. It provides both physical and logical inline bypass, with automatic failover protection in case of power failure. The ports can operate at 1Gb/10Gb speeds and can be configured as any port type
- Bypass modules with six inline network port pairs each, supporting multi-mode SR and single-mode LR, respectively, at 1Gb/10Gb speeds (BPS-HC1-D25A60 (SX/SR) and BPS-HC1-D35C60 (LX/LR)).

All these modules are hot-swappable and provide physical bypass protection, ensuring network continuity in the event of a power failure or module removal. The BPS modules provide optical bypass, while the TAP-HC1-G10040 provides copper bypass functionality.

Physical bypass is also supported on the following copper TAP modules on GigaVUE-HC1 and GigaVUE-HC1-Plus:

- On the GigaVUE-HC1 copper TAP module, TAP-HC1-G10040
- On the GigaVUE-HC1-Plus copper TAP module with 8 x 1Gb ports for either four TAPs or bypass copper (BPC) port pairs (TAP-HC1-G10040). It can be used as a copper bypass module, providing protected inline networks for copper ports. It requires GigaVUE-OS version 6.4 or higher.

The GigaVUE-HC3, GigaVUE-HC1 and GigaVUE-HC1-Plus offer physical and logical inline bypass. Physical bypass provides automatic failover protection in the case of a power failure. On the GigaVUE HC Series nodes, the bypass combo modules provide the physical bypass function. Or on GigaVUE-HC1, the copper TAP modules provide the physical bypass function.

The physical bypass function, as it applies to a single pair of inline network ports, is as follows:

- When the module is not powered, (either the entire node is powered down or the module is removed from the node), the inline network port pair is in the physical bypass mode. That means that traffic is exchanged directly between network Port A and network Port B of the inline network pair.
- When the module is powered, the mode (inline or bypass) of the inline network port pair is controlled through software. In the physical bypass mode, the inline network port pair behaves exactly as if the module was not powered. In the inline mode, the inline network port pair behaves as any other inline network port pair configured for working with an inline tool.

For information on the bypass combo module on GigaVUE-HC3, refer to the *GigaVUE-HC3 Hardware Installation Guide*. This document also contains procedures for installing, removing, and replacing modules in the GigaVUE-HC3.

For information on the bypass combo module or the TAP-HC1-G10040 module on GigaVUE-HC1, refer to the *GigaVUE-HC1 Hardware Installation Guide*. This document also contains procedures for installing, removing, and replacing modules in the GigaVUE-HC1.

For information on the bypass combo module on GigaVUE-HC1-Plus, refer to the *GigaVUE-HC1-Plus Hardware Installation Guide*. This document also contains procedures for installing, removing, and replacing modules in the GigaVUE-HC1-Plus Hardware Installation Guide.

Refer also to Figure 14Physical Bypass Protection and Protected Inline Network.

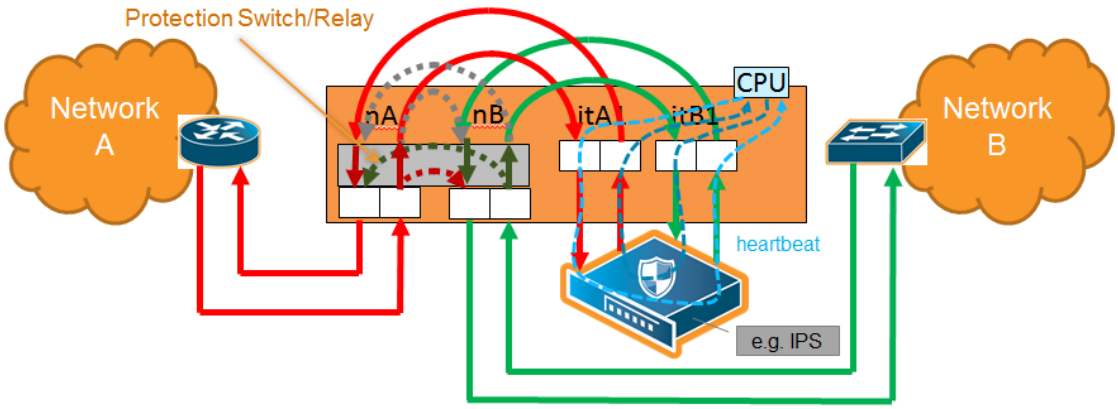


Figure 14 *Physical Bypass Protection*

Configure Inline Bypass Solutions

Inline bypass solutions are arrangements of inline networks that act to guide traffic through one or more inline tools attached to a GigaVUE node.

Configure inline bypass solutions using the following software constructs:

- Inline Network.

  For configuring inline networks, which are pairs of network ports attached to the sides of the network, refer to Inline Networks.
- Inline network groups

  For configuring inline network groups, which are multiple inline networks sharing the same inline tool or tools, refer to Inline Network Groups.
- Inline tool

  For configuring inline tools, which are pairs of inline tool ports plus the inline tools attached to the ports, refer to Inline Tools.

- Heartbeat profile

  For configuring monitoring of inline tools using heartbeat and/or negative heart beat packets or both, refer to Heartbeats.
- Inline tool group

  For configuring multiple inline tools, which are arrangements of multiple inline tools to which traffic is distributed based on hashing, refer to Inline Tool Groups.
- Inline serial tools

  For configuring inline tools in a series, in which traffic is guided through the inline tools in a serial fashion. Refer to Inline Serial Tools.

Inline Networks

Currently configured in-line networks are displayed on the Inline Networks page, which is opened by selecting **Inline Bypass > Inline Networks**.

Click **New** to open the **Inline Network** configuration page to configure an inline network, which is a pair of network ports arranged for inline monitoring. The arrangement facilitates access to a bidirectional link between the two networks (two far-end network devices) that need to be linked through an inline tool.

An inline network consists of inline network ports, always in pairs, running at the same speed, on the same medium (either fiber or copper). The inline network ports must be on the same GigaVUE-HC3, GigaVUE-HC1 or GigaVUE-HC1-Plus node.

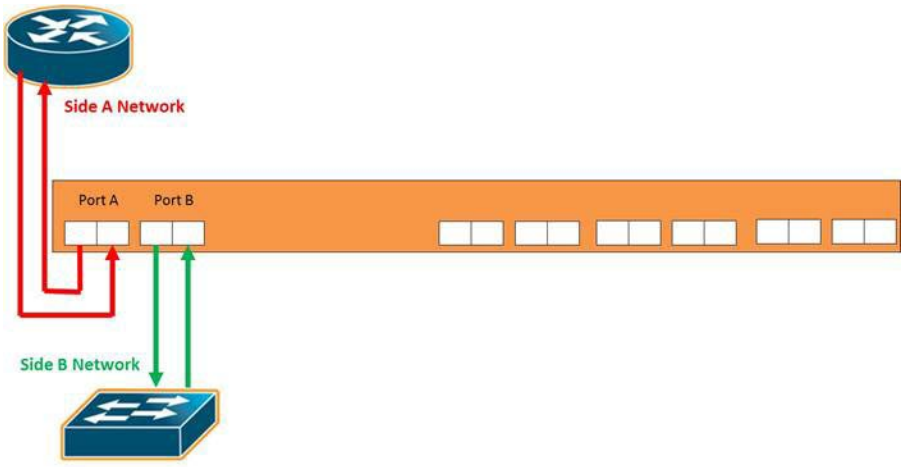
The types of inline networks are as follows:

- Unprotected Inline Network
- Protected Inline Network
- Mix of Protected and Unprotected

Unprotected Inline Network

An unprotected inline network consists of two ports of the inline-network type, which facilitate access to a bidirectional link between two networks (or more precisely, two far-end network devices).

Any available network type ports on a GigaVUE HC Series node can be configured to be inline-network type ports and combined to form an unprotected inline network. An unprotected inline network is shown in the following figure.



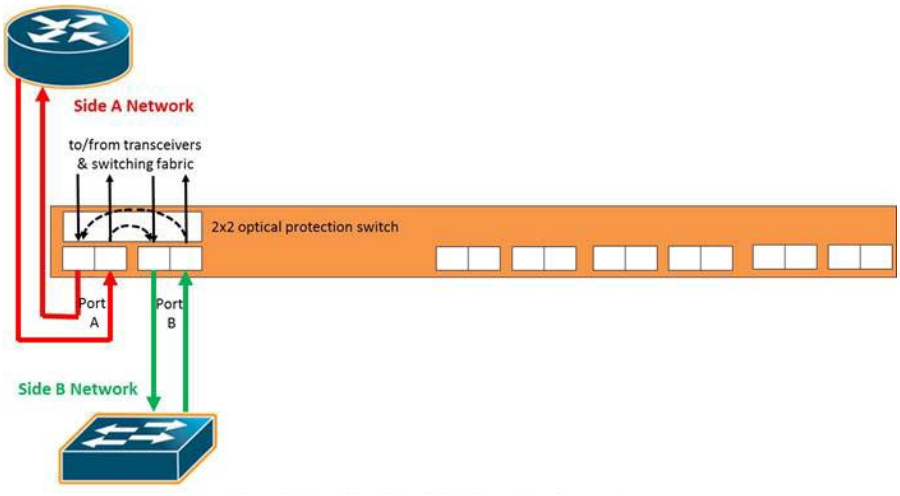
Protected Inline Network

A protected inline network uses the port pairs associated with optical protection switches located on bypass combo modules on GigaVUE-HC3, GigaVUE-HC1 and GigaVUE-HC1-Plus. The specialized hardware triggers a bypass when power is lost to the node. On GigaVUE-HC1 and GigaVUE-HC1-Plus, protected inline networks for copper are supported on the TAP-HC1-G10040 module.

Physical bypass provides protection against failure of the Gigamon node due to a loss of power. When physical bypass is activated, traffic flows from one side of the network to the other, but without monitoring.

Because of the specialized hardware, when a bypass combo module initializes, the system creates inline-network type ports for each of the protected port pairs on the module. This does not happen automatically on the TAP-HC1-G10040 module on the GigaVUE-HC1.

A protected inline network is shown in the following figure.



Mix of Protected and Unprotected

Any combination of protected and unprotected inline networks is supported in inline bypass solutions.

The GigaVUE-HC3 supports a 100Gb bypass combo module and starting in software version 5.2, the GigaVUE-HC3 100Gb bypass combo module supports dual speeds: 100Gb or 40Gb.

NOTE: You do not need to create the inline network ports for the protected port pairs on the bypass combo modules because they are created automatically when the bypass combo module initializes.

Other ports on the bypass combo module or on other line cards or modules can be designated as inline-network type ports. These ports will automatically be unprotected.

You also do not need to create the protected inline networks on the bypass combo modules because they are created automatically when the bypass combo module initializes. In addition, you cannot delete the protected inline networks on the bypass combo modules.

On GigaVUE-HC1, inline networks can also be configured on the copper TAP-HC1-G10040 module.

Once an inline network is created, it is fully configurable regardless of whether or not it is associated with an inline tool or an inline tool group through a map or map passall.

Use the Inline Networks page to display the configuration of inline networks as well as the current state of the inline bypass solution. The overall state of the inline bypass solution consists of the inline networks involved in the solution as well as the associated inline maps and inline tools. For details, refer to the following:

- Display Current State of Inline Bypass Solution

- How to Use SNMP Polling to Obtain Inline Network State
- SNMP Notification of Forwarding State Change
- How to Use Syslog to Obtain Inline Network State

For details on the parameters of inline networks, refer to the following:

- Network Port Link Status Propagation Parameter
- Traffic Path Parameter
- Physical Bypass Parameter.

Network Port Link Status Propagation Parameter

One of the parameters of inline networks is link status propagation, which controls the behavior of the link status for the inline network ports involved in a given inline network. The default is enabled.

When enabled, an inline network link failure on one side of the inline network will be propagated to the other side. For example, when the link is lost on one side of the network such that traffic cannot be sent to the inline tools, the link on the opposite side of the network is also brought down.

When the link is restored to the side that originally went down, the link will automatically be restored to the other side of the network. The GigaVUE node will not forward packets to the inline tools until the link is restored on both sides.

Link status propagation is enabled by selecting **Link Failure Propagation** when configuring an inline network port.

Traffic Path Parameter

One of the parameters of inline networks is **Traffic Path**, which specifies the path of the traffic received at an inline network port. The values are as follows:

- **Drop**—no traffic is exchanged through the inline network ports. All traffic to these ports is dropped. No traffic is forwarded to or from the inline tool or tools. No traffic is passed from inline network port A to inline network port B or from inline network port B to inline network port A. Refer to Figure 15Traffic Path of Drop.
- **ByPass**—there are two cases for bypass as follows, depending on the inline map configuration:
  - If there are no inline maps associated with the inline network or if the set of inline maps associated with the inline network guarantees that no traffic is dropped when the traffic path is set to **To Inline Tool**, then setting the traffic path to **Bypass** leads to the following: all traffic arriving at the side A inline network port is forwarded to the side B inline network port and all traffic arriving at the side B inline network port

is forwarded to the side A inline network port through a logical bypass. Refer to Figure 16Traffic Path of Bypass for a simple scenario involving a map-passall in which there is no possibility of traffic drops.

- If the set of inline maps associated with the inline network involves some traffic drop when the traffic path is set to **To Inline Tool**, then setting the traffic path to **Bypass** leads to the following: all traffic arriving at the side A inline network port that would not have been dropped with traffic path set to **To Inline Tool** is forwarded to the side B inline network port and all traffic arriving at the side B inline network port that would not have been dropped with traffic path set to **To Inline Tool** is forwarded to the side A inline network port through a logical bypass. In other words, packets that were dropped when the traffic path was set to **To Inline Tool** will also be dropped when set to **Bypass**.

In either of these two bypass cases, no traffic is forwarded to the inline tool or tools.

- **ByPass with Monitoring**—there are two cases for monitoring as follows, depending on the inline map configuration:
  - If there are no inline maps associated with the inline network or if the set of inline maps associated with the inline network guarantees that no traffic is dropped when the traffic path is set to **To Inline Tool**, then setting the traffic path to **ByPass with Monitoring** leads to the following: all traffic is forwarded as for bypass, but a copy of the traffic is forwarded to the inline tool or tools according to the configured maps between the inline network and the inline tool or tools. Refer to Figure 17Traffic Path of Monitoring for a simple scenario involving a map-passall in which there is no possibility of traffic drops.
  - If the set of inline maps associated with the inline network involves some traffic drop when the traffic path is set to **To Inline Tool**, then setting the traffic path to **ByPass with Monitoring** leads to the following: all traffic that would not have been dropped with traffic path set to **To Inline Tool** is forwarded as for bypass, but a copy of the traffic is forwarded to the inline tool or tools according to the configured maps between the inline network and the inline tool or tools.

  In either of these two monitoring cases, no traffic is taken from the inline tools.
- **To Inline Tool**—traffic received at the inline network ports is forwarded according to the following factors:
  - the configured maps between the inline network and the inline tools
  - the failover actions of the inline tool or tools
  - the health state of the inline tool or tools

  Refer to Figure 18Traffic Path of To-Inline-Tool.

The default is **Bypass**. This avoids any traffic loss when first configuring an unprotected inline network or when disabling physical bypass on a protected inline network.

Figure 15Traffic Path of Drop to Figure 18Traffic Path of To-Inline-Tool show the traffic path for a simple inline bypass solution with inline network ports (nA and nB), inline tool ports (tA and tB), and a map passall.

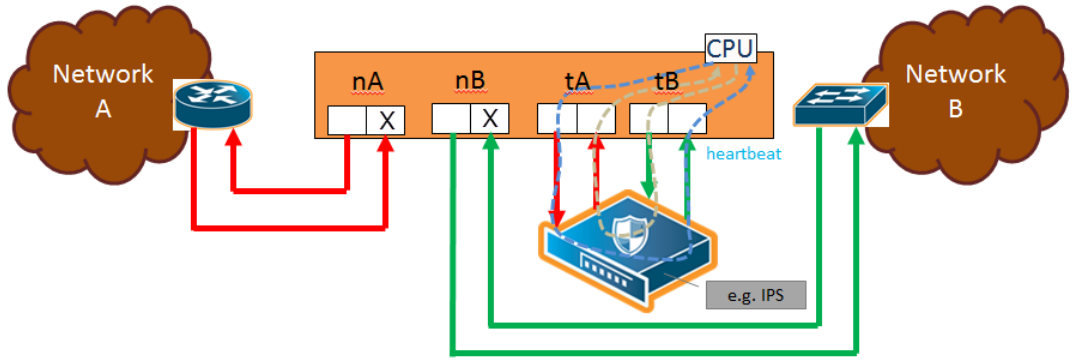


Figure 15 *Traffic Path of Drop*

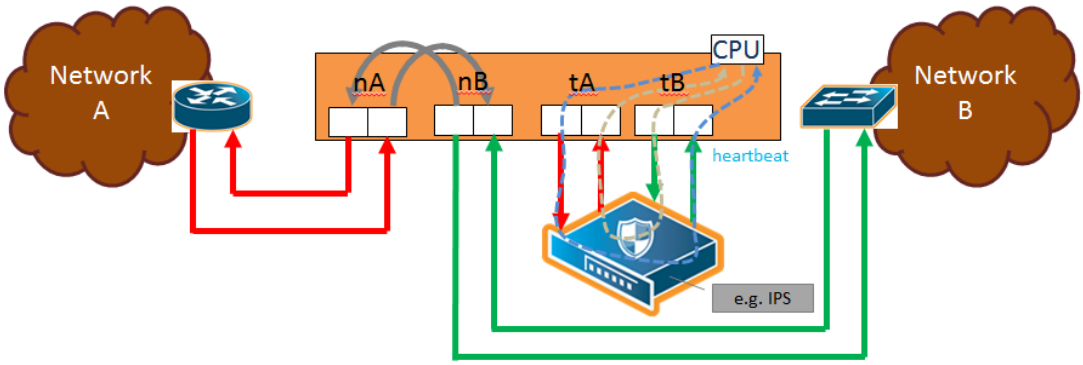


Figure 16 *Traffic Path of Bypass*

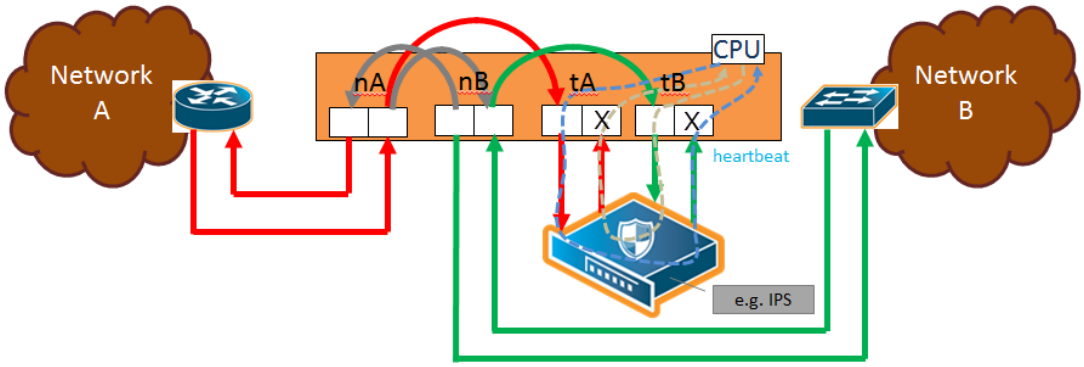


Figure 17 *Traffic Path of Monitoring*

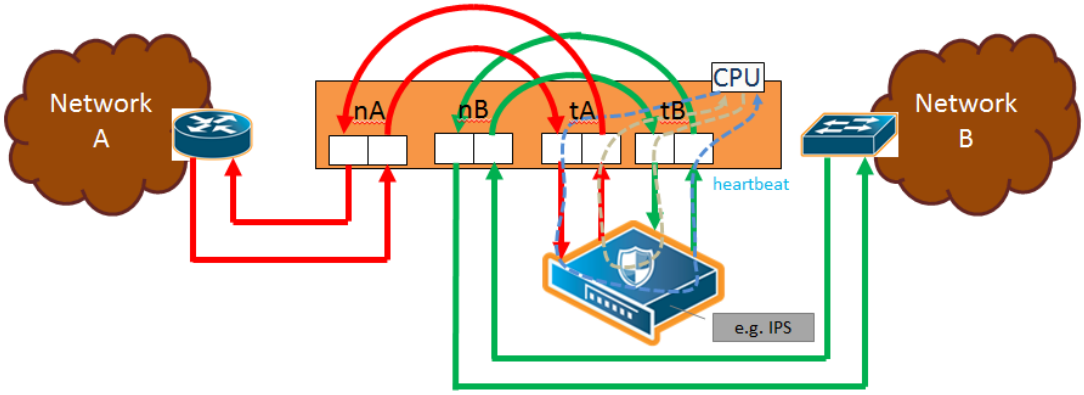


Figure 18 *Traffic Path of To-Inline-Tool*

Physical Bypass Parameter

One of the parameters of inline networks is physical bypass, which controls the state of the optical protection switch on the bypass combo module or copper TAP module when the module is powered on. The optical protection switch can have one of the following states:

- close—the fiber connected to the side A network port is passively coupled with the fiber connected to the side B port without any transceivers or switching fabric. Therefore, any traffic coming in on these fibers is exchanged between the two inline network ports without being noticed by the system.
- open—the fiber connected to the inline network ports is coupled through transceivers with the switching fabric that is under software control. Therefore, any traffic coming in on these fibers is subject to the traffic forwarding rules imposed by the current configuration as well as the current state of the inline tools.

When the bypass combo module or copper TAP modules is powered off, the optical protection switch is always in the close state.

When the bypass combo module or copper TAP modules is powered on, the state of the optical protection switch is as follows:

- the close state if **Physical Bypass** is set to enable (that is, selected on the configuration page)
- the open state if physical bypass is set to disable (that is, not selected on the configuration page)

Physical Bypass is enabled by default.

NOTE: Physical bypass only applies to protected inline networks.

Redundancy Profile Parameter

One of the parameters of inline networks is redundancy profile. A redundancy profile is used to configure Gigamon Resiliency for Inline Protection (GRIP). For more detailed information about GRIP, refer to theConfigure Gigamon Resiliency for Inline Protection.

To created a redundancy profile, doing the following:

1. Select **Inline Bypass > Redundancies**.
2. Click New to open the Add Redundancy Profile page.
3. Type a name for the profile in the Alias field to help identify the profile.
4. Select a stack port by clicking the **Signaling Port** field and selecting an available port.

The Signaling Port field specifies the used to signal the state of the two GigaVUE-HC3 nodes to each other. The ports provide the mechanism to detect loss of power in one of the GigaVUE-HC3 nodes.

5. Click in the **Protection Role** field and select a role.

   The protection role specifies the role of the GigaVUE-HC3, as primary, secondary, or suspended. The default is suspended. When suspended, the protection role is on hold. Changing a GigaVUE-HC3 from the primary role to the suspended role can be used to manually force one of the GigaVUE-HC3nodes to become active. The suspended role is also used when performing maintenance.

6. Click **Save**.

To add the redundancy profile to the inline network, select the profile from the Redundancy Profile drop-down list on the Inline Network configuration page.

Display Current State of Inline Bypass Solution

Inline networks, inline tools, and inline maps work together to form an inline bypass solution. The inline bypass solution has an overall state, which can change in response to hardware conditions and user configuration.

Several factors make up the overall state of an inline bypass solution, as follows:

- The physical bypass configuration of the inline network is protected.
- The inline network configuration, in particular, if physical bypass is enabled or disabled, what traffic path is configured, and if link failure propagation is enabled or disabled.
- The inline tool configuration, in particular, if the state of the inline tool is enabled or disabled, if there is a heartbeat profile configured and if the heartbeat is enabled or disabled, and what failover actions are configured.
- The inline tool group configuration, in particular, if the state of the inline tool group is enabled or disabled, what failover mode is configured, what failover action is configured, and what number of minimum healthy tools is configured.
- The status of the links attached to the inline network ports and inline tool ports.

Forwarding States and Determining Factors describes each forwarding state and the factors determining that state.

Whenever link failure propagation is configured as false (disabled), the inline network port status reflects the status of the respective far-end ports. Only when link failure propagation is configured as true (enabled) does this behavior change. Refer to the note under the Forwarding states for DISABLED and DISCONNECTED in Forwarding States and Determining Factors.

Table 16: Forwarding States and Determining Factors

physical-bypass	traffic-path	Status of far-end ports connected to inline network ports	Status of inline tool side	Description
enable	drop, bypass, monitoring, or to-inline-tool	any status	any status	Forwarding state—PHYSICAL BYPASS If physical bypass is enabled, all traffic is exchanged directly between side A network and side B network without any monitoring by the GigaVUE node. Only applies to protected inline networks.
disable	drop	any status	any status	Forwarding state—DISABLED If the inline network is configured with a traffic path of drop, no traffic is exchanged between side A network and side B network because all packets coming to the inline network ports are dropped. **Note**: If the inline network is configured with link failure propagation set to true (enabled), the status of the inline network ports will be determined by the status of the far-end ports connected to the inline network ports. If both far-end ports are up, then both inline network ports will be up. If any far-end ports are down, then both inline network ports will be down.
disable	bypass, monitoring, or to-inline-tool	at least one far-end port is down	any status	Forwarding state—DISCONNECTED When one of the inline network ports is down due to a link down caused by a far-end device, no traffic is exchanged between side A network and side B network. **Note**: If the inline network is configured with link failure propagation enable as true, the status of both inline network ports will be down. That is, if only one inline

physical-bypass	traffic-path	Status of far-end ports connected to inline network ports	Status of inline tool side	Description
				network port is down, the other will be brought down.
disable	bypass	both far-end ports are up	any status	Forwarding state—FORCED BYPASS If the inline network is configured with a traffic path of bypass, all traffic that would not have been dropped when the traffic path is set to to-inline-tool is exchanged directly between side A network and side B network through the switching fabric.
disable	monitoring	both far-end ports are up	any status	Forwarding state—FORCED BYPASS WITH MONITORING If the inline network is configured with a traffic path of monitoring, all traffic that would not have been dropped when the traffic path is set to to-inline-tool is exchanged directly between side A network and side B network through the switching fabric. If there is any inline map configured for the inline network, a copy of the respective traffic is directed to the respective inline tools according to the configured inline maps.
disable	to-inline-tool	both far-end ports are up	all inline tools involved in the inline bypass solution are operating as expected	Forwarding state—NORMAL If all inline tools involved in the inline bypass solution are enabled, all inline tool ports are up, and the inline tools operating with heartbeat protocol enabled have a heartbeat status of up, traffic flows as desired according to the configuration of the inline maps. **Note**: If there are no inline maps, setting the traffic path of a protected fiber inline network to to-inline-tool results in a NORMAL forwarding state, but the traffic sent to the source inline network ports will be dropped

physical-bypass	traffic-path	Status of far-end ports connected to inline network ports	Status of inline tool side	Description
				because there are no inline maps specifying destination tool ports to which to forward the traffic.
disable	to-inline-tool	both far-end ports are up	at least one inline tool involved in the inline bypass solution is disabled or associated with any inline tool port that is down or associated with the heartbeat protocol in a down state	Forwarding state—FAILURE-INTRODUCED BYPASS All traffic is exchanged directly between side A network and side B network through the switching fabric as a result of an inline tool port failure or heartbeat failure or inline tool being disabled that led to the network-level bypass due to the configured failover actions.
disable	to-inline-tool	both far-end ports are up	at least one inline tool involved in the inline bypass solution is disabled or associated with any inline tool port that is down or associated with the heartbeat protocol in a down state	Forwarding state—FAILURE-INTRODUCED DROP No traffic is exchanged between side A network and side B network. All packets coming to the inline network ports are dropped as a result of an inline tool port failure or heartbeat failure or inline tool being disabled that led to the network-level drop due to the configured failover actions.
disable	to-inline-tool	both far-end ports are up	at least one inline tool involved in the inline bypass solution is disabled or	Forwarding state—NETWORK PORTS FORCED DOWN No traffic is exchanged between side A network and side B network. The inline network ports are forced down as a result of an inline tool port failure

physical-bypass	traffic-path	Status of far-end ports connected to inline network ports	Status of inline tool side	Description
			associated with any inline tool port that is down or associated with the heartbeat protocol in a down state	or heartbeat failure or inline tool being disabled that led to the network ports being forced down due to the configured failover actions.
disable	to-inline-tool	both far-end ports are up	at least one inline tool involved in the inline bypass solution is disabled or associated with any inline tool port that is down or associated with the heartbeat protocol in a down state	Forwarding state—ABNORMAL Any condition of the inline bypass solution that does not meet the criteria of the forwarding states listed in this table. At least one inline tool passes traffic as desired. There are many scenarios that can lead to the abnormal forwarding state, for example, when one inline tool member of an inline tool group has failed, but the number of remaining healthy tools is still above the minimum required number of healthy tools.

How to Use SNMP Polling to Obtain Inline Network State

The overall inline network state can also be obtained through SNMP polling using an SNMP-compliant network management application or a MIB browser. The names of the MIB files that need to be loaded in order to poll the inline network state are: GIGAMON-COMMON-SMI and GIGAMONINLINEBYPASS.

The inline network forwarding states are described in Forwarding States and Determining Factors. Some of the states have slightly different names in SNMP.

SNMP Notification of Forwarding State Change

Use the following steps to configure a notification that will be sent when the inline bypass forwarding state changes:

1. Select **Settings > SNMP Traps**.
2. Click **Trap Settings**. The Edit SNMP Trap Settings page displays.
3. On the Edit SNMP Trap Settings page, select **Enable** for **Inline Bypass Forwarding State Change**.
4. Click **Save**.

How to Use Syslog to Obtain Inline Network State

The overall inline network state can also be obtained through syslog.

When the inline network state change is triggered by a configuration change, a sample syslog message is as follows:

```
IBFE STATE CHANGE inline network default_inline_net_1_1_1 status NORMAL changed
to PHYSICAL BYPASS triggered by config
```

When the inline network state change is triggered by a link status change or heartbeat status change, a sample syslog message is as follows:

```
IBFE STATE CHANGE inline network inNet1 status NORMAL changed to FAILURE
INTRODUCED BYPASS triggered by 1/x10
```

In the syslog messages, IBFE indicates the Inline Bypass Failover Engine.

The inline network forwarding states are described in Forwarding States and Determining Factors.

Inline Network Groups

Use the **Inline Network Group** configuration page to configure an inline network group. An inline network group is an arrangement of multiple inline networks that share the same inline tool or tools. Use this page to specify the list of inline networks in the inline network group.

The inline network ports that make up the inline networks participating in the inline network group are always in pairs, running at the same speed, on the same medium (fiber

or copper). All inline network ports of the inline network group must be on the same GigaVUE-HC3, or GigaVUE-HC1 or GigaVUE-HC1-Plus node. The inline networks participating in the inline network group can be different speeds and different mediums.

An inline network group can share an inline tool, or tools in an inline tool group or inline tool series. Many-to-one means from an inline network group to an inline tool and is shown in Figure 19Inline Network Group Many-to-One. Many-to-many means from an inline network group to an inline tool group or inline tool series.

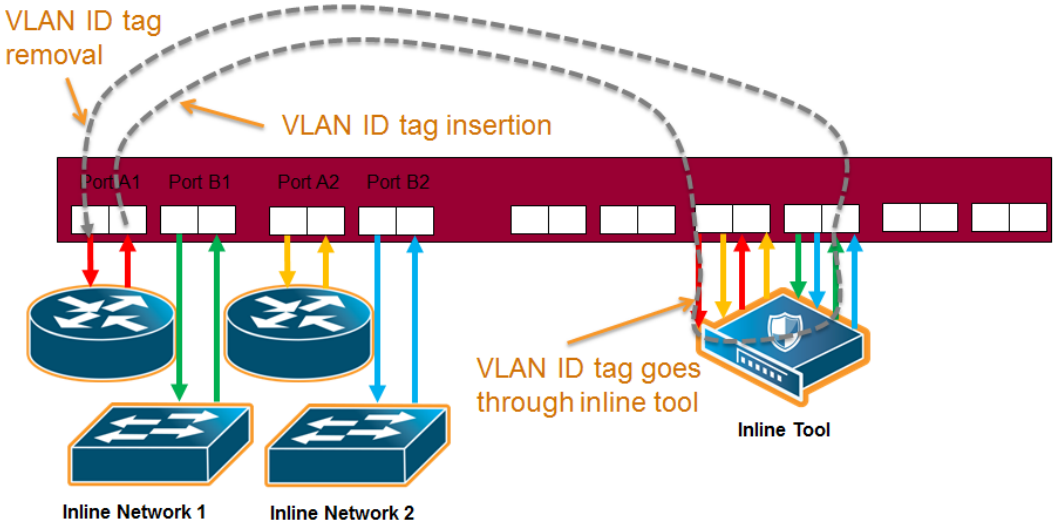


Figure 19 *Inline Network Group Many-to-One*

When an inline tool or inline tool group is configured between several pairs of networks, each pair operates independently of the others. The traffic coming from the inline tool or tools must be segregated into individual substreams according to the traffic source.

For example, the traffic that came into the GigaVUE node on inline network port net_A_3 must be sent through inline network port net_B_3 (with net_A_3 and net_B_3 belonging to the same inline tool that is the third member of the inline tool group). Refer to Figure 20Inline Tool Inserted in Multiple NetA-NetB Links.



Figure 20 *Inline Tool Inserted in Multiple NetA-NetB Links*

To accomplish the segregation, packets are tagged coming from an individual side A inline network port before they are sent to the inline tool or tools. When tagged packets come back from the inline tool or tools, the tag identifies the respective side B inline network port through which the packets should be sent. The tags are removed before the packets are sent through the side B inline network ports.

Traffic is guided to a particular inline network through internal VLAN ID tagging. This VLAN tagging affects packets only on their way from inline tool ports to the attached inline tool and back from the attached inline tool to the inline tool ports. The packets sent out from inline network ports remain untagged. Refer to Figure 19Inline Network Group Many-to-One.

NOTE: Internal VLAN ID tagging creates hidden VLAN ID tags. Explicit VLAN ID tagging is not needed, however starting in software version 4.6, explicit VLAN tagging is also available. Refer to Configurable VLAN Tagging.

Inline Tool Sharing

Inline network groups require inline tool sharing to be enabled on the inline tool or the members of the inline tool group or inline tool series specified in the inline map.

When shared is enabled (true), the inline tool can receive traffic from multiple sources (the inline networks in the inline network group) and can be used in a map in which the source is an inline network group.

An inline tool group or inline series does not have its own shared setting. The shared setting is derived from the inline tools. Therefore all the members in an inline tool group or inline series must have the same setting. For example, if an inline tool group has three inline tool members, the shared setting of all three inline tools must be the same.

Configurable VLAN Tagging

Explicit VLAN tagging for inline network groups can be configured. For example, you can use VLAN tags for managing policies. A mixture of internal and explicit VLAN tags is supported.

The VLAN tags are configured on the ports of inline networks. They can be configured at any time, but are only applied when the inline networks are part of an inline network group. Across the inline network group, the VLAN tags must be unique; however, both ports of an inline network can have the same VLAN tag. Refer to Tools in Bridge Mode.

An error message is displayed if the same VLAN tag is used for more than one inline network in an inline network group.

Refer also to Ingress and Egress VLAN.

NOTE: For inline SSL decryption, the inline network group does not support ingress VLAN tagging on the member links.

For out-of-band maps from inline network group ports or inline tool ports mapped from an inline network group, the out-of-band tool ports will receive the following:

- tagged packets, if they originally come from an inline network port with an ingress VLAN tag configured
- untagged packets, if they originally come from an inline network port without an ingress VLAN tag configured

Add VLAN Tag

The following are the steps for adding a VLAN Tag.

1. Go to **System > Ports > Ports > All Ports**.
2. Select the port to configure as an inline network port and click Edit.
3. Set the following parameters to configure an inline network port with VLAN tagging:
   - Select **Inline Network** or **Network** for Type.
   - Enter a VLAN ID in the **VLAN Tag** field.
4. Click **Save**.

Tools in Bridge Mode

The same VLAN tag can be assigned to both ports in an inline network port pair.

The following example configures the same ingress port VLAN tag on the net-a and net-b ports of an inline network. When the net-a and net-b ports have the same VLAN tag, an inline tool will send packets back to the network from which it came.

1. Go to **System > Ports > Ports > All Ports**.
2. Configure the net-a port.
   a. Select the port (for example, 1/1/x17) and click **Edit**.
   b. Enter inline-network-port-a in the **Alias** field.
   c. Select **Inline Network** for **Type**.
   d. Enter 123 in the VLAN Tag field.
   e. Click **Save**.
3. Configure the net-b port.
   a. Select the port (for example, 1/1/x18) and click **Edit**.
   b. Enter inline-network-port-a in the **Alias** field.

c. Select **Inline Network** for **Type**.

d. Enter 123 in the VLAN Tag field.

e. Click **Save**.

4. Configure the Inline Network.

a. Select Inline **Bypass > Inline Networks**.

b. Click **New**.

c. Enter an alias in the **Alias** field.

d. Select port 1/1/x17 for Port A.

e. Select port 1/1/x18 for Port B.

f. Click Save.

Inline Tools

There are two meanings to the term inline tool. The inline tool software construct consists of a pair of inline tool ports plus the inline tool attached to the ports. The software construct has attributes that are configured on the GigaVUE® HC Series node.

The term inline tool also refers to the pass-through device itself that performs packet inspection and selective forwarding, such as an Intrusion Protection System (IPS). This is a physical device, external to the GigaVUE node.

Use the Inline Tool page to configure the inline tool software construct. An inline tool consists of inline tool ports, always in pairs, running at the same speed, on the same medium (fiber or copper). The inline tool ports must be on the same GigaVUE-HC3, or GigaVUE-HC1 or GigaVUE-HC1-Plus node. The inline tool ports must also be on the same GigaVUE-HC3, or GigaVUE-HC1 or GigaVUE-HC1-Plus node as the inline network ports.

Inline tool ports are ports to which inline tools are attached. Inline tool ports are always in pairs, and must have the same line rate. Inline tool ports must be on the same GigaVUE-HC1 node as the inline network ports.

An inline tool consists of a pair of inline tool ports plus the inline tool attached to the ports. Refer to Figure 21Inline Tool and Inline Tool Ports.

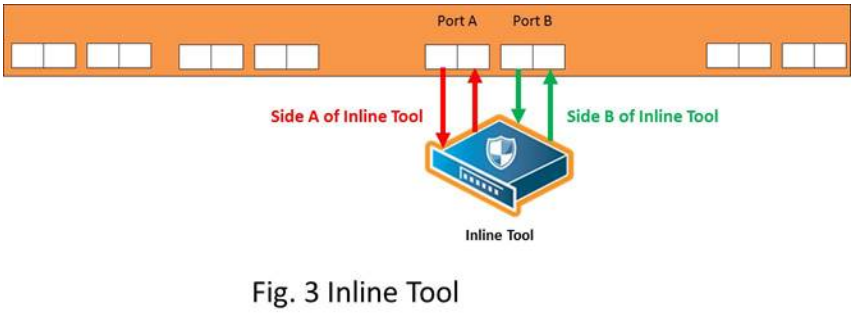


Figure 21 *Inline Tool and Inline Tool Ports*

Use the Inline Tools page to display of configuration of inline tools.

Inline Tool Failover Action

One of the parameters of inline tools is failover action, which controls the action taken when a tool is unhealthy or in response to a failure of an inline tool. You can configure one of the following failover actions:

- **ToolBypass**—when the inline tool fails, the traffic that normally was directed to the inline tool is redirected to the bypass path. Use this failover action for configurations involving multiple inline tools associated with an inline network or inline network group using rule-based maps. For configurations using map passalls, tool-bypass is the same as network-bypass.
- **NetworkBypass**—when the inline tool fails, all traffic that would not have been dropped when the inline network or networks had a NORMAL forwarding state is directed to the bypass path. That is, all such traffic arriving at the side A inline network port or ports is forwarded to the side B inline network port or ports and all traffic arriving at the side B inline network port or ports is forwarded to the side A inline network port or ports.
- **ToolDrop**—when the inline tool fails, the traffic that normally was directed to the inline tool is dropped. Use this failover action for configurations involving multiple inline tools associated with an inline network or inline network group using rule-based maps. For configurations using map passalls, tool-drop is the same as network-drop.
- **NetworkDrop**—when the inline tool fails, all traffic coming to the respective inline network (or inline network group) is dropped.
- **NetworkPortForcedDown**—when the inline tool fails, the inline network ports of the respective inline network (or inline network group) are forced down.

  NOTE: However, after the device reload, the inline network ports become operationally up, whereas the inline tool is still down. You need to manually force down the inline network ports.

The default is **ToolBypass**.

The bypass path is between side A and side B of the inline network ports.

Inline Tool Failover Action with Inline Flow Mapping®

When the inline bypass solution uses inline flow mapping, the failover actions of inline tools are as follows:

- **ToolBypass**—when the inline tool fails, the traffic that normally was directed to the inline tool is redirected to the bypass path. The traffic going to the healthy inline tools (through rule-based maps) remains unchanged.
- **ToolDrop**—when the inline tool fails, the traffic that normally was directed to the inline tool is dropped. The traffic going to the healthy inline tools (through rule-based maps) remains unchanged.

Inline Tool Recovery Mode

An inline tool detects failures in the traffic path between port pairs and automatically diverts traffic away to avoid disruption. After an inline tool goes down, the following modes specify how to bring it back up after it has recovered:

- **automatic**—Specifies automatic recovery, which redirects traffic back to the inline tool as soon as it has recovered from all faulty conditions.
- **manual**—Specifies manual recovery, which lets you control when to put an inline tool back into service after the tool has recovered. For example, you may wait for a maintenance window to return the inline tool to service.

The default is automatic.

By selecting the tool and selecting the Recover button, users can set the recovery of the inline tools to manual or automatic.

Refer to Figure 22Automatic and Manual Inline Tool Recovery from Failover for automatic and manual inline tool recovery from failover.

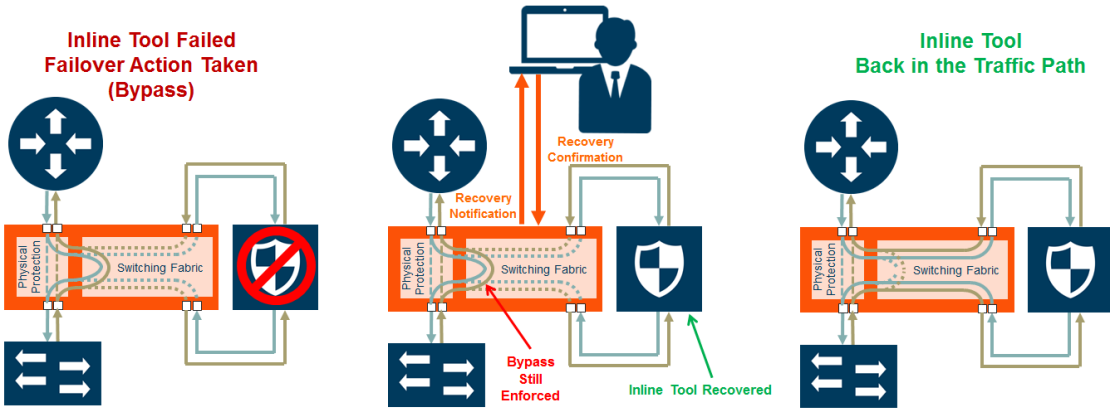


Figure 22 *Automatic and Manual Inline Tool Recovery from Failover*

The left side of Figure 22Automatic and Manual Inline Tool Recovery from Failover shows an inline tool that has failed and the bypass failover action has been executed.

Automatic recovery is shown on the right side of Figure 22Automatic and Manual Inline Tool Recovery from Failover. When the inline tool recovers, traffic is automatically directed back to it.

Manual recovery is shown in the center of Figure 22Automatic and Manual Inline Tool Recovery from Failover. When the recovery mode is configured as manual, an SNMP notification, if enabled, will send a notification when the inline tool is ready to be put back into service. The failover action, in this case, bypass, will be enforced until you manually put the inline tool back into service.

When the recovery mode is configured as manual, an SNMP notification, when enabled, will notify you when the inline tool is ready to be put back into service.

Use the following steps to configure notification that will be sent when the inline tool is ready:

1. Select **Settings > SNMP Traps**.
2. Click **Trap Settings**. The Edit SNMP Trap Settings page displays.
3. On the Edit SNMP Trap Settings page, select **Enable** for **Inlinetool Recovery**.
4. Click **Save**.

The default for **Inlinetool Recovery** is disabled.

Use the following steps to put an inline back in service when the recovery mode is manual and the inline tool has an operational state of ready.

1. Select Inline **Bypass > Inline Tools**.
2. On the Inline Tools page, select the inline tool.
3. Click **Recover**.

NOTE: Also use **Recover** after the GigaVUE® HC Series node is reloaded or rebooted (even though the inline tool has not failed). Issue the **Recover** on all the inline tools that are configured with manual recovery after a reload or **Actions > Shut Down** from the Chassis page on a selected card in Chassis Table View followed by **Actions > Start Up**. Refer to the *"Chassis"* section in the *GigaVUE Administration Guide* for more details.

NOTE: In certain circumstances while using classic inline, if the recovery mode is set to "manual" and the inline-tool is in "ready" state, traffic may still be forwarded to the inline-tool ports . This occurs because the recovery state shifts from "ready" to "up"

when "inline-tool disable/enable" is used. When the recovery mode for the inline-tool is "manual" and is in "ready" state, do not disable/enable the inline-tool . Instead, disable/enable the port.

The **Inline Tool** page displays the operational state of each inline tool as up, down, disable, or ready. Refer to Aggregate Inline Tool States for detailed descriptions of the states.

Inline Tool Sharing Mode

Inline tool sharing mode specifies how an inline tool is going to be shared as follows:

- **Enable**—Specifies that the inline tool is going to be shared by different sources.
- **Not enabled**—Specifies that the inline tool will not be shared by different sources.

The default is **not enabled**.

When sharing is enabled, the inline tool can receive traffic from multiple sources (the inline networks in the inline network group) and can be used in a map in which the source is an inline network group.

An inline tool group or inline series does not have its own shared setting. The shared setting is derived from the inline tools. Therefore all the members in an inline tool group or inline series must have the same setting. For example, if an inline tool group has three inline tool members, the shared setting of all three inline tools must be the same.

When an inline tool has sharing mode enabled, the traffic will be VLAN tagged. The connected inline device is expected to receive VLAN tagged packets. When an inline tool does not have sharing mode enabled, the extra VLAN tag is not added.

Go to **Inline Bypass > Inline Tools** and select an inline tool or click **New**. Under Configuration, Inline tool sharing mode, select Enable.

How to Use SNMP Polling to Obtain Inline Tool State

The inline tool state can also be obtained through SNMP polling using an SNMP-compliant network management application or a MIB browser. The names of the MIB files that need to be loaded in order to poll the inline tool state are: GIGAMON-COMMON-SMI and GIGAMONINLINEBYPASS.

The inline tool states are described in Table 17: Aggregate Inline Tool States. They are an aggregate of the inline tool port statuses and the heartbeat status.

Table 17: Aggregate Inline Tool States

Status	Description
up	The heartbeat is up and all the inline tool ports are up. The inline tool is operational and is forwarding traffic to the tool.
down	Either the heartbeat is down or one or more of the inline tool ports are down or disabled. The inline tool is in a failed state and is not ready to recover. The tool is not receiving any traffic.
disable	The inline tool is disabled.
ready	The inline tool is in a failed state but is ready to recover. The tool is not receiving any traffic.

Heartbeats

When heartbeat packets are sent to an inline tool, they are expected to be received back when the inline tool is healthy. Negative heartbeat packets complement heartbeat packets to verify the health of inline tools. When negative heartbeat packets are sent to an inline tool, they are not expected to be received back when the inline tool is healthy.

When some inline tools begin to fail, they allow packets though that should have been dropped. A negative heartbeat detects such a failure by sending a packet that should not pass through an inline tool. If the negative heartbeat packet passes through an inline tool, the tool is deemed to have failed. Therefore, a negative heartbeat packet received back from an inline tool indicates a tool failure.

Heartbeat packets and negative heartbeat packets can be used in any combination: heartbeat only, negative heartbeat only, neither, or both.

Rules and Notes

Keep in mind the following rules and notes when you configure the heartbeat mechanism between GigaVUE nodes:

- The heartbeat mechanism is supported in GigaVUE-HC1, GigaVUE-HC3 and GigaVUE-HC1-Plus devices.
- Negative heartbeat profiles are supported.
- Both standard and custom heartbeat packets are supported.
- Both protected and unprotected inline network pairs are supported.
- The heartbeat mechanism is not supported on inline netlag and inline network group components. However, it is supported on the underlying inline networks.

- The heartbeat mechanism is not supported for classic inline bypass solution.
- The minimum timeout for heartbeat sessions is 200 milliseconds, if the tool type is configured as GigaVUE Node.
- The size of a custom heartbeat packet must be less than 128 bytes.

Heartbeat Profiles

A heartbeat profile supports health monitoring of inline tools and inline networks. A heartbeat profile is a group of attributes applied to an inline tool or inline network to configure its heartbeat operation. Multiple inline tools can share a heartbeat profile. However, for inline network, heartbeat is supported only within an inline network pair; it is not supported on inline netlag or inline network groups.

To display the configured heartbeat profiles, select **Inline Bypass > Heartbeats** to open the Heartbeats page. An example is shown in Figure 23Heartbeats Page with Heartbeat Profiles.

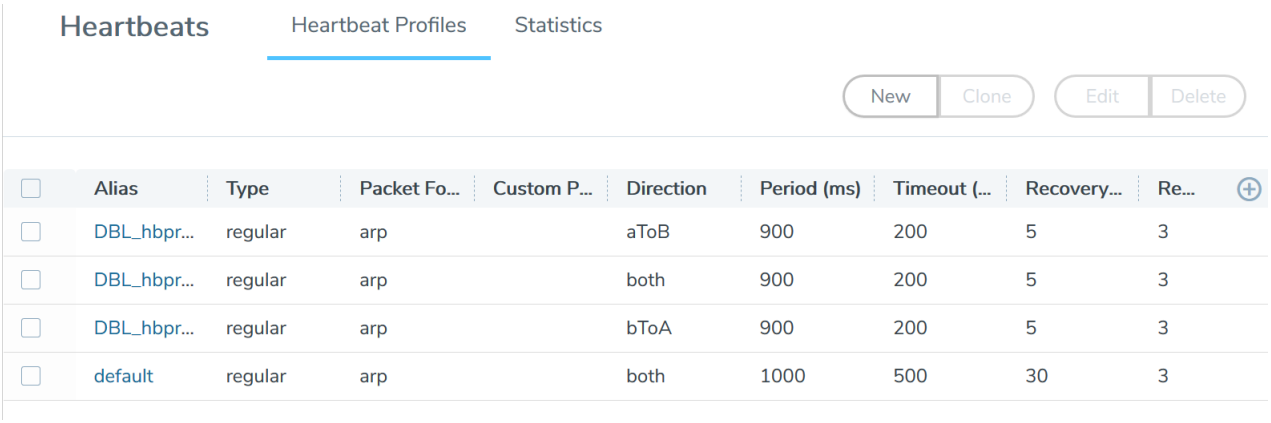


Figure 23 *Heartbeats Page with Heartbeat Profiles*

The Heartbeats page includes a default heartbeat profile that has the following settings:

- Alias—default
- Type—Regular
- Packet Format—arp or custom
- Custom Packet—URL from which a PCAP file can be imported
- Direction—bi-directional
- Period—1000 milliseconds
- Timeout—500 milliseconds
- Recovery Period—30 seconds
- Retries—3

The highest frequency heartbeat that can be configured is as follows:

- period—30 milliseconds
- timeout—20 milliseconds
- retry-count—0

The heartbeat mechanism supports the maximum number of inline tools, at the highest frequency, which is 48 on the GigaVUE-HC3, 16 on the GigaVUE-HC1 , and 48 in GigaVUE-HC1-Plus.

To display the heartbeat profile associated with an inline tool and the heartbeat status, open the Inline Tools page by selecting Inline **Bypass > Inline Tools**. There is also a combined heartbeat status, which combines heartbeat and negative heartbeat statuses and indicates the tool health used for inline tool failover actions or SNMP traps. The combined heartbeat status is the combination of both heartbeat statuses. If both are configured and one is down, the combined status will be down.

Standard Heartbeat

The standard heartbeat is a packet sent by the GigaVUE node that passes through the inline tool to verify that it is passing traffic, even if the link is *up*. If the packet is not passed through the tool, the tool is considered to have failed and a bypass action is triggered.

Even when the tool is considered down, heartbeat packets continue to be sent so that the bypass action can be reversed when the tool is healthy again.

Heartbeats are sent bidirectionally to the inline tool.

Standard or Custom Heartbeat Packet

The format of the heartbeat packet can be the standard ARP packet or a custom packet. For a custom packet, you must provide a URL from which a PCAP file can be imported. If the PCAP file contains several packets, the first packet present in the file is taken as the heartbeat packet. The size of a custom heartbeat packet must be less than 128 bytes.

NOTE: The system will overwrite the MAC address portion of the custom heartbeat packet.

If the inline tool through which the heartbeat packets are passed is expecting IPv6 traffic exclusively, you must select a custom heartbeat packet.

Custom heartbeat packets are needed in situations in which inline tools do not reliably pass standard ARP packets. For example, if an inline tool is configured to pass only IPv6 traffic, an ICMPv6 ARP packet might be appropriate.

If a custom heartbeat packet is specified, the **Heartbeats** page displays the name of the PCAP file from which it was imported.

Detect Inline Tool Failure

The health of the inline tool is critical to the proper handling of traffic. An inline tool is determined to have failed if:

- link is lost to the tool
- inline heartbeat fails

When the tool fails in one direction, it is considered to have failed in both directions. For example, if the heartbeat stops flowing in the northbound direction, neither northbound or southbound packets are sent to the tool.

Negative Heartbeat Profiles

A negative heartbeat profile is a group of attributes applied to an inline tool to configure its negative heartbeat operation. Multiple inline tools can share a negative heartbeat profile. The content of a negative heartbeat is configurable using the same PCAP file mechanism as for a custom heartbeat packet.

Use the Add Heartbeats page shown in Heartbeats to configure a negative heartbeat profile. A negative heartbeat profile can be created by selecting **Negative** in the **Type** field. The profile will have the following settings:

- **PacketFormat** set to **Custom**
- **Direction** set to **Bi-directional**
- **Period** set to **1000** (period is specified in milliseconds)
- **Recover Time** set to **30** (recovery is set in seconds)

You must provide a valid PCAP file when Packet Format is set to Custom before the negative heartbeat profile can be applied to an inline tool.

When a negative heartbeat is configured, the system will send packets specified by the **Custom Packet Format**, in the time specified by **Period**, in the direction specified by **Direction**. The inline tool absorbs the negative heartbeat packets until the number of

seconds specified by **Recovery Time** has passed. **Recovery Time** specifies the number of seconds of not receiving negative heartbeat packets in order for the inline tool to be declared healthy.

The negative heartbeat mechanism supports the maximum number of inline tools, which is 48 on the GigaVUE-HC3, 16 on the GigaVUE-HC1 , and 48 on the GigaVUE-HC1-Plus.

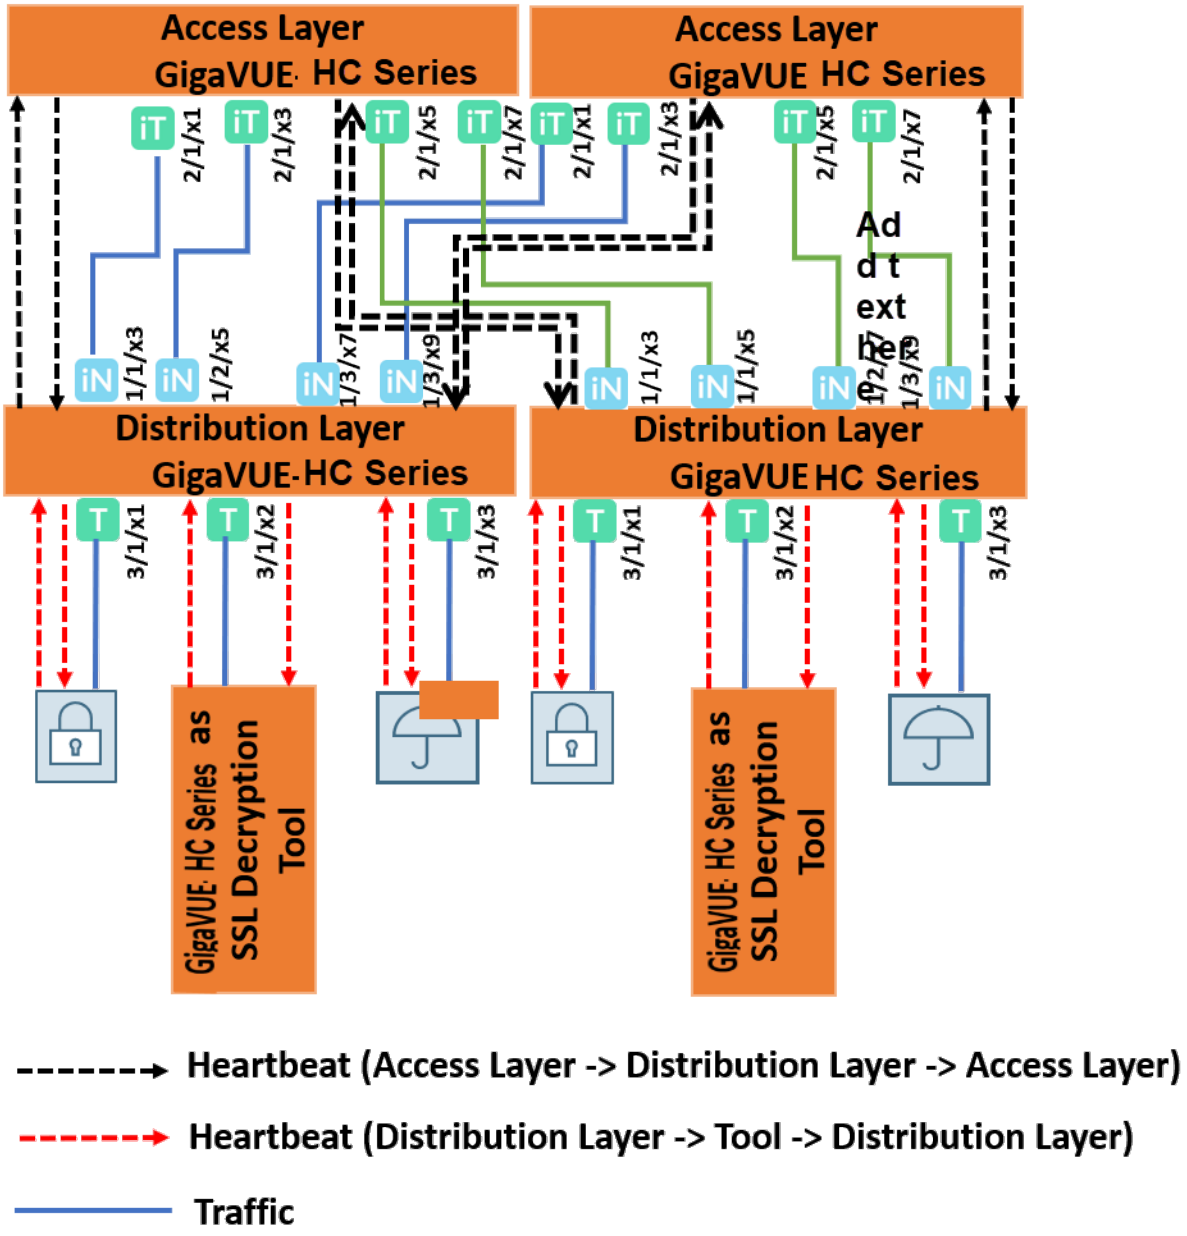
Use the Heartbeat page to display the configured negative heartbeat profiles. Use the Inline Tool page to display the negative heartbeat profile associated with an inline tool, the negative heartbeat status, and the counters of received packets in each direction. There is also a combined heartbeat status, which combines heartbeat and negative heartbeat statuses and indicates the tool health used for inline tool failover actions or SNMP polling. The combined heartbeat status is the combination of both heartbeat statuses. For example, if both are enabled and both are up, the combined status is up. If both are enabled and one is down, the combined status is down.

Heartbeat Support Between GigaVUE Nodes

The heartbeat mechanism focuses on providing extended heartbeat capability to monitor the following types of devices when the devices are connected to the inline-tool pair of ports as a tool:

- GigaVUE nodes
- GigaVUE nodes with GigaSMART operations configured

Following figure illustrates an example of a topology with GigaVUE nodes placed at three different layers.



GigaVUE HC Series means it can be either GigaVUE-HC3, GigaVUE-HC1 or GigaVUE-HC1-PLUS

The GigaVUE node at the access layer accesses the network traffic, gets the traffic processed by the tools at the tool layer, and transmits the processed traffic back to the network.

The GigaVUE node at the distribution layer distributes the traffic from the access layer to the tool layer.

The GigaVUE node at the tool layer acts as the SSL decryption tool.

In this topology, heartbeats are essential to monitor the traffic integrity at the distribution layer and to ensure automatic failover in case of a tool failure. In the access layer device, the ports that are connected to the distribution layer device are configured as inline tool ports because they face the tool side. In the distribution layer device, the ports that are connected to the access layer device are configured as inline network ports because they face the network side of the topology. The heartbeat packets will be sent from the inline tool port pair of the access layer device to the inline network port pair of the distribution layer device. If the forwarding state of the inline network pair is normal, the heartbeat packet is sent back to the inline tool port pair of the access layer device. Else, the packet is dropped.

The heartbeat mechanism is extended to support the GigaVUE node at the distribution layer to monitor the GigaVUE node that acts as a tool at the tool layer. In the distribution layer device, the ports that are connected to the tool layer device are configured as inline tool ports. In the tool layer device, the ports that are connected to the distribution layer device are configured as inline network ports. The heartbeat packets that are sent from the distribution layer device to the tool layer device will monitor the availability of both, the tool layer device and its GigaSMART engines.

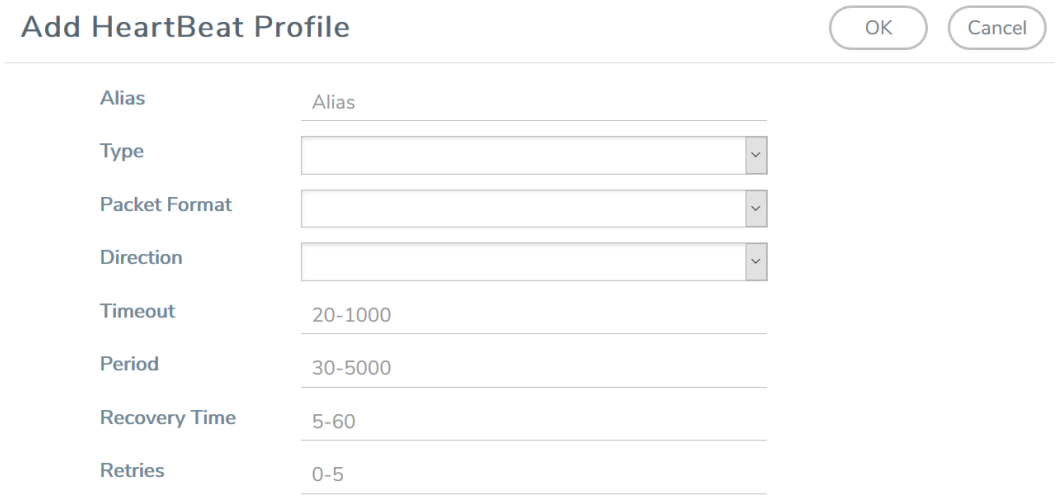
Troubleshoot Heartbeat Failures

In your topology, you may sometimes notice heartbeat packets getting dropped at the inline network or the heartbeat status of the inline tool is not up. For details about how to isolate and troubleshoot such failures, refer to the following sections:

- Heartbeat Status of Inline Tool is Not Operationally Up
- Heartbeat Packets are Dropped

Configure a Heartbeat Profile

Use the Add Heartbeats page shown in the below figure to configure a regular heartbeat profile or a negative heartbeat profile by selecting **Inline Bypass > Heartbeats > Heartbeats**, and then clicking **New**.



View Heartbeat Profile Statistics

Statistics about heartbeat profiles are displayed on the heartbeats statistics page. To open the Statistics page, select **Inline Bypass > Heartbeats > Statistics**. The page shows the following information:

Column	Description
Alias	The alias of the inline tool or inline network to which the heartbeat profile is associated with.
Inline Type	The inline type to which the heartbeat profile is associated with. The valid values are Inline Tool or Inline Network.
Heartbeat Profile	The alias of the heartbeat profile
Heartbeat Type	Indicates the type of heartbeat profile: Regular or Negative
A to B Packets	The number of packets sent/received from port A to port B of the inline tool.
B to A Packets	The number of packets sent/received from port B to port A of the inline tool.
A to B Packet Drops	The number of packets dropped from port A to port B of the inline tool or inline network.
B to A Packet Drops	The number of packets dropped from port B to port A of the inline tool or inline network.

You can clear the statistics for a specific heartbeat profile or all heartbeat profile. To clear the statistics for a heartbeat profile, select the required row, and then click **Clear**.

To clear the statistics for all heartbeat profiles, select **Clear All**, and the choose one of the following options:

- Select **Clear All Inline Network Heartbeat Stats** to clear the heartbeat statistics for all inline networks.
- Select **Clear All Inline Tools Heartbeat Stats** to clear the heartbeat statistics for all inline tools.
- Select **Clear All Negative Heartbeat Stats** to clear only the statistics for negative heartbeat profiles. The statistics for regular heartbeat profiles will remain.

Heartbeat Status after System Reload

Following a reload, there is a 5-minute delay for the system to stabilize before heartbeat packets are sent or received. During this delay, the heartbeat status is down.

Inline Tool Groups

Use the **Inline Tool Groups** configuration page to configure an inline tool group, which is an arrangement of multiple inline tools to which traffic is distributed based on hashing. In an inline tool group, traffic is shared. Each inline tool in the group receives a portion of the traffic. The distribution mechanism includes a way of dealing with failures of individual tools through traffic redistribution to the remaining healthy tools.

The inline tool ports that make up the inline tools participating in the inline tool group are always in pairs, running at the same speed, on the same medium (fiber or copper). All inline tool ports of the inline tool group must be on the same GigaVUE-HC3, or GigaVUE-HC1 and GigaVUE-HC1-Plus node, but can be on different modules on the node. The inline tool ports must also be on the same GigaVUE-HC3, or GigaVUE-HC1 and GigaVUE-HC1-Plus node as the inline network ports.

Inline tool groups can be configured as follows:

- non-redundant—multiple inline tools with no spare inline tool. Refer to Figure 24Inline Tool Group With No Spare, Non-Redundant.
- 1+1 redundancy—single inline tool with a spare inline tool. Refer to Figure 25Inline Tool Group With Spare, Redundant 1+1 Scenario.
- N+1 redundancy—multiple inline tools that are considered active, with a standby inline tool that is only used if one of the active inline tools fails. Refer to Figure 26Inline Tool Group With Spare, Redundant N+1 Scenario.

With 1+1 redundancy, an inline tool is paired with a standby tool. When there is a loss of link or a heartbeat failure on an active tool, the traffic will be sent to the standby tool with no loss. In addition, if the standby tool fails, you can configure what happens to the traffic in that case, such as drop it or forward it.

With N+1 redundancy, one tool is added to a group of N distributed inline tools. When any one of the N tools fails, the traffic from that tool is sent to the standby (or spare) tool with no loss. In addition, if the standby tool fails, you can configure what happens to the traffic in that case, such a redistribute it, send it to another tool, or declare a failure on the tool group.

Refer to Figure 24Inline Tool Group With No Spare, Non-Redundant for a non-redundant inline tool group.

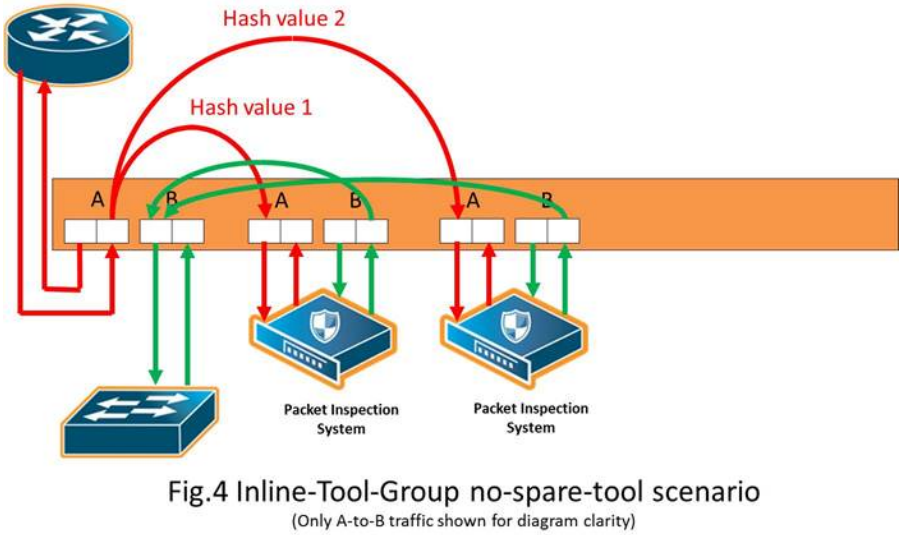


Figure 24 *Inline Tool Group With No Spare, Non-Redundant*

Refer to Figure 25Inline Tool Group With Spare, Redundant 1+1 Scenario for an inline tool group with a single inline tool and a spare inline tool configured. This is also referred to as 1+1 redundancy or N+1 redundancy where N equals 1. In Figure 25Inline Tool Group With Spare, Redundant 1+1 Scenario, traffic is only shown from A-to-B.

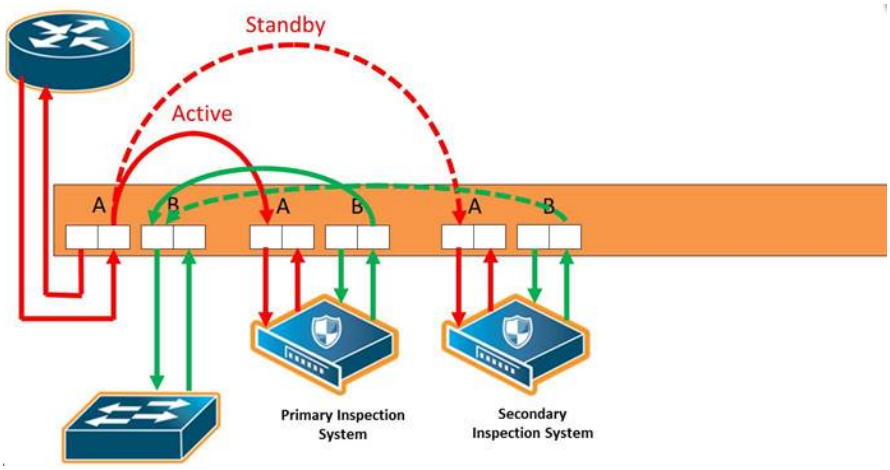


Figure 25 *Inline Tool Group With Spare, Redundant 1+1 Scenario*

Refer to Figure 26Inline Tool Group With Spare, Redundant N+1 Scenario for an inline tool group in an N+1 redundant scenario, in which N is greater than 1. In Figure 26Inline Tool Group With Spare, Redundant N+1 Scenario, traffic is only shown from A-to-B.

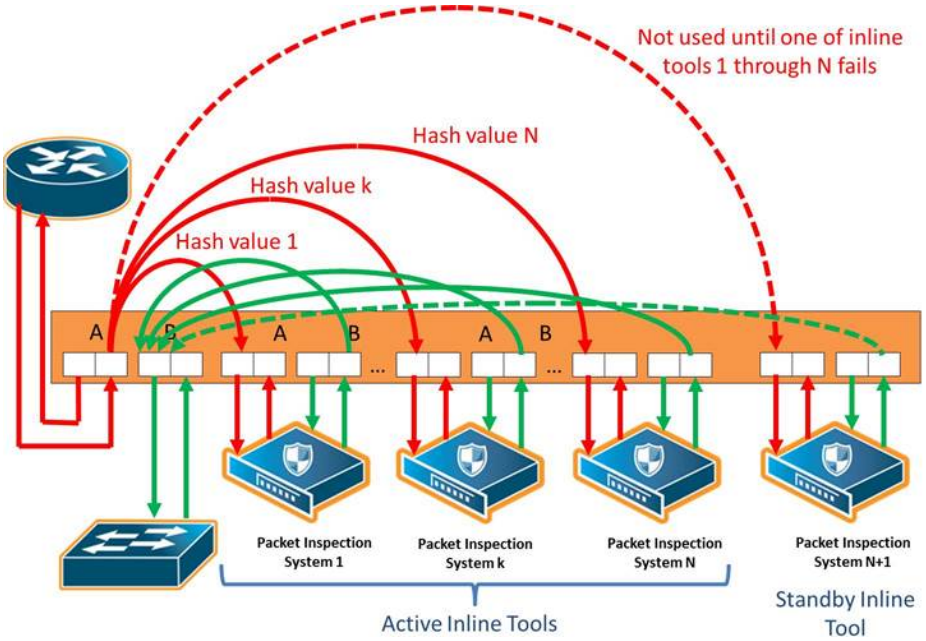


Figure 26 *Inline Tool Group With Spare, Redundant N+1 Scenario*

For details on the parameters of inline tool groups, refer to the following:

- Inline Tool Group Failover Action
- Inline Tool Group Spare Inline Tool
- Symmetrical and Asymmetrical Hashing.
- Resilient Weighted Hashing

Inline Tool Group Failover Action

One of the parameters of inline tool groups is the failover action, taken in response to a failure of an inline tool group, when the number of healthy inline tools in the inline tool group (including the spare inline tool, if configured) falls below the configured minimum. You can configure one of the following failover actions:

- **ToolBypass**—when the inline tool group fails, the traffic that normally was directed to the inline tool group is redirected to the bypass path. Use this failover action for configurations involving multiple inline tools or inline tool groups associated with an inline network or inline network group using rule-based maps. For configurations using map passalls, tool-bypass is the same as network-bypass.
- **NetworkBypass**—when the inline tool group fails, all traffic that would not have been dropped when the inline network or networks had a NORMAL forwarding state is directed to the bypass path. That is, all such traffic arriving at the side A inline network port or ports is forwarded to the side B inline network port or ports and all traffic arriving at the side B inline network port or ports is forwarded to the side A inline network port or ports.
- **NetworkPortForcedDown**—when the inline tool group fails, the inline network ports of the respective inline network (or inline network group) are forced down.
- **ToolDrop**—when the inline tool group fails, the traffic that normally was directed to the inline tool group is dropped. Use this failover action for configurations involving multiple inline tools or inline tool groups associated with an inline network or inline network group using rule-based maps. For configurations using map passalls, tool-drop is the same as network-drop.
- **NetworkDrop**—when the inline tool group fails, all traffic coming to the respective inline network (or inline network group) is dropped.

The default is tool-bypass.

The bypass path is between side A and side B of the inline network ports.

The failover action of all the inline tools specified by the inline tool list is overwritten by the failover mechanism of the inline tool group. This means that when a given inline tool specified by the inline tool list fails, the traffic originally directed to this inline tool is redirected to the spare inline tool (if one is configured and available) or handled according to the failover mode of the active tools, so long as the total number of healthy inline tools in the inline tool group is not smaller than the minimum required number of healthy inline tools.

When the total number of healthy inline tools in the inline tool group drops below the minimum required number of healthy inline tools, the failover action of the inline tool group determines the action to be taken.

Inline Tool Group Spare Inline Tool

One of the parameters of inline tool groups is a spare inline tool. If a spare is configured, the inline tool group becomes a redundant arrangement of inline tools. When the first failure occurs in a set of active inline tools, traffic will be forwarded to the spare with no loss, thus the spare will replace the failed tool in the active set.

The inline tools in the inline tool list are considered to be active inline tools. The traffic is hash-distributed over the active inline tools as long as all the inline tools are healthy, When one of the active inline tools fails, the spare inline tool takes the place of the failed inline tool and the new set operates as a new active set. If another inline tool fails, the traffic is redistributed according to the failover mode, as if there was no spare.

When the number of failed inline tools is such that the number of healthy inline tools is less than the minimum-group-healthy-size, the group heals itself by re-spreading the traffic over the healthy tools. When the number of healthy tools falls below the minimum-group-healthy-size, the failover action of the inline tool group takes place, while the failover action of the member inline tools is ignored.

The spare inline tool works with another parameter called release-spare-if-possible. When the inline tool that had been replaced with the spare inline tool recovers, the release-spare-if-possible parameter determines if the recovering inline tool is included in the active set of inline tools or if it becomes the new spare inline tool.

The default of the release-spare-if-possible parameter is disabled. Disabled means that even if the original inline tool recovers, the spare that replaced it will remain in the active set of inline tools. Enabled means that after the original inline tool recovers, the spare that replaced it will be released, if possible, from the active set of tools to become the spare again.

Configure the minimum-group-healthy-size and release-spare-if-possible parameters at the same time you configure the spare inline tool.

Symmetrical and Asymmetrical Hashing

One of the parameters of inline tool groups is hashing, which is used for distributing packets across a group of inline tools belonging to the inline tool group. The values for the hash parameter are as follows:

- **advanced**—Specifies symmetrical hashing, which is derived from the combination of packet fields based on the criteria selected for the advanced-hash algorithm configured under the **Ports > Port Groups > GigaStreams > Advanced Hash Settings** page.

For inline bypass applications, the most common choice of criteria for the advanced-hash algorithm is the combination of source IP and destination IP addresses. This produces a hash value that sends all traffic associated with the same session to the same inline tool in the inline tool group.

- **a-srcip-b-dstip**—Specifies asymmetrical hashing, which is derived from the source IP address for side A of the network and the destination IP address for side B of the network. This produces a hash value that sends all traffic associated with the same source address residing on side A to the same inline tool in the inline tool group, regardless of destination or session.
- **b-srcip-a-dstip**—Specifies asymmetrical hashing, which is derived from the destination IP address for side A of the network and the source IP address for side B of the network. This produces a hash value that sends all traffic associated with the same source address residing on side B to the same inline tool in the inline tool group, regardless of destination or session.

The default is **advanced**.

Use asymmetrical hashing if all traffic exchanged between a particular node on one side of the network and any nodes on the other side of the network that communicate with that node need to go to the same inline tool. The asymmetrical hashing options involve only source IP address (srcip) in one direction and only destination IP address (dstip) in the opposite direction. Bi-directional traffic, such as between a given user and all the Internet sites visited by the user, will be sent to the same inline tool in the group.

NOTE: When asymmetric hashing is configured, the portsrc and portdst packet fields are not included in the advanced-hash calculation for any GigaStream and inline tool groups across the GigaVUE node.

With symmetrical hashing, the inline network traffic path parameter can be configured to different values on the inline networks. With asymmetrical hashing, there is a restriction. Refer to Asymmetrical Hashing Restrictions.

Refer to Figure 27Asymmetrical Hashing for asymmetrical hashing.

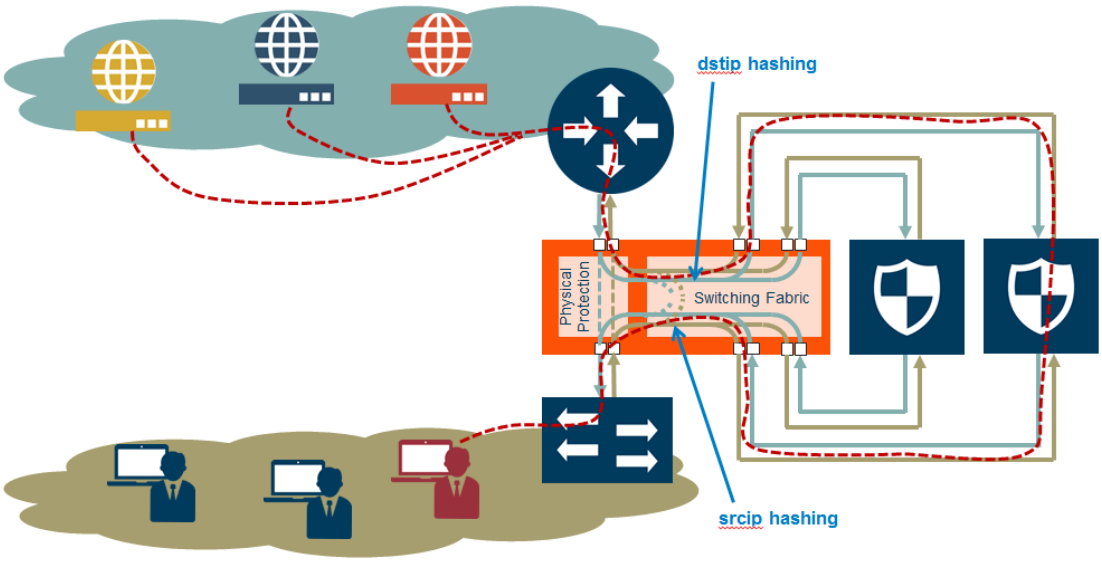


Figure 27 *Asymmetrical Hashing*

Use the hashing option **a-srcip-b-dstip** if the node is on side A of the network and the Internet is on side B. For example, Node A has IP address A. Traffic from Node A (from side A) will have IP address A. Traffic from side B (the Internet) destined for Node A, will have a destination of IP address A. This traffic will go to the same inline tool in the group.

If the network has the Internet on side A and the node on side B, use the hashing option **b-srcip-a-dstip**.

Asymmetrical Hashing Restrictions

The following are restrictions for asymmetrical hashing:

- If asymmetrical hashing is configured for the inline tool group, only rule-based maps or shared collector maps can be used to send traffic to the inline tool group. Inline map passalls cannot be used to send traffic to the inline tool group.
- For inline networks belonging to an inline network group, mapped to an inline tool group using asymmetrical hashing, the **Traffic Path** must be configured to the same value on all the inline networks, one of **Drop**, **Bypass**, **To Inline Tool**, or **ByPass with Monitoring**.

  NOTE: For the inline networks belonging to an inline network group, mapped to an inline tool group using symmetrical hashing, the traffic path parameter can be configured to different values on the inline networks.
- If an inline network is involved in an inline map to an inline tool group configured with asymmetrical hashing, the inline network ports of the inline network cannot be used as the **Source** in any out-of-band maps. Also, if the traffic path parameter for the inline network is configured to **ByPass with Monitoring**, there will not be any bypass traffic. All traffic will be forwarded to the inline tool group.

- When an inline tool group is included as a member of an inline series, asymmetrical hashing is not supported.
- If an inline network is involved in a flex-inline map and the inline tool used in the map is a part of an inline tool group; the traffic path parameter for the inline network or the flex traffic path of the inline tool from the inline tool group is set to **ByPass with Monitoring**, there will not be any bypassed traffic irrespective of the hashing. All traffic will be forwarded to the inline tool used from the inline tool group.

Resilient Weighted Hashing

One of the parameters of inline tool groups is weighting that provides you the ability to distribute traffic to the inline tools by assigning either an equal weighting or a custom weighting to the inline tools. You can assign custom weight in percentage or ratio. If an inline tool in a group goes down and the group maintains the **Minimum Healthy Group Size** that is defined for the group, the traffic is redistributed to the remaining tools based on the equal weighting or the custom weighting assigned to the tools. If the inline tool group does not meet the **Minimum Healthy Group Size** defined for the group, the traffic is redistributed based on the **Failover Action** defined for the group.

NOTE: Resilient hashing is not supported for classic inline maps.

The values for the weighting parameter are as follows:

- **Equal**—Traffic is distributed equally to all the inline tools in the inline tool group.
- **Relative**—Traffic is distributed to the inline tools in the inline tool group based on the relative weight or ratio assigned to the respective inline tools. The valid range is 1–256.
- **Percentage**—Traffic is distributed to the inline tools in the inline tool group based on the percentage assigned to the respective inline tools. The valid range is 1–100.

If you select **Relative** or **Percentage** as the weighting option, enter the hash weights for the inline tools that appear in the table below the **Weighting** drop-down list. Ensure that you assign a hash weight for each inline tool in the inline tool group.

Inline Serial Tools

Use the **Inline Serial Tools Groups** configuration page to configure inline tools in a series, in which the traffic from one side of the inline network is guided through the members of the inline tool series before it is sent out the other side of the inline network. With inline tools and inline tools groups arranged in a series, the traffic from one inline tool or inline tool group flows to the next, so all tools see the same traffic.

The inline tool ports that make up the inline tools and inline tool groups participating in the inline tool series are always in pairs, running at the same speed, on the same medium (fiber or copper). All inline tool ports of the inline tool series must be on the same GigaVUE-HC3, or GigaVUE-HC1 and GigaVUE-HC1-Plus node. The inline tool ports and inline tool groups must also be on the same GigaVUE-HC3, or GigaVUE-HC1 and GigaVUE-HC1-Plus node as the inline network ports.

To view the currently configured inline tool series, select **Inline Bypass > Inline Serial Tools** to open the Inline Serial Tool Groups page shown in the following figure.

Refer to Figure 28Inline Serial Tools for an illustration of an inline tool series. In Figure 28Inline Serial Tools, traffic is only shown from A-to-B.

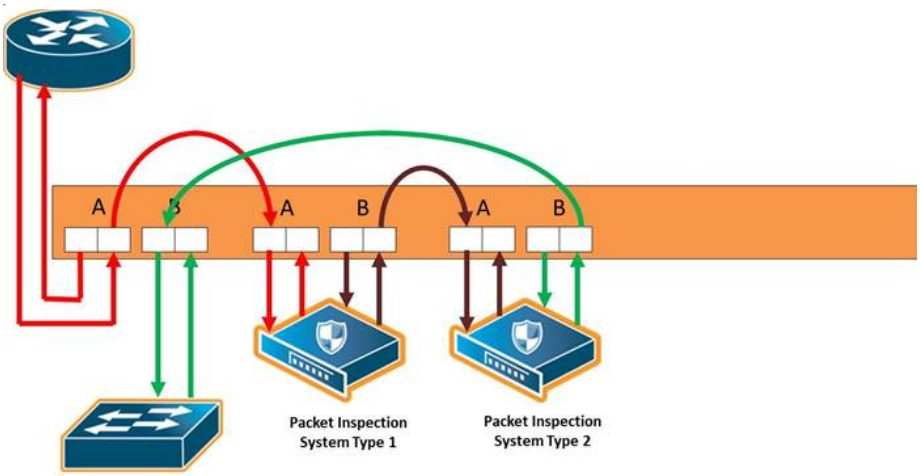


Figure 28 *Inline Serial Tools*

Refer to Figure 29Inline Tool Series, Including Inline Tool Group for an inline tool series that includes an inline tool group in addition to individual inline tools. The per-direction-order is set to forwarding. In Figure 29Inline Tool Series, Including Inline Tool Group, the inline tool group is placed as the middle member of the series, but it could be placed as the first member of the series or the last member of the series as well. In the inline tool group, traffic is shared.

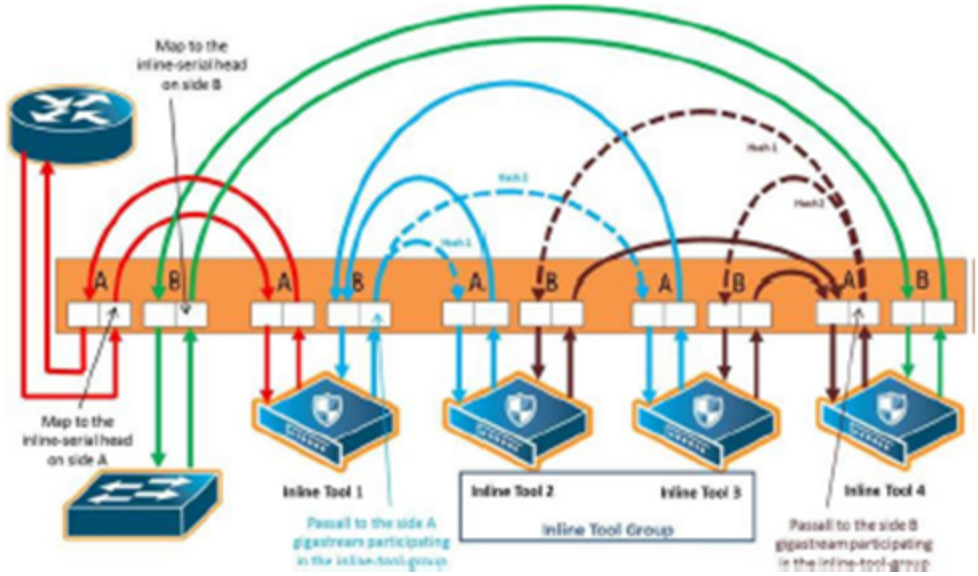


Figure 29 *Inline Tool Series, Including Inline Tool Group*

NOTE: For inline SSL decryption, the inline tool series does not support an inline tool group in the series.

The number of inline tools and inline tool groups in the inline series is limited only by the number of ports available for creating the inline networks and inline tools participating in the inline bypass solution.

When an inline tool group is included as a member of a inline series:

- a spare inline tool can be configured on the inline tool group
- inline maps to individual members of an inline tool group are not supported
- asymmetrical hashing is not supported, which means that the hash options, ascrip-b-dstip and b-scrip-a-dstip, are not allowed on the inline tool group.

To configure an inline serial tool group, do the following:

1. Open the **Inline Serial Tool Group** configuration page by selecting **Inline Bypass > Inline Serial Tools** from the main navigation pane, and then clicking **New**.
2. Enter a name for the inline serial tool group in the **Alias** field to identify the group and an optional description in the **Description** field.
3. Select and order the inline tools for the inline tool group.
   a. Click in the **Inline Tools** fields. The drop-down list shows the available inline tools
   b. Select the inline tools or in or inline tool group to add to the inline serial tool group.

      The inline tools are displayed in the order that they are selected. To change the order, click the up and down arrows.
4. Click **Save**.

NOTE: An inline serial tool group cannot be edited after it is saved.

Inline Serial Tools Global Failover Action

One of the parameters of inline tool series is the failover action taken in response to a failure of the inline tool series as a whole. This is referred to as the global failover action.

Each inline tool or inline tool group in the series can also have its own failover action. This is referred to as the local failover action. Refer to Inline Tool Series Local Failover Action for details.

For global failover actions, an inline tool series is declared to be in a failure condition as soon as any of its member inline tools goes into a failure condition. An inline tool series recovers from a failure condition after all the member inline tools recover from their failure conditions. The failover action attributes of the individual inline tools participating in an inline tool series are ignored. Instead, the **Failover action** configured on the Inline Serial Tool Group page for the inline tool series is executed. The values for global failover actions are as follows:

- **ToolBypass**—when the inline serial tools fails, the traffic that normally was directed to the inline tool is redirected to the bypass path. Use this failover action for configurations involving an inline serial tools that is associated with an inline network using rule-based maps. For configurations using map passalls, tool-bypass is the same as network-bypass. Refer to Figure 30Inline Tool Series Global Bypass Failover Action.
- **NetworkBypass**—when the inline serial tools fails, all traffic that would not have been dropped when the inline network or networks had a NORMAL forwarding state is directed to the bypass path. That is, all such traffic arriving at the side A inline network port or ports is forwarded to the side B inline network port or ports and all traffic arriving at the side B inline network port or ports is forwarded to the side A inline network port or ports. Refer to Figure 30Inline Tool Series Global Bypass Failover Action.
- **NetworkPortForcedDown**—when the inline serial tools fails, the inline network ports of the respective inline network are forced down. Refer to Figure 31Inline Tool Series Global Network Ports Forced Down Failover Action.
- **ToolDrop**—when the inline serial tools fails, the traffic that normally was directed to the inline tool is dropped. Use this failover action for configurations involving an inline serial tools that is associated with an inline network using rule-based maps. For configurations using map passalls, tool-drop is the same as network-drop. Refer to Figure 32Inline Tool Series Global Drop Failover Action.
- **NetworkDrop**—when the inline serial tools fails, all traffic coming to the respective inline network (or inline network group) is dropped. Refer to Figure 32Inline Tool Series Global Drop Failover Action.

The bypass path is between side A and side B of the inline network ports.

The default is **ToolBypass**.

Any failure of any member leads to the failure of the inline serial tools, hence the failover action for an inline serial tools will overwrite the failover action of the inline tool members of the series.

Figure 30Inline Tool Series Global Bypass Failover Action to Figure 32Inline Tool Series Global Drop Failover Action show the global failover actions for the inline series when any individual inline tool in the series fails.

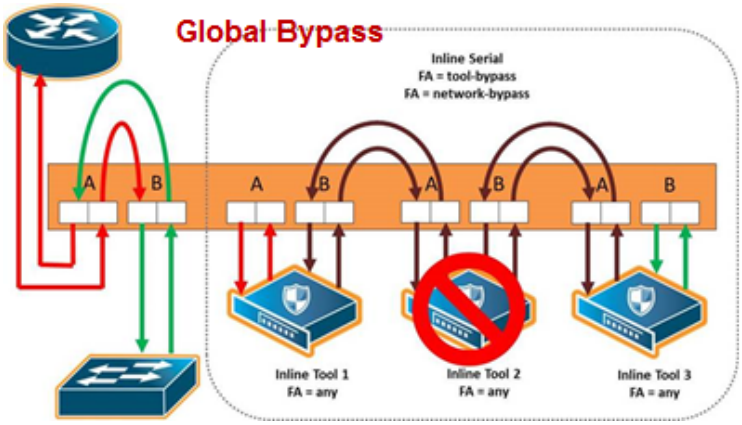


Figure 30 *Inline Tool Series Global Bypass Failover Action*

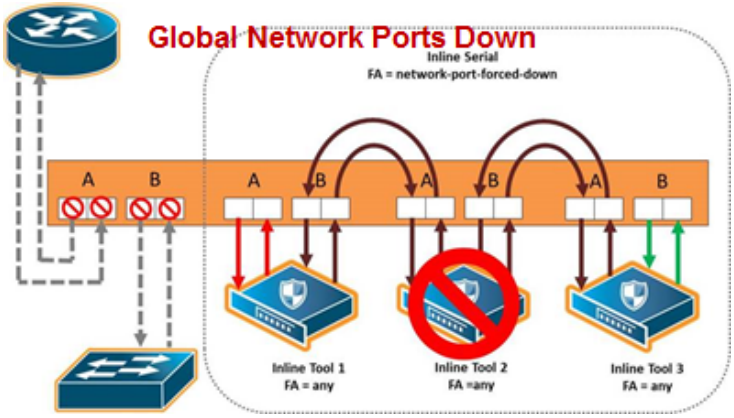


Figure 31 *Inline Tool Series Global Network Ports Forced Down Failover Action*

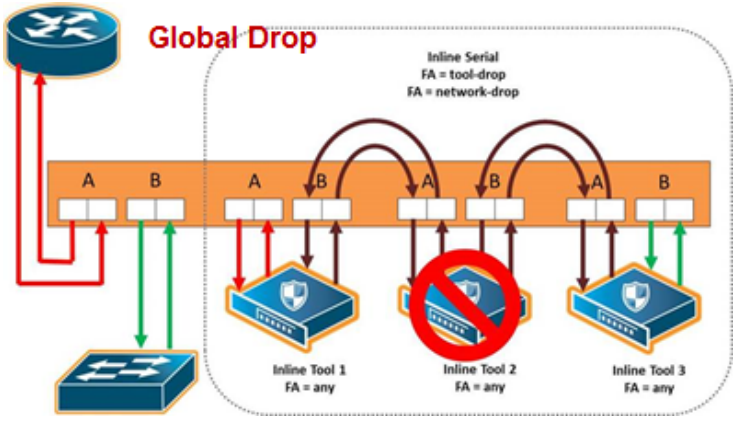


Figure 32 *Inline Tool Series Global Drop Failover Action*

Inline Tool Series Local Failover Action

Each inline tool in the series can have its own local failover action. When an individual inline tool or inline tool group in the series fails, the action taken depends on the failover action of the individual inline tool.

To configure local failover actions, configure a failover action of **Per Tool** for the series as a whole. Then the individual failover action for each inline tool in the series, as configured with the **inline Tool failover action**, takes effect. For details on the values, refer to Inline Tool Failover Action.

For example, if the failover action of the inline series is configured as **Per Tool** and the failover action of an individual inline tool in the series is configured as **ToolBypass**, when that tool in the series fails, the traffic will skip over the failed tool.

NOTE: For inline SSL decryption, the **Per Tool** failover action is not supported.

The values for local failover actions are as follows:

- **ToolBypass**—when the inline tool fails, the traffic bypasses the failed tool. That is, the traffic originally coming to the inline tool is diverted to the next inline tool in the series or to the appropriate inline network port if the inline tool is the last in the series. Refer to Figure 33Inline Tool Series Local Tool Bypass Failover Action.

  NOTE: When all the inline tools in a series are configured as **ToolBypass** and they all fail, this is the same as the failover action of **ToolBypass** for the series.

- **ToolDrop**—when the inline tool fails, traffic to this inline tool stops being forwarded. Effectively, this has the same result as the failover action of **ToolDrop** for the series as a whole, although the healthy members of the series will still receive traffic in one of the directions. Refer to Figure 34Inline Tool Series Local Tool Drop Failover Action.
- **NetworkBypass**—when the inline tool fails, a bypass is established between the inline network ports. Refer to Figure 35Inline Tool Series Local Network Bypass Failover Action.
- **NetworkDrop**—when the inline tool fails, traffic is dropped at the inline network ports. Refer to Figure 36Inline Tool Series Local Network Drop Failover Action.
- **NetworkPortForcedDown**—when the inline tool fails, the links for the inline network ports are brought down. Refer to Figure 37Inline Tool Series Local Network Ports Forced Down Failover Action.

Figure 33Inline Tool Series Local Tool Bypass Failover Action to Figure 37Inline Tool Series Local Network Ports Forced Down Failover Action show the local failover actions when an individual inline tool in a series fails

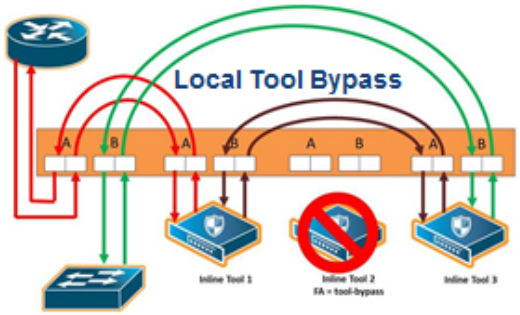


Figure 33 *Inline Tool Series Local Tool Bypass Failover Action*

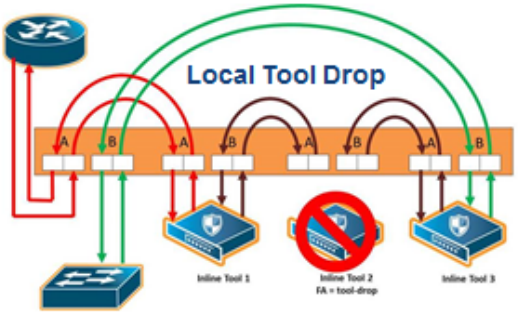


Figure 34 *Inline Tool Series Local Tool Drop Failover Action*

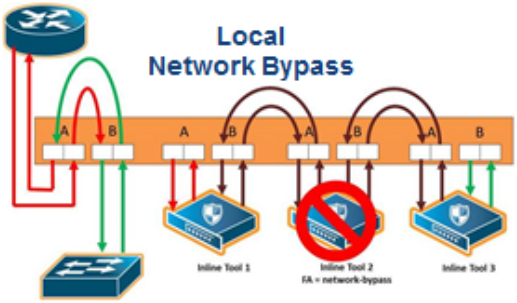


Figure 35 *Inline Tool Series Local Network Bypass Failover Action*

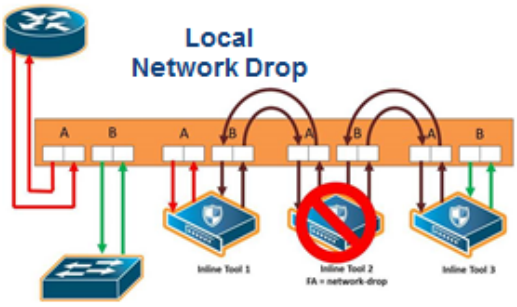


Figure 36 *Inline Tool Series Local Network Drop Failover Action*

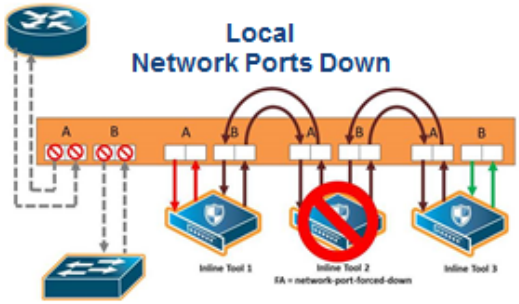


Figure 37 *Inline Tool Series Local Network Ports Forced Down Failover Action*

Figure 38Inline Tool Series Local Failover of Two Tools shows the failure of two individual inline tools in a series with different configured failover actions.

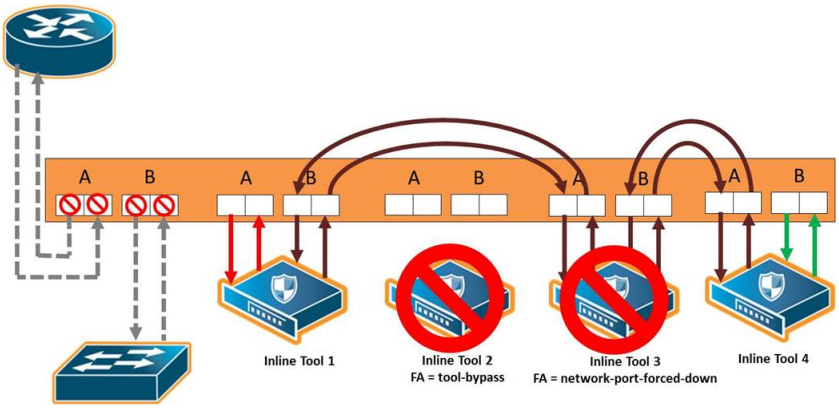


Figure 38 *Inline Tool Series Local Failover of Two Tools*

Inline Tool Series Per-Direction Order

One of the parameters of inline tool series is the per-direction order of the inline tool series. This parameter configures the traffic direction order of side B traffic with respect to the inline tool list, that is, the direction of the return traffic.

The **Return Direction** options on the Inline Serial Tool Group page specify the per-direction order of the side B traffic of the inline tool series as follows:

- **Reverse** specifies that the traffic from network B will flow through the inline tool list in reverse order, for example, from the third tool, to the second tool, to the first tool. This specifies the reverse order of inline tools for both directions.
- **Forward** specifies that the traffic from network B will flow through the inline tool list in the order it which it is defined, for example, from the first tool, to the second tool, to the third tool. This specifies the same order of inline tools for both directions of traffic.

The default is **Reverse**.

Figure 39Inline Tool Series Per-Direction Order: Simplified Reverse and Forward shows a simplified view of the flow through the tools:

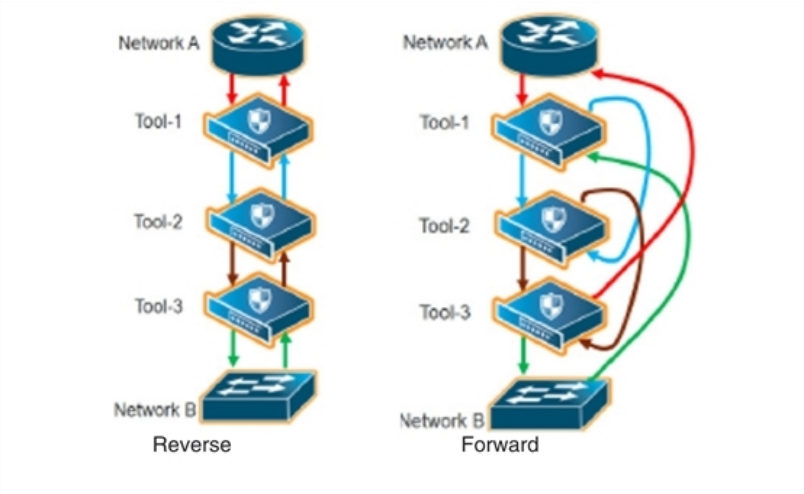


Figure 39 *Inline Tool Series Per-Direction Order: Simplified Reverse and Forward*

In Figure 39Inline Tool Series Per-Direction Order: Simplified Reverse and Forward, **Reverse** is on the left and **Forward** is on the right. Traffic from network side A to network side B for both reverse and forward flows from the first tool, to the second tool, to the third tool. But traffic from network side B to network side A with reverse, flows from the third tool, to the second tool, to the first tool, whereas traffic from network side B to network side A with forward, flows from the first tool, to the second tool, to the third tool.

Figure 40Inline Tool Series Per-Direction: Reversed shows the reverse direction with the tools connected to the GigaVUE node.

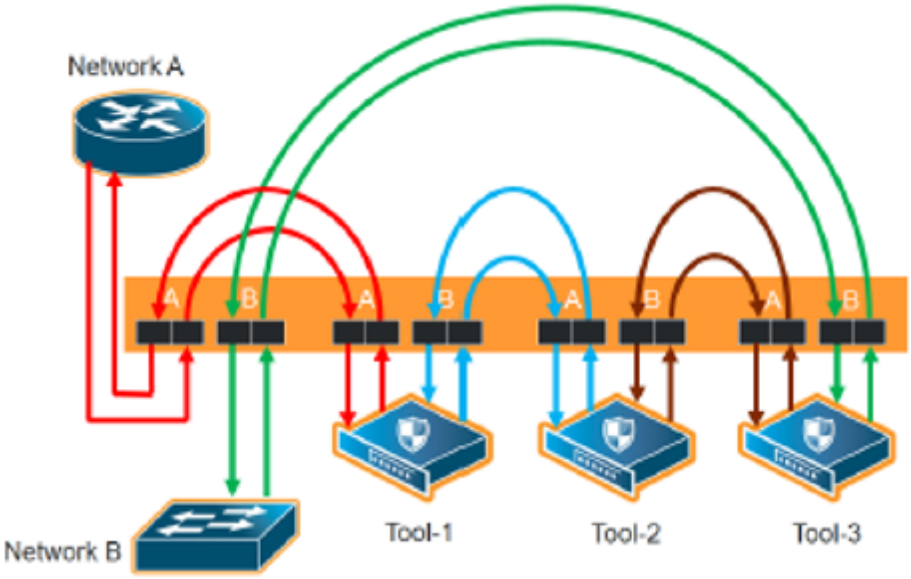


Figure 40 *Inline Tool Series Per-Direction: Reversed*

Figure 40Inline Tool Series Per-Direction: Reversed shows the forward direction with the tools connected to the GigaVUE node.

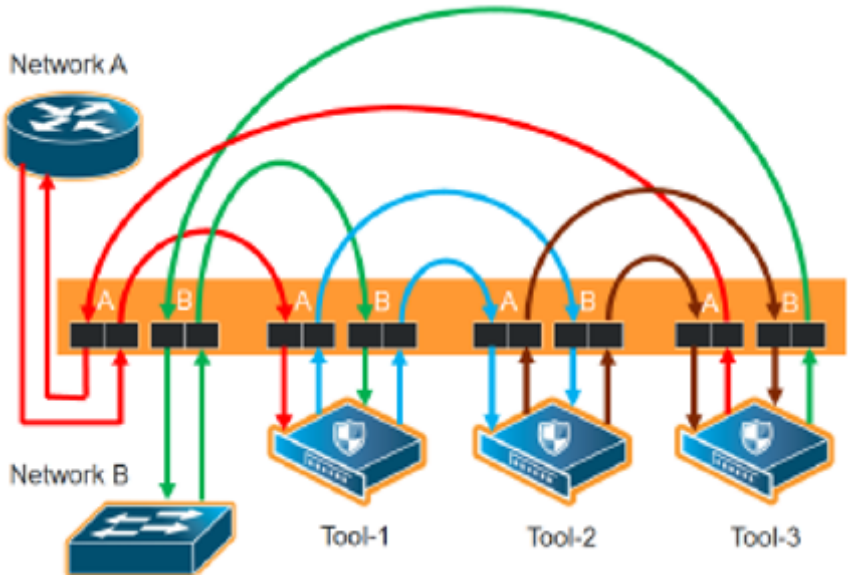


Figure 41 *Inline Tool Series Per-Direction: Forward*

Associate Inline Networks with Inline Tools Using Inline Maps

Inline networks and inline network groups are associated with inline tools, inline tool groups, and inline serial tools through inline maps. An inline map is a regular map, but the **Source** and **Destination** fields specify inline software constructs instead of port lists.

On the **Edit Map** or **New Map** page, the **Source** field specifies the inline network alias or inline network group alias. The **Destination** field specifies an inline tool alias, an inline tool group alias, an inline tool series alias, or an inline bypass. An inline bypass is a special construct that is used as a pseudo-inline tool to allow a portion of traffic to bypass any inline tools.

Maps can associate a network with multiple inline tools or they can associate multiple inline networks with the same inline tool or with multiple inline tools.

With inline maps, only the traffic that meets the map rules is sent to the tools, to bypass, or to shared collectors. For example, you can send traffic to tools for which they are specialized and send the rest to bypass. Or, if there is a type of traffic in which the tools are not interested or do not understand, that traffic can be sent to a shared collector.

NOTES:

- When an inline network is mapped to an inline tool or inline tool group, a second inline network cannot be mapped to the same inline tool. (In other words, an inline tool can be used in only one map.) However, when there are multiple inline networks, use an inline network group to map to the same inline tool.
- If an inline tool is already specified in a map, that tool cannot be included in an inline tool group (unless the map is first deleted).
- Inline network ports, inline tool ports, and out-of-band tool ports that are used in map configuration must all be configured on the same GigaVUE-HC3, or GigaVUE-HC1 or GigaVUE-HC1-Plus node. Even if nodes are in a cluster, the inline ports cannot be on different nodes.

Refer to the following sections:

- Inline Map Passall
- Inline Map
- Inline Map Shared Collector
- Inline Maps to Individual Members of an Inline Tool Group
- Out-of-Band (OOB) Map
- Symmetric and Asymmetric Maps

Inline Map Passall

Use the **Pass All** subtype to configure a type of inline map that passes all traffic. Map-passalls facilitate the sending of traffic between inline network ports (through inline tools or bypass). All of the inline network ports and inline tool ports involved in an inline map must be located on the same Gigamon node.

Use the **Source** of the **Pass All** to specify a single inline network.

Use the **Destination** of the **Pass All** to specify an inline tool alias, an inline tool group alias, an inline tool series alias, or an inline bypass. An inline bypass is a special construct that is used as a pseudo inline tool to allow a portion of traffic to bypass any inline tools.

An inline bypass, using the physical bypass option, is only valid if the **Source** is an inline network port. This applies only to the **Pass All** subtype for asymmetrical scenarios.

Inline Map

Use the **New Map** page to configure a type of inline map that uses rules to direct traffic. These inline maps are referred to as rule-based maps. All of the inline network ports and inline tool ports involved in an inline map must be located on the same GigaVUE node.

Use the **Source** field to specify an inline network alias or an inline network group alias.

Support for rule-based maps is limited to symmetric scenarios, which means that the **Destination** of rule-based inline maps can only be aliases of inline networks or inline network groups (not individual ports).

Use the **Destination** to specify an inline tool alias, an inline tool group alias, an inline tool series alias, or an inline bypass. An inline bypass is a special construct that is used as a pseudo inline tool to allow a portion of traffic to bypass any inline tools.

For rule-based maps, the **Destination** can be configured to inline bypass with no restrictions, so long as the **Source** specifies either an inline network or an inline tool.

Use the **Priority** to order inline maps by priority. For example, you can specify the highest priority map to be for encrypted traffic and lowest priority map to be for a shared collector. You can also place inline maps before or after one another using the **Priority**.

Use the **Type** to specify any map rule.

Inline Map Shared Collector

Use the **Collector** subtype to configure a shared collector to which to send any packets that do not match the map rules in the inline maps. Use a shared collector with one or more rule-based inline maps. All of the inline network ports and inline tool ports involved in an inline shared collector map must be located on the same GigaVUE node.

Use the Source argument to specify an inline network alias or an inline network group alias.

Use the Destination argument to specify an inline tool alias, an inline tool group alias, an inline tool series alias, or an inline bypass. An inline bypass is a special construct that is used as a pseudo inline tool to allow a portion of traffic to bypass any inline tools.

For shared collector maps, the **Source** field can be configured to inline bypass with no restrictions, so long as the **Destination** field specifies either an inline network or an inline tool.

Support for shared collector inline maps is limited to symmetric scenarios, which means that the **Destination** field of rule-based inline maps can only be aliases of inline networks or inline network groups (not individual ports).

Inline Maps to Individual Members of an Inline Tool Group

Prior to software version 4.4, the **Map** page was used to configure an inline map that directed traffic to the inline tool group as a whole. The map could be rule-based or passall. There was only one type of hashing available, which distributed the traffic across the tools in the inline tool group.

Starting in software version 4.4, the **Map** page can also be used to configure inline maps to the individual members of an inline tool group. The maps must be rule-based. There are also more hashing options that can be specified for traffic that does not match any of the map rules. The hashing options are described in Symmetrical and Asymmetrical Hashing.

The rule-based maps are defined with the inline tool group sharing the same source, either an inline network or an inline network group in the **Source** field of the map. The map destinations (the **Destination** field of the map) are the individual inline tools in the group. Traffic not matching any of the map rules is sent to a shared collector to be distributed according to the specified hashing value.

The shared collector must also have the same source as the maps to the individual members of the inline tool group. The destination for the shared collector is the inline tool group. The shared collector map is a mandatory part of the configuration.

Both configurations are available: either a single map to the inline tool group as a whole, or multiple rule-based maps to the individual members of the inline tool group plus a shared collector; however, they cannot both be configured at the same time.

Refer to Figure 42Rule-Based Maps to Individual Tools in an Inline Tool Group for the rule-based maps to the individual members of the inline tool group. In Figure 42Rule-Based Maps to Individual Tools in an Inline Tool Group, traffic is only shown from A-to-B.

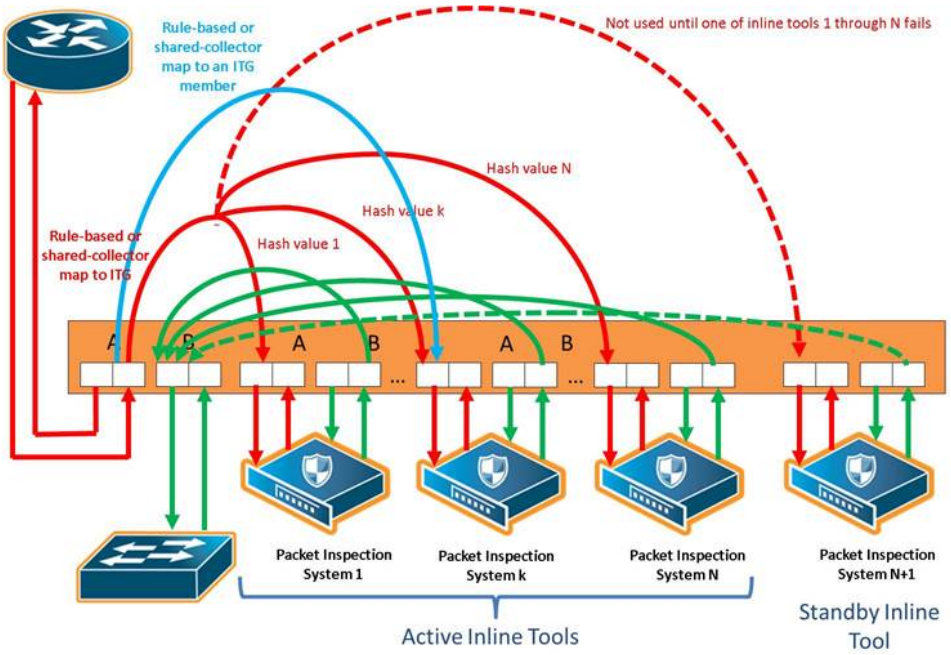


Figure 42 *Rule-Based Maps to Individual Tools in an Inline Tool Group*

Map, Inline Tool, and Inline Tool Group Configuration Restrictions

The following are map, inline tool, and inline tool group configuration restrictions for the rule-based maps to the individual members of the inline tool group:

- If there is a map to the inline tool group as a whole, there cannot also be rule-based maps to the individual inline tools in the group.
- Maps to the individual inline tools in a group must be rule-based. Map passalls to the individual tools cannot be configured.
- The source of rule-based maps to the individual inline tools in a group must be the same (either the same inline network or the same inline network group). The shared collector must have the same source as well.
- If there is a map configured to the individual inline tools in a group, the inline networks must have their traffic path set to **to-inline-tool**. This applies to individual inline networks as well as to inline networks involved in an inline network group.
- For the individual inline tools in a group, the recovery mode of the individual inline tools must be configured as **automatic**. A recovery mode of **manual** cannot be configured.
- For the inline tool group, the failover action must be either **NetworkBpass** or **NetworkDrop**. A failover action of either **ToolBypass** or **ToolDrop** for the inline tool group cannot be configured.
- Only one inline shared collector map can be configured (among the set of inline maps).
- Maps must be created in a specific order. The shared collector map must be configured last. For example, if there are three inline tools in the group, configure the three maps to the individual members of the group first, then configure the shared collector map.

- Maps must be deleted in a specific order. The shared collector map must be deleted first. Then the maps to the individual members of the group can be deleted.
- Once the shared collector map is configured, any changes to the maps to the individual members of the group are restricted. Only the map rules can be edited.

  NOTE: All the rules in a map cannot be deleted. All maps must have attributes for **Source**, **Destination**, and at least one **rule** configured.
- When an inline tool group is included as a member of an inline series, inline maps to individual members of an inline tool group are not supported.

- If one of the inline tools is disabled in an inline-network-group to inline-tool-group map that has the same ingress VLAN tag on the ports of an inline-network pair then the traffic will be looped back to the same network. This behavior is seen only with passall maps. In above scenarios use rule-based maps instead of passall maps.

Inline Tool Failures and Failover Actions

An inline tool group has a failover action for the group as a whole. The failover action is taken in response to a failure when the number of healthy inline tools in the inline tool group (including the spare inline tool, if configured) falls below the configured minimum. In addition, the individual inline tools in the group have failover actions.

When there are maps to individual inline tool members of the group, an inline tool has both a group failover action and an individual failover action.

Refer to Failover Actions When an Individual Tool in an Inline Tool Group Fails for the failover actions when an individual inline tool in an inline tool group fails.

Table 18: Failover Actions When an Individual Tool in an Inline Tool Group Fails

Inline Tool Group Spare	Number of Healthy Tools	When Individual Inline Tool Fails:	
		Traffic to Inline Tool Group	Traffic to Failed Tool
no spare	equal to or greater than the configured minimum healthy size	is redistributed among the remaining healthy tools in the group	no action—the failover action of the individual inline tool is ignored. If there is an inline map to the failed inline tool, the connectivity arrangements supporting the individual map are maintained as if nothing happened.
	less than the configured minimum healthy size	fails over according to the failover action for the inline tool group	no action—the failover action of the individual inline tool is ignored. If there is an inline map to the failed inline tool, the connectivity arrangements supporting the individual map are maintained as if nothing happened.

Inline Tool Group Spare	Number of Healthy Tools	When Individual Inline Tool Fails:	
		Traffic to Inline Tool Group	Traffic to Failed Tool
spare	equal to or greater than the configured minimum healthy size	is redistributed among the remaining healthy tools	no action—the failover action of the individual inline tool is ignored. If there is an inline map to the failed inline tool, the connectivity arrangements supporting the individual map are maintained as if nothing happened.
	less than the configured minimum healthy size	fails over according to the failover action for the inline tool group	no action—the failover action of the individual inline tool is ignored. If there is an inline map to the failed inline tool, the connectivity arrangements supporting the individual map are maintained as if nothing happened.
failed spare	N/A	no action, if there is no inline map configured to the spare	no action—the failover action of the individual inline tool (in this case, the spare) is ignored. If there is an inline map to the failed inline tool, the connectivity arrangements supporting the individual map are maintained as if nothing happened.

Maps That May Lead to Selective Traffic Drops

With inline bypass solutions based on inline flow mapping, the use of rule-based maps can lead to selective traffic drops. Traffic drops can occur as follows:

- if a shared collector from the inline network or inline network group has not been configured. Packets not matching the criteria specified by the rules in the configured rule-based maps will be dropped.
- if drop rules have been included in the rule-based maps configured from the inline network or inline network group

In most inline flow mapping solutions, all traffic exchanged between the two end nodes of a given inline network are expected to be processed by the inline tool or tools associated with this inline network through the configured maps. Therefore, it is recommended to always configure a shared collector and to not include drop rules in the rule-based maps.

Out-of-Band (OOB) Map

All inline network ports and inline tool ports can be subject to monitoring by listen-only tools. This means that an inline network port or inline tool port can be listed in the **Source** field in which the **Destination** field is an arbitrary tool type of port located anywhere in the system and not limited to the same node.

Inline network ports and inline tool ports involved in rule-based inline maps can be used as network ports for monitoring (or out-of-band) maps.

Out-of-band (OOB) maps are supported as follows:

- If the inline bypass solution use passalls, the OOB arrangements can use any rule-based maps or map shared collectors, or passalls. Refer to Figure 43Out-of-Band Rule-Based Maps.
- If the inline bypass solution use rule-based maps or map shared collectors, the OOB arrangements can use only map passalls. Refer to Figure 44Out-of-Band Map Passalls.
- When the source port of an OOB map is associated with an inline network, a list of inline ports is supported in the port list (the **Source** field of the Map page).

The following restrictions apply to OOB maps:

- When the source port of an OOB map is associated with an inline network group, only a single inline port (network or tool) is supported in the port list. In this case, multiple OOB maps are needed because each OOB map only accepts one inline port (network or tool) in the **Source** field on the Map page.
- OOB maps from inline network ports of inline networks involved in maps to inline tool groups configured with asymmetrical hashing are not allowed. If an inline network is involved in an inline map to an inline tool group configured with asymmetrical hashing, the inline network ports of the inline network cannot be used as the **Source** attribute in any out-of-band maps.

Prior to software version 4.4, if an inline network was part of an inline network group, sending traffic to an out-of-band tool was not allowed.

Starting in software version 4.4, out-of-band maps from inline ports involved in inline network groups are supported. You can configure OOB maps originating from inline network ports or inline tool ports when these ports are involved in an inline network group, except for the following:

- GigaSMART operations
- tool ports located on a different node

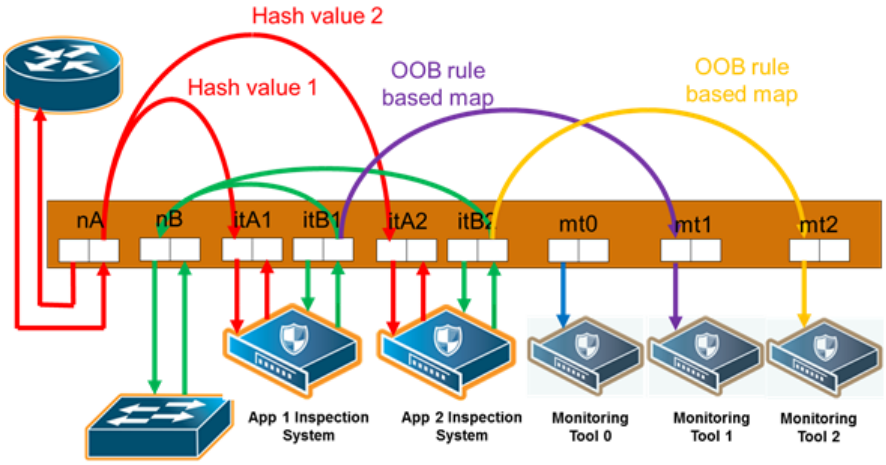


Figure 43 *Out-of-Band Rule-Based Maps*

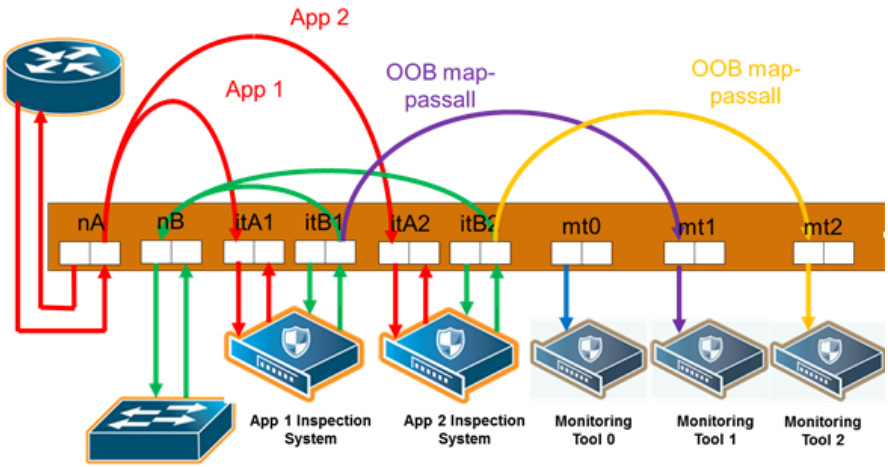


Figure 44 *Out-of-Band Map Passalls*

Symmetric and Asymmetric Maps

In a symmetric map configuration, the southbound and northbound forwarding to the tools is the same. For example, traffic from A to B goes through an inline tool, as does traffic from B to A. Rule-based maps are limited to symmetric configurations.

In an asymmetric map configuration, the southbound traffic is distributed to the inline tools, but northbound traffic can be sent through uninspected. For example, traffic from A to B goes through an inline tool, but traffic from B to A bypasses it.

Asymmetric configurations are only supported with map passalls. Traffic can come from individual inline networks, be sent to individual inline tools or inline tool groups, or to bypass.

Symmetrical combinations of two asymmetrical arrangements (that is, with both side A and side B pointing to the inline tool or to bypass) are not allowed.

Some IPSs do not need to inspect northbound traffic. For example, those focused on preventing Denial of Service (DoS) attacks, may not need to keep track of session flows and may only be concerned about southbound traffic.

Conversely, data loss prevention systems, those that are more concerned about what sensitive data is leaving the protected network than what is coming in, may focus solely on northbound traffic.

Configuration Steps

The configuration steps in summary for an inline bypass solution are as follows:

1. Configure inline network ports. (Optional for protected inline network.)

2. Configure inline network. (Optional for protected inline network.)
3. (Optional) Configure inline network group.
4. (Optional) Configure heartbeat or negative heartbeat profile.
5. Configure inline tool ports.
6. Configure inline tool.
7. (Optional) Configure inline tool group.
8. (Optional) Configure inline tool series.
9. Configure inline maps, either passall, map (rule-based), map shared collector, or bypass.
10. Configure non-default values for parameters of the inline networks or inline tools.

The summary steps are shown in Figure 45Configuration Steps for Inline Bypass Solutions.

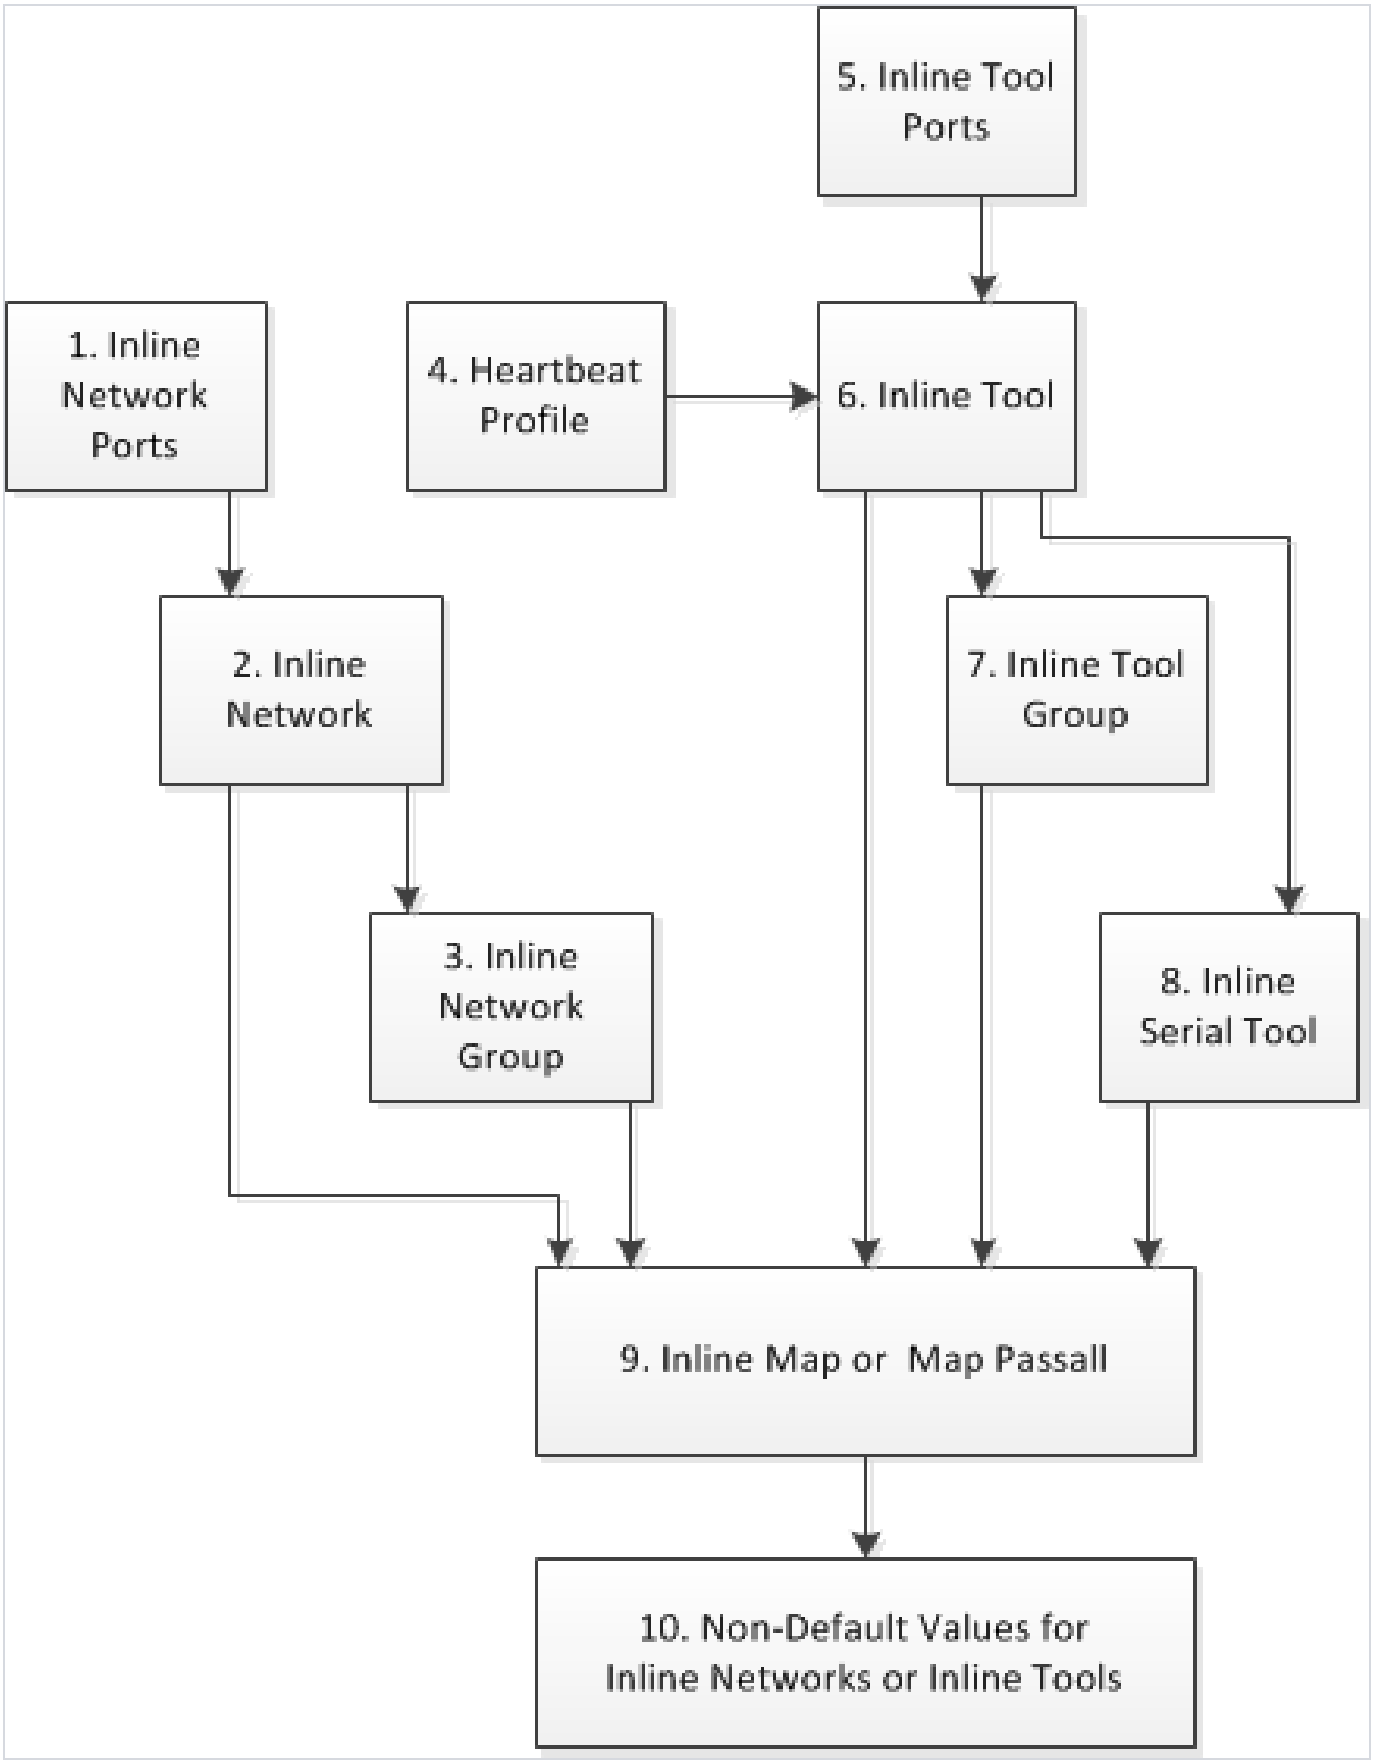

5. Inline Tool Ports
1. Inline Network Ports
4. Heartbeat Profile
6. Inline Tool
2. Inline Network
7. Inline Tool Group
3. Inline Network Group
8. Inline Serial Tool
9. Inline Map or Map Passall
10. Non-Default Values for Inline Networks or Inline Tools

Figure 45 *Configuration Steps for Inline Bypass Solutions*

Configuration Step Details

This section provides detailed steps for configuring inline bypass.

Configure Inline Network Ports

Use the following procedure to create inline network ports:

1. Go to **System > Ports > Ports > All Ports.**
2. Open Quick Port Editor by clicking the **Port Editor** button.
3. Use the Quick search field to Find the ports to configure.
4. In th an Alias field enter a name to help identify the inline port.
5. For Type, select **Inline Network**.

   An Inline Network (unprotected) is a software arrangement of two network-type ports allocated to facilitate access to a bidirectional link between two networks (far end network devices) that are linked to inline tool ports.
6. Click **OK**.

NOTE: Any available network-type ports on a Gigamon node can be used to form an unprotected inline network.

Configure Inline Network (Unprotected)

An Inline Network Group is an arrangement of multiple inline network ports to which traffic is distributed based on calculated hash values used by a Gigamon node.

Perform the steps to configure an inline network:

1. After configuring the Inline Network ports, select **Inline Bypass > Inline Networks**.
2. Click **New**.
3. In the **Alias** field, enter an alias for the inline network to help identify the inline network.
4. From the **Port A** drop-down list select an inline network port.
5. An inline network port is automatically selected for **Port B**. To select a different port, select one from the **Port B** drop-down list if there is more than one inline network port.
6. Select a traffic path from the **Traffic** Path drop-down list. The types of traffic paths are:

- **Bypass** — All traffic arriving at the Port A inline network port is directly forwarded to the Port B inline network port and all traffic arriving at the Port B inline network port is directly forwarded to the Port A inline network port.
- **Drop** — No traffic is exchanged through the inline network ports (all traffic coming to these ports is dropped).
- **ByPass with Monitoring** — All traffic is forwarded as a forced bypass value and a copy of the traffic is also forwarded to the inline tools. A traffic map must first be configured between the inline network and the inline tool to have the traffic forwarded with no traffic taken from the inline tools.
- **To Inline Tool** — The traffic received at the inline network ports is forwarded based on:

1. The traffic map between the inline network and the respective inline tools.
2. The failover action attributes of the inline tools
3. The health state of the inline tools.

7. Select the **Link Propagation** check box to enable whether the inline network link on one side of the inline network gets propagated to the other side.
8. Click **Save**.

Configure Inline Network Group

An Inline Network Group is an arrangement of multiple inline network ports to which traffic is distributed based on calculated hash values used by Gigamon node.

Perform the following steps to configure an Inline Network Group:

1. After configuring the Inline Network Ports, select **Inline Bypass > Inline Network Groups**.
2. Click **New**.
3. In the **Alias** field, enter a name for the network group.
4. From the **Inline Network** drop-down list, select the Inline Network ports.
5. Click **Save**.

Create Heartbeat Profile

The Heartbeat Profile is a data structure that contains the heartbeat attributes that are applied to an Inline Tool for configuring its heartbeat. The Create Heartbeat Profile wizard allows you to apply attribute values for the heartbeat profile. Use the following procedure:

1. Select **Inline Bypass > Heartbeats**.
2. Click **New**.
3. In the **Alias** field, enter a name for your heartbeat profile.

4. In the **Type** field, select **Regular**.

   For details about regular heartbeats, refer to Standard Heartbeat.

5. Use **Packet Format** drop-down list to select a packet type. The formats are:
   - **ARP**—This protocol (Address Resolution Protocol) is used for resolution of network layer address into link layer addresses, which is critical for multiple-access network operation. ARP is the default.
   - **Custom**—This format is a binary packet content associated with a packet capture (pcap) file. For details about custom packet format, refer to Standard or Custom Heartbeat Packet.

     When you select **Custom**, a **Custom Format** field displays with a **Browse** button. Use the **Browse** button to upload the pcap file.

6. Use the **Direction** drop-down menu to select the direction for sending heartbeat. The directions are:
   - **A to B**—From Port A to Port B of the inline tool.
   - **B to A**—From Port B to Port A of the inline tool.
   - **Bi-directional**—Both directions.

7. In the **Timeout** field, enter a number in milliseconds to indicate a timeout period for heartbeat packets between sending and receiving. The acceptable range is 20 to 1000 milliseconds. The default is 500 milliseconds.

8. In the **Period** field, enter a number in milliseconds for sending subsequent heartbeat packets. The acceptable range is 30 to 5000 milliseconds. The default is 1000 milliseconds.

9. In the **Recovery Time** field, enter a number in seconds to indicate that the inline tool is declared up with successfully receiving packets. The acceptable range is 5 to 60 seconds. The default is 30 seconds.

10. In the **Retries** field, enter the number for consecutive timed-out heartbeat packets at which the system will trigger (retry) a fail over condition.

11. Click **Save**.

    The heartbeat profile appears in the Heartbeat Profile table.

    NOTE: Highlight the heartbeat profile and click **Edit** to modify the parameters, if needed.

Create Negative Heartbeat Profile

The Negative Heartbeat Profile is a data structure that contains the negative heartbeat attributes that are applied to an Inline Tool for configuring its negative heartbeat. The Create Heartbeat Profile wizard allows you to apply attribute values for the negative heartbeat profile. Use the following procedure:

1. Select **Inline Bypass > Heartbeats**.
2. Click **New**.
3. In the **Alias** field, enter a name for your heartbeat profile.
4. In the **Type** field, select **Negative**

   This is a negative heartbeat profile. For details about negative heartbeats, refer to Negative Heartbeat Profiles.
5. Use **Browse** button in the **Custom Format** field to upload binary packet content associated with a packet capture (pcap) file. For details about the custom format, refer to Standard or Custom Heartbeat Packet.
6. Use the **Direction** drop-down menu to select the direction for sending heartbeat. The directions are:
   - **A to B**—From Port A to Port B of the inline tool.
   - **B to A**—From Port B to Port A of the inline tool.
   - **Bi-directional**—Both directions.
7. In the **Period** field, enter a number in milliseconds for sending subsequent negative heartbeat packets. The acceptable range is 30 to 5000 milliseconds. The default is 1000 milliseconds.
8. In the **Recovery Time** field, enter the minimum number of seconds since the last negative heartbeat packet is received to declare that the inline tool is up.

   The inline tool is up from the standpoint of the negative heartbeat if the negative heartbeats sent are *not* received. When a tool is declared down, sent heartbeats should not be received for a number of seconds in order to declare the tool as being up.

   The acceptable range for the Recovery Time field is 5 to 60 seconds. The default is 30 seconds.
9. Click **Save**.

   The heartbeat profile appears in the Heartbeat Profile table.

   NOTE: Highlight the heartbeat profile and click **Edit** to modify the parameters, if needed.

Configure Inline Tool Ports

Use the following procedure to create inline tool ports:

1. Go to **System > Ports > Ports > All Ports.**
2. Open Quick Port Editor by clicking the **Port Editor** button.
3. Use the Quick search field to Find the ports to configure.
4. In th an Alias field enter a name to help identify the inline port.

5. For **Type**, select **Inline Tool**.

   An Inline Tool represents a pair of inline tool ports.

6. Click **OK**.

Configure Inline Tool

An Inline Tool represents a pair of inline tool ports. To configure an Inline Tool, do the following:

1. Select **Inline Bypass > Inline Tools**.
2. Click **New**.
3. If needed, click **Port Editor** to open the Quick Port Editor to configure the inline tool ports.
4. Select the inline tool ports for the inline tool.
   - For **Port A**, select an inline tool port configured in the previous procedure, Configure Inline Tool Ports.
   - For **Port B**, select another inline tool port configured in the previous procedure, Configure Inline Tool Ports.
5. Select **Enabled** to set the inline tool ports as enabled for inline bypass traffic.
6. For **Failover action**, select one of the following:
   - **Tool Bypass** — When the inline tool fails all traffic coming to the respective inline tool is directed via the bypass path.
   - **Network Bypass** — When the inline tool fails the traffic is directed to multiple inline tools associated with an inline network or inline network group using rule-based inline maps.
   - **Tool Drop** — When the inline tool fails all traffic coming to the respective inline tool is dropped.
   - **Network Drop**—When the inline tool fails all traffic coming to the respective inline tool is dropped.
   - **Network Port Forced Down**—When the inline tool fails the inline network ports of the respective inline network are forced as "down".
7. Select the **Recovery Mode**. The recovery mode can be one of the following:
   - **automatic**—Specifies automatic recovery, which redirects traffic back to the inline tool as soon as it has recovered from all faulty conditions.
   - **manual**—Specifies manual recovery, which lets you control when to put an inline tool back into service after the tool has recovered. For example, you may want to wait for a maintenance window to return the inline tool to service.
8. Select **Enable Heartbeat** to set the heartbeat.

   If the heartbeat is enabled, do the following:

a. From **Profile**, select the desire heartbeat profile. The available heartbeat profiles are existing profiles from other inline tools. Once a heartbeat profile is selected, its attributes are displayed.

b. In the **HB IP Address A** field, enter the server's IP address to send the heartbeat packets.

c. In the **HB IP Address B** field, enter the server's IP address to send a second heartbeat packet.

9. Click **Save**.

Create Inline Tool Group

An Inline Tool Group is an arrangement of multiple Inline Tools to which traffic is distributed based on calculated hash values used by Gigamon node.

To configure an Inline Tool Group, do the following:

1. Use the **Quick Port Editor** to configure the inline tool ports.
2. Select **Inline Bypass > Inline Tool Groups**.
3. Click **New**.
4. In the **Alias** field, enter a name to help identify the inline tool group.
5. For **Inline Tools**, select the inline tool ports for the inline tool group.
6. (Optional) For the **Inline Spare Tool**, select another inline tool port.

   If a spare is selected, the inline tool group becomes a redundant arrangement of inline tools. When the first failure occurs in a set of active inline tools, traffic will be forwarded to the spare with no loss, thus the spare will replace the failed tool in the active set.
7. Select **Enable** to enable the inline tool group.
8. Select the **Failover Action**. The failover actions are:
   - **Tool Bypass**—When the inline tool group fails all traffic coming to the respective inline network is directed via the bypass path.
   - **Tool Drop**—When the inline tool group fails all traffic coming to the respective inline network is dropped.
   - **Network Bypass**—When the inline tool group fails all traffic coming to the respective inline network is directed to the inline tool group via the bypass path.
   - **Network Drop**—When the inline tool group fails all traffic coming to the respective inline network group is dropped.
   - **Network Port Forced Down**—When the inline tool group fails the inline network ports of the respective inline network are forced as "down".
9. For **Minimum Healthy Group Size**, select a number that represents the minimal amount of inline tools that are required to have a state of Normal.

10. Select the **Hash** for the inline tool group.

    Hashing, which is used for distributing traffic across the inline tools in an inline tool group. The values for the hash parameter are as follows:

    - **advanced**—Specifies symmetrical hashing, which is derived from the combination of packet fields based on the criteria selected for the advanced-hash algorithm. For inline bypass applications, the most common choice of criteria for the advanced-hash algorithm is the combination of source IP and destination IP addresses. This produces a hash value that sends all traffic associated with the same session to the same inline tool in the inline tool group.
    - **a-srcip-b-dstip**—Specifies asymmetrical hashing, which is derived from the source IP address for side A of the network and the destination IP address for side B of the network. This produces a hash value that sends all traffic associated with the same source address residing on side A to the same inline tool in the inline tool group, regardless of destination or session.
    - **b-srcip-a-dstip**—Specifies asymmetrical hashing, which is derived from the destination IP address for side A of the network and the source IP address for side B of the network. This produces a hash value that sends all traffic associated with the same source address residing on side B to the same inline tool in the inline tool group, regardless of destination or session.

11. Click **Save**.

Configure Inline Tool Series

To configure an Inline Tool Series, do the following:

1. Select **Inline Bypass > Inline Serial Tools**.
2. Click **New**.
3. In the **Alias** field, enter a name to help identify the inline tool series.
4. For **Inline Tools**, select the inline tool ports for the inline tool series.
5. Select **Enabled** to enable this configuration.
6. Select the **Failover action**.

   For details about the failover actions, refer to Inline Tool Series Local Failover Action
7. Select the **Per Direction Order**.

   For details about per-direction order, refer to Inline Tool Series Per-Direction Order.

Configure When GigaVUE® HC Series Modules are Operationally Up

Ensure that the GigaVUE® HC Series modules are in the operationally *up* state before configuring them. Configuration changes done when a module is operationally *down* are not supported.

Also, when an inline tool or inline tool group is in the operationally *down* state, do not modify the current failover action of that inline tool or inline tool group until the tool has recovered from the failover state.

Avoid Over subscription

In general, traffic received at inline network ports is delivered to the destination ports according to the inline maps and the out-of-band maps regardless of whether the destination ports have the capacity to absorb all the traffic or not.

> **NOTE:** When an inline network is involved in an inline map or an out-of-band map to a destination port (tool port or inline tool port), when there is temporary oversubscription, some packets arriving at the inline network port will be dropped. This can happen when the traffic path is set to bypass or monitoring.

Ensure that destination ports of maps originating from inline network ports have enough capacity to absorb the amount of traffic coming to the inline network ports.

Configure Gigamon Resiliency for Inline Protection

Gigamon Resiliency for Inline Protection (GRIP)™ is an Inline Bypass solution that connects two GigaVUE nodes together so that one node provides high availability to the other node when there is a loss of power. This redundant arrangement of two GigaVUE nodes maintains traffic monitoring by inline tools when one of the nodes is down.

GRIP makes use of the bypass protection switch relays for protected inline networks on GigaVUE-HC3, GigaVUE-HC1 nodes. The following modules are required to provide physical protection:

- bypass combo modules (BPS), for a protected pair of optical inline network ports on GigaVUE-HC3, GigaVUE-HC1.

- TAP-HC1-G10040 module, for a protected pair of copper inline network ports on GigaVUE-HC1

NOTE: GRIP is supported on GigaVUE-HC3 only when there are other modules installed in the node that can provide the stack link. The GRIP solution synchronizes the nodes through a signaling link using a stack link between two stack ports.

In the GRIP solution, two GigaVUE nodes are cabled so that traffic is guided through one GigaVUE node, acting in the primary role, while the other GigaVUE node is on standby, acting in a secondary role. If the primary node fails, the bypass protection switch relays on the modules switch the traffic over from the primary node to the secondary node. The failover or recovery process takes anywhere between 0-10 seconds depending on your configurations and devices. Contact Gigamon channel partner or Gigamon sales representatives for more information.

Using the physical protection for either copper or fiber, traffic is guided through inline tools by one of the GigaVUE nodes. The GigaVUE node with the open bypass protection switch relays is the one through which traffic flows. The traffic only flows through one GigaVUE node or the other.

To configure the GRIP solution for copper, use two TAP-HC1-G10040 modules on GigaVUE-HC1. The capacity will be 1Gb.

To configure the GRIP solution for fiber, use the following:

- two BPS-HC1-D25A24 or BPS-HC1-D25A60 or BPS-HC1-D35C60 modules on GigaVUE-HC1. The capacity will be 10Gb.
- two BPS-HC3-Q35C2G or BPS-HC3-C35C2G or BPS-HC3-C25F2G modules on GigaVUE-HC3. The capacity will be either 100Gb or 40Gb, depending on the configured port speed of the inline network port pairs

A stack port should be cabled between the two GigaVUE nodes. Also, two inline tools are needed for the GRIP solution. Refer to Configure Gigamon Resiliency for Inline Protection, Configure Gigamon Resiliency for Inline Protection, and Configure Gigamon Resiliency for Inline Protection.

How to Cable GigaVUE Nodes

To cable two GigaVUE nodes, as shown in Configure Gigamon Resiliency for Inline Protection with the primary on the left and the secondary on the right:

- Connect the network shown at the top of Configure Gigamon Resiliency for Inline Protection to inline network port A on the primary node.
- Connect inline network port B on the primary node to inline network port A on the secondary node.

- Connect inline network port B on the secondary node to the network shown at the bottom of the Configure Gigamon Resiliency for Inline Protection.
- Connect the signaling port on the primary node to the signaling port on the secondary node.

How to Handle Recovery

In the scenario in Configure Gigamon Resiliency for Inline Protection, after traffic is flowing through the secondary node, at some point the primary node will come back up. The primary node will establish the configured inline traffic paths, bring the signaling link up, and open its relays. Traffic will then flow through the primary node again.

Both Nodes Go Down and Only Secondary Comes Up

In the GRIP solution, if both primary and secondary nodes are powered down or if there is a power outage causing both primary and secondary nodes to go down, powering up the secondary alone without the primary ever coming up will cause network traffic to be bypassed instead of being sent to inline tools.

It is not recommended to power up/recover only the secondary node without the primary. The recommendation is to eventually bring the primary up also.

If the primary node is prone to failures or frequent power outages, another recommendation is to change the role of the secondary node to the primary.

How to Cable GigaVUE Nodes

To cable two GigaVUE nodes, as shown in Configure Gigamon Resiliency for Inline Protection with the primary on the left and the secondary on the right:

- Connect the network shown at the top of Configure Gigamon Resiliency for Inline Protection to inline network port A on the primary node.
- Connect inline network port B on the primary node to inline network port A on the secondary node.
- Connect inline network port B on the secondary node to the network shown at the bottom of the Configure Gigamon Resiliency for Inline Protection.
- Connect the signaling port on the primary node to the signaling port on the secondary node.

Configure GRIP Solution Software

To configure the GRIP solution in software:

1. On the left navigation pane, go to **Inventory** > **Nodes.**Click on the Left Node in which the configuration needs to be done.
2. Go to **System > Ports > >Ports>All** Ports. Select the port that would act as the signaling port and click **Edit**.
3. Select **Enable for Admin**.
4. Select Type **Stack**.
5. Click **OK**.
6. Repeat Steps 1 through 5 on the right node to complete the signaling port type configuration.
7. Go to > **Physical > Orchestrated Flows > Inline Flows > Configuration Canvas Inline Flows**, and then click **Configuration Canvas** to create a new Flexible Inline Canvas.
8. In the Flexible Inline Canvas that appears, select the required device and click the '+' icon next to the **Redundancy** option.
9. Enter a name for the profile in the Alias field.
10. Click on the Signaling Port field and select the stack port configured in Step 4. Select **Primary** for Protection Role.
11. Click **OK**.
12. Repeat Steps 8 through 11 for the right node to complete the Redundancy Profile Configuration but select Protection Role as **Secondary**.

The redundancy profile specifies the following:

- **Signaling Port**—specifies the ports used to signal the state of the two GigaVUE nodes to each other. The ports provide the mechanism to detect loss of power in one of the GigaVUE nodes.
- **Protection Role**—specifies the role of the GigaVUE node, as primary, secondary, or suspended. The default is suspended. When suspended, the protection role is on hold. Changing a GigaVUE node from the primary role to the suspended role can be used to manually force the primary node down so the secondary node can become active. The suspended role is also used when performing maintenance. Refer to Limitation for Suspended Role and How to Use Suspended Role for Maintenance.

The link between the signaling ports on the two GigaVUE HC Series nodes is for synchronization. When the node acting in the primary role is up, the signaling link is up, and the node acting in the secondary role sees the link as up. When the primary node loses power, the signaling link is brought down, and the secondary node sees the link as down and takes over.

The redundancy profile combines the protection role with the signaling port. The same redundancy profile is applied to the inline networks, so they have the same properties. If multiple inline networks on each GigaVUE node share the signaling link, they must be configured with the same protection role.

The primary and secondary roles on the two GigaVUE nodes do not change. That is, the role of the primary node stays the same and the role of the secondary node stays the same. The secondary always watches the state of the signaling port for whether the link is up or down.

For example, in Configure Gigamon Resiliency for Inline Protection, after the primary node recovers, it will open its bypass protection switch relays. Through the signaling port, the primary node will indicate that it is ready to receive traffic by setting the link state to up. The secondary node will notice that the link is up and will close its bypass protection switch relays. After recovery, the primary node automatically goes back into service.

Limitation for Suspended Role

Though GRIP is supported in a cluster, there is a limitation when the suspended protection role is used on the standby node in the cluster. The recommendation is to either switch the standby to the leader or apply the suspended role in the redundancy profile to the leader.

Configure Synchronization

You must synchronize the configuration of the two GigaVUE nodes involved in the GRIP solution. The configuration items that must be synchronized are as follows:

- the signaling ports, as dictated by the signaling link cabling
- the inline networks, as dictated by the network path cabling between the two GigaVUE nodes
- the redundancy profiles. The redundancy profile of each GigaVUE nodes needs to have the same signaling port as well as a redundancy role that is compatible with the redundancy role on the other GigaVUE node. For example, one is configured with the primary role and one is configured with the secondary role.
- the inline tools
- the inline maps

For a configuration example of two GigaVUE-HC3 nodes, refer to Example: Gigamon Resiliency for Inline Protection. In the example, the configuration is the same on both nodes, except for the protection role (primary or secondary).

Display Redundancy Control State

To display the Redundancy Control State, go to the Inline Networks page and click on the alias of the Inline Network for which you want to display the redundancy control state. The state is displayed on the Quick View under Configuration.

Table 19: Redundancy Control States

State	Description
Neutral	No redundancy profile is configured.
Suspended	The protection role is configured as suspended.
Primary Forwarding	The protection role is configured as primary. The node is acting in the primary role. Traffic flows through this node.
Secondary Bypass	The protection role is configured as secondary. The node is acting in the secondary role. Traffic bypasses this node.
Secondary Forwarding	The protection role is configured as secondary. The node is acting in the primary role due to a loss of power on the primary node. Traffic flows through this node.

How to Use Suspended Role for Maintenance

Use the suspended protection role to perform maintenance activities on the primary and secondary nodes. Maintenance activities include: bringing up a module, shutting down a module, or swapping a module.

For example, to remove a module on one of the GigaVUE-HC3 nodes (Primary node), use the following steps on that module:

1. Select **Physical > Orchestrated Flows > Inline Flows > Configuration Canvas > Redundancy**.
2. On the Redundancies page, for **Protection Role**, select **suspended**, and then click**OK**.
3. Once this is configured, the Primary node will be moved to 'Suspended' and the Secondary node will be moved to 'Secondary Forwarding' state. All the traffic will now be forwarded to the Secondary node and the Inline Tool inspection takes place.
4. Perform the maintenance activity in Primary node, like bringing up a new module, shutting down a module, swapping the modules, replacing the external inline tool.
5. Once the maintenance is done, revert the **Protection Role** in Primary Node back to **'primary'**. This will move the Redundancy Control State back to the **Primary Forwarding** and traffic will start flowing via the Primary Node.

In case of a maintenance activity (chassis, card, external inline-tool) required in Secondary node follow the below steps:

1. Set the Protection Role to '**suspended**' in the redundancy profile.
2. Once this is set, the Secondary node will be moved to **'Suspended'** and the Primary node will remain in the **'Primary Forwarding'** state and will handle the traffic.
3. Do the required maintenance activity in secondary node.
4. Once the maintenance is done, revert the Protection Role in Secondary Node back to **'secondary'**.
5. This will move the Redundancy Control State back to the "**Secondary Bypass**".

Upgrade Procedure Recommendations

There are no specific procedures for upgrading the nodes in the GRIP deployment, but we would recommend the below steps to be on the safer side. Even when the upgrading device goes into an issue state, traffic will be inspected in any one of the nodes.

- Refer the *"Supported Upgrade Path—Standalone Nodes"* to know about the order of build version in which the GigaVUE-OS needs to be upgraded.
- Save the configuration and take the backup of both Primary and Secondary node.
- First upgrade the secondary node so that traffic will get inspected in the Primary node.
- Once the Secondary node upgrade gets completed then go for the Primary node upgrade.
- When the Primary node upgrade is in progress, traffic will be inspected in the Secondary node as the failover kicks in, with the help of the signaling port.
- As soon as the Primary node upgrade is complete, the signaling port will come up, and the Primary Node will start inspecting the traffic.

Rules and Notes

Keep in mind the following rules and notes when you work with the Gigamon Resiliency for Inline Protection feature:

1. The signaling port type should be a stack port, and only one port should be used.
2. All the inline components should be located in the same box within the cluster.
3. Adding the Inline Networks in the Inline Network Bundle and deploying the Solutions is recommended, which will be easy to export and import across the GRIP Nodes.
4. It is recommended to create a separate redundancy profile for each solution's sources (Inline Network) deployed from GigaVUE-FM Canvas.
5. Before deleting any Solution/Map, it is recommended to move the respective redundancy profile protection role to suspended and revert back to the respective protection role once the solution is redeployed.

6. Post-reload or Power Cycle, the signaling port link stays down when the redundancy profile is attached to the inline network and no maps have been configured. A map should be configured to bring the signaling port up. If a map exists, the signaling port will appear without any issues.
7. Link Failure Propagation is not recommended for copper ports (TAP card ports). Fabric ports support LFP only in a single path (a-to-b only) is available. In all other cases, it is best to leave LFP enabled.
8. Gigamon Resiliency for Inline Protection (GRIP™) is not supported in GigaVUE-HCT devices.
9. GRIP is not supported in other GigaVUE TA Series devices due to the absence of BPS modules.
10. Refer to Flexible Inline Arrangements—Rules and Notes, which also apply to GRIP.

Limitations

- Link failure propagation is not recommended when inline network ports involve copper ports (TAP card ports) or fabric ports with only a single available path (a-to-b only). In all other cases, enabling LFP is recommended.

Troubleshoot

If any of the below issues occur, kindly follow the given steps to troubleshoot the issue.

Signaling Ports Down

1. Check if the power is proper from the optics; if not, try replacing the optics with new ones in a maintenance window

2. If the issue occurs after reloading, check if the maps are configured on the inline network to which the redundancy profile is attached.

3. If maps are not configured, deploying a map will bring up the signaling ports.

4. If a map is available and the signaling ports are still down, contact Gigamon Support for assistance.

Traffic outage in Inline Tool

1. Check if the Redundancy Control State is set to Primary Forwarding in the Primary node or Secondary Forwarding in case traffic is handled in the Secondary Node.

2. Check if the Inline Tool Flex Traffic Path is configured as “Bypass” by mistake. If so, revert to To-Inline-Tool to recover the traffic.

3. Check if the Inline Tool ports are in downstate and failover kicked into tool bypass (the default failover for Inline Tool). If so, correct the optics power, cable, and external Inline Tool faultiness.

4. If the issue persists, contact Gigamon Support for assistance.

Network Traffic Outage

1. Check if any of the Inline Tool’s Flex Traffic Paths is set to “Drop” by mistake. If so, revert the Flex Traffic Path to To-Inline-Tool.

2. Check if any Inline Tool ports are down and failover kicked into tool-drop or network-drop. If so, check the optics power, cable, and External Inline Tool Faultiness and correct the same.

3. Check if the traffic hits the map and drops in any Inline Component ports using 'show port stats port-list <port-alias>.'

4. Try redeploying the Flex Inline Solution from GigaVUE-FM and check if the traffic resumes. If not, contact Gigamon Support for assistance.

FAQs

This section answers frequently asked questions when configuring the Gigamon Resiliency for Inline Protection feature.

1. **How many members can we have in the GRIP configuration?**

   It is must have two nodes to do the GRIP configuration where one node acts as the Primary and the other acts as the Secondary.

2. **What happens if the primary node becomes unresponsive?**

   Even if the Primary node becomes unresponsive, traffic should not be impacted as the unresponsiveness is from Software, whereas traffic will be forwarded to the Hardware level. Also, unresponsiveness in the Primary Node will not trigger the failover switch to the Secondary.

3. **Is GRIP Supported in Flexible Inline Maps?**

   Yes, GRIP is supported in both Classic and Flexible Inline Maps.

4. **What happens if we delete a map on the primary?**

   When we delete a map on the primary node, irrespective of the inline-network traffic path, the traffic is switched to the secondary node.

5. **Will the configuration between the Primary and Secondary nodes get synchronized automatically?**

   .The user needs to ensure that all the configuration changes made in the Primary node are also made in the Secondary node. Refer to Configuration Synchronization for more details.

Migrate Inline Bypass(Classic) Solution from version 6.5 to 6.12

This section explains how to migrate Inline Bypass (Classic) solutions from version 6.5.xx or earlier to version 6.12.xx.

Migration Scenarios

Based on your device model, you may need to migrate your solution or upgrade your device. Use the following table to determine the required action.

Device Model	Current Version	Migrate or Upgrade	Notes
GigaVUE-HC2	6.5.xx or earlier	Migrate the configurations to GigaVUE-HC1-Plus running on version 6.12	GigaVUE-HC2 is end of support.
GigaVUE-HC1-Plus	6.5.xx or earlier	Migrate the configurations to GigaVUE-HC1-Plus running on version 6.12	Support for Classic inline bypass solution is reintroduced in version 6.12.xx.
GigaVUE-HC1 or GigaVUE-HC3	6.5.xx or earlier	Upgrade the device	Supported upgrade path: 6.5 > 6.9 > 6.12 Refer to Supported Upgrade Path—Standalone Nodes

Migration Considerations

When migrating Inline Bypass configurations from GigaVUE-HC2 to GigaVUE-HC1-Plus, review the port differences carefully:

- Inline network that uses protected ports for Inline bypass configuration on GigaVUE-HC2 must be mapped to the available protected ports on GigaVUE-HC1-Plus.
- Inline tools that used unprotected ports on GigaVUE-HC2 must be mapped to available unprotected ports on GigaVUE-HC1-Plus.
- If breakout ports are in use, configure the same breakout mode on GigaVUE-HC1-Plus.

- The following table provides some examples of port mapping between GigaVUE-HC2 and GigaVUE-HC1-Plus.

Mapping Examples	GigaVUE-HC2 Port (Type/Role)	GigaVUE-HC1-Plus Port (Role/Type)
Inline Network (Protected) Ports pair	1/1/c1 - 1/1/c2 (protected, from BPS card)	1/2/x1 ↔ 1/2/x2 (protected ports)
Inline Tool (Unprotected) Ports	1/2/x1 (unprotected, from BPS card)	1/1/x1 (onboard fixed port)

Migrate Classic Inline Bypass Solution using GigaVUE-OS CLI

Follow these steps to migrate configurations from GigaVUE-HC2 or GigaVUE-HC1-Plus running 6.5 or earlier to GigaVUE-HC1-Plus running version 6.12:

1. Prepare the Source Configuration
   a. Generate the running configuration.
   b. Save the configuration as a text file and upload it to a remote server using the following command:
   configuration text generate active running save <file name> upload <remote server path>
2. Edit the Configuration File

   1. Download the configuration file form the remote server and update it to match the GigaVUE-HC1-Plus device. Modify the following parameters:
      - IP address – If static, update with the new GigaVUE-HC1-Plus IP. DHCP addresses will not appear in the file.
      - Hostname – Optional.
      - Logging configurations – No changes required.
      - Local user accounts – Optional. Update username and password as needed.
      - Chassis and card configurations – Update to match the GigaVUE-HC1-Plus device.
      - Port configurations – Update to align with GigaVUE-HC1-Plus ports. See Migration Considerations
      - Inline networks, Inline tools, and map configurations – Update to reflect port configuration changes.
3. Ensure your GigaVUE-HC1-Plus is running on version 6.12.xx .
4. Apply the configuration to GigaVUE-HC1-Plus running on version 6.12.xx:
   a. Upload the edited configuration text file to your remote server.
   b. On the GigaVUE-HC1-Plus device, fetch the configuration file using the following command:
   config text fetch <path-to-file-in-remote-server>

c. Apply the configuration using the following command:
config text file <filename> apply fail-continue

NOTE: If mismatches occur, the system skips invalid entries and displays details of the failed configurations.

5. Verify your configurations
Run the following **show** commands to confirm that the configurations were applied successfully.

   a. **show inline-network**

   b. **show inline-tool**

   c. **show map**

Migrate Classic Inline Bypass Solution using GigaVUE-FM

Follow these steps to migrate configurations from GigaVUE-HC2 or GigaVUE-HC1-Plus running version 6.5.xx or earlier to GigaVUE-HC1-Plus running version 6.12.xx:

1. Download running configuration of your GigaVUE-HC2 or GigaVUE-HC1-Plus in version 6.5.xx to the remote server. Follow the below steps:

   a. On the left navigation pane, click  , and then select **Physical>Nodes**. Select your GigaVUE-HC2 or GigaVUE-HC1-Plus.

   b. Go to **Settings > Config and Licenses > Configurations**.The Configurations page appears.

   c. Select **Actions > Save as Text** to save the currently running configuration as text file.

   d. Select the running configuration file and click **Actions > Download**.

   e. Upload the downloaded configuration text file to your remote server.

2. Edit the Configuration File

   1. Download the configuration file from the remote server and update it to match the GigaVUE-HC1-Plus device. Modify the following parameters:
      - IP address – If static, update with the new GigaVUE-HC1-Plus IP. DHCP addresses will not appear in the file.
      - Hostname – Optional.
      - Logging configurations – No changes required.
      - Local user accounts – Optional. Update username and password as needed.
      - Chassis and card configurations – Update to match the GigaVUE-HC1-Plus hardware.

- Port configurations – Update to align with GigaVUE-HC1-Plus ports. See Migration Considerations
- Inline networks, inline tools, and map configurations – Update to reflect port configuration changes.

3. Ensure your GigaVUE-HC1-Plus is running on version 6.12.xx.
4. Apply the configuration to GigaVUE-HC1-Plus running on version 6.12 as follows:
   a. On the left navigation pane, click , and then select **Physical>Nodes**.
   b. Select the GigaVUE-HC1-Plus where you want to migrate the Inline Bypass configuration.
   c. Select **Settings > Config and Licenses > Configurations**.The Configurations page appears.
   d. Click **Actions > Upload**. Upload the edited configuration files from your local drive.

      If you have any other map configurations in your device clear them before you apply your configuration text file.
   e. Select the configuration text file and click **Actions >Apply**. In the Apply Configurations dialog box select *Fail Continue* and then click **Apply** to apply the configurations.

NOTE: If mismatches occur, the system skips invalid entries and displays details of the failed configurations.

5. Verify the Configurations
   Confirm that the configurations were applied successfully in the following pages:
   a. Inline Networks: On the left navigation pane, click , and then select **Physical>Nodes>System > Inline Bypass> Inline Networks**.
   b. Inline Tools: On the left navigation pane, click , and then select **Physical>Nodes>System > Inline Bypass> Inline Tools**.
   c. Maps: On the left navigation pane, click , and then select **Physical>Nodes>Traffic > Maps**.

Inline Bypass Solution Examples

The following sections provide examples of inline bypass solutions. The solutions are presented in an order from simple to complex. Refer to the following:

- Example 1: Unprotected Inline Bypass with an Inline Tool Group
- Example 2: Unprotected Inline Bypass with Default Heartbeat

- Example 3: Protected Inline Bypass Using Combo Modules
- Example 4: Gigamon Resiliency for Inline Protection

Example 1: Unprotected Inline Bypass with an Inline Tool Group

Example 1 is a simple, unprotected inline bypass solution. In the example, aliases are used for inline network ports (iN1 and iN2), inline tool ports (iT1 and iT2), inline network (inNet), inline tool (inTool), and inline map (inMap).

On GigaVUE-HC3, an unprotected inline bypass solution can be configured on the bypass combo module with the inline networks and inline tools on ports 1/1/x1..x16 or on ports c1..c4, or on any other module on the GigaVUE-HC3 node.

On GigaVUE-HC1-Plus, an unprotected inline bypass solution can be configured with the inline networks and inline tools on ports 1/1/x1..x16 or on ports x17..x24, or on any other module on the GigaVUE-HC1-Plus node. Refer to Figure 46Logical Bypass Without Bypass Combo Module which shows a GigaVUE-HC1-Plus.



Figure 46 *Logical Bypass Without Bypass Combo Module*

On GigaVUE-HC1-Plus, an unprotected inline bypass solution can be configured on the base module, with the inline networks and inline tools on ports 1/1/x1..x12 and 1/1/g1..g4, or on the bypass combo module on ports x1..x4.

Task	Description	UI Steps
1.	Configure inline network aliases, port type (inline-network), and administratively enable inline network ports.	1. Go to **System > Ports > Ports > All Ports.** 2. Click Quick Port Editor. 3. Use Quick search to find the ports to configure. In

Task	Description	UI Steps
		this example, the ports are 3/1/x1 and 3/1/x2 **4.** Set port 3/1/x1 to Type Inline-Network and select Enable. Enter iN1 for the port alias. **5.** Set port 3/1/x2 to Type Inline-Network and select Enable. Enter iN2 for the port alias. **6.** Make sure Enable is selected for Admin on the ports. **7.** Click OK.
2.	Configure inline network.	**1.** Select Inline Bypass > Inline Networks. **2.** Click New. **3.** In the Alias field, type InNet. **4.** For Port A, select iN1. **5.** For Port B, select iN2. **6.** Click Save.
3.	Configure inline tool ports, port type (inline-tool), and administratively enable inline tool ports.	**1.** Go to **System > Ports > Ports > All Ports**. **2.** Click Quick Port Editor. **3.** Use Quick search to find the ports to configure. In this example, the ports are 3/1/x3 and 3/1/x4 **4.** Set port 3/1/x4 to Type Inline-Tool and select Enable. Enter iT1 for the port alias. **5.** Set port 3/1/x4 to Type Inline-Tool and select Enable. Enter iT2 for the port alias. **6.** Make sure Enable is selected for Admin on the ports. **7.** Click OK.
4.	Configure inline tool, and enable it. Also enable the default heartbeat profile.	**1.** Select Inline Bypass > Inline Tools. **2.** Click New. **3.** In the Alias field, type InTool. **4.** For Port A, select iT1. **5.** For Port B, select Ii2. **6.** Under Configuration: **a.** Select Enabled **b.** Select Enabled Heartbeat and set Profile to default. **7.** Click Save.
5.	Configure map passall, from inline network to inline tool.	**1.** Select Maps > Maps. **2.** Click New. **3.** In the Alias field, type InMap. **4.** Select Inline for Type and Pass All for Subtype. **5.** For Source, select InNet

Task	Description	UI Steps
		6. For Destination, select InTool. **7.** Click Save.
6.	Configure the path of the traffic to inline tool.	**1.** Select Inline ByPass > Inline Networks. **2.** Select the Inline Network InNet and click Edit **3.** Under Configuration, set the Traffic Path field to To Inline Tool. **4.** Click Save.

Example 2: Unprotected Inline Bypass with Default Heartbeat

Example 2 adds the default heartbeat profile to the unprotected inline bypass solution in Example 1.

Task	Description	UI Steps
1.	Configure inline network aliases, port type (inline-network), and administratively enable inline network ports.	**1.** Go to **System > Ports > Ports > All Ports**. **2.** Click Quick Port Editor. **3.** Use Quick search to find the ports to configure. In this example, the ports are 3/1/x1 and 3/1/x2 **4.** Set port 3/1/x1 to Type Inline-Network and select Enable. Enter iN1 for the port alias. **5.** Set port 3/1/x2 to Type Inline-Network and select Enable. Enter iN2 for the port alias. **6.** Make sure Enable is selected for Admin on the ports. **7.** Click OK.
2.	Configure inline network.	**1.** Select Inline Bypass > Inline Networks. **8.** Click New. **9.** In the Alias field, type InNet. **10.** For Port A, select iN1. **11.** For Port B, select iN2. **12.** Click Save.
3.	Configure inline tool ports, port type (inline-tool), and administratively enable inline tool ports.	**1.** Go to **System > Ports > Ports > All Ports**. **13.** Click Quick Port Editor. **14.** Use Quick search to find the ports to configure. In this example, the ports are 3/1/x3 and 3/1/x4 **15.** Set port 3/1/x4 to Type Inline-Tool and select Enable.

Task	Description	UI Steps
		Enter iT1 for the port alias. **16.** Set port 3/1/x4 to Type Inline-Tool and select Enable. Enter iT2 for the port alias. **17.** Make sure Enable is selected for Admin on the ports. **18.** Click OK.
4.	Configure default heartbeat profile.	In GigaVUE-HVUE, the default heartbeat profile is already configured. To view the profile, select Inline Bypass > Heartbeats.
5.	Configure inline tools, and enable it. Also enable the default heartbeat profile.	**1.** Select Inline Bypass > Inline Tools. **19.** Click New. **20.** In the Alias field, type InTool1. **21.** For Port A, select iT1. **22.** For Port B, select iT2. **23.** Under Configuration: **a.** Select Enabled **b.** Select Enabled Heartbeat and set Profile to default. **24.** Click Save.
6.	Configure map passall, from inline network to inline tool.	**1.** Select Maps > Maps. **25.** Click New. **26.** In the Alias field, type InMap. **27.** Select Inline for Type and Pass All for Subtype. **28.** For Source, select InNet **29.** For Destination, select InTool. **30.** Click Save.
7.	Configure the path of the traffic to inline tool.	**1.** Select Inline ByPass > Inline Networks. **31.** Select the Inline Network InNet and click Edit **32.** Under Configuration, set the Traffic Path field to To Inline Tool. **33.** Click Save.

Example 3: Protected Inline Bypass Using Combo Modules

Example 3 is a protected inline bypass solution using bypass combo modules on GigaVUE-HC3. It also configures heartbeat and negative heartbeat profiles.

Protected inline networks are based on the pairs of ports associated with the physical protection switches located on the bypass combo modules. Unlike the unprotected examples, you do not need to configure inline network ports because they are created automatically. On GigaVUE-HC3, the port pairs are numbered for example: 2/2/x17 and 2/2/x18, 2/2/x19and 2/2/x20, 2/2/x21 and 2/2/x22, 2/2/x23 and 2/2/x24.

You do not need to configure inline networks because they are also created automatically on bypass combo modules. The aliases of the default inline networks are: default_inline_net_2_2_1, default_inline_net_2_2_2, default_inline_net_2_2_3, default_inline_net_2_2_4.

On GigaVUE-HC3, protected inline bypass can be configured on the bypass combo module on ports c1..c4.

On GigaVUE-HC1, protected inline bypass can be configured on the bypass combo module. It can also be configured on the TAP-HC1-G10040 module placed in either bay 2 or bay 3, so the ports will be 1/2/g1..g8 or 1/3/g1..g8.

NOTE: The default value of the physical-bypass attribute of protected inline networks is set to enable, which means that the fibers attached to ports net-a and net-b of the inline network are optically coupled and the traffic is exchanged between end nodes without coming to the switching fabric. of the GigaVUE node. As shown in Example 4, after configuring the inline tool and the map passall, the physical-bypass attribute is set to disable in order to activate the Inline Bypass solution.

Task	Description	UI Steps
1.	Configure inline too aliases, port type (inline-tool), and administratively enable inline network ports.	**1.** Go to **System > Ports > Ports > All Ports**. **1.** Click Quick Port Editor. **2.** Use Quick search to find the ports to configure. In this example, the ports are 2/2/x11 and 2/2/x12 **3.** Set port 2/2/x11 to Type Inline Tool and select Enable. Enter iT1 for the port alias. **4.** Set port 2/2/x12 to Type Inline Tool and select Enable. Enter iT2 for the port alias. **5.** Click OK.
2.	Configure a heartbeat profile.	**1.** Select Inline Bypass > Heartbeats. **6.** Click New. **7.** In the Alias field, type hb2. **8.** For Type, select Regular. **9.** Click Save.
3.	Configure negative heartbeat profile alias and PCAP file	**1.** Select Inline Bypass > Heartbeats. **10.** Click New.

Task	Description	UI Steps
		11. In the Alias field, type nhb1. **12.** For Type, select Negative. **13.** Click the Browse button for Custom Format and upload the pcap file; for example, hnb.pcap. **14.** Click Save.
4.	Configure inline tools, and enable them. Also specify the heartbeat profile.	**1.** Select Inline Bypass > Inline Tools. **15.** Click New. **16.** In the Alias field, type InTool1. **17.** For Port A, select iT1. **18.** For Port B, select iT2. **19.** Under Configuration: **a.** Select Enabled **b.** Select Enabled Regular Heartbeat and set Profile to hb2. **c.** Select Enabled Negative Heartbeat and set Negative Heartbeat Profile to nhb1. **20.** Click Save. Configure the second inline tool. **1.** Select Inline Bypass > Inline Tools. **21.** Click New. **22.** In the Alias field, type InTool2. **23.** For Port A, select iT3. **24.** For Port B, select iT4. **25.** Under Configuration: **a.** Select Enabled **b.** Select Enabled Heartbeat and set Profile to hb_custom. **26.** Click Save.
5.	Configure the inline tool group and enable it.	**1.** Select Inline Bypass > Inline Tool Group **27.** Click New. **28.** In the Alias field, type inToolGroup. **29.** For Inline Tools, select InTool1 and InTool2. **30.** Under Configuration, select Enabled. **31.** Click Save.
6.	Configure map passall, from inline network to inline tool.	**1.** Select Maps > Maps. **32.** Click New. **33.** In the Alias field, type InMap. **34.** Select Inline for Type and Pass All for Subtype.

Task	Description	UI Steps
		35. For Source, select default_inline_net_2_2_1 **36.** For Destination, select InTool1 **37.** Click Save.
7.	Configure the path of the traffic to inline tool and disable the physical bypass	**1.** Select Inline ByPass > Inline Networks. **38.** Select the Inline Network default_inline_net_2_2_1 and click Edit **39.** Under Configuration, set the Traffic Path field to To Inline Tool. **40.** Make sure Physical Bypass is not selected. **41.** Click Save.

Example 4: Gigamon Resiliency for Inline Protection

Example 5 is an inline bypass solution for GRIP using TAP-HC0-G100C0 modules on the GigaVUE-HC2, with copper ports.

First, configure the GigaVUE-HC2 with the primary role, then configure the GigaVUE-HC2 with the secondary role. The configuration is the same (is synchronized) on both nodes, except for step 3, in which the protection role (primary or secondary) is specified.

Note that in this example, link fail propagation (LFP) is disabled to reduce inline network recovery time after failover. When a primary to secondary failover occurs and LFP is enabled for copper inline bypass links, network service recovery may take several seconds because of Ethernet link renegotiation. Optical links failover faster and typically recover service much faster. For inline networks where only one path is available, this is a consideration. When GRIP is deployed with high availability networks where a second path is present, it is a best practice to leave LFP enabled.

You can use the Chassis page to view the chassis and modules.

Configure Primary Role GigaVUE-HC2

Task	Description	UI Steps
1.	Configure ports on the TAP-HC0-G100C0 module as passive (in	**1.** Go to **System > Ports > Ports > All Ports**. **1.** Select a port of the TAP-HC0-G100C0 module.

Task	Description	UI Steps
	passive mode, relays are closed). Also configure ports, port type (inline-network).	**2.** Click to open the Ports page. **3.** Select Passive for TapTX **4.** Click Save. **5.** Repeat steps 2 through 6 for each port on the TAP-HC0-G100C0 module **6.** Configure Inline Network ports **a.** Select the port. **b.** Click Quick Port Editor. **g.** Select Inline Network for Type. **Note**: You can use the Chassis page to locate the position of the module in the chassis and identify port IDs.
2.	Configure stack port (for signaling port/link) and enable it.	**1.** Select the port and click Edit. **8.** Select Enable for Admin. **9.** Select Stack for Type. **10.** Click OK.
3.	Create the redundancy profile by giving it a name and configuring parameters for the redundancy profile such as the signaling port and protection role (primary).	**1.** Select Inline Bypass > Redundancies **11.** Click New. **12.** Enter a name for the profile in the Alias field. For example, RP_001. **13.** Click in the Signaling Port field and select the stack port configured in Task 2. **14.** Select Primary for Protection Role. **15.** Click Save.
4.	Configure inline network. This step associates the redundancy profile to the inline network and also disables link fail propagation on the inline network.	**1.** Select Inline Bypass > Inline Tools. **16.** Click New. **17.** Enter a name for the inline network in the Alias field. For example, IN_001. **18.** Click in the Port A field and select inline network port for network A. **19.** Click in the Port B field and select an inline network port f. **20.** Select Bypass for Traffic Path. (This is the default setting.) **21.** Make sure Link Failure Propagation option is NOT checked. (It is enabled by default.) **22.** Click in the Redundancy Profile field and select the profile created in step 3. For example RP_001.
5.	Configure inline tool ports, port type (inline-tool), and administratively enable them.	**1.** Go to **System > Ports > Ports > All Ports**. **23.** Select the first port (for example, 1/4/x1) to configure

Task	Description	UI Steps
		as an inline-tool port. **24.** Click Quick Port Editor. **25.** Select Inline Network for Type and select Enable for Admin. **26.** Click OK. **27.** Select the second port (for example, 1/4/x2) and repeat steps 3 through 5.
6.	Configure inline tool and failover action. then enable inline tool.	**1.** Select Inline Bypass > Inline Tools **28.** Enter a name for the inline tool in the Alias field. For example, IT_001 **29.** Click in the Port A field and select the first inline tool port configured in Task 5. **30.** Click in the Port B field and select the second inline tool port configured in Task 5. **31.** Make sure Enable is selected. (It is enabled by default **32.** Select NetworkBypass for Failtover action. **33.** Click Save.
7.	Configure map passall, from inline network to inline tool. **Note**: When you delete a map on the primary node, irrespective of the inline-network traffic-path, the traffic is switched to the secondary node. The port utilization must be 0% on the primary node and active on the secondary node.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Enter a name for the map in the Alias field. For example, INtoIT. **4.** Select Regular for Type. **5.** Select Pass All for Subtype. **6.** Select the inline network created in Task 4 for Source. For example, IN_001. **7.** Select the inline tool created in Task 6. For example, IT_001.

Configure Secondary Role GigaVUE-HC2

Task	Description	UI Steps
1.	Configure ports on the TAP-HC0-G100C0 module as passive (in passive mode, relays are closed). Also configure ports, port type (inline-network).	**1.** Go to **System > Ports > Ports> All Ports**. **8.** Select a port of the TAP-HC0-G100C0 module. **9.** Click to open the Ports page. **10.** Select Passive for TapTX **11.** Click Save. **12.** Repeat steps 2 through 6 for each port on the TAP-HC0-G100C0 module **13.** Configure Inline Network ports

Task	Description	UI Steps
		a. Select the port. **b.** Click Quick Port Editor. **n.** Select Inline Network for Type. **Note**: You can use the Chassis page to locate the position of the module of the module in the chassis and identify port IDs
2.	Configure stack port (for signaling port/link) and enable it.	**1.** Select the port and click Edit. **15.** Select Enable for Admin. **16.** Select Stack for Type. **17.** Click OK.
3.	Create the redundancy profile by giving it a name and configuring parameters for the redundancy profile such as the signaling port and protection role (secondary).	**1.** Select Inline Bypass > Redundancies **18.** Click New. **19.** Enter a name for the profile in the Alias field. For example, RP_001. **20.** Click in the Signaling Port field and select the stack port configured in Task 2. **21.** Select Secondary for Protection Role. **22.** Click Save.
4.	Configure inline network. This step associates the redundancy profile to the inline network and also disables link fail propagation on the inline network.	**1.** Select Inline Bypass > Inline Tools. **23.** Click New. **24.** Enter a name for the inline network in the Alias field. For example, IN_001. **25.** Click in the Port A field and select inline network port for network A. **26.** Click in the Port B field and select an inline network port f. **27.** Select Bypass for Traffic Path. (This is the default setting.) **28.** Make sure Link Failure Propagation option is not checked. (It is enabled by default.) **29.** Click in the Redundancy Profile field and select the profile created in step 3. For example RP_001.
5.	Configure inline tool ports, port type (inline-tool), and administratively enable them.	**1.** Go to **System > Ports > Ports > All Ports.** **30.** Select the first port (for example, 1/4/x1) to configure as an inline-tool port. **31.** Click Quick Port Editor. **32.** Select Inline Network for Type and select Enable for Admin. **33.** Click OK. **34.** Select the second port (for example, 1/4/x2) and

Task	Description	UI Steps
		repeat steps 3 through 5.
6.	Configure inline tool and failover action. then enable inline tool.	**1.** Select Inline Bypass > Inline Tools **35.** Enter a name for the inline tool in the Alias field. For example, IT_001 **36.** Click in the Port A field and select the first inline tool port configured in Task 5. **37.** Click in the Port B field and select the second inline tool port configured in Task 5. **38.** Make sure Enable is selected. (It is enabled by default **39.** Select NetworkBypass for Failtover action. **40.** Click Save.
7.	Configure map passall, from inline network to inline tool.	**1.** Select Maps > Maps > Maps. **41.** Click New. **42.** Enter a name for the map in the Alias field. For example, INtoIT. **43.** Select Regular for Type. **44.** Select Pass All for Subtype. **45.** Select the inline network created in Task 4 for Source. For example, IN_001. **46.** Select the inline tool created in Task 6. For example, IT_001.
8.	Configure the path of the traffic to the inline tool, disabling physical bypass on the inline network to open the relay on the node with the primary role.	**1.** Select Inline Bypas > Inline Networks. **47.** Select the Inline Network configured in Task 4 and click Edit. **48.** Select To Inline Tool for Traffic Path **49.** Make sure Physical Bypass is NOT checked. **50.** Click Save.

Timestamps

This chapter describes how to use the timestamp feature. The timestamp feature provides accurate time information to measure application and network performance.

Required License: Advanced Feature License

Refer to the following sections for details:

- About Timestamps
- Using Timestamps
- Why PTP?
- Using PTP to Timestamp Packets
- Enabling Timestamp on a Port
- Viewing PTP Details
- PTP and Timestamp—Rules and Notes

About Timestamps

Timestamps provide high-definition timing information for installations where packet timing is critical, particularly in high frequency trading, financial, and service provider environments. Timestamps are easily correlated across worldwide installations, improving your ability to monitor latency, Quality of Service (QoS), jitter, and end-to-end network performance across multiple instrumentation points.

This feature provides you the ability to append timestamps to ingress or egress packets. Timestamps are used in conjunction with Precision Time Protocol (PTP) for synchronization with a UTC reference clock.

You can choose to append timestamps to ingress packets, or to egress packets, or to both ingress and egress packets. Figure 47Timestamping illustrates how timestamps are added to ingress and egress packets.

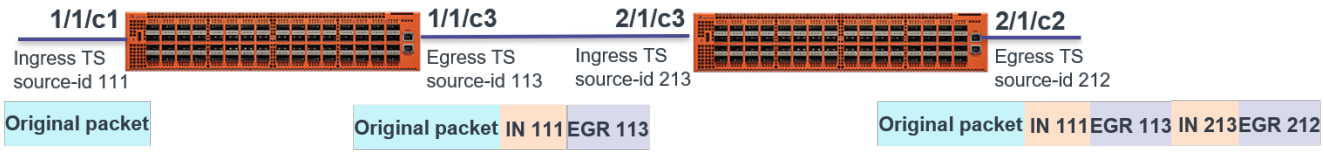


Figure 47 *Timestamping*

In this figure, the original packet is ingress timestamped on the 1/1/c1 port in a GigaVUE-TA200 device and egress timestamped on the 1/1/c3 port. The packet is then forwarded to the next device, where ingress and egress timestamps are again added at the 2/1/c3 and 2/1/c2 ports respectively. Thus, the packet has both ingress and egress timestamps with unique source IDs for each port, appended to its header. The source IDs provide port identification where timestamp is applied. You can add a maximum of two timestamps for a packet in a device.

Using Timestamps

You must perform the following tasks to use the timestamp feature:

S.No	Task	Using GigaVUE-FM, Refer to:	Using GigaVUE-OS CLI, Refer to:
1.	Select PTP as the time synchronization source for your device. Using PTP, you can synchronize distributed clocks with nanoseconds accuracy.	Configuring PTP Globally	*"ptp"* command section in the *"GigaVUE-OS CLI Reference Guide"*
2.	After you enable PTP on the required network port, check the port's PTP clock state, PTP configuration, and the synchronization status.	Viewing the PTP Configuration on a Port	*"Related Commands" section under the "ptp"* command section in the *"GigaVUE-OS CLI Reference Guide"*
3.	Enable timestamp on the required ports. You can add timestamps to ingress or egress packets, or to both the packets.	Enabling Timestamp on a Port	*"port"* command section in the *"GigaVUE-OS CLI Reference Guide"*
4.	View the PTP statistics for troubleshooting purposes.	Viewing PTP Details	*"Related Commands" section under the "ptp"* command section in the *"GigaVUE-OS CLI Reference Guide"*

Why PTP?

PTP is a time synchronization protocol defined in IEEE 1588 version 2 (ITU-T G.8275.1) for devices across a network. Using PTP, you can synchronize distributed clocks with nanoseconds accuracy. The clocks in the PTP-enabled devices follow a source-receiver hierarchy. The receivers are synchronized to their sources. The top-level source is called the primary source. The primary source clock is synchronized with a Global Positioning System (GPS).

The source will send a synchronization message to the receiver. In turn, the receiver sends a delay request to the source and the source reverts with a delay response. Thus, the delay is measured and the clock of the receiver is synchronized with the source. This is also called as the one-step synchronization.

PTP-enabled devices can have the following clock modes:

- Ordinary clock— A clock that has only one PTP-enabled port and maintains the timescale used in the domain. It can serve either as a source or a receiver.
- Boundary clock—A clock that has multiple PTP-enabled ports and maintains the timescale used in the domain. The ports can serve either as a source or a receiver.

The Best Master Clock (BMC) algorithm that is run on every clock is used to determine the best clock in a distributed network. The algorithm compares the attributes from two different clocks to determine the data that describes the better clock. The algorithm is used to determine which of the clocks described in the announce messages received by a local clock port is the best clock. It is also used to determine whether a newly discovered clock, which is a foreign source, is better than a local clock. For more information on the BMC algorithm, refer to the ITU-T G.8275.1 standard.

PTP Deployment shows an example of a distributed network, in which the PTP-enabled devices work in a source-receiver hierarchy.

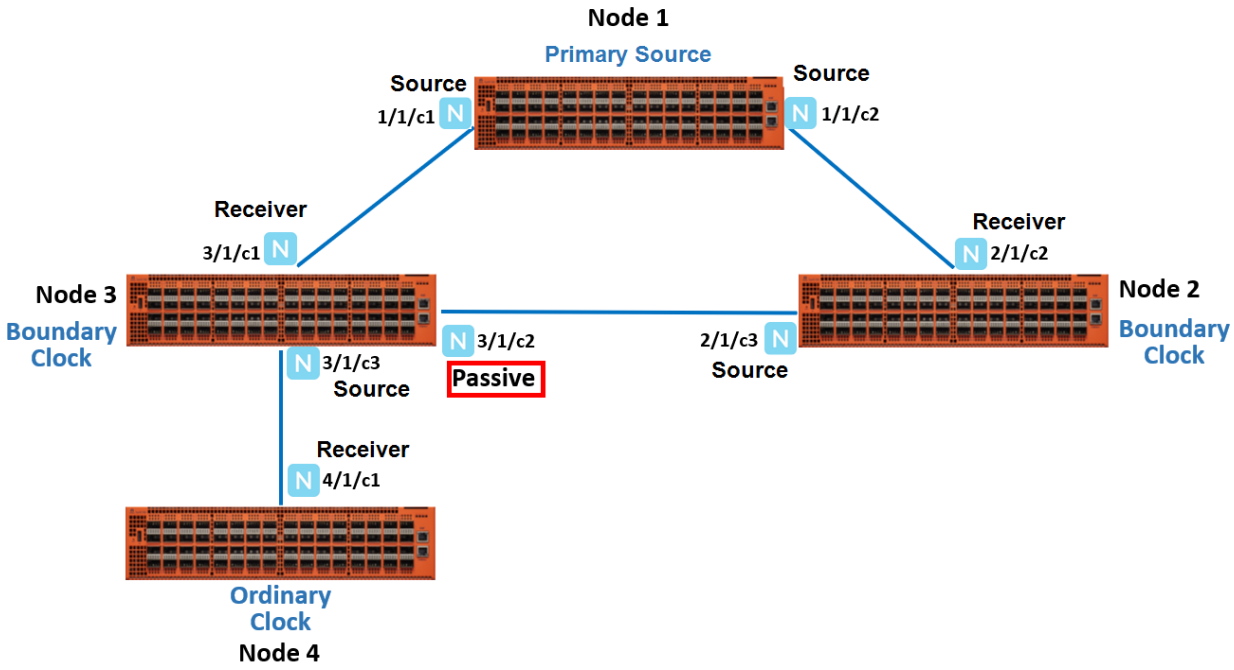


Figure 48 *PTP Deployment*

In this example, you have node 1, which is a top-level device that acts as a primary source. This device is linked to node 2 and node 3, which are boundary clocks. Node 3, in turn is linked to node 4, which is an ordinary clock. All these devices are synchronized in a source-receiver hierarchy. Node 3 is synchronized with the primary source through the network port, 3/1/c1, which acts as a receiver. The other port 3/1/c2 on node 3 is in passive state. When the link between the primary source and node 3 goes down, the network port, 3/1/c2, which

was in the passive state transitions to receiver state. This will ensure that node 3 as well as node 4 remain synchronized with the primary source through node 2 (refer to Figure 49PTP State Transition).

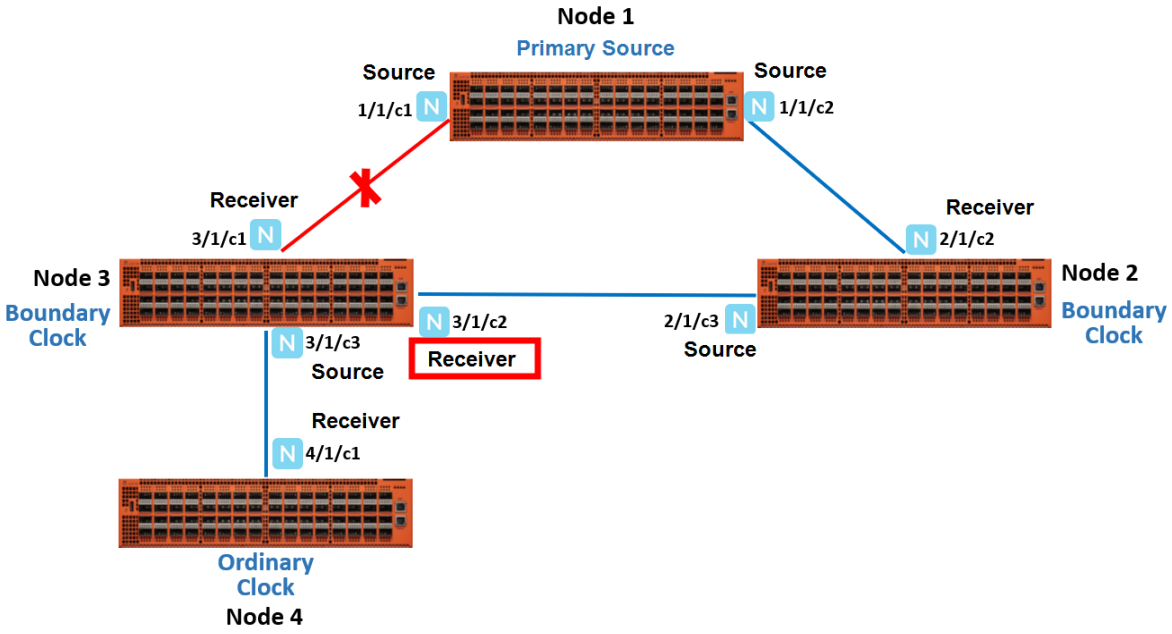


Figure 49 *PTP State Transition*

Limitations

The following restrictions are applicable for the above feature:

- A two step PTP clock mode is not supported.
- One step GigaVUE-TA200 clock cannot synchronize time with a two step PTP Primary Source.

Synchronization of the PTP and Local System Clock

From 5.13.01 version, PTP clock will be synchronized with the local system clock to provide accurate UTC timestamping. This process takes over in the absence of a Primary source clock and when the node is acting as a source clock. In this scenario, if there is a change in the local system clock due to NTP or manual configuration the same change will be synced to the PTP clock.

Recommendation: All local clocks that are running on PTP domain should share the same time value. This can be ensured through NTP server.

Using PTP to Timestamp Packets

Refer to the following sections about how to configure and enable PTP on the network ports to timestamp packets:

- Configuring PTP Globally
- Enabling PTP on a Network Port

Configuring PTP Globally

You must configure PTP globally on a device. You can configure the PTP domain, clock mode, and priority at the device level. Based on these PTP parameters, the primary source will be selected from the clocks in a network.

To configure PTP globally:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**..
2. In the Physical Nodes page, click the required cluster ID for which you want to configure PTP.
3. On the left-navigation pane, go to **System > Settings > Date and Time > PTP**.
4. Click **New** to open the PTP page.
5. In the **Alias** field, enter a unique name for the PTP template that you are configuring globally on the device. Ensure that you provide a unique name for each device in the cluster. You can also choose to leave the field blank and a default alias will be added automatically. For example, ptp_default_<box_id>.
6. From the **Box ID** drop-down list, select the box ID of the device for which you want to configure PTP.
7. In the **Domain** field, enter a PTP domain number that will be used for this clock. The valid range includes 0 and the numbers between 24 and 43.
   - Domain 0 for IEEE 1588 default profile.
   - Domain 24 - 43 for ITU-T G.8275.1 Telecom profile.
8. From the **Mode** drop-down list, select **Boundary** or **Ordinary** as the clock mode. The default is **Boundary**.
9. In the **Priority 2** field, enter a priority value that will be used to determine the primary source in a network. For example, if there are two clocks in a network that match the default criteria, the clock that has the lower priority value will be selected as the primary source. The valid range is between 0 and 255. The default is 128.

10. In the **Local Priority** field, enter a priority value that will be used to determine the PTP source when the BMC algorithm chooses more than one clock as the source. This value overrides the default criteria for the BMC selection. The clock with the lower priority value will be selected.

    For example, on a particular device, the local priority of a port is compared with the local priority of the clock. Based on the following criteria, the source is determined:
    - If the local priority of the port is greater than the local priority of the clock, the clock becomes the source.
    - If the local priority of the port is lesser than the local priority of the clock, the port becomes the source.
    - If the local priority of the port is equal to the local priority of the clock, the clock identity of the port and the clock is compared and the one that has the lower value becomes the source.

    The valid range is between 1 and 255. The default is 128.
11. Click **OK**.

Enabling PTP on a Network Port

After you configure PTP globally on a device, you must enable PTP on the required network ports. You can also use the fanout ports to configure PTP. Once you have enabled PTP on a port, you cannot use the port for any other services such as maps. To configure PTP on a network port:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**..
2. In the Physical Nodes page, click the required cluster ID for which you want to configure timestamp.
3. On the left-navigation pane, go to **Ports > Ports > All Ports**.
4. Select a network port on which you want to enable PTP, and then click **Edit**.
5. Scroll down to the PTP section, and then select the **Enable** check box.
6. In the **Local Priority** field, enter a priority value that will be used to determine the PTP source when the BMC algorithm chooses more than one clock as the source. This value overrides the default criteria for the BMC selection. The clock with the lower priority value will be selected.

    For example, on a particular device, the local priority of a port is compared with the local priority of the clock. Based on the following criteria, the source is determined:
    - If the local priority of the port is greater than the local priority of the clock, the clock becomes the source.
    - If the local priority of the port is lesser than the local priority of the clock, the port becomes the source.

- If the local priority of the port is equal to the local priority of the clock, the clock identity of the port and the clock is compared and the one that has the lower value becomes the source.

The valid range is between 1 and 255. The default is 128.

7. From the **Role** drop-down list, select one of the following options:
   - **Source**—The port will act as the source with which all the connected ports with receiver role will be synchronized.
   - **Follower**—The port will act as the receiver and will be synchronized with the source port to which it is connected.
   - **Standard**—The BMC algorithm will determine the port's role as source or receiver.

   NOTE: The default is **Standard**.
8. In the **Announce Interval** field, enter the time interval, in seconds on a logarithmic scale, between PTP announce messages in the port.
   - For Domain 0, the valid range is between -2 and 4 log seconds. The default is 1.
   - For Domain 24-43, the valid range is between -3 and 4 log seconds. The default is -3.
9. In the **Delay Request Interval** field, enter the minimum interval allowed between PTP delay request messages. The valid range is between -7 and 0 log seconds.
   - For Domain 0, the default is -5.
   - For Domain 24-43, the default is -4.
10. In the **Sync Interval** field, enter the time interval between PTP synchronization messages on the port. The valid range is between -7 and 0 log seconds.
    - For Domain 0, the default is -5.
    - For Domain 24-43, the default is -4.
11. In the **VLAN** field, enter the PTP VLAN value for the port. The valid range is between 2 and 4000. The VLAN ID would be allocated based on availability of the VLAN resource.

    Hover over the help icon to know about the available VLAN resources that can be allocated.
12. Click **OK** to save the configuration.

Enabling Timestamp on a Port

You can configure ingress timestamp, egress timestamp, or both ingress and egress timestamps on a port. Refer to Configuring PTP Globally for the ports that are supported for timestamps.

To configure timestamp on a port:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**..

2. In the Physical Nodes page, click the required cluster ID for which you want to configure timestamp.
3. Go to **System > Ports > Ports > All Ports**.
4. Select a port on which you want to configure timestamp, and then click **Edit**.
5. Scroll down to the Timestamp section.
6. Select the **Enable** check box corresponding to the **Ingress** or **Egress** fields.
7. In the **Source ID** field, enter a unique ID for the ingress or egress timestamp. The packets that ingress or egress the port will be timestamped with this unique source ID. The source ID enables you to identify the port on which the timestamp is applied.
8. Click **OK** to save the configuration.



Figure 50 *Timestamp Configuration*

After you enable timestamp on the various ports, you can view the timestamp details for a list of ports in the Timestamp page. To view this page, go to **Settings > Date and Time > Timestamp**.

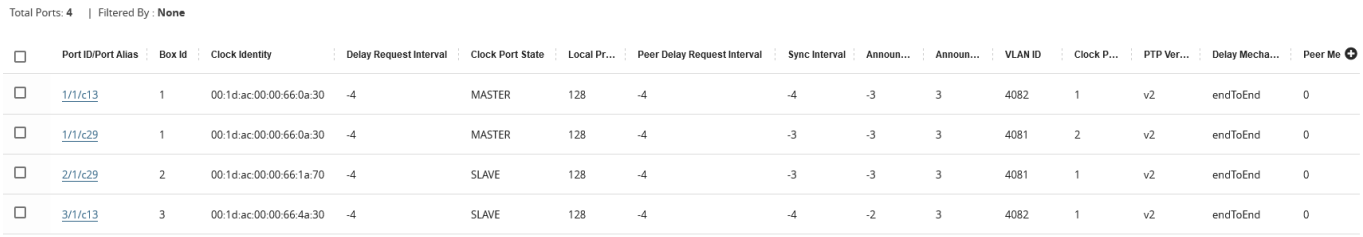
Viewing PTP Details

Following table provides information about the different PTP details that you can view:

To view:	Refer to:
The PTP configuration on a port	Viewing the PTP Configuration on a Port
The PTP clock details such as the local PTP clock, parent clock state, and the time property for a device	Viewing the PTP Clock Details
The PTP details of the source that is learned by a port	Viewing the PTP Foreign Source Details
The PTP statistics	Viewing PTP Statistics

Viewing the PTP Configuration on a Port

After you enable PTP on a network port, you can view the port's clock state and other details of PTP configuration in the PTP Port State page. To view the page, go to **Settings > Date and Time > PTP Port State**. Figure 51PTP Port State shows the PTP configuration details.



Total Ports: **4** | Filtered By : **None**

Port ID/Port Alias	Box Id	Clock Identity	Delay Request Interval	Clock Port State	Local Pr...	Peer Delay Request Interval	Sync Interval	Announ...	Announ...	VLAN ID	Clock P...	PTP Ver...	Delay Mecha...	Peer Me
1/1/c13	1	00:1d:ac:00:00:66:0a:30	-4	MASTER	128	-4	-4	-3	3	4082	1	v2	endToEnd	0
1/1/c29	1	00:1d:ac:00:00:66:0a:30	-4	MASTER	128	-4	-3	-3	3	4081	2	v2	endToEnd	0
2/1/c29	2	00:1d:ac:00:00:66:1a:70	-4	SLAVE	128	-4	-3	-3	3	4081	1	v2	endToEnd	0
3/1/c13	3	00:1d:ac:00:00:66:4a:30	-4	SLAVE	128	-4	-4	-2	3	4082	1	v2	endToEnd	0

Figure 51 *PTP Port State*

Following table lists the PTP port states and their descriptions:

PTP Port State	Description
Port ID/Port Alias	The identifier of the port on which PTP is enabled.
Box ID	The box identifier of the device.
Clock Identity	The clock identity of the device. It is a 64-bit global identifier (EUI-64) as defined by the IEEE standard.
Delay Request Interval	The minimum interval allowed between PTP delay request messages when the port is in the source state.
Clock Port State	Based on the role defined for the port, the clock state will be Master or Slave. Following are the clock port states: • Initializing—The port is initializing it's datasets, hardware, and communication facilities. The port does not place any PTP messages on it's communication path during this state. If one port on the boundary clock is in the Initializing state, all the other ports on the boundary clock will also be in the same state. • Listening—Only signaling messages or management messages that are responses to another management message are placed on the port's communication path. • Faulty—The fault state of the protocol. Only signaling messages or management messages that are responses to another management message are placed on the port's communication path. In a boundary clock, any activity on the port with the Faulty clock port state does

PTP Port State	Description
	not affect the other ports in the device. If any activity on the faulty port is not restricted to the port, all other ports in the device changes to the Faulty state. • Disabled—The port does not place any messages on it's communication path. In a boundary clock, the port's activity does not affect the activity in other ports that are part of the boundary clock. A port in this state discards all PTP received messages except the management messages. • Pre-master—The port acts like a source but does not place any messages on it's communication path, except for signaling or management messages. • Master—The port becomes the source. • Passive—The port does not place any messages on it's communication path. • Uncalibrated—One or more source ports are detected in the domain. The appropriate source port has been selected, and the local port is preparing to synchronize to the selected source port. This is a transient state to initiate the servo synchronization, datasets updates, and other implementation-specific activities. • Slave—The port is synchronizing to the selected source port.
Local Priority	The priority value used to determine the PTP source when the BMC algorithm chooses more than one clock as the source. This value overrides the default criteria for the BMC selection. The clock with the lower priority value will be selected. The valid range is between 1 and 255. For example, on a particular device, the local priority of a port is compared with the local priority of the clock. Based on the following criteria, thesource is determined: • If the local priority of the port is greater than the local priority of the clock, the clock becomes the source. • If the local priority of the port is lesser than the local priority of the clock, the port becomes the source. • If the local priority of the port is equal to the local priority of the clock, the clock identity of the port and the clock is compared and the one that has the lower value becomes the source.
Peer Delay Request Interval	The minimum interval allowed between PTP delay request messages between peer ports.
Sync Interval	The time interval between PTP synchronization messages on the port.

PTP Port State	Description
Announce Interval	The time interval, in log seconds, between PTP announce messages in the port.
Announce Receipt Timeout	The maximum number of consecutive announce messages that the port fails to receive. When the number of consecutive announce messages that the port failed to receive is greater than the value specified in this field, the port's clock state changes to source.
VLAN ID	The identifier of the VLAN interface.
Clock Port ID	The port identifier assigned by the PTP clock.
PTP Version	The PTP version that is enabled in the port.
Delay Mechanism	The mechanism used to determine the delay between a source and a receiver. Following is the delay mechanism that Gigamon supports: • End-to-End—The source and the receiver send delay requests and delay responses back and forth between themselves to measure the delay. Once the delay is measured, the receiver synchronizes the time with the source.
Peer Mean Path Delay	Reserved for future use.

Viewing the PTP Clock Details

Use the PTP Clock page to view the local PTP clock, parent clock state, and the time property for the devices for which you have configured PTP. To view this page, go to **Settings > Date and Time > PTP Clock**. Figure 52PTP Clock Details shows the PTP clock details.

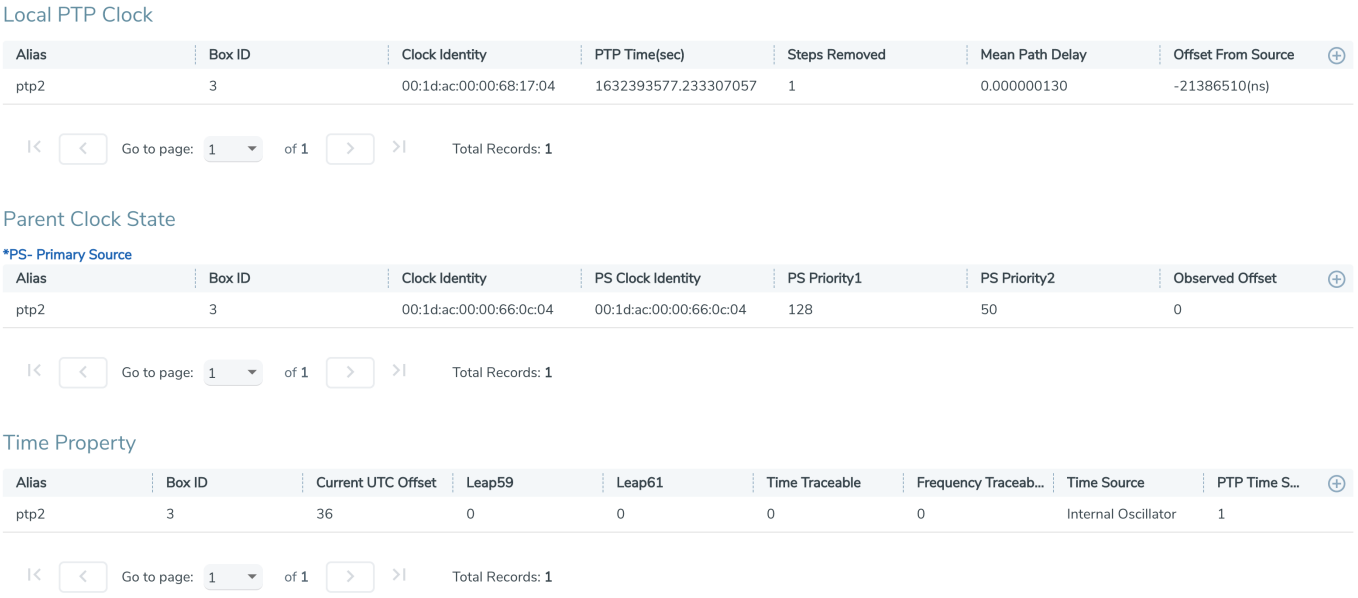


Figure 52 *PTP Clock Details*

Following table lists the PTP clock attributes and their descriptions:

PTP Clock Attribute	Description
Local PTP Clock	
Alias	The unique name for the PTP template that you configured globally for the device.
Box ID	The box identifier of the device.
Clock Identity	The clock identity of the device. It is a 64-bit global identifier (EUI-64) as defined by the IEEE 1588 standard.
Ptp Time (sec)	The actual PTP time with nanosecond resolution.
Mean Path Delay	The forward and reverse path delay in steady state is used to calculate the mean path delay.
Steps Removed	The number of boundary clocks between the local clock and the foreign source clock.
Domain	The network within which PTP operates, that is all the clocks within a domain are in synchronization.
Offset From Source	The time difference between the source and the receiver, measured in nanoseconds. Once this is computed, the receiver will readjust its clock to align with the source.

PTP Clock Attribute	Description
Priority2	The priority value used to determine the primary source in a network. For example, if there are two clocks in a network that match the default criteria, the clock that has the lower priority value will be selected as the primary source.
Mode	The clock modes. The values are: • Ordinary clock—A clock that has only one PTP-enabled port and maintains the timescale used in the domain. It can serve either as a source or a receiver. • Boundary clock—A clock that has multiple PTP-enabled ports and maintains the timescale used in the domain. The ports can serve either as a source or a receiver.
Port Ptp Count	Number of ports on which the PTP template is enabled. Gigamon supports a maximum of 10 boundary clock ports in a device.
Clock Quality Accuracy	The clock quality announced by the source. This value is used to select the best clock source between two PTP source clocks and switches dynamically to the clock that has greater accuracy.
Step Type	The timestamp mode that is used by PTP. The values is: • One-step mode—PTP source does not send any follow up message because the synchronization message itself contains the transmit timestamp as well as the correction filed to obtain the precise transmit timestamp.
Local Clock Time	The actual local clock time on the device.
Local Priority	The priority value used to determine the PTP source when the BMC algorithm chooses more than one clock as the source. This value overrides the default criteria for the BMC selection. The clock with the lower priority value will be selected.
Clock Quality Class	Attribute of an ordinary or boundary clock that denotes the traceability of the time or frequency distributed by the primary source clock.
Clock Quality Offset	The time difference in clock quality between the source and the receiver, measured in nanoseconds.
Parent Clock State	
PS Clock Identity	The clock identity of the primary source. It is a 64-bit global identifier (EUI-64) as defined by the IEEE 1588 standard.
PS Priority1	The priority value of the primary source clock. Lower value takes precedence.

PTP Clock Attribute	Description
PS Priority2	The priority value of the primary source that is used when **PS Priority1** value is the same for different sources in a network.
Observed Offset	The time difference that is observed between the source and the receiver.
Clock Port ID	The port identifier assigned by the PTP clock.
Stats	Indicates whether the values of parentDS.observedParentOffsetScaledLogVariance and parentDS.observedParentClockPhaseChangeRate have been measured and are valid.
Observed Phase Change Rate	An estimate of the parent clock's phase change rate as measured by the receiver clock.
PS Clock Quality Accuracy	The expected accuracy of the primary source clock.
PS Clock Quality Class	The primary source clock's traceability of the distributed time or frequency.
PS Clock Quality Offset	The time difference in clock quality between the primary source and the receiver, measured in nanoseconds.
Time Property	
Current UTC Offset	The time difference between International Atomic Time (TAI) and Universal Coordinated Time (UTC).
Leap59	The last minute of the current UTC day has only 59 seconds (instead of the 60 SI seconds).
Leap61	The last minute of the current UTC day has only 61 seconds (instead of the 60 SI seconds).
Time Traceable	The timescale and the UTC offset are traceable to a primary reference.
Frequency Traceable	The frequency that determines the timescale is traceable to a primary reference.
Time Source	The time source that is external to PTP. The time source is traceable to the international standards laboratories maintaining clocks that form the basis for the TAI and UTC timescales. Examples of these are GPS, NTP, and National Institute of Standards and Technology (NIST) timeservers.
Ptp Time Scale	The time scale to use when advertising time for PTP.

Viewing the PTP Foreign Source Details

Use the PTP Foreign Source page to view the PTP details of the source that is learned by a port. The foreign source record will be available only when a port is in the receiver state. To view this page, go to **Settings > Date and Time > PTP Foreign Source**.

In Figure 52PTP Clock Details, you can see that there is no foreign sourcerecord for the port 1/1/c13 because the port is in the source state. Whereas, the foreign source record appears for the port 1/1/c29, which is in the receiver state.

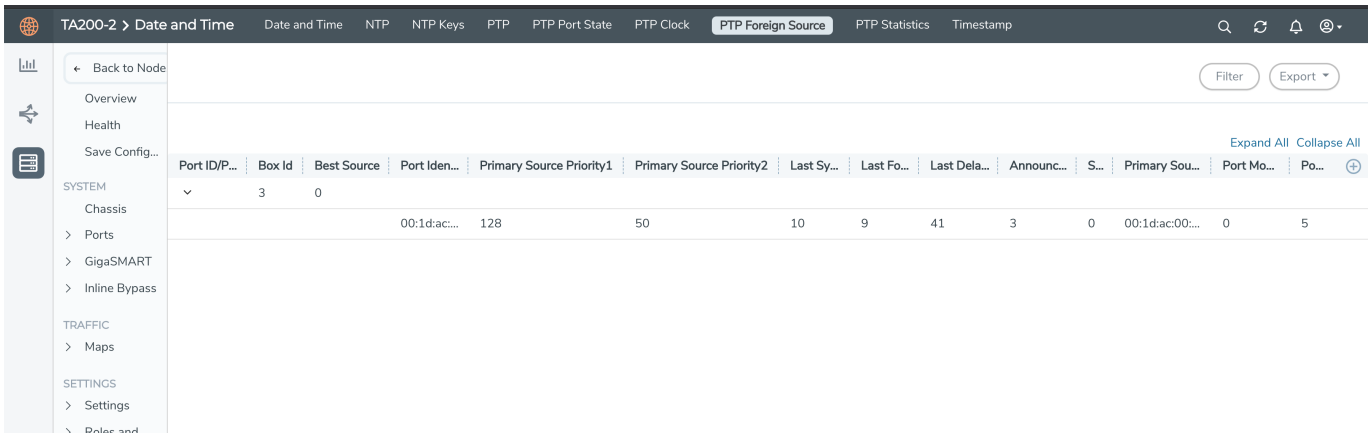


Figure 53 *PTP Foreign Source Details*

Following table lists the PTP foreign source attributes and their descriptions:

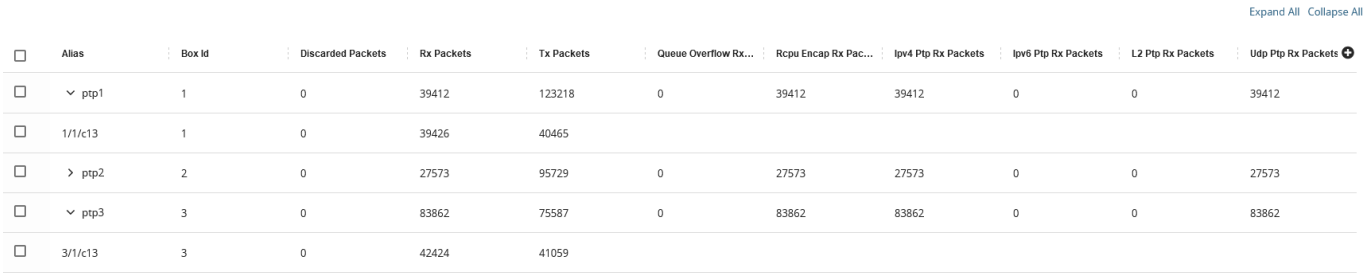
PTP Foreign Source Attribute	Description
Port ID / Port Alias	The port identifier.
Box ID	The box identifier of the device.
Best Source	The best source clock for the device.
Port Identity	The clock identity of the foreign source that the port has learned.
Primary Source Priority2	The priority value of the primary source that is used when **Primary source Priority1** value is the same for different sources in a network.
Announce Msg Count	The total number of announce messages that were sent and received by the port.

PTP Foreign Source Attribute	Description
Steps Removed	The number of boundary clocks between the local clock and the foreign source clock.
Primary Source Clock Identity	The clock identity of the primary source. It is a 64-bit global identifier (EUI-64) as defined by the IEEE standard.
Port Module Number	The module number of the physical port.
Port Number	The physical port identifier.
PTP Protocol	The protocol used by the PTP is IPv4 Ethernet.
Clock Address	The IP address that is derived based on the clock ID.
Primary Source Priority1	The priority value of the primary source clock. Lower value takes precedence.
Last Sync	The time when the last synchronization message was sent by the foreign source to the receiver.
Last Follow Up	The time when the last follow-up message was sent by the receiver to the foreign source.
Last Delay Response	The time when the last delay response was sent by the foreign source to the receiver.
Mtsd Scaled Avar	MTSD-scaled Allan variance (nsec-sq).
Quality Class	The traceability of the time or frequency distributed by the foreign source clock.
Quality Offset	The log variance field of the announce message.
Quality Accuracy	The clock accuracy indicates the expected accuracy of the foreign source clock.

Viewing PTP Statistics

Use the PTP Statistics page to view the PTP statistical data such as the received, transmitted, and discarded packets, queue overflow, and so on for a port on a device. You can also verify if there are any errors in the PTP configuration. You can choose to filter the data by box ID and PTP alias.

To view this page, go to **Settings > Date and Time > PTP Statistics**. Figure 54PTP Statistics shows the PTP statistics details.



Expand All Collapse All

	Alias	Box Id	Discarded Packets	Rx Packets	Tx Packets	Queue Overflow Rx...	Rcpu Encap Rx Pac...	Ipv4 Ptp Rx Packets	Ipv6 Ptp Rx Packets	L2 Ptp Rx Packets	Udp Ptp Rx Packets
☐	ptp1	1	0	39412	123218	0	39412	39412	0	0	39412
☐	1/1/c13	1	0	39426	40465						
☐	ptp2	2	0	27573	95729	0	27573	27573	0	0	27573
☐	ptp3	3	0	83862	75587	0	83862	83862	0	0	83862
☐	3/1/c13	3	0	42424	41059						

Figure 54 *PTP Statistics*

Following table lists the PTP statistical attributes and their descriptions:

PTP Statistical Attribute	Description
Alias	The unique name for the PTP template that you configured globally for the device.
Box ID	The box identifier for the device.
Discarded Packets	The number of packets that were discarded by the port because of mismatch in the domain or VLAN.
Rx Packets	The number of packets that were received by the port.
Tx Packets	The number of packets that were transmitted by the port.
Queue Overflow Rx Packets	The queue overflow of the received packets.
Rcpu Encap Rx Packets	The RCPU-encapsulated packets received.
Ipv4 Ptp Rx Packets	The IPv4 PTP packets received.
Ipv6 Ptp Rx Packets	The IPv6 PTP packets received.
L2 Ptp Rx Packets	The L2 PTP packets received.
Udp Ptp Rx Packets	The UDP PTP packets received.

PTP and Timestamp—Rules and Notes

Keep in mind the following rules and notes when working with PTP and timestamp:

- PTP is supported only on the GigaVUE-TA200 and GigaVUE-TA200Edevices.
- Timestamp is supported only on the GigaVUE-TA200 device.
- Timestamp is not supported on circuit, stack, and hybrid ports.
- PTP supports only one-step synchronization.

- PTP requires a dedicated front-end port and can be enabled only on network ports. You cannot use the PTP-enabled port for any other GigaVUE-OS features.
- Ensure that the VLAN ID and the domain for the PTP are the same for both the nodes between which the session needs to be established. However, the VLAN IDs of the ports within a node must be unique.
- All port-related PTP intervals such as, announce, sync, and delay request must be the same between nodes.
- Transparent clocks and peer-to-peer delay request mechanism are not supported.
- Priority 1 configurations are not supported based on the ITU-T G.8275.1 standard.
- Once you have configured PTP and timestamps are being captured, it is recommended that you do not change the PTP configurations so that the timestamps are accurate.
- If you want to upgrade to release version 5.13.00 with a VLAN configured PTP enabled port, it is recommended to remove the configurations before the upgrade and re-apply as required after the upgrade.

Fabric Maps

Fabric maps simplify how you create flow mapping across multiple clusters. Using GigaVUE-FM, you can create and manage fabric maps. You can provide cross-cluster flow mapping parameters by creating a fabric map. GigaVUE-FM then allocates the required circuit ID resources, along with generating and deploying cluster specific maps and fabric paths (refer About Circuit-ID Tunnels). Fabric paths can be defined as the link between the devices through which the traffic flows. With a successfully deployed fabric map, you can save time and effort because you no longer need to replicate rules at each hop on the network to manage cross-cluster traffic.

Refer Load Share Flow-Based Traffic to Identical Tools in a Fabric Tool Group Using Fabric Maps for more detailed information.

Supported Topologies

Fabric map support is available for the following topologies:

- Multihop Topology
- Leaf and Spine Topology
- Dual Multipath Leaf and Spine Topology

Refer to the following rules and notes:

- Source port of a fabric map can be either a network port, hybrid port or a vport.

- Destination port of a fabric map can be a tool port, hybrid port, GigaStream, Fabric Port Group or a vport.

  NOTE: Fabric port groups cannot be used as destination for subscriber intelligence solutions with load-balancing application.

- For topologies with leaf and spine clusters:
  - Do not configure the spine nodes as source or destination nodes. Source and destination must be configured either in the leaf nodes or in the node/clusters that are connected to the leaf nodes.
  - You can connect standalone nodes or clusters to all spine nodes of leaf-spine cluster.
  - You can connect the spine nodes to its leaf node. Spine nodes can also be connected to spine nodes of another leaf-spine cluster.
- Do not configure fabric maps in topologies that are not supported. Though GigaVUE-FM allows such configurations, it might lead to unpredictable results.

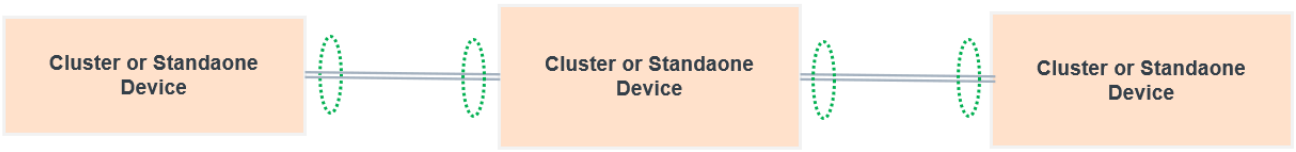
Multihop Topology

In this topology, multiple standalone nodes and/or non-leaf and spine clusters are connected using circuit GigaStreams.

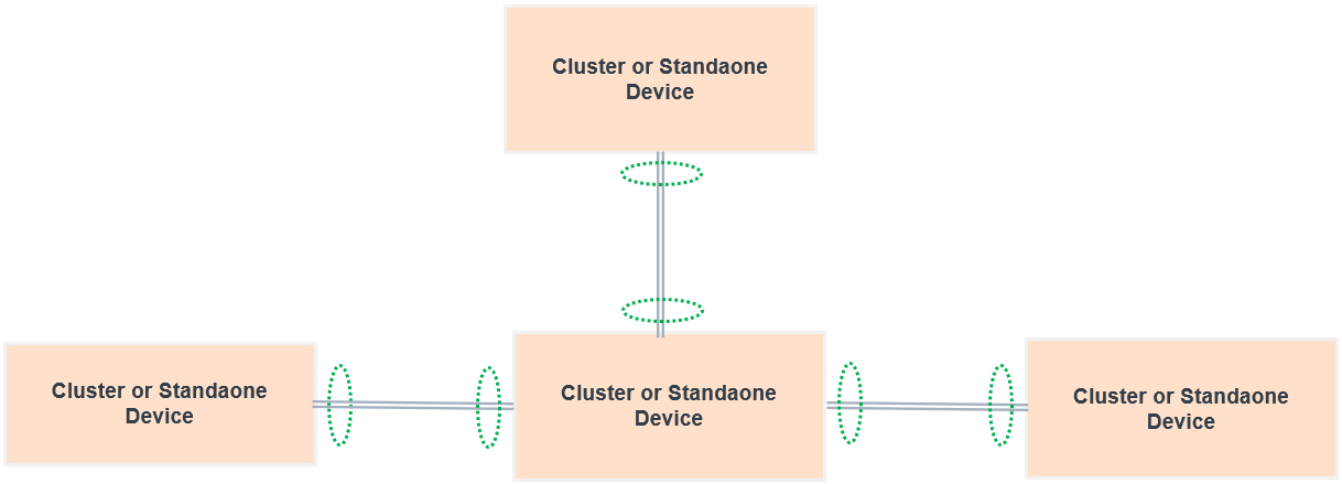
Refer to the following notes:

- Source and destination nodes can be configured anywhere in the topology.
- Redundant links must not be configured between the nodes, that is, there should not be more than one path from source to destination.

The Standalone nodes or clusters can be configured as follows:

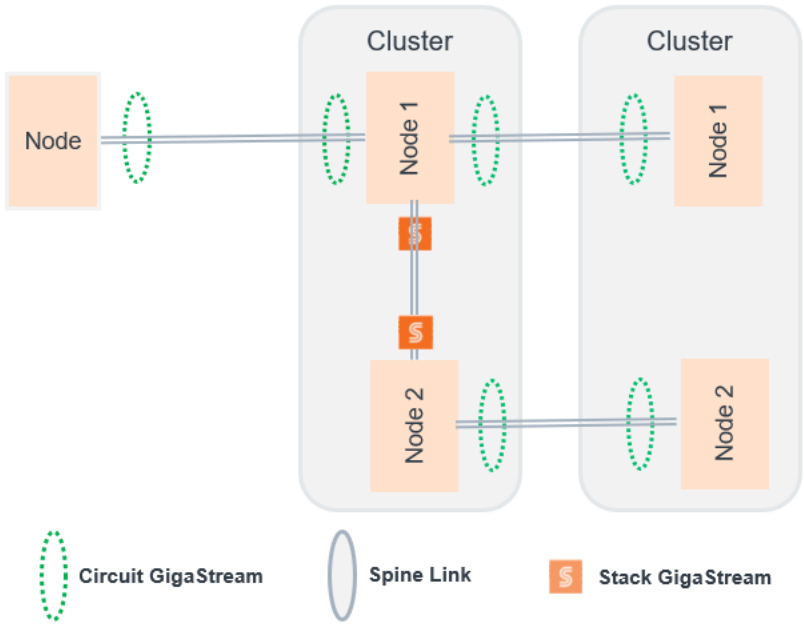


The standalone nodes or clusters can also be configured like a tree topology. Refer to the following figure.



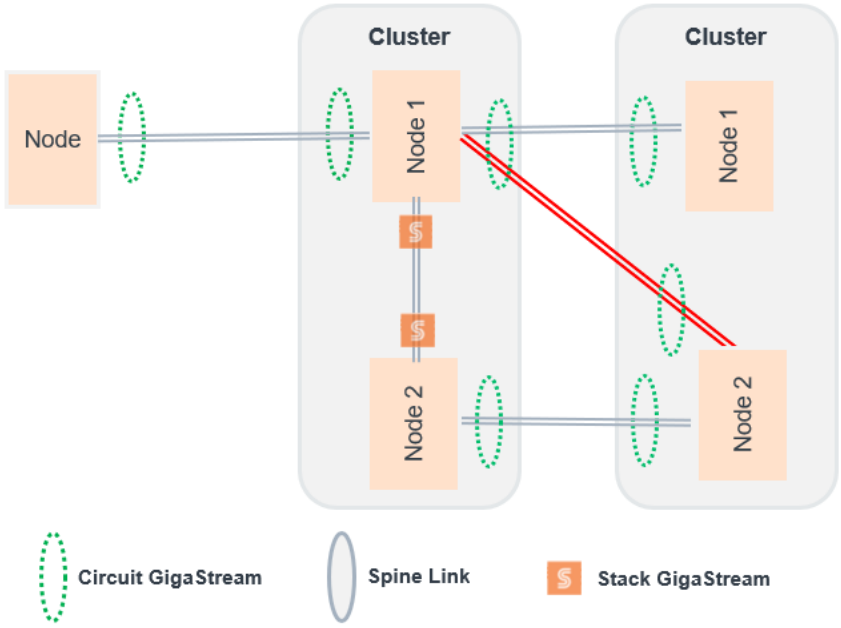
Supported Multihop Topology

The following multihop topology has only one path from source to destination.



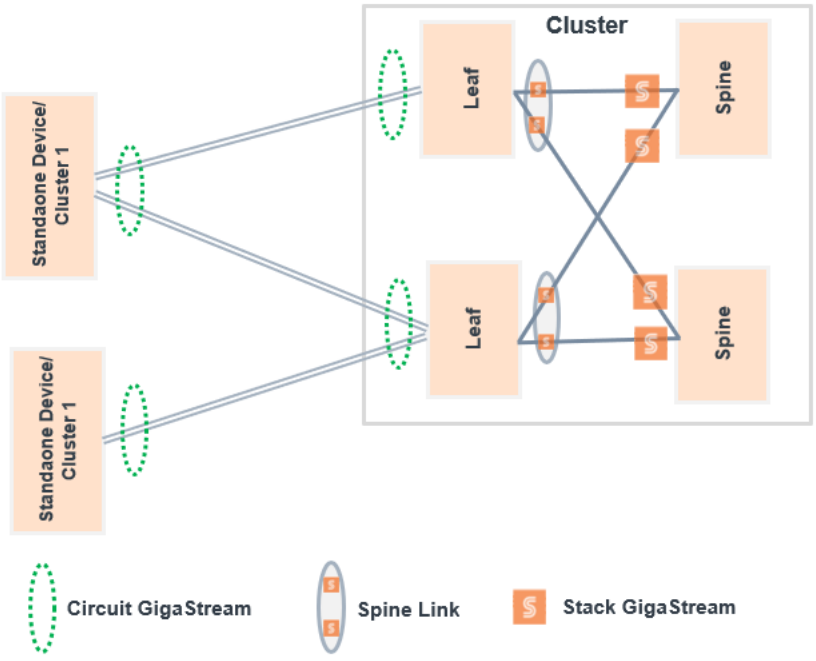
Unsupported Multihop Topology

The following multihop topology has more than one path from source to destination and therefore it is not supported.



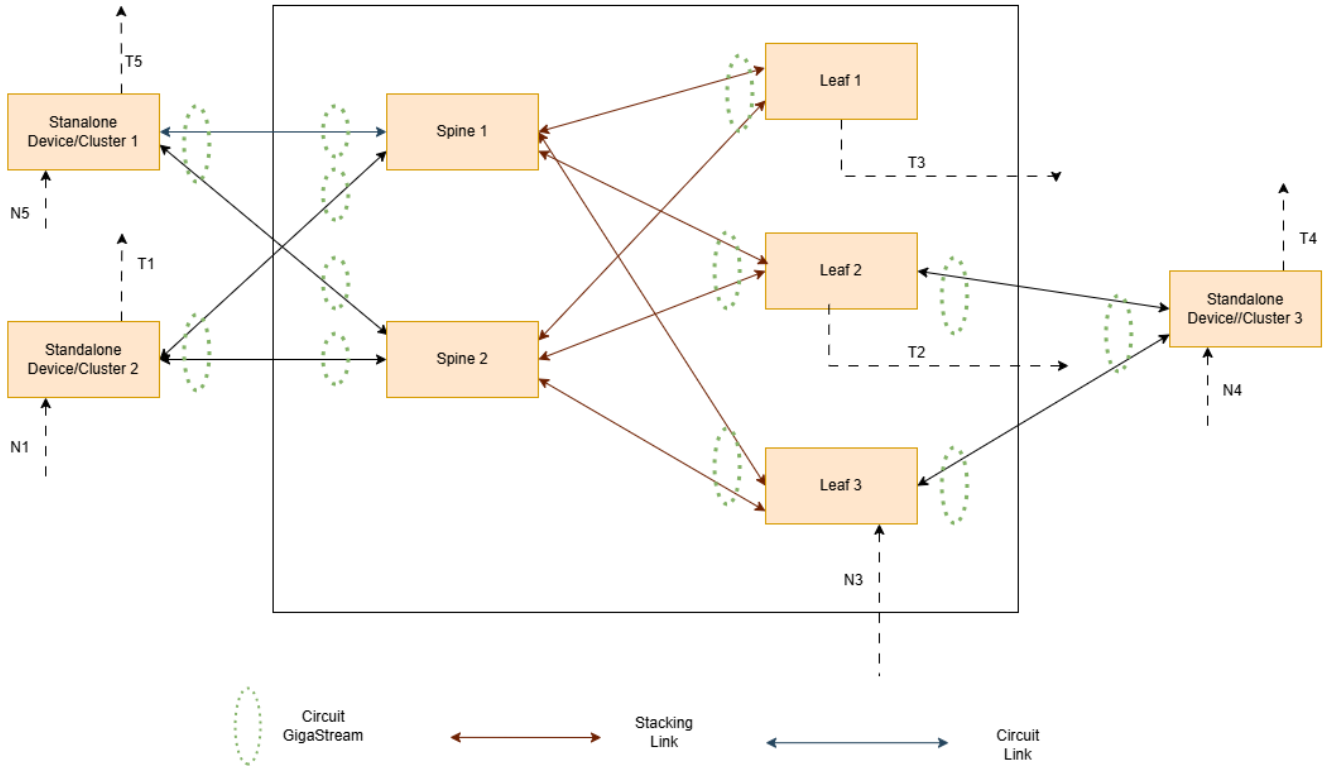
Leaf and Spine Topology

In this topology, a standalone node or a cluster is connected to any one of the leaf nodes in a leaf-spine cluster. It can also be load balanced to multiple leaf nodes of a leaf spine cluster. Refer to the Multi-Path Leaf and Spine section for details.



Spine Source Map Mode

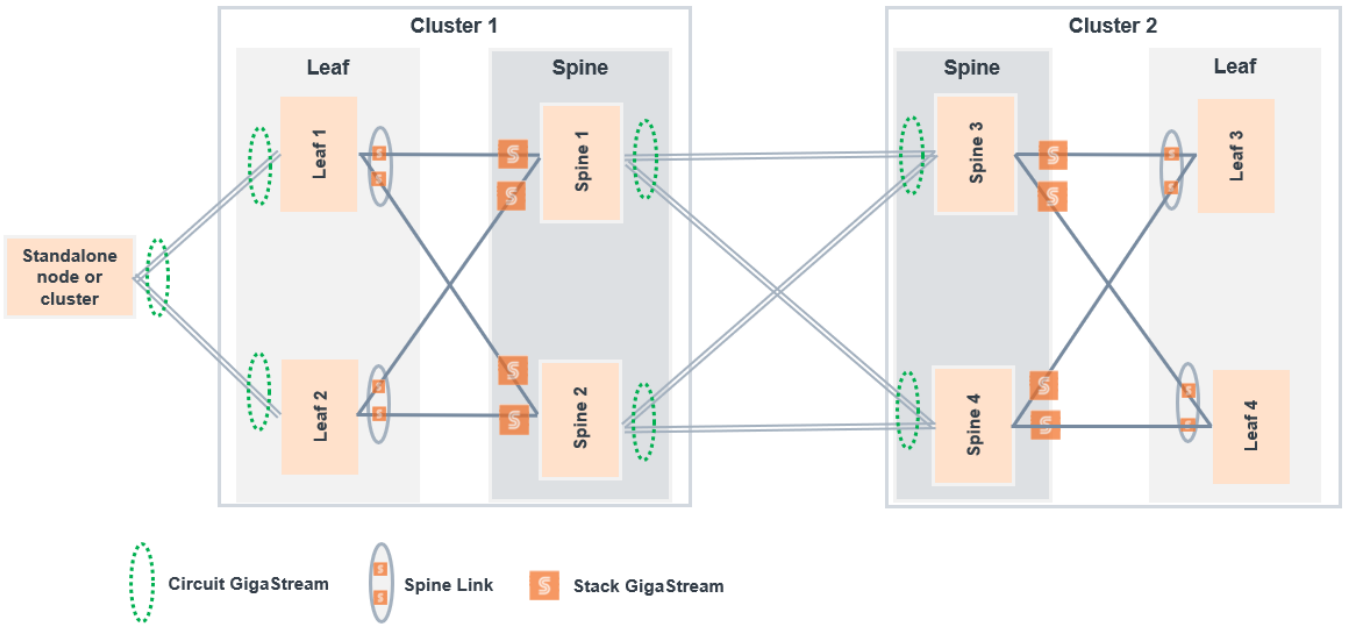
The Spine Source Map mode is a subset of leaf-and-spine topology that defines how Spine sources are visualized, either Separate or Combined within a leaf-and-spine fabric map.



The COMBINE setting of this topology optimizes fabric map deployments so that the cluster can use a single logical map to forward the traffic to the destination, even when traffic arrives via multiple decap circuit GigaStreams. This mode introduces a Spine Source Map Mode setting which is global in the FM. Refer to Configure Fabric Map Settings for more information on how to configure spine source map mode.

Dual Multipath Leaf and Spine Topology

In this topology, a leaf-spine cluster is connected to another leaf-spine cluster by configuring circuit GigaStreams across the spine nodes. Each spine node in a cluster is load-balanced with spine nodes in another cluster.

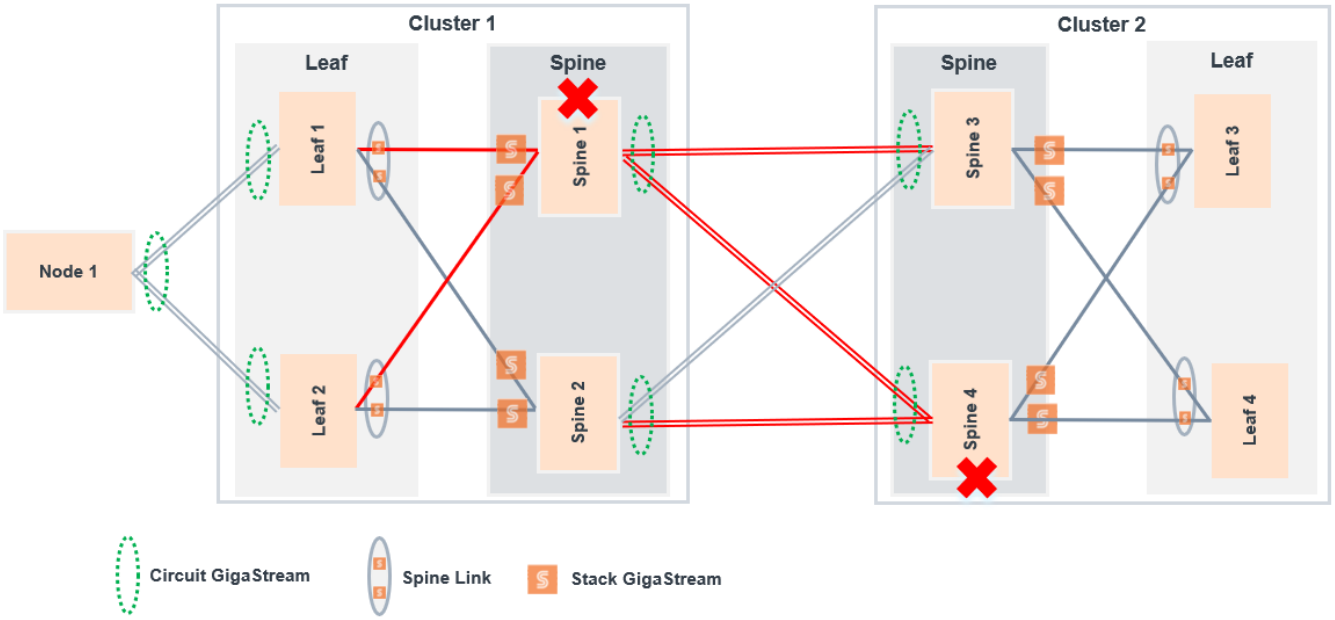


Supported Failovers

In a dual multipath leaf and spine cluster topology, traffic from one cluster to another cluster will be routed in case of the following failover scenarios:

- **Scenario 1**: At least one of the spine nodes in both clusters is in active state.
- **Scenario 2**: At least one circuit link connecting the spines across the cluster is active.

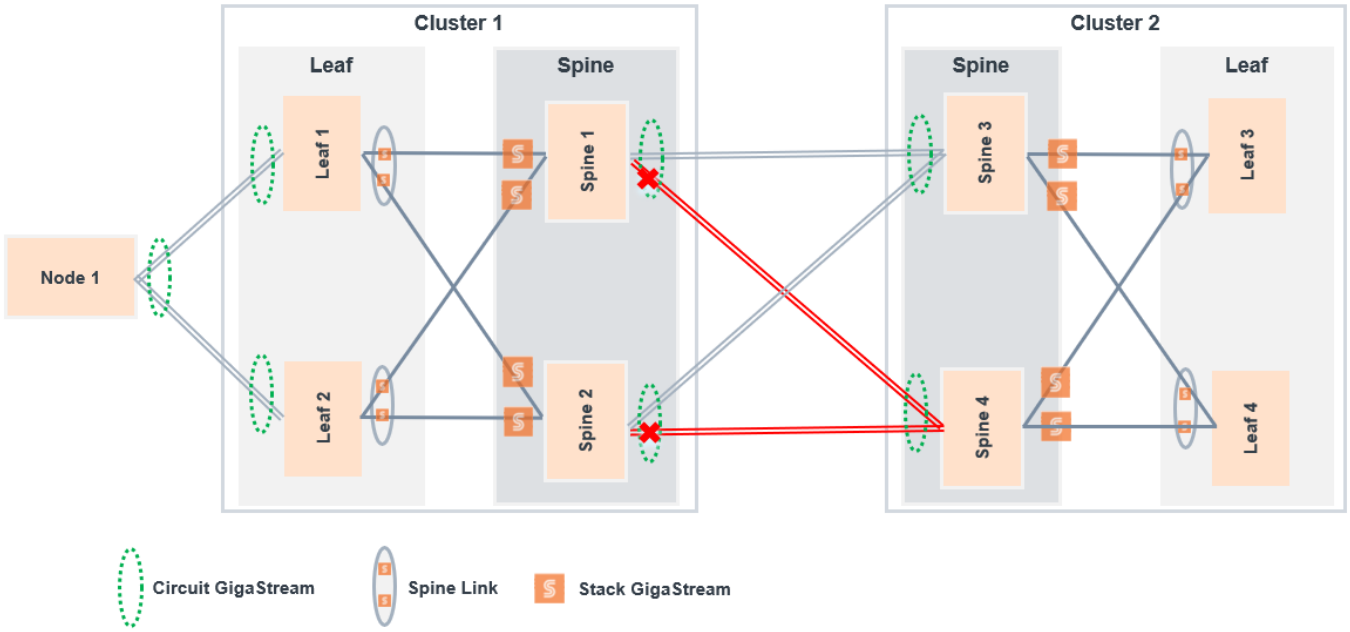
Scenario 1



In the above topology:

- Spine nodes 1 and 4 are down.
- Spine nodes 2 and 3 are up.
- Traffic is routed from leaf nodes on cluster 1 to leaf nodes on cluster 2 even though spine 1 and spine 4 are down.
- Traffic flows through the green links between spine 2 and spine 3.

Scenario 2



In the above topology:

- All spine nodes are active.
- Circuit GigaStream from Spine 1 to Spine 4 is down.
- Circuit GigaStream from Spine 2 to Spine 4 is down.
- Circuit GigaStream from Spine 1 to Spine 3 is up.
- Circuit GigaStream from Spine 2 to Spine 3 is up.
- Traffic from leaf nodes in cluster 1 are routed to leaf nodes in cluster 2 through circuit links from spine 1 and spine 2 to spine 3.

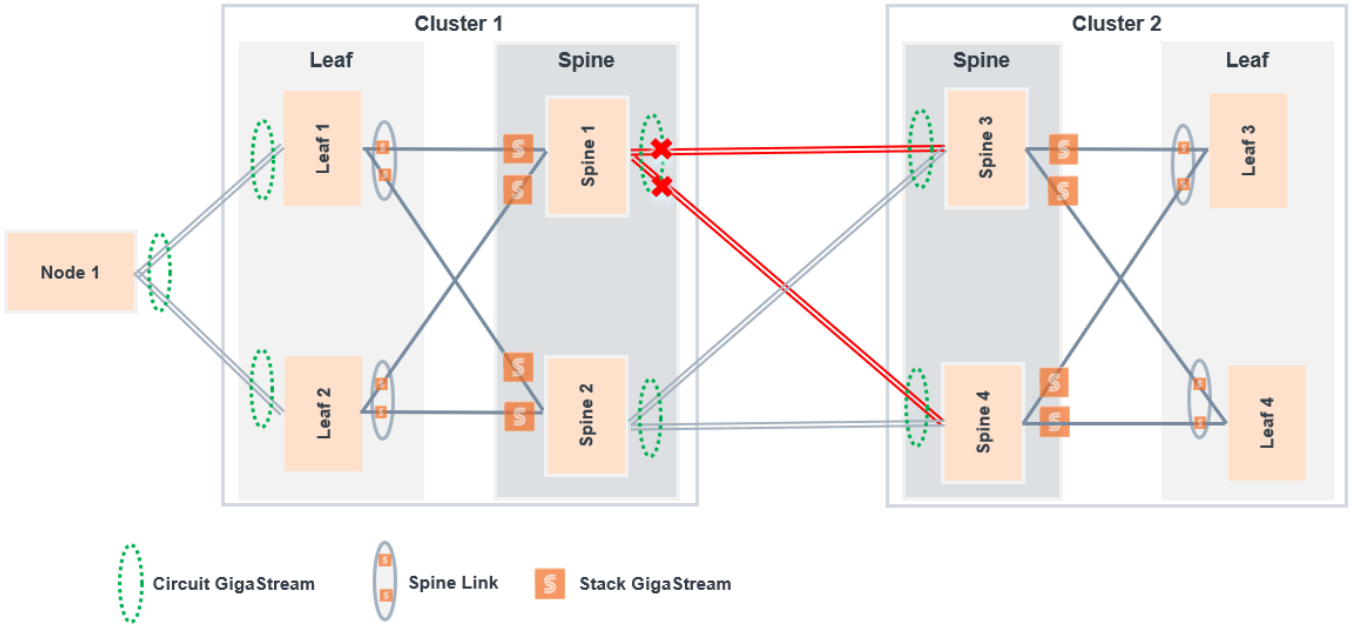
Thus, even if one of the links is up, traffic is routed across the leaf nodes in the clusters.

UnSupported Failovers

In a dual multipath leaf and spine cluster topology, traffic from one cluster to another cluster will not be routed in case of the following failover scenarios:

- **Scenario 1**: Circuit links connecting the spine nodes across cluster is down
- **Scenario 2**: Spine links connecting the leaf nodes to spine node is down.

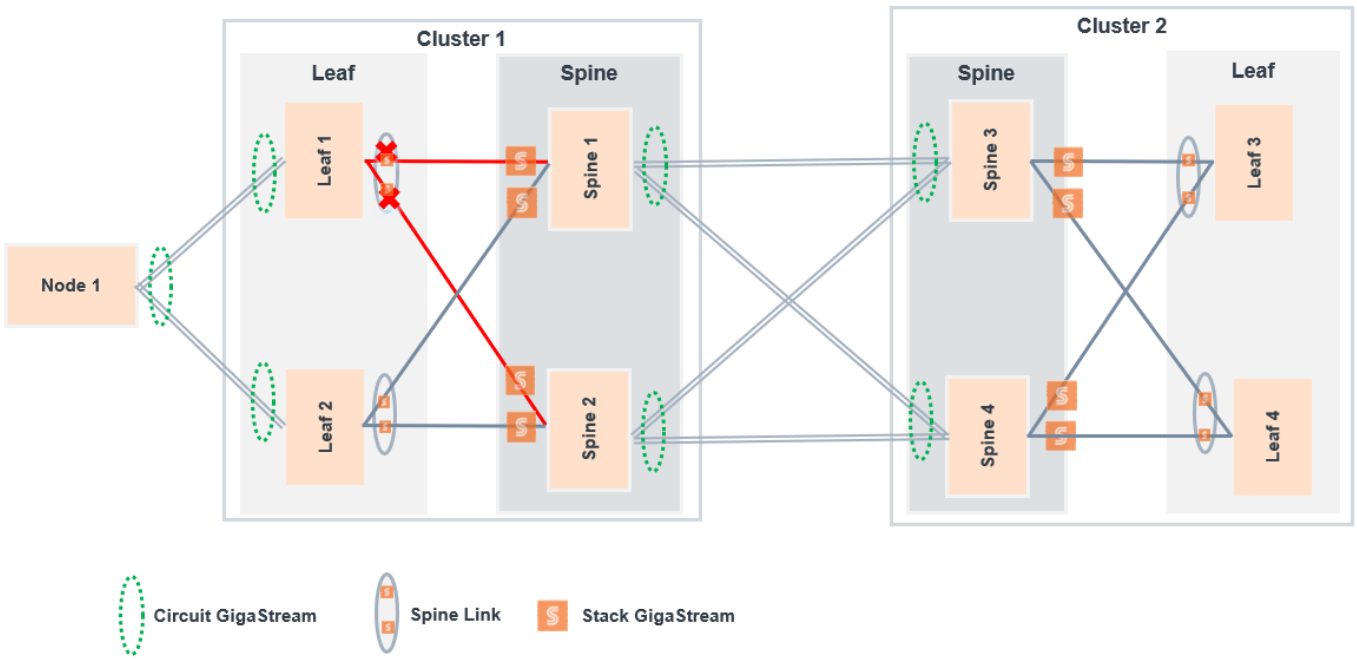
Scenario 1



In the above topology:

- Traffic from Leaf 1 and Leaf 2 is load-balanced to spine node 1.
- Traffic received on Spine 1 will not be delivered to the adjacent cluster as the following circuit GigaStreams are down:
  - Circuit GigaStream from Spine 1 to Spine 3
  - Circuit GigaStream from Spine 1 to Spine 4.

Scenario 2



In the above topology, traffic received on leaf 1 will be discarded as the following spine links are down:

- Spine link connecting leaf 1 to spine 1
- Spine link connecting leaf 1 to spine 2

Fabric Maps Prerequisites

To deploy fabric maps, you must complete the following prerequisites:

- Designate inter-cluster connection ports as circuit ports. Refer to Configure Ports
- Configure a regular circuit GigaStream with the required circuit ports. Refer to Configure Regular GigaStream.
- Create links between clusters. Refer to Create Links between Clusters.

Note: Keep in mind the following notes about circuit-id tunnels when configuring circuit ports:

Circuit-ID tunnels support the following:

- From Port in Regular Map
- GigaStream
- Tool port in regular map and collector map

- Port filter
- Port groups

Circuit-id tunnels does not support the following. However, you are not restricted from configuring these resources as circuit ports, but the functionality does not work as expected.

- Ingress ports with VLAN tag
- Port Pair
- Tool Mirror
- Inline Tool
- Pass-all/collector with from port as circuit-port

For more information about circuit-id tunnels, refer to Circuit Tunnels.

Using GigaVUE-FM, you can create connections between devices using manual topology links, or you can have the devices discovered through the Gigamon Discovery (GDP). After you create physical links between devices, the device configuration needs to be added into GigaVUE-FM, so that Fabric Maps can use it. Once the circuit and GigaStream ports are created and connected to the devices, you can start creating Fabric Maps. See Enable Gigamon Discovery on Chassis for more information on GDP.

Notes on Circuit Ports and Circuit GigaStream in Fabric Maps

As mentioned in the pre-requisites section, Circuit ports and Circuit GigaStreams are required for creating Fabric Maps. Refer to the following notes:

- You cannot specify the Circuit Ports and Circuit GigaStreams that are to be used in a fabric map. You can only define the Circuit Ports and Circuit GigaStreams.
- Circuit Ports and the Circuit GigaStream for the fabric path are selected as follows:
  - Circuit GigaStreams at the first clusters are selected randomly for efficient load balancing.
  - For hop beyond second hops, previously used Circuit Gigastreams may be selected if 'Resource Sharing' is set to enabled mode. This allows to reuse the circuit tunnels previously created. Refer to Configure Fabric Map Settings.
- For Circuit GigaStreams on the Circuit links that connect two clusters, it is recommended to not have more than one ingress GigaStream in any of the devices of the cluster.

- Do not configure redundant Circuit GigaStream between non leaf-spine clusters as this may lead to various issues such as traffic duplication, source port configuration and other such issues.

Consider the following leaf-spine topology. Fabric map is created between the two leaf nodes and the spine node. The second Circuit GigaStream in the leaf node may be used when the circuit ports are of different speeds.

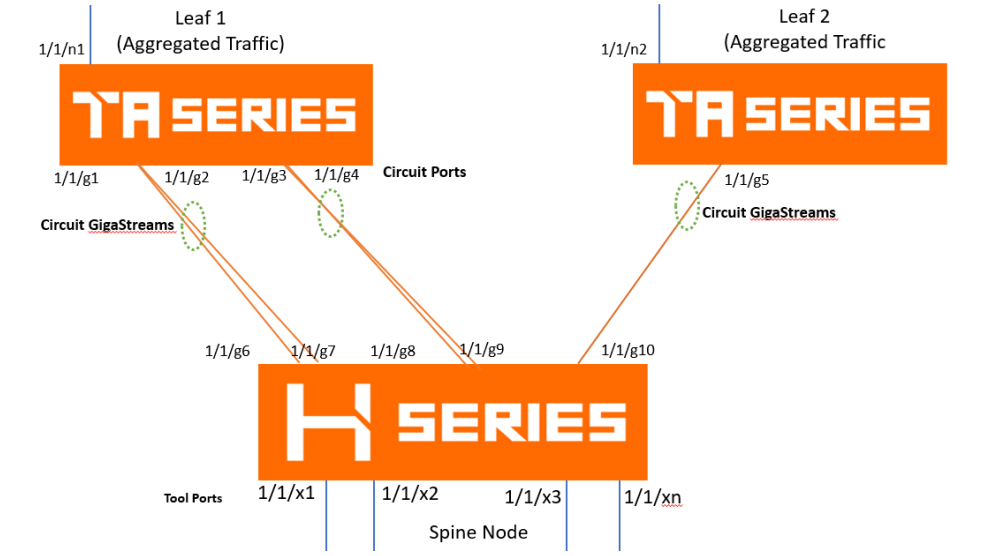


Figure 55 *Circuit GigaStreams in Fabric Maps*

Notes for Circuit ID

Keep in mind the following points while configuring the Circuit ID:

- Circuit ID is used internally to pass traffic from one cluster to another. You can configure Circuit ID ranges from 2 to 4000, or set your own custom range, with each topology able to reuse ID ranges (e.g., 2-513) across multiple topologies.
- Circuit ID allocation is managed globally, beginning from the lower limit of the defined range. This ensures that IDs are allocated efficiently across different topologies. It is allocated automatically.
- To enable the Circuit ID optimal allocation feature in your currently deployed maps, it is necessary to delete the existing maps. After deletion, create a new map that includes the optimal allocation feature, as it cannot be applied to the existing maps directly.
- Each cluster can use up to 512 Circuit IDs from the global pool (default range is 2–4000), allowing for 512 unique destination sets per cluster or standalone node. Multiple fabric maps can also be created, each with 512 unique destination sets.
- Circuit IDs are allocated on a per-hop and per-direction basis.
- The Circuit ID optimal allocation is not applicable for second-level maps with Fabric Port Group destination.

- The circuit tunnel is not shared for first-level maps even in Shared mode.

Create Links between Clusters

You can either connect the devices manually or you can have them automatically discovered by Gigamon Discovery Protocol (GDP). If a user's ports are already physically connected, then the link between those ports will be displayed if GDP is enabled at both ends and the physical links are up. GDP must be enabled in both the source and destination ports to have the devices displayed in the topology map.

- Create Manual Links
- Create Gigamon Discovery Protocol (GDP) Based Links

Create Manual Links

To create links manually, do the following.

1. On the left navigation pane, go to **Inventory** select **Physical >Topology**.
2. Select **Add Link(s)** option from the **Add** tab drop-down menu on the Topology menu bar.
3. Select the **Source** device and the circuit port which you already created.
4. Select the **Destination** device and the circuit port which you already created.
5. Click **Submit** to connect the devices.

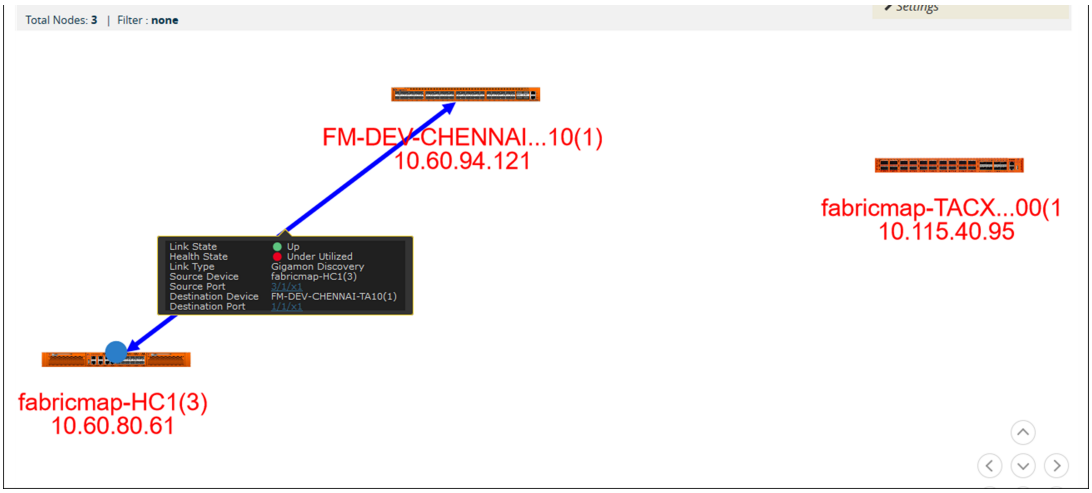


Figure 56 Connecting Devices

6. Click the Topology link from left-navigation menu to view the connection.

After you have successfully connected your devices, the topology map is displayed with the connections. If you click on the connected devices link, you can view the connection details of the ports you created.

Create Gigamon Discovery Protocol (GDP) Based Links

To create Gigamon Discovery Protocol (GDP) based links:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**..
2. Select the node (Source or Destination).
3. Select **Ports**. On the Ports screen, select the circuit port.
4. Click **Edit** from the top menu bar. The Port screen is displayed for the device you selected.

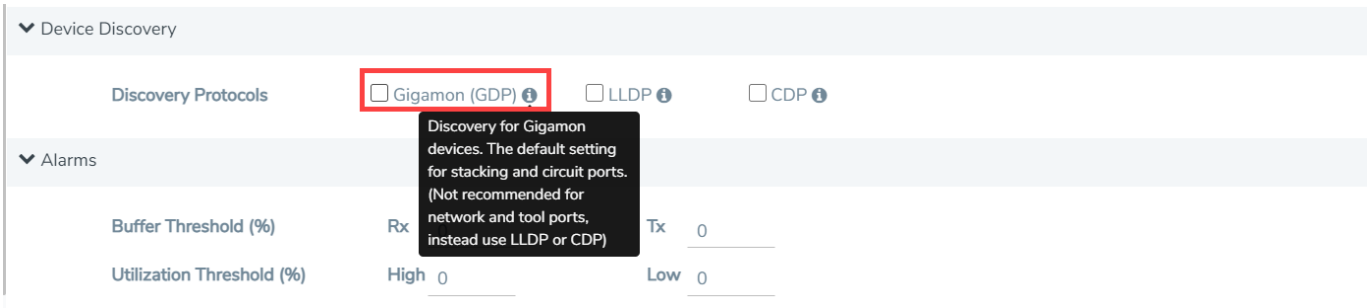
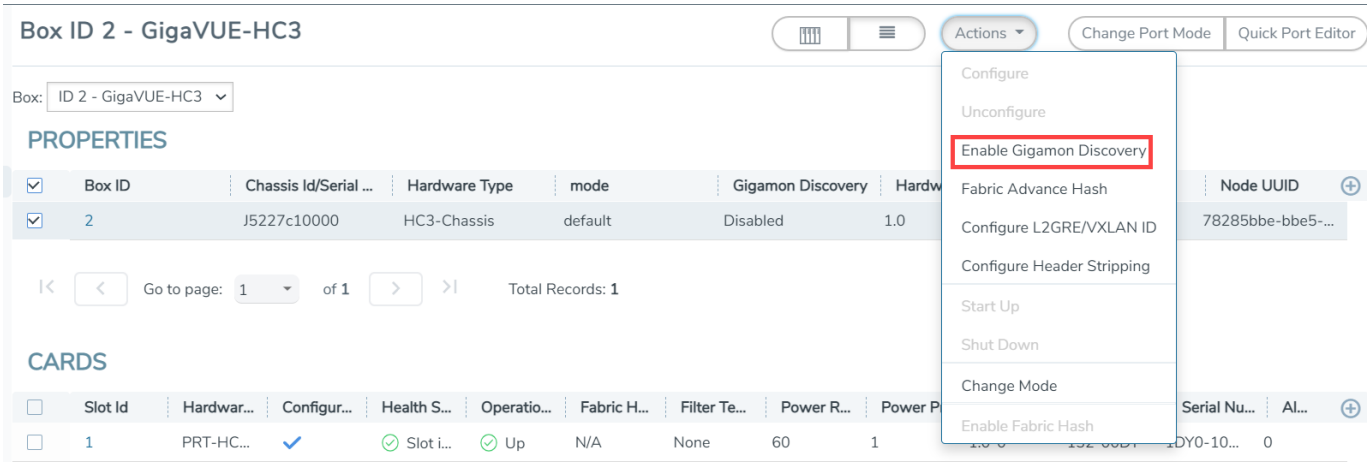


Figure 57 Gigamon Discovery Protocol

5. Enable Gigamon GDP. Click **OK** to submit configuration.
6. You must then enable GDP on the device chassis as follows:
   a. Select > **Chassis**. The chassis screen is displayed showing the Gigamon hardware.
   b. Click the **List View** icon located at the top of the page. Details about the chassis is displayed in a list view format.
   c. Select **Enable Gigamon Discovery** from the **Actions** menu at the top of the screen.



NOTE: This action must be done for both the source and the destination devices.

This completes the prerequisites steps for creation of Fabric Maps. Refer to section Create Fabric Maps for details on creating Fabric Maps.

Create Fabric Maps

To create Fabric Maps:

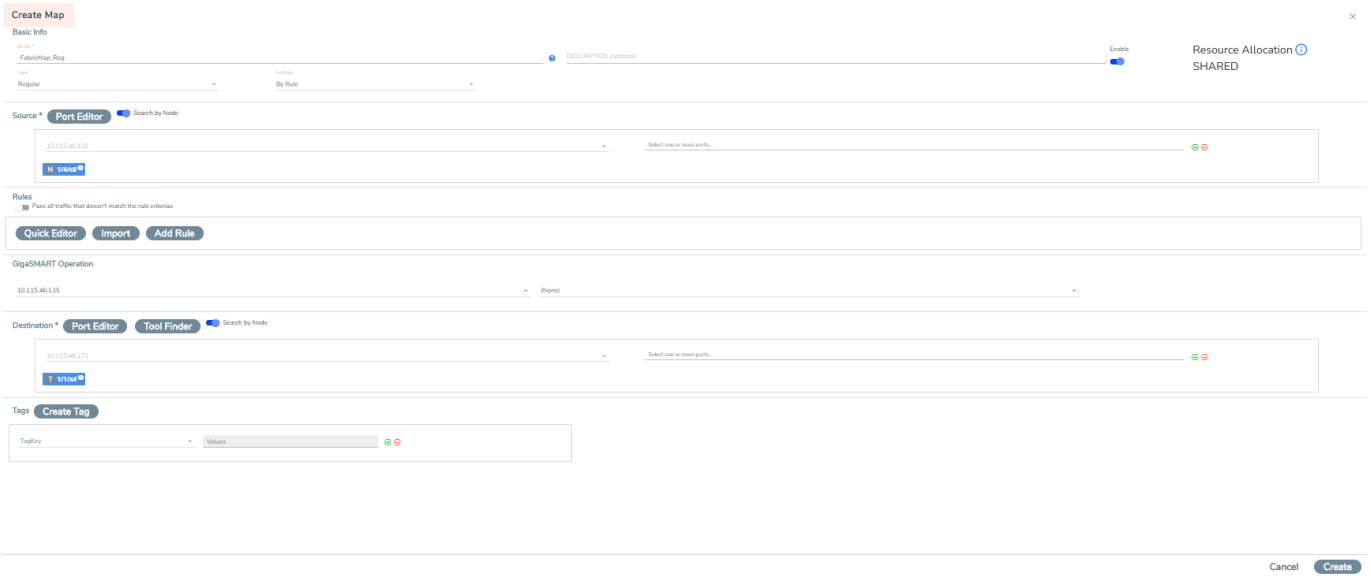
1. Go to **Traffic** and under **Physical > Orchestrated Flows> Fabric Maps,** click **Create**.

   The following table describes the parameters displayed in the fabric map list view.

Field	Description
Alias	Alias name of the fabric map. **Note**: GigaVUE-FM does not allow you to configure fabric maps alias to start with *_FM_*, *_FM-*, *FMAuto-*. These prefixes are used in auto-generated fabric map alias.
Enabled	Indicates if fabric map is enabled – Yes/No
Source	Number of source ports associated with fabric map
Destination	Number of destination ports associated with fabric map
Rules	Rules for fabric map
GigaSMART Alias	Alias name of the GigaSMART operation

Field	Description
Type	Port type
Status	Health status of the fabric map

1. Click **Create**.The create map screen is displayed.



2. Enter map parameters as described in the following table:

Field	Description
Basic Info	
Alias	Alias name of the fabric map
Description	Fabric map description
Enable	Toggle option that lets you toggle between: ▪ **Enable**: Traffic passes through the fabric map that is created. ▪ **Disable:** Traffic is temporarily prevented from passing through fabric map if 'Disable' is selected. **Note**: You cannot disable the second-level maps (option is disabled in the UI). Alternatively, you must disable the corresponding first-level map.
Resource Allocation	Displays whether it uses shared or non-shared circuit id resources.
Type	Map type. Options are: ▪ Regular ▪ First level

Field	Description
	▪ Second level
Subtype	Map subtype. Options are: ▪ By Rule ▪ Pass All ▪ Collector
Source	Source port. Use the **Search by Node** toggle bar to toggle between the following two options ▪ Enter the port alias/port id (Respective Cluster Id is displayed for selecting the required ports), or ▪ Specify the source node and source ports. Use the +/- icon to add new nodes or remove source ports. **Note**: You cannot use ports already used as source ports in the orchestrated configuration.
Rules	Enable Non-Matching Rules for map type 'Regular'. Use the toggle bar to either Pass all traffic or not. The following options are available: ▪ **Quick Editor**: To select the port number or numbers to add for a pass or drop rule or both. ▪ **Import:** To import a map template ▪ **Add a Rule**: To add a description and select one or more conditions. Use this option to add rules one at a time. Click the +/- icon to add or delete rules and conditions. ▪ You can also associate tags to the map rules. For detailed information about map rules, refer to Map Rules. You can configure Inner Header qualifiers and MPLS Header qualifiers for GIgaVUE-TA400 device. Refer to Inner Header and MPLS Header Filtering **Note**: When you associate a tag to a rule in Fabric map, GigaVUE-FM associates the same tag to the internally created flow maps. The visualizations displayed in the Map Rule Statistics dashboard is based on the tags associated with the internal flow maps . Refer to the Add Tags to Map Rules section for more details.
GigaSMART Operation	Select a GigaSMART node and operation.
Destination	Destination port. You can either: ▪ Enter the port alias/port id, or ▪ Specify the destination node and destination ports. Use the +/- icon to add new nodes or remove destination ports. Use the toggle bar to toggle between these two options. Click the Tool Finder option to search for the available destination ports.

Field	Description
	• You can also use Fabric Port Group as destination for the second level maps to allow load-balanced traffic. You must enter the name of the FPG created. (For details, refer to Fabric Port Group section). **Note**: Starting in software version 5.7, you can also select the **Null Port** option for the second level maps if you do not want to send the traffic to any physical port. Refer to the *"Create Application Filtering Intelligence by Selecting Applications from Dashboard"* section for the use of null port in application intelligence solution.
Tags	Select the tag key and the tag value to which the Fabric Maps must be associated to. **Note**: The tag key and the associated tag values must be created in advance in GigaVUE-FM. Refer to the "Tags" and "Role Based Access Control" sections in the GigaVUE Administration Guide for more details You can only view tags that are permitted for your role. Refer to the "Tags" section in the GigaVUE Administration Guide for more details.

3. After you enter all the fabric map parameters, click **Create**. The new fabric map is added to the list view.

To create a copy of the existing Fabric Map:

1. Select the required fabric map.
2. Click the copy icon on the top.

A copy of the selected map is created.

NOTE: In GigaVUE-OS nodes, when using an API script, you may experience a discrepancy in the GigaVUE-FM GET API response for a brief period. This can occur when the script is updating some device components via the GigaVUE-FM API while configuration synchronization is also in progress. In such cases, the configuration synchronization may overwrite the latest data (updated during the GigaVUE-FM API call) with older data (synced during the configuration refresh) in the database. GigaVUE-FM will automatically update to reflect the latest data after subsequent configuration synchronization.

Edit and Delete Fabric Maps

You can edit and delete the flow maps . Refer to the following section for details:

- Edit Fabric Maps
- Delete Fabric Maps

Edit Fabric Maps

To edit Fabric Maps:

1. Select a fabric map. Each fabric map has a summary page.

- The top half of the summary page displays the fabric map component interactions and traffic flow.
- The bottom half of the screen displays the following tabs:
  - **Summary**: Summary of ports and maps
  - **Rules**: List of rules associated with the fabric map
  - **Statistics**: Fabric Map statistics.

2. Click on any component in the fabric map to display more details about the component.

NOTE: If the *Show Auto Generated* maps option is enabled, then the Fabric Maps page lists the auto generated Fabric Maps that have the prefix as *FmAuto-*, *_FM_*, *_FM-* .

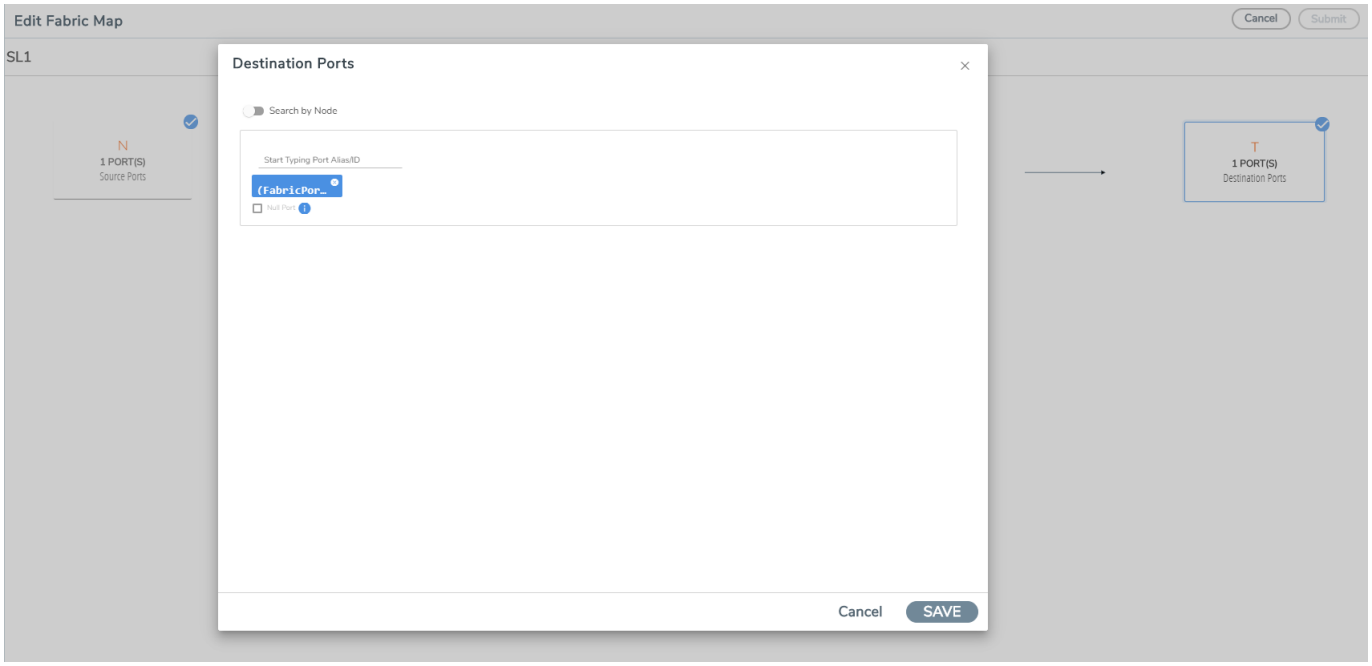
Edit Fabric Map Component

1. Click the **FabricMapcomponent** icon. The fabric map component summary screen appears.
2. Click the option menu and select **edit**.
3. Edit fabric map parameters.
4. Click **Save** to update the fabric map components.

NOTE:

- You cannot edit or delete a map if another person is working on the same map. The error message is available in the audit log page. For information on the audit log page, refer to Audit Logs.
- You can edit Fabric Maps with Fabric Port Group configured as destination in your second level maps. You can replace the FPG with another FPG or replace it with tool ports, as required. Refer to the following screenshot.

- In a cluster, the same source ports are often used across multiple fabric maps. When you need to delete the source ports in one fabric map within the cluster, you must delete the corresponding source ports from all related fabric maps to avoid any overlapping of sources. For example, if a cluster has two source ports—1/1/x1 and 1/1/x2—used across several fabric maps, and you decide to remove source port 1/1/x1 from one fabric map, GigaVUE-FM will also remove source port 1/1/x2 from that fabric map. This automatic removal helps to prevent source overlapping within the cluster.



Delete Fabric Maps

1. Select the fabric map you want to delete from the Fabric Maps summary page.
2. Click the **Trash Can** icon.

   A confirmation message appears to confirm whether to delete the selected Fabric Maps.
3. In the message, do any one of the following as per your requirement:
   - Type **DELETE** if you have selected multiple maps, and click **OK**.
   - Type **DELETE ALL** if you have selected all the maps, and click **OK**.

You can view the error message in the audit logs description field for all the failed attempts. For more information on the audit log, refer toDashboard

NOTE: You cannot select and delete the auto-generated Fabric Maps, and you cannot select the auto-generated fabric maps using the select-all option.

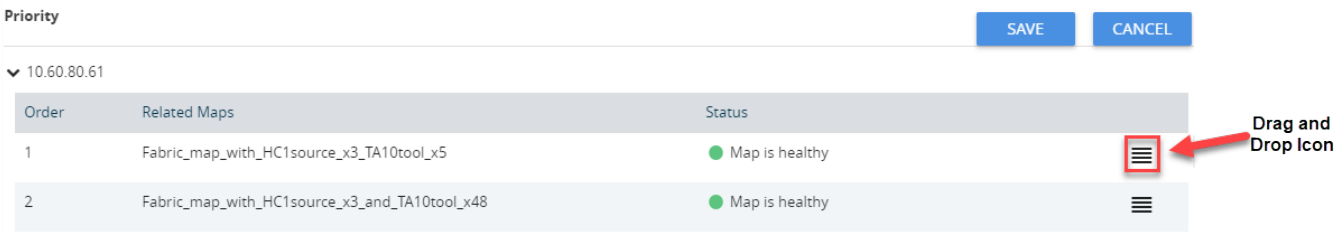
After the fabric map is deleted you will see an acknowledgment on the screen. Refresh the fabric map summary page to confirm the deletion.

Prioritize Fabric Maps

You can prioritize the traffic flowing through the flow maps. The priority of the fabric map is implemented through cluster level map priorities and can be adjusted at the cluster level.

To prioritize the traffic flow through flow maps:

1. On the left navigation pane, go to **Inventory** and select **Physical** > **Fabric Maps**.
2. Select a fabric map and click **View Details** either from the **Task** drop-down menu or from the fabric map selected.
3. Scroll down to view the current priority setting of the flow maps.
4. Click **Change FabricMap Priority**.
5. Use the drag and drop icon to change the order of priority.



6. Click **Save** to save the changed priority.

NOTE: The priority list includes both flow maps and cluster maps, which are grouped and prioritized based on the cluster ID.

Configure Fabric Map Settings

Share Circuit ID Resource

You can share the circuit id resources for the Fabric Maps created for the same destination. To do this:

1. Select the **Settings** icon on the top navigation bar to open the Share Resources dialog.
2. Enter the **Circuit Id Range** (Min and Max values). Use the +/- icon to increment/decrement the circuit id range.
3. Click **Save** to save the settings.

The Share Resources option is **enabled** by default. Use the toggle bar to disable this option.

When you change the option from shared to non-shared, a note message appears. Click **OK** to proceed and type **Yes** in the confirmation message to change from Shared to Non-Shared. You can use this option only if you who have Read and Write privileges in both the Traffic Control Management and GigaVUE-FM Security Management.

Spine Source Map Mode

You can configure this global setting in FM that controls how cluster-local maps are created in a leaf-and-spine cluster when the source is at the spine (i.e., a decap circuit GigaStream in a leaf-spine cluster). This setting is applicable for creation of the first map that uses the spine mode. Refer to Supported Topologies topic for more information on topology of Spine Source Map mode.

The following image shows the Spine Source Map Mode in the SETTINGS dialog.



When a leaf-and-spine cluster acts as a middle or destination cluster in a fabric map, its cluster-level map source can be one or more circuit GigaStreams.

If these circuit GigaStreams are on the cluster's spine nodes, you can create the map in the leaf-and-spine cluster in one of two ways:

- **COMBINE** - A single map that combines the circuit GigaStreams from all spine nodes
- **SEPARATE** - Multiple maps, with one map per spine node's circuit GigaStream. This is the default mode.

NOTE: Use SEPARATE mode when the map destination is on nodes that are connected to all spine nodes.
Use COMBINE mode to optimize fabric map deployments so that the cluster can use a single logical map to forward the traffic to the destination, even when traffic arrives via multiple decap circuit GigaStream on the spine nodes.

Limitations of Spine Source Map Mode

- If existing maps are in SEPARATE mode, the future maps from the same circuit GigaStream continue to be SEPARATE even if the global mode is switched to COMBINE.
- If existing maps are in COMBINE mode, the future maps from the same circuit GigaStream to be COMBINE even if the global mode is switched to SEPARATE.
- To truly change mode for the circuit GigaStream, you must delete the existing maps and recreate them after changing the global setting.

Fabric Maps Statistics

GigaVUE-FM provides the ability to view detailed information about flow maps configured between the devices including packets received and transmitted by the following ports:

- Network ports
- Tool ports
- Hybrid ports
- Circuit ports
- Inline-network ports
- Inline-tool ports
- GigaSMART engine ports

Display Fabric Map Statistics

Use the fabric map statistics page to view the statistics associated with the port types. If there are packet drops in the ports, you can use the statistics page to investigate the cause of the packet drops. Fabric map statistical data is generated from the rules you specify when you create a fabric map.

Display Fabric Map Details

To display the fabric map details:

1. On the left navigation pane, go to **Inventory** and select **Physical> Fabric Maps.**
2. Click on a **Fabric Map** from the main page. The fabric map list view is displayed.
3. Click the **View Statistics** from the options menu. The statistics screen is displayed. Using this screen, you can view fabric map traffic data rates intervals.

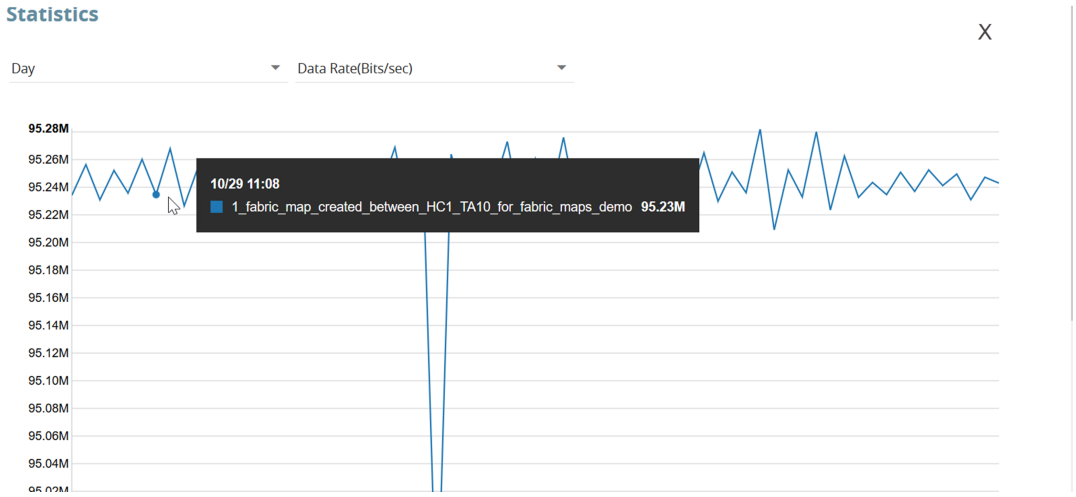


Figure 58 Fabric Maps Statistics

NOTE: Statistics and data rates can be displayed per hour, day, week, or month.

4. Click the current statistics label, (Hour, Day, Week or Month), to display the options menu where you can select a different statistics view.

Filter Fabric Maps List View

Using the filter function, you can search and narrow down the flow maps you want to be displayed on the fabric map list view page.

1. To use the filtering functionality, click the **Filter** icon. The Filter quick view dialog is displayed.
2. Enter the parameters to specify the data you want to display:

Fabric map filter parameter options:

Criteria	Description
Alias	Alias name of the Fabric Map
Source	Source port(s) of the Fabric Map
Destination	Destination port(s) of the Fabric Map
Status	Displays health of the fabric map. Options are: Healthy, Unhealthy, and Warning
Show Auto Generated Fabric Maps	Toggle bar that allows you to show/hide the auto-generated flow maps when flow maps page is loaded.

Fabric Port Group

GigaSMART Load Balancing functionality distributes outgoing traffic to multiple tool ports or tool port groups within a cluster (refer to GigaSMART Load Balancing section for detailed information). However, it is not possible to load balance the traffic if the tool ports and the GigaSMART operations are configured across clusters, as in the case of Fabric Maps.

The Fabric Port Group (FPG) feature extends the GigaSMART load balancing functionality to fabric maps using circuit ids and circuit tunnels, thus allowing load balancing between tool ports that reside on different clusters. However, you can also use FPG with in a cluster.

Refer to the following sections for more details:

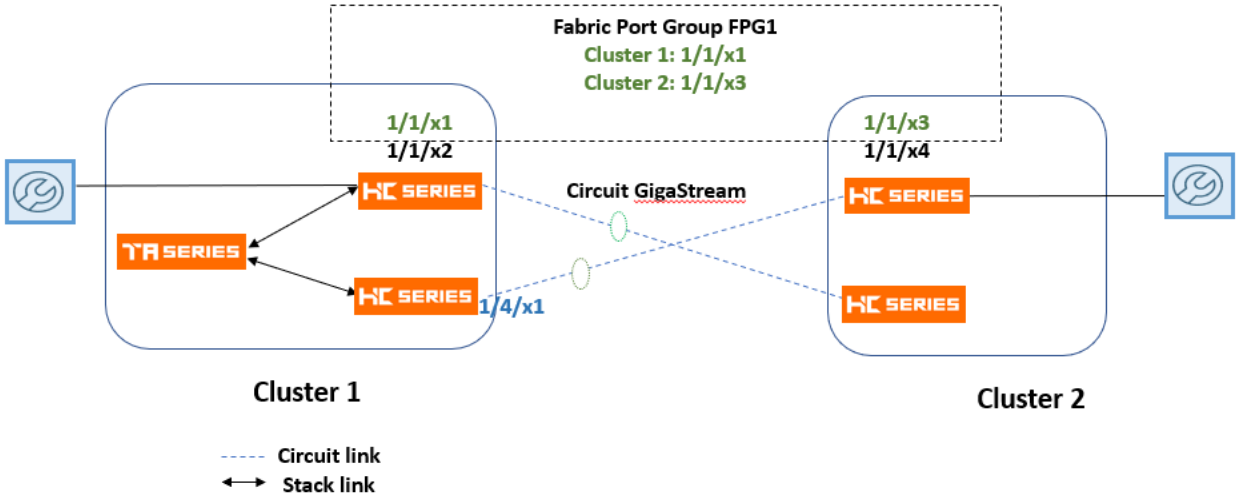
- Fabric Port Group
- Rules and Notes
- Create Fabric Port Group
- Configure Fabric Map Using Fabric Port Group

How Fabric Port Groups Work?

A Fabric Port Group may be defined as a logical collection of tool ports that reside on different clusters. Each of the tool port in the Fabric Port Group receives load-balanced traffic.

Consider a GigaVUE-FM instance with the following clusters configured with tool ports:

- Cluster 1: Configured with Port 1 (1/1/x1), Port 2 (1/1/x2)
- Cluster 2: Configured with Port 3 (1/1/x3), Port 4 (1/1/x4)



FM1: First Level: Source: 1/4/x1 Destination: vPort of Cluster 1
FM2: Second Level: Source: vPort of Cluster 1 , Destination: FPG

Create a Fabric Port Group FPG1 using port 1/1/x1 of cluster 1 and port 1/1/x3 of cluster 2. Define the port weight for each of the ports. For the Fabric Port Group that is created, GigaVUE-FM:

- Creates cluster specific GigaSMART port groups with unique global circuit IDs. These circuit IDs route the traffic tagged by the source vport to the GigaSMART port groups of the destination cluster.
- Creates circuit tunnels with matching circuit ids along the fabric path. Based on the circuit id rules, the maps generated at the cluster level route the traffic to the next hop towards the final cluster level GigaSMART port groups.

NOTE: If any of the ports in the fabric port group goes down, the traffic from the port is redistributed to other healthy ports in the fabric port group.

Rules and Notes

Keep in mind the following rules and notes when you work with the Fabric Port Group feature:

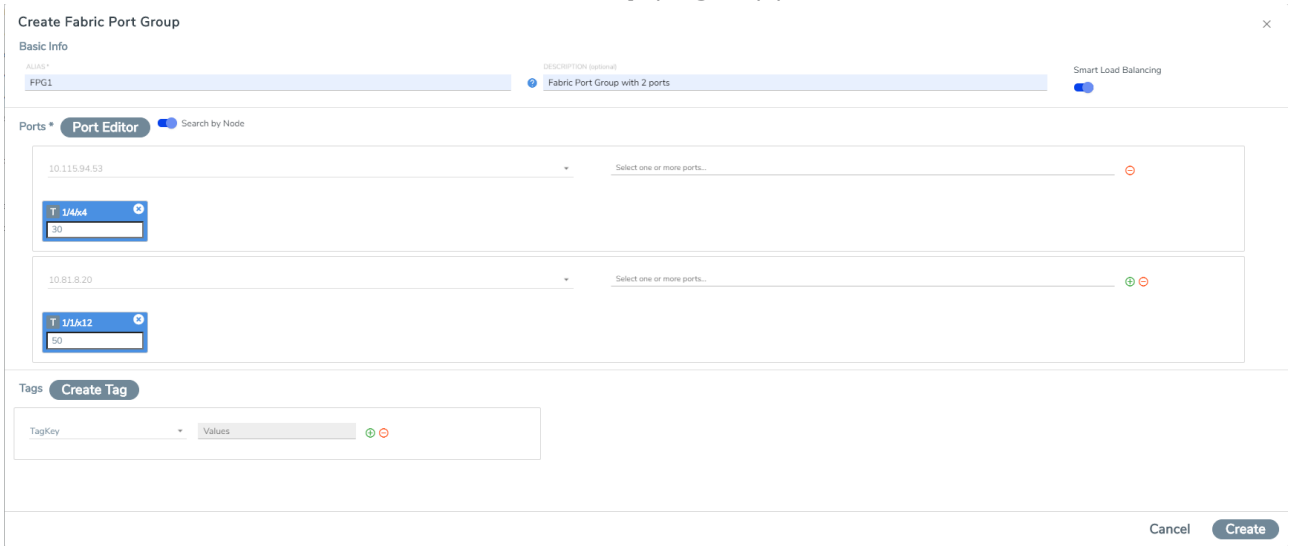
- You can create a FPG using ports of types: Hybrid or Tool. Ports in a single FPG must be of the same port type.

- You can use FPG as destination ports in second-level maps that use load balanced GigaSMART operation in Fabric Maps. You cannot use FPG in the first-level, RegularbyRule, PassAll, or the collector maps.
- You cannot delete or edit an FPG that is deployed in a fabric map. You must first delete the fabric map that deploys the FPG and then delete/edit the FPG.
- You cannot use the individual tool ports of an FPG in another FPG.
- You cannot change the port type and mode of a tool port that is used in the FPG.
- You must not use the tool ports of a FPG in device or cluster-level configurations such as GigaStream, Port Pair, Cluster-level port group (though GigaVUE-FM does not restrict this).
- You cannot use two different load balancing GigaSMART operations for the same FPG.

Create Fabric Port Group

To create a Fabric Port Group:

1. Go to **Traffic** fromPhysicalselect**Orchestrated Flows >Fabric Maps> Fabric Port Group.**
2. Click **Create**. The **Create Fabric Port Group** page appears.



.

3. In the **Basic Info** section, enter the name and description for the Fabric Port Group.
4. In the ports section:
   - Enter the port alias/port id, or
   - Specify the node and ports. Use the +/- icon to add new nodes or remove the ports

5. Define the weight for each of the ports used in FPG. The weight of the individual ports must be less than 100. The combined value of the ports can be greater than 100, as the actual load balancing ratio is computed with individual values divided by the combined value.

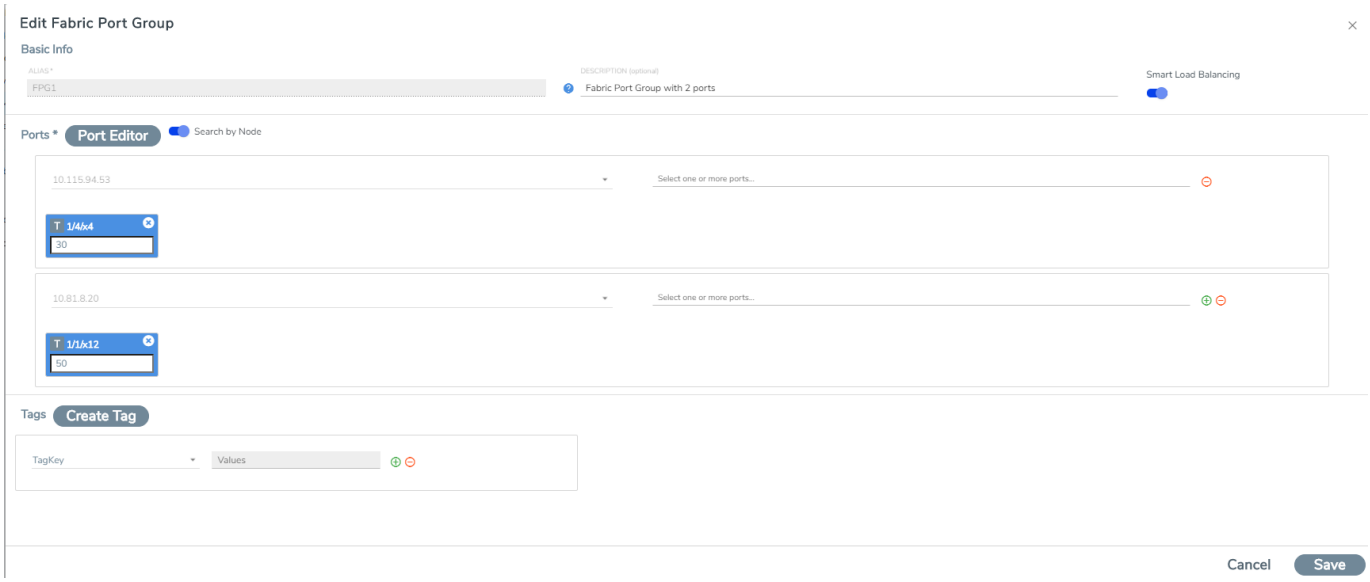
   NOTE: Smart Load Balancing option is always enabled, and you cannot disable this option. You can use Fabric Port Group in Fabric Maps only if Smart Load Balancing is enabled.

6. Select the required tag key and tag value combination to associate the FPG to tags.

   NOTE: The tag key and the associated tag values must be created in advance in GigaVUE- FM. Refer to the "Tags" and "Role Based Access Control" sections in the GigaVUE Administration Guide for more details

7. Click **Create**. The fabric port group is created and is added to the list view.

To edit an FPG, select the FPG that you want to edit from the list view and click the edit icon. Make the required edits to the FPG and click **Save**.



Configure Fabric Map Using Fabric Port Group

To configure the Fabric Port Group in Fabric Maps perform the following steps:

Task	Description	UI Steps
1	Create a Fabric Port Group with alias FPG1.	Refer to Create Fabric Port Group
2	Create a first level fabric map.	Refer to Create Fabric Maps for the steps. **Map Alias**: FM_Firstlevel

Task	Description	UI Steps
		Map Type: First Level **Source Port(s)**: Select a single port or multiple source ports from different clusters in a given topology. **Destination Port(s):** Select virtual port 'vport1' for the destination. **Rule**: Add the required rules.
3	Create a second level fabric map.	**Map Alias**: FM_Secondlevel **Map Type:** Second Level **Source Port(s)**: Select the virtual port 'vport1' (defined as destination in Task 2). **Destination Port(s):** To select the FPG1 created in Task 1: Enter the first three characters of the FPG created to view the FPG for selection. **Rule**: Add the required rules.
4	Click **Create**.	First-level and second-level fabric maps are created, with the FPG1 deployed in the second-level fabric map.

Troubleshooting

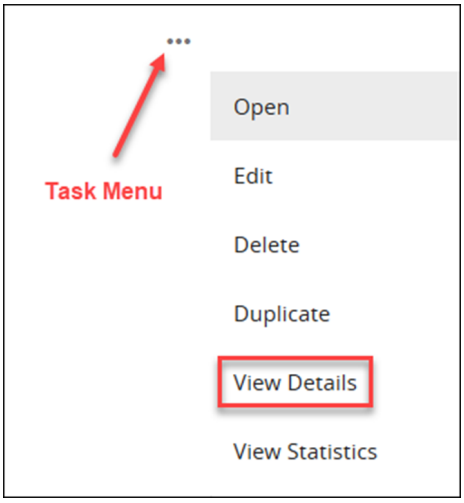
You may face situations when fabric fails because of a configuration error or you need to investigate an unhealthy GigaStream or port associated with fabric map. Fabric map maintains the configuration status and health state.

The configuration statuses are SUCCESS, FAILED, and PARTIAL_SUCCESS. You can investigate issues related to flow maps and once issues are corrected, you can run resubmit to push changes. Health state is consolidated from generated components healthStates.

Fabric Maps health states consist of the cluster level maps, (ports), and circuit tunnels (ports and links). After configuration and health state issues are corrected you can sync the Fabric configStatus and healthStates.

To use the troubleshooting feature,

1. On the left navigation pane, go to **Inventory** under **Physical > Fabric Maps**.
2. Select a **Fabric Map** to troubleshoot.
3. From the **Task** drop-down menu, Select **View Details.**



The Details screen displays fabric map component details.

4. Select the **More Info** link to display the fabric map details.
5. Select the **Troubleshoot** link to display troubleshooting options.

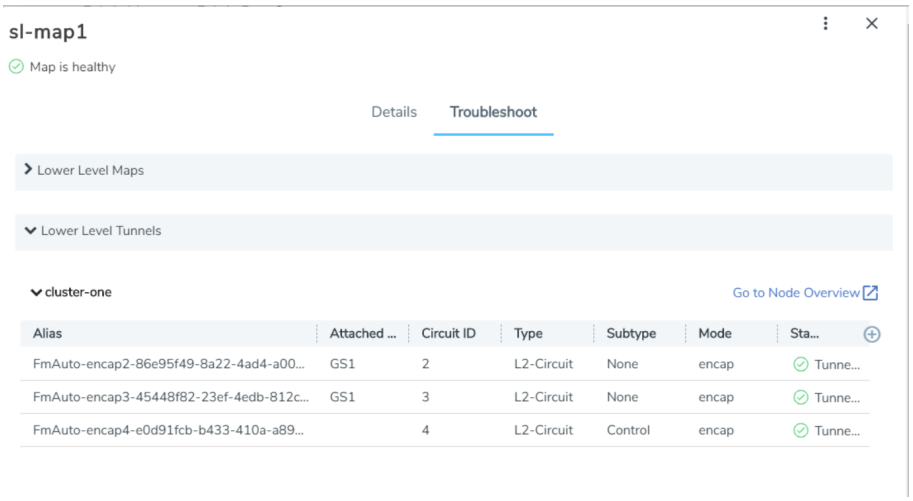


Figure 59 Troubleshooting

6. Select the **> arrow** next to the fabric map to display the map details.
   Use this screen to investigate configuration issues related to your fabric map.

Fabric Map Configuration Status

The following table helps to troubleshoot a failed or partially successful fabric map configuration.

Configuration Status	Definition	Suggested Action
SUCCESS	Fabric map was accepted at GigaVUE-FM level and traffic flow deployed correctly.	• No action required
FAILED	Fabric map was not configured correctly and not accepted or deployed at GigaVUE-FM level.	• Go to the topology page and ensure there is a link between the cluster. • Make sure the circuit GigaStreams are successfully created. • If the clusters are no longer connected, connect the clusters and reapply the configuration. • If the circuit GigaStreams are configured correctly, make the necessary changes and reapply the configuration.
PENDING	• This means no fabric path exists between 2 nodes.or • A failed operation on the device and further deployment of cluster level maps are not allowed in that device.	• Create a fabric path and reapply fabric map.
PARTIAL_ SUCCESS	This means a portion of the fabric map was accepted at GigaVUE-FM level, but devices cannot provision the generated components.	• For details, see How to Troubleshoot Partial Success Errors

Fabric Map Health State

The following table helps to troubleshoot the health state of your fabric map.

Health State	Definition	Suggested Action
Healthy	All fabric map components (nodes, ports, GigaStreams) are all healthy.	• No action required.
Unhealthy	An unhealthy state indicates that a port is down or packets dropped.	• Make sure no ports are down and no packets dropped. • Correct any configuration connection issues, and re-sync the fabric map.

How to Troubleshoot Partial Success Errors

If your fabric map displays a red status light on the list view page, this means a portion of the fabric map was accepted at GigaVUE-FM level, but devices cannot provision the generated components.

The error message associated with fabric map appears when you hover over the fabric map status column. You can use the error message to help you troubleshoot and identify the components that are in conflict or mis-configured.

Example 1:

In the following example, the status error message indicates that the ports are mis-configured. You can use the following workflow as a best practice for investigating and correcting map configuration problems.

1. Hover over the fabric map status column displaying a red status light to view the error message. The error message provides information on the possible issues present with fabric map components.



Figure 60 Troubleshooting

This error indicates that the input network ports overlapped with the existing map, and therefore it is required to correct this issue and remove the overlapping ports. Use this screen to investigate the exact ports causing the problem.

2. Click **View Details** from the options menu. The fabric map detail screen is displayed.

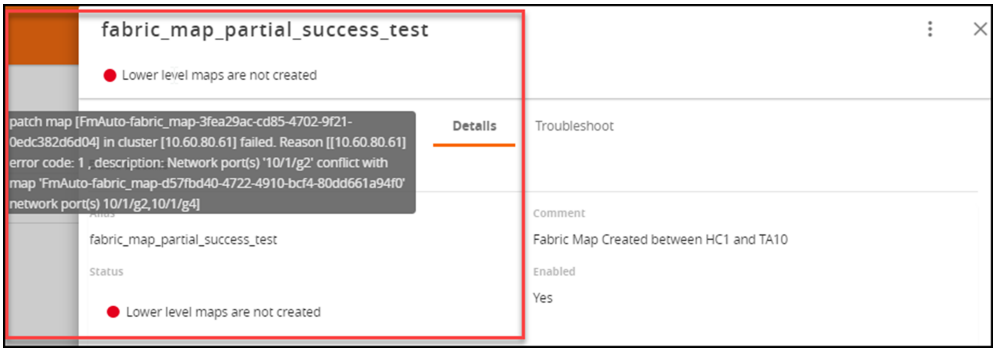


Figure 61 Fabric Maps Details

3. Click the **Troubleshooting** link. Scroll down the screen to investigate the Source and Destination ports configurations.

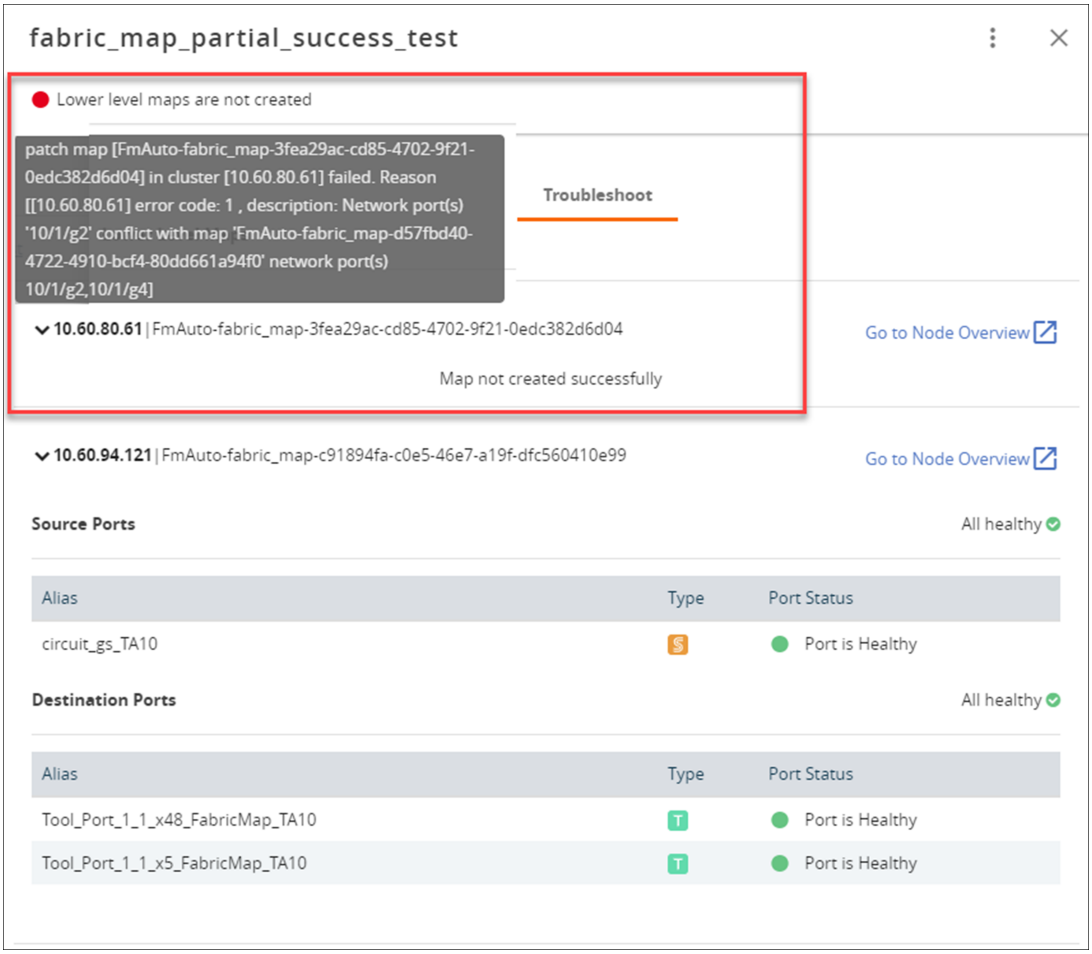


Figure 62

The network ports that you tried to add to your fabric map were not available and contributed to the error in the fabric map configuration. Next, you need to edit the network ports.

4. Click **Edit** from the options menu to edit the fabric map.

Figure 63 Editing Fabric Maps

5. Select the **Source Ports** component.

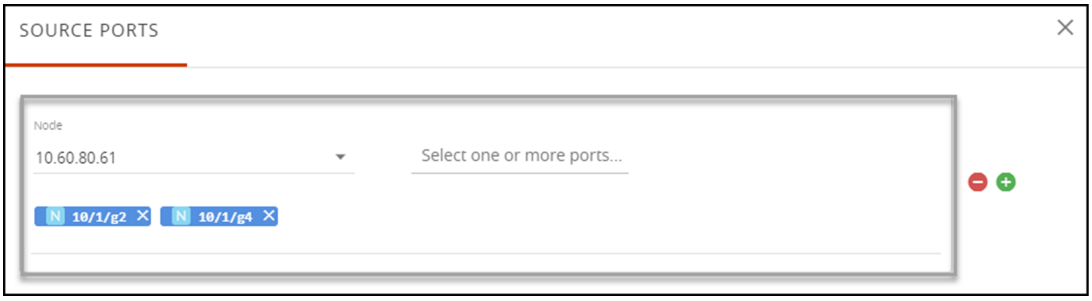


Figure 64 Source Port Component

6. The Source Ports screen appear.
7. Edit the source ports with correct information and save your configuration.
8. Click **Submit** on the **Edit Fabric Map** page.

After you have successfully updated the fabric map component causing the problem, the fabric map list view displays a healthy status symbol.

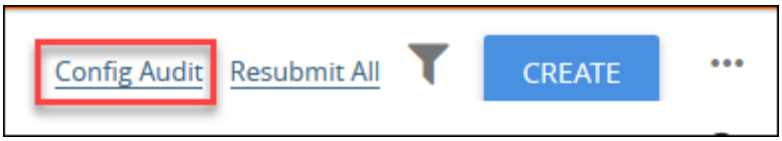
If a problem with the fabric map still exists, repeat the steps above to identify and track down the problem.

Config Audit

Use the Config Audit option to conduct an audit of your fabric map configuration. This process audits the configuration including the components of your fabric map. Config audit verifies that the fabric map configuration is active, the devices are connected, and traffic is flowing through the fabric map with no issues.

To start the configuration audit process:

1. On the left navigation pane, go to **Inventory** under **Physical > Fabric Maps.**
2. Click the **Config Audit** link at the top of the screen.



After the config audit is complete, a notification appears with the audit results.

If the config audit is not successful, check the fabric map configuration error message(s) for help in resolving issues.

Limitations of Fabric Maps

The following are the limitations of Fabric Maps:

- Circuit ports, circuit GigaStreams and manual topology links (or GDP must be enabled on circuit ports to discover the connected clusters) must be configured before creating fabric map.
- In first-level type fabric map, one vport is required for the destination. Fabric Maps can have multiple tool/hybrid ports from multiple clusters in the destinations in addition to the vport, but the map type remains as the first-level map type.
- In multi-hop scenarios, it is not recommended to create a single level fabric map with GSOP. You must create a vport associated with GSOP, and then create the first and second level maps with the vport to perform the GSOP operations.
- Fabric map configurations are saved in GigaVUE-FM and not in the individual devices. Therefore, if you change the GigaVUE-FM instance that manages the devices, then fabric map configurations cannot be rediscovered from the devices. In this scenario, the GigaVUE-FM configuration that has been saved and backed-up from the original GigaVUE-FM instance must be restored to the new GigaVUE-FM instance to continue.
- To avoid FAILED or PENDING state, or no traffic stats found during deployment of flow maps, you must ensure that the associated devices are connected to GigaVUE-FM and also interconnected for traffic paths.
- Fabric Map or a map with an encapsulation circuit-id cannot be configured from a source port if that source port is used by a passall-map. If you try to configure this from the CLI, the following error message is displayed: **Cannot have encap on a map with a map-passall or port-pair on the source ports.**
- To enable the Circuit ID optimal allocation feature in your currently deployed maps, it is necessary to delete the existing maps. After deletion, create a new map that includes the optimal allocation feature, as it cannot be applied to the existing maps directly.
- A single circuit GigaStream can source and decap traffic from up to 1024 Fabric maps. Multiple decap circuits GigaStreams may exist within a cluster.
- Fabric maps do not support ingress VLAN port tagging. Circuit-ID tunnels used for inter-cluster connections in fabric maps do not support ingress ports with VLAN tags; configuring ingress VLAN tagging on these ports may result in unexpected behavior and is not supported.
- Circuit IDs are allocated on a per-hop and per-direction basis.
- The Circuit ID optimal allocation is not applicable for second-level maps with Fabric Port Group destination.
- The circuit tunnel is not shared for first-level maps even in Shared mode.

- When you create a regular fabric map with GSOP, the GSOP can be placed anywhere along the path from the source to the destination cluster or node, including within the source or destination cluster itself. However, it must not be positioned beyond either the source or destination cluster.
- If existing maps are in SEPARATE mode, the future maps from the same circuit GigaStream continue to be SEPARATE even if the global mode is switched to COMBINE.
- If existing maps are in COMBINE mode, the future maps from the same circuit GigaStream to be COMBINE even if the global mode is switched to SEPARATE.
- To truly change mode for the circuit GigaStream, you must delete the existing maps and recreate them after changing the global setting.

Backup and Restore Fabric Maps and Orchestrated Configurations

You must backup the devices and GigaVUE-FM at the same time. Ensure that the devices are not undergoing configuration edits during backup or restore. For more information, refer to the *"Restore Devices and GigaVUE-FM for Traffic Management Solutions"* section in the *"GigaVUE Administration Guide"*.

Orchestrated Configurations

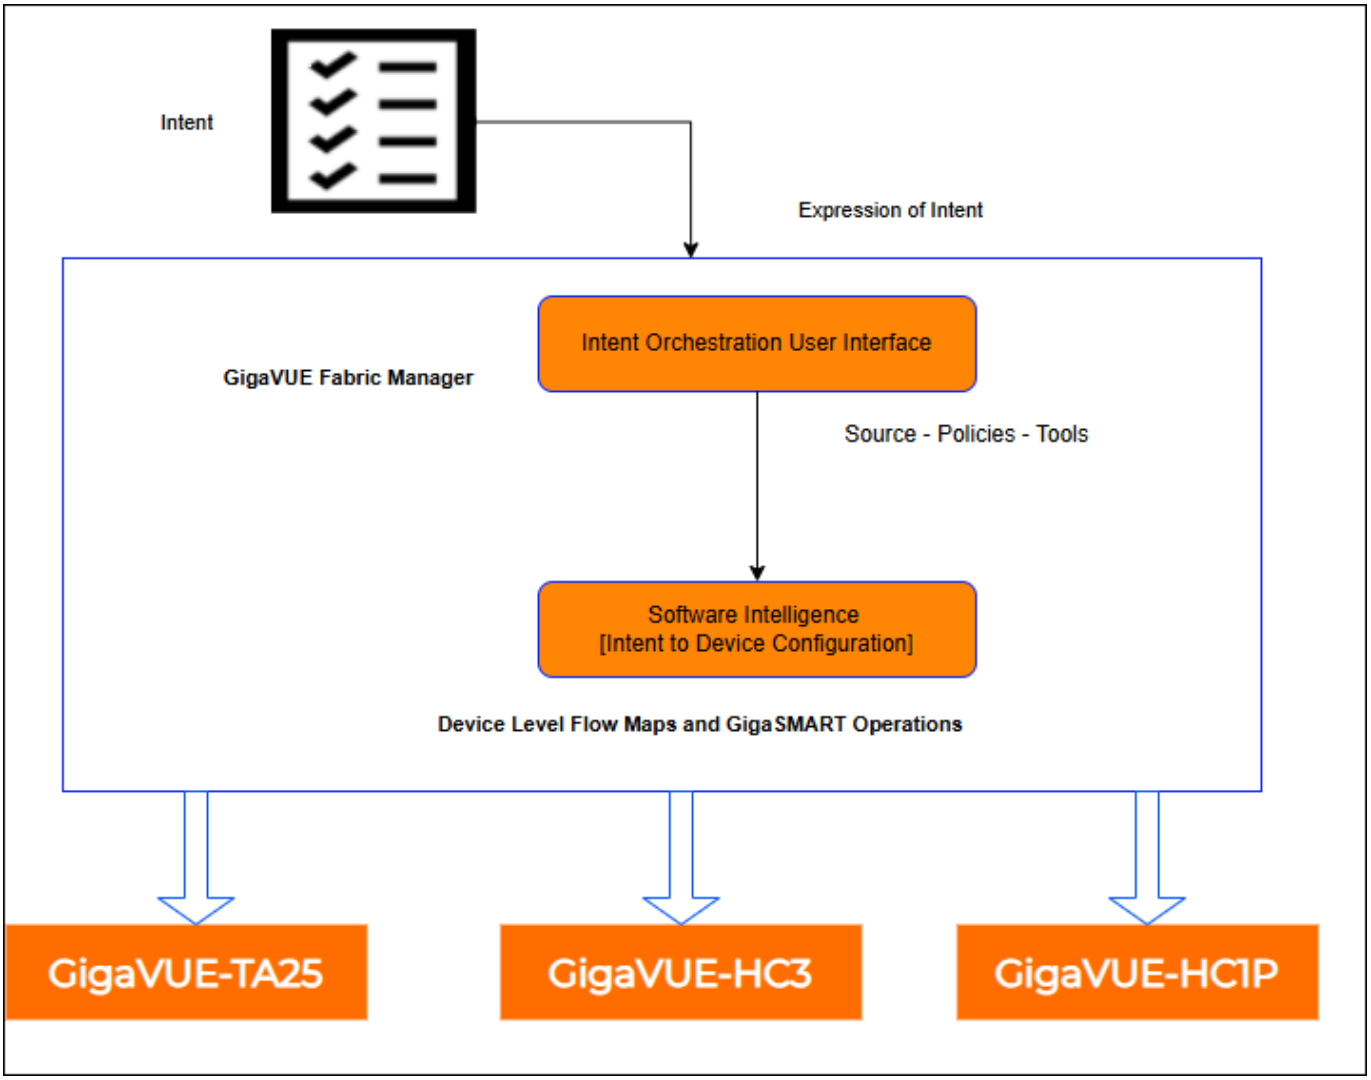
This chapter provides an overview of the configurations that can be automated from GigaVUE-FM and describes how to use the **Orchestrated Flows> Out-of-Band Flows** option in the GigaVUE-FM GUI to define traffic flow through out-of-band visibility fabric.

Featured Content:

- About Intent Based Orchestrated Configurations
- Benefits of Orchestrated Configurations
- Supported Topologies
- Orchestrated Configuration: Examples
- Rules and Notes for Orchestrated Configurations
- How to Create a Policy
- Drop Rules
- Import and Export Orchestrated Policies
- Egress Filters for Additional Filtering Capabilities
- Glossary

About Intent Based Orchestrated Configurations

Intent Based Orchestration (IBO) is a new approach that enables you to automate the configuration tasks in GigaVUE-FM, thereby reducing the complexity and the manual labor involved in those tasks. IBO configurations (or 'Orchestrated Configurations') leverage the intelligence of GigaVUE-FM to detect the intention of the user, and accordingly translates the intentions into capabilities for configuring end-to-end flow maps.



With Orchestrated Configurations, you can:

- Create ***policies***, which are user-defined prescriptions for what to do with the traffic. The policies are translated internally into device-level maps that have the information required for the flow of traffic.

- Use filters and advanced mapping rules in policies, such as the following:
  - L2-L4 filters
  - Advanced filters such as GTP, ERSPAN, VXLAN, VNTAG, and MPLS
  - **NOTE:** Refer to the GigaVUE-FM User's Guide and GigaVUE-OS CLI Reference Guide for detailed information about the map rules and filters.
- Use packet transformation options along with the filters mentioned above to deliver the traffic to the desired tools. Supported packet transformation options include:
  - De-duplication
  - Masking
  - Slicing
  - Header Stripping
  - Header Addition
  - **NOTE:** Tunneling and flow-based operations (GTP, NetFlow, Application Filtering Intelligence) will be supported in future releases.

Benefits of Orchestrated Configurations

Orchestrated configurations provides the following advantages:

Priority-free Policies

The device-level flow maps created using the standard GigaVUE-FM GUI operations are priority-based maps. As a result some of the traffic may not reach the destination tools due to priority. With orchestrated configuration, you can create priority-free policies. These priority free policies are translated into internal device-level maps that direct the traffic to the configured tools. Refer to Priority-Free Policies section for detailed information.

Overlapping Sources and Overlapping Rules

Using the standard GigaVUE-FM operations (CLI and GigaVUE-FM), you cannot create maps with overlapping source ports. However, with orchestrated configuration you can create policies with overlapping sources and also policies with overlapping rules. Refer to Overlapping Sources in Policies for detailed information.

Usability and Ease of Deployment

Without the Orchestration option, if you have to guide the tapped-and-aggregated traffic to the available traffic-monitoring and analysis tools, you must perform the following manual configurations using the standard GigaVUE-FM GUI operations:

- Create and group engine ports.
- Bind engine ports with the required application profiles to create GigaSMART Operations.
- Incorporate GigaSMART operations into flow maps.
- Use hybrid ports, if required.

Using Orchestrated Configurations you can overcome the above restrictions by simply specifying your intention for the traffic and the software intelligently creates the required device-level maps. You need not perform any of the required manual configurations mentioned above.

Orchestrated Configuration: Examples

This section provides examples for orchestrated configurations:

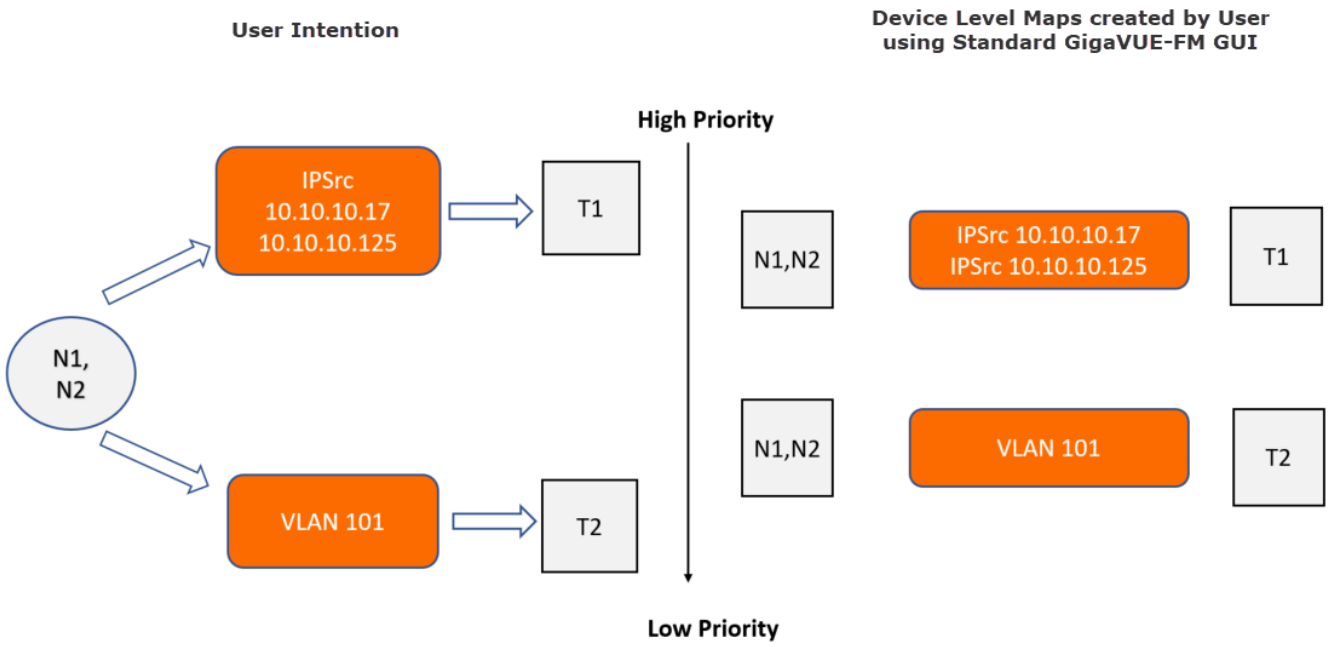
- Priority-Free Policies
- Overlapping Sources in Policies

Priority-Free Policies

In this example traffic is tapped and sent through network port sources N1 and N2 to tools T1 and T2. The intention is to send IP traffic of specific source addresses (10.10.10.17 and 10.10.10.125) to tool T1 and VLAN traffic of a specific VLAN value to tool T2.

If you use the standard approach from the GigaVUE-FM GUI to create maps, then you have to configure the following device-level maps:

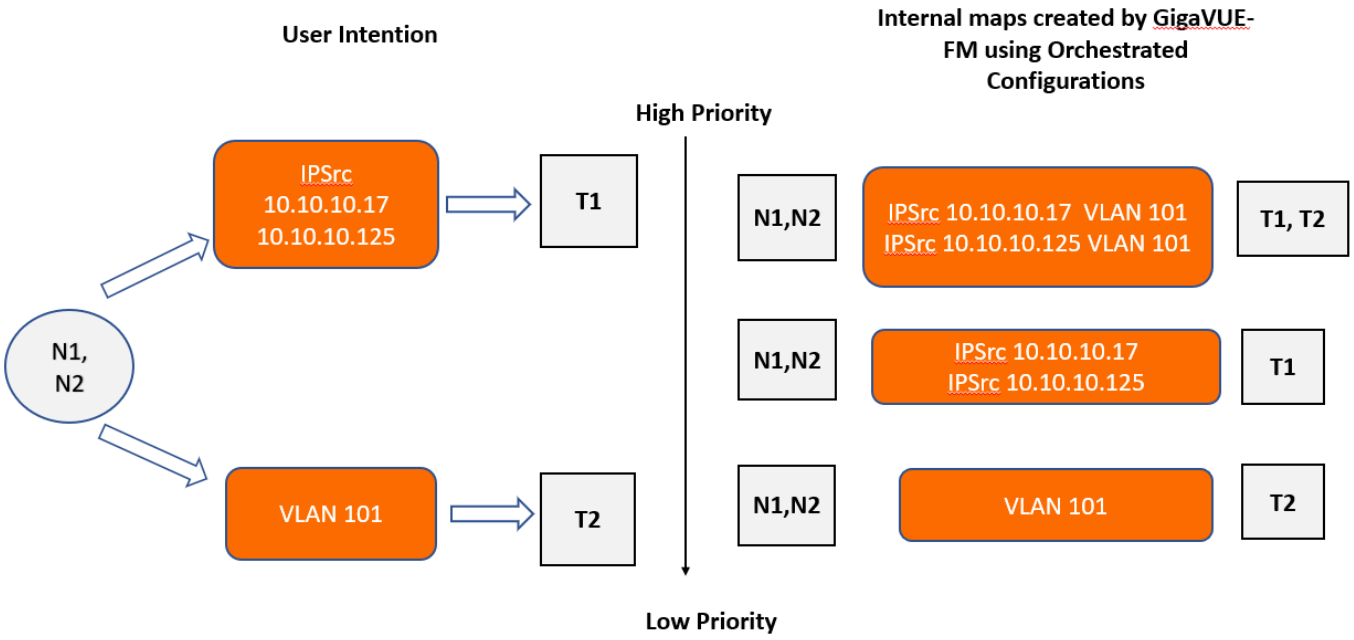
- Map from source ports N1, N2 with a rule to send traffic with IP source 10.10.10.17 and 10.10.10.125 to tool T1.
- Map from source ports N1, N2 with a rule to send VLAN 101 traffic to tool T2.



By default, the first map configured always has the highest priority. If a traffic has both IP and VLAN sources, IP traffic is sent to tool T1 but VLAN traffic is not delivered to tool T2.

With Orchestrated configuration, you only have to specify your intention for the traffic in GigaVUE-FM GUI and create a policy. The software intelligently creates the following internal maps:

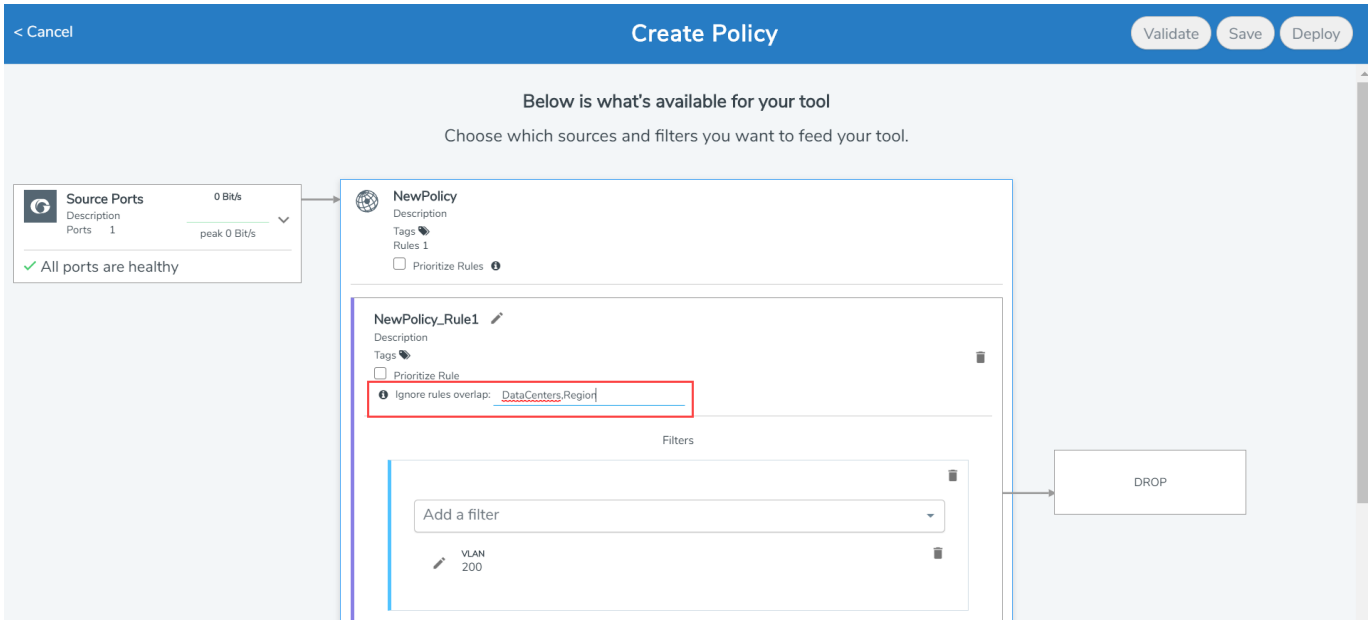
- Map from Sources N1, N2 to send both IP and VLAN traffic to tools T1 and T2.
- Map from Sources N1, N2 to send IP traffic to tool T1.
- Map from Sources N1, N2 to send VLAN traffic to tool T2.



NOTE: You can create priority-based policies by enabling the Prioritize Rules checkbox in the Orchestrated Configuration GUI. Refer to How to Create a Policy for details.

Overlapping Rules in Policies

With priority free policies, internal maps with overlapping rules will be created. You can use the **Ignore rules overlap** option to restrict creating overlapping rules across policies, thereby saving filter resources. However, to ensure that the designated traffic reaches the destination tools in case the traffic matches multiple rules, use a comma separated list of tags to the ignore matching rules.

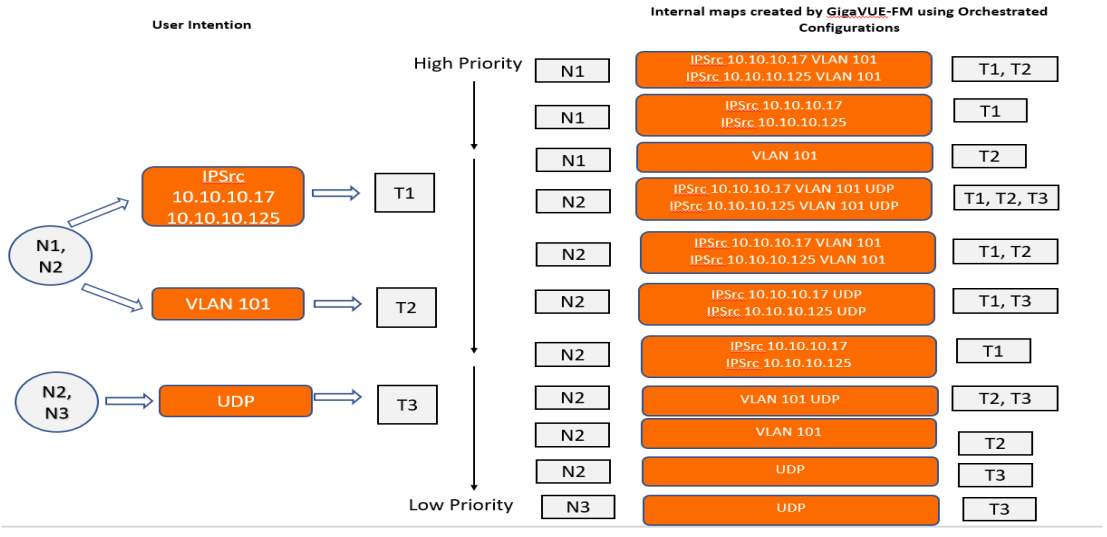


Overlapping Sources in Policies

In this example, traffic is tapped and sent through network port sources N1, N2, and N3 to tools T1, T2 and T3. The intention is to send:

- IP traffic from source ports N1, N2 to tool T1
- VLAN Traffic from source ports N1, N2 to tool T2
- UDP traffic from source ports N2, N3 to tool T3

You cannot create a map with overlapping sources using the standard GigaVUE-FM operations (CLI and GUI). However, with orchestrated configurations, you can create priority-free policies with overlapping sources as shown in the following figure.



Supported Topologies

Refer to Supported Topologies

Rules and Notes for Orchestrated Configurations

Refer to the following rules and notes for Orchestrated Configurations.

- The orchestrated configurations co-exist with the configurations created and maintained using the standard GigaVUE-FM interface options. However, you cannot edit the orchestrated configurations using the standard GigaVUE-FM operations and vice-e-versa.
- You can view the orchestrated policies and maps by enabling the *Show Auto Generated* maps option. GigaVUE-FM does not allow you to edit the auto generated maps (maps with prefix FMAuto/_FMAuto).
- When creating a policy, ensure to not include the following special characters: **. @ \ " % * : ; / , ?** in the policy name.
- The priority-free policies are translated into multiple device-level map configurations, which in turn can result in exhaustion of memory resources. However, the software does not allow you to deploy a policy that may result in exhaustion of memory resources. Appropriate warning messages that describe the reason for the policy not getting deployed, will be displayed.
- Intent Based Orchestration feature is integrated with the GigaVUE-FM Tagging and RBAC infrastructure. Refer to the Tagging and RBAC Support section for detailed information.
- GigaVUE-FM triggers alarms to track the changes in the health status of the policies. Refer to the *"Alarms"* section in the *GigaVUE Administration Guide* for more details.
- In orchestrated configurations, you can drop the traffic instead of sending the traffic to a specified tool port. Refer to Drop Rules in IBO section for details related to drop rules.
- You can use hybrid ports in a policy as both source ports and destination ports. You must configure the hybrid ports prior to using them in a policy. Refer to Work with Hybrid Ports section for details about configuring the hybrid ports.
  You can use hybrid ports for policy destination as follows:
  - Tool object
  - Hybrid port
  - GigaStream comprising of hybrid ports

  Hybrid port that are part of the destination in a policy can also be used as source in another policy.
- Ports used in a policy will not be released until you delete the policy.
- When you upgrade to software version 5.11.00, the policies created previously will not be retained in the system and will not be listed. You must create the policies again.

- Orchestration is not supported in GigaVUE-HB1 device (both as a stand-alone device and also when it exists within a cluster).

Create Orchestrated Policies

Go to **Traffic** and under **Physical > Orchestrated Flows> Out-of-Band Flows** option in the GigaVUE-FM GUI to create and view policies, go to **Traffic >**. The following tabs are available:

- **Policies:** Displays the list of configured policies. Click on a policy to view the following details:
  - Health Status
  - Deployment Status

  You can deploy, undeploy, edit, delete, and create policies from this page.
- **Tools:** Displays the configured tools. You can associate new tools using the **Create New Tool** option. You can also group the tools to create a tool GigaStream. You can apply egress filters in tool ports. Refer to the Egress Filters section for more details.
- **Sources:** Displays the configured ports. You can edit the ports and create port groups from sources. Do not use fan out port as source port in policy.
  - **Filters:** Displays the available filters. You can create a new template from the Filters option. Starting in software version 5.9, you can use advanced filters for creating the filter rules.
    - **Packet Transformations:** Displays the supported packet transformation options. You can also create a template with the required Packet Transformation option.
    - **Visualize**: You can select multiple policies and click **Visualize** to compare the policies.

NOTE: You can create policies by navigating to **Actions > Create Policy** from all of the tabs listed above. You can view the health status of the ports, tools and policies when creating the policy.

Prerequisites

You must have the following licenses:

- GigaSMART Masking
- GigaSMART Packet Slicing
- GigaSMART De-duplication
- GigaSMART Header Addition
- GigaSMART Header Stripping

How to Create a Policy

Refer to the Rules and Notes for Orchestrated Configurations section before creating the policy.

To create orchestrated policies:

1. Go to **Traffic** and under **Physical > Orchestrated Flows> Out-of-Band Flows**.
2. Click **Create Policy**. The Create Policy wizard is displayed. You can specify a name for the policy.
3. **NOTE:** You can configure the sources, tools and policies in any sequence you want.
4. Select **Sources** to select the available source ports or port groups.
5. **NOTE:** If a port (port type: network or hybrid) is already used as a source port in any of the following standard GigaVUE-FM configurations, then those ports will not be listed in the drop-down:
   MapsFabric MapsApplication Intelligence 5G CUPS solutionHowever, you can view all the ports in the Ports page.
6. Click **Select Tools** to select the required tool ports (port type: tool port or hybrid). You cannot select a port that is already used. You can also associate a new tool using the **Create New Tool** option.
7. **NOTE:** You can select **Drop** to drop the packets without selecting the tools. Refer to the Drop Rules in IBO section for details related to drop rules.
8. Create the policy with the required rules. You can define the required filters and criterion in the rule.
9. **NOTE:** You can create a policy with multiple rules. With in each rule, you can configure multiple criteria and multiple filters and use them together with the packet transformation options.If you have configured multiple filters in a rule, then the traffic will be filtered only if all the filter rules are true.If you have configured multiple criterion in a rule, then the traffic will be filtered even if one of the criteria is true.
10. Select the required packet transformation option. You can combine multiple packet transformation options with in a single rule. The GigaSMART packet transformation operations are performed in parallel.
11. Click on the edit option next to the **Processing Engine** option to select the required GigaSMART engine ports for your rule. If you do not select the engine ports, then they will be automatically selected.

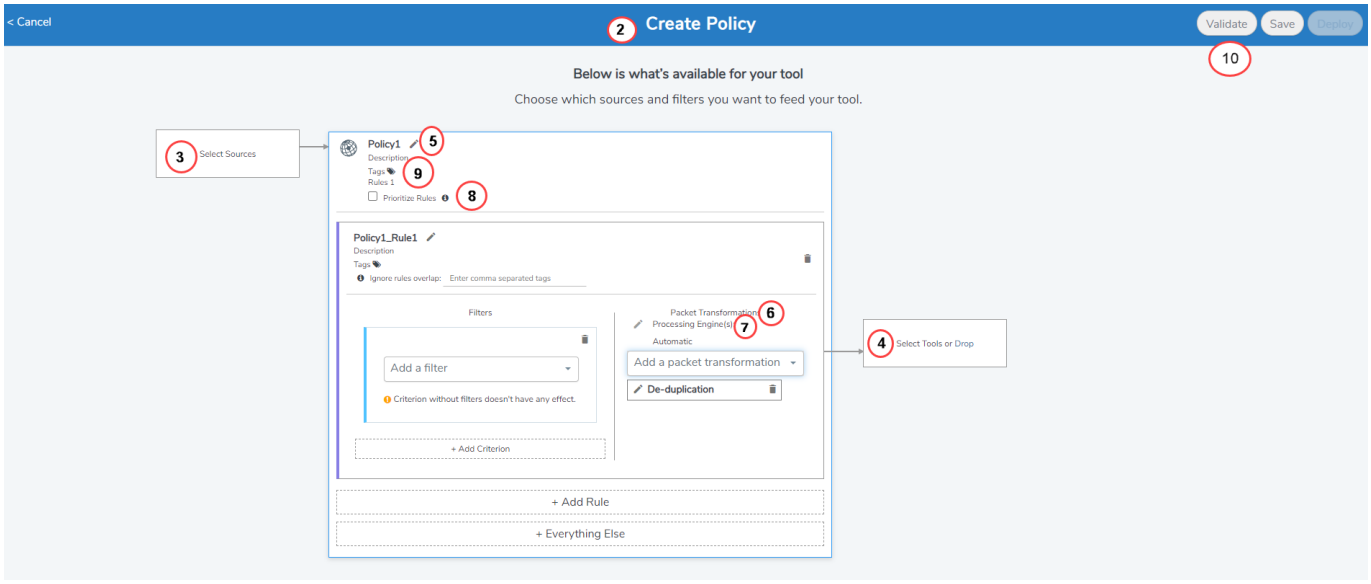
12.

> **NOTE:** Engine ports that have not been used in standard GigaSMART operations will only be available for selection. You cannot edit or delete an engine once a policy is deployed. However, you can remove the packet transformation and add the packet transformation option with a different engine.
> For policies created based on multicluster topology, you must manually select the engine ports, as automatic allocation is not supported.
> For policies created based on single cluster topology, automatic allocation of engine ports is supported. However, only engine ports that have not been already used will be allocated.

13. Enable the **Prioritize Rules** check box to create priority-based policy.
14. Click the **Tags** option to associate the policy to tags. Refer to Tagging and RBAC Support section for more details.
15. Click **Validate** to validate the policy. You can either:
    - Click **Save** to save the policy. You can deploy the saved policies later.
    - Click **Deploy** to deploy the policy.

> **NOTE:** Use the *Everything Else* option to configure the shared collector. Refer to the *GigaVUE Fabric Management Guide* for details on Shared Collector.

The policy thus created is listed together with the list of policies. You can edit, delete, deploy and undeploy a policy, as required.



Tagging and RBAC Support

Role Based Access Control (RBAC) in GigaVUE-FM controls the access privileges of users and restricts users from either modifying or viewing unauthorized data. In Orchestrated Configurations, RBAC controls the accessibility of the users to the policies based on the tags. Tags can be either RBAC tags or aggregation tags. Multiple tags can be assigned to policies and rules.

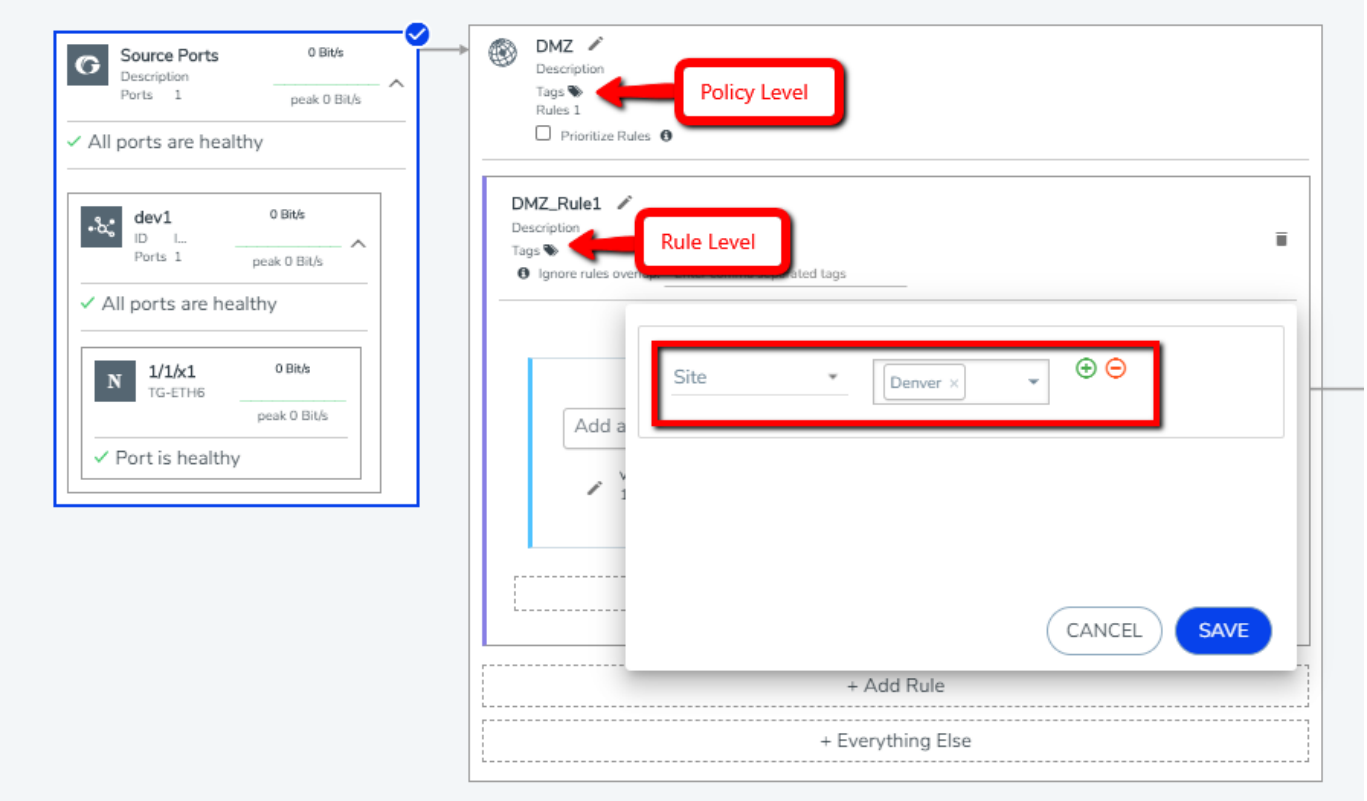
NOTE: RBAC tags are supported only at the policy level. The orchestration wizard allows RBAC tags to be applied at the rule level as well, though it is not supported in this release.

To associate policies to tags:

1. Click the **Tags** option at the policy level. You can associate tags either at the time of creating a policy or edit an existing policy.
2. Select the required tag keys and tag values.

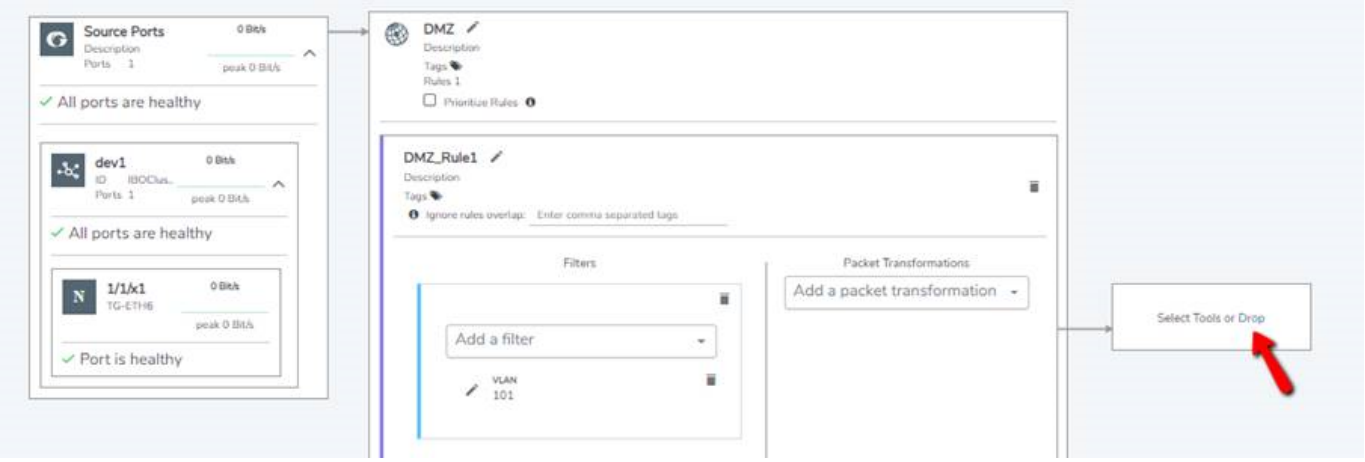
   NOTE: The tag key and the associated tag values must be created in advance in GigaVUE- FM. Refer to the "Tags" and "Role Based Access Control" sections in the GigaVUE Administration Guide for more details
3. Click **Save**.

   Once you associate the policy to a tag, only users with fm_super_admin role or users with read/write access to the Traffic Control Management category can access the policies.



Drop Rules

You can use the **Drop** option in the 'Create Policy' wizard as the destination to drop a packet. For a drop rule to pass a packet, it must be used together with a pass rule or with the 'Everything Else' rule, such that the packets that do not match the drop criteria will be forwarded to the tools.

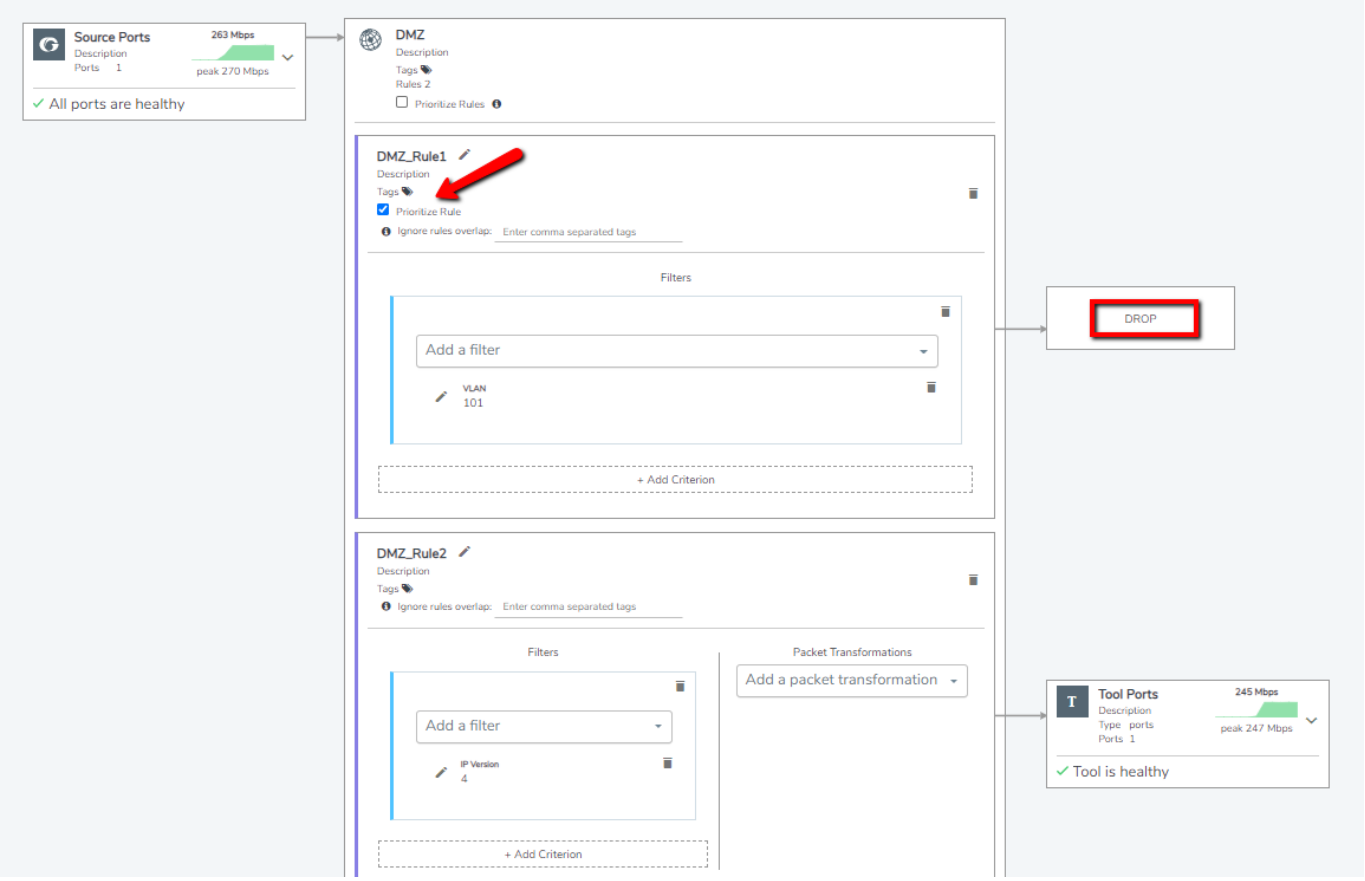


In orchestrated configurations, there are no rule priorities for pass rules. However, the following is the order of priority for prioritized and unprioritized drop rules together with the pass rules:

1. Prioritized drop rules
2. Pass rules
3. Unprioritized drop rules
4. Default collector

A drop rule, by default, has lower priority than the pass rule and higher priority than 'Everything else'. Therefore, a drop rule will drop the traffic before the traffic gets passed to the destination of the 'Everything Else' rule. However, a prioritized drop rule has the highest priority amongst all the rules in a policy and will drop matching traffic before any other policy rule can process that traffic.

Consider a scenario in which the intent is to drop packets with VLAN 101 and pass packets with IP version 4. Based on the default settings, if there is only IPv4 traffic on the source side, packets with VLAN Id 101 will not be dropped because pass rule has higher priority than the drop rule. To choose 'Drop' as high priority within the policy, you must prioritize the drop rule by checking the 'Prioritize Rule' option.



NOTE: To prioritize the drop rules in deployed policies, you must first undeploy the policy.

Import and Export Orchestrated Policies

Starting in software version 5.11.00, you can import and export an orchestrated policy in YAML format. This allows you to bulk deploy the policies.

The following are the advantages of importing and exporting policies:

- Retrieve a policy that was deleted unintentionally.
- Deploy the policy in another device.
- Re-deploy a policy in the device after GigaVUE-FM is upgraded to a new version (in case of issues in the existing solution).

The YAML file contains the following information:

- Policy
- Sources
- Rules
- Packet Transformation
- Destination

The below is the Import and Export Policy sample template of a YAML file. You can use the following sample template as a reference while building a policy in YAML format.

```
name: <PolicyExample>
priority: false
deployed: false
source:
- <IP address of the device>:<Port Number>
- <Network Name>
rules:
- ruleName: <Rule Name>
tags: chicago
type: pass
highPriorityDrop: false
noExpansionTags:
criteria:
- criteria1:
filters:
- type: ipSrc
addresses:
```

```
- <IP address of the device>
- <IP address of the device>
- type: portSrc
addresses: 5
- type: ipVer
addresses: 4
- criteriaWithBinding
packetTransformation:
- type: dedup
parameters:
action: drop
timer: 50000
ipTclass: include
vlan: ignore
tcpSeq: include
ipTos: include
- Slicing-gsApp
tools:
- Tool1

# it is practical to have all the tool bindings separate from the policy body in
order to leave the policy body as an abstract policy template
toolBindings:
- toolName: Tool1
# in the future we may have also metadata receivers
receiverType: packets
groupingType: replicate
outputs:
- type: port
ports:
- Generic2:22/3/x1
# possibly more tool bindings...
sourceBindings:
- sourceName: ChicagoServers
tags:
inputs:
# in the future we will have various types including types that reflect
tunnelling arrangements
- type: port
clusterName: <ClusterName>
ports:
- cluster1:22/2/x12
- cluster2:22/2/x11
criteriaBindings:
- criteriaName: CriteriaWithBinding
filters:
- type: ipDest
addresses:
- <IP address of the device>
- <IP address of the device>
- type: port
addresses:
- 5
```

Rules and Notes

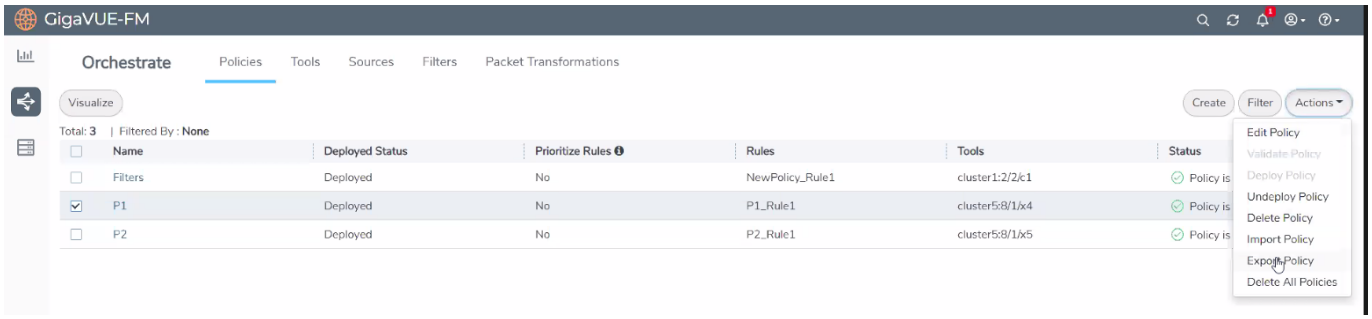
Refer to the following rules and notes:

- You can export policies that either deployed or undeployed. However, when you import a policy, it will be in undeployed status. You must manually deploy the imported policy.
- You cannot import or export the egress port filters in a policy. You must manually apply them.
- When you export a policy that has GigaStream , port groups and other such groups, you must ensure that those groupings exist when you import the same policy. This also applies to tools that were created through the tool wizard.
- Port types involved in a policy that is exported must not be changed when trying to import the same policy.
- If templates were used to create a policy that was exported, then those templates must remain when importing a policy.
- You can import and export several policies in one operation.

Import and Export Orchestrated Configuration

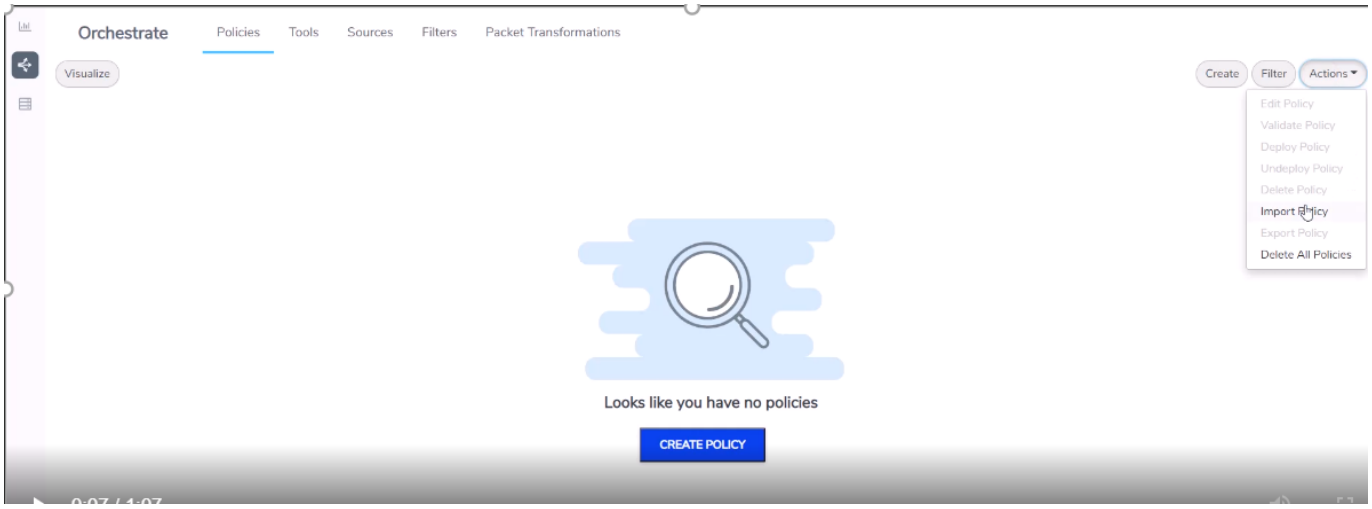
To export a policy:

1. In the Policies page, select the policy that must be exported.
2. Select **Actions** and click **Export Policy**. The policy is downloaded as an YAML file.
3. Save the file to the required location.



To import a policy:

- In the Policies page, select Actions and click **Import Policy.**
- Browse to the folder that has the required policy file in YAML format.
- Select the file and click **Open**. The **Import** button gets activated.
- Click **Import** and refresh the page.



Egress Filters for Additional Filtering Capabilities

Egress filters provide additional filtering capabilities when applied on tool or hybrid ports in a policy. Egress filters are used to pass or drop the traffic and can be combined logically using 'AND' and 'OR' operators.

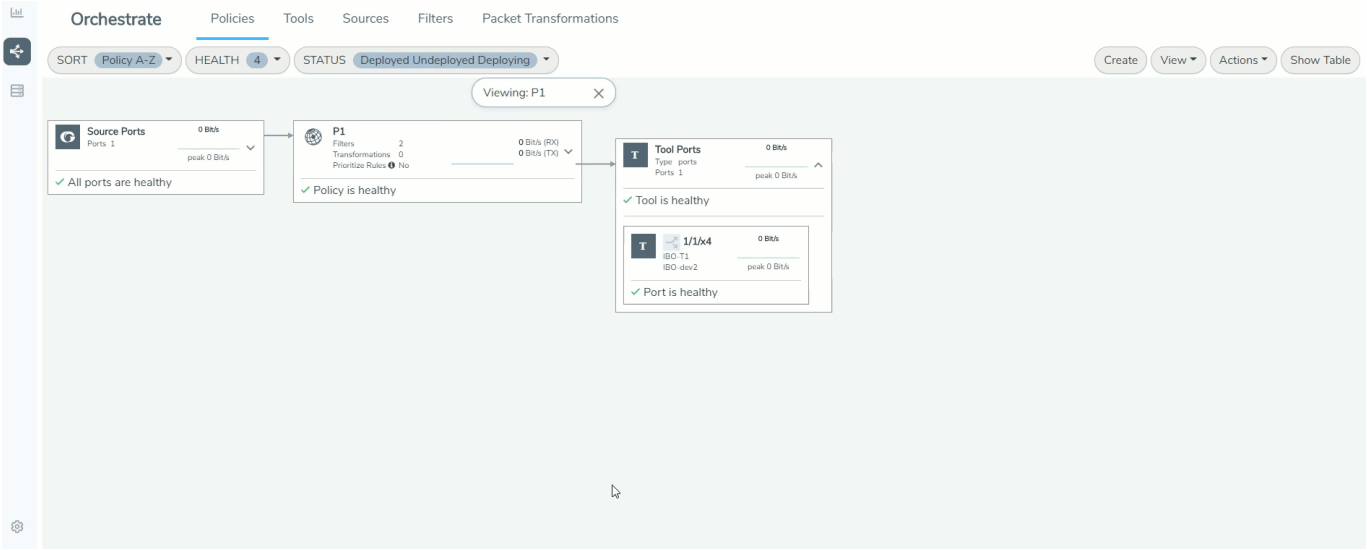
NOTE: The number of egress filters vary depending on the platform. Refer to the Port Filters section for the number of filters allowed for each platform.

To apply egress filter for a tool port in a policy:

1. Select the tool port or the hybrid port for which you need to add filters.
2. Click the **Egress filter** icon and click **Edit.**
3. Click **Add a Rule.**
4. Select the required conditions.
   - **To create an egress filter with logical AND:** In a single rule, create multiple conditions. That is, create *Rule 1* with conditions for filtering IPv4 traffic and source port as '443'. Traffic will get filtered to the tool only if both of these conditions are true.
   - **To create an egress filter with logical OR**: Create multiple rules with required conditions. That is create Rule 1 to filter IPv4 traffic and Rule 2 to filter traffic from source port ''443'. Traffic will get filtered to the tool even if one of the condition is true.
5. Define a name to each of the rules created and click **Save**.
6. You can click **Pass** or **Drop** against each of the rules to either pass the traffic to the destination or to drop the traffic. Refer to the following notes for pass and drop rules:

- Drop rules are applied first.
- If there are only drop rules, then all the traffic except that specified in the drop rule(s) will be passed.
- If there are only pass rules, only the traffic specified in the pass rule(s) will be passed and all other traffic will be dropped.
- If there are both drop and pass rules, only the traffic specified in the pass rule(s) that does not also match the drop rule(s) will be passed.

NOTE: Use the edit icon to edit the egress filter directly from the canvas. If you edit the egress filter in edit mode, the **Deploy** button is disabled.



Glossary

- ***Policy***: A user defined instruction for what to do with the traffic
- ***Packet transformation***: GigaSMART Operations
- ***Everything else***: Shared Collector
- ***Filters and Criteria***: Map rule criteria to filter the traffic

Traffic Policy

Traffic Policy provide you with a unified way to configure how network traffic flows through your monitoring and security tools. It enables you to define rules and actions that determine how traffic received from s

ource ports is processed and forwarded to the appropriate tool ports.

With the unified Traffic Policy interface, previously separate configuration models are brought together into a single workflow:

- When you select source and destination ports within the same cluster, Traffic Policy acts as the replacement for flow maps, providing a consistent way to filter, process, and forward traffic within a single cluster.
- When you select source and destination ports across clusters, Traffic Policy replaces Fabric Maps, allowing you to configure end-to-end traffic distribution across multiple clusters using the same policy-driven approach.
- When configuring Application Intelligence solutions, you define a single policy that captures your intent for how to filter, process, and distribute traffic. GigaVUE-FM then creates and maintains the required underlying objects, such as virtual ports (vPorts), GigaSMART operations (GSOPS), and different map levels.

Map Components in a Traffic Policy

Every Traffic Policy uses a small set of map components. These appear as blocks GigaVUE-FM. These components define how traffic enters the policy, how it is filtered and processed, and where it finally goes.

- **Source** - The entry point for traffic into the policy.
  - Can be network ports, port groups, GigaStream Ports, inline networks, or inline network groups, depending on the use case.
  - The ports can be from a standalone device, within a cluster, or across clusters.
  - When you select the source port from a standalone node or within a cluster, all applications are available (subject to your license). When you select source ports across clusters, only the My GSOP application is available.
  - You can add multiple Source blocks in a traffic policy to ingest traffic from different ingress locations into the same Application Intelligence workflow, allowing you to apply a single, consistent policy to traffic coming from different parts of the network.

  NOTE: The ports can be searched using regEx pattern in the **Quick Port Editor** while configuring the source.

- **Source Ruleset** – A set of L2–L4 header rules that filter traffic at the source. These are used in regular, first-level, and inline map types. These are not used in second-level maps as the maps rely on application-level rules instead. The Source Ruleset has the following three rule types:
  - **By Rule** – Use this when you want to filter traffic based on specific L2–L4 conditions. For example, VLAN, IP subnet, protocol, DSCP, or TCP/UDP port, among others. You add one or more conditions (for example, IPv4/IPv6 source or destination, VLAN range, protocol, DSCP, TCP/UDP ports, pattern match), and GigaVUE-FM treats them as an OR/AND combination. Rules cannot be empty for this rule type. If no rules are defined, traffic from the source ports is sent to the collector map (if one is configured). If no collector map exists, the traffic is discarded.
  - **Pass All** – Use this when you want to accept all traffic from the selected source ports without any filtering. You cannot define any rules for this rule type because the intent is to match all traffic and send it to the destination.
  - **Collector** – Use this when the map acts as a collector for traffic forwarded from other maps with the same set of source ports, instead of directly filtering incoming packets. A Collector ruleset usually does not define match conditions. Similar to Pass All, you cannot define any rules for this rule type because the map receives whatever its upstream maps forward.

  > **NOTE:** You can create only one collector map for a set of source ports through CLI/API. However, GigaVUE-FM allows you to create and manage multiple collector maps for Traffic Policy from the same set of source ports. During deployment, GigaVUE-FM merges the collector maps and deploy them as a single collector map in the device. It is recommended to create only one shared collector though GigaVUE-FM allows to create multiple collectors.

- **Application Ruleset** – A set of application-level rules that filter traffic based on application attributes or advanced matching, such as gsrule patterns or AFI rules, among others.
- **Applications** – The Applications component lets you attach one or more GigaSMART operations and Application Intelligence functions to the selected traffic. You can choose from several options and chain them when needed.
  - **My GSOP** – Use this when you want to apply a custom chain of GigaSMARToperations that you define in a GigaSMART group. You can combine multiple operations in a GSOP profile. For details, refer to How to Combine GigaSMART Operations.
  - **Deduplication** – Use this when you want to remove duplicate packets from mirrored traffic before it is processed by AppViz, AFI, or AMI, so that each flow is analyzed only once and your dashboards and metadata exports reflect accurate traffic volumes. You do not need to create a separate GSOP for this. GigaVUE-FM automatically creates and manages the required underlying objects based on your Traffic Policy intent.
  - **Application Filtering** – Use this when you want to pass or drop traffic based on applications instead of only IP or port values. You can select individual applications, application families, or tags and mark them as Pass or Drop. You can configure

multiple Destination Traffic Priorities so that different sets of applications are directed to specific tools, with a clear priority order for each destination. You can add a maximum of five AFIs in a Traffic Policy.

- **Application Metadata** – Use this when you want to export flow-level records for monitored traffic to external collectors. Application Metadata Intelligence (AMI) generates unsampled NetFlow/IPFIX/CEF records with standard information elements, for example, 5tuple, bytes, packets, and basic attributes.
- **Application Visualization** – Use this when you want to see which applications are present in the traffic and how much bandwidth they use before you add filtering or metadata export. It identifies thousands of applications and groups them by application, family, and tags. It displays traffic statistics in bytes, packets, and flows, and shows top applications over a selected time range.

NOTE: AppViz is not mandatory for configuring Traffic Policy for Application Intelligence solutions. You can create and deploy Traffic Policy that uses only Application Filtering Intelligence (AFI) and/or Application Metadata Intelligence (AMI) without enabling AppViz, if you do not require dashboard-level application statistics.

- **NetFlow** – Use this when you want to generate flow records directly from maps. It filters the traffic at the packet level using map rules, while exporting summarized flow records or rich application metadata to analytics tools instead of, or in addition to, raw packets. You can add applications to the Traffic Policy based on the source ports you select and the licenses you purchase. When you select the source port from a standalone node or within a cluster, all applications are available (subject to your license). When you select source ports across clusters, only the My GSOP application is available.

- **Destination** – The egress point for the processed traffic. Can be tool ports, tool GigaStream, null-ports, tunnels (L2GRE, VXLAN), inline tool ports, inline tool groups, inline serial tool, fabric port groups, or vPorts, to name a few depending on the use case.

How Traffic Policy Choose Map Types

When you create a Traffic Policy, GigaVUE-FM derives the map type from the map components you use and where you place them in the flow.

- **Regular map** – Connects a source directly to a destination with optional GSOP applied, but does not support application-level rules.
  - Source is a network port or a hybrid port, and the destination is a tool port.
  - Rules reside in the Source Ruleset component and match on L2–L4 header rules.
  - Typical for out-of-band monitoring, where line-rate traffic from TAPs or SPANs is filtered and sent to tools. Suitable for simple “tap traffic and send to tools” use cases.
- **First-level map** – Typically performs coarse filtering before handing traffic to a second-level map.

  - Source is a network port or a hybrid port, and the destination is a vPort or tool port. At least one vport is required for a first-level map.
  - Rules reside in the Source Ruleset component. The first-level map forwards traffic from physical sources to vPorts, often as a preparation step for application processing in a second-level map.
- **Second-level map** – Processes traffic based on application-level rules and GigaSMART operations together.
  - Source is a vPort and the destination is a tool port, fabric port group, tool/hybrid port group, tool/hybrid circuit GigaStream.
  - You cannot select a Source Ruleset component here. You can only use the Application Ruleset component.
  - Second-level maps often fan out to multiple destinations and can be prioritized when needed.
- **Transit-level map** – When you chain multiple GSOPs between vPorts because some applications cannot run in a single GSOP.
  - Source and destination are vPorts. The map just performs GSOP and passes traffic on to the next vPort.
  - From these vPorts, you can attach multiple (prioritized) second-level maps to branch traffic to different tools.
- **Inline map** – When you want to route the traffic through tools in series, not just out-of-band to them.
  - Source is an inline network port or inline network group and destination is an inline tool port, inline tool group or an inline serial tool.
- **Inline first-level map** – When you want to route the traffic through tools in series in a first-level map.
  - Source is an inline network port or inline network group and destination is a vPort.
- **Inline second-level map** – When you want to route the traffic through tools in series in a second-level map.
  - Source is a vPort and destination is an inline tool port, inline tool group or an inline serial tool..

Note: Intent-based map selection only applies to Traffic Policy created for the Application Intelligence solution. For configuration of intent-based map, refer to Configure Traffic Policy for Application Intelligence solutions.

Map Priorities in Traffic Policy

When you have more than one Traffic Policy using the same source, map priority decides which policy sees the traffic first. Traffic Policy follow the same idea as classic flow maps and fabric maps. When rules overlap, the higher-priority map handles matching packets before the lower-priority maps. The Pass All and collector maps do not have priority. This behavior aligns with the existing implementation for flow maps and fabric maps, and it is preserved when those maps are migrated into Traffic Policies.

Priority Options

When a Traffic Policy shares a source with other Traffic Policy, you can change its priority using one of these options:

- **Highest** – Move the policy to the top of the list for that source.
- **Lowest** – Move the policy to the bottom of the list for that source.

You can create a map with the required priority. By default, maps are created with the lowest priority. You can change a map's priority from lowest to highest at the policy level when you create or edit the policy.

How Traffic Policy Prioritize Maps

Traffic Policy keeps the order of policies per source in a structure called a map chain.

For each Traffic Policy, GigaVUE-FM tracks:

- The Traffic Policy alias.
- The Source Ruleset for each Source block.
- The cluster and source ports used by that Source block.

For a given source on a cluster, GigaVUE-FM keeps an ordered list of traffic flows. This is the map chain. When you change priority, GigaVUE-FM updates this ordering and then updates the device-level map order to match it.

You do not manage map chains yourself. You change priority at the policy level, andGigaVUE-FM applies the changes to the underlying maps.

How Map Priority Affects Traffic Flow

When a packet arrives at a source that is used by multiple Traffic Policy, GigaVUE-FM follows this flow:

1. It finds all deployed Traffic Policy that use the source on the same cluster.
2. It orders them using the map chain (priority order).
3. For each policy in order:
    - It applies the Source Ruleset and, if present, the Application Ruleset.
    - If a rule matches and the map has a Pass action, GigaVUE-FMforwards the packet to that Traffic Policy's destinations.
    - If a rule matches and the map has a Drop action, the packet is dropped for that Traffic Policy.
    - If the packet does not match any rule in that policy, GigaVUE-FM moves on to the next Traffic Policy in the chain.

Best Practices for Map Priority

Use these guidelines to avoid surprises with overlapping traffic:

- **Check which Traffic Policy shares a source**
  Before you change priority, review which Traffic Policy use the same source ports on each cluster.
- **Place specific Traffic Policy before broad ones**
  Put Traffic Policys with narrow match conditions (for example, one subnet or one VLAN) above Traffic Policy that match a wide range of traffic.
- **Use Pass All with care**
  A high-priority Pass All map can prevent lower-priority policies from ever seeing traffic on that source. Use Pass All only when you want a catch-all map at that point in the chain.
- **Validate changes with statistics**
  After you change priority, review map statistics and policy level NRT data to confirm that the right policy is now handling the expected traffic.

Following these practices helps you control how overlapping Traffic Policy behave and keeps packet delivery predictable across your deployment.

Source Overlap Behavior in Traffic Policy

In a cluster, the same source ports are often used across multiple Traffic Policies. To avoid overlapping sources, when you remove a source port from one policy, you must also remove the corresponding source ports from all related policies. For example, if a cluster uses source ports 1/1/x1 and 1/1/x2 across several Traffic Policies and you remove 1/1/x1 from one policy, GigaVUE-FM also removes 1/1/x2 from that policy to keep the source set consistent and prevent source overlap within the cluster.

Traffic Policy Life Cycle

Each Traffic Policy follows a clear life cycle that gives you flexibility and control:

- **Create and draft** – You can create a Traffic Policy and save it as a draft while you refine the policy and its components.
- **Edit rules** – You can add, edit, or remove rules in the draft. You can sort, filter, and paginate rules to manage large configurations.
- **Deploy** – When you deploy a policy,GigaVUE-FM pushes the required configuration to the devices and keeps the policy definition.
- **Update and redeploy** – You can update a deployed policy, save the changes, and deploy again when needed.
- **Undeploy and delete** – You can undeploy a policy to remove its configuration from the devices while keeping it in GigaVUE-FM, or delete it when you no longer need it.

Best Practices for Traffic Policy

Use the following best practices when creating or working with Traffic Policy:

- Before you create a map that spans multiple clusters, ensure that inter-cluster connectivity is already in place: configure circuit ports, circuit GigaStream, and manual topology links, or enable GDP on the circuit ports so that GigaVUE-FM can discover all connected clusters.

Important Notes

Keep in mind the following notes while working with Traffic Policy

- When you create a regular multi-hop map with GSOP, you can place the GSOP at any point along the path between the source and destination cluster or node, including within the source or destination cluster. However, it cannot be placed beyond the source or destination cluster.
- Circuit ID allocation and optimization for maps behave as follows (for both multi-hop maps and regular maps):
  - The limit that a single circuit GigaStream can source and decap traffic is 1024 multi-cluster maps, and multiple decap circuit GigaStreams can be configured within a cluster.
  - You can use the Circuit ID optimal allocation feature when you create new maps, it cannot be applied retroactively to existing maps.
  - Multi-hop maps do not support ingress VLAN port tagging on ports that participate in inter-cluster circuit ID tunnels. For any map that uses circuit ID tunnels between clusters, do not configure ingress VLAN tagging on the associated ingress ports. This configuration is unsupported and can lead to unexpected behavior.
  - Circuit IDs are allocated per hop and per direction.
  - Circuit ID optimal allocation is not applicable to second-level maps that use a fabric port group as the destination.
  - Circuit tunnels are not shared between first-level maps, even when the maps are configured in Shared mode.

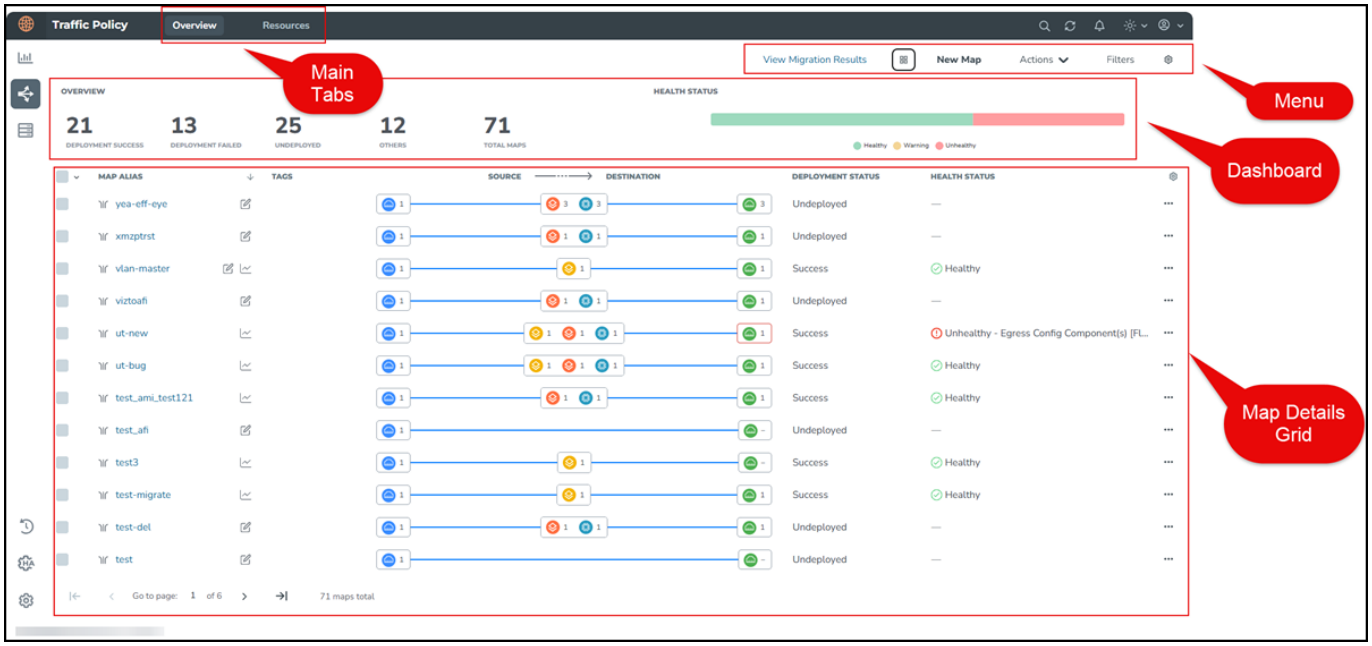
Limitations

Review the following limitations when you work with Traffic Policy

- In multi-hop scenarios, the configuration is stored only in GigaVUE-FM, not on the individual devices. If you move the devices to a different GigaVUE-FM instance, these map configurations cannot be rediscovered from the devices. To retain them, back up the configuration from the original instance and restore it on the new instance.
- You cannot configure a map that uses encapsulation circuit ID on a source port that is already used in a passall map. If you try this from the CLI, GigaVUE-FM displays the following error message: "Cannot have encap on a map with a map-passall or port-pair on the source ports."
- You can deploy or undeploy only one map at a time.

Explore Traffic Policy Interface

In the GigaVUE-FM Homepage, go to **Traffic > Traffic Policy** to view the Traffic Policy landing page as shown below.



The Traffic Policy landing page includes two main tabs, namely:

- Overview
- Resources

In the **Overview**, you can view the details of the existing maps, their health status, deployment status and you can access the maps to edit, delete, clone, deploy, and undeploy, and so on. In the **Resources**, you can create fabric port groups and can export or import tool templates.

Overview

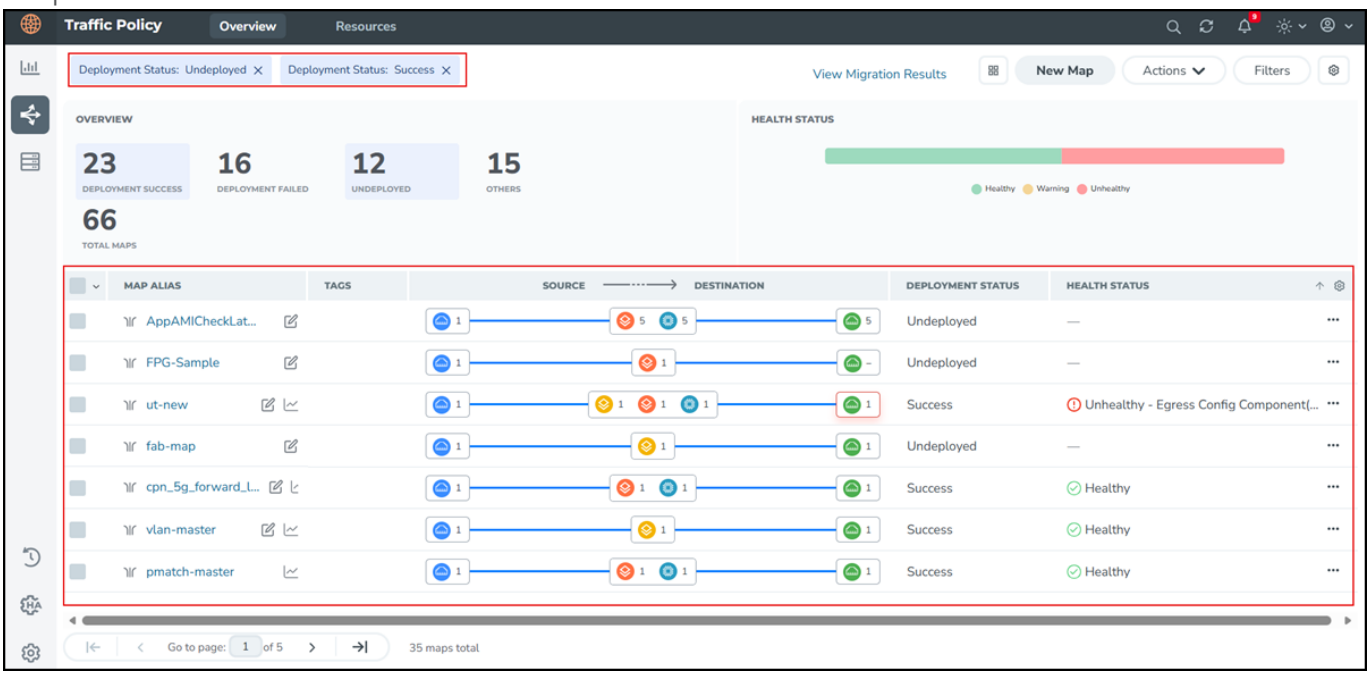
The Overview includes the Dashboard, Map Details Grid, and Menu including actions that can be performed in the map. The following sections explain these in detail.

- Dashboard
- Map Details Grid
- Menu

Dashboard

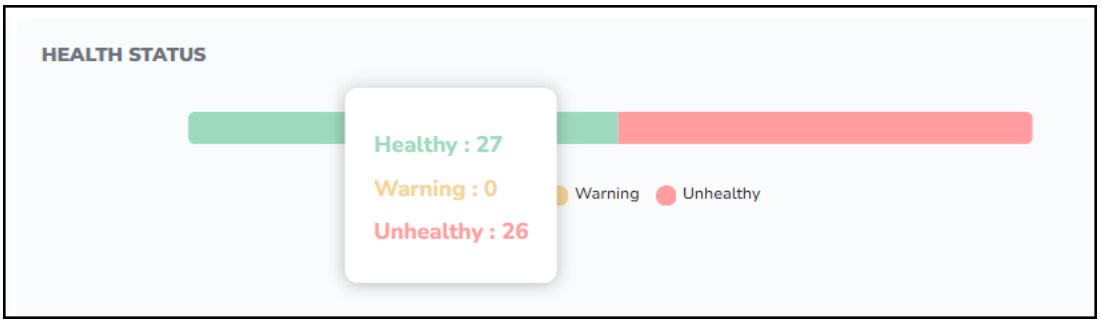
The Dashboard provides the overview of the traffic policies and their health status. You can hide or unhide the Dashboard using the Dashboard icon in the Menu. The Dashboard section includes the Overview section and Health Status section.

- **Overview section**
  - Displays the total number of maps, number of maps based on the deployment status such as Success, Failed, Undeployed, and Others (status including Pending, Partial Success, and Damaged).
  - Sorts the maps in the Map Details Grid based on your selection of items in the Overview section. For example, if you click the Success item, the maps with Success status are listed in the grid followed by other statuses. This section displays the selected item labels at the top of the Dashboard. The Map Details Grid displays the maps in the sorted order of selected items in the Overview section.



- **Health Status section**
  - Represents the overall health of all traffic policies created in the Traffic Policy interface in a color-coded horizontal bar with segments for Healthy (green), Warning (Amber), and Unhealthy (red).
  - Displays the statuses with the count, when you hover over the horizontal bar as shown below.

- Sorts the traffic policies in the Map Details Grid based on the health status you select.



The following table lists the health status and its descriptions.

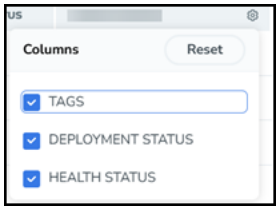
Color	Health Status	Description
Green	Healthy	All components in the Traffic Policy (sources, applications, destinations) are up and passing traffic as expected.
Amber	Warning	One or more components are in a partially available state. Traffic may be impacted, but some traffic can still flow.
Red	Unhealthy	One or more components are down or misconfigured, and traffic is likely not flowing as expected for this policy.

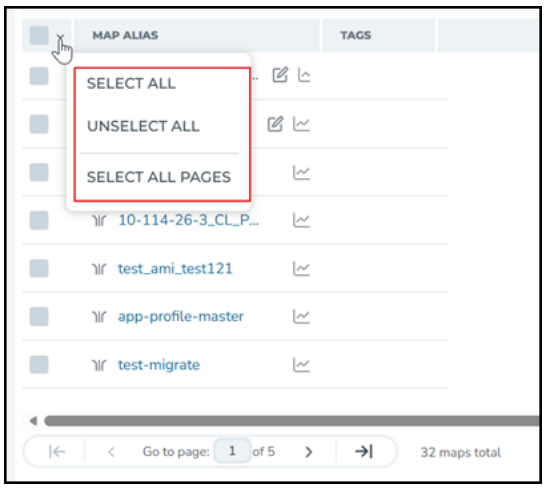
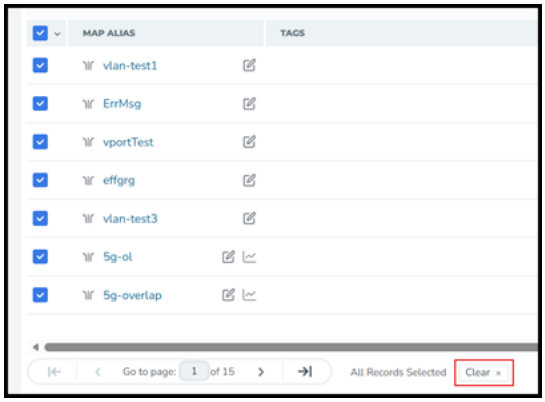
Map Details Grid

The Map Details Grid displays the maps that are created in the Traffic Policy with the tags, deployment status, health status, and so on.

The following table displays the different columns in the Map Details Grid and their description.

Grid Columns	Description
Map Alias	Name of the Traffic Policy.
Tags	Tags linked to the Traffic Policy. To know how to create, work with and add tags to maps, refer to Tags.
Source -> Destination	Displays a visual representation of the map components that are used in the Traffic Policy.

Grid Columns	Description
	- Application - Application Ruleset - By Rule - Collector - Destination - Pass All - Source
Deployment Status	Displays the deployment status of the Traffic Policy. Valid values are Success, Failed, Undeployed, Pending, Partial Success, and Damaged.
Health Status	Displays the health status of the Traffic Policy. Valid values are Healthy, Warning, and Unhealthy.
Settings	Enables you to add or remove columns from the grid view. 
Select All	Enables you to select the traffic polices in the current page, select the traffic policies in all pages, or unselect all the traffic policies.

Grid Columns	Description
	 When you click the **Select All Pages** option, all the traffic policies in the interface are selected and **Clear** option is given as shown below. 

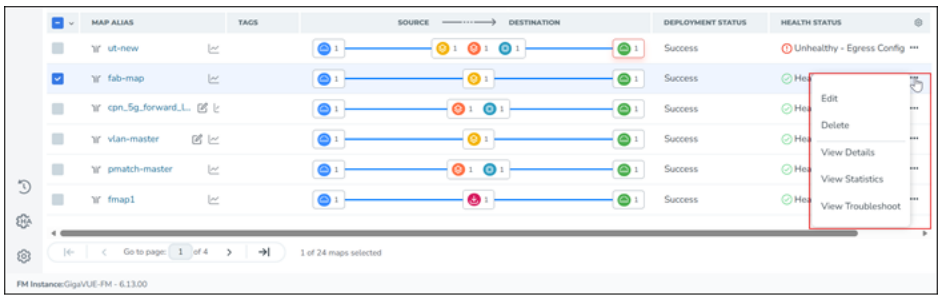
The following table lists the possible deployment statuses for the Traffic Policy and suggested actions to be done for each status.

Deployment Status	Definition	Suggested Action
Success	The Traffic Policy was accepted at GigaVUE-FM level and traffic flow was deployed correctly.	No action required. Optionally verify the device-side maps and statistics if you want to confirm counters are incrementing as expected.
Failed	The Traffic Policy was not configured correctly and not accepted or deployed at	View Troubleshooting details for the policy to review the error message. Check device

Deployment Status	Definition	Suggested Action
	GigaVUE-FM level.	reachability and configuration (ports, GigaSMART operations, licenses). Fix the reported issue and deploy the Traffic Policy again.
Pending	The Traffic Policy deployment request is in progress or waiting for device feedback. GigaVUE-FM has started the deployment but the final status from all devices is not yet available.	Wait for the deployment to complete. If the status remains Pending for an extended period, verify device connectivity and job status in GigaVUE-FM, then redeploy the Traffic Policy if needed. To redeploy the Traffic Policy, use Resubmit All option from the Actions menu.
Partial Success	The Traffic Policy was deployed successfully on some devices or components, but failed on others—for example, some maps or logical components (vPorts, GigaSMART operation chains) could not be created or synchronized.	View Troubleshooting details for the policy to review the error message. Review per-device or per-component errors, and correct the failing resources (such as fixing ports, GigaSMART operation resources, or missing maps). Redeploy the Traffic Policy and confirm it moves to Success. To redeploy the Traffic Policy, use Resubmit All option from the Actions menu.
Undeployed	The Traffic Policy configuration exists in GigaVUE-FM but is not currently deployed on any devices, or it has been explicitly undeployed from the devices. No traffic is being steered by this policy.	This status applies to policies that have never been deployed which means the policy configuration is saved in GigaVUE-FM but is not applied on the device. If you want the policy to be active, select the Traffic Policy and choose Deploy from the Actions menu or from the canvas. If the policy should remain inactive, no action is required.
Damaged	GigaVUE-FM detected that the deployed state of the Traffic Policy does not fully match the intended configuration—for example, some underlying maps or components are missing or	Use View Troubleshoot and Deployment Status to identify missing or inconsistent components. Resolve the underlying issues (for example, recreate missing maps or fix device problems) and redeploy the Traffic Policy. If the policy remains Damaged and is no longer needed, you can delete it and recreate a clean policy.

Deployment Status	Definition	Suggested Action
	out of sync during/after deployment or configuration sync.	

The Map Details Grid includes a Three-dots menu that enables you to edit, delete, view map details, view map statistics, and troubleshoot the Traffic Policy if there are any errors. Refer to Manage Traffic Policy for more information.



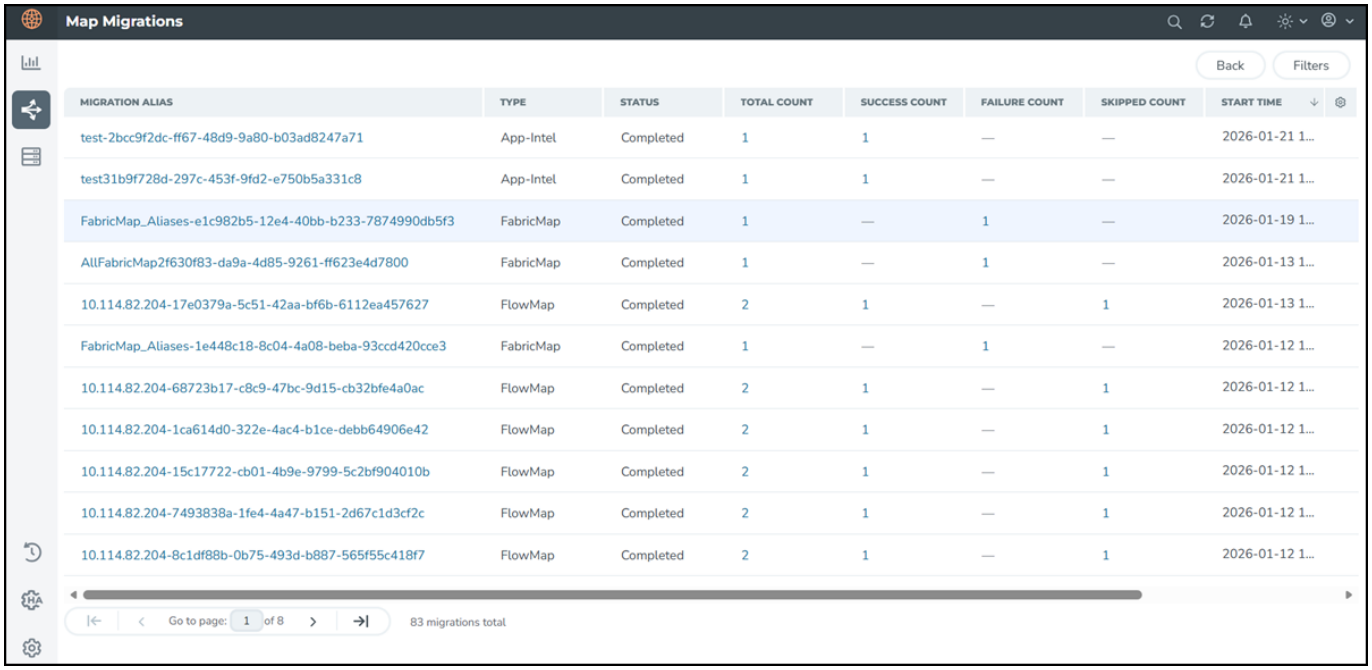
Menu

The Menu includes the following functions:

- View Migration Results
- New Map
- Actions
- Filters
- Settings

View Migration Results

Select this to view the migration status of the existing fabric maps to traffic policies.



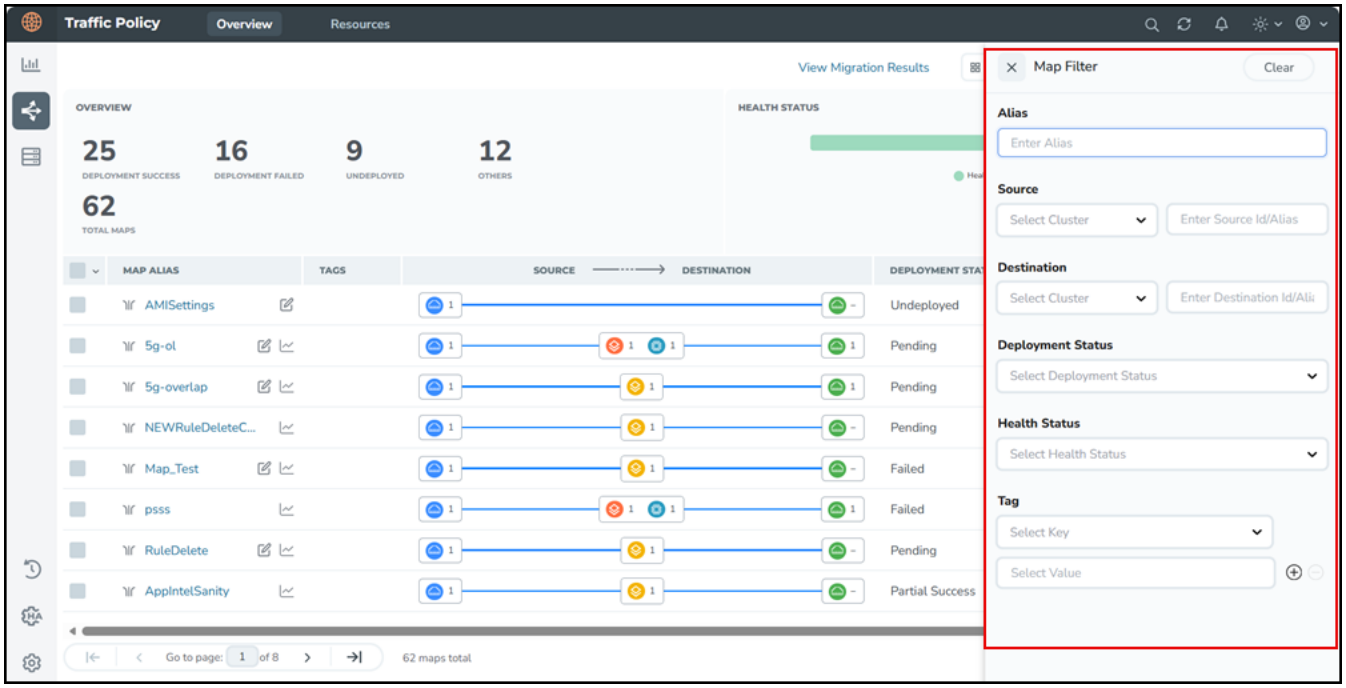
Refer to Migrate Maps to Traffic Policy for more information.

Filters

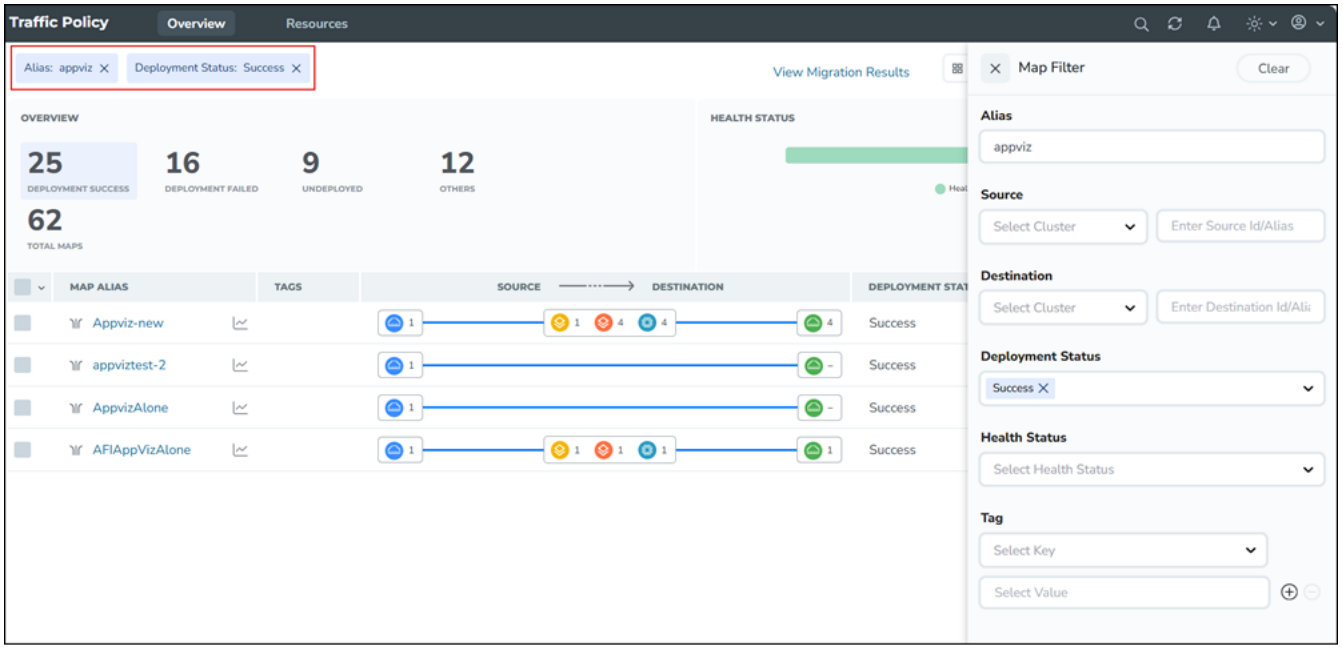
You can filter the maps based on the Alias, deployment status, health status, and tags to search for a specific map. The

To create a map filter:

1. Click **Filters** in the top right corner of the landing page. The Map Filters page appears.



2. Enter or select the required fields to filter the maps. The landing page displays the maps based on the filter criteria, and the selected filter criteria in the top left corner as shown below.



NOTE: The Alias field supports multiple entries separated by comma, and supports regex-based filtering of traffic policies.

3. Select **Clear** to reset the filter.

NOTE:
If you create a filter and log out of the GigaVUE-FM, then when you login back, the Traffic Policy landing page displays the existing maps based on the filter.
If you want to filter the traffic policy of a specific cluster, then enter the cluster ID in the Source or Destination field.

Settings

You can configure the map settings to control how resources, spine source maps, configuration synchronization, and circuit ID ranges are handled for traffic flows. Refer to Configure Map Settings topic for more information.

Actions

Select the map that you want to edit. delete, clone, deploy, undeploy, and so on and click the Actions dropdown to view the applicable options for the Traffic Policy.

Config Audit

This option enables you to conduct an audit of the map configuration. This process audits the configuration including the components of the Traffic Policy. Config audit verifies that the map configuration is active, the devices are connected, and traffic is flowing through the map with no issues.

To configure audit:

1. Select the map that you want to audit.
2. Go to **Actions > Config Audit**. Once the audit is completed, a success message displays as shown below.

Resubmit All

Resubmit All is a bulk retry option used for all outstanding config operations in the unified Traffic Policy framework. This option is used after you fix any deployment issues (for example, device unreachable, missing circuit/GigaStream, validation errors), to push or retry earlier failed or pending changes to the devices without changing the policy definitions themselves.

Edit

You can edit a map and it's components including the configurations of source, source ruleset, application, application ruleset and destination based on the requirement before deploying the map.
You can edit a map in two ways:

1. From the Menu, click **Actions > Edit** or
2. Select the map to edit then click the **Three-dots Menu** ... **> Edit**.

> **NOTE:** If you do not deploy the edited version of the map, and come back to the landing page, the Map Details Grid displays an edit icon in the Map Alias column as shown below.
>
> 
>
>
> You can always edit either the deployed policy or the policy in draft mode.

Delete

You can delete a map and it's components including the configurations of source, source ruleset, application, application ruleset and destination based on the requirement before deploying the map. The delete operation removes the configuration from the device and GigaVUE-FM.
You can delete a map in two ways:

1. From the Menu, click **Actions > Delete** or

2. Select the map to delete then click the **Three-dots Menu** ... **> Delete**.
3. When you select multiple maps to delete, a Delete Map dialog appears.
4. Enter Delete in the Type Delete field to delete the selected maps.

Clone

You can clone a map and it's components based on the requirement.
You can clone a map in two ways:

1. From the Menu, click **Actions > Clone** or
2. Select the map to delete then click the **Three-dots Menu** ... **> Clone**. The Clone Map page appears.
3. Edit and rename the map in the **Alias** field.
4. Click **Save**. The cloned map opens in the map canvas.

Deploy

To deploy a map

1. From the Menu, click **Actions > Deploy** or
2. Select the map to deploy. The selected map appears in the map canvas with the **Deploy** button.
3. Click **Deploy**. The map is deployed successfully.

Undeploy

The undeploy operation removes the configuration only from the device and the configuration exists in GigaVUE-FM.

To undeploy a map

- From the Menu, click **Actions > Undeploy** The map is undeployed successfully.

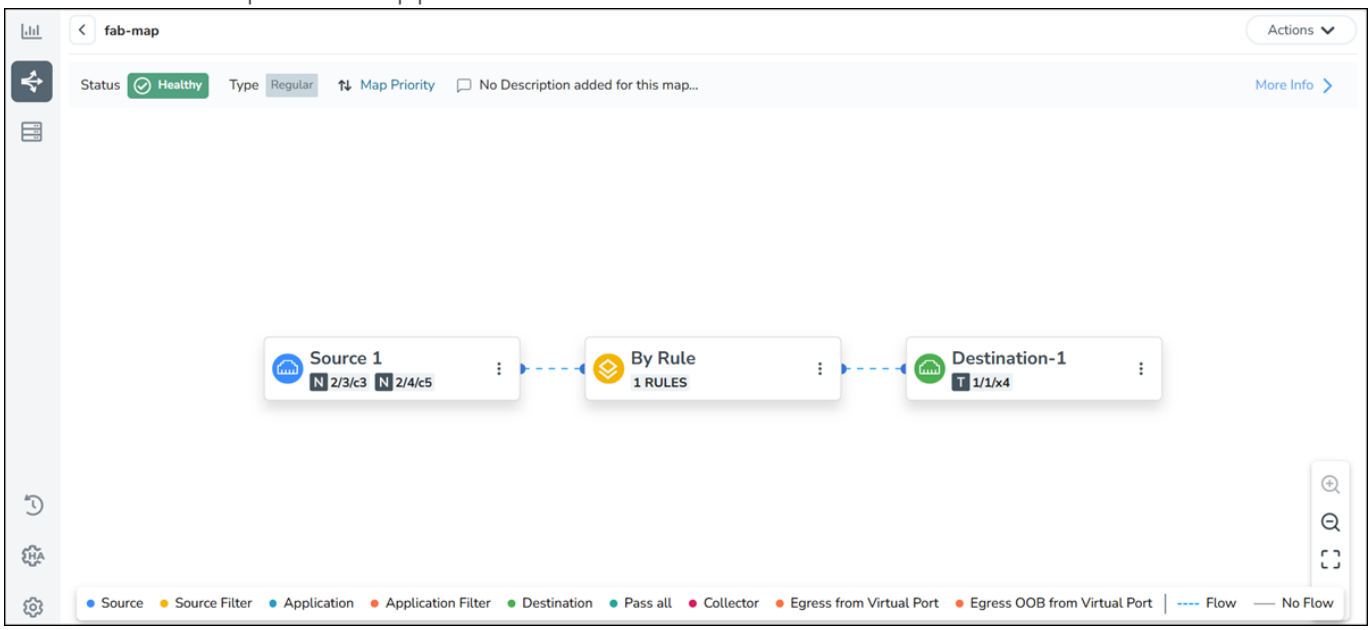
View Details

You can view the details of the Traffic Policy to view the map components, configuration, status, map type, map priority, and tags associated with the policy.

To view the details:

1. Click the map alias or

2. Select the map you want to view the details, click **Three-dots menu** ⋯ **> View Details**
   The selected map details appears as shown below.



3. Click **More Info** to view the details of the map components in the policy as shown below.



The Overview tab displays the general map information, such as Alias, description, and timestamps. You can select the other tabs namely Sources, Ruleset, Applications, Application Ruleset, and Destinations to view their details.

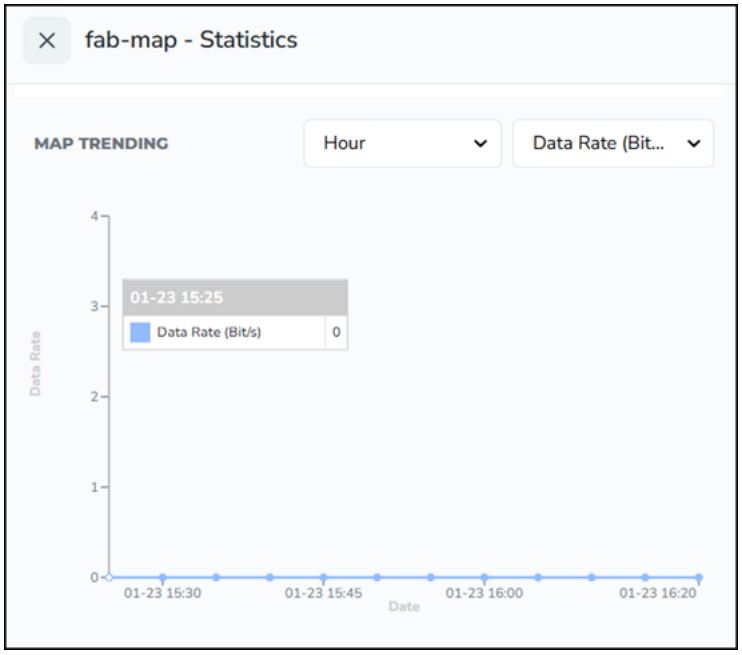
NOTE:
In the View mode, you can select **Edit** from the **Actions** dropdown to edit the Traffic Policy if required. When you click the Map Alias to view the map, the deployed version of the map appears.

View Statistics

This option helps you to quickly verify that a Traffic Policy is actively passing traffic and to spot drops, spikes, or anomalies without going to device-level CLIs or separate tools.

To view the statistics information:

1. Select the map and click **Three-dots menu** ... **> View Statistics** The Traffic Policy statistics graph appears as shown below.



2. Hover over the statistics graph to view the Data Rate/Packet rate values at any point of time.
3. You can change the X-axis and Y-axis values of statistics graph from the list as required. The statistics graph displays the Data Rate/Packet rate values at the selected time.



> **NOTE:** You can also select the Statistics , to view the statistics information of the Traffic Policy. This is not applicable to maps with deployment status as Undeployed.

View Troubleshoot

Click **Three-dots menu > View Troubleshoot** to view the troubleshooting details of the Traffic Policy. The selected traffic policy appears in the troubleshooting mode. Select the map to troubleshoot. Refer to Troubleshoot Traffic Policy for more information on troubleshooting steps.

> **NOTE:** You can view this option for all the traffic policies with deployment status other than Undeployed.

Resources Tab

The Resources tab of the Traffic Policy includes Fabric Port Group and Tool Templates sections. You can create and configure fabric port groups, export or import tool templates in the Resources tab. The following sections explain these in detail:

- Fabric Port Groups
- Tool Templates

Fabric Port Groups

A Fabric Port Group (FPG) can be defined as a logical collection of tool ports that reside on different clusters. Each of the tool port in the Fabric Port Group receives load-balanced traffic. Thus, allowing load balancing between tool ports that reside on different clusters. However, you can also use FPG within a cluster.

To know more information on creating, managing and configuring fabric port groups for a Traffic Policy refer to Create and Manage Fabric Port Group.

Tool Templates

Tool templates are pre-built profiles that define a set of applications including the metadata attributes to export to your monitoring tools. You can select the required tool template when configuring a Traffic Policy for an Application Intelligence solution, and reuse it across multiple exporters and devices.

To know more information on creating customized tool templates, refer to Create and Manage Custom Tool Template section.

Create and Manage Fabric Port Group

Create Fabric Port Group

To create fabric port groups:

1. Go to **Traffic** > **Traffic Policy** > **Resources**. The **Resources** tab opens.
2. Click **New** in the top right of the Fabric Port Group section. The New Fabric Port Group page appears.
3. **Alias** - Enter a name (FPG1) for the port group.
4. (Optional) **Description** - Enter a description for the port group.
5. **Smart Load Balancing** option is enabled by default.

   NOTE: You cannot disable the Smart Load Balancing option. You can use fabric port group in fabric maps only if this option is enabled.
6. Select the required ports in the Port Editor pane to create the port group.
7. Enter the weights for each port.
8. (Optional) **Tags** - Select the required tag key and tag value combination from the existing tags to associate the FPG to tags.
9. Click **Create Tags** to create a new tag. The Create Tags pane appears.
10. Do the following in the Create Tags pane. Refer to Tags to know more about tag creation, tag types, and working with tags.
    a. **Name** - Enter a name for the tag.
    b. (Optional) **Description** - Enter a description for the tag.
    c. **Multi Valued** - Turn on the toggle to create a multi-valued tag. If a tag is multi-valued, then the resources can be assigned multiple tag values.
    d. **Tag Type** - Select the required type from the list. Valid values are Rbac and Aggregation.
    e. **Hierarchy** - This option is enabled by default.

    NOTE: If Hierarchy is enabled, for any node or cluster that is tagged, the ports and slots belonging to the node or cluster will automatically inherit the tags.

    f. **Tag Values** - Enter the tag values separated by comma.

g. Click **Save**. The created tag appears in the New Fabric Port Group page.

11. Click **Apply**. The created fabric port group appears in the Fabric Port Group section.

Edit Fabric Port Groups

To edit the fabric port groups:

1. Select the port group that you want to edit, click **Actions > Edit**. The Edit Fabric Port Group page appears.
2. Make the required changes.
3. Click **Apply**. The fabric port group is updated successfully.

Delete Fabric Port Groups

To delete the fabric port groups:

1. Select the port group that you want to delete, click **Actions > Delete**. The confirmation dialog appears.
2. Click **Ok**. The fabric port group is deleted successfully.

> **NOTE:** You cannot edit or delete a fabric port group that is already in use by a Traffic Policy.

Export Fabric Port Groups

To export fabric port groups:

1. To export all the port groups, click **Export > Export All > CSV/XLSX**. The fabric port groups are exported in the CSV or XLSX format.
2. To export the selected port groups, select the required port groups and click **Export > Export Selected > CSV/XLSX**. The selected fabric port groups are exported in CSV or XLSX format.

Configure Traffic Policy using Fabric Port Groups

To configure Traffic Policy using fabric port groups, create fabric port groups based on the instructions in the previous section.

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.

2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure a First-Level Map:

To configure a first-level map:

1. Access Traffic Policy
   a. In the Traffic Policy landing page, click **New Map**. The New Map pane appears.
   b. Create a new map with **Alias** as map1.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network ports 1/1/x3 and 1/1/x4 from different clusters.

   > **NOTE:** To configure or modify the port type, use the **Quick Port Editor**.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.

a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.

b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle to pass the IPv4 traffic.

d. Select IP Version from the **Condition** list, and then select V1 in the **Version** field.

e. Click **Save**. The added rule appears in the **Source Ruleset** page.

f. Click **Close**. The By Rule block appears on the canvas with the configured rule.

4. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears.

b. Select the configured virtual port vp1.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

5. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

6. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure a Second-Level Map:

To configure a second-level map:

1. Access Traffic Policy

a. In the Traffic Policy landing page, click **New Map**. The New Map pane appears.

b. Create a new map with **Alias** as map2.

2. Select Source

a. Click **Sources**. The **Source 1** page appears with the list of available ports.

b. Select the virtual port vp1.

c. Click **Add**. The Source block appears on the canvas with the selected virtual port.

3. Select Application

a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group (gs-grp) you created earlier.

f. Select Load Balancing from the **Applications** list. By default, the application you select is enabled.

g. In the Load Balancing section, select Stateful option, and type as GTP

h. Select Weighted Least Bandwidth in the **Metric** field.

i. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

4. Select Application Ruleset.

a. Expand the **Application Ruleset** component and then drag and drop a **By Rule**. The **Application Ruleset** page appears.

b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle to pass the configured traffic.

d. Select IP Version from the **Condition** list, and then Version 4 from the **Version** list.

e. Click **Save**. The **Application Ruleset** page appears with the configured rule.

f. Click **Close**. The By Rule block appears on the canvas with the configured application rule.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the fabric port group FPG1.

> **NOTE:** To configure or modify the port type, use the **Quick Port Editor**.

c. Click **Add**. The Destination block appears on the canvas with the selected tool ports.

6. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Manage Rulesets across Traffic Policies

Duplication of existing rulesets across policies simplifies operations and improves ease of use. Manage Rulesets provides the ability to copy rules from one policy (source) and apply them to another policy (target). This feature copies rules within the same ruleset type (source-to-source or application-to-application), not across types. You can copy one, multiple, or all rules from a policy. When pasting rules into a target policy, the system automatically detects duplicate rules. Duplicate rules are not replaced; instead, you are prompted to either cancel the paste operation or proceed by pasting only the non-duplicate rules. Pasted rules are saved to the system but are not deployed automatically. You must explicitly deploy the policy to apply the changes to the device.

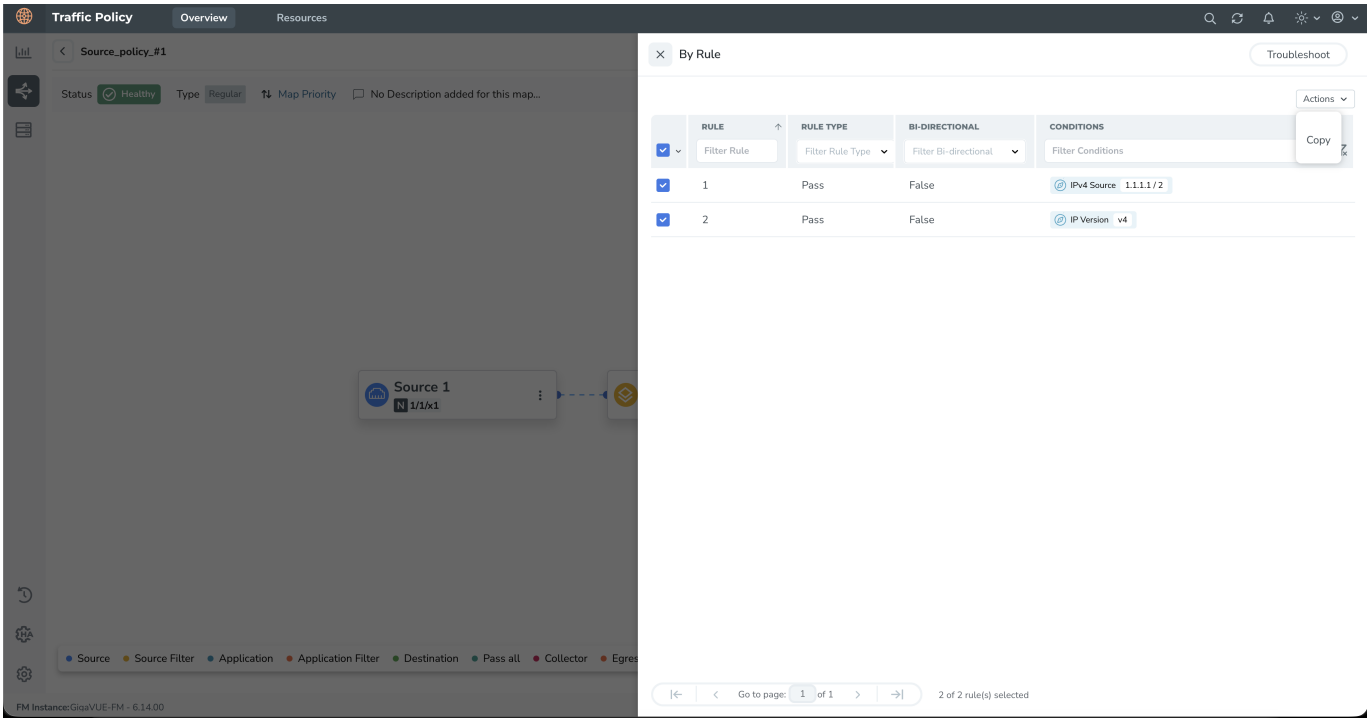
Copy Rules from a Source Policy

This section describes copy rules from a source policy, perform the following:

1. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
2. Click the policy alias to navigate to View mode.

NOTE: View mode allows the configured policies to be seen in the Traffic Policies section by selecting them, and they can also be displayed in the canvas.

3. Select the **Policy > Actions > View Details**.
4. To edit the map, click the **Three-dots Menu** ⋯ **> Edit**.
5. Select the Source Ruleset or Application Ruleset that already exists to copy the rules.
6. Use the checkboxes to select rules: choose one or more specific rules, or use the Select All option (if available) to select the entire ruleset.
7. From the **Actions** menu, click **Copy**.



8. The selected rules are stored and made available for pasting into one or more target policies.

Paste Rules into a Target Policy

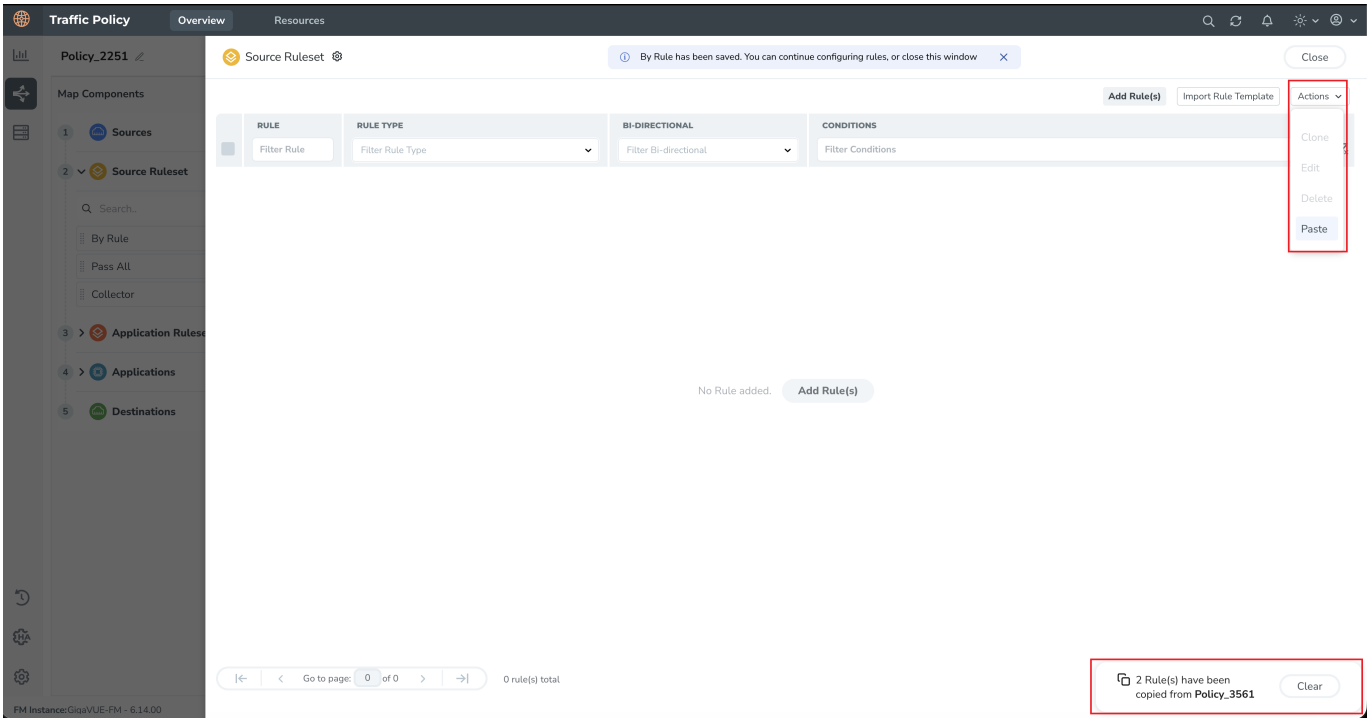
This section describes paste rules into a target policy, perform the following:

1. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears, open a target Traffic Policy where you want to reuse the copied rules.
2. To edit the map, click the **Three-dots Menu** ••• **> Edit**.



NOTE: Open a ruleset of the same type in the target policy (Source or Application) as the one you copied from.

3. From the **Actions** menu, click **Paste**.



4. If any duplicate rules are detected in the target ruleset:
   a. A dialog lists the existing rule IDs and the duplicate rule entries from the copied set.
   b. Choose one of the following:
      - **Cancel** – cancel the paste operation; no rules are added.
      - **Proceed** – skips the duplicate rules and pastes only the non-duplicate rules.
5. After you proceed, the system:
   - Inserts the non-duplicate rules into the target ruleset.
   - Allocates new rule IDs for all pasted rules in the target policy.
6. Repeat steps 1–5 for any additional target policies into which you want to paste the same copied rules.
7. Click **Deploy** to deploy the updated policies to the relevant devices.

Manage Rulesets across Traffic Policies — Rules and Notes

1. Rule-level MAC and IP rewrite options are copied, but map-level MAC and IP rewrite settings are not copied.
2. If you modify or delete rules in the source policy after copying, those changes do not affect the already copied rules.
3. The paste operation cannot be undone once the rules have been pasted.

4. You cannot copy/paste PassAll rules or Collector rules.
5. You cannot copy/paste AppIntel AFI/AMI rules.
6. You can copy rules only from a single ruleset at a time.
7. Duplicate detection covers exact duplicates only. For example, suppose the target policy already has a source rule that matches VLAN 100 and source IP 10.1.1.1/32. If you paste a rule that also matches VLAN 100 and 10.1.1.1/32 (or the equivalent mask 10.1.1.1 255.255.255.255), the system treats it as an exact duplicate and lists it in the duplicate rules dialog. However, if the target policy has a rule for VLAN range 100 – 200 and you paste a rule that matches only VLAN 150, the rules are considered overlapping but not exact duplicates, so they are not flagged by duplicate detection.
8. Copied rules persist until you clear them; they remain after logout, session timeout, or FM restart. You can use the Clear option at the bottom right of the pop-up to remove them. Copied rules are stored at the user level. If a user account is deleted, any rules copied by that user are also deleted.
9. Platform-specific qualifiers may be rejected at deployment time if the target device does not support them. For example, inner-header qualifiers supported on GigaVUE-TA400/GigaVUE-TA400E.

Configure Traffic Policy Without GigaSMART Operations

Use a regular map without GSOP when you want to forward traffic from network ports to tools without any GigaSMART processing. You filter traffic using source rules only.

Use cases

- Send all traffic from SPAN or TAP ports to a monitoring tool.
- Send specific VLANs or IP subnets to a tool without changing packets.

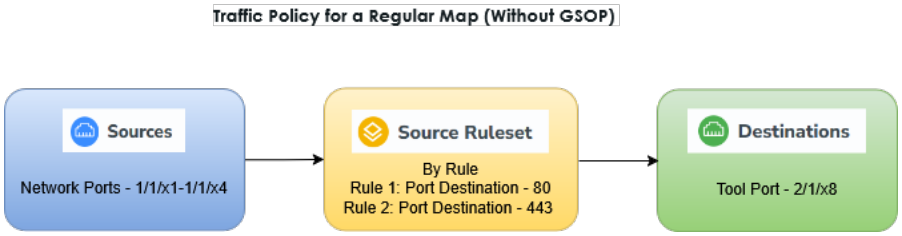
Example

Send all HTTP and HTTPS traffic from a set of SPAN ports to a packet recorder tool.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
Source	Network (SPAN) ports - 1/1/x1 – 1/1/x4
Source Ruleset	Match TCP destination ports 80 and 443
Application Ruleset	N/A
Applications	N/A
Destination	Tool port - 2/1/x8



Steps:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The New Map dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
2. Select Source.
   a. Click **Sources**. The Source 1 page appears.
   b. Select the network ports 1/1/x1 - 1/1/x4.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset

a. Expand the Source Ruleset component and then drag and drop **By Rule** component. The **Source Ruleset** page appears.

b. Click **Add Rules**. The **Add Ruleset** page appears.

c. Turn on **Pass** toggle to pass the HTTP traffic.

d. Select **Port Destination** from the **Condition** list.

e. Enter **80** in the **Min** field.

f. Click **+** icon next to the Rule Summary to add an OR rule.

g. Turn on **Pass** toggle to pass the HTTPS traffic.

h. Select **Port Destination** from the **Condition** list.

i. Enter 443 in the **Min** field.

j. Click **Save**. The configured rules appear in the **Source Ruleset** page.

k. Click **Close**. The By Rule block appears on the canvas with the configured rules.

4. Select Destination.

a. Click **Destinations**. The Destination 1 page appears with the list of available ports.

b. Select the tool port 2/1/x8.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Destination block appears on the canvas with the selected tool port.

5. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to implement the map configuration in the device.

6. Verify deployment.

a. **Status:** In the Traffic Policy landing page, confirm that the map status is deployed and healthy.

b. (Optional) **Implemented Device:** View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure VLAN Manipulation

VLAN manipulation can be configured based on map rules and maps:

- **Rule-Based**
- **Map-Based**

Refer to the VLAN Manipulation.

NOTE: If you have configured both map level and rule level VLAN manipulation functionality in the same map, then rule-based configuration takes priority.

NOTE: From 6.8 Release, VLAN manipulation add is supported in standalone mode. From 6.10 Release, VLAN manipulation add is supported in Cluster mode.

NOTE: Starting from 6.13 Release, VLAN manipulation Modify, VLAN manipulation Strip is supported only in standalone mode. It is not supported in cluster mode.

Map-Level VLAN Manipulation

The below VLAN Manipulation operations can be configured on map-level:

- **Map-Level Add VLAN Manipulation**
- **Map-Level Modify VLAN Manipulation**
- **Map-Level Strip VLAN Manipulation**

Refer to the VLAN Manipulation.

To configure VLAN manipulation on map-level,

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   d. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
2. Select Source.
   a. Click **Sources**. The Source 1 page appears.
   b. Select the network ports 1/1/x1 and 1/1/x2.

   NOTE: To configure or modify the port type, use the **Quick Port Editor**.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.

3. Select **Source Ruleset**.
   a. Expand the Source Ruleset component and then drag and drop **By Rule/Collector** component. The **Source Ruleset** page appears.
   b. Click  near the **Source Ruleset**. The **Map Configuration** page appears.
   c. If you want to add map-level VLAN manipulation configuration, go to the **VLAN** section, and select one of the below options.
   - Select **Add**, from the **VLAN Action** list.
     - Enter the **VLAN ID** value between 1 and 4095.
     - In the **Tag Protocol Id** drop-down list, select the **TPID** value for the VLAN Tag. The default value is 0x8100, but you can also select the other supported values 0x9100 and 0x88a8.
   - Select **Modify**, from the **VLAN Action** list.
     - Enter the **VLAN ID** value between 1 and 4095.
     - In the **Tag Protocol Id** drop-down list, select the **TPID** value for the VLAN Tag. The default value is 0x8100, but you can also select the other supported values 0x9100 and 0x88a8.
   - Select **Strip**, from the **VLAN Action** list.
4. Click **Save** to complete the map level VLAN manipulation configuration.
5. Click **Add Rule(s)**. The **Add Ruleset** page appears.
   d. For example, select **IPv4 Source** from the **Conditions** list.
   e. Enter the **IPv4 Address** and **Cidr** or **Net Mask** values.
   f. If you want to add rule level VLAN manipulation configuration, click  **Settings**.
      - In the Rule Configuration dialog, in the VLAN section, select **VLAN Manipulation Add/VLAN Manipulation Modify** from the VLAN Action list.
      - For Add/Modify you need to specify the VLAN ID and TPID.
      - Click **Save**.
   g. After adding the rule conditions, click **Save** in the Add Ruleset page.
   h. Click **Close**. The By Rule block appears on the canvas with the configured source ruleset.
6. Select Destination.
   a. Click **Destination**. The Destination 1 page appears with the list of available ports.
   b. Select the tool port 2/1/x10.

> **NOTE:** To configure or modify the port type, use the **Quick Port Editor**.

c. Click **Add**. The Destination block appears on the canvas with the selected tool port.

7. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
8. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure Traffic Policy for Traffic Intelligence Solutions

Traffic intelligence solutions are a set of GigaSMART operations that optimize and streamline Layer 2–7 network traffic before it reaches your security, monitoring, and observability tools. They filter, de-duplicate, slice, and intelligently distribute traffic so tools see only the right data at the right volume, improving performance and helping contain tool costs while maintaining complete visibility.

Refer to the following sections for instructions on how to configure Traffic Policy for traffic intelligence solutions with real-time examples:

- GigaSMART Adaptive Packet Filtering-(APF)
- GigaSMART Load Balancing
- GigaSMART De-Duplication
- GigaSMART Masking
- GigaSMART NetFlow Generation
- GigaSMART Slicing
- GigaSMART Tunneling Operations

GigaSMART Adaptive Packet Filtering (APF)

Required License: Adaptive Packet Filtering

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus, GigaVUE-HC1 Plus Gen3, and GigaVUE-HCT Gen3.

This section explains the Traffic Policy configurations of APF for different real time examples.

- Identify Social Security Numbers in User-Level Transactions
- Mask Social Security Numbers
- Filter on Fiber Channel over Ethernet (FCOE) Traffic
- Filter on Subscriber Device IP (User-Endpoint IP or UE-IP)
- Filter on Inner Layer 2-4 Parameters for Unrecognized Headers
- GTP Tunnel ID-Based Filtering
- ERSPAN Tunneling
- Distribute Traffic Based on Inner IP Addresses and Inner TCP Port Values
- MPLS Label Based Filtering
- Conditional Header Stripping
- Facilitate Overlapping Rules

Identify Social Security Numbers in User-Level Transactions

This topic describes the workflow required to identify the Social Security Numbers in user level transactions using APF.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port - 1/1/x1.
Source Ruleset	By Rule • Condition - IP Version. • Version - Version 4.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port - vp1
Second-Level Map	
Source	Virtual port - vp1

Map Components	Sample Configuration
Source Ruleset	N/A
Application	• My GSOP - Adaptive Packet Filtering
Application Ruleset	• Condition - Pattern Match • Type - Regex and Value-\d{3}-?\d{2}-?\d{4} . • Offset Start - 40 and Offset End- 8
Destination	Tool ports- 1/1/x3 and 1/1/x4

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.

   NOTE: For APF, no GigaSMART parameters are required unless combined with other GigaSMART operations.

7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure Traffic Policy for a First-Level Map

To create a first-level map to forward traffic from network port 1/1/x3 to virtual port vp1:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The traffic policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the IPv4 traffic.
   d. Select IP Version from the **Condition** list, and then select Version 4 from the **Version** list.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears.

b. Select the configured virtual port vp1.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

5. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

6. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure Traffic Policy for a Second-Level Map

To create a second-level map to forward traffic from the virtual port vp1 to GigaSMART for pattern matching:

1. Access Traffic Policy.

a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.

c. In the **New Map** dialog, do the following:

- **Alias** - Enter map2.
- (Optional) **Description** - Enter the description for the Traffic Policy.
- (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
- (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.

d. Click **Save**. The traffic policy canvas appears.

2. Select Source.

a. Click **Sources**. The **Source 1** page appears with the list of available ports.

b. Select the virtual port vp1.

c. Click **Add**. The Source block appears on the canvas with the selected virtual port.

3. Select Application.

a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. You can either select an existing GSOP or create a new GSOP by clicking **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group (gs-grp) you created earlier.

f. Select Adaptive Packet Filtering from the **Applications** list. By default, the application you select is enabled.

g. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

4. Select Application Ruleset.

a. Expand the **Application Ruleset** component and then drag and drop a **By Rule**. The **Application Ruleset** page appears.

b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle to pass the configured traffic.

d. Select Pattern Match from the **Condition** list, and then select Regex in the **Type** list.

e. Enter \d{3}-?\d{2}-?\d{4} in the **Value** field.

f. Enter **Offset Start** value as 40 and **Offset End** value as 8.

> **NOTE:** The Offset Start and End values are applicable only for Gen 2 GigaSMART modules.

g. Click **Save**. The **Application Ruleset** page appears with the configured rule.

h. Click **Close**. The By Rule block appears on the canvas with the configured application rule.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool ports 1/1/x3 and 1/1/x4.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Destination block appears on the canvas with the selected tool ports.

6. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Mask Social Security Numbers

This topic describes the workflow required to mask the Social Security Numbers in user level transactions using APF.

NOTE: In the following pattern matching example, IPv4 packets contain Social Security Numbers (SSNs) in the format xxx-xx-xxxx. If the SSNs are between offset 40 and 80, they will be replaced with zeros. This example is not applicable to Gen 3 GigaSMART module.

Ensure that a GigaSMART Operation using the correct GigaSMART generation (Gen2 or Gen3) is created and associated with the vPort you will use in this Traffic Policy. The APF pattern-match fields available in Traffic Policy are validated against the GSOP generation attached to the vPort.

NOTE: In Traffic Policy, APF pattern-match (regex) fields are validated against the GSOP generation (Gen2 or Gen3) associated with the vPort, not just the source vPort selection. Attach the appropriate GSOP to the vPort before you create the Application Ruleset and APF pattern-matching rule.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port- 1/1/x1
Source Ruleset	By Rule • Condition - IP Version and Version - Version 4
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port - vp1
Second-Level Map	
Source	Virtual port - vp1
Source Ruleset	N/A

Map Components	Sample Configuration
Applications	My GSOP - Adaptive Packet Filtering .
Application Ruleset	By Rule • Condition - Pattern Match, Type - Regex and Value - \d{3}-?\d{2}-?\d{4} • Offset Start- 40 and Offset End - 80. • Mask - 0
Destination	Tool port - 1/1/x6

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.

   NOTE: For APF, no GigaSMART parameters are required unless combined with other GigaSMART operations.

7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

NOTE: If you later change the Source from one vPort to another after configuring this APF pattern-matching rule, the existing pattern-match configuration is not automatically reset for the new vPort or GSOP. Review and update or recreate the APF pattern-matching rule before you deploy, to ensure the selected fields are valid for the new GSOP generation and to avoid deployment-time validation errors.

Configure Traffic Policy for a First-Level Map

To create a first-level map to forward traffic from network port 1/1/x1 to virtual port gsTraffic:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset
   a. Expand the Source Ruleset component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the IPv4 traffic.
   d. Select IP Version from the **Conditions** list, and then select Version 4 from the **Version** list.

e. Click **Save**. The added rule appears in the **Source Ruleset** page.

f. Click **Close**. The By Rule block appears on the canvas with the configured rule.

4. Select Destination

a. Click **Destinations**. The **Destination 1** page appears.

b. Select the configured virtual port vp1.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

5. Deploy Traffic Policy

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

6. Verify deployment

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure Traffic Policy for a Second-Level Map

To create a second-level map to forward traffic from the virtual port gsTraffic to GigaSMART with pattern matching:

1. Access Traffic Policy.

a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

b. Create a new map with **Alias** as map2.

2. Select Source.

a. Click **Sources**. The **Source 1** page appears with the list of available ports.

b. Select the virtual port vp1.

c. Click **Add**. The Source block appears on the canvas with the selected virtual port.

3. Select Application.

a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group (gs-grp) you created earlier.

f. Select Adaptive Packet Filtering from the **Applications** list. By default, the application you select is enabled.

g. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

4. Select Application Ruleset.

a. Expand the Application Ruleset component and then drag and drop a **By Rule**. The **Application Ruleset** page appears.

b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle to pass the configured traffic.

d. Select Pattern Match from the **Conditions** list, and then select Regex in the **Type** list.

e. Enter \d{3}-?\d{2}-?\d{4} in the **Value** field.

f. Enter **Offset Start** value as 40 and **Offset End** value as 80.

> **NOTE:** The Offset Start and End values are applicable only for Gen 2 GigaSMART modules.

g. Enter 0 in the **Mask** field.

h. Click **Save**. The **Application Ruleset** page appears with the configured rule.

i. Click **Close**. The By Rule block appears on the canvas with the configured application rule.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool port 1/1/x6.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Destination block appears on the canvas with the selected tool ports.

6. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Filter on Fiber Channel over Ethernet (FCOE) Traffic

This topic describes the workflow required to filter FCOE traffic using APF. In this example, regular expression-based filters are used for filtering on the source address in a Fiber Channel header. The flexibility offered by regular expression-based filters can be used as an infrastructure to classify traffic streams with protocol headers that are typically unsupported on traditional TAP/SPAN aggregation devices.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port- 1/1/x1
Source Ruleset	By Rule • Condition - IP Version and Version - Version 4
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port - vp1
Second-Level Map	
Source	Virtual Port - vp1
Source Ruleset	N/A
Applications	My GSOP - Adaptive Packet Filtering
Application Ruleset	By Rule • Condition - Pattern Match, Type - Regex and Value - txff\xff\xfe • Offset Start - 0 and Offset End-29
Destination	Tool port - 1/1/x2

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.

2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.

   NOTE: For APF, no GigaSMART parameters are required unless combined with other GigaSMART operations.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure Traffic Policy for a First-Level Map

To configure a first-level map to forward traffic from network port 1/1/x1 to virtual port vp1:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.

- (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
- (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.

d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the FCOE traffic.
   d. Select Ether Type from the **Condition** list, and then enter 8906 in the **Value** field.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears.
   b. Select the configured virtual port vp1.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure Traffic Policy for a Second-Level Map

To configure a second-level map to forward traffic from the virtual port to GigaSMART with pattern matching:

1. Access Traffic Policy
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Create a new map with **Alias** as map2.
2. Select Source
   a. Click **Sources**. The **Source 1** page appears with the list of available ports.
   b. Select the virtual port vp1.
   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
   e. Select the GigaSMART group (gs-grp) you created earlier.
   f. Select Adaptive Packet Filtering from the **Applications** list. By default, the application you select is enabled.
   g. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.
4. Select Application Ruleset.
   a. Expand the **Application Ruleset** component and then drag and drop a **By Rule**. The **Application Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the configured traffic.
   d. Select Pattern Match from the **Condition** list, and then select Regex in the **Type** list.
   e. Enter txff\xff\xfe in the **Value** field.
   f. Enter **Offset Start** value as 0 and **Offset End** value as 29.

> **NOTE:** The Offset Start and End values are applicable only for Gen 2 GigaSMART modules.

g. Click **Save**. The **Application Ruleset** page appears with the configured rule.

h. Click **Close**. The By Rule block appears on the canvas with the configured application rule.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool port 1/1/x2.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Destination block appears on the canvas with the selected tool ports.

6. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Filter on Fiber Subscriber Device IP (User-Endpoint IP or UE-IP)

This topic describes the workflow required to filter subscriber device IP using APF.
In this example, we are:
- Identifying and forwarding traffic from / to a UE-IP of 1.1.1.1 to a monitoring tool connected to 1/1/x1
- Identifying and forwarding traffic from / to a UE-IP of 1.1.1.2 to a different monitoring tool connected to tool port 1/1/x4

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map1	
Source	Network port- 1/1/x1
Source Ruleset	By Rule • Condition - Port Source, Value - 2152.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port (vp1)
First-Level Map2	
Source	Network port- 1/1/x3.
Source Ruleset	By Rule • Condition - Port Source, Value - 2152.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port (vp1)
Second-Level Map1	
Source	Virtual Port (vp1)
Source Ruleset	N/A
Applications	My GSOP - Adaptive Packet Filtering
Application Ruleset	By Rule • Condition - Port Destination and Min-1.1.1.1. • Net Mask - 255.255.255.255 and Position - 2
Destination	Tool ports- 1/1/x1 and 1/1/x4
Second-Level Map2	
Source	Virtual Port (vp1)
Source Ruleset	N/A
Applications	My GSOP - Adaptive Packet Filtering
Application Ruleset	By Rule • Condition - Port Source and Min-1.1.1.2. • Net Mask - 255.255.255.255 and Position - 2
Destination	Tool port - 1/1/x4

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.

3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure Traffic Policy for First-level Maps

To create a first-level map to forward GTP-u traffic to virtual port vp1:
Map 1

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.

d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   > **NOTE:** To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle and select **Bi-directional** option.
   d. Select Port Source from the **Conditions** list, and then enter 2152 in the **Value** field.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears.
   b. Select the configured virtual port vp1.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

To create a first-level map to forward GTP-c traffic to tools:

Follow the same steps in above topic to create another first-level map with the following changes:

1. Create the map with

- Source - Network Port 1/1/x1
- Source Ruleset - By Rule
  - Condition - Port Source and Value - 2152
- Destination - Tool Ports 1/1/x3 and 1/1/x4

2. Deploy and verify the Traffic Policy.

Configure Traffic Policy for Second-Level Maps

To create a second-level map to filter on source and destination IP (bi-directional)

Map 1:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Create a new map with **Alias** as map1.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears with the list of available ports.
   b. Select the virtual port vp1.
   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application.
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
   e. Select the GigaSMART group (gs-grp) you created earlier.
   f. Select Adaptive Packet Filtering from the **Applications** list. By default, the application you select is enabled.
   g. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.
4. Select Application Ruleset.
   a. Expand the Application Ruleset component and then drag and drop a **By Rule**. The **Application Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle to pass the configured traffic.

d. Select IPv4 Destination from the **Condition** list, and then enter 1.1.1.1 in the **Min** field.

e. Enter 255.255.255.255 in the **Net Mask** field and 2 in the **Position** field.

f. Click **Save**. The **Application Ruleset** page appears with the configured rule.

g. Click **Close**. The By Rule block appears on the canvas with the configured application rule.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears.

b. Select the tool port 1/1/x3 and 1/1/x4.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

6. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the map with
   - Source port - Virtual Port vp1
   - Destination port - Tool Port 1/1/x4
   - Application ruleset - By Rule
     - Condition - IPv4 Source, IPv4 Address - 1.1.1.2
     - Net Mask - 255.255.255.255, Position - 2
2. Deploy and verify the Traffic Policy.

Filter on Inner Layer 2-4 Parameters for Unrecognized Headers

This topic describes the workflow required to filter on Inner Layer 2-4 parameters for unrecognized headers using APF.

In this example, we are:

- Identifying and forwarding traffic from/to ip 1.1.1.1 in the inner / original packet to monitoring tool connected to tool port 1/1/x1
- Identifying and forwarding traffic from/to ip 1.1.1.2 in the inner / original packet to monitoring tool connected to tool port 1/1/x4

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port - 1/1/x1.
Source Ruleset	By Rule • Condition - Port Source, Value-2152.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port - vp1
Second-Level Map	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	• My GSOP - Adaptive Packet Filtering.
Application Ruleset	• Condition - IPv4 Source, Min-1.1.1.1. • Net Mask-255.255.255.255, Position-1.
Destination	Tool ports - 1/1/x3 and 1/1/x4

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.

2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.

   NOTE: For APF, no GigaSMART parameters are required unless combined with other GigaSMART operations.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure Traffic Policy for a First-Level Map

To create a first-level map to forward fabric path packets to the virtual port:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.

- (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
- (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.

d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

> **NOTE:** To configure or modify the port type, use the **Quick Port Editor**.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle and select **Bi-directional** option.
   d. Select Ether Type from the **Conditions** list, and then enter 8903 in the **Value** field.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears.
   b. Select the configured virtual port vp1.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure Traffic Policy for Second-Level Maps

Map 1:

To create a second-level map to filter on source and destination IP (bi-directional):

1. Access Traffic Policy
    a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
    b. Create a new map with **Alias** as map1.
2. Select Source.
    a. Click **Sources**. The **Source 1** page appears with the list of available ports.
    b. Select the virtual port vp1.
    c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application.
    a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
    b. Select either a standalone node or a cluster from the **Select Node** list.
    c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
    d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
    e. Select the GigaSMART group (gs-grp) you created earlier.
    f. Select Adaptive Packet Filtering from the **Applications** list. By default, the application you select is enabled.
    g. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.
4. Select Application Ruleset.
    a. Expand the Application Ruleset component and then drag and drop a **By Rule**. The **Application Ruleset** page appears.
    b. Click **Add Rules**. The **Add Ruleset** pane appears.
    c. Turn on **Pass** toggle to pass the configured traffic.
    d. Select IPv4 Source from the **Conditions** list, and then enter 1.1.1.1 in the **Min** field.
    e. Enter 255.255.255.255 in the **Net Mask** field and 1 in the **Position** field.
    f. Click **Save**. The **Application Ruleset** page appears with the configured rule.
    g. Click **Close**. The By Rule block appears on the canvas with the configured application rule.
5. Select Destination.
    a. Click **Destinations**. The **Destination 1** page appears.

b. Select the tool port 1/1/x3.

> **NOTE:** To configure or modify the port type, use the **Quick Port Editor**.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

6. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
7. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the map (map2) with
   - Source - Virtual port (vp1)
   - Destination - Tool port (1/1/x4)
   - Application ruleset - By Rule
     - Condition - IPv4 Destination and Min - 1.1.1.2
2. Deploy and verify the Traffic Policy.

GTP Tunnel ID-Based Filtering

The following example demonstrates filtering and forwarding traffic based on tunnel IDs included as part of the GTP user-plane messages using APF. It also illustrates the concept of a shared collector to which traffic not matching any of the configured filters can be optionally sent. GTP control sessions are forwarded to all the monitoring tools leveraging the power of flow mapping by filtering on Layer-4 UDP port 2123.

For GTP-u:

- Filter and forward teid ranges 0x001e8480..0x001e8489 to a monitoring tool
- Filter and forward teid ranges 0x001e8490..0x001e8499 to another monitoring tool
- Forward the rest of the traffic to a shared collector

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map1	
Source	Network port- 1/3/x9
Source Ruleset	By Rule • Condition - Port Source and Value-2152.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port - vp1
First-Level Map2	
Source	Network port- 1/3/x9
Source Ruleset	By Rule • Condition - Port Source and Value-2123.
Application Ruleset	N/A
Applications	N/A
Destination	Tool ports - 1/3/x13, 1/3/x15
Second-Level Map1	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	My GSOP - Adaptive Packet Filtering
Application Ruleset	By Rule • Condition - GTP-U TEID, Min - 0x001e8480, Max - 0x001e8489, and Subset - None.
Destination	Tool port- 1/3/x15
Second-Level Map2	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	My GSOP - Adaptive Packet Filtering
Application Ruleset	By Rule • Condition - GTP-U TEID, Min - 0x001e8490, Max - 0x001e8489, and Subset - None.
Destination	Tool port- 1/3/x15
Second-Level Map3	
Source	Virtual port - vp1
Source Ruleset	N/A

Map Components	Sample Configuration
Applications	N/A
Application Ruleset	Collector
Destination	Tool port- 1/3/x14

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.

   NOTE: For APF, no GigaSMART parameters are required unless combined with other GigaSMART operations.

7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure Traffic Policy for a First-Level Map

Map 1

To create a first-level map to forward GTP-u traffic to virtual port vp1

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter GTP-u.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/3/x9.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass GTP-u traffic.
   d. Select Ether Type from the **Conditions** list, and then enter 2152 in the **Value** field.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears.

b. Select the configured virtual port vp1.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

5. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

6. Verify Deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2

To create a first-level map to forward GTP-c traffic to tools:

Follow the same steps in above topic to create another first-level map with the following changes:

1. Create the map (GTP-c) with
   - Source port - Network port (1/3/x9)
   - Destination ports - Tool ports 1/3/x13 and 1/3/x15
   - Source Ruleset - By Rule
     - Condition - Port Source and Value - 2123
2. Deploy and verify the Traffic Policy.

Configure Traffic Policy for Second-Level Maps

Map 1

To create a second-level map to filter GTP-u based on Tunnel ID rule match

1. Access Traffic Policy.

a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

b. Create a new map with **Alias** as m1.

2. Select Source.

a. Click **Sources**. The **Source 1** page appears with the list of available ports.

b. Select the virtual port vp1.

c. Click **Add**. The Source block appears on the canvas with the selected virtual port.

3. Select Application.
    a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
    b. Select either a standalone node or a cluster from the **Select Node** list.
    c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
    d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
    e. Select the GigaSMART group (gs-grp) you created earlier.
    f. Select Adaptive Packet Filtering from the **Applications** list. By default, the application you select is enabled.
    g. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.
4. Select Application Ruleset.
    a. Expand the Application Ruleset component and then drag and drop a **By Rule**. The **Application Ruleset** page appears.
    b. Click **Add Rules**. The **Add Ruleset** pane appears.
    c. Turn on **Pass** toggle to pass the configured traffic.
    d. Select GTP-U TEID from the **Condition** list, and then enter 0x001e8480 in the **Min** field and 0x001e8489 in the **Max** field.
    e. Set **Subset** to None.
    f. Click **Save**. The **Application Ruleset** page appears with the configured rule.
    g. Click **Close**. The By Rule block appears on the canvas with the configured application rule.
5. Select Destination.
    a. Click **Destinations**. The **Destination 1** page appears.
    b. Select the tool port 1/3/x15.

    > **NOTE:** To configure or modify the port type, use the Quick Port Editor.

    c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
6. Deploy Traffic Policy
    a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
7. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2

To create a second-level map to filter GTP-u based on Tunnel ID rule match

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (m2) with the following components and configurations:
   - Source port - Virtual Port (vp1)
   - Destination port - Tool port (1/3/x15)
   - Application ruleset
     - Condition - GTP-U TEID, Min value - 0x001e8490, Max value - 0x001e8499, and Subset - None
2. Deploy and verify the Traffic Policy.

Map 3

To create a second-level map to collect unmatched traffic

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (scoll) with the following components:
   - Source port - Virtual port (vp1)
   - Application Ruleset - Collector
   - Destination port - Tool port (1/3/x14)

ERSPAN Tunneling

This topic describes the workflow required to filter packets based on ERSPAN ID using APF. In this example, the ERSPAN header is not removed from the packet.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port- 1/1/x1
Source Ruleset	By Rule • Condition - MAC Source, MAC Address - 00:a0:d1:e1:02:01 and MAC mask - 0000.0000.0000
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port - vp1
Second-Level Map	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	My GSOP-Adaptive Packet Filtering.
Application Ruleset	By Rule • Condition - MAC Source, MAC Address - 00:a0:d1:e1:02:01 and MAC mask - 0000.0000.0000
Destination	Tool port - 1/1/x3.

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.

3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure Traffic Policy for a First-Level Map

To create a first-level map to forward traffic from network port 1/1/g3 to virtual port vp:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle.

d. Select MAC Source from the **Conditions** list, and then enter 00:a0:d1:e1:02:01 in **Mac Address** field.

e. Enter MAC Mask as 0000.0000.0000.

f. Click **Save**. The added rule appears in the **Source Ruleset** page.

g. Click **Close**. The By Rule block appears on the canvas with the configured rule.

4. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears.

b. Select the configured virtual port vp1.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

5. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

6. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure Traffic Policy for a Second-Level Map

To create a second-level map to forward traffic from the virtual port vp to tool port 1/1/g1:

1. Access Traffic Policy.

a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

b. Create a new map with **Alias** as test1b.

2. Select Source.

a. Click **Sources**. The **Source 1** page appears with the list of available ports.

b. Select the virtual port vp1.

c. Click **Add**. The Source block appears on the canvas with the selected virtual port.

3. Select Application.

a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. You can either select an existing GSOP or create a new GSOP by clicking **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group (gs-grp) you created earlier.

f. Select Adaptive Packet Filtering from the **Applications** list. By default, the application you select is enabled.

g. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

4. Select Application Ruleset.

a. Expand the **Application Ruleset** component and then drag and drop a **By Rule**. The **Application Ruleset** page appears.

b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle to pass the configured traffic.

d. Select MAC Source from the **Conditions** list, and then enter 00:a0:d1:e1:02:01 in **Mac Address** field.

e. Enter MAC Mask as 0000.0000.0000.

f. Click **Save**. The **Application Ruleset** page appears with the configured rule.

g. Click **Close**. The By Rule block appears on the canvas with the configured application rule.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears.

b. Select the tool port 1/1/x3.

> **NOTE:** To configure or modify the port type, use the **Quick Port Editor**.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

6. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Distribute Traffic Based on Inner IP Addresses and Inner TCP Port values

In the following example, traffic is distributed based on inner IP addresses and inner TCP port values as follows:

- Packets from VLAN 20 with GTP inner IP 65.128.7.21 and 98.43.132.70, inner TCP port 80 is forwarded to one tool port
- Packets from VLAN 20 with GTP inner IP 65.128.7.21 and 98.43.132.70, inner TCP port 443 is forwarded to a second tool port
- All packets not matching these rules is forwarded to a third tool port

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port- 1/1/x3.
Source Ruleset	By Rule • Rule 1: Condition - VLAN and Min - 20 • Rule 2: Condition - Protocol, Protocol - UDP, and Value - 17 • Rule 3: Condition - Port Destination and Min - 2152
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port - vp1
Second-Level Map1	
Source	Virtual Port - vp1
Source Ruleset	N/A
Applications	My GSOP-Adaptive Packet Filtering.
Application Ruleset	By Rule Rule 1: • Condition - IPv4 Destination, Min value - 65.128.7.21, Net Mask - 255.255.255.255, and Position - 2 • Condition - Protocol, and Protocol - UDP, and Value - 17 • Condition - Port Destination and Min value - 80, and Position - 2 Rule 2: • Condition - IPv4 Destination, Min value - 98.43.132.70, Net Mask - 255.255.255.255, and Position - 2

Map Components	Sample Configuration
	• Condition - Protocol, and Protocol - UDP, and Value - 17 • Condition - Port Destination and Min value - 80, and Position - 2
Destination	Tool ports - 1/1/x10
Second-Level Map2	
Source	Virtual Port - vp1
Source Ruleset	N/A
Applications	My GSOP-Adaptive Packet Filtering.
Application Ruleset	By Rule Rule 1: • Condition - IPv4 Destination, Min value - 65.128.7.21, Net Mask - 255.255.255.255, and Position - 2 • Condition - Protocol, and Protocol - UDP, and Value - 17 • Condition - Port Destination and Min value - 443, and Position - 2 Rule 2: • Condition - IPv4 Destination, Min value - 98.43.132.70, Net Mask - 255.255.255.255, and Position - 2 • Condition - Protocol, and Protocol - UDP, and Value - 17 • Condition - Port Destination and Min value - 443, and Position - 2
Destination	Tool ports - 1/1/x10
Second-Level Map3	
Source	Virtual Port - vp1
Source Ruleset	N/A
Applications	N/A
Application Ruleset	Collector
Destination	Tool ports - 1/1/x12

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.

> **NOTE:** For APF, no GigaSMART parameters are required unless combined with other GigaSMART operations.

7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure Traffic Policy for a First-Level Map

To create a first-level map to forward GTP-u traffic to virtual port vp1:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.

b. Select the network port 1/1/x1.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Source block appears on the canvas with the selected network port.

3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the configured traffic.
   d. Add a rule with the following conditions and values:
      - **Condition** - VLAN and **Min** value - 20
      - **Condition** - Protocol and **Protocol** - UDP, Value - 17
      - **Condition** - Port Destination and **Min** value- 2152
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears.
   b. Select the configured virtual port vp1.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure Traffic Policy for Second-Level Maps

Map 1

To create a second-level map to filter traffic based on inner IP address and TCP values

1. Access Traffic Policy.

a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

b. Create a new map with **Alias** as map2.

2. Select Source.

a. Click **Sources**. The **Source 1** page appears with the list of available ports.

b. Select the virtual port vp1.

c. Click **Add**. The Source block appears on the canvas with the selected virtual port.

3. Select Application.

a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group (gs-grp) you created earlier.

f. Select Adaptive Packet Filtering from the **Applications** list. By default, the application you select is enabled.

g. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

4. Select Application Ruleset.

a. Expand the Application Ruleset component and then drag and drop a **By Rule**. The **Application Ruleset** page appears.

b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle to pass the configured traffic.

d. Add a rule with the following conditions and values:

- **Condition** - IPv4 Destination, **Min** value - 65.128.7.21, **Net Mask** - 255.255.255.255, and **Position** - 2
- **Condition** - Protocol, and **Protocol** - UDP, and **Value** - 17
- **Condition** - Port Destination and **Min** value - 80, and **Position** - 2

e. Add another rule with the following conditions and values:

- **Condition** - IPv4 Destination, **Min** value - 98.43.132.70, **Net Mask** - 255.255.255.255, and **Position** - 2
- **Condition** - Protocol, and **Protocol** - UDP, and **Value** - 17

- **Condition** - Port Destination and **Min** value - 80, and **Position** - 2

f. Click **Save**. The **Application Ruleset** page appears with the configured rules.

g. Click **Close**. The By Rule block appears on the canvas with the configured application rule.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears.

b. Select the tool port 1/1/x10.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

6. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2

To create a second-level map to filter traffic based on inner IP address and TCP values

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (map3) with the following components and configurations:
   - Source port - Virtual Port (vp1)
   - Destination port - 1/1/x11
   - Application ruleset - Add two rules with the following conditions and values:
   - Rule1
     - **Condition** - IPv4 Destination, **Min** value 65.128.7.21, **Net Mask** - 255.255.255.255, and **Position** - 2
     - **Condition** - Protocol, and **Protocol** - UDP, and **Value** - 17
     - **Condition** - Port Destination and **Min** value - 443, and **Position** - 2
   - Rule2
     - **Condition** - IPv4 Destination, **Min** value 98.43.132.70, **Net Mask** - 255.255.255.255, and **Position** - 2

- **Condition** - Protocol, and **Protocol** - UDP, and **Value** - 17
- **Condition** - Port Destination and **Min** value - 443, and **Position** - 2

2. Deploy and verify the Traffic Policy.

Map 3

To create a second-level map to collect unmatched traffic

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (mapclin) with the following components:
   - Source port - Virtual port (vp1)
   - Source Ruleset - Collector
   - Destination port - Tool port (1/1/x12)

MPLS Label-Based Filtering

This topic describes the workflow required to filter traffic based on MPLS label using APF.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port- 1/1/x1.
Source Ruleset	By Rule • Rule 1: Condition - IP Version and Version - Version 4 • Rule 2: Condition - MAC Source, MAC Address - 00:a0:d1:e1:02:01 and MAC mask - 0000.0000.0000
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port - vp1
Second-Level Map1	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	• My GSOP-Adaptive Packet Filtering.
Application Ruleset	By Rule • Condition - MPLS Label, Min - 2, and Position - 1

Map Components	Sample Configuration
Destination	Tool port- 1/1/x3
Second-Level Map2	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	• My GSOP-Adaptive Packet Filtering.
Application Ruleset	By Rule • Condition - MPLS Label, Min - 3, and Position - 1
Destination	Tool port - 1/1/x4

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.

   NOTE: For APF, no GigaSMART parameters are required unless combined with other GigaSMART operations.

7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.

6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure Traffic Policy for a First-Level Map

To configure a first-level map to forward fabric path packets to the virtual port:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Add two rules with the following conditions:
   d. Add Rule1:
      - Turn on **Pass** toggle to pass the configured traffic.
      - Select IP Version as **Condition** and **Version** as Version 4.
   e. Add Rule2:

- Turn on **Pass** toggle and select **Bi-directional** option.
- Select Mac Source as **Condition**, Mac Address as 00:a0:d1:e1:02:01, and Mac Mask - 0000.0000.0000

f. Click **Save**. The added rules appear in the **Source Ruleset** page.

g. Click **Close**. The By Rule block appears on the canvas with the configured rules.

4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears.
   b. Select the configured virtual port vp1.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure Traffic Policy for Second-Level Maps

To create a second-level map to filter based on MPLS label:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Create a new map with **Alias** as map1.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears with the list of available ports.
   b. Select the virtual port vp1.
   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application.
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.

c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group (gs-grp) you created earlier.

f. Select Adaptive Packet Filtering from the **Applications** list. By default, the application you select is enabled.

g. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

4. Select Application Ruleset.

a. Expand the Application Ruleset component and then drag and drop a **By Rule**. The **Application Ruleset** page appears.

b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle to pass the configured traffic.

d. Select MPLS Label from the **Condition** list and enter 2 in the **Min** field.

e. Select **Position** to 1.

f. Click **Save**. The **Application Ruleset** page appears with the configured rule.

g. Click **Close**. The By Rule block appears on the canvas with the configured application rule.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears.

b. Select the tool port 1/1/x3.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

6. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

To create a second-level map to filter based on MPLS label

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (map2) with the following components and configurations:
   - Source port - Virtual Port (vp1)
   - Destination port - Tool port 1/1/x4
   - Application ruleset
     - Condition - MPLS Label, Min value - 3, and Position - 1
2. Deploy and verify the Traffic Policy.

Conditional Header Stripping

Another use-case that can be addressed leveraging the flexibility of APF would be the capability to header strip packets based on specific contents found across the packet including the inner packet contents. Since the APF rules are enforced before any other GigaSMART operation, operators can filter based on encapsulation protocol values and /or encapsulated (original) packet contents and apply conditional Header Stripping operations.

The following example shows how an end-user can filter and strip out outer VXLAN headers for a subset of the traffic based on inner IP addresses, while sending the rest of the traffic “as-is” to monitoring tools that need the VXLAN headers for traffic analysis, as follows.

- Identifying and forwarding traffic from/to ip 1.1.1.1 in the inner / encapsulated packets to monitoring tool connected to tool port 1/1/x1 after Header Stripping VXLAN.
- Identifying and forwarding traffic from/to ip 1.1.1.2 in the inner / encapsulated packets to monitoring tool connected to tool port 1/1/x4 without stripping the VXLAN header.

NOTE: This can be applied to any GigaSMART operation. While this example shows filtering based on inner packet contents, conditional SMART operations can be applied by filtering on encapsulation headers as well.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port- 1/1/x3.
Source Ruleset	By Rule • Condition - Port Source and Value - 8472

Map Components	Sample Configuration
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port - vp1
Second-Level Map1	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	My GSOP - Adaptive Packet Filtering and Header Stripping Header Stripping: Type - VXLAN and VXLAN ID - 0-16777216
Application Ruleset	By Rule • Rule 1: Condition - IPv4 Source, Min - 1.1.1.1, Net Mask - 255.255.255.255, and Position - 2 • Rule 2: Condition - IPv4 Destination and Min - 1.1.1.2, Net Mask - 255.255.255.255, and Position - 2
Destination	Tool port - 1/1/x1
Second-Level Map2	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	My GSOP - Adaptive Packet Filtering.
Application Ruleset	By Rule • Rule 1: Condition - IPv4 Source, Min - 1.1.1.1, Net Mask - 255.255.255.255, and Position - 2 • Rule 2: Condition - IPv4 Destination and Min - 1.1.1.2, Net Mask - 255.255.255.255, and Position - 2
Destination	Tool port - 1/1/x4

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Create a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure Traffic Policy for a First-Level Map

To create a first-level map to forward fabric path packets to the virtual port:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x3.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.

3. Select Source Ruleset.
    a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
    b. Click **Add Rules**. The **Add Ruleset** pane appears.
    c. Turn on **Pass** toggle to pass the VXLAN traffic.
    d. Select Port Source from the **Condition** list, and then enter 8472 in the **Value** field.
    e. Click **Save**. The added rule appears in the **Source Ruleset** page.
    f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.
    a. Click **Destinations**. The **Destination 1** page appears.
    b. Select the configured virtual port vp1.
    c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
    a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
    a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
    b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure Traffic Policy for Second-Level Maps

Map 1:

To create a second-level map to filter on source and destination IP (bi-directional):

1. Access Traffic Policy.
    a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
    b. Create a new map with **Alias** as map1.
2. Select Source.
    a. Click **Sources**. The **Source 1** page appears with the list of available ports.
    b. Select the virtual port vp1.
    c. Click **Add**. The Source block appears on the canvas with the selected virtual port.

3. Select Application.
    a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
    b. Select either a standalone node or a cluster from the **Select Node** list.
    c. Click **New** from the **GSOP** field, the **New GSOP** pane appears with the selected node.
    d. Enter a gsfil_vxlanhs for the GigaSMART operation in the **GSOP Alias** field.
    e. Select the GigaSMART group (gs-grp) you created earlier.
    f. Select Adaptive Packet Filtering and Header Stripping from the **Applications** list. By default, the applications you select are enabled.
    g. Select VXLAN from the Header Stripping list.
    h. Enter **VXLAN ID** in the range 0-16777216.
    i. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operations.
4. Select Application Ruleset.
    a. Expand the **Application Ruleset** component and then drag and drop a **By Rule**. The **Application Ruleset** page appears.
    b. Click **Add Rules**. The **Add Ruleset** pane appears.
    c. Add Rule1:
        - Turn on **Pass** toggle to pass the configured traffic.
        - Select IPv4 Source from the **Condition** list, and then enter 1.1.1.1 in the **Min** field.
        - Enter 255.255.255.255 in the **Net Mask** field and 2 in the **Position** field.
    d. Add Rule2:
        - Turn on **Pass** toggle to pass the configured traffic.
        - Select IPv4 Destination from the **Conditions** list, and then enter 1.1.1.2 in the **Min** field.
        - Enter 255.255.255.255 in the **Net Mask** field and 2 in the **Position** field.
    e. Click **Save**. The **Application Ruleset** page appears with the configured rule.
    f. Click **Close**. The By Rule block appears on the canvas with the configured application rule.
5. Select Destination.
    a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.
    b. Select the tool port 1/1/x1.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Destination block appears on the canvas with the selected tool port.

6. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
7. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the map (map2) with
   - Source - Virtual port (vp1)
   - Destination - Tool port (1/1/x4)
   - Application - Adaptive Packet Filtering
   - Application ruleset - By Rule
     - Rule1: Turn on **Pass** toggle to pass the configured traffic.
     - Select IPv4 Source from the **Condition** list, and then enter 1.1.1.2 in the **Min** field.
     - Enter 255.255.255.255 in the **Net Mask** field and 2 in the **Position** field.
     - Rule2: Turn on **Pass** toggle to pass the configured traffic.
     - Select IPv4 Destination from the **Conditions** list, and then enter 1.1.1.2 in the **Min** field.
     - Enter 255.255.255.255 in the **Net Mask** field and 2 in the **Position** field.
2. Deploy and verify the Traffic Policy.

Facilitate Overlapping Rules

Because APF is implemented as a second level map operation, APF can also be leveraged for implementing basic overlapping rules. For the same incoming input stream, a copy of the traffic can be sent out to a group of monitoring tools while a refined subset of the traffic stream can be sent to a different set of monitoring tools. Typically overlapping rules would be implemented by combining APF with the patented Flow Mapping® technology.

Note that Role-Based Access control in the case of APF is applied at the gsgroup / e port.

In the following example, for the same input stream:

- HTTP traffic is identified and distributed to a monitoring tool connected to tool port 1/1/x1.
- At the same time, the same stream of HTTP packets are being sent out after slicing unwanted packet contents to a different monitoring tool connected to tool port 1/1/x4.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port- 1/1/x3.
Source Ruleset	By Rule • Condition - Port Source and Value - 2152
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port - vp1
Second-Level Map1	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	My GSOP - Adaptive Packet Filtering and Slicing. • Type - None and Offset - 150.
Application Ruleset	By Rule • Condition - IP Version, Version - Version 4, and Position - 1.
Destination	Tool port- 1/1/x4
Second-Level Map2	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	My GSOP - Adaptive Packet Filtering.
Application Ruleset	By Rule • Condition - IP Version, Version - Version 4, and Position - 4.
Destination	Tool port- 1/1/x1

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.

2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.

   NOTE: For APF, no GigaSMART parameters are required unless combined with other GigaSMART operations.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure Traffic Policy for a First-Level Map

To create a first-level map to forward fabric path packets to the virtual port:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.

- (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
- (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.

d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x3.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the configured traffic.
   d. Select Port Source from the **Condition** list, and then enter 2152 in the **Value** field.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears.
   b. Select the configured virtual port vp1.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure Traffic Policy for Second-Level Maps

Map 1:

To create a second-level map to filter on source and destination IP (bi-directional):

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Create a new map with **Alias** as map1.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears with the list of available ports.
   b. Select the virtual port vp1.
   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application.
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. You can either select an existing GSOP or create a new GSOP by clicking **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
   e. Select the GigaSMART group (gs-grp) you created earlier.
   f. Select Adaptive Packet Filtering and Slicing from the **Applications** list. By default, the applications you select are enabled.
   g. Select None from the **Slicing** list and enter 150 in the **Offset** field.
   h. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operations.
4. Select Application Ruleset.
   a. Expand the **Application Ruleset** component and then drag and drop a **By Rule**. The **Application Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Select IP Version from the **Condition** list and Version 4 from the **Version** list.
   d. Select **Position** as 1.
   e. Click **Save**. The **Application Ruleset** page appears with the configured rule.

f. Click **Close**. The By Rule block appears on the canvas with the configured application rule.

5. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.
   b. Select the tool port 1/1/x4.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Destination block appears on the canvas with the selected tool port.
6. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
7. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the map (map2) with
   - Source - Virtual port (vp1)
   - Destination - Tool port (1/1/x1)
   - Application - Adaptive Packet Filtering
   - Application Ruleset - By Rule
     - Condition - IP Version, Version - 4, and Position - 4.
2. Deploy and verify the Traffic Policy.

Combine APF with GigaSMART Operations

APF can also be combined with other GigaSMART functions including Header Stripping, Packet Slicing or Masking, De-duplication and FlowVUE. This provides network administrators and operators to perform a second layer of filtering in combination with the GigaSMART tool optimization and packet manipulation operations.

In the following example, operators can distribute traffic to monitoring tools based on decapsulated contents, more specifically, after Header stripping VXLAN:

- Identifying and forwarding traffic from/to ip 1.1.1.1 from the decapsulated packets to monitoring tool connected to tool port 1/1/x1
- Identifying and forwarding traffic from/to ip 1.1.1.2 in the decapsulated packets to monitoring tool connected to tool port 1/1/x4

NOTE: This can be applied to any protocol that is supported through header-stripping, for example: GTP, VXLAN, ISL, MPLS, MPLS+VLAN, VLAN, VN-Tag, fabric-path. This is also supported for Gigamon tunnel decapsulation.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port- 1/1/x1
Source Ruleset	By Rule • Condition - Port Source and Value - 8472
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port - vp1
Second-Level Map1	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	• My GSOP-Adaptive Packet Filtering, Header Stripping • Header Stripping: Type - VXLAN ID, Value - 0-16777216
Application Ruleset	By Rule • Rule 1: Condition - IPv4 Source, Min - 1.1.1.1, Net Mask - 255.255.255.255, and Position - 2 • Rule 2: Condition - IPv4 Destination, Min - 1.1.1.2, Net Mask - 255.255.255.255, and Position - 2
Destination	Tool port- 1/1/x3
Second-Level Map2	
Source	Virtual Port - vp1
Source Ruleset	N/A
Applications	• My GSOP-Adaptive Packet Filtering.
Application Ruleset	• Rule 1: Condition - IPv4 Source, Min - 1.1.1.1, Net Mask - 255.255.255.255, and Position - 2 • Rule 2: Condition - IPv4 Destination, Min - 1.1.1.2, Net Mask - 255.255.255.255, and Position - 2
Destination	Tool port- 1/1/x4

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The Add GigaSMART Group page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**.

   NOTE: For APF, no GigaSMART parameters are required unless combined with other GigaSMART operations.

7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure Traffic Policy for a First-Level Map

To configure a first-level map to forward fabric path packets to the virtual port:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.

c. In the **New Map** dialog, do the following:

- **Alias** - Enter map1.
- **Alias** - Enter GTP- Control.
- (Optional) **Description** - Enter the description for the Traffic Policy.
- (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
- (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.

d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source.

a. Click **Sources**. The Source 1 page appears.

b. Select the network port 1/1/x1.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Source block appears on the canvas with the selected network port.

3. Select Source Ruleset.

a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The Source Ruleset page appears.

b. Click **Add Rules**. The Add Ruleset pane appears.

c. Turn on **Pass** toggle to pass the VXLAN traffic.

d. Select Port Source from the **Condition** list, and then enter 8472 in the **Value** field.

e. Click **Save**. The added rule appears in the **Source Ruleset** page.

f. Click **Close**. The By Rule block appears on the canvas with the configured rule.

4. Select Destination.

a. Click **Destinations**. The Destination 1 page appears.

b. Select the configured virtual port vp1.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

5. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

6. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure Traffic Policy for Second-Level Maps

Map 1:

To configure a second-level map to filter on source and destination IP (bi-directional):

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Create a new map with **Alias** as map1.
2. Select Source.
   a. Click **Sources**. The Source 1 page appears with the list of available ports.
   b. Select the virtual port vp1.
   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application.
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The GSOP page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. You can either select an existing GSOP or create a new GSOP by clicking **New** from the **GSOP** field. The New GSOP pane appears with the selected node.
   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
   e. Select the GigaSMART group (gs-grp) you created earlier.
   f. Select Adaptive Packet Filtering and Header Stripping from the **Applications** list. By default, the applications you select are enabled.
   g. Select VXLAN from the Header Stripping list.
   h. Enter **VXLAN ID** in the range 0-16777216.
   i. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operations.
4. Select Application Ruleset.

a. Expand the **Application Ruleset** component and then drag and drop a **By Rule**. The Application Ruleset page appears.

b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Add Rule1:

- Turn on **Pass** toggle to pass the configured traffic.
- Select IPv4 Source from the **Condition** list, and then enter 1.1.1.1 in the **Min** field.
- Enter 255.255.255.255 in the **Net Mask** field and 2 in the **Position** field.

d. Add Rule2:

- Turn on **Pass** toggle to pass the configured traffic.
- Select IPv4 Destination from the **Conditions** list, and then enter 1.1.1.2 in the **Min** field.
- Enter 255.255.255.255 in the **Net Mask** field and 2 in the **Position** field.

e. Click **Save**. The Application Ruleset page appears with the configured rule.

f. Click **Close**. The By Rule block appears on the canvas with the configured application rule.

5. Select Destination.

a. Click **Destinations**. The Destination 1 page appears with the list of available ports.

b. Select the tool port 1/1/x3.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Destination block appears on the canvas with the selected tool port.

6. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

To configure a second-level map to filter on source and destination IP (bi-directional)

Follow the same steps in the previous section to configure another second-level map with the following changes:

1. Create the map (map2) with

- Source - Virtual Port vp1
- Destination - Tool Port 1/1/x4
- Application ruleset - By Rule
  - Rule1:
  - Turn on **Pass** toggle to pass the configured traffic.
  - Select IPv4 Source from the **Condition** list, and then enter 1.1.1.2 in the **Min** field.
  - Enter 255.255.255.255 in the **Net Mask** field and 2 in the **Position** field.
  - Rule2:
  - Turn on **Pass** toggle to pass the configured traffic.
  - Select IPv4 Destination from the **Condition** list, and then enter 1.1.1.2 in the **Min** field.
  - Enter 255.255.255.255 in the **Net Mask** field and 2 in the **Position** field.

2. Deploy and verify the Traffic Policy.

GigaSMART Load Balancing

Load balancing distributes GigaSMART outgoing traffic to multiple tool ports or multiple tunnel endpoint destinations. In this way, traffic processed is shared. Stateful load balancing distributes GigaSMART processed traffic to multiple tool ports or tunnel endpoints based on GigaSMART application-specific flow sessions whereas Stateless load balancing distributes GigaSMART processed traffic based on hash values generated from predefined protocol fields in the packet.

Refer to the following sections for instructions on how to configure Traffic Policy for load balancing operation with real-time examples:

- Stateful Load Balancing
- Stateless Load Balancing

Stateful Load Balancing

This topic describes the workflow required to perform stateful load balancing with real time examples..

Example 1:

Configure stateful load balancing of GigaSMART GTP traffic among tool ports 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4 based on bandwidth with different weights for each port. The same subscriber (imsi) traffic will be forwarded to the same tool port. GTP-c packets are replicated to all tool ports.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Tool port- 1/1/x1
Source Ruleset	By Rule • Rule1: Condition-Port Destination, Min-2123 • Rule2: Condition-Port Destination, Min-2152
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port -vp1
Second-Level Map	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	My GSOP - Flow Filtering and Load Balancing • Load Balancing - Stateful, Type - GTP, Metric- Hashing
Application Ruleset	Flow Filter • Condition - GTP IMSI, IMSI - 234567*, and Version - Any
Destination	Tool port group - portgrp1

Configure a Port Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new port group. The New Port Group page appears.
4. Enter portgrp1 in the **Alias** field and select **Tool** option.
5. Enable **Load Balancing**.
6. Choose the **Tool Ports** you want to attach to the port group, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
7. Specify the weights for each port as follows:
    - Weight 1/1/x6 - 5
    - Weight 1/1/x7 - 10
    - Weight 1/2/x3 - 20

- Weight 1/2/x4 - 10

8. Click **Apply**. The created port group appears in the port group table.

Configure a GigaSMART Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (GS-GRP-CHAIN) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Expand the **GigaSMART Parameters**, and go the **Load Balance** section- Select **Replicate GTP-c** option in.

   NOTE: You are enabling **Replicate GTP-c** so that every tool port in the load balancing port group receives a copy of all GTP control-plane (GTP-c) packets, even though the user-plane (GTP-u) traffic is being load-balanced across those ports.

8. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure a Traffic Policy for First-Level Map

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules** to add Rule 1. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle.
   d. Select Port Destination from the **Condition** list.
   e. Enter 2123 in the **Min** field as port value.
   f. Click **Save**. The added rule appears in the **Source Ruleset** page.
   g. Click **Add Rules** to add Rule 2. The **Add Ruleset** pane appears.
   h. Select Port Destination from the **Condition** list.
   i. Enter 2152 in the **Min** field as port value.

   NOTE: In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.

j. Click **Save**. The added rule appears in the **Source Ruleset** page.

k. Click **Close**. The By Rule block appears on the canvas with the configured rules.

4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears.
   b. Select the virtual port vp1.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify Deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure a Traffic Policy for Second-Level Map

To create a second-level map to forward traffic from the virtual port vp1 to GigaSMART for pattern matching

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Create a new map with **Alias** as map1.
2. Select Source.
   a. Click **Sources**. The Source 1 page appears with the list of available ports.
   b. Select the virtual port vp1.
   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application Ruleset.
   a. Expand the **Application Ruleset** then drag and drop the **Flow Filter**. The **Flow Filter** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the configured traffic.
   d. Select GTP IMSI from the **Condition** list.

e. Enter 234567* In the **IMSI** field.

f. Select Version in the **Type** field. Select Any in the **Version** field.

g. Click **Save**. The **Application Ruleset** page appears with the configured rule.

h. Click **Close**. The **Flow Filter** block appears on the canvas with the configured application rule.

4. Select Application.

a. Expand the **Application** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

d. Select the GigaSMART group (gs-grp) that you created earlier.

e. Select Flow Filtering and Load Balancing from the **Applications** list. By default, the applications you select are enabled.

f. In the Load Balancing section, select Stateful option and type as GTP.

g. Select Weighted Least Bandwidth from the **Metric** list.

h. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool port group portgrp1.

c. Click **Add**. The Destination block appears on the canvas with the selected tool port group.

6. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify Deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Example 2:

Configure stateful load balancing of GigaSMART GTP traffic among tool ports 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4 based on hashing of the imei value. The same device ID (imei) traffic will be forwarded to the same tool port. GTP-c packets are replicated to all tool ports.

Applicable Map Components

Map Components	Sample Configuration
First-Level Map	
Source	Tool port- 1/1/x1
Source Ruleset	By Rule • Rule1: Condition-Port Destination, Min-2123 • Rule2: Condition-Port Destination, Min-2152
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port -vp1
Second-Level Map	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	My GSOP - Flow Filtering and Load Balancing • Load Balancing - Stateful, Type - GTP, Metric- Hashing
Application Ruleset	Flow Filter • Condition - GTP IMEI, IMEI - 234567*, and Version - Any
Destination	Tool port group - portgrp1

Configure a Port Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter a portgrp1 in the **Alias** field for the new Port Group and click on **Tool** button.
5. Enable **Load Balancing**.
6. Choose the **Tool Ports** you want to attach to this **Port Group**, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
7. Specify the weights for each port as follows-
   - Weight 1/1/x6 5
   - Weight 1/1/x7 10
   - Weight 1/2/x3 20
   - Weight 1/2/x4 10
8. Click **Apply** to save. The created Port group appears in the table.

Configure a GigaSMART Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Expand the **GigaSMART Parameters**, and go the **Load Balance** section- Select **Replicate GTP-c** option in.

NOTE: You are enabling **Replicate GTP-c** so that every tool port in the load-balancing port group receives a copy of all GTP control-plane (GTP-c) packets, even though the user-plane (GTP-u) traffic is being load-balanced across those ports.

8. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure Traffic Policy for a First-Level Map

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.

c. In the **New Map** dialog, do the following:

- **Alias** - Enter map1.
- (Optional) **Description** - Enter the description for the Traffic Policy.
- (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
- (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.

d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source.

a. Click **Sources**. The **Source 1** page appears.

b. Select the network port 1/1/x1.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Source block appears on the canvas with the selected network port.

3. Select Source Ruleset.

a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.

b. Click **Add Rules** to add Rule 1. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle.

d. Select Port Destination in the **Condition** field.

e. Enter 2123 in the **Min** field.

f. Click **Save**. The added rule appears in the **Source Ruleset** page.

g. Click **Add Rules** to add Rule 2. The **Add Ruleset** pane appears.

h. Select Port Destination in the **Condition** field.

i. Enter 2152 in the **Min** field.

> **NOTE:** In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.

j. Click **Save**. The added rule appears in the **Source Ruleset** page.

k. Click **Close**. The By Rule block appears on the canvas with the configured rules.

4. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears.

b. Select the virtual port vp1.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

5. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

6. Verify Deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure Traffic Policy for a Second-Level Map

To create a second-level map to forward traffic from the virtual port vp1 to GigaSMART for pattern matching:

1. Access Traffic Policy.

a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

b. Create a new map with **Alias** as map1.

2. Select Source.

a. Click **Sources**. The Source 1 page appears with the list of available ports.

b. Select the virtual port vp1.

c. Click **Add**. The Source block appears on the canvas with the selected virtual port.

3. Select Source.

a. Click **Sources**. The **Source 1** page appears with the list of available ports.

b. Select the virtual port vp1.

c. Click **Add**. The Source block appears on the canvas with the selected virtual port.

4. Select Application Ruleset.

a. Expand the **Application Ruleset** then drag and drop the **Flow Filter**. The **Flow Filter** page appears.

b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle to pass the configured traffic.

d. Select GTP IMEI from the **Condition** list.

e. Enter 234567* In the **IMEI** field.

f. Select Version in the **Type** field.

g. Select Any from the **Version** list.

h. Click **Save**. The **Application Ruleset** page appears with the configured rule.

i. Click **Close**. The **Flow Filter** block appears on the canvas with the configured application rule.

5. Select Application.

a. Expand the **Application** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

d. Select the GigaSMART group (gs-grp) that you created earlier.

e. Select Flow Filtering and Load Balancing from the **Applications** list. By default, the applications you select are enabled.

f. In the Load Balancing section, select Stateful option and type as GTP.

g. Select Weighted Least Bandwidth from the **Metric** list.

h. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

6. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool port group portgrp1.

c. Click **Add**. The Destination block appears on the canvas with the selected tool port group.

7. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

8. Verify Deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Stateless Load Balancing

This topic describes the workflow required to perform stateless load balancing with real time examples..

Example 1

Configure stateless load balancing of traffic among tool ports 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4 after slicing the packet to an offset of 70 bytes.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port- 1/1/x1
Source Ruleset	By Rule • Condition - IP Version and Version - Version 4
Application Ruleset	N/A
Applications	My GSOP - Load Balancing Load Balancing - Type - Stateless, Metric - IP Only and Type - Outer
Destination	Tool port group - portgrp1

Configure a Port Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new port group. The New Port Group page appears.
4. Enter portgrp1 in the **Alias** field and select **Tool** option.
5. Enable **Load Balancing**.
6. Choose the **Tool Ports** you want to attach to the port group, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
7. Specify the weights for each port as follows:
    - Weight 1/1/x6 - 5
    - Weight 1/1/x7 - 10
    - Weight 1/2/x3 - 20
    - Weight 1/2/x4 - 10

8. Click **Apply**. The created port group appears in the port group table.

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Traffic Policy for Regular Map

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.

3. Select Source Ruleset.
    a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
    b. Click **Add Rules**. The **Add Ruleset** pane appears.
    c. Turn on **Pass** toggle to pass the configured traffic.
    d. Select IP Version from **Condition** list, and select Version 4 from the **Version** list.
    e. Click **Save**. The added rule appears in the **Source Ruleset** page.
    f. Click **Close**. The By Rule block appears on the canvas with the configured rules.
4. Select Application.
    a. Expand the **Application** component and then drag and drop **My GSOP**. The **GSOP** page appears.
    b. Select either a standalone node or a cluster from the **Select Node** list.
    c. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
    d. Select the GigaSMART group (gs-grp) that you created earlier.
    e. Select Load Balancing from the **Applications** list. By default, the application you select is enabled.
    f. In the Load Balancing section, select Stateless option.
    g. Select IP Only from the **Metric** list and select **Outer** option.
    h. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.
5. Select Destination.
    a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.
    b. Select the tool port group portgrp1.
    c. Click **Add**. The Destination block appears on the canvas with the selected tool port group.
6. Deploy Traffic Policy.
    a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
7. Verify deployment.
    a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
    b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Example 2

Configure stateless load balancing of GTP traffic among tool ports 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4. Data packets with the same GTP-u tunnel ID will be forwarded to the same tool port.

Applicable Map Components

Map Components	Sample Configuration
First-Level Map	
Source	Network port- 1/1/x1.
Source Ruleset	By Rule • Condition - IP Version and Version - Version 4
Application Ruleset	N/A
Applications	My GSOP - Load Balancing Load Balancing - Type - Stateless, Metric - IP Only and Type - Outer
Destination	Tool port group - portgrp1

Configure a Port Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new port group. The New Port Group page appears.
4. Enter portgrp1 in the **Alias** field and select **Tool** option.
5. Enable **Load Balancing**.
6. Choose the **Tool Ports** you want to attach to the port group, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
7. Specify the weights for each port as follows:
   - Weight 1/1/x6 - 5
   - Weight 1/1/x7 - 10
   - Weight 1/2/x3 - 20
   - Weight 1/2/x4 - 10
8. Click **Apply** to save. The created port group appears in the port group table.

Configure a GigaSMART Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.

3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1 and 1/3/e2.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure Traffic Policy for a First-Level Map

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the configured traffic.
   d. Configure Rule 1: Select Protocol from the **Condition** list.
   e. Select UDP from the **Protocol** list and enter 17 in the **Value** field.

f. Configure Rule 2: Select Port Destination from the **Condition** list.

g. Enter 2123 in the **Min** field.

h. Click **Save**. The added rule appears in the **Source Ruleset** page.

i. Click **Close**. The By Rule block appears on the canvas with the configured rules.

4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.
   b. Select the tool port group portgrp1.
   c. Click **Add**. The Destination block appears on the canvas with the selected tool port group.
5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Source - Virtual Port (vp1)
2. Source Ruleset - By Rule
   - Turn on **Pass** toggle to pass the configured traffic.
   - Rule 1: **Condition** - Protocol, **Protocol**- UDP and **Value**-4.
   - Rule 2: **Condition** - Port Destination and **Min** - 2152.
3. Application - Load Balancing
   - Type - Stateless and Metric - GTP-TEID
4. Destination - Tool Port Group (portgrp1)

Example 3

Configure stateless load balancing of HTTP on GTP traffic among tool ports 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4. Data packets with the same inner IP will be forwarded to the same tool port.

Configure a Port Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter portgrp1 in the **Alias** field for the new Port Group and click **Tool** option.
5. Enable **Load Balancing**.
6. Choose the **Tool Ports** you want to attach to the port group, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
7. Click **Apply**. The created port group appears in the port group table.

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.

7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure Traffic Policy for a First-Level Map

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the configured traffic.
   d. Configure Rule 1: Select Protocol from the **Condition** list.
   e. Select UDP from the **Protocol** list and enter 17 in the **Value** field.
   f. Configure Rule 2: Select Port Destination from the **Condition** list.
   g. Enter 2123 in the **Min** field.
   h. Click **Save**. The added rule appears in the **Source Ruleset** page.
   i. Click **Close**. The By Rule block appears on the canvas with the configured rules.

4. Select Destination.
    a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.
    b. Select the tool port group portgrp1.
    c. Click **Add**. The Destination block appears on the canvas with the selected tool port group.
5. Deploy Traffic Policy.
    a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
    a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
    b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

For map2 in first level, following the above procedure

Map 2:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Source - Virtual Port (vp1)
2. Source Ruleset - By Rule
    - Turn on **Pass** toggle to pass the configured traffic.
    - Rule 1: **Condition** - Protocol, **Protocol**- UDP and **Value**-17.
    - Rule 2: **Condition** - Port Destination and **Min** - 2152.
3. Application - Load Balancing
    - Type - Stateless and Metric - GTP-TEID
4. Destination - Tool Port Group (portgrp1)

Configure Traffic Policy for a Second-Level Map

To create a second-level map to forward traffic from the virtual port vp1 to GigaSMART for pattern matching

1. Access Traffic Policy.
    a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.

c. Create a map with **Alias** as map2.

2. Select Source Ruleset.

a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.

b. Click **Add Rules** to add Rule 1. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle to pass the configured traffic.

d. Select Port Destination from the **Condition** list

e. Enter 80 in the **Min** field.

f. Click **Save**. The added rule appears in the **Source Ruleset** page.

g. Click **Close**. The By Rule block appears on the canvas with the configured rules.

3. Select Application.

a. Expand the **Application** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

d. Select the GigaSMART group (gs-grp) that you created earlier.

e. Select Load Balancing from the **Applications** list. By default, the application you select is enabled.

f. In the Load Balancing section, select Stateless option.

g. Select IP Only from the **Metric** list and select **inner** option.

h. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

4. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool port group portgrp1.

c. Click **Add**. The Destination block appears on the canvas with the selected tool port group.

5. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

6. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Enhanced Load Balancing

Applicable Map Components

Map Components	Sample Configuration
First-Level Map	
Source	
Source Ruleset	
Application Ruleset	N/A
Applications	
Destination	
Second-Level Map	
Source	
Source Ruleset	
Applications	
Application Ruleset	
Destination	

GigaSMART De-Duplication

Required License:De-duplication

Supported Devices: GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus.

Example

De-duplicate traffic from multiple SPAN ports and send the result to an Intrusion Detection System (IDS) tool.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
Source	Network (SPAN) ports - 1/1/x6, 1/1/x7
Source Ruleset	By Rule • Condition - Port Source, Min - 0 and Max - 443.
Application Ruleset	N/A
Applications	My GSOP - De-duplication
Destination	IDS Tool port - 1/1/x3

Configure Traffic Policy

To configure a Traffic Policy to drop duplicate packets using the GigaSMART De-duplication:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network ports 1/1/x6 and 1/1/x7.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle to pass the traffic from the selected source ports.

d. Select the **Bi-directional** checkbox.

e. Select Port Source from the **Condition** list, and then set the **Min** to 0 and **Max** to 443.

f. Click **Save**. The added rule appears in the **Source Ruleset** page.

g. Click **Close**. The By Rule block appears on the canvas with the configured rule.

4. Select Application.

a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. You can either select an existing GSOP or create a new GSOP by clicking **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field. For example, gsop-dedup.

e. You can either select an existing GigaSMART group or create a new group by clicking **New** in the **GigaSMART Group** field and the **Create GigaSMART Group** pane appears.

f. Enter a name for the GigaSMART group in the **Alias** field. For example, gs2port1.

g. Select a GigaSMART engine port from the **Port List**.

h. Expand the **GigaSMART Parameters**, and go to the **Dedup** section.

i. Configure the following parameters:

- **Action**: Drop
- **IP Tclass**: Include
- **IP TOS**: Ignore
- **TCP Sequence**: Ignore
- **VLAN**: Ignore
- **Timer** (µs): 500000

j. Click **Save**. The **New GSOP** pane appears.

k. Select the created GigaSMART group from the **GigaSMART Group** list.

> **NOTE:** Enter the GigaSMART operation name in the **GSOP Alias** field, only after creating the GigaSMART group, if you are creating a new group.

l. Select De-duplication from the **Applications** list. By default, the application you select is enabled.

m. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

5. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.
   b. Select the tool port 1/1/x3 port.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Destination block appears on the canvas with the selected tool port.
6. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
7. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

GigaSMART Masking

Required License: Base
Supported Devices: GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus.

GigaSMART masking operations overwrite a defined portion of a packet with a specified one-byte pattern. Each masking operation is defined by an offset, a length, and a pattern.

This topic describes the workflow required to perform masking.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port- 1/1/x1
Source Ruleset	By Rule • Condition - IP Version and Version - 4
Application Ruleset	N/A
Applications	My GSOP - Masking • Protocol - IPv4 • Offset - 77, Pattern - FF, and Length - 180
Destination	Tool Port - 1/1/x3

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The Add GigaSMARTGroup page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Traffic Policy for Masking

To configure a Traffic Policy for masking:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.

c. In the **New Map** dialog, do the following:

- **Alias** - Enter map1.
- (Optional) **Description** - Enter the description for the Traffic Policy.
- (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
- (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.

d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source

a. Click **Sources**. The **Source 1** page appears.

b. Select the network port 1/1/x1.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Source block appears on the canvas with the selected network port.

3. Select Source Ruleset

a. Expand the Source Ruleset component and then drag and drop **By Rule**. The **Source Ruleset** page appears.

b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle to pass the IPv4 traffic.

d. Turn on **Pass** toggle to pass the configured traffic.

e. Select IP Version from **Condition** list and select Version 4 for **Version** list.

f. Click **Save**. The By Rule block appears on the canvas with the configured rule.

4. Select Application.

a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group (gs-grp) that you created earlier.

f. Select Masking from the **Applications** list. By default, the application you select is enabled.

g. In the Masking section, select IPv4 protocol.

- Enter **Offset** - 0, **Pattern** - FF, and **Length** - 180.

h. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool port 1/1/x3.

> **NOTE:** To configure or modify the port type, use the **Quick Port Editor**.

c. Click **Add**. The Destination block appears on the canvas with the selected tool port.

6. Deploy Traffic Policy

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

GigaSMART NetFlow Generation

Required License: NetFlow GenerationRequired License for NetFlow with Second Level Maps: Adaptive Packet Filtering (APF)
Supported Devices: GigaVUE-HC1 Gen 2, GigaVUE-HC3 Gen 2.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network Port - 1/1/g3
Source Ruleset	By Rule • Condition - MAC Source, MAC Address - 00:a0:d1:e1:02:01, and MAC mask - 0000.0000.0000

Map Components	Sample Configuration
Applications	N/A
Application Ruleset	N/A
Destination	Virtual Port - vp1
Second-Level Map	
Source	Virtual Port - vp1
Source Ruleset	N/A
Applications	Adaptive Packet Filtering
Application Ruleset	By Rule • Condition - IP Version and Version - Version 4
Destination	Tool Port - 1/1/g1

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The Overview page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (Netflow-gsgrp) for the GigaSMART group in the **Alias** field.
6. Select the NetFlow from the GigaSMART Operations (GSOP) list.
7. By default, NetFlow is Enabled.
8. Select a GigaSMART engine port (1/5/e1) from the **Port List** .

   NOTE:
   The GigaSMART Group can contain multiple GigaSMART engine ports.
   Only one NetFlow Generation Monitor can be configured per GigaSMART Group.
   Once a GigaSMART group is created, the Alias and port List values are no longer modifiable. You will need to recreate the group to change these values.

9. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a NetFlow Exporter

To configure a NetFlow Exporter:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The Overview page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > NetFlow / IPFIX Generation > Exporter**. The **NetFlow Exporter** page appears.
4. Click **New**. The **New NetFlow Exporter** page appears.
5. Enter a name (NetFlow-Exporter) for the NetFlow in **Alias** field.
6. (Optional) Enter a description of the NetFlow Exporter.
7. Select the format as either **NetFlow or CEF**.
    a. If you select NetFlow, select the version as either **NetFlow-v9, NetFlow-v5,** or **IPFIX**. After each template-refresh-interval, the record template is sent to the collector. Also, the option template is sent. Enter the value from 1 - 216000 secs for Template Refresh Interval.
    b. If you select CEF, the version is 23, and enter the value from 1 - 216000 secs for Template Refresh Interval.
8. Enable the **SNMP** for the SNMP packet support on IP interfaces associated with the NetFlow Exporter.
9. By default, the UDP port of the collector is selected for the transport protocol.
10. Select either the **IPv4 or IPv6** of the destination IP. Default is set as IPv4.
11. Enter the IP address of the NetFlow/IPFIX collector. Default is set as 0.0.0.0.
12. Enter the **Port** range from 1 - 65535 for the destination IP. Default is set as 2055.
13. Enter the DSCP priority of the packet range from 0 - 63. Default is set as 0.
14. Enter the Time to Live (TTL) of the packet range from 1 - 255. Default is set as 64.
15. Select a rule, from the Condition list.
16. Click **Apply**. The **NetFlow Exporter** page appears with the configured.

Configure an IP Interface

To configure an IP Interface:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
    a. In the left menu, go to **System > Ports > Ports**. The **All Ports** page appears.
    b. Click the **Quick Port Editor** button to open the Quick Port Editor.

c. In the Quick Port Editor select the port to use for the IP interface, provide an alias to help identify the port (for example, NetFlow_Tunnel_Port), select **Tool** for the port type, and select **Enable**.

d. Click **OK**.

2. Select **Ports > IP Interfaces**.
3. Click **New**.
4. In the IP Interface page, perform the following:
   a. Enter the **Alias** and **Description** for the IP interface.
   b. Select the tool port from the **Port** list, that you configured in Step 1.
   c. Select the type of IP interface as either **IPv4** or **IPv6**.
   d. Enter the I**P Address, IP Mask, Gateway address,** and **MTU** value. (For example, 10.1.1.10, 255.255.255.0, 10.1.1.1)
   e. Select the GigaSMART group from the **GigaSMART Group** list you created in Step 1.
   f. Select the NetFlow exporter from the **Exporter** list you created in Step 2.

Configure a NetFlow Record

To configure a NetFlow Record:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > NetFlow / IPFIX Generation > Record**.
4. Click **New**. The **New NetFlow Record** page appears.
5. In the NetFlow Record page, perform the following:
   a. Enter a name (NetFlow-record) for the NetFlow Record in **Alias** field.
   b. (Optional) Enter a description of the NetFlow record.
   c. The Sampling Rate is multi-rate only, and is specified as 1 in N, where N is the packet count. The packet count can be a number from 1 to 16000. By default, the Sampling Rate is disabled.
   d. Select the **Exporter** that you want from the Exporters.
   e. Select the version either **NetFlow-v9 or IPFIX**. The NetFlow version must be configured with the same version of the Exporter and the Record. By default, NetFlow-v9 is the selected. NetFlow-v9 and IPFIX let you configure Match/Key and

Collect/Non-Key elements. Ensure that you configure the NetFlow version prior to configuring the match and collect parameters because the subsequent parameters depend on the NetFlow version configured.

f. The parameters that identify unique flows. The available Match/Key fields are based on the configured NetFlow version.

g. The parameters that identify you want to collect for the unique flows. The number of Collect/Non-Key elements in a record can be up to 32. Select the following option:

- **Exporter-** You can collect IPv4 or IPv6 switch management interface address by including additional collects in their configuration. You can select any of the following options:
  - **IPv4 address**—Adds new switch or router management IPv4 address.
  - **IPv6 address**—Adds new switch or router management IPv6 address.

Configure a NetFlow Monitor

Configure a NetFlow Generation Monitor and associate the NetFlow Generation Record to the specified NetFlow Generation Monitor by doing the following:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > NetFlow / IPFIX Generation > Monitor**.
4. Click **New**. The **Monitor** page appears.
5. In the NetFlow Monitor page, perform the following:

   a. Enter a name (NetFlow-monitor) for the NetFlow Record in **Alias** field.

   b. (Optional) Enter a **description** of the NetFlow monitor.

   c. Configure the Cache parameters.

   - **Cache Type**—Set as Normal.
   - **Cache Timeout Active**—Even if the flow is still active, it is "flushed out" to the Exporter once the timeout (in seconds) expires.
   - **Cache Timeout Inactive**—Inactive flows are flushed out to the Exporter once the timeout expires (in seconds) expires.
   - **Cache Timeout Event**—Applies to TCP flows. A flow is "flushed out" to the Exporter when a FIN or RST is detected.
   - **Mode**—Select the sampling mode:
     - No sampling
     - Multi rate
     - Single rate

- **Single Sampling Rate**—*<Need Input>*

d. Configure the Sampling parameters.

6. Click **Save**.

Add the NetFlow Monitor to GigaSMART Group

To Add the NetFlow Monitor to GigaSMART Group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART Groups**.
4. Select the GigaSMART Group configured, and then click **Edit**.
5. Under GigaSMART Parameters, go to NetFlow. Select the **NetFlow monitor** from the Monitor list.
6. Click **Save**.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

NOTE: For a single NetFlow GigaSMART Operation, make sure that you create a Regular By Rule map. When combining with APF or De-duplication, use First Level or Single Level map types. Make sure that the other combining GigaSMART Operations are configured before creating maps using NetFlow. When combining NetFlow with APF or De-duplication, create virtual ports to use with the second level maps. The destination tool port must be the IP interface with tool port identified when IP Interface is configured.

Configure Traffic Policy for a First-Level Map

To create a first-level map to forward traffic from network port 1/1/g3 to virtual port vp:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/g3.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle.
   d. Select MAC Source from the **Conditions** list, and then enter 00:a0:d1:e1:02:01 in the **Min** field.
   e. Enter MAC Mask as 0000.0000.0000.
   f. Click **Save**. The added rule appears in the **Source Ruleset** page.
   g. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears.

b. Select the configured virtual port vp1.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

5. Deploy Traffic Policy.

   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

6. Verify deployment.

   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure Traffic Policy for a Second-Level Map

To create a second-level map to forward traffic from the virtual port vp to tool port 1/1/g1:

1. Access Traffic Policy.

   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

   b. Create a new map with **Alias** as test1b.

2. Select Source.

   a. Click **Sources**. The **Source 1** page appears with the list of available ports.

   b. Select the virtual port (vp1).

   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.

3. Select Application.

   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.

   b. Select either a standalone node or a cluster from the **Select Node** list.

   c. You can either select an existing GSOP or create a new GSOP by clicking **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

   e. Select the GigaSMART group (gs-grp) you created earlier.

   f. Select Adaptive Packet Filtering from the **Applications** list. By default, the application you select is enabled.

g. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

4. Select Application Ruleset.
   a. Expand the **Application Ruleset** component and then drag and drop a **By Rule**. The **Application Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the configured traffic.
   d. Select IP Version from the **Condition** list and Version as Version 4.
   e. Click **Save**. The **Application Ruleset** page appears with the configured rule.
   f. Click **Close**. The By Rule block appears on the canvas with the configured application rule.
5. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.
   b. Select the tool port 1/1/g1.
   c. Click **Add**. The Destination block appears on the canvas with the selected tool ports.
6. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
7. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

GigaSMART Packet Slicing

Required License: Base
Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus.

GigaSMART operations with a Slicing selected truncate packets after either a specified header/layer and offset (a relative offset) or at a specific offset. Slicing operations are typically configured to preserve specific packet header information, allowing effective network analysis without the overhead of storing full packet data.

Packets can have multiple variable-length headers, depending on where they are captured, the different devices that have attached their own headers along the way, and the protocols in use (for example, IPv4 versus IPv6). Because of this, slicing operations with a hard-coded offset will not typically provide consistent results.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port- 1/1/x1
Source Ruleset	By Rule • Condition - IP Version and Version - 4
Application Ruleset	N/A
Applications	My GSOP - Slicing • Protocol - IPv4 and Offset - 77
Destination	Tool Port - 1/1/x3

Configure a Traffic Policy for Packet Slicing.

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Source block appears on the canvas with the selected network port.

3. Select Source Ruleset
   a. Expand the Source Ruleset component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the IPv4 traffic.
   d. Select IP Version from **Condition** list and select Version 4 for **Version** list.
   e. Click **Save** to apply the configured Ruleset. The Ruleset block appears in the canvas.
4. Select Application.
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
   e. Select the GigaSMART group (gs-grp) that you created earlier.
   f. Select Slicing from the **Applications** list. By default, the application you select is enabled.
   g. In the Slicing section, select IPv4 protocol and enter 77 in the **Offset** field.
   h. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.
5. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.
   b. Select the tool port 1/1/x3.

   > **NOTE:** To configure or modify the port type, use the **Quick Port Editor**.

   c. Click **Add**. The Destination block appears on the canvas with the selected tool port.
6. Deploy Traffic Policy
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
7. Verify deployment
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

GigaSMART Tunneling Operations

The following sections explain the Traffic Policy configurations of GigaSMART tunneling operations for different real-time example.

- GigaSMART IP Encapsulation/Decapsulation
- GigaSMART Layer-2 GRE Tunnel Encapsulation/Decapsulation
- GigaSMART VXLAN Tunnel Decapsulation
- GigaSMART ERSPAN Tunnel Decapsulation
- GigaSMART Custom Tunnel Decapsulation

GigaSMART IP Encapsulation/Decapsulation

Required License for IP Decapsulation: Base (GigaVUE-HC3), Tunneling (GigaVUE-HC1)
Required License for IP Encapsulation: Advanced Tunneling (GigaVUE-HC3), Tunneling (GigaVUE-HC1)

Use GigaSMART encapsulation and decapsulation operations to send traffic arriving on one GigaSMART-enabled node over the Internet to a second GigaSMART-enabled node. There, the traffic is decapsulated and made available to local tool ports.

This feature is useful when instrumenting remote data centers – you can tunnel selected portions of the traffic from the remote GigaSMART-enabled node to tools in a central location. Traffic is encapsulated at the sending end of the tunnel and decapsulated at the receiving end.

GigaSMART drops packets when it receives packets with multiple encapsulation headers that are added due to unpredictable looping in the network.

IP fragmentation and reassembly are supported. Refer to IP Fragmentation and Reassembly on L2GRE and GMIP Tunnels.

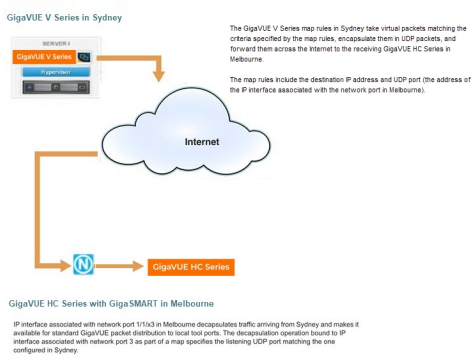
The source of the GigaSMART tunnel can be any of the following:

GigaSMART-Enabled GigaVUE H Series Node

- Standalone GigaVUE-HC3 node with SMT-HC3-C05 modules installed.
- Standalone GigaVUE-HC1 nodes.
- Any GigaVUE HC Series node operating in a cluster with the previous node types.
- GigaVUE V Series node or a GigaVUE-VM

Example: GigaSMART Encapsulation/Decapsulation (GigaVUE V Series

The following figures demonstrate how to create a sample IP tunnel between a sending GigaVUE V Series node in Sydney and a receiving GigaVUE HC Series in Melbourne. First, the overall tunnel is summarized, followed by configuration descriptions for the sending and receiving ends.



Configure Sending End of Tunnel: GigaVUE V Series map rules in Sydney

A GigaVUE V Series node in this location is configured with map rules that will send data over the Internet to the IP interface associated with a network port on a GigaVUE HC Series.

Map rules are created in the GigaVUE-FM user interface – Step 2 in the Create Map wizard includes **Tunnel Traffic To** settings that specify where matching traffic should be sent:

Create "Tunnel Traffic To" Option	Setting
UDP IP	This is the destination IP address for the IP interface associated with network port on the GigaVUE H Series in Melbourne. We will set it to 10.150.68.222
UDP Source Port	This is the UDP source port from which tunneled packets will be sent. We will set this to 5000.
UDP Destination Port	This is the listening port on the receiving GigaVUE H Series IP interface associated with network port. We will set this to 10000.

Configure Receiving End of Tunnel: GigaVUE HC Series with GigaSMART in Melbourne

Now we need to configure the receiving end of the tunnel with an IP interface associated with network port. The GigaVUE H Series in this location will have an IP interface associated with network port configured on network port 1/1/3 with an IP address of 10.150.68.222 and a GigaSMART decapsulation operation that listens on UDP port 10000.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port (IP Interface) - 1/1/x3
Source Ruleset	By Rule • Condition - IP Source, IPv4 Address - 10.10.10.10. • Net Mask - 255.255.255.255
Application Ruleset	N/A
Applications	Tunnel Decapsulation • GMIP Port - 10000.
Destination	Tool port - 1/1/x4

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure an IP Interface

To configure IP interface:

1. In GigaVUE-FM, go to **Inventory > Nodes**.

2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > Ports > IP Interfaces**. The **IP Interfaces** page appears.
4. Click **New.** The **New IP Interface** pane appears.
5. Configure the following fields as follows:
   - **Alias** - Tunnel_IP
   - **Port** - 1/1/x3
   - **Type**: IPv4
   - **IP Address** - 10.150.68.222
   - **IP Mask** - 255.255.255.255
   - **Gateway** - 10.150.68.1
   - **MTU** - 9400
   - **GigaSMART Group** - gs-grp
6. Click **Apply**. The created IP interface appears in the **IP Interfaces** page.

Configure a Traffic Policy

To create a Traffic Policy:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x3 configured at the IP interface as source.
   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset

a. Expand the Source Ruleset component and then drag and drop **By Rule**. The **Source Ruleset** page appears.

b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle to pass the IPv4 traffic.

d. Select IP Source from the **Condition** list, and then enter 10.10.10.10 in the **IPv4 Address** field.

e. Enter 255.255.255.255 in the **Net Mask** field.

f. Click **Save**. The added rule appears in the **Source Ruleset** page.

g. Click **Close**. The By Rule block appears on the canvas with the configured rule.

4. Select Application.

a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group (gs-grp) you created earlier.

f. Select Tunnel Decapsulation from the **Applications** list. The Tunnel Decapsulation section appears with the type list.

g. Select GMIP from the list and enter 10000 in the **GMIP Port** field.

h. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool port 1/1/x4.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Destination block appears on the canvas with the selected tool port.

6. Deploy Traffic Policy

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation

Required License for L2GRE Decapsulation: Base (GigaVUE-HC3), Tunneling (GigaVUE-HC1)Required License for L2GRE Encapsulation: Advanced Tunneling (GigaVUE-HC3), Tunneling (GigaVUE-HC1)

Supported Devices: GigaVUE-HC1 Gen 2, GigaVUE-HC3 Gen 2, GigaVUE-HC1-Plus, GigaVUE-HC1 Gen 3

GigaVUE-HC1 Gen 3 module supports IPv6 for L2GRE tunneling.

Use GigaSMART Layer 2 (L2) Generic Routing Encapsulation (GRE) tunnel encapsulation to send traffic from one GigaSMART node over the Internet to a second GigaSMART node using L2GRE encapsulation. Use GigaSMART L2GRE tunnel decapsulation at the second GigaSMART node to decapsulate the traffic before sending it to local tool ports.

GigaSMART Layer 2 GRE tunnel encapsulation/decapsulation provides the following:

- L2GRE tunnel initiation and encapsulation on the tool port at the sending end of the tunnel (for example, at a remote site)
- L2GRE tunnel termination and decapsulation on the network port at the receiving end of the tunnel (for example, at a main office site)

The GigaSMART at the remote site encapsulates the filtered packets, adds an encapsulation header, and routes it to the main office site. The encapsulation protocol is GRE and the delivery protocol is IP or IPv6, so the encapsulation header consists of Ethernet + IP + GRE or Ethernet + IPv6 + GRE headers.

The parameters of the encapsulated header are user-configurable, such as the IPv4 address of the IP interface on the destination GigaSMART node and the GRE key that identifies the source of the tunnel.

At the remote end, packets are decapsulated, the L2GRE header is stripped off, and packets are sent to the specified tool port.

IP fragmentation and reassembly are supported.

The below Figure shows the remote site encapsulating the filtered traffic and routing it to the main office from the remote end.

The encapsulated packet is sent out of the tool port, which is connected to the public network (the Internet). This packet is routed in the public network to reach the main office site. It ingresses at the routed network port of the GigaVUE node at the main office.

The ingress encapsulated packet is then sent to the GigaSMART at the main office, where the packet is decapsulated and sent to the tool port. The received packet's destination IP is checked against the source IP/IPv6 configured for the network port. If they match, decapsulation is applied. The Ethernet + IP + GRE or Ethernet + IPv6 + GRE header is stripped and the remaining packet is sent to the tool port.

NOTE: IPv6 addresses are not supported on SMT-HC1-S (Generation 3 GigaSMART module on GigaVUE-HC1.

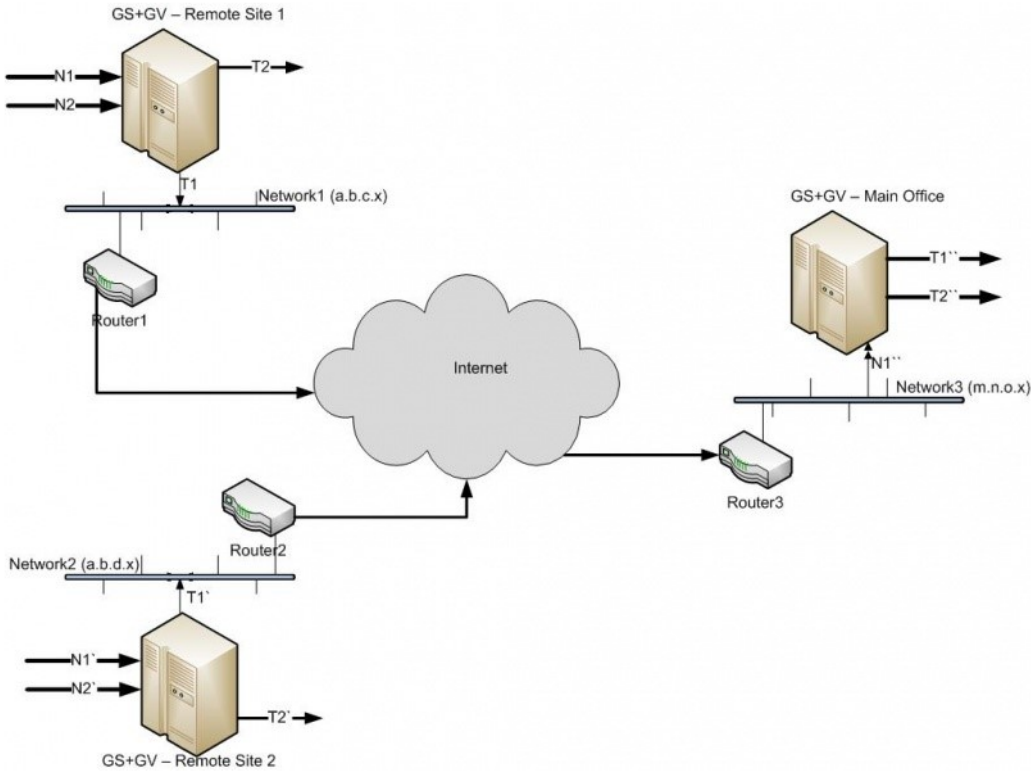


Figure 1 *L2GRE Tunnel Encapsulation/Decapsulation*

Layer 2 GRE Header Length

The L2GRE header length is as follows:

Header	Length in Bytes
With Key	42 bytes consisting of 14 Ethernet + 20 IP + 4 GRE + 4 GRE Key.
Without Key	38 bytes consisting of 14 Ethernet + 20 IP + 4 GRE.

GigaSMART L2GRE Tunnel Encapsulation

This topic explains the workflow required to configure a Traffic Policy for L2GRE tunnel encapsulation.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	IP Interface - 1/1/x3
Source Ruleset	By Rule • Condition - IP Version, Version - Version 4.
Application Ruleset	N/A
Applications	Tunnel Encapsulation. • Type - L2GRE, Destination - IPv4
Destination	Tool port - 1/1/x3

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.

6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure an IP Interface

To configure IP interface:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > Ports > IP Interfaces**. The **IP Interfaces** page appears.
4. Click **New.** The **New IP Interface** pane appears.
5. Configure the following fields as follows:
   - **Alias** - Tunnel_IP
   - **Port** - 1/1/x3
   - **Type**: IPv4
   - **IP Address** - 10.150.68.222
   - **IP Mask** - 255.255.255.255
   - **Gateway** - 10.150.68.1
   - **MTU** - 9400
   - **GigaSMART Group** - GS2
6. Click **Apply**. The created IP interface appears in the **IP Interfaces** page.

Configure a Traffic Policy

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source
    a. Click **Sources**. The **Source 1** page appears.
    b. Select the network port 1/1/x3 configured at the IP interface as source.
    c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset
    a. Expand the Source Ruleset component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
    b. Click **Add Rules**. The **Add Ruleset** pane appears.
    c. Turn on **Pass** toggle to pass the IPv4 traffic.
    d. Select IP Version from the **Condition** list, and then select Version 4 from the **Version** list.
    e. Click **Save**. The added rule appears in the **Source Ruleset** page.
    f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Application.
    a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
    b. Select either a standalone node or a cluster from the **Select Node** list.
    c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
    d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
    e. Select the GigaSMART group (gs-grp) you created earlier.
    f. Select Tunnel Encapsulation from the **Applications** list. The Tunnel Encapsulation section appears.
    g. Select L2GRE as the type and select IPv4 for the **Destination** option.
    h. Enter **Destination IP** address and **Key** value.
    i. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.
5. Select Destination.
    a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.
    b. Select the tool port 1/1/x4.

    NOTE: To configure or modify the port type, use the **Quick Port Editor**.

c. Select **Encapsulation Tunnel** button in the top right corner of the Destination 1 page. The Add Encapsulation dialog appears.

d. Select the node from the **Select Node** list and select the tunnel from the **Encapsulation Tunnel** list.

> **NOTE:** Source port and destination port should be from the same cluster to apply the tunnel.

e. Click **Save**. The **Destination 1** page appears.

f. Click **Cancel**. The Destination block appears on the canvas with the selected tool port.

6. Deploy Traffic Policy

   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

GigaSMART L2GRE IPv6 Tunnel Encap/Decap with Load Balancing

This topic explains the workflow required to configure a Traffic Policy for L2GRE tunnel encapsulation.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	IP Interface - 1/3/x8
Source Ruleset	By Rule • Condition - IP Version, Version - 6
Applications	Tunnel Encapsulation and Load Balancing • Tunnel Encapsulation: Type - L2GRE, Destination - Port Group

Map Components	Sample Configuration
	• Load Balancing: Stateful, Metric - Round Robin
Application Ruleset	N/A
Destination	Tool port group

Configure a Port Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter portgrp1 in the **Alias** field for the new Port Group and click **Tool** option.
5. Enable **Load Balancing**.
6. Choose the **Tool Ports** you want to attach to the port group, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
7. Specify the weights for each port as follows:
   - Weight 1/1/x6 - 5
   - Weight 1/1/x7 - 10
   - Weight 1/2/x3 - 20
   - Weight 1/2/x4 - 10
8. Click **Apply** to save. The created port group appears in the port group table.

Configure GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure an IP Interface

To configure an IP interface:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > Ports > IP Interfaces**. The **IP Interfaces** page appears.
4. Click **New.** The **New IP Interface** pane appears.
5. Configure the following fields as follows:
   - **Alias** - Tunnel_IP
   - **Port** - 1/3/x8
   - **Type**: IPv6
   - **IP Address** - 10.1.1.10
   - **IP Mask** - 255.255.255.0
   - **Gateway** - 10.1.1.1
   - **MTU** - 1500
   - **GigaSMART Group** - GS2
6. Click **Apply**. The created IP interface appears in the **IP Interfaces** page.

Configure a Traffic Policy

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.

b. Select the network port 1/3/x8 configured at the IP interface as source.

c. Click **Add**. The Source block appears on the canvas with the selected network port.

3. Select Source Ruleset

a. Expand the Source Ruleset component and then drag and drop **By Rule**. The **Source Ruleset** page appears.

b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle to pass the IPv4 traffic.

d. Select IP Version from the **Condition** list, and then select Version 6 from the **Version** list.

e. Click **Save**. The added rule appears in the **Source Ruleset** page.

f. Click **Close**. The By Rule block appears on the canvas with the configured rule.

4. Select Application.

a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group (GS2) you created earlier.

f. Select Tunnel Encapsulation and Load Balancing from the **Applications** list. The Tunnel Encapsulation and Load Balancing sections appear.

- In the Tunnel Encapsulation section, select L2GRE as the type and select Port Group for the **Destination** option.
- Enter the **Destination IP** address and **Key** value.
- In the Load Balancing section, select Stateful option.
- Select Tunnel from the **Application** list and Round Robin from the **Metric** list.

g. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool port group that you created earlier.

c. Click **Cancel**. The Destination block appears on the canvas with the selected tool port.

6. Deploy Traffic Policy
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
7. Verify deployment
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

GigaSMART L2GRE Tunnel Encapsulation Stateless Load Balancing

This topic explains the workflow required to configure a Traffic Policy for L2GRE tunnel encapsulation with stateless load balancing.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	IP Interface - 1/1/x3
Source Ruleset	By Rule • Condition - IP Version, Version - Version 4.
Application Ruleset	N/A
Applications	Tunnel Encapsulation. • Type - L2GRE, Destination - Port Group • Load Balancing: Stateless, Metric - Five Tuple
Destination	Tool port group

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.

3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure an IP Interface

To configure an IP interface:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > Ports > IP Interfaces**. The **IP Interfaces** page appears.
4. Click **New.** The **New IP Interface** pane appears.
5. Configure the following fields as follows:
   - **Alias** - Tunnel_IP
   - **Port** - 1/1/x3
   - **Type**: IPv4
   - **IP Address** - 10.150.68.222
   - **IP Mask** - 255.255.255.255
   - **Gateway** - 10.150.68.1
   - **MTU** - 9400
   - **GigaSMART Group** - GS2
6. Click **Apply**. The created IP interface appears in the **IP Interfaces** page.

Configure a Tunnel Endpoint

To configure tunnel endpoints:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > Ports >Tunnel Endpoints**.
4. Click **New**. The **New Tunnel Endpoint** page appears.
5. Enter the ID in **Tunnel Endpoint ID** field.
6. Select IPv4 or IPv6 for the **Destination** option.

7. Enter the IPv4 or IPv6 address based on the select destination.
8. Click Apply. The **Tunnel Endpoints** page appears with the configured endpoint.

Configure a Port Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter portgrp1 in the **Alias** field for the new Port Group and click **Tool** option.
5. Enable **Load Balancing**.
6. Choose the **Tool Ports** you want to attach to the port group, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
7. Specify the weights for each port as follows:
   - Weight 1/1/x6 - 5
   - Weight 1/1/x7 - 10
   - Weight 1/2/x3 - 20
   - Weight 1/2/x4 - 10
8. Select the configured tunnel endpoint from the **Tunnel Endpoints** list.
9. Click **Apply** to save. The created port group appears in the port group table.

Configure a Traffic Policy

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.

a. Click **Sources**. The **Source 1** page appears.

b. Select the network port 1/1/x3 configured at the IP interface as source.

c. Click **Add**. The Source block appears on the canvas with the selected network port.

3. Select Source Ruleset

a. Expand the Source Ruleset component and then drag and drop **By Rule**. The **Source Ruleset** page appears.

b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle to pass the IPv4 traffic.

d. Select IP Version from the **Condition** list, and then select Version 4 from the **Version** list.

e. Click **Save**. The added rule appears in the **Source Ruleset** page.

f. Click **Close**. The By Rule block appears on the canvas with the configured rule.

4. Select Application.

a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group (GS2) you created earlier.

f. Select Tunnel Encapsulation and Load Balancing from the **Applications** list. The Tunnel Encapsulation and Load Balancing sections appear.

- In the Tunnel Encapsulation section, select L2GRE as the type and select Port Group for the **Destination** option.
- Enter the **Destination IP** address and **Key** value.
- In the Load Balancing section, select Stateless option.
- Select Five Tuple from the **Metric** list.

g. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool Port 1/1/x3.

> **NOTE:** To configure or modify the port type, use the **Quick Port Editor**.

c. Select **Encapsulation Tunnel** button in the top right corner of the Destination 1 page. The Add Encapsulation dialog appears.

d. Select the node from the **Select Node** list and select the tunnel from the **Encapsulation Tunnel** list.

> **NOTE:** Source port and destination port should be from the same cluster to apply the tunnel.

e. Click **Save**. The **Destination 1** page appears.

f. Click **Cancel**. The Destination block appears on the canvas with the selected tool port.

6. Deploy Traffic Policy

   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment

   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

GigaSMART L2GRE Tunnel Decapsulation

This topic explains the workflow required to configure a Traffic Policy for L2GRE tunnel encapsulation.

In this example, an IP interface is configured on the network port. A GigaSMART operation for tunnel decapsulation is configured to decapsulate the filtered packets. A map is configured that uses the L2GRE tunnel decapsulation GigaSMART operation, which receives packets from the remote site over the Internet to the main office using the IP interface associated with a tool port and then forwards packets over the tool port. Staring with software version 5.4 GigaSMART L2GRE Tunnel Decapsulation provides support for IPv6 with load-balancing.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port (IP Interface) - 1/1/x3
Source Ruleset	By Rule • Condition - IP Version, Version - Version 4.
Application Ruleset	N/A
Applications	My GSOP - Tunnel Decapsulation. • Type - L2GRE, Key - GRE
Destination	Tool port - 1/1/x3

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure an IP Interface

To configure IP interface:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > Ports > IP Interfaces**. The **IP Interfaces** page appears.
4. Click **New.** The **New IP Interface** pane appears.
5. Configure the following fields as follows:
    - **Alias** - Tunnel_IP
    - **Port** - 1/1/x3
    - **Type**: IPv4
    - **IP Address** - 10.150.68.222

- **IP Mask** - 255.255.255.255
- **Gateway** - 10.150.68.1
- **MTU** - 9400
- **GigaSMART Group** - GS2

6. Click **Apply**. The created IP interface appears in the **IP Interfaces** page.

Configure a Traffic Policy

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x3 configured at the IP interface as source.
   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset
   a. Expand the Source Ruleset component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the IPv4 traffic.
   d. Select IP Version from the **Condition** list, and then select Version 4 from the **Version** list.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Application.

a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group (GS2) you created earlier.

f. Select Tunnel Decapsulation from the **Applications** list. The Tunnel Decapsulation section appears.

g. Select L2GRE as the type and enter the GRE key value in the **Key** field.

h. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool port 1/1/x3.

> **NOTE:** To configure or modify the port type, use the **Quick Port Editor**.

c. Click **Cancel**. The Destination block appears on the canvas with the selected tool port.

6. Deploy Traffic Policy

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

GigaSMART L2GRE Tunnel Encapsulation Load Balancing

This topic explains the workflow required to configure a Traffic Policy for L2GRE tunnel encapsulation with stateful load balancing.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port (IP Interface) - 1/1/x3
Source Ruleset	By Rule • Condition - IP Version, Version - Version 4.
Application Ruleset	N/A
Applications	My GSOP - Tunnel Encapsulation and Load Balancing • Tunnel Encapsulation: Type - L2GRE, Key - GRE • Load Balancing: Stateful, Metric - Round Robin
Destination	Tool port group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Create IP Interface

To configure IP interface:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > Ports > IP Interfaces**. The **IP Interfaces** page appears.
4. Click **New.** The **New IP Interface** pane appears.
5. Configure the following fields as follows:
    - **Alias** - Tunnel_IP
    - **Port** - 1/1/x3
    - **Type**: IPv4

- **IP Address** - 10.150.68.222
- **IP Mask** - 255.255.255.255
- **Gateway** - 10.150.68.1
- **MTU** - 9400
- **GigaSMART Group** - GS2

6. Click **Apply**. The created IP interface appears in the **IP Interfaces** page.

Configure a Tunnel Endpoint

To configure tunnel endpoints:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > Ports >Tunnel Endpoints**.
4. Click **New**. The **New Tunnel Endpoint** page appears.
5. Enter the ID in **Tunnel Endpoint ID** field.
6. Select IPv4 or IPv6 for the **Destination** option.
7. Enter the IPv4 or IPv6 address based on the select destination.
8. Click Apply. The **Tunnel Endpoints** page appears with the configured endpoint.

Configure a Port Group

To create a port group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new port group. The New Port Group page appears.
4. Enter portgrp1 in the **Alias** field for the new Port Group and click **Tool** option.
5. Enable **Load Balancing**.
6. Choose the **Tool Ports** you want to attach to the port group, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
7. Specify the weights for each port as follows:
   - Weight 1/1/x6 - 5
   - Weight 1/1/x7 - 10
   - Weight 1/2/x3 - 20
   - Weight 1/2/x4 - 10

8. Select the configured tunnel endpoint from the **Tunnel Endpoints** list.
9. Click **Apply** to save. The created port group appears in the port group table.

Configure Traffic Policy

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x3 configured at the IP interface as source.
   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset
   a. Expand the Source Ruleset component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the IPv4 traffic.
   d. Select IP Version from the **Condition** list, and then select Version 4 from the **Version** list.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Application.
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group (gs-grp) you created earlier.

f. Select Tunnel Encapsulation and Load Balancing from the **Applications** list. The Tunnel Encapsulation and Load Balancing sections appear.

- In the Tunnel Encapsulation section, select L2GRE as the type and select Port Group for the **Destination** option.
- Enter the **Destination IP** address and **Key** value.
- In the Load Balancing section, select Stateful option.
- Select Tunnel from the **Application** list and Round Robin from the **Metric** list.

g. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool port group.

> **NOTE:** To configure or modify the port type, use the **Quick Port Editor**.

c. Select **Encapsulation Tunnel** button in the top right corner of the Destination 1 page. The Add Encapsulation dialog appears.

d. Select the node from the **Select Node** list and select the tunnel from the **Encapsulation Tunnel** list.

> **NOTE:** Source port and destination port should be from the same cluster to apply the tunnel.

e. Click **Save**. The **Destination 1** page appears.

f. Click **Cancel**. The Destination block appears on the canvas with the selected tool port.

6. Deploy Traffic Policy

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

GigaSMART VXLAN Tunnel Decapsulation

Required License for VXLAN Decapsulation: Base (GigaVUE-HC3), Tunneling (GigaVUE-HC1)

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus.

This example includes VXLAN tunnel termination. You can configure a GigaSMART-enabled node to act as the receiving end of a VXLAN tunnel by configuring a GigaSMART tunnel decapsulation operation with type set to VXLAN.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port (IP Interface)- 1/2/x1
Source Ruleset	By Rule • Condition - IP Version, Version - 4
Application Ruleset	N/A
Applications	My GSOP - Tunnel Encapsulation • Tunnel Encapsulation: Type - VXLAN
Destination	Tool Port - 1/4/x2

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.

4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure an IP Interface

To configure IP interface:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > Ports > IP Interfaces**. The **IP Interfaces** page appears.
4. Click **New.** The **New IP Interface** pane appears.
5. Configure the following fields as follows:
   - **Alias** - Tunnel_IP
   - **Port** - 1/2/x1
   - **Type**: IPv4
   - **IP Address** - 10.115.9.5
   - **IP Mask** - 255.255.255.0
   - **Gateway** - 10.115.9.1
   - **MTU** - 1500
   - **GigaSMART Group** - gs-grp
6. Click **Apply**. The created IP interface appears in the **IP Interfaces** page.

Configure a Traffic Policy

To configure a Traffic Policy:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.

- (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
- (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.

d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/2/x1 configured at the IP interface as source.
   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset
   a. Expand the Source Ruleset component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the IPv4 traffic.
   d. Select IP Version from the **Condition** list, and then select Version 4 from the **Version** list.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Application.
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
   e. Select the GigaSMART group (gs-grp) you created earlier.
   f. Select Tunnel Decapsulation from the **Applications** list. The Tunnel Decapsulation section appears with the type list.
   g. Select VXLAN as the type in the Tunnel Decapsulation section.
   h. Enter the values for **Source Port**, **Destination Port**, and **VNI** as required.
   i. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.
5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool port 1/4/x2.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Destination block appears on the canvas with the selected tool port.

6. Deploy Traffic Policy

   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment

   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

GigaSMART ERSPAN Tunnel Decapsulation

Required License for ERSPAN Decapsulation: Advanced Tunneling (GigaVUE-HC1-HC1 Gen3, GigaVUE-HC1-HC3, GigaVUE-HC1-HC1-Plus, and GigaVUE-HC1-HCT Gen 3), Tunneling (GigaVUE-HC1, and GigaVUE-HC1-HCT Gen 3).
Supported Devices: GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus, and GigaVUE-HCT Gen 3.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port (IP interface) - 1/1/g2
Source Ruleset	By Rule • Rule 1: IPV4 Protocol, Value - GRE
Application Ruleset	N/A
Applications	My GSOP - Tunnel Decapsualtion • FlowID - 0
Destination	Tool Port Group - PG-Sample

Configure an IP Interface

To configure IP interface:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > Ports > IP Interfaces**. The **IP Interfaces** page appears.
4. Click **New.** The **New IP Interface** pane appears.
5. Configure the following fields as follows:
   - **Alias** - 1_1_x3
   - **Port** - 1/1/g2
   - **Type**: IPv4
   - **IP Address** - 10.10.10.10
   - **IP Mask** - 255.255.225.0
   - **Gateway** - 0.10.10.1
   - **MTU** - 1500
   - **GigaSMART Group** - gs-grp
6. Click **Apply**. The created IP interface appears in the **IP Interfaces** page.

Configure a Port Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter PG-Sample in the **Alias** field for the new port group and click **Tool** option.
5. Choose the **Tool Ports** you want to attach to the port group, For example 1/1/g1.

Configure a GigaSMART Group

6. Click **Apply** to save. The created port group appears in the port group table.

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.

4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Traffic Policy for Regular Map.

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/g2 configured at the IP interface as source.
   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules** to add Rule 1. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the configured traffic.
   d. Select IPv4 Protocol from the **Condition** list and select GRE for **Value**.
   e. Click **Save** to apply the configured Ruleset.
4. Select Application.
   a. Expand the **Application** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

d. You can either select an existing GigaSMARTgroup or create a new GSOP by clicking **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

e. Enter gsop1 in the **Alias Field** to name the GSOP. Select the (gs-grp) from the GigaSMART Groups list..

f. Select Tunnel Decapsulation from the **GigaSMART Operations** list. Select ERSPAN for the type.

g. Enter 0 in the **Flow ID field**.

h. Click **Save**.

i. Select the created GSOP from the **GSOP** list.

j. Click **Add**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

5. Select Destination.

   a. Click **Destinations**. The **Destination** page appears with the list of available ports.

   b. Select tool port group for the Destination, that you have created earlier.

   c. Click **Save** to apply. The Destination block appears on the canvas with the selected tool port.

6. Deploy Traffic Policy.

   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

GigaSMART Custom Tunnel Decapsulation

Required Licenses for Custom Decapsulation: Base (GigaVUE-HC3), Tunneling (GigaVUE-HC1),and Header Stripping

Use custom tunnel termination to terminate a custom tunnel header that is received at the IP interface that is associated with a network port, but is not known to GigaSMART. The destination IP and MAC addresses must match the IP and MAC addresses of the network tunnel. The packets that are successfully received at GigaSMART on a custom tunnel can be stripped, after some validations are performed, or can be sent to tools. The existing generic Header Stripping operation can be leveraged to remove the tunnel header if required. The reassembly of fragmented IP packets is also supported.

You can configure a GigaSMART-enabled node to act as the receiving end of a VXLAN tunnel by configuring a GigaSMART tunnel decapsulation operation with type set to custom and Layer 4 (L4) source and destination ports.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	IP interface - 1/1/g1
Snource Ruleset	By Rule • Condition - IP Version, Version - 4
Application Ruleset	N/A
Applications	My GSOP - Tunnel Decapsulation • Type - custom
Destination	1/1/x3

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure an IP Interface

To configure IP interface:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > Ports > IP Interfaces**. The **IP Interfaces** page appears.
4. Click **New.** The **New IP Interface** pane appears.
5. Configure the following fields as follows:
   - **Alias** - Tunnel_IP
   - **Port** - 1/1/g1
   - **Type**: IPv4
   - **IP Address** - 10.115.9.5
   - **IP Mask** - 255.255.255.255
   - **Gateway** - 10.115.9.1
   - **MTU** - 1500
   - **GigaSMART Group** - gs-grp
6. Click **Apply**. The created IP interface appears in the **IP Interfaces** page.

Configure a Traffic Policy

To create a Traffic Policy:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source

a. Click **Sources**. The **Source 1** page appears.

b. Select the network port 1/1/g1 configured at the IP interface as source.

c. Click **Add**. The Source block appears on the canvas with the selected network port.

3. Select Source Ruleset

a. Expand the Source Ruleset component and then drag and drop **By Rule**. The **Source Ruleset** page appears.

b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle to pass the IPv4 traffic.

d. Select IP Version from the **Condition** list, and then select Version 4 from the **Version** list.

e. Click **Save**. The added rule appears in the **Source Ruleset** page.

f. Click **Close**. The By Rule block appears on the canvas with the configured rule.

4. Select Application.

a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group (gs-grp) you created earlier.

f. Select Tunnel Decapsulation from the **Applications** list. The Tunnel Decapsulation section appears with the type list.

g. Select Custom as the type in the Tunnel Decapsulation section.

h. Enter the values for **Source Port** and **Destination Port** as required.

i. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool port 1/1/x3.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Destination block appears on the canvas with the selected tool port.

6. Deploy Traffic Policy

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment

   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure Traffic Policy for Application Intelligence Solutions

Required Licenses:

- Application Filtering Intelligence — Application Filtering Intelligence with Application Visualization capability.
- Application Metadata Intelligence — Application Metadata Intelligence with Application Visualization capability.
- NetFlow/IPFIX — NetFlow/IPFIX generation for IPv4 and IPv6. This allows to export only the standard elements and not application metadata.
- NetVUE
- NetVUE PLUS
- SecureVUE
- SecureVUE PLUS
- Zero Trust Architecture

NOTE: These licenses are valid for a year. After the expiration date, the system does not allow you to add new configuration or make any configuration updates but allows traffic flow for two weeks.

Supported Devices: GigaVUE-HC1, GigaVUE-HC1-Plus, GigaVUE-HC3, GigaVUE-HCT.

When you create a Traffic Policy for an Application Intelligence solution, you describe your intent for the monitored traffic, such as which applications to observe, which to pass or drop, and which to export as metadata, and GigaVUE-FM automatically builds the required map levels for you. You attach one or more Application Intelligence functions to sources (for example, network ports or vPorts) and Traffic Policy applies them as a unified, application-aware pipeline.

- Application Visualization (App Viz)—Identifies which applications are present in the traffic and how much bandwidth they consume, showing traffic statistics and top applications over time so you can understand the application mix before enforcing policies or exporting records.

- Application Filtering Intelligence (AFI)—Filters traffic by application, application family, or tags so you can pass critical applications and drop or divert irrelevant ones, ensuring that only policy-relevant traffic reaches your tools.
- Application Metadata Intelligence (AMI)—Generates and exports detailed flow and application metadata (for example, NetFlow/IPFIX/CEF) for selected applications to external collectors and observability tools, providing rich context for performance analysis and threat detection.

Together, AppViz, AFI, and AMI let you use Traffic Policy to start with simple visibility, then layer on application-aware filtering and metadata export as your monitoring and compliance needs evolve.

If you need a broader conceptual overview of how Application Intelligence solutions work, refer to Application Intelligence Solutions.

How Traffic Policy Choose Deployment Type for Application Intelligence Solutions

Intent-Based Map: When you create a Traffic Policy for the Application Intelligence solution. In this case, you do not design the individual maps yourself. Instead, you describe what you want the Application Intelligence solution to do with the traffic, and GigaVUE-FM builds the required map levels for you.

- When you create a Traffic Policy for an Application Intelligence solution:
  - You choose one or more Application Intelligence functions such as:
  - Application Visualization (AppViz) – see which applications are present.
  - Application Filtering Intelligence (AFI) – pass or drop selected applications.
  - Application Metadata Intelligence (AMI) – export flow and metadata records.
  - NetFlow - generate Layer 2-4 flow data.

  NOTE: This configuration is applicable only when using NetFlow License on Gen 3 cards.
- You attach these Application Intelligence functions to sources, for example, network ports or vPorts, by defining Application Intelligence intents:
  - For example, monitor all applications on the traffic arriving at specific source ports or vPorts, then filter traffic based on applications, and export metadata for selected applications.
- Based on this intent, Traffic Policy chooses the underlying map types automatically:
  - If you only include an Application, for example, Application Visualization, without Application Ruleset, GigaVUE-FM creates regular maps with GSOP, where possible. This keeps performance close to classic regular maps.

- If you include Application Ruleset such as GigaSMART rules or AFI rules, GigaVUE-FM creates a first level and second level map pair:
- The first level map sends L2–L4 hardwarefiltered traffic from the source to a vPort.
- The second level map takes traffic from the vPort, applies the Application Ruleset, for example, GigaSMART rules or AFI rules, and then runs the Application Intelligence functions (AFI/AMI).

- If you chain multiple Application Intelligence solutions that cannot run in one GigaSMART operation, GigaVUE-FM inserts transit-level maps between vPorts:
  - Each transit level map performs one GigaSMART operation and forwards traffic to the next vPort.
  - From these vPorts, GigaVUE-FM can create multiple prioritized second-level maps to branch traffic to different tools or exporters.

Application Intelligence – Rules and Notes

Keep in mind the following rules and notes when you use Application Intelligence in a Traffic Policy:

- You can deploy only one Application Intelligence solution (Application Visualization, Application Filtering Intelligence, and/or Application Metadata Intelligence) on a Gen 3 GigaSMART card or V Series node.
- When you back up and restore configurations that include traffic policies for Application Intelligence, you must back up both the devices and GigaVUE-FM, then restore the backed-up data on both.
- Application Intelligence in a Traffic Policy supports both asymmetric (unidirectional) and symmetric (bidirectional) traffic. However, for the most accurate classification and metadata, use symmetric traffic whenever possible.
- After you deploy or update a Traffic Policy that uses Application Intelligence, GigaVUE-FM can take up to 10 minutes to fully populate the Application Intelligence dashboards.
- On Gen 2 GigaSMART cards, when both AMI and AFI licenses are installed, packets that do not have a complete 5-tuple are dropped and are not forwarded to AMI for attribute extraction.
- The AFI license is optional for AMI. If you enable AMI without explicitly configuring AFI for that traffic, an AFI pass-all (“No-Rule-Match” pass) behavior is enabled by default so that AMI still receives the traffic.
- In the second-level maps of Application Intelligence, you cannot use pattern-matching (regex/gsrule) rules and AFI application filter rules together in the same Application Ruleset.
- Application Intelligence does not detect or classify ARP requests on either Gen 2 or Gen 3 devices. Packets that do not meet an Adaptive Packet Filtering (APF) match are dropped and not processed further by Application Intelligence.

- You can configure up to 120 user-defined applications per GigaVUE-FM instance. These user-defined applications can be used across one or more Application Intelligence sessions or traffic policies.
- You can configure a maximum of 8 rules per application. You can configure a maximum of 3 protocols per rule.
- When the AMI session table/cache size is set to 10M, support for long-duration TLS flows depends on the traffic profile and the enabled attributes. The recommended limit is 10K CPS with 5M active flows for GTP traffic and 14K CPS with 7M active flows for non-GTP traffic. If these limits are exceeded, some exported flow records might not include attributes such as TLS SNI. In some traffic conditions, Flow-ID reset can also occur.

Refer to the following sections for instructions on how to configure Traffic Policy for Application Intelligence solutions with real time examples:

- Use Case: Reduce Tool Overload
- Use Case: Enhance Context for Incident Response
- Use Case: Maintain Visibility Whilst Filtering Traffic

Create and Manage Custom Tool Template

A tool template pre-defines a list of applications and its attributes, which you can choose as per your requirements while configuring an Application Metadata Intelligence solution. A template once created can be used by multiple exporters to export the attributes in the specified format to the destination tools.

You can use the tool templates while creating an Application Metadata Intelligence session.

Create Custom Tool Template

To create a tool template

1. In the Traffic Policy landing page, go to **Resources > Tool Templates**. The Tool Templates page appears.
2. Click **New**. The New Tool Template pane appears. This includes Tool Template Info section, Application section, and Advanced Settings section.
3. In the **Tool Template Info** section, do the following:
   a. **Template Name** - Enter a name for the template.
   b. (Optional) **Description** - Enter a description for the template.

c. (Optional) **SNMP** - Turn on the toggle to enable this option.

> **NOTE:** Enable this option to respond to SNMP requests per exporter. To configure SNMP globally, go to the device's Global Settings.

d. **GS Version** - Select the GS Version from the list.

e. (Optional) **Fast Mode** - Turn on the toggle to enable this option.

> **NOTE:** Enable this option to improve the performance for monitoring HTTP and DNS traffic. It supports exporting only limited attributes. Once enabled, Fast mode cannot be disabled. When a solution is Fast Mode-enabled, you can only select the tool templates that has all the attributes supported in the fast mode. The other tool templates that are listed will be greyed out. Refer "Create an Application Intelligence Session" topic for information on fastmode.

f. **Format** - Select the format in which the application details need to be exported to the tool. Valid formats are CEF and Netflow. Select the Version and the Template Refresh Interval for the Netflow format.

g. **Record / Template Type** - Select the required type.

- Segregated - The application-specific attributes and the generic attributes will be exported as individual records to the tool.
- Cohesive- The application-specific attributes and the generic attributes will be combined as a single record and exported to the tool.

h. **Active Timeout** - Enter the value in seconds. This is the time Interval after which Application Metadata Intelligence exports flow record to ensure metadata is regularly updated for long-running flows.

i. **Inactive Timeout** - Enter the value in seconds. This is the time interval after which Application Metadata Intelligence expires flow entry and exports flow record when no packets are received.

j. **Maximum Packet Size Mode** - Turn on the toggle to configure the maximum packet size of the records that can be exported.

4. In the **Application** section, do the following:

a. **Application Editor** - Click this to open the App Editor page.

- In the App Editor page, search for and select the applications such as antivirus, webmail that needs to be filtered from the traffic. You can also select multiple application families.
- Application Tag - Select the application tags that need to be filtered from the traffic.
- Click **Apply**. The New Tool Template page appears.

b. **Application Name** - Turn on the toggle to export Application Name Attribute in Application Metadata Intelligence records.

> **NOTE:** Enable this when only Application Name attribute needs to be exported along with network attributes.

c. **Application Family** - Turn on the toggle to export Application Family Attribute in Application Metadata Intelligence records.

d. **Application Tag**- Turn on the toggle to export Application Tags Attribute in Application Metadata Intelligence records.

5. In the **Advanced Settings** section, do the following:

a. **Collects** - Select the required option. Enter or select the following parameters based on the option selected in this field.

- Counter - Select the Bytes, Packets and Inner Bytes.
- IPv4 - Select or enter the required information in the fields such as Source Address, Destination Address, Fragmentation, and Header Length.
- IPv6 - Select and enter the required information in the fields such as Source Address, Destination Address, Fragmentation, and Header Length.
- Transport - Select and enter the required information in the fields such as Source Port, Destination Port, TCP, ICMP, and UDP.
- Data Link - Select any one of the parameters such as Source Mac, Destination Mac or VLAN.
- Timestamp - Select the required timestamp such as System Uptime First, System Uptime Last, Flow Start, and Flow End.
- Flow - Select the End Reason checkbox if required.
- Interface - These options are supported only in standalone deployments (GigaVUE -HC1, GigaVUE-HC3, GigaVUE-HCT, and GigaVUEGigaVUE-HC1P) and legacy cluster deployments. Select any one of the following parameters.
- Input Physical - Select the Input Physical checkbox to export the ingress interface as one of the fields sent in the NetFlow record. It also allows exporting the interface index in the NetFlow record. Under Input Physical Width, choose 2 bytes or 4 bytes. A width of 4 bytes is recommended for both v9 and IPFIX protocols, while v5 supports only 2 bytes. CEF supports exporting the Input interface index with a width of 2B (default) or 4B.
- Output Physical - Select the Output Physical checkbox to export the egress interface as one of the fields sent in the NetFlow record. It also allows exporting the interface index in the NetFlow record. Under Output Physical Width, choose 2 bytes or 4 bytes. A width of 4 bytes is recommended for both v9 and IPFIX protocols, while v5 supports only 2 bytes. CEF supports exporting the Output interface index with a width of 2B (default) or 4B.

> **NOTE:** when Input/Output Physical interface width is set to 2B, only the lower order bytes of the interface index are exported.

- Input Name - Select the Input Name `to export the interface name.
- Input Name Width - Enter a value between 1 and 32 bytes. The default value is 16 bytes. The total character limit for the interface name is 128 characters.

6. Click **Apply** to save the configuration. The Tool Templates page appears with the created template.

You can edit, delete, clone, export, and view details of the tool templates in the Actions dropdown.

Edit Tool Template

To edit a tool template

1. Go to **Actions > Edit**. The Edit Tool Template page appears.
2. Make the required changes.
3. Click **Apply**. The Tool Templates page appears with the updated template.

Delete Tool Template

To delete a tool template

1. Go to **Actions > Delete**. The Delete Tool Template confirmation dialog appears.
2. Click **Ok** to delete the selected tool template.

Clone Tool Template

To edit a tool template

1. Go to **Actions > Clone**. The Clone Tool Template page appears.
2. Make the required changes.
3. Click **Apply**. The Tool Templates page appears with the updated template.

Export Tool Template

To export a tool template

1. Go to **Actions > Export**. The Export Tool Template page appears.
2. Enter a name for the template.

3. Click **Export**. The Tool Template is exported to your local system in JSON format.

Import Tool Template

To import tool template

1. In the Tool Templates page, click **Import**. The Import New Template in Json Format.
2. Enter a name in the **Template Name** field.
3. Click Choose File in the **Browse File** and upload the template from your local system.
4. Click **Import**. The imported template appears in the Tool Templates page.

Application Intelligence Field References

The following table lists all the fields utilized while configuring an Application Intelligence, Application Filtering Intelligence and Application Metadata Intelligence sessions.

Application Intelligence Session

Field	Description
Name	Enter a unique name for the Application Intelligence session.
Description	(Optional) Enter a description for the session.
Environment	Select Physical to configure the session for GigaVUE HC Series devices.
Node	Select the node on which you want to deploy the session.
Monitoring	Enables or disables Application Visualization (AppViz) for this session. Monitoring is enabled by default and can be disabled when not required to improve performance.
Export Interval	Displays the interval at which AppViz statistics are exported to GigaVUE-FM. The value is fixed at 300 seconds (5 minutes) and is not user-configurable from version 6.4 onwards.
Fast Mode	When enabled, uses light parsers to increase DPI performance for HTTP and DNS. Only a limited set of HTTP attributes is supported and GigaVUE-FM automatically exposes only the supported attributes. Fast mode can be enabled or disabled only when creating a new session.
GigaSMART Group	Select an existing GigaSMART group, or create a new group by specifying an Alias and selecting one or more engine ports in the Port List.
Application Session	Configures the buffer size used by Application Session Filtering in the GigaSMART group.

Field	Description
Filtering Buffer Size	
Metadata Export Buffer Size	Configures the buffer size used for exporting metadata in the GigaSMART group.
User Defined Applications	Associates a User Defined Applications profile with the session so that the DPI engine can classify proprietary or internal applications based on user-defined signatures.
IP Interface Type	Select IPv4 or IPv6 as the IP version for the AppViz export interface.
IP Address / IP Mask / Gateway / MTU	Configure the IP interface used to export application statistics. The destination IP address must use the same IP version as the selected interface.
Destination IP Address	Specifies the IP address of GigaVUE-FM (or other collector) that receives AppViz statistics. By default, the GigaVUE-FM management IP is displayed.
Source Ports	Select one or more network ports or port groups that provide the traffic to be monitored by the session. Ports already used as sources in intent-based orchestrated solutions are not listed.
L2–L4 Rules	Defines Layer 2–Layer 4 conditions for each Source Traffic block. For each rule, click Select Conditions, choose a qualifier (such as IP version, IP source/destination, VLAN, protocol, or port), set the value, and select Pass or Drop and Bidirectional or unidirectional.
+ (Add a Rule)	Adds another L2–L4 rule to the current Source Traffic block.
+ New Source Traffic	Adds another Source Traffic entry with its own source ports and L2–L4 rules.

Application Filtering Intelligence (AFI)

Application selection and destination traffic

Field	Description
Selected Applications	Lists the applications chosen from Total Applications on the Application Intelligence dashboard for creating an AFI policy. You can select one or more applications.
Pass	Marks the selected application to be passed to the tool ports configured in the associated Destination Traffic Priority block.
Drop	Marks the selected application to be dropped at the tool ports configured in the associated Destination Traffic Priority block.
Destination Traffic Priority	Represents a second-level AFI map that forwards filtered traffic to tool ports. You can configure up to five Destination Traffic Priorities per AFI solution and reorder them to change priority.
Tool Ports (Select ports ...)	Select the tool ports or tool port groups that receive traffic for this Destination Traffic Priority. If the required port is not visible, use Port Editor to mark it as a Tool port.
Port Editor	Opens an inline editor to change the Type of a physical port to Tool, after which it appears in the tool port selection list.

Field	Description
Pass All	When enabled for a Destination Traffic Priority, passes all applications that do not match any explicit AFI rules in that priority. No-rule-match traffic can be monitored using map rule counters.
Advanced Rules	Opens the advanced rule configuration for the selected Destination Traffic Priority, allowing you to define non-application criteria (for example IP, VLAN, DSCP, pattern match) that must be met before traffic is passed or dropped.
+ Add New (Destination Traffic Priority)	Adds another Destination Traffic Priority (second-level map). A maximum of five priorities can be configured per AFI instance.
Filter to	Applies the AFI configuration and binds selected applications and rules to the Destination Traffic Priority.

AFI Settings

Field	Description
Bidirectional	Configures whether AFI maintains flow state and applies rules to both directions of a session. When disabled, each direction is treated independently. Enabled by default.
Timeout	Configures the idle timeout for AFI sessions. The default is 15 seconds; the valid range is 10–120 seconds. Flows that remain silent beyond this interval are flushed.
Buffer	Configures the number of initial packets that AFI buffers before applying filtering rules. The default is 20; the valid range is 3–20. Lower values may cause drops if more packets are needed for DPI classification.
Protocol	Selects the transport protocols for which AFI maintains sessions. Options include TCP only, UDP only, TCP-UDP, TCP UDP and SCTP, and SCTP only. Traffic that does not match the selected protocol set is dropped unless explicitly handled elsewhere.
Session Slicing	When enabled, allows you to mention the **Packet Count**. The Packet Count limits the number of packets passed per flow to a configurable value in the range 2–100. Useful when tools only need the first few packets, such as TLS handshakes.
Session Fields	Configures the protocol fields used to uniquely identify flows. The default 5tuple uses inner source and destination IP addresses, source and destination ports, and protocol; VLAN can optionally be included.

AFI Advanced Rules

Field	Description
Select Options	For each advanced rule, select a qualifier from the drop-down list, such as DSCP, ERSPAN ID, EtherType, GRE Key, GTP-U TEID, IP Fragmentation, IP Version, IPv4 Source/Destination, IPv6 Source/Destination, MAC Source/Destination, MPLS Label, Port Source/Destination, TCP Control, VLAN, VN-Tag VIF IDs, VXLAN ID, or Pattern Match.
Pattern Match Type	When Pattern Match is selected, choose String or Regex and provide the search pattern to match data in the packet payload.
Value	Enter the value appropriate for the selected qualifier (for example VLAN ID, IP address, DSCP

Field	Description
	number, or string/regular expression pattern).
Pass / Drop	Specifies whether packets matching this advanced rule are passed to the destination or dropped.

Application Metadata Intelligence (AMI)

Application and attribute selection

Field	Description
Selected Applications	Lists the applications chosen from Total Applications on the Application Intelligence dashboard for exporting metadata.
Attributes	For each selected application, enables you to select the attributes to be exported. Each exporter can export metadata for up to 256 applications, and up to 64 attributes per application.
Export / Export To	Marks the selected applications for export and associates them with the configured exporters.

Destination Traffic – Exporters

Field	Description
Tool Name	Configures an alias for the metadata collector (for example Splunk, NDR-1).
IP Interface	Selects the IP interface on the Gigamon device that connects to the collector.
Tool IP Address	Specifies the IP address of the tool that receives exported records.
Template	Selects a pre-defined tool template that defines applications, attributes, and export format, such as SplunkMetadataTemplate, SecurityPosture, RogueActivity, or UnmanagedAssets.
L4 Source Port	Specifies the source UDP port used by the exporter.
L4 Destination Port	Specifies the UDP destination port on the collector. When the export format is CEF, the default is 514; when the format is NetFlow, the default is 2055.
SNMP	Enables or disables processing of SNMP packets (UDP 161) on the exporter interface. All exporters sharing an IP interface must have SNMP configured consistently.
Application ID	When enabled, exports the Application Name for all applications identified by the DPI engine. Requires AMI/SVP/ZTA license.
Application List	Limits application metadata export to the selected applications or protocols on this exporter. Packet attributes (Collects) are still exported for all applications unless traffic is filtered using AFI or flow maps.
Format	Selects the export format: NetFlow or CEF.
Version	For NetFlow format, selects the NetFlow/IPFIX version: v5, v9, or IPFIX (IPFIX is the default).
Template	Configures how often template records are exported for IPFIX. The default is 60 seconds; the

Field	Description
Refresh Interval	valid range is 1–216000 seconds.
Record Type	Selects how network and application metadata are exported. Segregated exports separate records for network and application metadata. Cohesive exports a consolidated record comprising both. The default depends on the configured Flow Behavior (Unidirectional or Bidirectional).
Active Timeout	Configures the interval for exporting interim records for long-lived flows. The default is 60 seconds; the valid range is 1–604800 seconds.
Inactive Timeout	Configures the idle timeout for marking flows as inactive and exporting their records. The default is 15 seconds; the valid range is 1–604800 seconds.

Advanced Settings – Collects

Field	Description
Counter	Enables export of Bytes and Packets counters and configures the counter width (32- or 64-bit). NetFlow v5 supports only 32-bit counters.
IPv4 Collects	Enables export of IPv4 attributes such as Source Address, Destination Address, Protocol, DSCP, TOS, TTL, header length, payload length, and fragmentation details.
IPv6 Collects	Enables export of IPv6 attributes such as Source Address, Destination Address, Next Header, Flow Label, Traffic Class, DSCP, and fragmentation details.
Transport Collects	Enables export of transport-layer fields such as Source Port, Destination Port, TCP flags, TCP sequence and acknowledgment numbers, MSS, Aggregate Window Size, and Zero Window statistics.
ICMP Collects	Enables export of ICMP Type and Code for IPv4 and IPv6 ICMP traffic.
Data Link Collects	Enables export of Source MAC, Destination MAC, and ingress VLAN ID.
Timestamp Collects	Enables export of System Uptime First, System Uptime Last, Flow Start, and Flow End timestamps.
Flow Collects	Enables export of the flow end reason (for example TCP FIN, RST, or inactive timeout).
Interface Collects	Enables export of input and output interface index (2- or 4-byte widths) and input interface name, with configurable name width. Supported for standalone and legacy cluster deployments.

Application Metadata Settings

Field	Description
Events	Controls when records are exported for TCP flows. Transaction End exports records as soon as TCP connections terminate; None exports records based only on inactive timeout.
Flow Direction / Behavior	Selects Unidirectional or Bidirectional flow behavior. In Unidirectional mode, one record is exported per direction. In

Field	Description
	Bidirectional mode, a single record is exported for both directions of a flow.
Timeout	Configures the cache timeout for flows maintained in the AMI cache. The default is 300 seconds; the valid range is 1–604800 seconds.
Cache Size	Sets the maximum number of concurrent flows that can be tracked. Valid ranges depend on platform and generation; from 6.12 onwards, the maximum supported value is used as the default for new configurations.
Multi-Collect	When enabled, allows multiple values to be exported for attributes that can have more than one value (for example DNS host addresses). IPFIX supports up to five values per attribute; CEF has no limit.
Data Link	When enabled, exports Source and Destination MAC and VLAN ID and treats flows with identical 5-tuples and different VLAN IDs as the same flow.
Observation Domain ID	Configures a user-defined identifier in the range 0–255 that, combined with the GS engine slot, forms the 4-byte Observation Domain ID used by collectors to distinguish sessions.
DPI Packet Limit	Restricts the number of packets per session sent to the DPI engine to improve performance. Recommended values are between 20–50, since the first packets in a flow typically contain the most significant attributes.
Aggregate Mode	When enabled (Gen3 only), exports minimum, maximum, and mean RTT and TCP loss bytes per export interval for supported protocols such as TCP, HTTP, SSH, TELNET, ICMP, and ICMPv6.
Recurring Attributes	When enabled, exports TLS Server Name for the lifetime of new TLS flows. When disabled, the value is exported only once at the beginning of a flow. Configurable from the Traffic Policy workflow supported on GigaVUE HC Series Gen3 GigaSMART module. .

De-duplication

Field	Description
De-duplication	When enabled for an Application Intelligence solution, inserts the GigaSMART De-duplication operation into the AFI and/or AMI path to remove duplicate packets before they reach the tools. Requires a De-duplication license.

Use Case: Reduce Tool Overload

This use case focuses on optimizing tool capacity by filtering out high-volume, low-risk applications such as multimedia streaming and general audio/video traffic. Since platforms like Netflix and YouTube generate a significant share of internet bandwidth yet pose minimal security risk, excluding this traffic from analysis helps reduce noise and allows monitoring tools to focus on higher-value data.

To enable this capability, you must configure Application Visualization and Application Filtering. These features classify traffic based on application tags (such as Instant Messaging, Multimedia Conferencing, and Multimedia Streaming) and application families like Audio/Video, allowing you to easily identify and filter bandwidth-heavy, low-priority applications.

Applicable Map Components

The following table lists the map components you must configure for this use case.

Map Components	Sample Configuration
Source	Network ports - 1/1/x1
Source Ruleset	N/A
Application Ruleset	N/A
Applications	• Application Filtering: Drop Netflix and Youtube traffic and pass the rest of the traffic. o Application Family - audio-video o Applications - youtube-tv, netflix, and netflix-video o Traffic Action - Drop • Application Visualization
Destination	Tool port - 1/1/x7

To configure a Traffic Policy with the above sample configurations:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.

- **Alias** - Enter GTP- Control.
- (Optional) **Description** - Enter the description for the Traffic Policy.
- (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
- (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.

d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Applications.
   a. Expand the **Applications** component and then drag and drop **Application Filtering**. The GigaSMART Resource page appears.
   b. You can either select an existing GigaSMARTgroup or create a new group by clicking **New** in the **GigaSMART Group** field. The **Create GigaSMART Group** pane appears.
   c. Enter a name for the GigaSMART group in the **Alias** field. For example, gs2port1.
   d. Select a GigaSMART engine port from the **Port List** field, and then click **Save**. The GigaSMART Resource page appears.
   e. Click **Proceed**. The AFI pane appears.
   f. Select the **Pass Applications** check box to allow all applications to be passed to the tool, except those you explicitly set to **Drop** (for example, youtube-tv, netflix, netflix-video). If you set the **Traffic Action** of any application to **Pass**, the **Pass Applications** check box becomes read-only and cannot be selected. This is because choosing specific applications to Pass indicates that your intent is to forward only those selected applications, not all applications.
   g. Search **Application Family** by audio-video, and then select it.

      By default, all applications under the audio-video Application Family is selected and the **Traffic Action** for all the applications are selected as Pass All. You can choose which application traffic you want to pass or drop.
   h. Search **Application** by youtube-tv, netflix, and netflix-video, and then select the **Traffic Action** for these applications as Drop.
   i. Click **Save**. The **GS-GRP** and the **AFI** blocks appear on the canvas.

j. From the **Applications** component, drag and drop **Application Visualization**. The AppViz page appears.

k. In the **Exporter Interface** field, select one of the following options:

- Management Interface—Use the GigaVUE-FM management interface to export application statistics.
- Tool IP Interface—Use the IP interface configured on the tool port to export application statistics.

l. Click Save. The Application Visualization dialog box appears.

m. Click **Continue**. The App Viz block appears on the canvas.

4. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool port 1/1/x7.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Destination block appears on the canvas with the selected tool port.

5. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

6. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.



Even if you have not configured any rules within the Traffic Policy, you notice that a **By Rule** block automatically appears on the canvas. This is expected behavior. Based on your intent during policy creation, GigaVUE-FM automatically generates a default By Rule to ensure

that IPv4 traffic passing through the associated network port is included in the policy flow. This auto-created rule helps establish the foundational traffic path so you can continue building or refining your Traffic Policy as needed.

Use Case: Enhance Context for Incident Response

This use case helps security teams accelerate investigation and response by enriching incident data with deep network visibility. When an alert provides an IP address and timestamp, you can query your NDR solution for network-forensic details, such as DNS queries, HTTP activity, and other enriched metadata, to understand what was happening at the exact moment of the incident.

To enable this level of contextual insight, you must configure Application Metadata Intelligence (AMI) and Application Visualization. These components work together to extract, classify, and present application-layer metadata, allowing analysts to correlate alert information with detailed traffic behavior during incident triage.

Applicable Map Components

The following table lists the map components you must configure for this use case.

Map Components	Sample Configuration
Source	Network ports - 1/1/x5
Source Ruleset	N/A
Application Ruleset	N/A
Applications	• Application Visualization • Application Metadata: o Exporter Details: Create the Tool IP interface to which metadata should be sent for visualization. This is the destination interface on the monitoring or analytics tool that receives enriched metadata for processing. Selecting the correct Tool IP ensures that application insights and traffic attributes are forwarded accurately, enabling effective visibility and analysis on the downstream platform. ▪ Tool IP Interface - IP-1 ▪ Tool Name - Exp1 ▪ Tool Port - 1/1/x11 ▪ IP Address - 1.1.1.2 ▪ Tool Destination IP - 1.1.1.1 ▪ Format / Version - CEF

Map Components	Sample Configuration
	o Application Attributes: Export the specific metadata your tools need for deep investigation. These metadata attributes enriches each flow record with high-value context while avoiding unnecessary fields, helping downstream tools improve detection accuracy and reduce storage and processing overhead. ▪ Application Family - network-service; Application - dns; Attributes - query, reply-code, response-time ▪ Application Family - web; Application - http; Attributes - accept, code, response-ts ▪ Application Family - encrypted; Application - ssl; Attributes - cipher-suite-id, protocol-version, validity-not-after o Advanced Settings: Use timestamps to capture the exact start and end time of each flow, enabling precise correlation with alert timestamps and detailed timeline reconstruction. Use the IPv4 source and destination address prefixes to export flows per host IP, rather than as aggregated subnets, so downstream tools can accurately identify the specific endpoints involved in an incident. ▪ Timestamp - Flow Start (milliseconds) and Flow End (milliseconds) ▪ IPv4 - Source Address and Destination Address Prefix set to /32.
Destination	N/A

To configure a Traffic Policy with the above sample configurations:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - **Alias** - Enter GTP- Control.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Source block appears on the canvas with the selected network port.

2. Select Applications.

a. Expand the **Applications** component and then drag and drop **Application Metadata**. The GigaSMART Resource page appears.

b. You can either select an existing GigaSMARTgroup or create a new group by clicking **New** in the **GigaSMART Group** field. The **Create GigaSMART Group** pane appears.

c. Enter a name for the GigaSMART group in the **Alias** field. For example, gs2port1.

d. Select a GigaSMART engine port from the **Port List** field, and then click **Save**. The GigaSMART Resource page appears.

e. Click **Proceed**. The Exporter 1 pane appears.

> **NOTE:** You can create up to a maximum of 5 exporters in an Application Metadata Intelligence solution.

f. On the Exporter Details tab, enter Exp1 in the Tool Name field.

g. You can either select an existing Tool IP Interface or create a new IP interface by clicking **New** in the **Tool IP Interface** field. The **New Tool IP Interface** pane appears.

h. Enter an **Alias** and **Description** for the Tool IP Interface.

i. Select the tool port 1/1/x11 and then **Type** as IPv4.

j. Enter 1.1.1.2 in the **IP Address** field, /24 in the **IP Mask** field, and then 1.1.1.1 in the **Gateway** field.

k. Leave the **MTU** to the default value, and then click **Save**. The Exporter 1 pane appears.

l. Enter 1.1.1.1 in the **Tool IP Address** field. This is the IP Address of the Gateway you configured in your Tool IP Interface.

m. Leave the L4 Source Port and L4 Destination port to the default values.

n. Select the tool template you want to use for the exporter. For instructions on how to create a tool template, refer to Create Custom Tool Template.

o. Under the **Export Format** section, select CEF in the **Format** field and Segregated in the **Record Type** field.

p. On the **Application Attributes** tab, select the following:

- Application Family—network-service; Application—dns; Attributes—query, reply-code, response-time
- Application Family—web; Application—http; Attributes—accept, code, response-ts
- Application Family—encrypted; Application—ssl; Attributes—cipher-suite-id, protocol-version, validity-not-after

NOTE:
The AMI exporter exports application metadata for the selected applications but exports packet attributes for all applications. To export packet attributes only for the selected applications, filter the required traffic by using a flow map or AFI. Flow Direction/Behaviour under Application Metadata Settings must be configured as **Bidirectional** for exporting application Metadata. Each exporter can be configured to export application metadata for up to 256 applications, and up to 64 attributes can be exported for each application. GigaVUE-FM allows selecting metadata for applications that support exporting metadata. Applications which do not support exporting metadata are not listed in the GigaVUE-FM. By default, application name is exported for all such applications.

q. On the **Advanced Settings** tab, add Timestamp and IPv4 from the **Add L2-L4 attributes**. The Timestamp and IPv4 sections appear.

r. Under the Timestamp section, select **Flow Start (milliseconds)** and **Flow End (milliseconds)**.

s. Under IPv4 section, enter /32 as the Prefix for the **Source Address** and **Destination Address** fields.

t. Click **Save**. The Application Metadata pane appears.

u. Click **Close**. The AMI block appears on the canvas.

v. Search **Application** by youtube-tv, netflix, and netflix-video, and then select the **Traffic Action** for these applications as Drop.

w. Click **Save**. The **GS-GRP** and the **AFI** blocks appear on the canvas.

x. From the **Applications** component, drag and drop **Application Visualization**. The AppViz pane appears.

y. Select Tool IP Interface in the **Exporter Interface** field, and then select the Tool IP Interface you created while configuring the Exporter for AMI.

z. Click **Save**. In the Application Visualization dialog box that appears, click **Continue**. The App Viz block appears on the canvas.

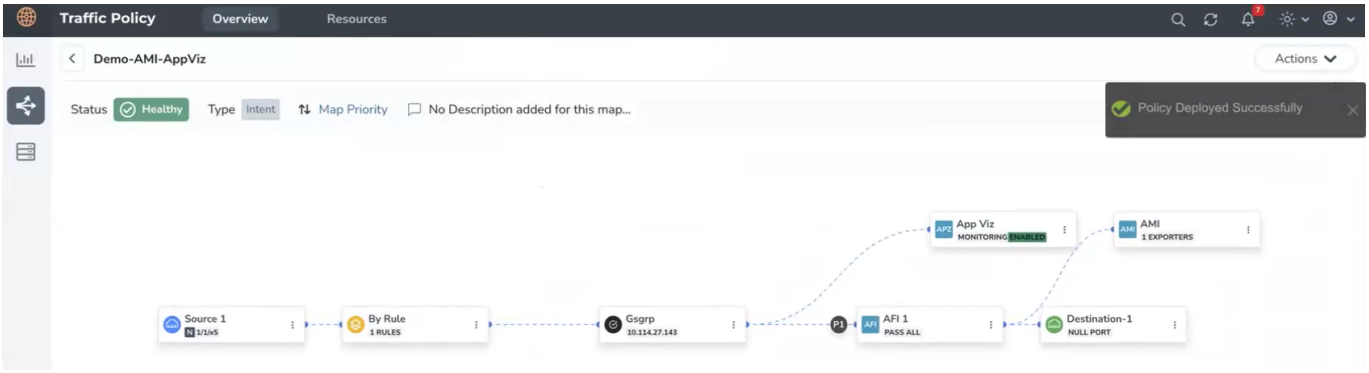
3. Deploy Traffic Policy.

   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

4. Verify the deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.



After deploying the Traffic Policy, you notice that GigaVUE-FM automatically displays the following blocks on the canvas even if you have not configured them. This is expected behaviour. :

- By Rule block— GigaVUE-FM automatically generates a default By Rule to ensure that IPv4 traffic passing through the associated network port is included in the policy flow. This auto-created rule helps establish the foundational traffic path so you can continue building or refining your Traffic Policy as needed.
- AFI block—GigaVUE-FM places an AFI block on the canvas to represent a default pass-all filtering map. This placement helps you clearly understand where application filtering fits within the traffic flow.
- Destination (Null Port)— GigaVUE-FM displays a Null port as the default destination to keep the traffic path complete from ingress to egress on the canvas and prevents breaks in the visual workflow.

Use Case: Maintain Visibility Whilst Filtering Traffic

This use case focuses on maintaining complete visibility into network activity while selectively filtering out non-critical applications. Critical workloads, such as IoT and Artificial Intelligence (AI) applications, continue to be monitored in full, while high-volume, low-priority traffic like multimedia streaming and general audio/video applications (for example, Netflix and YouTube) is excluded from tool processing. This approach helps reduce unnecessary load on monitoring systems while ensuring that essential application traffic remains fully observable.

To enable this capability, you must configure Application Visualization (AppViz), Application Filtering Intelligence (AFI), and Application Metadata Intelligence (AMI). Together, these features identify application types, filter out non-essential traffic, and export metadata for all traffic, ensuring 100% visibility across your environment even when certain traffic is not forwarded for deep analysis.

Applicable Map Components

The following table lists the map components you must configure for this use case.

Map Components	Sample Configuration
Source	Network ports - 1/1/x1
Source Ruleset	N/A
Application Ruleset	N/A
Applications	• Application Visualization • Application Filtering: Drop Netflix and Youtube traffic and pass the Artificial Intelligence and IoT traffic. o Application Family - audio-video; Applications - youtube-tv, netflix, and netflix-video; Traffic Action - Drop. o Application Tag - Artificial Intelligence and IoT (Internet of Things); Traffic Action - Pass. • Application Metadata: o Exporter Details: Create the Tool IP interface to which metadata should be sent for visualization. This is the destination interface on the monitoring or analytics tool that receives enriched metadata for processing. Selecting the correct Tool IP ensures that application insights and traffic attributes are forwarded accurately, enabling effective visibility and analysis on the downstream platform. ▪ Tool IP Interface - IP-1 ▪ Tool Name - Exp1 ▪ Tool Port - 1/1/x11 ▪ IP Address - 1.1.1.2 ▪ Tool Destination IP - 1.1.1.1 ▪ Format / Version - CEF o Application Attributes: Export the specific metadata your tools need for deep investigation. These metadata attributes enriches each flow record with high-value context while avoiding unnecessary fields, helping downstream tools improve detection accuracy and reduce storage and processing overhead. ▪ Application Family - network-service; Application - dns; Attributes - query, reply-code, response-time ▪ Application Family - web; Application - http; Attributes - accept, code, response-ts ▪ Application Family - encrypted; Application - ssl; Attributes - cipher-

Map Components	Sample Configuration
	suite-id, protocol-version, validity-not-after o Advanced Settings: Use timestamps to capture the exact start and end time of each flow, enabling precise correlation with alert timestamps and detailed timeline reconstruction. Use the IPv4 source and destination address prefixes to export flows per host IP, rather than as aggregated subnets, so downstream tools can accurately identify the specific endpoints involved in an incident. ▪ Timestamp - Flow Start (milliseconds) and Flow End (milliseconds) ▪ IPv4 - Source Address and Destination Address Prefix set to /32.
Destination	N/A

To configure a Traffic Policy with the above sample configurations:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - **Alias** - Enter GTP- Control.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Configure Application Filtering.
   a. Expand the **Applications** component and then drag and drop **Application Filtering**. The GigaSMART Resource page appears.

b. You can either select an existing GigaSMARTgroup or create a new group by clicking **New** in the **GigaSMART Group** field. The **Create GigaSMART Group** pane appears.

c. Enter a name for the GigaSMART group in the **Alias** field. For example, gs2port1.

d. Select a GigaSMART engine port from the **Port List** field, and then click **Save**. The GigaSMART Resource page appears.

e. Click **Proceed**. The AFI pane appears.

f. Search **Application Family** by audio-video, and then select it.

g. Search **Application** by youtube-tv, netflix, and netflix-video, and then select the **Traffic Action** for these applications as Drop.

h. Search **Application Tag** by Artificial Intelligence and IoT (Internet of Things), and then select the **Traffic Action** for these applications as Pass.

i. Click **Save**. The **GS-GRP** and the **AFI** blocks appear on the canvas.

4. Configure Application Visualization.

a. From the **Applications** component, drag and drop **Application Visualization**. The AppViz page appears.

b. In the **Exporter Interface** field, select one of the following options:

- Management Interface—Use the GigaVUE-FM management interface to export application statistics.
- Tool IP Interface—Use the IP interface configured on the tool port to export application statistics.

c. Click Save. The Application Visualization dialog box appears.

d. Click **Continue**. The App Viz block appears on the canvas.

5. Configure Application Metadata.

a. Expand the **Applications** component and then drag and drop **Application Metadata**. The GigaSMART Resource page appears.

b. You can either select an existing GigaSMARTgroup or create a new group by clicking **New** in the **GigaSMART Group** field. The **Create GigaSMART Group** pane appears.

c. Enter a name for the GigaSMART group in the **Alias** field. For example, gs2port1.

d. Select a GigaSMART engine port from the **Port List** field, and then click **Save**. The GigaSMART Resource page appears.

e. Click **Proceed**. The Exporter 1 pane appears.

> **NOTE:** You can create up to a maximum of 5 exporters in an Application Metadata Intelligence solution.

f. On the Exporter Details tab, enter Exp1 in the Tool Name field.

g. You can either select an existing Tool IP Interface or create a new IP interface by clicking **New** in the **Tool IP Interface** field. The **New Tool IP Interface** pane appears.

h. Enter an **Alias** and **Description** for the Tool IP Interface.

i. Select the tool port 1/1/x11 and then **Type** as IPv4.

j. Enter 1.1.1.2 in the **IP Address** field, /24 in the **IP Mask** field, and then 1.1.1.1 in the **Gateway** field.

k. Leave the **MTU** to the default value, and then click **Save**. The Exporter 1 pane appears.

l. Enter 1.1.1.1 in the **Tool IP Address** field. This is the IP Address of the Gateway you configured in your Tool IP Interface.

m. Leave the L4 Source Port and L4 Destination port to the default values.

n. Select the tool template you want to use for the exporter. For instructions on how to create a tool template, refer to Create Custom Tool Template.

o. Under the **Export Format** section, select CEF in the **Format** field and Segregated in the **Record Type** field.

p. On the **Application Attributes** tab, select the following:

- Application Family—network-service; Application—dns; Attributes—query, reply-code, response-time
- Application Family—web; Application—http; Attributes—accept, code, response-ts
- Application Family—encrypted; Application—ssl; Attributes—cipher-suite-id, protocol-version, validity-not-after

q. On the **Advanced Settings** tab, add Timestamp and IPv4 from the **Add L2-L4 attributes**. The Timestamp and IPv4 sections appear.

r. Under the Timestamp section, select **Flow Start (milliseconds)** and **Flow End (milliseconds)**.

s. Under IPv4 section, enter /32 as the Prefix for the **Source Address** and **Destination Address** fields.

t. Click **Save**. The Application Metadata pane appears.

u. Click **Close**. The AMI block appears on the canvas.

6. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.
   b. Select the tool port 1/1/x7.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Destination block appears on the canvas with the selected tool port.
7. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
8. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Even if you have not configured any rules within the Traffic Policy, you notice that a **By Rule** block automatically appears on the canvas. This is expected behavior. Based on your intent during policy creation, GigaVUE-FM automatically generates a default By Rule to ensure that IPv4 traffic passing through the associated network port is included in the policy flow. This auto-created rule helps establish the foundational traffic path so you can continue building or refining your Traffic Policy as needed.

Configure the Traffic Policy for Subscriber Intelligence Solutions

Subscriber Intelligence is a suite of GigaSMART applications for service providers that gives coherent visibility into subscriber voice, video, and data traffic across fast-growing 3G/4G/5G networks. It correlates control- and user-plane traffic and provides subscriber-aware flow sampling and metadata export, so monitoring and security tools can see complete subscriber sessions (by subscriber, device, RAN, or slice) while processing only an optimized subset of traffic.

In high-scale Subscriber Intelligence deployments that use GigaSMART GTP and CUPS correlation with stateful load balancing, you can configure load-balanced port groups (LBPGs) with up to 64 GigaStream members when the LBPG members are GigaStreams on supported SMT-HC3-C08 modules. For details, see the GigaSMART GTP and CUPS Correlation and SMT-HC3-C08 module documentation.

Refer to the following sections for instructions on how to configure Traffic Policy for Subscriber Intelligence solutions with real time examples:

- GigaSMART GTP and CUPS Correlation
- GigaSMART Rotational Sampling
- GigaSMART 4G RAN Correlation
- GigaSMART SIP/RTP Correlation
- GTP Overlap Flow Sampling
- GigaSMART GTP Whitelisting and GTP Flow Sampling
- 4G/5G Traffic Monitoring using UPN
- GigaSMART 5G CUPS
- GTP Scaling

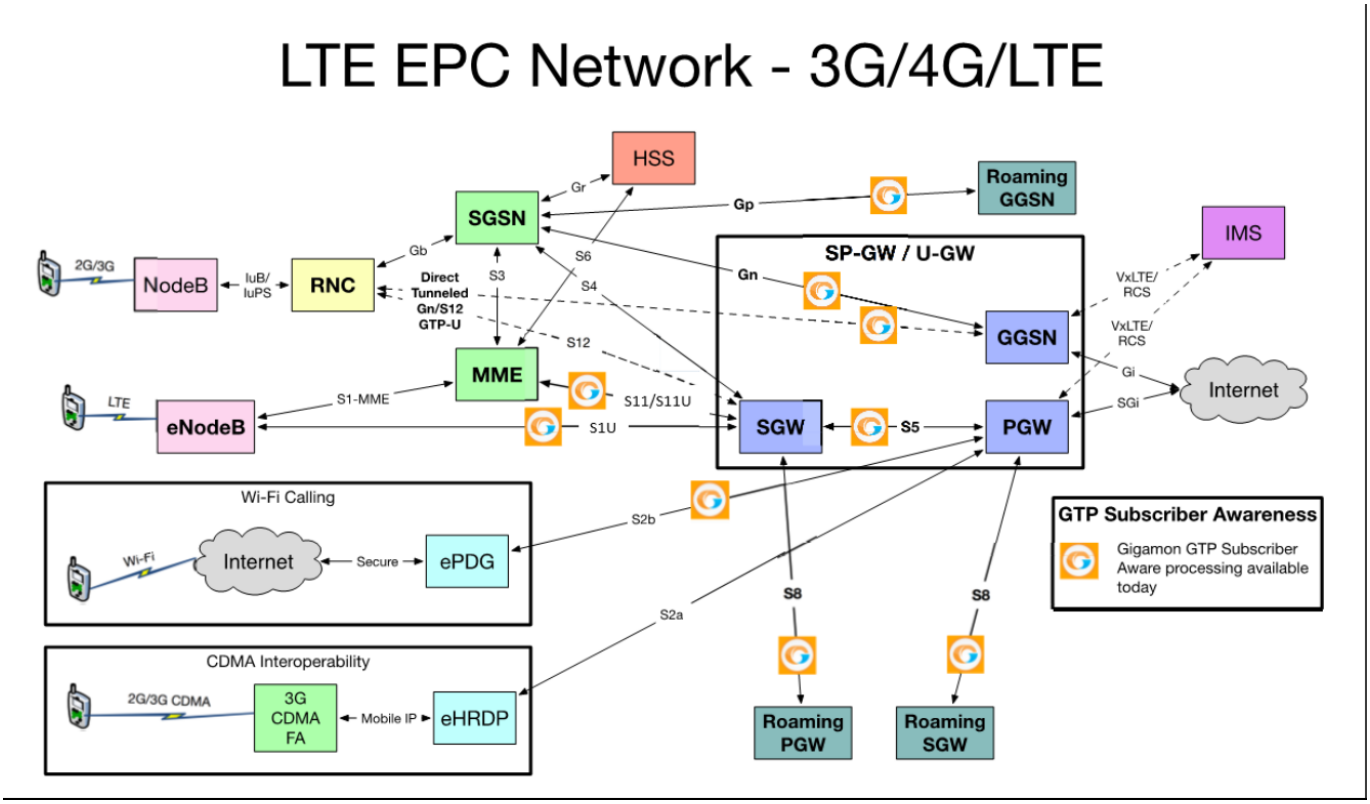
GigaSMART GTP and CUPS Correlation

Refer to the following sections for instructions on how to configure Traffic Policy for GigaSMART GTP and CUPS Correlation with real time examples:

- Identifying High Value and/or Roaming Subscribers Based on IMSIs
- Identifying GTP Versions
- Same Subscriber Filter on Different Versions
- Same Subscriber Filter on Different Interfaces
- EPC Filtering in Scenario1
- EPC Filtering in Scenario2

Identifying High-Value and/or Roaming Subscribers Based on IMSIs

Use GTP correlation to identify high value subscribers based on an IMSI or group of IMSIs. GTP correlation keeps track of the IMSIs that you are interested in monitoring. It correlates them to the corresponding data/user-plane sessions for the subscriber and/or group of subscribers. Filtering on subscriber ID (IMSI) limits the amount of traffic that is sent to monitoring tools.



In Example 1, filter rules are configured to identify and forward all the traffic related to subscribers identified by an IMSI prefix. All traffic specific to the filtered IMSIs 22222222222223*, including GTP-c and GTP-u, is forwarded to a monitoring tool. A shared collector is configured to which traffic not matching the filters is sent.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port 1/1/x1
Source Ruleset	By Rule • Rule 1: Condition - Port Source, Min - 2123 • Rule 2: Condition - Port Source, Min - 2152 • Rule 3: Condition - IP Fragmentation, Value - All Frag No First
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port - vp1
Second-Level Map1	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	My GSOP - Flow Filtering • Type - GTP
Application Ruleset	By Rule • Condition - GTP IMSI, IMSI - 22222222222223*
Destination	Tool port - 1/1/x4
Second-Level Map2	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	N/A
Application Ruleset	Collector
Destination	Tool port - 1/1/x6

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The GigaSMART Groups page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name for the GigaSMART group in the **Alias** field.

6. Select a GigaSMART engine port from the **Port List**.
7. In the **GigaSMART Parameters** section, go to **GTP Flow** to configure the Timeout. The default is value is 48, which is 480 minutes.
8. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure a First-Level Map

To configure a first-level map that directs GTP traffic from physical network port/s to the virtual port vp1:

1. Access Traffic Policy.
    a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
    b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
    c. In the **New Map** dialog, do the following:
        - **Alias** - Enter GTP- Control.
        - (Optional) **Description** - Enter the description for the Traffic Policy.
        - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
        - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
    d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.

a. Click **Sources**. The Source 1 page appears.

b. Select the network port 1/1/x4.

> **NOTE:** To configure or modify the port type, use the **Quick Port Editor**.

c. Click **Add**. The Source block appears on the canvas with the selected network port.

3. Select Source Ruleset.

a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.

b. Configure Rule 1: Click **Add Rules**. The Add Ruleset pane appears.

c. Turn on **Pass** toggle and **Bi-directional** to pass GTP-u traffic.

d. Select Port Source in the **Condition** list.

- Enter 2123 in the **Min** field.

e. Configure Rule 2: Click **+**.

f. Turn on **Pass** toggle and **Bi-directional** to pass GTP-u traffic.

g. Select Port Source in the **Condition** list.

- Enter 2152 in the **Min** field.

h. Configure Rule 3: Click **+**.

i. Turn on **Pass** toggle and **Bi-directional** to pass GTP-u traffic.

j. Select **IP Fragmentation** from the Condition list.

k. Select **All Frag No First** from the list.

l. Click **Save**. The added rule appears in the **Source Ruleset** page.

m. Click **Close**. The By Rules block appears on the canvas with the configured rules.

4. Select Destination.

a. Click **Destinations**. The Destination 1 page appears.

b. Select the configured virtual port vp1.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

5. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

6. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure Second-Level Maps

Map 1

To configure a second level map that takes traffic from the virtual port, applies the GigaSMART operation, matches IMSIs specified by the flow rule, and sends matching traffic to physical tool ports:

1. Access Traffic Policy
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Create a new map with **Alias** as GTP-User.
2. Select Source
   a. Click **Sources**. The Source 1 page appears with the list of available ports.
   b. Select the virtual port vp1.
   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The GSOP page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
   e. Select the GigaSMART group you created earlier.
   f. Select Flow Filtering from the **Applications** list.
   g. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.
4. Select Application Ruleset
   a. Expand the Application Ruleset component and then drag and drop a **Flow Filter**. The Application Ruleset page appears.
   b. Click **Add Rules**. The Add Ruleset pane appears.

c. Turn on **Pass** toggle to pass the configured traffic.

d. Select GTP IMSI from the **Condition** list.

e. Enter **22222222222223*** in the IMSI field.

f. Click **Save**. The Application Ruleset page appears with the configured rule.

g. Click **Close**. The By Rule block appears on the canvas with the configured application rule.

5. Select Destination

   a. Click **Destinations**. The Destination 1 page appears.

   b. Select the tool port 1/1/x4.

   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

6. Deploy Traffic Policy

   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment

   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2

To configure a second-level map to collect unmatched traffic:

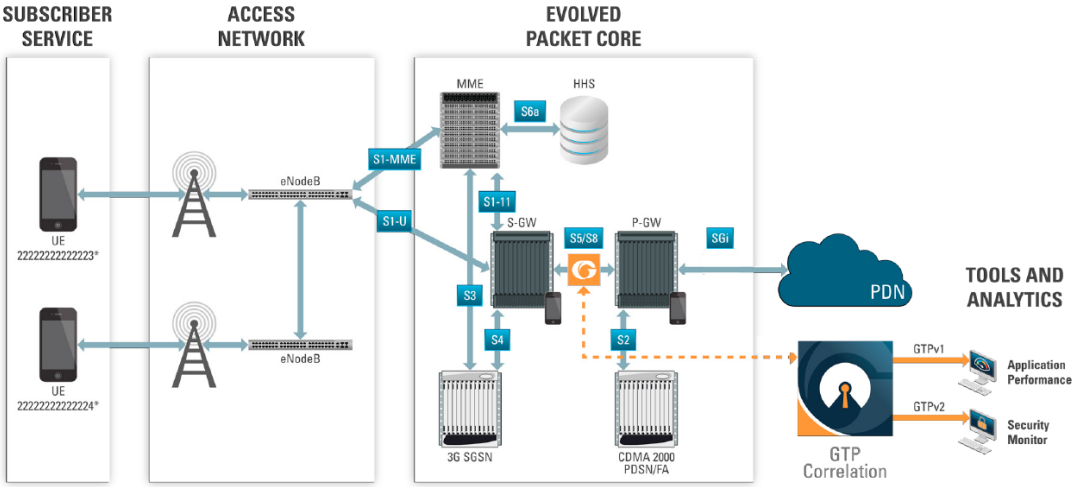
Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (scoll) with the following components:
   - Source port - vp1
   - Application Ruleset - Collector
   - Destination port - 1/1/x6

Identifying GTP Versions

As part of GTP correlation, GigaVUE nodes also provide the flexibility to identify GTPv1 and GTPv2 messages. GTP version information is typically exchanged only as part of the control

sessions. By correlating the control and user-plane sessions, GigaVUE nodes can identify, filter, and forward all sessions specific to a GTPv1 or v2 to one or more monitoring/analytic tools.



In Example 2, EMEI traffic is distributed based on GTP versions as follows:

- Filter and forward GTPv1 to a tool port
- Filter and forward GTPv2 to another tool port

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port 1/1/x4
Source Ruleset	By Rule • Rule 1: Condition - Port Source, Min - 2123 • Rule 2: Condition - Port Source, Min - 2152 • Rule 3: Condition - IP Fragmentation, Value - All Frag No First
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port - vp1
Second-Level Map1	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	My GSOP - Flow Filtering • Type - GTP
Application Ruleset	Flow Filter • Condition - GTP IMSI

Map Components	Sample Configuration
	• IMSI - * and Version - V1
Destination	Tool port - 1/1/x4
Second-Level Map2	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	My GSOP - Flow Filtering • Type - GTP
Application Ruleset	Flow Filter • Condition - GTP IMSI • IMSI - * and Version - V2
Destination	Tool port - 1/3/x15

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The GigaSMART Groups page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**.
7. In the GTP Flow set the Timeout. The default is value is 48, which is 480 minutes.
8. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.

5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure First-Level Maps

Map 1

To configure a first-level map that directs GTP traffic from physical network ports to the virtual port vp1:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter GTP- Control.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The Source 1 page appears.
   b. Select the network port 1/1/x4.

   NOTE: To configure or modify the port type, use the **Quick Port Editor**.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The Source Ruleset page appears.
   b. Configure Rule 1: Click **Add Rules**. The Add Ruleset pane appears.
   c. Turn on **Pass** toggle and **Bi-directional** to pass GTP-c traffic.
   d. Select Port Source in the **Condition** list.

- Enter 2123 in the **Min** field.

e. Configure Rule 2: Click **+**.

f. Turn on **Pass** toggle and **Bi-directional** to pass GTP-u traffic.

g. Select Port Source in the **Condition** list.

- Enter 2153 in the **Min** field.

h. Configure Rule 3: Click **+**.

i. Turn on **Pass** toggle to pass the configured traffic.

j. Select IPv4 Fragmentation from the **Condition** list.

k. Select All Frag No First from the **Value** field.

l. Click **Save**. The added rule appears in the **Source Ruleset** page.

m. Click **Close**. The By Rules block appears on the canvas with the configured rules.

4. Select Destination.
   a. Click **Destinations**. The Destination 1 page appears.
   b. Select the configured virtual port vp1.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure Second-Level Maps

Map 1

To configure a second level map that takes traffic from the virtual port, applies the GigaSMART operation, matches IMEIs specified by the flow rule, and sends matching traffic to a tool port

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

b. Create a new map with **Alias** as m1.

2. Select Source.
   a. Click **Sources**. The Source 1 page appears with the list of available ports.
   b. Select the virtual port vp1.
   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application.
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The GSOP page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
   e. Select the GigaSMART group you created earlier.
   f. Select Flow Filtering from the **Applications** list. By default, the application you select is enabled.
   g. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.
4. Select Application Ruleset.
   a. Expand the Application Ruleset component and then drag and drop a **Flow Filter**. The Application Ruleset page appears.
   b. Click **Add Rules**. The Add Ruleset pane appears.
   c. Turn on **Pass** toggle to pass the configured traffic.
   d. Select GTP IMSI from the **Condition** list.
   e. Enter * in the **IMSI** field and select V1 from the **Version** list.
   f. Click **Save**. The Application Ruleset page appears with the configured rule.
   g. Click **Close**. The Flow Filter block appears on the canvas with the configured application rule.
5. Select Destination.
   a. Click **Destinations**. The Destination 1 page appears.
   b. Select the tool port.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

6. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
7. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2

To create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, matches IMEIs specified by the flow rule, and sends matching traffic to another tool port

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (m2) with the following components and configurations:
   - Source port - Virtual Port (vp1)
   - Destination port - 1/3/x15
   - Application Ruleset - By Rule
     - Condition - GTP IMSI
     - IMSI - *, and Version - V2
2. Deploy and verify the Traffic Policy.

Same Subscriber Filter on Different Versions

In this example, traffic from the same subscriber is forwarded to two different load balancing groups based on version. GTP version 1 traffic is sent to one load balancing group and GTP version 2 traffic is sent to another load balancing group.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port - 1/1/x4

Map Components	Sample Configuration
Source Ruleset	By Rule • Condition - MAC Destination o MAC Address - 00:a0:d1:e1:02:01) o MAC mask - 0000.0000.0000
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port - vp1
Second-Level Map1	
Source	Virtual Port - vp1
Source Ruleset	N/A
Applications	My GSOP - Flow Filtering and Load Balancing • Load Balancing: Type - Stateful, Metric - Hashing, and Hashing Key - IMSI.
Application Ruleset	Flow Filter • Condition - GTP IMSI • IMSI - *, Version - V1, and Subset - None.
Destination	Tool port group portgrpv1
Second-Level Map2	
Source	Virtual Port - vp1
Source Ruleset	N/A
Applications	My GSOP - Flow Filtering and Load Balancing • Load Balancing: Type - Stateful, Type - GTP, Metric - Hashing, and Hashing Key - IMSI.
Application Ruleset	Flow Filter • Condition - GTP IMSI • IMSI - *, Version - V2, and Subset - None.
Destination	Tool port group portgrpv2

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure Port Groups

To configure port groups:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter portgrpv1 in the **Alias** field for the new Port Group and click **Tool** option.
5. Enable **Load Balancing**.
6. Choose the **Tool Ports** you want to attach to the port group.
7. Enter the weights for each tool port.
8. Click **Apply**. The created port group appears in the port group table.
9. Repeat the steps 1 - 8 and create another port group with **Alias** as portgrpv2.

Configure an Ingress First-Level Map

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.

- (Optional) **Description** - Enter the description for the Traffic Policy.
- (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
- (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.

d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source.
   a. Click **Sources**. The Source 1 page appears.
   b. Select the network port 1/1/x4.

   NOTE: To configure or modify the port type, use the **Quick Port Editor**.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The Source Ruleset page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the configured traffic.
   d. Select MAC Destination from the **Condition** list.
   e. Enter **MAC Address** as 00:a0:d1:e1:02:01 and **MAC Mask** as 0000.0000.0000.
   f. Click **Save**. The added rule appears in the **Source Ruleset** page.
   g. Click **Close**. The By Rules block appears on the canvas with the configured rules.
4. Select Destination.
   a. Click **Destinations**. The Destination 1 page appears.
   b. Select the configured virtual port vp1.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment. Deploy Traffic Policy.On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure Traffic Policy for Second-Level Maps

Map 1

To create a second-level map for version1:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Create a new map with **Alias** as map2_1.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears with the list of available ports.
   b. Select the virtual port vp1.
   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application.
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
   e. Select the GigaSMART group (gs-grp) you created earlier.
   f. Select Flow Filtering and Load Balancing from the **Applications** list. By default, the application you select is enabled.
   g. In the Load Balancing section, select Stateful option and Type as GTP.
   h. Select Hashing in the **Metric** list.
   i. Select IMSI option in the **Hashing Key** field.
   j. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.
4. Select Application Ruleset.
   a. Expand the Application Ruleset component and then drag and drop a **Flow Filter**. The **Application Ruleset** page appears.

b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle to pass the configured traffic.

d. Select GTP IMSI from the **Condition** list, and enter the following fields:

- IMSI - *, Type - Version, Version - V1.

e. Click **Save**. The **Application Ruleset** page appears with the configured rule.

f. Click **Close**. The Flow Filter block appears on the canvas with the configured application rule.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears.

b. Select the tool port group portgrpv1.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

6. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2

To create a another second-level map for version2:

Follow the same steps in the previous section to configure another second-level map with the following changes:

1. Create the map (map2_2) with
   - Source - Virtual Port vp1
   - Application ruleset - Flow Filter
     - **Condition** - GTP IMSI
     - **IMSI** - *, **Type** - Version, **Version** - V2.
   - Application - Flow Filtering and Load Balancing
     - Load Balancing: Type - Stateful, **Metric** - Hashing, **Hashing Key** - IMSI
   - Destination - Tool Port Group portgrpv2.
2. Deploy and verify the Traffic Policy.

Same Subscriber Filter on Different Interfaces

In this example, traffic from the same subscriber is forwarded to two different load balancing groups based on interface. In this example, VLANs 1601 and 1602 are from S5/S8 interface and VLANs 1611 and 1612 are from S11/S1-U interface. The first level maps split the VLAN traffic to different virtual ports. The second level maps send the traffic to different load balancing groups.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map1	
Source	Network port - 1/1/x4
Source Ruleset	By Rule • Condition - VLAN, Min - 1601, and Max - 1602
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port - vp1
First-Level Map2	
Source	Network port - 1/1/x4
Source Ruleset	By Rule • Condition - VLAN, Min - 1611, and Max - 1612
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port - vp2
Second-Level Map1	
Source	Virtual Port - vp1
Source Ruleset	N/A
Applications	My GSOP - Flow Filtering and Load Balancing • Load Balancing: Type - Stateful, Metric - Hashing, and Hashing Key - IMSI.
Application Ruleset	Flow Filter • Condition - GTP IMSI • IMSI - *, Version - V1, and Subset - None.
Destination	Tool port group portgrpv1
Second-Level Map2	
Source	Virtual Port - vp2

Map Components	Sample Configuration
Source Ruleset	N/A
Applications	My GSOP - Flow Filtering and Load Balancing • Load Balancing: Type - Stateful, Type - GTP, Metric - Hashing, and Hashing Key - IMSI.
Application Ruleset	Flow Filter • Condition - GTP IMSI • IMSI - *, Version - V2, and Subset - None.
Destination	Tool port group portgrpv2

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The GigaSMARTGroups page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.
8. Repeat steps 1- 8 and create another virtual port with **Alias** as vp2.

Configure Port Groups

To configure port groups:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter portgrpv1 in the **Alias** field for the new Port Group and click **Tool** option.
5. Enable **Load Balancing**.
6. Choose the **Tool Ports** you want to attach to the port group.
7. Enter the weights for each tool port.
8. Click **Apply**. The created port group appears in the port group table.
9. Repeat the steps 1 - 8 and create another port group with **Alias** as portgrpv2.

Configure First-Level Maps

Map 1

To configure a first level map for vport1:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The Source 1 page appears.
   b. Select the network port 1/1/x4.

   NOTE: To configure or modify the port type, use the **Quick Port Editor**.

c. Click **Add**. The Source block appears on the canvas with the selected network port.

3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The Source Ruleset page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the configured traffic.
   d. Select VLAN from the **Condition** list.
   e. Enter 1601 in the **Min** field and 1602 in the **Max** field.
   f. Click **Save**. The added rule appears in the **Source Ruleset** page.
   g. Click **Close**. The By Rules block appears on the canvas with the configured rules.
4. Select Destination.
   a. Click **Destinations**. The Destination 1 page appears.
   b. Select the configured virtual port vp1.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2

To configure another first level map for vport2:

Follow the same steps in the previous section with the following changes:

1. Create the map (map2) with
   - Source - Network port (1/1/x4)
   - Source Ruleset - By Rule
     - Turn on **Pass** toggle
     - **Condition** - VLAN, **Min** - 1611, and **Max** - 1612
   - Destination - Virtual port (vp2)
2. Deploy and verify the Traffic Policy.

Configure Second-Level Maps

Map 1

To configure a second level map using vport1

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Create a new map with **Alias** as map2_1.
2. Select Source.
   a. Click **Sources**. The Source 1 page appears with the list of available ports.
   b. Select the virtual port vp1.
   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application.
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
   e. Select the GigaSMART group (gs-grp) you created earlier.
   f. Select Flow Filtering and Load Balancing from the **Applications** list. By default, the application you select is enabled.
   g. In the Load Balancing section, select Stateful option and Type as GTP.
   h. Select Hashing in the **Metric** list.
   i. Select IMSI option in the **Hashing Key** field.
   j. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.
4. Select Application Ruleset.
   a. Expand the Application Ruleset component and then drag and drop a **Flow Filter**. The Application Ruleset page appears.
   b. Click **Add Rules**. The Add Ruleset pane appears.
   c. Turn on **Pass** toggle to pass the configured traffic.

d. Select GTP IMSI from the **Condition** list, and enter the following fields:
   - IMSI - *, Type - Version, Version - V1.

e. Click **Save**. The Application Ruleset page appears with the configured rule.

f. Click **Close**. The Flow Filter block appears on the canvas with the configured application rule.

5. Select Destination.
   a. Click **Destinations**. The Destination 1 page appears.
   b. Select the tool port group portgrpv1.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
6. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
7. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2

To create another second level map using vport2

Follow the same steps in the previous section to configure another second-level map with the following changes:

1. Create the map (map2_2) with
   - Source - Virtual Port (vp2)
   - Application ruleset - Flow Filter
     - **Condition** - GTP IMSI
     - **IMSI** - *, **Type** - Version, **Version** - V2.
   - Application - Flow Filtering and Load Balancing
     - Load Balancing: Type - Stateful, **Metric** - Hashing, **Hashing Key** - IMSI
   - Destination - Tool Port Group portgrpv2.
2. Deploy and verify the Traffic Policy.

EPC Filtering in Scenario 1

In this example, traffic for all subscribers on interfaces S11/S1-U and Gn/Gp is sent to the same load balancing group. All other traffic is dropped.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port - 1/1/x4
Source Ruleset	By Rule • Rule 1: Condition - Port Source, Min - 2123 • Rule 2: Condition - Port Source, Min - 2152
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port - vp1
Second-Level Map	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	My GSOP - Flow Filtering and Load Balancing • Load Balancing: Type - Stateful, Type - GTP, Metric - Hashing, and Hashing Key - IMSI.
Application Ruleset	Flow Filter • Condition - GTP IMSI • IMSI - *, Type - Interface, and Interface - Gn/Gp
Destination	Tool port group - portgrp1

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The GigaSMARTGroups page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name for the GigaSMART group in the **Alias** field.

6. Select a GigaSMART engine port from the **Port List**.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure a Port Group

To configure a port group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter portgrp1 in the **Alias** field for the new Port Group and click **Tool** option.
5. Enable **Load Balancing**.
6. Choose the **Tool Ports** you want to attach to the port group.
7. Enter the weights for each tool port.
8. Click **Apply**. The created port group appears in the port group table.

Configure a First-Level Map

To configure an ingress (first level) map:

1. Access Traffic Policy.
    a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.

c. In the **New Map** dialog, do the following:

- **Alias** - Enter map1.
- (Optional) **Description** - Enter the description for the Traffic Policy.
- (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
- (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.

d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source.

a. Click **Sources**. The Source 1 page appears.

b. Select the network port 1/1/x4.

> **NOTE:** To configure or modify the port type, use the **Quick Port Editor**.

c. Click **Add**. The Source block appears on the canvas with the selected network port.

3. Select Source Ruleset.

a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The Source Ruleset page appears.

b. Click **Add Rules**. The Add Ruleset pane appears.

c. Configure Rule 1:

- Turn on **Pass** toggle and **Bi-directional** to pass GTP-u traffic.
- Select Port Source from the **Condition** list.
- Enter 2123 in the **Min** field.

d. Configure Rule 2: Click **+**.

- Turn on **Pass** toggle and **Bi-directional** to pass GTP-u traffic.
- Select Port Source in the **Condition** list.
- Enter 2152 in the **Min** field.

e. Click **Save**. The added rule appears in the **Source Ruleset** page.

f. Click **Close**. The By Rule block appears on the canvas with the configured rules.

4. Select Destination.

a. Click **Destinations**. The Destination 1 page appears.

b. Select the configured virtual port vp1.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure a Second-Level Map

To create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, matches IMSIs specified by the flow rules, and sends matching traffic to physical tool ports:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Create a new map with **Alias** as m1.
2. Select Source.
   a. Click **Sources**. The Source 1 page appears with the list of available ports.
   b. Select the virtual port vp1.
   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application.
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The GSOP page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
   e. Select the GigaSMART group you created earlier.
   f. Select Flow Filtering and Load Balancing from the **Applications** list. By default, the application you select is enabled.
   g. In the Load Balancing section, select Stateful option and Type as GTP.

h. Select Hashing in the **Metric** list.

i. Select IMSI option in the **Hashing Key** field.

j. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

4. Select Application Ruleset.

   a. Expand the Application Ruleset component and then drag and drop a **Flow Filter**. The Application Ruleset page appears.

   b. Click **Add Rules**. The Add Ruleset pane appears.

   c. Configure Rule 1: Turn on **Pass** toggle to pass the configured traffic.

   d. Select GTP IMSI from the **Condition** list, and enter the following fields:

      - **IMSI** - *, **Type** - Interface, **Interface** - Gn/Gp.

   e. Select GTP IMSI from the **Condition** list, and enter the following fields:

      - **IMSI** - *, **Type** - Interface, **Interface** - S11/S1U.

   f. Click **Save**. The Application Ruleset page appears with the configured rule.

   g. Click **Close**. The Flow Filter block appears on the canvas with the configured application rule.

5. Select Destination.

   a. Click **Destinations**. The Destination 1 page appears.

   b. Select the tool port group portgrp1.

   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

6. Deploy Traffic Policy.

   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

EPC Filtering in Scenario 2

In this example, traffic for all subscribers from all interfaces except S5/S8 is sent to the same load balancing group. Traffic from the S5/S8 interface is dropped.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network port - 1/1/x4
Source Ruleset	By Rule • Rule 1: Condition - Port Source, Min - 2123 • Rule 2: Condition - Port Source, Min - 2152
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port - vp1
Second-Level Map	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	My GSOP - Flow Filtering and Load Balancing • Load Balancing: Type - Stateful, Type - GTP, Metric - Hashing, and Hashing Key - IMSI.
Application Ruleset	Flow Filter - Drop • Condition - GTP IMSI • IMSI - *, Type - Interface, and Interface - S5/S8
Destination	Tool port group - portgrp1

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.

> **NOTE:** For APF, no GigaSMART parameters are required unless combined with other GigaSMART operations.

7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure a Port Group

To configure a port group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter portgrp1 in the **Alias** field for the new Port Group and click **Tool** option.
5. Enable **Load Balancing**.
6. Choose the **Tool Ports** you want to attach to the port group.
7. Enter the weights for each tool port.
8. Click **Apply**. The created port group appears in the port group table.

Configure a First-Level Map

To configure an ingress (first level) map:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.

c. In the **New Map** dialog, do the following:

- **Alias** - Enter map1.
- **Alias** - Enter GTP- Control.
- (Optional) **Description** - Enter the description for the Traffic Policy.
- (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
- (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.

d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source.

a. Click **Sources**. The Source 1 page appears.

b. Select the network port 1/1/x4.

> **NOTE:** To configure or modify the port type, use the **Quick Port Editor**.

c. Click **Add**. The Source block appears on the canvas with the selected network port.

3. Select Source Ruleset.

a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The Source Ruleset page appears.

b. Click **Add Rules**. The Add Ruleset pane appears.

c. Configure Rule 1:

- Turn on **Pass** toggle and **Bi-directional** to pass GTP-u traffic.
- Select Port Source from the **Condition** list.
- Enter 2123 in the **Min** field.

d. Configure Rule 2: Click **+**.

- Turn on **Pass** toggle and **Bi-directional** to pass GTP-u traffic.
- Select Port Source in the **Condition** list.
- Enter 2152 in the **Min** field.

e. Click **Save**. The added rule appears in the **Source Ruleset** page.

f. Click **Close**. The By Rule block appears on the canvas with the configured rules.

4. Select Destination.

a. Click **Destinations**. The Destination 1 page appears.

b. Select the configured virtual port vp1.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

5. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

6. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure a Second-Level Map

To create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, matches IMSIs specified by the flow rules, and sends matching traffic to physical tool ports:

1. Access Traffic Policy.

a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

b. Create a new map with **Alias** as m1.

2. Select Source.

a. Click **Sources**. The Source 1 page appears with the list of available ports.

b. Select the virtual port vp1.

c. Click **Add**. The Source block appears on the canvas with the selected virtual port.

3. Select Application.

a. Expand the **Applications** component and then drag and drop **My GSOP**. The GSOP page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group you created earlier.

f. Select Flow Filtering and Load Balancing from the **Applications** list. By default, the application you select is enabled.

g. In the Load Balancing section, select Stateful option and Type as GTP.

h. Select Hashing in the **Metric** list.

i. Select IMSI option in the **Hashing Key** field.

j. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

4. Select Application Ruleset.

   a. Expand the Application Ruleset component and then drag and drop a **Flow Filter**. The Application Ruleset page appears.

   b. Click **Add Rules**. The Add Ruleset pane appears.

   c. Configure Rule 1: Turn on **Drop** toggle to drop the configured traffic.

   d. Select GTP IMSI from the **Condition** list, and enter the following fields:

      - **IMSI** - *, **Type** - Interface, **Interface** - S5/S8.

   e. Select GTP IMSI from the **Condition** list, and enter the following fields:

      - **IMSI** - *, **Type** - Interface, **Interface** - S11/S1U.

   f. Click **Save**. The Application Ruleset page appears with the configured rule.

   g. Click **Close**. The Flow Filter block appears on the canvas with the configured application rule.

5. Select Destination.

   a. Click **Destinations**. The Destination 1 page appears.

   b. Select the tool port group portgrp1.

   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

6. Deploy Traffic Policy.

   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

GigaSMART Rotational Sampling

Rotational Sampling samples sessions by International Mobile Subscriber ID (IMSI), using a rotating or sliding IMSI range, that rotates automatically on a configurable interval, without disrupting the treatment of existing sessions. It rotates through the entire pool of subscribers over time, such that all subscribers can be sampled in. The IMSI range size is configurable, the length of time for each period or window is configurable, and as is the length of time for each period or window. Additionally, the amount it shifts or slides every period is also configurable. This capability supports both network security and performance monitoring. Sampling should account for all active subscribers and be deterministic, so it is always evident whether a subscriber is included.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
Second-Level Map	
Source	Virtual port (vp1)
Source Ruleset	N/A
Applications	Application - Flow Sampling, Type - Flow Sampling - GTP
Application Ruleset	Type - Flow Sample GTP Enable Rotational Flow Sampling • Offset (%) - 1 - 99 • Timer (mins) - 15 - 4000 Condition - GTP, Enable Periodic Recalculation • Percentage - 50
Destination	Tool port 1/1/x4.

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The GigaSMARTGroups page appears.
4. Click **New**. The **Add GigaSMART Groups** page appears.
5. Enter a name for the GigaSMART group in the **Alias** field.

6. Select a GigaSMART engine port from the **Port List**.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Enable **GTP Overlap** checkbox.
8. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure a Second-Level Map

To configure Rotational Sampling in the second level map of GTP Flow Sample and GTP Flow Sample Overlap:

1. Access Traffic Policy.
    a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
    b. Create a new map with **Alias** as Flowsample-map1.
2. Select Source.
    a. Click **Sources**. The Source 1 page appears with the list of available ports.
    b. Select the virtual port vp1.
    c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application Ruleset.
    a. Expand the Application Ruleset component and then drag and drop a Flow Sample GTP. The **Application Ruleset** page appears.
    a. Click the **Settings** icon. The Map Configuration pane appears.
    b. Enable the **Rotational Flow Sampling** checkbox.

NOTE: In GigaVUE-FM, enabling Rotational flow sampling means it selects the flow sampling type (fstype) as rotational. Disabling the checkbox means the flow sampling type is default. The flow sampling type is only applicable for second level flow sample and second level flow sample overlap map types.

c. Enter the percentage of sessions to be sampled within the given interval in the **Offset** field.

d. Enter the value range from 15 to 45000 minutes in the **Timer** field.

e. Click **Save**. The Map Configuration pane closes.

f. Click **Add Rules**. The Add Ruleset pane appears.

g. Turn on **Pass** toggle to pass the configured traffic.

h. Select GTP from the **Condition** list.

i. Select the **Periodic Recalculation** checkbox, to enable dynamic session sampling.

NOTE: In GigaVUE-FM, you can edit the Periodic Recalculation flag only for the existing map rules.

j. Enter the required value in the **Percentage** field.

k. Click **Save**. The Application Ruleset page appears with the configured rule.

l. Click **Close**. The Flow Sample GTP block appears on the canvas with the configured application rule.

4. Select Application.

a. Expand the **Applications** component and then drag and drop **My GSOP**. The GSOP page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group you created earlier.

f. Select Flow Sampling, from the **Applications** list.

g. Select the type as **Flow Sampling - GTP**.

h. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

5. Select Destination.

a. Click **Destinations**. The Destination 1 page appears.

b. Select the tool port 1/1/x4.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

6. Deploy Traffic Policy.

   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

GigaSMART 4G RAN Correlation

A Radio Access Network (RAN) is part of a telecommunications system that connects individual radio devices to the core mobile network through radio connections.

RAN uses GTP correlation. It exposes RAN data fields from control plane (S11), for example:

- E-UTRAN Cell Global Identifier (ECGI)
- Tracking Area Identity (TAI)
- Tracking Area Code (TAC)

The above identifiers are extracted from the User location Information (ULI) and sends them to flow sampling and whitelisting.

The RAN correlation for 4G CUPS flow sampling filters require the ECGI, TAC and PLMN-ID values.

For 4G RAN based whitelisting requires ECGI includes the PLMN Id and the E-UTRAN Cell ID.

RAN data fields are available for filtering. GigaVUE-FM enables you to configure RAN based correlation for both 4G CUPS and 5G CUPS in the Flow Sample and Whitelist Maps. In CUPS solution, you can configure the RAN Correlation properties while creating and editing the properties in Flow Sample and Whitelist maps in Global policies for both CPN and UPN.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
Second-Level Map	
Source	Virtual Port (vp1)
Source Ruleset	N/A
Applications	• Applications - Load Balancing and Flow Sampling • Flow Sampling, type - GTP. • Load Balancing - Stateful, type - GTP.
Application Ruleset	Application Ruleset - Flow Sample GTP • Condition - GTP o Percentage - 100 and IMSI - 1234*, and PLMN ID - 310.210
Destination	Circuit GigaStream.

Configure Port Groups

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter a name (pg1) in the **Alias** field.
5. Select **Tool** option.
6. Enable **Load Balancing**.
7. Select the tool ports that you want to attach to the port group, For example 1/1/x9 - 1/1/x12.
8. Specify the weights for each port.
9. Click **Apply** to save. The created port group appears in the Port Groups page.

Create a GigaSMART Group

To create a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.

6. Select four GigaSMART engine ports from the **Port List**. For example, 3/1/e1, 23/1/e2, 23/3/e1, and 23/3/e2.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure a Second-Level Map

To create a second level map for Flow Sample GTP using RAN Correlation

1. Access Traffic Policy
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Create a new map with **Alias** as Map1.
2. Select Source
   a. Click **Sources**. The **Source 1** page appears with the list of available ports.
   b. Select the virtual port vp1.
   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
   e. Select the GigaSMART group (gs-grp) you created earlier.

f. Select Flow Sampling and Load Balancing from the **Applications** list. By default, the application you select is enabled.

g. In the Flow Sampling section, select GTP as the type.

h. In the Load Balancing section, select Stateful option and type as GTP.

i. Select Hashing from the **Metric** list.

j. Select IMSI option in the **Hashing Key** field.

k. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

4. Select Application Ruleset.

   a. Expand the Application Ruleset component and then drag and drop a Flow Sample GTP. The **Application Ruleset** page appears.

   b. Click **Add Rules**. The **Add Ruleset** pane appears.

   c. Select GTP from the **Condition** list and enter the given values:

      - Percentage - 100 and IMSI - 1234*, and PLMN ID - 310.210

   d. Click **Save**. The **Application Ruleset** page appears with the configured rule.

   e. Click **Close**. The By Rule block appears on the canvas with the configured application rule.

5. Select Destination.

   a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

   b. Select the port group pg1 that you created earlier.

   c. Click **Add**. The Destination block appears on the canvas with the selected GigaStream.

6. Deploy Traffic Policy.

   a. On the Traffic Policy canvas, click **Deploy** to implement the map configuration in the device.

7. Verify deployment.

   a. Status: In the Traffic Policy landing page, confirm that the map status is deployed and healthy.

   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

GigaSMART SIP/RTP Correlation

This is a load balancing configuration example of SIP/RTP.

SIP/RTP Load Balancing Example

This is a load balancing configuration example of SIP/RTP.

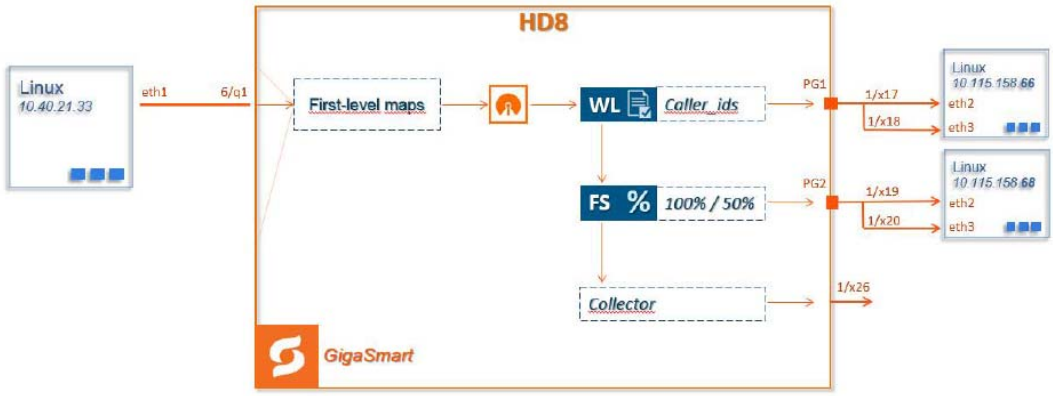


Figure 2 *SIP/RTP UDP/TCP Support*

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network Port (1/1/x1)
Source Ruleset	Condition - Port Destination, Value - 2123.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port-vp1

Configure a Port Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter PG-Sample in the **Alias** field.
5. Select **Tool** option.

6. Enable **Load Balancing**.
7. Choose the **Tool Ports** you want to attach to the port group, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
8. Specify the weights for each port as follows:
    - Weight 1/1/x6 - 5
    - Weight 1/1/x7 - 10
    - Weight 1/2/x3 - 20
    - Weight 1/2/x4 - 10
9. Click **Apply** to save. The created port group appears in the port group table.

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The GigaSMARTGroups page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**.
7. In the **SIP** section, go to **SIP Ports**, enter the values for SIP Port and RTP Port (Min and Max).
8. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.

7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Create a SIP Whitelist

To create SIP Whitelist:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Forward List**. The Forward List page appears.
4. Click **New** to add a Forward List. The Add Forward List pane appears.
5. Enter a name (MyIMSI) for the forward list in **Alias** field.
6. Select **Type** as SIP.
7. Select **Bulk Entry Operation** for uploading the IMSI list.
8. Select the **Import Type** as Upload from URL.
9. Select the **Operation Type** as Append.
10. Enter the URL (Example: http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx) in the **Remote URL** field.
11. Click **Apply**. The Forward List page appears with the created Forward List list.

Associate the SIP Whitelist to GigaSMART Group

To associate the SIP Whitelist with the GigaSMART group:

1. Go to **System > GigaSMART > GigaSMART Groups**. The GigaSMART Groups page appears.
2. Select the GigaSMART group (gs-grp) you crated earlier.
3. In Actions dropdown, select **Edit**. The Edit GigaSMART Group page appears.
4. In the **GigaSMART Parameters** section, go to SIP Forward List.
5. Select the forward list you created earlier from the **SIP Forward List Alias**.
6. Click **Apply**. The GigaSMART Groups page appears with the updated GigaSMART group.

Configure Traffic Policy

Create First-Level Maps

Map 1:

To create a first-level map:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle and select **Bi-directional** option to pass the configured traffic.
   d. Select Port Destination from the **Condition** list, and then enter 2123 in the **Value** field.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.

4. Select Destination.
    a. Click **Destinations**. The **Destination 1** page appears.
    b. Select the configured virtual port vp1.
    c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
    a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
    a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
    b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the first-level map (map2) with the following components and configurations:
    - Source port - Network Port (1/1/g1)
    - Source Ruleset - By Rule
        - Condition - Port Destination, Value - 2152
    - Destination port - Virtual Port (vp1)

Create Second-Level Maps

Map 1:

To create a second-level map for SIP Flow Sampling

1. Access Traffic Policy
    a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
    b. Create a new map with **Alias** as FlowSampe-SIP.
2. Select Source
    a. Click **Sources**. The **Source 1** page appears with the list of available ports.
    b. Select the virtual port vp1.
    c. Click **Add**. The Source block appears on the canvas with the selected virtual port.

3. Select Application
    a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
    b. Select either a standalone node or a cluster from the **Select Node** list.
    c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
    d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
    e. Select the GigaSMART group (gs-grp) you created earlier.
    f. Select Forward List and Load Balancing from the **Applications** list. By default, the applications you select are enabled.
    g. In the Forward List section, select SIP option.
    h. In the Load Balancing section, select Stateful option and type as GTP.
    i. Select Hashing from the **Metric** list and IMSI from the **Hashing Key** field.
    j. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.
4. Select Application Ruleset.
    a. Expand the Application Ruleset component and then drag and drop a Flow Sample SIP. The **Application Ruleset** page appears.
    b. Click **Add Rules**. The **Add Ruleset** pane appears.
    c. Rule 1: Select SIP from the **Condition** list and then select Caller ID from the **Type** list.
        - Enter 408* in the **Caller ID** field, and **Percentage** as 50.
    d. Rule 2: Select SIP from the **Condition** list and then select Caller ID from the **Type** list.
        - Enter 6501234* in the **Caller ID** field, and **Percentage** as 50.
    e. Click **Save**. The Application Ruleset page appears with the configured rules.
    f. Click **Close**. The Flow Sample SIP block appears on the canvas with the configured application rules.
5. Select Destination.
    a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.
    b. Select the tool port group (PG-Sample).
    c. Click **Add**. The Destination block appears on the canvas with the selected tool ports.
6. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

To create another second-level map for Flow Whitelist SIP

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (Flow Whitelist-SIP) with the following components and configurations:
   - Source port - Virtual Port (vp1)
   - Destination port - Tool port group (PG-Sample)
   - Application - Flow Sampling and Load Balancing
     - In the Flow Sampling section, select Flow Sampling - SIP as the **Type**.
     - In the Load Balancing section, select Stateful option and type as SIP.
     - Select the **Metric** as Hashing and IMSI from the **Hashing Key** options.
   - Application ruleset - Flow Sample GTP
     - **Condition** - SIP, **Type** - Caller ID
     - **Percentage** - 50, and **Caller ID** - 408*
2. Deploy and verify the Traffic Policy.

GTP Overlap Flow Sampling

This topic explains the workflow required to configure a Traffic Policy for GTP Overlap Flow Sampling.

Traffic from a single network port goes to a single first level map (mapLevel1-GTP) which directs GTP-Control, and GTP-User traffic to a virtual port (vp1). Traffic from vp1 is replicated to two GTP whitelisting maps (WLMAP1 and WLMAP2) and two GTP flow sampling maps (FSMAP1 and FSMAP2), which is then forwarded to the final port-group destinations, pg1 and pg2, for load balancing.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	
Source Ruleset	Source Ruleset - By Rule • Rule 1: Condition - Port Destination, Value - 2123 • Rule 2: Condition - Port Destination, Value - 2152
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port (vp1)
Second-Level Map1	
Source	Virtual Port (vp1)
Source Ruleset	N/A
Applications	Application - Forward List and Load Balancing. • Forward List: Type - Flow Sampling - GTP • Load Balancing: Stateful, Type - GTP, Metric - Hashing, and Hashing Key - IMSI
Application Ruleset	Application Ruleset - Flow Sample Overlap • Condition - GTP. o APN - mobile.com o Version - V2
Destination	Tool port group (pg1).
Second-Level Map2	
Source	Virtual Port (vp1)
Source Ruleset	N/A
Applications	Applications - Flow Sampling and Load Balancing. • Flow Sampling: type - GTP • Load Balancing: Stateful, type - GTP. • Metric - Hashing, Hashing Key - IMSI.
Application Ruleset	Application Ruleset - Flow Sample Overlap • Condition - GTP. o IMSI - 3102609835*, IMEI - 35609507*, and Percentage - 20
Destination	PortGroup - (pg2)
Second-Level Map3	
Source	Virtual port (vp1)
Source Ruleset	N/A
Applications	Application - Forward List and Load Balancing. • Forward List: Type - Flow Sampling - GTP • Load Balancing: Stateful, Type - GTP, Metric - Hashing, and Hashing Key - IMSI

Map Components	Sample Configuration
Application Ruleset	Application Ruleset - Flow Sample Overlap • Condition - GTP. o APN - mobile.com o Version - V2
Destination	Port Group (pg2)
Second-Level Map4	
Source	Virtual port (vp1)
Source Ruleset	N/A
Applications	Application - Flow Sampling and Load Balancing. • Flow Sampling: Type - Flow Sampling - GTP • Load Balancing: Stateful, Type - GTP, Metric - Hashing, and Hashing Key - IMSI
Application Ruleset	Application Ruleset - Flow Sample Overlap • Condition - GTP. o IMSI - 3102609835*, IMEI - 35609507*, and Percentage - 20
Destination	Port Group (pg2)

Configure GigaStream

To configure GigaStream:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.Go
3. Go to **Ports > Port Groups**> **GigaStream**.
4. Click **New**. The New GigaStream pane appears.
5. Enter a name (gs1-gtp) in the **Alias** field.
6. Select Tool GigaStream as the **Type**.
7. Select the ports 1/1/x16 and 1/1/x17 from the **Ports** list.
8. Select the **Weighting** as required.
9. Click **Apply**. The GigaStream page appears with the created GigaStream.
10. Follow the steps 1 - 8 to create another GigaStream with **Alias** as gs2-gtp.

Configure Port Groups

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .

3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter pg1 in the **Alias** field.
5. Select **Tool** option.
6. Enable **Load Balancing**.
7. Choose the **Tool Ports** you want to attach to the port group, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
8. Specify the weights for each port as follows:
    - Weight 1/1/x6 - 5
    - Weight 1/1/x7 - 10
9. Select gs1 from the **GigaStream** field.
10. Click **Apply** to save. The created port group appears in the Port Groups page.
11. Follow the steps 1 - 10 and create another port group (pg2) with:
    - Ports - 1/1/x8 and 1/1/x10
    - GigaStream - gs2

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Create a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.

5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Create a GTP Whitelist

To create GTP Whitelist:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Forward List**. The Forward List page appears.
4. Click **New** to add a Forward List. The Add Forward List pane appears.
5. Enter a name (MyIMSI) for the forward list in **Alias** field.
6. Select **Type** as GTP.
7. Select **Bulk Entry Operation** for uploading the IMSI list.
8. Select the **Import Type** as Upload from URL.
9. Select the **Operation Type** as Append.
10. Enter the URL (Example: http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx) in the **Remote URL** field.
11. Click **Apply**. The Forward List page appears with the created Forward List list.

Associate the GTP Whitelist to GigaSMART Group

To associate the GTP Whitelist with the GigaSMART group:

1. Go to **System > GigaSMART > GigaSMART Groups**. The GigaSMART Groups page appears.
2. Select the GigaSMART group (gs-grp) you crated earlier.
3. In Actions dropdown, select **Edit**. The Edit GigaSMART Group page appears.
4. In the **GigaSMART Parameters** section, go to GTP Forward List.
5. Select the forward list you created earlier from the **GTP Forward List Alias**.
6. Click **Apply**. The GigaSMART Groups page appears with the updated GigaSMART group.

Configure Traffic Policy

Create a First-Level Map

To create a first-level map:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter GTP- Control.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle and select **Bi-directional** option to pass the configured traffic.
      - Rule 1 - Select Port Destination from the **Condition** list, and then enter 2123 in the **Value** field.
      - Rule 2 - Select Port Destination from the **Condition** list, and then enter 2152 in the **Value** field.
   d. Click **Save**. The added rule appears in the **Source Ruleset** page.
   e. Click **Close**. The By Rule block appears on the canvas with the configured rule.

4. Select Destination.
    a. Click **Destinations**. The **Destination 1** page appears.
    b. Select the configured virtual port vp1.
    c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
    a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
    a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
    b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Create Second-Level Maps

Map 1:

To create a second-level map to forward the traffic that matches the IMSI whitelist for GTP Version 1 traffic to the tool port group

1. Access Traffic Policy
    a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
    b. Create a new map with **Alias** as wlmap1.
2. Select Source
    a. Click **Sources**. The **Source 1** page appears with the list of available ports.
    b. Select the virtual port vp1.
    c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application
    a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
    b. Select either a standalone node or a cluster from the **Select Node** list.
    c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
    d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group (gs-grp) you created earlier.

f. Select Forward List and Load Balancing from the **Applications** list. By default, the applications you select are enabled.

g. In the Forward List section, select GTP option.

h. In the Load Balancing section, select Stateful option and type as GTP.

i. Select Hashing from the **Metric** list and IMSI from the **Hashing Key** field.

j. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

4. Select Application Ruleset.

   a. Expand the Application Ruleset component and then drag and drop a Flow Whitelist Overlap. The **Application Ruleset** page appears.

   b. Click **Add Rules**. The **Add Ruleset** pane appears.

   c. Select GTP from the **Condition** list and enter the given values:

      - APN - mobile.com
      - Version - V1

   d. Click **Save**. The Add Ruleset pane appears with the configured rule.

   e. Click **Close**. The Flow Whitelist Overlap block appears on the canvas with the configured application rules.

5. Select Destination.

   a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

   b. Select the tool port group (pg1).

   c. Click **Add**. The Destination block appears on the canvas with the selected tool ports.

6. Deploy Traffic Policy.

   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

To create a second-level map for GTP flow sampling, that samples the traffic flow and forward to the tool port group if there is not a match to an IMSI in the whitelist,

1. Create a second level map (fsmap1) with the following changes:
   a. Source - Virtual Port (vp1)
   b. Application - Flow Sampling and Load Balancing
      - Flow Sampling: Type - Flow Sampling - GTP
      - Load Balancing: Stateful, Type - GTP, Metric - Hashing, and Hashing Key - IMSI
   c. Application Ruleset - Flow Sample Overlap with **Condition** as GTP.
      - IMSI - 3102609835*, IMEI - 35609507*, and Percentage - 20
   d. Destination - Port Group (pg2)

Map 3:

To create a second-level map to forward the traffic that matches the IMSI whitelist for GTP Version 2 traffic to the tool port group

1. Create a second level map (wlmap2) with the following changes:
   a. Source - Virtual Port (vp1)
   b. Application - Forward List and Load Balancing
      - Forward List: Type - GTP
      - Load Balancing: Stateful, Type - GTP, Metric - Hashing, and Hashing Key - IMSI
   c. Application Ruleset - Flow Whitelist Overlap with **Condition** as GTP.
      - APN - mobile.com
      - Version - V2
   d. Destination - Port Group (pg2)

Map 4:

To create a second-level map for GTP flow sampling, that samples the traffic flow and forward to the tool port group if there is not a match to an IMSI in the whitelist,

1. Create a second level map (fsmap2) with the following changes:
   a. Source - Virtual Port (vp1)
   b. Application - Flow Sampling and Load Balancing
      - Flow Sampling: Type - Flow Sampling - GTP
      - Load Balancing: Stateful, Type - GTP, Metric - Hashing, and Hashing Key - IMSI
   c. Application Ruleset - Flow Sample Overlap with **Condition** as GTP.
      - IMSI - 3102609835*, IMEI - 35609507*, and Percentage - 20

d. Destination - Port Group (pg2).

GigaSMART GTP Whitelisting and GTP Flow Sampling

Required License: GTP Filtering & Correlation and FlowVUE
Supported Devices : GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3 ,GigaVUE-HC1-Plus.

Use GTP forward listing and GTP flow sampling to provide subsets of GTP correlated flows to tools. GTP forward listing selects specific subscribers based on IMSI, while GTP flow sampling uses map rules to select subscribers. Starting in software version 4.8, GigaSMART supports GTP overlap mapping, which combines both forward listing and flow sampling maps as part of a map group.

The following sections explain the Traffic Policy configurations of GTP Whitelisting and GTP Flow Sampling for different real-time examples.

- GTP Whitelisting
- GTP Whitelisting with Multiple Maps
- GTP Flow Sampling
- GTP Whitelisting, GTP Flow Sampling, and Load Balancing
- GTP Flow Sampling with Multiple Maps
- APN for Forward List and GTP Flow Sampling
- APN for FTP Whitelisting, APN and QCI for GTP Flow Sampling

GTP Whitelisting

This topic explains the workflow required to configure a Traffic Policy for GTP Whitelisting. The traffic from network ports are forwarded to the three first level maps (GTP-Control, GTP-User, and Fragments-Not_First) and then to the virtual port (vp1). If there is a match to an IMSI in the whitelist (MyIMSIs), it is forwarded to the tool port.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map1	
Source	Network Port - 1/1/x1
Source Ruleset	Condition - Port Destination, Value - 2123.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port-vp1
First-Level Map2	
Source	1/1/x1
Source Ruleset	Condition - Port Destination, Value - 2152.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port - vp1
First-Level Map3	
Source	1/1/x1
Source Ruleset	Condition - IP Fragmentation and Value - All Frag No First.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port - vp1
Second-Level Map	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	• Applications - Forward List and Load Balancing. • Load Balancing - Stateful, type - GTP. • Metric - Hashing, Hashing Key - IMSI.
Application Ruleset	N/A
Destination	Tool Group - PG-Sample.

Configure a Port Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter PG-Sample in the **Alias** field.
5. Select **Tool** option.
6. Enable **Load Balancing**.

7. Select the tool ports you want to attach to the port group, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
8. Specify the weights for each port as follows:
    - Weight 1/1/x6 - 5
    - Weight 1/1/x7 - 10
    - Weight 1/2/x3 - 20
    - Weight 1/2/x4 - 10
9. Click **Apply** to save. The created port group appears in the Port Groups page.

Configure a GigaSMART Group

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Create a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Create a GTP Whitelist

To create GTP Whitelist:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Forward List**. The Forward List page appears.
4. Click **New** to add a Forward List. The Add Forward List pane appears.
5. Enter a name (MyIMSI) for the forward list in **Alias** field.
6. Select **Type** as GTP.
7. Select **Bulk Entry Operation** for uploading the IMSI list.
8. Select the **Import Type** as Upload from URL.
9. Select the **Operation Type** as Append.
10. Enter the URL (Example: http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx) in the **Remote URL** field.
11. Click **Apply**. The Forward List page appears with the created Forward List list.

Associate the GTP Whitelist to GigaSMART Group

To associate the GTP Whitelist with the GigaSMART group:

1. Go to **System > GigaSMART > GigaSMART Groups**. The GigaSMART Groups page appears.
2. Select the GigaSMART group (gs-grp) you crated earlier.
3. In **Actions** dropdown, select **Edit**. The Edit GigaSMART Group page appears.
4. In the **GigaSMART Parameters** section, go to GTP Forward List.
5. Select the forward list you created earlier from the **GTP Forward List Alias**.
6. Click **Apply**. The GigaSMART Groups page appears with the updated GigaSMART group.

Configure Traffic Policy

Create First-Level Maps

Map 1:

To create a first-level map:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter GTP- Control.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle and select **Bi-directional** option to pass the configured traffic.
   d. Select Port Destination from the **Condition** list, and then enter 2123 in the **Value** field.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears.
   b. Select the configured virtual port vp1.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the second first-level map (GTP-User) with the following source rule:
   - Source ruleset
     - Condition - Port Destination and Value -2152.
2. Deploy and verify the Traffic Policy.

Map 3:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the second first-level map (Fragments-Not-First) with the following source rule:
   - Source ruleset
     - Condition - IP Fragmentation and Value - All Frag No First.
2. Deploy and verify the Traffic Policy

Create a Second-Level Map

To create a second-level map to forward traffic that match the IMSI whitelist to the tool port

1. Access Traffic Policy
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Create a new map with **Alias** as GTP-Whitelist.
2. Select Source
   a. Click **Sources**. The **Source 1** page appears with the list of available ports.
   b. Select the virtual port vp1.

c. Click **Add**. The Source block appears on the canvas with the selected virtual port.

3. Select Application Ruleset.
   a. Expand the Application Ruleset component and then drag and drop a **Flow Whitelist GTP**. The **Application Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Select GTP from the **Condition** list, and select **Version** as V1.
   d. Click **Save**. The added rule appears in the **Application Ruleset** page.
   e. Click **Close**. The Flow Whitelist GTP block appears on the canvas with the configured application rule.
4. Select Application
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
   e. Select the GigaSMART group (gs-grp) you created earlier.
   f. Select Forward List and Load Balancing from the **Applications** list. By default, the applications you select are enabled.
   g. In the Load Balancing section, select Stateful option and type as GTP.
   h. Select Hashing from the **Metric** list.
   i. Select IMSI option in the **Hashing Key** field.
   j. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.
5. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.
   b. Select the tool port group PG-Sample.
   c. Click **Add**. The Destination block appears on the canvas with the selected tool ports.
6. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
7. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

GTP Whitelisting with Multiple Maps

This topic explains the workflow required to configure a Traffic Policy for GTP Whitelisting.

This example includes multiple GTP whitelisting maps, which provide a more granular selection of tool ports. The traffic from the network ports are forwarded to the three first level maps (GTP-Control, GTP-User, and Fragments-Not-First) and then to the virtual port (vp1). Two second level whitelist maps are configured. The first map specifies a rule for version 1 traffic and the second map specifies a rule for version 2 traffic.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map1	
Source	Network Ports - 8/1/x40 and 8/1/x6
Source Ruleset	Condition - Port Destination, Value - 2123.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port-vp1
First-Level Map2	
Source	1/1/x1
Source Ruleset	Condition - Port Destination, Value - 2152.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port - vp1
First-Level Map3	
Source	1/1/x1
Source Ruleset	Condition - IP Fragmentation and Value - All Frag No First.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port - vp1

Map Components	Sample Configuration
Second-Level Map1	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	Applications - Forward List and Load Balancing. • Forward List - GTP. • Load Balancing - Stateful, type - GTP. • Metric - Hashing, Hashing Key - IMSI.
Application Ruleset	Application Ruleset - Whitelist GTP o Conditions - GTP, Version - V1.
Destination	Tool Group - PG-Sample.
Second-Level Map2	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	N/A
Application Ruleset	Application Ruleset - Whitelist GTP • Conditions - GTP, Version - V2.
Destination	1/2/x24

Configure a Port Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter PG-Sample in the **Alias** field.
5. Select **Tool** option.
6. Enable **Load Balancing**.
7. Choose the **Tool Ports** you want to attach to the port group, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
8. Specify the weights for each port as follows:
    - Weight 1/1/x6 - 5
    - Weight 1/1/x7 - 10
    - Weight 1/2/x3 - 20
    - Weight 1/2/x4 - 10
9. Click **Apply** to save. The created port group appears in the port group table.

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Create a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Create a GTP Whitelist

To create GTP Whitelist:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Forward List**. The Forward List page appears.
4. Click **New** to add a Forward List. The Add Forward List pane appears.
5. Enter a name (MyIMSI) for the forward list in **Alias** field.
6. Select **Type** as GTP.
7. Select **Bulk Entry Operation** for uploading the IMSI list.

8. Select the **Import Type** as Upload from URL.
9. Select the **Operation Type** as Append.
10. Enter the URL (Example: http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx) in the **Remote URL** field.
11. Click **Apply**. The Forward List page appears with the created Forward List list.

Associate GTP Whitelist to GigaSMART Group

To associate the GTP Whitelist with the GigaSMART group:

1. Go to **System > GigaSMART > GigaSMART Groups**. The GigaSMART Groups page appears.
2. Select the GigaSMART group (gs-grp) you crated earlier.
3. In Actions dropdown, select **Edit**. The Edit GigaSMART Group page appears.
4. In the **GigaSMART Parameters** section, go to GTP Forward List.
5. Select the forward list you created earlier from the **GTP Forward List Alias**.
6. Click **Apply**. The GigaSMART Groups page appears with the updated GigaSMART group.

Configure Traffic Policy

Create First-Level Maps

Map 1:

To create a first-level map:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter GTP- Control.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.

- (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.

d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network ports 8/1/x40 and 8/1/x6.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle and select **Bi-directional** option to pass the configured traffic.
   d. Select Port Destination from the **Condition** list, and then enter 2123 in the **Value** field.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears.
   b. Select the configured virtual port vp1.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the second first-level map (GTP-User) with the following source rule:
   - Source ruleset
     - Condition - Port Destination and Value -2152.
2. Deploy and verify the Traffic Policy.

Map 3:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the third first-level map ((Fragments-Not-First) with the following source rule:
   - Source ruleset
     - Condition - IP Fragmentation and Value - All Frag No First.
2. Deploy and verify the Traffic Policy

Create Second-Level Maps

To create a second-level map to forward traffic that match the IMSI whitelist to the tool ports

Map 1:

1. Access Traffic Policy
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Create a new map with **Alias** as GTP-Whitelist-v1.
2. Select Source
   a. Click **Sources**. The **Source 1** page appears with the list of available ports.
   b. Select the virtual port vp1.
   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
   e. Select the GigaSMART group (gs-grp) you created earlier.

f. Select Forward List and Load Balancing from the **Applications** list. By default, the applications you select are enabled.

g. In the Forward List section, select GTP option.

h. In the Load Balancing section, select Stateful option and type as GTP.

i. Select Hashing from the **Metric** list.

j. Select IMSI option in the **Hashing Key** field.

k. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

4. Select Application Ruleset.

   a. Expand the Application Ruleset component and then drag and drop a **Whitelist GTP**. The **Application Ruleset** page appears.

   b. Click **Add Rules**. The **Add Ruleset** pane appears.

   c. Select GTP from the **Condition** list, and Version as V1.

   d. Click **Save**. The Application Ruleset page appears with the configured rule.

   e. Click **Close**. The Whitelist GTP block appears on the canvas with the configured application rule.

5. Select Destination.

   a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

   b. Select the tool port group PG-Sample.

   c. Click **Add**. The Destination block appears on the canvas with the selected tool ports.

6. Deploy Traffic Policy.

   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (GTP-Whitelist-v2) with the following components and configurations:

- Source port - Virtual Port (vp1)
- Destination port - 1/2/x24
- Application ruleset - Whitelist GTP
  - Condition - GTP, Version -V2

2. Deploy and verify the Traffic Policy.

GTP Flow Sampling

This topic explains the workflow required to configure a Traffic Policy for GTP Flow Sampling.

Traffic from the network ports are forwarded to the three first level maps (GTP-Control, GTP-User, and Fragments-Not_First) and then to the virtual port (vp1). The traffic flow is sampled based on the rules in one flow sampling map (GTP-Sample-01). The flow sampling rules specify IMSI, IMEI, and MSISDN numbers, as well as the percentage to sample. Packets are then accepted or rejected. Accepted packets are forwarded to a port. Rejected packets are dropped. Packets that do not match a rule will be passed to subsequent maps.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map1	
Source	Network Port - 1/1/x1
Source Ruleset	Condition - Port Destination, Value - 2123.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port-vp1
First-Level Map2	
Source	1/1/x1
Source Ruleset	Condition - Port Destination, Value - 2152.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port - vp1
First-Level Map3	
Source	1/1/x1
Source Ruleset	Condition - IP Fragmentation and Value - All Frag No First.
Application Ruleset	N/A
Applications	N/A

Map Components	Sample Configuration
Destination	Virtual Port - vp1
Second-Level Map	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	Applications - Flow Sampling and Load Balancing. • Load Balancing - Stateful, type - GTP. • Metric - Hashing, Hashing Key - IMSI.
Application Ruleset	Application Ruleset - Flow Sample GTP o Conditions - GTP o Percentage - 50 and IMEI - 01416800* o Percentage - 80 and IMSI - 46* o Percentage - 25, MSISDN - 1509* o Percentage - 15, IMSI - 01400* o Percentage - 20, IMSI - 31*, and MSISDN - 1909*
Destination	Tool Group - PG-Sample.

Configure a Port Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter PG-Sample in the **Alias** field.
5. Select **Tool** option.
6. Enable **Load Balancing**.
7. Choose the **Tool Ports** you want to attach to the port group, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
8. Specify the weights for each port as follows:
   - Weight 1/1/x6 - 5
   - Weight 1/1/x7 - 10
   - Weight 1/2/x3 - 20
   - Weight 1/2/x4 - 10
9. Click **Apply** to save. The created port group appears in the port group table.

Create a GigaSMART Group

To create a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Create a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Create a GTP Whitelist

To create GTP Whitelist:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Forward List**. The Forward List page appears.
4. Click **New** to add a Forward List. The Add Forward List pane appears.
5. Enter a name (MyIMSI) for the forward list in **Alias** field.
6. Select **Type** as GTP.
7. Select **Bulk Entry Operation** for uploading the IMSI list.

8. Select the **Import Type** as Upload from URL.
9. Select the **Operation Type** as Append.
10. Enter the URL (Example: http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx) in the **Remote URL** field.
11. Click **Apply**. The Forward List page appears with the created Forward List list.

Associate the GTP Whitelist to GigaSMART Group

To associate the GTP Whitelist with the GigaSMART group:

1. Go to **System > GigaSMART > GigaSMART Groups**. The GigaSMART Groups page appears.
2. Select the GigaSMART group (gs-grp) you crated earlier.
3. In Actions dropdown, select **Edit**. The Edit GigaSMART Group page appears.
4. In the **GigaSMART Parameters** section, go to GTP Forward List.
5. Select the forward list you created earlier from the **GTP Forward List Alias**.
6. Click **Apply**. The GigaSMART Groups page appears with the updated GigaSMART group.

Configure Traffic Policy

Create First-level Maps

Map 1:

To create a first-level map:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter map1.
      - **Alias** - Enter GTP- Control.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.

- (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.

d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle and select **Bi-directional** option to pass the configured traffic.
   d. Select Port Destination from the **Condition** list, and then enter 2123 in the **Value** field.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears.
   b. Select the configured virtual port vp1.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the second first-level map (GTP-User) with the following source rule:
   - Source ruleset
     - Condition - Port Destination and Value -2152.
2. Deploy and verify the Traffic Policy.

Map 3:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the third first-level map ((Fragments-Not-First) with the following source rule:
   - Source ruleset
     - Condition - IP Fragmentation and Value - All Frag No First.
2. Deploy and verify the Traffic Policy

Create a Second-level Map

To create a second-level map to forward traffic that match the IMSI whitelist to the tool port

1. Access Traffic Policy
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Create a new map with **Alias** as PG-Sample-01.
2. Select Source
   a. Click **Sources**. The **Source 1** page appears with the list of available ports.
   b. Select the virtual port vp1.
   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
   e. Select the GigaSMART group (gs-grp) you created earlier.
   f. Select Flow Sampling and Load Balancing from the **Applications** list. By default, the applications you select are enabled.
   g. In the Flow Sampling section, select the type as Flow Sampling - GTP.

h. In the Load Balancing section, select Stateful option and type as GTP.

i. Select Hashing from the **Metric** list and IMSI from the **Hashing Key** field.

j. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

4. Select Application Ruleset.

   a. Expand the Application Ruleset component and then drag and drop a Flow Sample GTP. The **Application Ruleset** page appears.

   b. Click **Add Rules**. The **Add Ruleset** pane appears.

   c. Select GTP from the **Condition** list and enter the given values:

      - Percentage - 50 and IMEI - 01416800*
      - Percentage - 80 and IMSI - 46*
      - Percentage - 25, MSISDN - 1509*
      - Percentage - 15, IMSI - 01400*
      - Percentage - 20, IMSI - 31*, and MSISDN - 1909*

   d. Click **Save**. The Application Ruleset page appears with the configured rule.

   e. Click **Close**. The Flow Sample GTP block appears on the canvas with the configured application rules.

5. Select Destination.

   a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

   b. Select the tool port group (PG-Sample).

   c. Click **Add**. The Destination block appears on the canvas with the selected tool ports.

6. Deploy Traffic Policy.

   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

GTP Whitelisting, GTP Flow Sampling, and Load Balancing

This topic explains the workflow required to configure a Traffic Policy for GTP Whitelisting, GTP Flow Sampling, and Load Balancing.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map1	
Source	Network Port - 1/1/x1
Source Ruleset	Condition - Port Destination, Value - 2123.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port-vp1
First-Level Map2	
Source	1/1/x1
Source Ruleset	Condition - Port Destination, Value - 2152.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port - vp1
First-Level Map3	
Source	1/1/x1
Source Ruleset	Condition - IP Fragmentation and Value - All Frag No First.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port - vp1
Second-Level Map1	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	Applications - Forward List and Load Balancing. • Load Balancing - Stateful, type - GTP. • Metric - Hashing, Hashing Key - IMSI.
Application Ruleset	Application Ruleset - Flow Whitelist GTP . • Condition - GTP, Version - V1.
Destination	Tool Group - PG-Sample.

Map Components	Sample Configuration
Second-Level Map2	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	Applications - Flow Sampling and Load Balancing. • Flow Sampling. type - GTP • Load Balancing - Stateful, type - GTP. • Metric - Hashing, Hashing Key - IMSI.
Application Ruleset	Application Ruleset - Flow Sample GTP • Conditions - GTP o Percentage - 50 and IMEI - 01416800* o Percentage - 80 and IMSI - 46* o Percentage - 25, MSISDN - 1509* o Percentage - 15, IMSI - 01400* o Percentage - 20, IMSI - 31*, and MSISDN - 1909*
Destination	Tool Group - PG-Sample.

The traffic from network ports are forwarded to the three first level maps (GTP-Control, GTP-User, and Fragments-Not_First) and then to the virtual port (vp1). If there is a match to an IMSI in the whitelist (MyIMSIs), it is forwarded to the tool port. If there is a match to an IMSI in the whitelist (MyIMSIs), it is forwarded to the port group (PG-Whitelist) for load balancing.

NOTE: The tool ports in the port group should be on the same node as the GigaSMART group and GigaSMART operation.

If there is not a match to an IMSI in the whitelist, the traffic flow is sampled based on the rules in the flow sampling map (GTP-Sample-01). The flow sampling rules specify IMSI, IMEI, and MSISDN numbers, as well as the percentage to sample. Packets are then accepted or rejected. Accepted packets are forwarded to the port group (PG-Sample) for load balancing. Rejected packets are dropped. Packets that do not match a rule will be passed to subsequent maps.

Configure a Port Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter PG-Sample in the **Alias** field.
5. Select **Tool** option.
6. Enable **Load Balancing**.

7. Choose the **Tool Ports** you want to attach to the port group, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
8. Specify the weights for each port as follows:
    - Weight 1/1/x6 - 5
    - Weight 1/1/x7 - 10
    - Weight 1/2/x3 - 20
    - Weight 1/2/x4 - 10
9. Click **Apply** to save. The created port group appears in the port group table.
10. Follow the Steps 1 - 9 to create another tool port group with PF-Sample as **Alias**.

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Create a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Create a GTP Whitelist

To create GTP Whitelist:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Forward List**. The Forward List page appears.
4. Click **New** to add a Forward List. The Add Forward List pane appears.
5. Enter a name (MyIMSI) for the forward list in **Alias** field.
6. Select **Type** as GTP.
7. Select **Bulk Entry Operation** for uploading the IMSI list.
8. Select the **Import Type** as Upload from URL.
9. Select the **Operation Type** as Append.
10. Enter the URL (Example: http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx) in the **Remote URL** field.
11. Click **Apply**. The Forward List page appears with the created Forward List list.

Associate the GTP Whitelist to GigaSMART Group

To associate the GTP Whitelist with the GigaSMART group:

1. Go to **System > GigaSMART > GigaSMART Groups**. The GigaSMART Groups page appears.
2. Select the GigaSMART group (gs-grp) you crated earlier.
3. In Actions dropdown, select **Edit**. The Edit GigaSMART Group page appears.
4. In the **GigaSMART Parameters** section, go to GTP Forward List.
5. Select the forward list you created earlier from the **GTP Forward List Alias**.
6. Click **Apply**. The GigaSMART Groups page appears with the updated GigaSMART group.

Optional - add a single IMSI to GTP whitelist

Configure Traffic Policy

Create First-level Maps

Map 1:

To create a first-level map:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter GTP- Control.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle and select **Bi-directional** option to pass the configured traffic.
   d. Select Port Destination from the **Condition** list, and then enter 2123 in the **Value** field.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.

4. Select Destination.
    a. Click **Destinations**. The **Destination 1** page appears.
    b. Select the configured virtual port vp1.
    c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
    a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
    a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
    b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the first-level map (GTP-User) with the following source rule:
    - Source ruleset
        - Condition - Port Destination and Value -2152.
2. Deploy and verify the Traffic Policy.

Map 3:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the first-level map (Fragments-Not-First) with the following source rule:
    - Source ruleset
        - Condition - IP Fragmentation and Value - All Frag No First.
2. Deploy and verify the Traffic Policy

Create a Second-level Map

Map 1:

To create a second-level map to forward traffic that match the IMSI whitelist to the tool port

1. Access Traffic Policy
    a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

b. Create a new map with **Alias** as GTP-Whitelist.

2. Select Source
   a. Click **Sources**. The **Source 1** page appears with the list of available ports.
   b. Select the virtual port vp1.
   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
   e. Select the GigaSMART group (gs-grp) you created earlier.
   f. Select Forward List and Load Balancing from the **Applications** list. By default, the applications you select are enabled.
   g. In the Forward List section, select GTP option.
   h. In the Load Balancing section, select Stateful option and type as GTP.
   i. Select Hashing from the **Metric** list.
   j. Select IMSI option in the **Hashing Key** field.
   k. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.
4. Select Application Ruleset.
   a. Expand the Application Ruleset component and then drag and drop a **Flow Whitelist GTP**. The **Application Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Select GTP from the **Condition** list, and Version as V1.
   d. Click **Save**. The Application Ruleset page appears with the configured rule.
   e. Click **Close**. The Flow Whitelist GTP block appears on the canvas with the configured application rule.
5. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.
   b. Select the tool port group (PG-Sample).

c. Click **Add**. The Destination block appears on the canvas with the selected tool ports.

6. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
7. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

To create a second-level map to flow sample the traffic if these is no match to the IMSI whitelist and to forward the accepted packets to the load balancing port group

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (GTP-Sample-01) with the following components and configurations:
   - Source port - Virtual Port (vp1)
   - Destination port - Too port group (PG-Sample)
   - Application - Flow Sampling and Load Balancing
     - In the Flow Sampling section, select Flow Sampling - GTP as the **Type**.
     - In the Load Balancing section, select Stateful option and type as GTP.
     - Select the **Metric** as Hashing and IMSI from the **Hashing Key** options.
   - Application ruleset - Flow Sample GTP
     - Select GTP from the **Condition** list and enter the given values:
     - Percentage - 50 and IMEI - 01416800*, and IMSI - 31*
     - Percentage - 80 and IMSI - 46*
     - Percentage - 25, MSISDN - 1509*
     - Percentage - 15, IMSI - 01400*, and IMSI - 31*
     - Percentage - 20, IMSI - 31*, and MSISDN - 1909*
2. Deploy and verify the Traffic Policy.

GTP Flow Sampling with Multiple Maps

This topic explains the workflow required to configure a Traffic Policy for GTP Flow Sampling with multiple maps.

Traffic from the network ports are forwarded to the three first level maps (GTP-Control, GTP-User, and Fragments-Not_First) and then to the virtual port (vport1). If there is a match to an IMSI in the whitelist (VoLTE_1MM), it is forwarded to the port-group (PG-Whitelist-1) for load balancing.

NOTE: The tool ports in the port group should be on the same node as the GigaSMART group and GigaSMART operation.

If there is not a match to an IMSI in the whitelist, the traffic flow is sampled based on the rules in the four flow sampling maps (GTP-Sample-1 to GTP-Sample-4).

The flow sampling rules in each map specify IMSI, IMEI, and MSISDN numbers, as well as the percentage to sample. Packets are then accepted or rejected. Accepted packets are forwarded to the port-group (PG-Sample-1 to PG-Sample-4) for load balancing. Rejected packets are dropped. Packets that do not match a rule will be passed to subsequent maps, in this example, to a shared collector.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map1	
Source	Network Port - 1/1/x1
Source Ruleset	Condition - Port Destination, Value - 2123.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port-vp1
First-Level Map2	
Source	1/1/x1
Source Ruleset	Condition - Port Destination, Value - 2152.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port - vp1

Map Components	Sample Configuration
First-Level Map3	
Source	1/1/x1
Source Ruleset	Condition - IP Fragmentation and Value - All Frag No First.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual port (vp1)
Second-Level Map1	
Source	Virtual port (vp1)
Source Ruleset	N/A
Applications	Applications - Forward List and Load Balancing. • Load Balancing - Stateful, type - GTP. • Metric - Hashing, Hashing Key - IMSI.
Application Ruleset	Application Ruleset - Flow Whitelist GTP . • Condition - GTP, Percentage - 20, Version - V1.
Destination	Tool Group (PG-Sample.)
Second-Level Map2	
Source	Virtual port - vp1
Source Ruleset	N/A
Applications	Applications - Flow Sampling and Load Balancing. • Flow Sampling. type - GTP • Load Balancing - Stateful, type - GTP. • Metric - Hashing, Hashing Key - IMSI.
Application Ruleset	Application Ruleset - Flow Sample GTP • Conditions - GTP o Percentage - 75, IMSI - 3182609833* and IMEI - 35609506* o Percentage - 10, IMSI - 3182609834* and IMEI - 3560950* o Percentage - 20, IMSI - 31826098350* and IMEI - 356095* o Percentage - 20, IMSI - 31826098351* and IMEI - 35609* o Percentage - 20, IMSI - 31826098352* and IMEI - 3560* o Percentage - 20, IMSI - 31826098353* and IMEI - 356* o Percentage - 20, IMSI - 31826098354* and IMEI - 35* o Percentage - 20, IMSI - 31826098355* and IMEI - 31* o Percentage - 20, IMSI - 31826098356* and IMEI - 356095* o Percentage - 20, IMSI - 31826098357* and IMEI - 3560* o Percentage - 20, IMSI - 31826098358* and IMEI - 35* o Percentage - 20, IMSI - 31826098359* and IMEI - 356095*
Destination	Tool Group (PG-Sample-1).
Second-Level Map3	

Map Components	Sample Configuration
Source	Virtual port (vp1)
Source Ruleset	N/A
Applications	Applications - Flow Sampling and Load Balancing. • Flow Sampling. type - GTP • Load Balancing - Stateful, type - GTP. • Metric - Hashing, Hashing Key - IMSI.
Application Ruleset	Application Ruleset - Flow Sample GTP • Conditions - GTP o Percentage - 50, IMSI - 3182609837* and IMEI - 356095062* o Percentage - 50, IMSI - 31826098380* and IMEI - 356095062* o Percentage - 50, IMSI - 31826098381* and IMEI - 35609506* o Percentage - 50, IMSI - 31826098382* and IMEI - 3560950* o Percentage - 50, IMSI - 31826098383* and IMEI - 356095* o Percentage - 50, IMSI -31826098384* and IMEI - 35* o Percentage - 50, IMSI - 31826098385* and IMEI - 356* o Percentage - 50, IMSI - 31826098386* and IMEI - 3560* o Percentage - 50, IMSI - 31826098387* and IMEI - 35609* o Percentage - 50, IMSI - 31826098388* and IMEI - 356095* o Percentage - 50, IMSI -31826098389* and IMEI - 3560950*
Destination	Tool Group (PG-Sample-2).
Second-Level Map4	
Source	Virtual port (vp1)
Source Ruleset	N/A
Applications	Applications - Flow Sampling and Load Balancing. • Flow Sampling. type - GTP • Load Balancing - Stateful, type - GTP. • Metric - Hashing, Hashing Key - IMSI.
Application Ruleset	Application Ruleset - Flow Sample GTP • Conditions - GTP o Percentage - 10, IMSI - 31826098390* and IMEI - 35609506* o Percentage - 10, IMSI - 31826098391* and IMEI - 35609506* o Percentage - 10, IMSI - 31826098392* and IMEI - 35609506* o Percentage - 10, IMSI - 31826098393* and IMEI - 35609506* o Percentage - 10, IMSI - 31826098394* and IMEI - 35609506*
Destination	Tool Group (PG-Sample-3)
Second-Level Map5	
Source	Virtual port (vp1)
Source Ruleset	N/A
Applications	Applications - Flow Sampling and Load Balancing. • Flow Sampling, type - GTP

Map Components	Sample Configuration
	• Load Balancing - Stateful, type - GTP. • Metric - Hashing, Hashing Key - IMSI.
Application Ruleset	Application Ruleset - Flow Sample GTP • Conditions - GTP o Percentage - 10, IMSI - 31826098390* and IMEI - 35609506*
Destination	Tool Group - (PG-Sample-4)
Second-Level Map6	
Source	Virtual port (vp1)
Source Ruleset	Collector
Applications	N/A
Application Ruleset	N/A
Destination	Tool Group (PG-Sample) 1/4/x6

Configure a Port Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter PG-Sample-1 in the **Alias** field.
5. Select **Tool** option.
6. Enable **Load Balancing**.
7. Choose the **Tool Ports** you want to attach to the port group, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
8. Specify the weights for each port as follows:
    - Weight 1/1/x6 - 5
    - Weight 1/1/x7 - 10
    - Weight 1/2/x3 - 20
    - Weight 1/2/x4 - 10
9. Click **Apply** to save. The created port group appears in the port group table.
10. Follow the Steps 1 - 9 to create three tool port groups with PG-Sample-2, PG-Sample-3, and PG-Sample-4 as **Alias**.

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.

2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Create a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Create a GTP Whitelist

To create GTP Whitelist:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Forward List**. The Forward List page appears.
4. Click **New** to add a Forward List. The Add Forward List pane appears.
5. Enter a name (MyIMSI) for the forward list in **Alias** field.
6. Select **Type** as GTP.
7. Select **Bulk Entry Operation** for uploading the IMSI list.
8. Select the **Import Type** as Upload from URL.

9. Select the **Operation Type** as Append.
10. Enter the URL (Example: http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx) in the **Remote URL** field.
11. Click **Apply**. The Forward List page appears with the created Forward List list.

Associate the GTP Whitelist to GigaSMART Group

To associate the GTP Whitelist with the GigaSMART group:

1. Go to **System > GigaSMART > GigaSMART Groups**. The GigaSMART Groups page appears.
2. Select the GigaSMART group (gs-grp) you crated earlier.
3. In Actions dropdown, select **Edit**. The Edit GigaSMART Group page appears.
4. In the **GigaSMART Parameters** section, go to GTP Forward List.
5. Select the forward list you created earlier from the **GTP Forward List Alias**.
6. Click **Apply**. The GigaSMART Groups page appears with the updated GigaSMART group.

Configure Traffic Policy

Create First-Level Maps

Map 1:

To create a first-level map:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter GTP- Control.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.

d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   > **NOTE:** To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle and select **Bi-directional** option to pass the configured traffic.
   d. Select Port Destination from the **Condition** list, and then enter 2123 in the **Value** field.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears.
   b. Select the configured virtual port vp1.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the second first-level map (GTP-User) with the following source rule:
   - Source ruleset

        - Condition - Port Destination and Value -2152.
2. Deploy and verify the Traffic Policy.

Map 3:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the third first-level map ((Fragments-Not-First) with the following source rule:
    - Source ruleset
        - Condition - IP Fragmentation and Value - All Frag No First.
2. Deploy and verify the Traffic Policy

Create Second-Level Maps

To create a second-level map to forward traffic that match the IMSI whitelist to the tool port

1. Access Traffic Policy
    a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
    b. Create a new map with **Alias** as GTP-Whitelist.
2. Select Source
    a. Click **Sources**. The **Source 1** page appears with the list of available ports.
    b. Select the virtual port vp1.
    c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application
    a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
    b. Select either a standalone node or a cluster from the **Select Node** list.
    c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
    d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
    e. Select the GigaSMART group (gs-grp) you created earlier.
    f. Select Forward List and Load Balancing from the **Applications** list. By default, the applications you select are enabled.
    g. In the Forward List section, select GTP option.
    h. In the Load Balancing section, select Stateful option and type as GTP.

i. Select Hashing from the **Metric** list.

j. Select IMSI option in the **Hashing Key** field.

k. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

4. Select Application Ruleset.

   a. Expand the Application Ruleset component and then drag and drop a **Flow Sample GTP**. The **Application Ruleset** page appears.

   b. Click **Add Rules**. The **Add Ruleset** pane appears.

   c. Select GTP from the **Condition** list, and enter 20 in the **Percentage** field and select **Version** as V1.

   d. Click **Save**. The added rule appears in the **Application Ruleset** page.

   e. Click **Close**. The Flow Sample GTP block appears on the canvas with the configured application rule.

5. Select Destination.

   a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

   b. Select the tool port group (PG-Sample-1).

   c. Click **Add**. The Destination block appears on the canvas with the selected tool ports.

6. Deploy Traffic Policy.

   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

To create a second-level map to flow sample the traffic if these is no match to the IMSI whitelist and to forward the accepted packets to the load balancing port group

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (GTP-Sample-1) with the following components and configurations:
   - Source port - Virtual Port (vp1)
   - Destination port - Tool port group (PG-Sample-1)

- Application - Flow Sampling and Load Balancing
  - In the Flow Sampling section, select Flow Sampling - GTP as the **Type**.
  - In the Load Balancing section, select Stateful option and type as GTP.
  - Select the **Metric** as Hashing and IMSI from the **Hashing Key** options.
- Application ruleset - Flow Sample GTP: Select GTP from the **Condition** list and add the rules with the following configurations:
  - Percentage 75, IMSI 3182609833*, IMEI 35609506*
  - Percentage 10, IMSI 3182609834*, IMEI 3560950*
  - Percentage 20, IMSI 31826098350*, IMEI 356095*
  - Percentage 20, IMSI 31826098351*, IMEI 35609*
  - Percentage 20, IMSI 31826098352*, IMEI 3560*
  - Percentage 20, IMSI 31826098353*, IMEI 356*
  - Percentage 20, IMSI 31826098354*, IMEI 35*
  - Percentage 20, IMSI 31826098355*, IMEI 31*
  - Percentage 20, IMSI 31826098356*, IMEI 356095*
  - Percentage 20, IMSI 31826098356*, IMEI 356095*
  - Percentage 20, IMSI 31826098357*, IMEI 3560*
  - Percentage 20, IMSI 31826098358*, IMEI 35*
  - Percentage 20, IMSI 31826098359*, IMEI 356095*

2. Deploy and verify the Traffic Policy.

Map 3:

To create a second-level map to flow sample the traffic if these is no match to the IMSI whitelist and to forward the accepted packets to the load balancing port group

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (GTP-Sample-2) with the following components and configurations:
   - Source port - Virtual Port (vp1)
   - Destination port - Tool port group (PG-Sample-2)
   - Application - Flow Sampling and Load Balancing
     - In the Flow Sampling section, select Flow Sampling - GTP as the **Type**.
     - In the Load Balancing section, select Stateful option and type as GTP.
     - Select the **Metric** as Hashing and IMSI from the **Hashing Key** options.
   - Application ruleset - Flow Sample GTP: Select GTP from the **Condition** list and add the rules with the following configurations:Percentage 30, IMSI 3182609836*, IMEI 35609506*

- Percentage 5, IMSI 3182609837*, IMEI 356095062*
- Percentage 50, IMSI 31826098380*, IMEI 356095062*
- Percentage 50, IMSI 31826098381*, IMEI 35609506*
- Percentage 50, IMSI 31826098382*, IMEI 3560950*
- Percentage 50, IMSI 31826098383*, IMEI 356095*
- Percentage 50, IMSI 31826098384*, IMEI 35*
- Percentage 50, IMSI 31826098385*, IMEI 356*
- Percentage 50, IMSI 31826098386*, IMEI 3560*
- Percentage 50, IMSI 31826098387*, IMEI 35609*
- Percentage 50, IMSI 31826098388*, IMEI 356095*
- Percentage 50, IMSI 31826098389*, IMEI 3560950*

2. Deploy and verify the Traffic Policy.

Map 4:

To create a second-level map to flow sample the traffic if these is no match to the IMSI whitelist and to forward the accepted packets to the load balancing port group

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (GTP-Sample-3) with the following components and configurations:
   - Source port - Virtual Port (vp1)
   - Destination port - Tool port group (PG-Sample-3)
   - Application - Flow Sampling and Load Balancing
     - In the Flow Sampling section, select Flow Sampling - GTP as the **Type**.
     - In the Load Balancing section, select Stateful option and type as GTP.
     - Select the **Metric** as Hashing and IMSI from the **Hashing Key** options.
   - Application ruleset - Flow Sample GTP: Select GTP from the **Condition** list and add the rules with the following configurations:
     - Percentage 10, IMSI 31826098390*, IMEI 35609506*
     - Percentage 10, IMSI 31826098391*, IMEI 35609506*
     - Percentage 10, IMSI 31826098392*, IMEI 35609506*
     - Percentage 10, IMSI 31826098393*, IMEI 35609506*
     - Percentage 10, IMSI 31826098394*, IMEI 35609506*
2. Deploy and verify the Traffic Policy.

Map 5:

To create a second-level map to flow sample the traffic if these is no match to the IMSI whitelist and to forward the accepted packets to the load balancing port group

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (GTP-Sample-4) with the following components and configurations:
    - Source port - Virtual Port (vp1)
    - Destination port - Tool port group (PG-Sample-4)
    - Application - Flow Sampling and Load Balancing
        - In the Flow Sampling section, select Flow Sampling - GTP as the **Type**.
        - In the Load Balancing section, select Stateful option and type as GTP.
        - Select the **Metric** as Hashing and IMSI from the **Hashing Key** options.
    - Application ruleset - Flow Sample GTP: Select GTP from the **Condition** list and add the rules with the following configurations:
        - Percentage 10, IMSI 31826098390*, IMEI 35609506*
2. Deploy and verify the Traffic Policy.

Map 6:

To create a second-level map to collect unmatched traffic

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (GTP-Collector) with the following components:
    - Source port - vp1
    - Destination port - Tool port (1/4/x6)
    - Source Ruleset - Collector

APN for GTP Forward Listing and GTP Flow Sampling

This example specifies APN patterns for GTP whitelisting and GTP flow sampling. Traffic from the network ports are forwarded to the two first level maps (GTP-Control and GTP-User) and then to the virtual port (vp1).

In the whitelist map, if there is a match to the APN pattern and if the IMSI is present in the whitelist (IMSI), packets are forwarded to a tool port. If there is not a match to an IMSI in the

whitelist, the traffic is flow sampled based on the APN pattern in the flow sampling map. Accepted packets are forwarded to the same tool port as specified in the whitelist map. Any unmatched traffic goes to a shared collector that sends it to a different tool port.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map1	
Source	Network Port - 1/1/x1
Source Ruleset	Source Ruleset - By Rule • Condition - Port Destination, Value - 2123.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port (vp1)
First-Level Map2	
Source	Virtual Port (vp1)
Source Ruleset	Source Ruleset - By Rule • Condition - Port Destination, Value - 2152. • Condition - IP Fragmentation and Value - All Frag No First
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port (vp1)
Second-Level Map1	
Source	Virtual Port (vp1)
Source Ruleset	N/A
Applications	Applications - Forward List and Load Balancing. • Load Balancing - Stateful, type - GTP. • Metric - Hashing, Hashing Key - IMSI.
Application Ruleset	Application Ruleset - Flow Whitelist GTP . • Condition - GTP, APN - mobile.com.
Destination	Tool Group - PG-Sample.
Second-Level Map2	
Source	Virtual Port (vp1)
Source Ruleset	N/A
Applications	Applications - Flow Sampling and Load Balancing. • Flow Sampling. type - GTP • Load Balancing - Stateful, type - GTP. • Metric - Hashing, Hashing Key - IMSI.
Application Ruleset	Application Ruleset - Flow Sample GTP

Map Components	Sample Configuration
	• Conditions - GTP o Percentage - 100, APN imsi*
Destination	Tool Group - PG-Sample.
Second-Level Map3	
Source	Virtual port (vp1)
Source Ruleset	Collector
Applications	N/A
Application Ruleset	N/A
Destination	Tool Group (PG-Sample)

Configure a Port Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter PG-Sample in the **Alias** field.
5. Select **Tool** option.
6. Enable **Load Balancing**.
7. Choose the **Tool Ports** you want to attach to the port group, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
8. Specify the weights for each port as follows:
    - Weight 1/1/x6 - 5
    - Weight 1/1/x7 - 10
    - Weight 1/2/x3 - 20
    - Weight 1/2/x4 - 10
9. Click **Apply** to save. The created port group appears in the port group table.

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.

6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Create a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Create a GTP Whitelist

To create GTP Whitelist:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Forward List**. The Forward List page appears.
4. Click **New** to add a Forward List. The Add Forward List pane appears.
5. Enter a name (MyIMSI) for the forward list in **Alias** field.
6. Select **Type** as GTP.
7. Select **Bulk Entry Operation** for uploading the IMSI list.
8. Select the **Import Type** as Upload from URL.
9. Select the **Operation Type** as Append.
10. Enter the URL (Example: http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx) in the **Remote URL** field.
11. Click **Apply**. The Forward List page appears with the created Forward List list.

Associate GTP Whitelist to GigaSMART Group

To associate the GTP Whitelist with the GigaSMART group:

1. Go to **System > GigaSMART > GigaSMART Groups**. The GigaSMART Groups page appears.
2. Select the GigaSMART group (gs-grp) you crated earlier.
3. In Actions dropdown, select **Edit**. The Edit GigaSMART Group page appears.
4. In the **GigaSMART Parameters** section, go to GTP Forward List.
5. Select the forward list you created earlier from the **GTP Forward List Alias**.
6. Click **Apply**. The GigaSMART Groups page appears with the updated GigaSMART group.

Configure Traffic Policy

Create First-level Maps

Map 1:

To create a first-level map:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter GTP- Control.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Source block appears on the canvas with the selected network port.

3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle and select **Bi-directional** option to pass the configured traffic.
   d. Select Port Destination from the **Condition** list, and then enter 2123 in the **Value** field.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears.
   b. Select the configured virtual port vp1.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the second first-level map (GTP-User) with the following source rule:
   - Source ruleset
     - Condition - Port Destination and Value -2152.
     - Condition - IP Fragmentation and Value - All Frag No First.
2. Deploy and verify the Traffic Policy.

Create a Second-level Map

Map 1:

To create a second-level map to forward traffic that match the IMSI whitelist to the tool port

1. Access Traffic Policy
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Create a new map with **Alias** as GTP-Whitelist.
2. Select Source
   a. Click **Sources**. The **Source 1** page appears with the list of available ports.
   b. Select the virtual port vp1.
   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
   e. Select the GigaSMART group (GS-SUBS-INTEL) you created earlier.
   f. Select Forward List and Load Balancing from the **Applications** list. By default, the applications you select are enabled.
   g. In the Forward List section, select GTP option.
   h. In the Load Balancing section, select Stateful option and type as GTP.
   i. Select Hashing from the **Metric** list.
   j. Select IMSI option in the **Hashing Key** field.
   k. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.
4. Select Application Ruleset.
   a. Expand the Application Ruleset component and then drag and drop a **Flow Whitelist GTP**. The **Application Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Select GTP from the **Condition** list, and enter mobile.com in the **APN** field..

d. Click **Save**. The Application Ruleset page appears with the configured rule.

e. Click **Close**. The Flow Whitelist GTP block appears on the canvas with the configured application rule.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool port group (PG-Sample).

c. Click **Add**. The Destination block appears on the canvas with the selected tool ports.

6. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

To create a second-level map to flow sample the traffic if these is no match to the IMSI whitelist and to forward the accepted packets to the load balancing port group

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (GTP-Sample-1) with the following components and configurations:
   - Source port - Virtual Port (vp1)
   - Destination port - Too port group (PG-Sample)
   - Application - Flow Sampling and Load Balancing
     - In the Flow Sampling section, select Flow Sampling - GTP as the **Type**.
     - In the Load Balancing section, select Stateful option and type as GTP.
     - Select the **Metric** as Hashing and IMSI from the **Hashing Key** options.
   - Application ruleset - Flow Sample GTP: Select GTP from the **Condition** list and add the rules with the following configurations:
     - Percentage 100, APN imsi**
2. Deploy and verify the Traffic Policy.

Map 3:

To create a second-level map to collect unmatched traffic

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (GTP-Collector) with the following components:
    - Source port - vp1
    - Destination port - Tool port (PG-Sample)
    - Source Ruleset - Collector

APN for GTP Whitelisting, APN and QCI for GTP Flow Sampling

This example specifies APN patterns for GTP whitelisting and GTP flow sampling. It also specifies QCI for GTP flow sampling.

Traffic from the network ports are forwarded to the two first level maps (GTP-Control and GTP-User) and then to the virtual port (vp1). In the whitelist map, if there is a match to the APN pattern and if the IMSI is present in the whitelist (IMSI), packets are forwarded to a tool port. If there is not a match to an IMSI in the whitelist, the traffic is flow sampled based on the APN pattern and QCI value in the flow sampling map. Accepted packets are forwarded to the same tool port as specified in the whitelist map. Only 50% of traffic with QCI 5 is sent to the tool port. Any unmatched traffic goes to a shared collector that sends it to a different tool port.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map1	
Source	Network Port - 1/1/x1
Source Ruleset	By Rule • Condition - Port Destination, Value - 2123.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port (vp1)
First-Level Map2	
Source	Network Port - 1/1/x1
Source Ruleset	By Rule • Condition - Port Destination, Value - 2152. • Condition - IP Fragmentation and Value - All Frag No First
Application Ruleset	N/A

Map Components	Sample Configuration
Applications	N/A
Destination	Virtual Port (vp1)
Second-Level Map1	
Source	Virtual Port (vp1)
Source Ruleset	N/A
Applications	My GSOP - Forward List and Load Balancing. • Load Balancing: Stateful, type - GTP. • Metric - Hashing, Hashing Key - IMSI.
Application Ruleset	Flow Whitelist GTP . • Condition - GTP, APN - mobile.com.
Destination	Tool Group - PG-Sample.
Second-Level Map2	
Source	Virtual Port (vp1)
Source Ruleset	N/A
Applications	My GSOP - Flow Sampling and Load Balancing. • Flow Sampling: type - GTP • Load Balancing: Stateful, type - GTP. • Metric - Hashing, Hashing Key - IMSI.
Application Ruleset	Flow Sample GTP • Condition - GTP o Rule 1 - Percentage 50, APN imsi*, QCI 5 o Rule 2 - Percentage 100, IMSI imsi*
Destination	Tool Port Group - PG-Sample.
Second-Level Map3	
Source	Virtual port (vp1)
Source Ruleset	Collector
Applications	N/A
Application Ruleset	N/A
Destination	Tool Port Group (1/4/x19)

Configure a Port Group

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter PG-Sample in the **Alias** field.
5. Select **Tool** option.

6. Enable **Load Balancing**.
7. Choose the **Tool Ports** you want to attach to the port group, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
8. Specify the weights for each port as follows:
    - Weight 1/1/x6 - 5
    - Weight 1/1/x7 - 10
    - Weight 1/2/x3 - 20
    - Weight 1/2/x4 - 10
9. Click **Apply** to save. The created port group appears in the port group table.

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Create a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Create a GTP Whitelist

To create GTP Whitelist:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Forward List**. The Forward List page appears.
4. Click **New** to add a Forward List. The Add Forward List pane appears.
5. Enter a name (MyIMSI) for the forward list in **Alias** field.
6. Select **Type** as GTP.
7. Select **Bulk Entry Operation** for uploading the IMSI list.
8. Select the **Import Type** as Upload from URL.
9. Select the **Operation Type** as Append.
10. Enter the URL (Example: http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx) in the **Remote URL** field.
11. Click **Apply**. The Forward List page appears with the created Forward List list.

Associate GTP Whitelist to GigaSMART Group

To associate the GTP Whitelist with the GigaSMART group:

1. Go to **System > GigaSMART > GigaSMART Groups**. The GigaSMART Groups page appears.
2. Select the GigaSMART group (gs-grp) you crated earlier.
3. In Actions dropdown, select **Edit**. The Edit GigaSMART Group page appears.
4. In the **GigaSMART Parameters** section, go to GTP Forward List.
5. Select the forward list you created earlier from the **GTP Forward List Alias**.
6. Click **Apply**. The GigaSMART Groups page appears with the updated GigaSMART group.

Configure Traffic Policy

Create First-level Maps

Map 1:

To create a first-level map:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter GTP- Control.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle and select **Bi-directional** option to pass the configured traffic.
   d. Add two rules:
   e. Rule 1 - Select Port Destination from the **Condition** list, and then enter 2123 in the **Value** field.
   f. Rule 2 - Select Port Destination from the **Condition** list, and then enter 2122 in the **Value** field.
   g. Click **Save**. The added rule appears in the **Source Ruleset** page.
   h. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears.

b. Select the configured virtual port vp1.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

5. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

6. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the second first-level map (GTP-User) with the following source rule:
   - Source ruleset
     - Condition - Port Destination and Value -2152.
     - Condition - IP Fragmentation and Value - All Frag No First.
2. Deploy and verify the Traffic Policy.

Create a Second-level Map

To create a second-level map to forward traffic that match the IMSI whitelist to the tool port

1. Access Traffic Policy

a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

b. Create a new map with **Alias** as GTP-Whitelist.

2. Select Source

a. Click **Sources**. The **Source 1** page appears with the list of available ports.

b. Select the virtual port vp1.

c. Click **Add**. The Source block appears on the canvas with the selected virtual port.

3. Select Application

a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group (gs-grp) you created earlier.

f. Select Forward List and Load Balancing from the **Applications** list. By default, the applications you select are enabled.

g. In the Forward List section, select GTP option.

h. In the Load Balancing section, select Stateful option and type as GTP.

i. Select Hashing from the **Metric** list.

j. Select IMSI option in the **Hashing Key** field.

k. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

4. Select Application Ruleset.

a. Expand the Application Ruleset component and then drag and drop a **Flow Whitelist GTP**. The **Application Ruleset** page appears.

b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Select GTP from the **Condition** list, and enter mobile.com in the **APN** field..

d. Click **Save**. The Application Ruleset page appears with the configured rule.

e. Click **Close**. The Flow Whitelist GTP block appears on the canvas with the configured application rule.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool port group (PG-Sample).

c. Click **Add**. The Destination block appears on the canvas with the selected tool ports.

6. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

To create a second-level map to flow sample the traffic if these is no match to the IMSI whitelist and to forward the accepted packets to the load balancing port group

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (GTP-Sample-1) with the following components and configurations:
    - Source port - Virtual Port (vp1)
    - Destination port - Too port group (PG-Sample)
    - Application - Flow Sampling and Load Balancing
        - In the Flow Sampling section, select Flow Sampling - GTP as the **Type**.
        - In the Load Balancing section, select Stateful option and type as GTP.
        - Select the **Metric** as Hashing and IMSI from the **Hashing Key** options.
    - Application ruleset - Flow Sample GTP: Select GTP from the **Condition** list and add the rules with the following configurations:
        - Rule 1 - Percentage 50, APN imsi*, QCI 5
        - Rule 2 - Percentage 100, IMSI imsi*
2. Deploy and verify the Traffic Policy.

Map 3:

To create a second-level map to collect unmatched traffic

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (GTP-Collector) with the following components:
    - Source port - vp1
    - Destination port - Tool port (1/4/x19)
    - Source Ruleset - Collector

4G/5G Traffic Monitoring using UPN

In 4G/5G traffic monitoring using User Processing Node (UPN), the UPN processes the subscriber information that is extracted from the PFCP packets and correlates with the GTP-u traffic without the CPN. When the traffic contains the user's field information such as IMSI, IMEI, you can use the 4G/5G traffic monitoring using UPN.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map1	
Source	Network Port - 1/1/x1
Source Ruleset	Source Ruleset - By Rule • Condition - Port Destination, Value - 2123.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port (vp1)
First-Level Map2	
Source	Network Port - 1/1/x1
Source Ruleset	Source Ruleset - By Rule • Condition - Port Destination, Value - 2152.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port (vp1)
First-Level Map3	
Source	Network Port - 1/1/x1
Source Ruleset	Source Ruleset - By Rule • Condition - IP Fragmentation and Value - All Frag No First
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port (vp1)
Second-Level Map1	
Source	Virtual Port - vp1
Source Ruleset	N/A
Applications	Applications - Load Balancing. Flow Sampling Flow Sampling • Type - Flow Sampling - Flow 5G Load Balancing - Stateful. • Type - GTP, Metric - Hashing, Hashing Key - SUPI.
Application Ruleset	Application Ruleset - Flow Sample 5G. • Condition - Flow 5G, Percentage - 100.
Destination	Tool Port Group - PG-Sample.
Second-Level Map2	
Source	Virtual Port - vp1
Source Ruleset	N/A

Map Components	Sample Configuration
Applications	Applications - Forward List, type - GTP.
Application Ruleset	Application Ruleset - Flow Sample GTP • Condition - GTP, Type -IMSI, and Whitelist - MyIMSI.
Destination	Tool Port Group - PG-Sample.

Configure a Port Group

To configure a port group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter PG-Sample in the **Alias** field.
5. Select **Tool** option.
6. Enable **Load Balancing**.
7. Choose the **Tool Ports** you want to attach to the port group, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
8. Specify the weights for each port as follows:
    - Weight 1/1/x6 - 5
    - Weight 1/1/x7 - 10
    - Weight 1/2/x3 - 20
    - Weight 1/2/x4 - 10
9. Click **Apply** to save. The created port group appears in the port group table.

Configure a GigaSMART Group

To create a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**.

7. In the GigaSMART Parameters section, go to 3GPP & CUPS.
8. Select User Plane from the Node Role list.
9. Select **Standalone Mode** checkbox.
10. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Create a GTP Whitelist

To create GTP Whitelist:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Forward List**. The Forward List page appears.
4. Click **New** to add a Forward List. The Add Forward List pane appears.
5. Enter a name (MyIMSI) for the forward list in **Alias** field.
6. Select **Type** as GTP.
7. Select **Bulk Entry Operation** for uploading the IMSI list.
8. Select the **Import Type** as Upload from URL.
9. Select the **Operation Type** as Append.
10. Enter the URL (Example: http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx) in the **Remote URL** field.
11. Click **Apply**. The Forward List page appears with the created Forward List list.

Associate GTP Whitelist to GigaSMART Group

To associate the GTP Whitelist with the GigaSMART group:

1. Go to **System > GigaSMART > GigaSMART Groups**. The GigaSMART Groups page appears.
2. Select the GigaSMART group (gs-grp) you crated earlier.
3. In Actions dropdown, select **Edit**. The Edit GigaSMART Group page appears.
4. In the **GigaSMART Parameters** section, go to GTP Forward List.
5. Select the forward list you created earlier from the **GTP Forward List Alias**.
6. Click **Apply**. The GigaSMART Groups page appears with the updated GigaSMART group.

Configure a Virtual Port

To configure a virtual port:

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure Traffic Policy

Create First-level Maps

Map 1:

To create a first-level map:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter GTP- Control.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.

b. Select the network port 1/1/x1.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Source block appears on the canvas with the selected network port.

3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the configured traffic.
   d. Select Port Destination from the **Condition** list, and then enter 8805 in the **Value** field.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears.
   b. Select the configured virtual port vp1.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the second first-level map (GTP-User) with the following source rule:
   - Source ruleset
     - Condition - Port Destination and Value -2152.
2. Deploy and verify the Traffic Policy.

Map 3:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the third first-level map ((Fragments-Not-First) with the following source rule:
   - Source ruleset
     - Condition - IP Fragmentation and Value - All Frag No First.
2. Deploy and verify the Traffic Policy

Create Second-Level Maps

To create a second-level map to forward traffic that match the IMSI whitelist to the tool ports

Map 1:

1. Access Traffic Policy
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Create a new map with **Alias** as CPN-FlowSample.
2. Select Source.
   a. Click **Sources**. The **Source 1** page appears with the list of available ports.
   b. Select the virtual port vp1.
   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application Ruleset.
   a. Expand the Application Ruleset component and then drag and drop a **Flow Sample GTP**. The **Application Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Select GTP from the **Condition** list.
   d. Enter the Percentage as 100.
   e. Click **Save**. The Application Ruleset page appears with the configured rule.
   f. Click **Close**. The Whitelist GTP block appears on the canvas with the configured application rule.
4. Select Application
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.

c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group (gs-grp) you created earlier.

f. Select Load Balancing and Flow Sampling from the **Applications** list. By default, the applications you select are enabled.

g. In the Flow Sampling section, select type as Flow Sampling - Flow 5G.

h. In the Load Balancing section, select Stateful option and type as Flow 5G.

i. Select Hashing from the **Metric** list.

j. Select SUPI option in the **Hashing Key** field.

k. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool port group PG-Sample.

c. Click **Add**. The Destination block appears on the canvas with the selected tool ports.

6. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (GTP-Whitelist) with the following components and configurations:
   - Source port - Virtual Port (vp1)
   - Destination port - 1/2/x24
   - Application ruleset - Flow Whitelist GTP
     - Condition - GTP, Type -IMSI, and Whitelist - MyIMSI
   - Application - Forward List

- Type - GTP

2. Deploy and verify the Traffic Policy.

GigaSMART 5G CUPS

Required License: 5G Correlation and CUPS
Supported Devices : GigaVUE-HC3 Gen 2 and GigaVUE-HC3 Gen 3

The 5G Correlation feature correlates the 5G Control and User packets to deliver it to different tool ports based on the filtering policies configured. The control and user packets are processed in Control Plane Network (CPN) and User Plane Network (UPN) located in same or different locations, and then the packets are sent to the tools.

NOTE: It is recommended that you unconfigure and reload the GigaSMART card when configuring a particular GS group with 5G and switching back to 4G or vice versa.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map1	
Source	Network Port - 1/1/x1
Source Ruleset	Source Ruleset - By Rule • Condition - Port Destination, Value - 8805.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port (vp1)
First-Level Map2	
Source	Virtual Port (vp1)
Source Ruleset	Source Ruleset - By Rule • Condition - Port Destination, Value - 2123.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port (vp1)
First-Level Map3	
Source	Virtual Port (vp1)
Source Ruleset	Source Ruleset - By Rule • Condition - IP Fragmentation and Value - All Frag No First

Map Components	Sample Configuration
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port (vp1)
First-Level Map4	
Source	Source - IP Interface port of UPN (1/2/x1).
Source Ruleset	Source Ruleset - By Rule • Condition -Port Destination and Value - 50002.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port (vp1)
First-Level Map5	
Source	Virtual Port (vp1)
Source Ruleset	Source Ruleset - By Rule • Condition - IP Fragmentation and Value - All Frag No First
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port (vp1)
First-Level Map6	
Source	Virtual Port (vp1)
Source Ruleset	Source Ruleset - By Rule • Condition - IP Fragmentation and Value - All Frag No First
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port (vp1)
Second-Level Map1	
Source	Network Port (1/1/x1)
Source Ruleset	N/A
Applications	Applications - Load Balancing. • Load Balancing - Stateful, type - GTP. • Metric - Hashing, Hashing Key - SUPI.
Application Ruleset	Application Ruleset - Flow Sample 5G. • Condition - Flow 5G, Percentage - 100.
Destination	Tool Group (PG-Sample).
Second-Level Map2	
Source	Virtual Port (vp1)
Source Ruleset	N/A

Map Components	Sample Configuration
Applications	Applications - Forward List, type - GTP.
Application Ruleset	Application Ruleset - Flow Sample GTP • Condition - GTP, Type -IMSI, and Whitelist - MyIMSI.
Destination	Tool Group (PG-Sample).

Configure a Port Group

To configure a port group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter PG-Sample in the **Alias** field.
5. Select **Tool** option.
6. Enable **Load Balancing**.
7. Choose the **Tool Ports** you want to attach to the port group, For example 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4.
8. Specify the weights for each port as follows:
    - Weight 1/1/x6 - 5
    - Weight 1/1/x7 - 10
    - Weight 1/2/x3 - 20
    - Weight 1/2/x4 - 10
9. Click **Apply** to save. The created port group appears in the port group table.

Configure a GigaSMART Group

To create a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp-upn) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**.

7. In the GigaSMART Parameters section, go to 3GPP & CUPS.
8. Select User Plane from the **Node Role** list.
9. Select **Standalone Mode** checkbox.
10. Configure GPFCP profile, refer to Command Line Interface - gpfcp profile.
11. Click **Apply**. The GigaSMART Groups page appears with the group you created.
12. Create another GigaSMART group with Alias as (gs-grp-cpn) with **Node Role** as Control Plane 5G.

Configure an IP Interface

To configure an IP Interface:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the IP interface. The Overview page of the selected node appears.
3. Go to **System > Ports > IP Interfaces**. The IP Interfaces page appears.
4. Click **New**. The New IP Interface pane appears.
5. Configure the following fields as follows:
    - **Alias** - CPN_IP
    - **Port** - 1/2/x1
    - **Type**: IPv4
    - **IP Address** - 10.202.24.2
    - **IP Mask** - 255.255.255.0
    - **Gateway** - 10.115.9.1
    - **MTU** - 1500
    - **GigaSMART Group** - gs-grp
6. Click **Apply**. The created IP interface appears in the IP Interfaces page.
7. Follow the steps 1 - 6 and create another IP interface with **Alias** as UPN_IP for UPN communication.

Configure TCP/IP Host

To configure a TCP/IP Host:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the host. The Overview page of the selected node appears.

3. Go to **System > GigaSMART > TCP/IP Host**. The New Exporter page appears.
4. Configure the following fields as follows:
    - **Alias** - Enter a name (cpn-ex) for the exporter.
    - **Application Type** - Mobility CUPS.
    - **GigaSMART Group** - gs-grp-cpn
    - **IP Interface** - CPN_IP
    - **L4 Port** - 50001
    - **L4 Protocol** - TCP
5. Click **Apply**. The created exporter appears in the Exporters page.
6. Follow the steps 1 - 5 and create another exporter for UPN with **Alias** as upn-ex, GigaSMART group as gs-grp-upn, IP Interface as UPN_IP and **L4 Port** as 50002.

Create a GTP Whitelist

To create GTP Whitelist:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Forward List**. The Forward List page appears.
4. Click **New** to add a Forward List. The Add Forward List pane appears.
5. Enter a name (MyIMSI) for the forward list in **Alias** field.
6. Select **Type** as GTP.
7. Select **Bulk Entry Operation** for uploading the IMSI list.
8. Select the **Import Type** as Upload from URL.
9. Select the **Operation Type** as Append.
10. Enter the URL (Example: http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx) in the **Remote URL** field.
11. Click **Apply**. The Forward List page appears with the created Forward List list.

Associate GTP Whitelist to GigaSMART Group

To associate the GTP Whitelist with the GigaSMART group:

1. Go to **System > GigaSMART > GigaSMART Groups**. The GigaSMART Groups page appears.
2. Select the GigaSMART group (gs-grp-upn) you crated earlier.

3. In Actions dropdown, select **Edit**. The Edit GigaSMART Group page appears.
4. In the **GigaSMART Parameters** section, go to GTP Forward List.
5. Select the forward list you created earlier from the **GTP Forward List Alias**.
6. Click **Apply**. The GigaSMART Groups page appears with the updated GigaSMART group.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select the GigaSMART group gs-grp-upn from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.
8. Follow the steps 1 - 7 and create another virtual port (vp2) assigned to GigaSMART group (gs-grp-cpn).

Configure Traffic Policy

Create First-level Maps

Map 1:

To create a first-level map:

1. Access Traffic Policy.
    a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
    b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
    c. In the **New Map** dialog, do the following:
        - **Alias** - Enter GTP- Control.
        - (Optional) **Description** - Enter the description for the Traffic Policy.

- (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
- (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.

d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x1.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle to pass the configured traffic.
   d. Select Port Destination from the **Condition** list, and then enter 8805 in the **Value** field.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears.
   b. Select the configured virtual port vp1.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the first-level map (GTP-User) with the following source rule:
    - Source ruleset
        - Condition - Port Destination and Value -2152.
    - Destination - Virtual Port (vp1)
2. Deploy and verify the Traffic Policy.

Map 3:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the first-level map ((Fragments-Not-First) with the following source rule:
    - Source ruleset
        - Condition - IP Fragmentation and Value - All Frag No First.
    - Destination - Virtual Port (vp1)
2. Deploy and verify the Traffic Policy

Map 4:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the first-level map (UPN-GPFCP) with the following source rule:
2. Source - IP Interface port of UPN (1/2/x1)
    - Source ruleset
        - Condition - Port Destination and Value - 50002
    - Destination - Virtual Port (vp1)
3. Deploy and verify the Traffic Policy

Map 5:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the first-level map (CPN-GPFCP) with the following source rule:
2. Source - IP Interface port of CPN (1/2/x1)
    - Source ruleset
        - Condition - Port Destination and Value - 50001
    - Destination - Virtual Port (vp2)
3. Deploy and verify the Traffic Policy

Map 6:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the first-level map (CPN-HTTP) with the following source rule:
2. Source - Network port of CPN (1/1/x1)
    - Source ruleset
        - Condition - Port Destination and Value - 80
    - Destination - Virtual Port (vp2)
3. Deploy and verify the Traffic Policy

Create Second-Level Maps

To create a second-level map to forward traffic that match the IMSI whitelist to the tool ports

Map 1:

1. Access Traffic Policy
    a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
    b. Create a new map with **Alias** as CPN-FlowSample.
2. Select Source
    a. Click **Sources**. The **Source 1** page appears with the list of available ports.
    b. Select the virtual port vp1.
    c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application
    a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
    b. Select either a standalone node or a cluster from the **Select Node** list.
    c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
    d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
    e. Select the GigaSMART group (gs-grp-upn) you created earlier.
    f. Select Load Balancing from the **Applications** list. By default, the applications you select are enabled.
    g. In the Load Balancing section, select Stateful option and type as Flow 5G.
    h. Select Hashing from the **Metric** list.

i. Select SUPI option in the **Hashing Key** field.

j. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

4. Select Application Ruleset.
   a. Expand the Application Ruleset component and then drag and drop a **Flow Sample GTP**. The **Application Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Select GTP from the **Condition** list.
   d. Enter the Percentage as 100.
   e. Enter 123.456 in **PLMN ID** field.
   f. Enter 5 in **5QI** field.
   g. Click **Save**. The Application Ruleset page appears with the configured rule.
   h. Click **Close**. The Flow Sample GTP block appears on the canvas with the configured application rule.
5. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.
   b. Select the tool port group PG-Sample.
   c. Click **Add**. The Destination block appears on the canvas with the selected tool ports.
6. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
7. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (GTP-Whitelist) with the following components and configurations:
   - Source port - Virtual Port (vp1)
   - Destination port - 1/2/x24
   - Application ruleset - Flow Whitelist GTP

        - Condition - GTP, Type -IMSI, and Whitelist - MyIMSI
    - Application - Forward List
        - Type - GTP
2. Deploy and verify the Traffic Policy.

Map 3:

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (FlowSample-5G) with the following components and configurations:
    - Source port - Virtual Port (vp2)
    - Destination port - 1/2/x22
    - Application ruleset - Flow Sample 5G
        - Condition - Flow 5G and Percentage - 100.
    - Application - Flow Sampling
        - Type - Flow Sampling - Flow 5G
2. Deploy and verify the Traffic Policy.

GTP Scaling

Required License: GTP Filtering & Correlation

A GigaSMART group (gsgroup) associated with GTP applications can have multiple GigaSMART engine port members. You can combine up to eight engine ports to form a single engine group. The engine group provides higher capacity to GTP applications by load balancing GTP user-data plane (GTP-u) traffic among the members of the group. Grouping multiple GigaSMART engine ports increases the effective throughput for GTP applications.

NOTE: A GigaSMART Engine Group supports 12 million GTP subscriber sessions for GigaVUE HC3 nodes.

GTP engine grouping is supported on GigaVUE-HC3 nodes and can be used by the following GTP applications: GTP flow filtering, GTP flow sampling, and GTP whitelisting.

The following sections explain the Traffic Policy configurations of GTP Engine Grouping for different real-time example.

- GTP Engine Grouping Configuration with 2 Engine Ports
- GTP Engine Grouping Configuration with 4 Engine Ports

GTP Engine Grouping Configuration with 2 Engine Ports

This topic explains the workflow required to configure a Traffic Policy for GTP engine group consisting of two engine ports on a GigaVUE-HC1 node and with GigaSMART operation as GTP flow filtering.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map1	
Source	Network Port - 1/1/x1
Source Ruleset	Condition - Port Destination, Value - 2123. • Enable Control Traffic
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port(vp1)
First-Level Map2	
Source	1/1/x1
Source Ruleset	Condition - Port Destination, Value - 2152.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port - vp1
Second-Level Map1	
Source	Virtual Port(vp1)
Source Ruleset	N/A
Applications	Applications - Flow Filtering.
Application Ruleset	Application Ruleset - Flow Filter • Condition - IMEI. • IMEI - * and V1 - Version.
Destination	
Second-Level Map2	
Source	Virtual port (vp1)
Source Ruleset	Collector
Applications	N/A
Application Ruleset	N/A
Destination	Tool Group (PG-Sample) 1/3/x14

Create a GigaSMART Group

To create a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select two GigaSMART engine ports from the **Port List**.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Create a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Create First-level Maps

Map 1:

To create a first-level map:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.

c. In the **New Map** dialog, do the following:

- **Alias** - Enter GTP- Control.
- (Optional) **Description** - Enter the description for the Traffic Policy.
- (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
- (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.

d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source.

a. Click **Sources**. The **Source 1** page appears.

b. Select the network port 1/1/x3.

> **NOTE:** To configure or modify the port type, use the Quick Port Editor.

c. Click **Add**. The Source block appears on the canvas with the selected network port.

3. Select Source Ruleset.

a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.

b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle and select **Bi-directional** option to pass the configured traffic.

d. Select Port Destination from the **Condition** list, and then enter 2123 in the **Value** field.

e. Click **Save**. The added rule appears in the **Source Ruleset** page.

f. Click the **Settings** icon. The Map Configuration pane appears.

g. Enter or select the required options.

h. Click **Save**. The Map Configuration pane closes.

i. Click **Close**. The By Rule block appears on the canvas with the configured rule.

4. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears.

b. Select the configured virtual port vp1.

c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.

d. After configuring the destination, click the By Rule block. The **Source Ruleset** page appears.

a. Click the **Settings** icon. The Map Configuration pane appears.

b. Select the Control Traffic checkbox.

c. Click **Save**. The Map Configuration pane closes.

5. Deploy Traffic Policy.

   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

6. Verify deployment.

   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the second first-level map (GTP-User) with the following source rule:
   - Source ruleset
     - Condition - Port Destination and Value -2152.
2. Deploy and verify the Traffic Policy.

Create a Second-level Map

To create a second-level map to forward traffic that match the IMSI whitelist to the tool port

Map 1:

1. Access Traffic Policy

   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

   b. Create a new map with **Alias** as GTP-Whitelist.

2. Select Source

   a. Click **Sources**. The **Source 1** page appears with the list of available ports.

   b. Select the virtual port vp1.

   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.

3. Select Application

   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.

b. Select either a standalone node or a cluster from the **Select Node** list.

c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.

d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.

e. Select the GigaSMART group (gs-grp) you created earlier.

f. Select Flow Filtering from the **Applications** list. By default, the application you select is enabled.

g. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.

4. Select Application Ruleset

a. Expand the Application Ruleset component and then drag and drop a Flow Filter. The **Application Ruleset** page appears.

b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Select GTP IMEI from the **Condition** list.

d. Enter * in the **IMEI** field and select V1 as **Version**.

e. Click **Save**. The Application Ruleset page appears with the configured rule.

f. Click **Close**. The Flow Filter block appears on the canvas with the configured application rule.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the tool port group PG-Sample.

c. Click **Add**. The Destination block appears on the canvas with the selected tool ports.

6. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (scoll) with the following components:
    - Source port - vp1
    - Source Ruleset - Collector
    - Destination port - 1/3/x14

GTP Engine Grouping Configuration with 4 Engine Ports

This topic explains the workflow required to configure a Traffic Policy for GTP engine group consisting of four engine ports on on two GigaSMART line cards on the same GigaVUE-HC3 node. The GigaSMART line cards are in slots 1 and 3.

The GigaVUE-HC3 node is the cluster leader of a two-node out-of-band cluster. A GigaVUE -HC1 is the standby node in the cluster. This example includes GigaSMART operations for GTP flow filtering with load balancing, GTP flow sampling with load balancing, and GTP whitelisting. The forward list must be associated with the GigaSMART group on the leader, the GigaVUE-HC3.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
First-Level Map	
Source	Network Port - 1/1/x1
Source Ruleset	Source Ruleset - By Rule • Condition - Port Destination, Value - 2123.
Application Ruleset	N/A
Applications	N/A
Destination	Virtual Port (vp1)
First-Level Map2	
Source	Virtual Port (vp1)
Source Ruleset	Source Ruleset - By Rule • Condition - Port Destination, Value - 2152. • Condition - IP Fragmentation and Value - All Frag No First
Application Ruleset	N/A
Applications	N/A

Map Components	Sample Configuration
Destination	Virtual Port (vp1)
Second-Level Map1	
Source	Virtual Port (vp1)
Source Ruleset	N/A
Applications	Applications - Forward List, type - GTP.
Application Ruleset	Condition - IP Version and Version - Version 4.
Destination	Port group (gs1-hc3)
Second-Level Map2	
Source	Virtual Port (vp1)
Source Ruleset	N/A
Applications	Applications - Flow Sampling. • Type - Flow Sampling - GTP.
Application Ruleset	Applications Ruleset - Flow Sample GTP Condition - GTP • Percentage - 60, IMSI - 5*, and Version - V2.
Destination	Port group (hc3-pg2)
Second-Level Map3	
Source	Virtual Port (vp1)
Source Ruleset	N/A
Applications	Applications - Flow Sampling. • Type - Flow Sampling - GTP.
Application Ruleset	Applications Ruleset - Flow Sample GTP Condition - GTP • Percentage - 60, IMSI - 5*, and Version - V1.
Destination	Port group (hc3-pg1)
Second-Level Map4	
Source	Virtual Port (vp1)
Source Ruleset	N/A
Applications	Applications - Flow Filtering. • Type - GTP.
Application Ruleset	Applications Ruleset - Flow Filter • Condition - GTP IMSI • IMSI - *, and Version - Any
Destination	Port group (hc3-pg3)
Second-Level Map5	
Source	Virtual port (vp1)
Source Ruleset	Collector

Map Components	Sample Configuration
Applications	N/A
Application Ruleset	N/A
Destination	Tool Group (PG-Sample)

Configure GigaStream

To configure GigaStream:

In GigaVUE-HC3 device:

1. In GigaVUE-HC3 device, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.Go
3. Go to **Ports > Port Groups**> **GigaStream**.
4. Click **New**. The New GigaStream pane appears.
5. Enter a name (gs1-hc3) in the **Alias** field.
6. Select Tool GigaStream as the **Type**.
7. Select five tool ports 1/1/x20 - 1/1/x24 from the **Ports** list.
8. Select the **Weighting** as required.
9. Click Advanced Hash Settings.
10. Click **Apply**. The GigaStream page appears with the created GigaStream.
11. Follow the steps 1 - 9 to create another GigaStream (gs2-hc3) with two tool ports 1/1/x26 and 1/1/x27.

In GigaVUE-HC1 device:

Follow the same steps as the previous section and configure GigaStream with the following values:

1. GigaStream 1
   - Alias - gs3-hc1
   - Tool ports - 1/1/x33 - 1/1/x37
2. GigaStream 2
   - Alias - gs4-hc1
   - Tool ports - 1/1/x41 and 1/1/x42

Configure a Stack Link

To configure a stack link:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. Go to **Ports > Stack Links**. The Stack Links page appears.
4. Click **New**. The New Stack Link pane appears.
5. Enter a name in the **Alias** field.
6. Select Stack GigaStream option in the **Configuration**.
7. Select gs2-hc3 for the GigaStream in Device 1 and gs4-hc1 for the GigaStream in Device 2, that you created earlier.
8. Click Apply. The Stack Links page appears with the created stack link.

Configure Port Groups

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click **Nodes** then, go to **Ports > Port Groups** .
3. Click **New** to create a new Port Group. The New Port Group page appears.
4. Enter a name (hc3-pg1) in the **Alias** field.
5. Select **Tool** option.
6. Enable **Load Balancing**.
7. Select four tool ports that you want to attach to the port group, For example 1/1/x9 - 1/1/x12.
8. Specify the weights for each port.
9. Click **Apply** to save. The created port group appears in the Port Groups page.
10. Follow the steps 1 - 9 and create two more port groups each having four ports with **Alias** as hc3-pg2 and hc3-pg3.

Create a GigaSMART Group

To create a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.

3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select four GigaSMART engine ports from the **Port List**. For example, 3/1/e1, 23/1/e2, 23/3/e1, and 23/3/e2.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Create a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Create a GTP Whitelist

To create GTP Whitelist:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Forward List**. The Forward List page appears.
4. Click **New** to add a Forward List. The Add Forward List pane appears.
5. Enter a name (MyIMSI) for the forward list in **Alias** field.
6. Select **Type** as GTP.
7. Select **Bulk Entry Operation** for uploading the IMSI list.
8. Select the **Import Type** as Upload from URL.
9. Select the **Operation Type** as Append.

10. Enter the URL (Example: http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx) in the **Remote URL** field.
11. Click **Apply**. The Forward List page appears with the created Forward List list.

Associate the GTP Whitelist to GigaSMART Group

To associate the GTP Whitelist with the GigaSMART group:

1. Go to **System > GigaSMART > GigaSMART Groups**. The GigaSMART Groups page appears.
2. Select the GigaSMART group (gs-grp) you crated earlier.
3. In **Actions** dropdown, select **Edit**. The Edit GigaSMART Group page appears.
4. In the **GigaSMART Parameters** section, go to GTP Forward List.
5. Select the forward list you created earlier from the **GTP Forward List Alias**.
6. Click **Apply**. The GigaSMART Groups page appears with the updated GigaSMART group.

Configure Traffic Policy

Create First-level Maps

Map 1:

To create a first-level map:

1. Access Traffic Policy.
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
   c. In the **New Map** dialog, do the following:
      - **Alias** - Enter GTP- Control.
      - (Optional) **Description** - Enter the description for the Traffic Policy.
      - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
      - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
   d. Click **Save**. The Traffic Policy canvas appears.

2. Select Source.
   a. Click **Sources**. The **Source 1** page appears.
   b. Select the network port 1/1/x3.

   NOTE: To configure or modify the port type, use the Quick Port Editor.

   c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
   a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.
   c. Turn on **Pass** toggle and select **Bi-directional** option to pass the configured traffic.
   d. Select Port Destination from the **Condition** list, and then enter 2123 in the **Value** field.
   e. Click **Save**. The added rule appears in the **Source Ruleset** page.
   f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.
   a. Click **Destinations**. The **Destination 1** page appears.
   b. Select the configured virtual port vp1.
   c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
   a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
   a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
   b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

Follow the same steps in the previous section to create another first-level map with the following changes:

1. Create the second first-level map (GTP-User) with the following source rule:
   - Source ruleset
     - Rule 1: Condition - Port Destination and Value -2152.

- Rule 2: Condition - IP Fragmentation and Value - All Frag No First.

2. Deploy and verify the Traffic Policy.

Create a Second-Level Map

Map 1:

To create a second level map for GTP whitelisting, the forward list map, that takes traffic from the virtual port, applies the forward listing GigaSMART operation, and sends traffic to the remote GigaVUE-HC1 node through a GigaStream

1. Access Traffic Policy
   a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
   b. Create a new map with **Alias** as GTP-Whitelist.
2. Select Source
   a. Click **Sources**. The **Source 1** page appears with the list of available ports.
   b. Select the virtual port vp1.
   c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application
   a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
   b. Select either a standalone node or a cluster from the **Select Node** list.
   c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
   d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
   e. Select the GigaSMART group (gs-grp) you created earlier.
   f. Select Forward List from the **Applications** list. By default, the application you select is enabled.
   g. In the Forward List section, select GTP as the type.
   h. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.
4. Select Application Ruleset.
   a. Expand the Application Ruleset component and then drag and drop a **By Rule**. The **Application Ruleset** page appears.
   b. Click **Add Rules**. The **Add Ruleset** pane appears.

c. Turn on **Pass** toggle to pass the configured traffic.

d. Select IP Version from the **Condition** list and select Version 4 from the **Version** list.

e. Click **Save**. The **Application Ruleset** page appears with the configured rule.

f. Click **Close**. The By Rule block appears on the canvas with the configured application rule.

5. Select Destination.

a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.

b. Select the GigaStream gs1-hc3 that you created earlier.

c. Click **Add**. The Destination block appears on the canvas with the selected GigaStream.

6. Deploy Traffic Policy.

a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.

7. Verify deployment.

a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.

b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Map 2:

To configure a second level map for GTP flow sampling that filters for version 2 and forwards traffic from the virtual port and applies the flow sampling GigaSMART operation. Traffic flow is sampled based on the flow sampling rule in this map and accepted packets are forwarded to load balancing port group.

Follow the same steps in the previous section to create another second-level map (GTP-FlowSampling) with the following changes:

1. Source - Virtual Port (vp1)
2. Application Ruleset - Flow Sample GTP
    - Condition - GTP
    - Percentage - 60, IMSI - 5*, and Version - V2
3. Application - Flow Sampling
    - Type - Flow Sampling - GTP
4. Destination - Port group (hc3-pg2)

Map 3:

To configure a second level map for GTP flow sampling that filters for version 1 and forwards traffic from the virtual port and applies the flow sampling GigaSMART operation. Traffic flow is sampled based on the flow sampling rule in this map and accepted packets are forwarded to load balancing port group.

Follow the same steps in the previous section to create another second-level map (GTP-FlowSampling) with the following changes:

1. Source - Virtual Port (vp1)
2. Application Ruleset - Flow Sample GTP
    - Condition - GTP
    - Percentage - 60, IMSI - 5*, and Version - V1
3. Application - Flow Sampling
    - Type - Flow Sampling - GTP
4. Destination - Port group (hc3-pg1)

Map 4:

To configure a second level map for GTP flow filtering that takes traffic from the virtual port, applies the flow filtering GigaSMART operation, matches IMSIs specified by the flow rule, and sends matching traffic to load balancing port group.

Follow the same steps in the previous section to create another second-level map (GTP-FlowFiltering) with the following changes:

1. Source - Virtual Port (vp1)
2. Application Ruleset - Flow Filter
    - Condition - GTP IMSI
    - IMSI - *, and Version - Any
3. Application - Flow Filtering
    - Type - GTP
4. Destination - Port group (hc3-pg3)

Map 5:

Follow the same steps in the previous section to create another second-level map with the following changes:

1. Create the second-level map (scoll) with the following components:
    - Source port - vp1
    - Application Ruleset - Collector
    - Destination port - 1/3/x14

Migrate Maps to Traffic Policy

Traffic Policies provide a unified way to configure all physical out-of-band maps (Flow Map, Fabric Map, App Intelligence solutions) through a single User Interface.

You can migrate existing Flow Maps, Fabric Maps, and App Intelligence solutions into Traffic Policies without recreating them manually. After migration, a given configuration exists as a Traffic Policy map. The underlying device configuration is not modified during migration; traffic forwarding continues to work as it is until you change the configuration using Traffic Policies.

Dry Run before Migration

Dry Run lets you validate whether a selected Traffic Intelligence , Application Intelligence or Subscriber Intelligence solutions can be migrated successfully

> **NOTE:** Running a *Dry Run* in Traffic Policy does not migrate or change your existing configuration. It only reports, for each map, whether migration would succeed or fail and the reason for failure.

When you start a dry run for Traffic Intelligence , Application Intelligence or Subscriber Intelligence pages, you can choose:

- Dry Run All- simulate migration and record results for all records.
- Dry Run Selected – simulate migration and record results for selected records.

After Dry Run completes, go to **Traffic** **>Physical>Traffic Policy > View Migration Results** and review the run:

- Ensure the Status is Completed.
- Check that the Failed and Skipped counts are 0, or that any skipped entries are acceptable.
- If there are failures, use the failure details shown in the Migration Result to correct configuration or data issues, or contact Gigamon Support with the Migration Result information.

We recommend running a Dry Run before performing an actual migration. This helps detect configuration issues in advance and increases confidence that the migration will succeed. This approach helps you avoid migration failures and unexpected errors during the actual migration.

NOTE: After migration, the selected maps are represented as Traffic Policies and can no longer be viewed or edited from the legacy Flow Maps, Fabric Maps, or Application Intelligence pages. If multiple maps share the same alias across clusters (Flow Map vs Flow Map) or between Flow Maps and Fabric Maps, the migrated Traffic Policy alias is adjusted by prefixing *ClusterId_CL_* to make it unique. If there is a collision between a migrated Flow Map, Fabric Map, or Application Intelligence item and an existing Traffic Policy alias, the migrated Traffic Policy alias is adjusted by appending *_TF_xxx*. If the final alias would exceed the supported length, that map is not migrated and appears under Failed in *View Migration Results*, with the reason shown in the details.

Migrating Traffic Intelligence and Subscriber Intelligence Solutions

Traffic Policy migration affects solution families differently:

- Traffic Intelligence solutions(for example, GigaSMART APF, de-duplication, masking, slicing, tunneling, NetFlow) are built on classic Flow Maps and Fabric Maps. You migrate these by migrating the underlying Flow Maps and Fabric Maps described in this section.
- Subscriber Intelligence solutions (for example, GTP and CUPS correlation, FlowVUE, SIP/RTP correlation, 4G/5G correlation) ,if you have Fabric Maps, that use Subscriber Intelligence applications, you can migrate those Fabric Maps using the standard Fabric Map migration flow.

When you migrate Flow Maps or Fabric Maps to Traffic Policy:

- The underlying device configuration is preserved; only the control plane representation moves to Traffic Policy. Existing Traffic Intelligence behavior continues unchanged.
- After migration, configuration and troubleshooting must be done from the Traffic Policy page; the entries in the legacy Flow Maps and Fabric Maps pages are no longer intended for user edits
- The corresponding entries on the legacy Flow Maps and Fabric Maps pages are treated as internal artifacts and are no longer intended for user edits.

Migrating Flow Maps

You can migrate existing Flow Map into Traffic Policies, which preserves their behavior while moving configuration and troubleshooting to the Traffic Policy page.

NOTE: Before migrating Flow Maps on a device, disconnect the corresponding node or cluster in GigaVUE-FM (**Inventory > Nodes > Actions > Disconnect Node**). GigaVUE-FM still completes any in-progress sync, but new configuration changes are not pushed while the node is disconnected.

1. Go to  **> Physical > flow maps**.
2. Click on Migrate button and either:
   - Migrate All- migrates all records
   - Migrate Selected-migrates only the selected records.
3. Click on OK to migrate your flow maps.
4. To view the progress of the migration use the *View Migration Results* page:
   - Success count matches the number of maps you intended to migrate.
   - Failed count is 0.
   - Skipped entries are acceptable if they correspond to auto-created / internal maps

After migration, reconnect the cluster/standalone so normal configuration sync resumes. The Flow Maps that were migrated will have their configuration represented in Traffic Policies.

NOTE: After you migrate a Flow Map to Traffic Policy, that Flow Map no longer appears in the legacy Flow Maps list and is treated as an internal artifact. You cannot view or edit it from the legacy Flow Maps page; manage the configuration only from the Traffic Policy page

Migrating Fabric Maps

You can migrate existing Fabric Maps into Traffic Policies, which preserves their behavior while moving configuration and troubleshooting to the Traffic Policy page.

5. Go to  **> Physical > Orchestrated Flows > Fabric Maps**.
6. Click on Migrate button and either:
   - Migrate All- migrates all records
   - Migrate Selected-migrates only the selected records.
7. Click on OK to migrate your Fabric Maps.
8. To view the progress of the migration use the *View Migration Results* page:
   - Success count matches the number of maps you intended to migrate.
   - Failed count is 0.
   - Skipped entries are acceptable if they correspond to auto-created / internal maps.

After the Migration your Fabric Map configuration can now only be managed via the Traffic Policy page. The migrated Fabric Map will not be editable from the Fabric Map page.

NOTE: After you migrate a Fabric Map to Traffic Policy, that Fabric Map no longer appears in the legacy Fabric Maps list and is treated as an internal artifact. You cannot view or edit it from the legacy Fabric Maps page; manage the configuration only from the Traffic Policy page

Migrating Application Intelligence Solutions

Migration converts existing Application Intelligence solutions (on the old App Intel page) into unified Traffic Policies, preserving behavior but moving management to the Traffic Policy page.

To Migrate Application Intelligence solutions do the following:

1. Go to **Traffic**  **> Solutions > Application Intelligence.**
2. Select the  Actions option of the session that you wish to migrate.
3. Click on Dry Run if you wish to see how the solution would be migrated before the actual migration.
4. Click on Migrate option to migrate your session to Traffic Policy
5. To view the progress of the migration use the *View Migration Results* page:
   - Success count matches the number of maps you intended to migrate.
   - Failed count is 0.
   - Skipped entries are acceptable if they correspond to auto-created / internal maps
   - Check that all expected sources, App Intel applications, and destinations are present.

NOTE: After you migrate a Fabric Map to Traffic Policy, that Fabric Map no longer appears in the legacy Fabric Maps list and is treated as an internal artifact. You cannot view or edit it from the legacy Fabric Maps page; manage the configuration only from the Traffic Policy page

Force rollback is not supported for Application Intelligence solutions. If you want to roll back the App Intel solution, you must contact customer support.

View Migration Results

Use the *View Migration Results* page to check the outcome of Dry Run and Migration actions.

1. In GigaVUE-FM, go to **Traffic > Traffic Policy**.

2. Click View Migration Results.

3. In the list, locate the run by its Start Time and Map Type (Flow Map, Fabric Map, or Application Intelligence solution).

4. Verify that:

- Success count matches the number of maps you intended to migrate.
- Failed count is 0.
- Skipped entries are acceptable if they correspond to auto-created / internal maps.

You can use the filters at the top of the page to narrow the view by:

- Map type – Show only Flow Maps, Fabric Maps, or Application Intelligence solutions.
- Status – Show only runs that are In Progress, Completed, Failed, or Skipped.

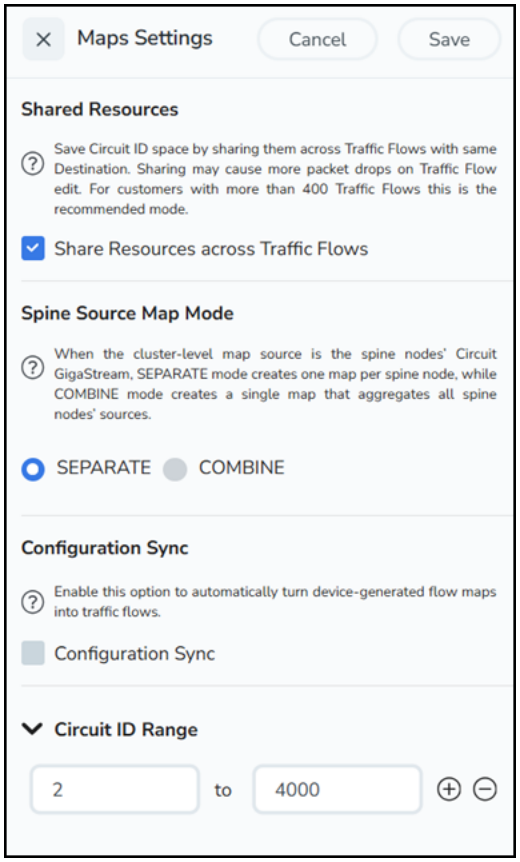
For each successful migration, confirm that a corresponding Traffic Policy exists and that its sources, applications, and destinations match the original configuration.

Migration converts the configuration stored in GigaVUE-FM for Flow Maps, Fabric Maps, and Application Intelligence solutions into Traffic Policies. Migrated configurations must be managed only from the Traffic Policies page. The corresponding entries in the legacy Flow Map / Fabric Map / Application Intelligence pages are treated as internal artifacts and are no longer intended for user edits.

Configure Map Settings

You can configure the map settings to control how resources, spine source maps, configuration synchronization, and circuit ID ranges are handled for traffic flows.

The following image shows the Maps Settings dialog in Traffic Policy.



Share Circuit ID Resource

You can share the circuit ID resources for the flow maps created for the same destination. To do this:

1. Select the **Settings** icon on the top right corner of the Traffic Policy landing page to open the **Map Settings** dialog.
2. Select **Share Resources Across Traffic Flows** option to share the circuit ID resources.
3. Enter the **Circuit ID Range** (Min and Max values). Use the +/- icon to add or remove the circuit ID ranges.
4. Click **Save** to save the settings.

When you change the option from shared to non-shared, a note message appears. Enter Yes in the confirmation message and click **OK** to change from Shared to Non-Shared. You can use this option only if you who have Read and Write privileges in both the Traffic Control Management and GigaVUE-FM Security Management.

Spine Source Map Mode

You can configure this global setting in FM that controls how cluster-local maps are created in a leaf-and-spine cluster when the source is at the spine (i.e., a decap circuit GigaStream in a leaf-spine cluster). This setting is applicable for creation of the first map that uses the spine mode. Refer to Supported Topologies topic for more information on topology of Spine Source Map mode.

When a leaf-and-spine cluster acts as a middle or destination cluster in a fabric map, its cluster-level map source can be one or more circuit GigaStreams. If these circuit GigaStreams are on the cluster's spine nodes, you can create the map in the leaf-and-spine cluster in one of two ways:

- **COMBINE** - A single map that combines the circuit GigaStreams from all spine nodes
- **SEPARATE** - Multiple maps, with one map per spine node's circuit GigaStream. This is the default mode.

NOTE: Use SEPARATE mode when the map destination is on nodes that are connected to all spine nodes.
Use COMBINE mode to optimize fabric map deployments so that the cluster can use a single logical map to forward the traffic to the destination, even when traffic arrives via multiple decap circuit GigaStream on the spine nodes.

Limitations of Spine Source Map Mode

- If existing maps are in SEPARATE mode, the future maps from the same circuit GigaStream continue to be SEPARATE even if the global mode is switched to COMBINE.
- If existing maps are in COMBINE mode, the future maps from the same circuit GigaStream to be COMBINE even if the global mode is switched to SEPARATE.
- To truly change mode for the circuit GigaStream, you must delete the existing maps and recreate them after changing the global setting.

Configuration Sync

You can enable this option to automatically turn device-generated flow maps into Traffic Policy.

Troubleshoot Traffic Policy Issues

This section describes how to use the troubleshooting capabilities in GigaVUE-FM to diagnose and resolve issues with Traffic Policy used for Traffic Intelligence, Application Intelligence, and Subscriber Intelligence solutions.

The Troubleshooting view for a deployed Traffic Policy provides:

- A consolidated health view of Traffic Policies and their components (sources, applications, destinations).
- Near real-time (NRT) statistics for supported map types, including:
  - PassAll and ByRule maps
  - Collector maps
  - Inline first-level and second-level maps
  - Maps used in Subscriber Intelligence topologies (e.g., GTP, 5G user/control plane).
- A canvas based visualization of the policy topology, with health indicators for each logical block and links to additional details.

NOTE: Draft or undeployed policies do not show the Troubleshoot option.

Use this view whenever a Traffic Policy does not behave as expected (for example, no traffic, degraded health, or migration related issues).

NOTE: If a Traffic Policy deployment fails and you later see "**No changes have been made**" when you click Deploy without changing the policy, this is expected. After fixing the underlying issue, either wait for the system to retry automatically or use *Dashboard> Actions >Resubmit All* to trigger a new attempt.

View the Troubleshooting Page

To open troubleshooting for a specific **Traffic Policy** :

- In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
- The first view in Troubleshoot is a topology canvas for the selected Traffic Policy:
  - For most Flow Map and Application Intelligence policies, the topology typically spans a single cluster.
  - For Fabric Map–style policies, the topology can span multiple clusters, showing how clusters and maps are connected.
- Locate the required Traffic Policy in the list view.
- Click on **View Troubleshoot** to view the topology in detailed manner.
- Confirm that the map topology, health indicators, and NRT statistics are visible for the selected policy.
- Under the **View Info** section the following details will be listed

- Source Ports - The Port IDs, Alias, Type and Status.
- Destination Ports- The Port IDs, Alias, Type and Status.
- Auto-refresh: when enabled, NRT stats refresh automatically; when disabled, the last snapshot is frozen.
- Clear Stats: clears existing counters so you can observe fresh statistics.

- You can also use the **View Map Info** page to view a detailed information on all components.

Troubleshooting Statistics and Health Indicators

The Troubleshooting view uses color-coded health states and messages to show the status of the policy and its components.

Policy Level Health

The overall policy health indicates whether any component in the Traffic Policy is unhealthy:

- Healthy (green): All monitored components (sources, applications, destinations) are up and passing traffic as expected.
- Warning (yellow/amber): A component is in a degraded or partially available state. Traffic may be impacted.
- Critical (red): One or more components are down or misconfigured, and traffic is likely not flowing as expected.

The policy health aggregates underlying component health from sources, destinations, and application elements (for example, GSOP-based GigaSMART chains).

Component Level Health

Each logical block on the Traffic Policy canvas has its own health state:

- Source and Destination components- Shows link state and configuration issues for network ports, tool ports, port groups, GigaStream, vPorts, and inline components. If a component's underlying device ports are down, the component is shown with a critical health state and a corresponding message.
- Application components (GSOP / App chains)- The health of GigaSMART applications (such as Dedup, Masking, APF, AMI/AFI, etc.) is reflected either:
  - On the canvas (for top-level health of the application block), and/or in the Applications tab, which provides additional metadata and health details for GSOP elements.
  - If application engines are down, the health state is marked critical, and the corresponding reason string includes GSOP metadata and vPort.
- Depending on the map type, the view shows:
  - Packet and Byte counters for:
    - PassAll maps

- ByRule maps (rule-based VLAN/IP match conditions)
- Collector maps
- Inline maps (first-level, second-level, and transit maps)

- Rule-level counters (for example, per pass rule or drop rule).

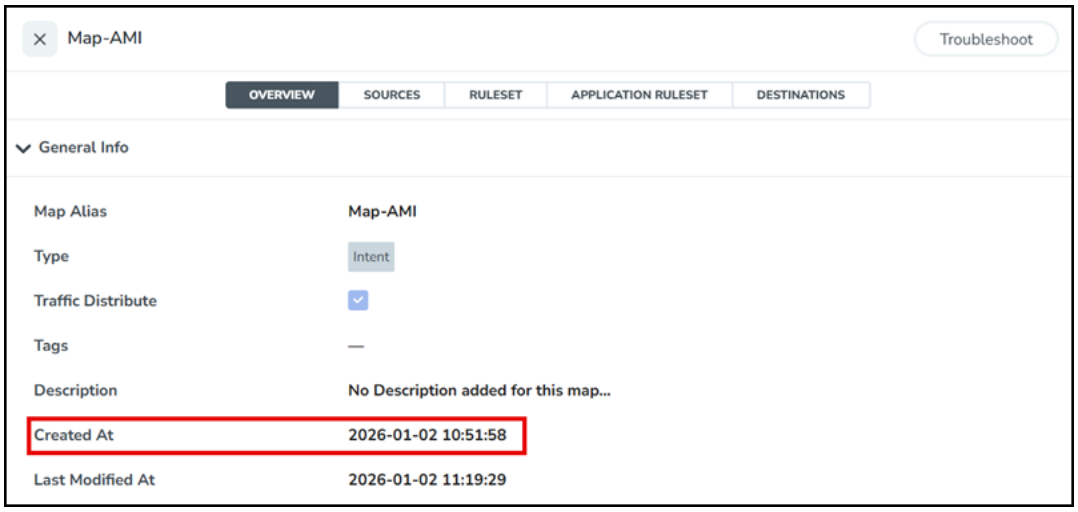
FAQ Traffic Policy

Why don't I see traffic in the dashboard after creating the Application Intelligence solution?

After you create an Application Intelligence solution, traffic should begin appearing in the dashboard within 5 to 7 minutes. If you do not see any traffic, first verify the creation time of the solution to ensure enough time has passed for traffic ingestion to start.

To view the creation time:

1. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
2. Select the required Traffic Policy you have configured for the Application Intelligence solution. The Traffic Policy canvas appears with the map components deployed.
3. Click **More Info**. The Traffic Policy details appear. You can view the creation time here.



If more than 7 minutes have elapsed and the dashboard still shows no traffic, the issue may be related to your environment or configuration. In that case, contact Customer Support for assistance.

What can be migrated?
Traffic Policy migration supports Traffic Intelligence , Subscriber Intellugence (legacy Flow Map–based solutions, user-created Fabric Map–based solutions0 and Application Intelligence solutions. Internally generated Fabric Maps that belong to other solutions (for example, FM-auto maps) are treated as solution artifacts and are not migrated.

What solutions cannot be migrated?
Flexible Inline and Intent-Based Orchestration (IBO) solutions cannot be migrated to Traffic Policy.

Where can I see the migration results?
Use **Traffic > Traffic Policy > View Migration Result**s. You can also open View Migration Results from the legacy Flow Maps or Fabric Maps pages where the migration was triggered.

What should I check in the migration results?
Verify that Total and Success counts match your expectations. If there are any Skipped or Failed maps, click the run and then the Skipped or Failed count to see which maps were affected and why.

What happens after migration?
After migration, the configuration is represented as Traffic Policies. Any further edits to those maps must be done from the Traffic Policy page. The corresponding entries on the legacy Flow Maps / Fabric Maps / App Intel pages are treated as internal artifacts and are no longer intended for user edits.

Will migration affect traffic on the nodes?
No. Migration does not change the existing data-plane configuration on the devices; it only changes how the configuration is managed in GigaVUE-FM.

What happens to statistics collected before migration?

Existing traffic trending statistics continue to work as they are. New dashboards and views created for Traffic Policy will show statistics starting from the time of migration.

What should I do if there is a failure in migration?

Open the migration result, click the Failed count, and review the reason shown for each map. After collecting a sysdump and the failure details, contact Gigamon Support for assistance.

Why do I see “Unable to acquire lock for cluster <clustername>” during migration even after disconnecting the node from FM?
This usually means a configuration sync was already in progress when you disconnected the node. The Disconnect option does not interrupt in-progress communication between GigaVUE-FM and the node. Wait 5–10 minutes and try again.

GigaSMART®

GigaSMART Operations

Use the **GigaSMART Operations** page to create GigaSMART® operations. **GigaSMART Operations** operations consist of a name and a supported combination of the available GigaSMART applications that you have licensed.

- Refer to How to Combine GigaSMART Operations for details on supported combinations of GigaSMART operations.
- Refer to Order of GigaSMART Operations for information on the order in which GigaSMART components are applied in a single operation.
- Refer to Supported GigaSMART Operations for information on supported GigaSMART operations on physical and virtual nodes.

The details of each GigaSMART operation are described in the following sections, organized by solution type:

Application Intelligence	These GigaSMART operations provide network and application visibility to increase the efficiency of security and analytic tools: • Application Intelligence • Application Session Filtering
Subscriber Intelligence	These GigaSMART operations provide network and subscriber visibility to increase the efficiency of network monitoring, security, and customer experience tools: • GTP Correlation • SIP/RTP Correlation • FlowVUE
GigaSMART TLS/SSL Decryption for Out-of-Band Tools	GigaSMART TLS/SSL Decryption for Inline and Out-of-Band Tools provide network visibility to increase the efficiency of network performance tools.
Traffic Intelligence	These GigaSMART operations provide network visibility to increase the efficiency of network performance tools: • Adaptive Packet Filtering • Advanced Load Balancing • De-duplication • Flow Masking • Header Stripping • Masking • NetFlow Generation • Slicing • SSL/TLS Decryption • PCAPng Application • Tunneling

For additional information about working with GigaSMART operations, refer to:

- About GigaSMART® Applications for devices that support GigaSMART.

- Work with GigaSMART Operations for rules, tips, and general guidance about working with GigaSMART Operations.
- Create GigaSMART Operations – A Summary to get started with GigaSMART.
- GigaSMART Logs to learn about GigaSMART application logs.

For information about the types of licenses available, how to add and activate GigaSMART licenses, refer to the *GigaVUE Administration Guide*.

Refer Efficiently chaining the GigaSMART Apps to Optimize Traffic Before Forwarding to the Tools for more detailed information.

About GigaSMART® Applications

GigaSMART applications are packet modification features available on the following GigaVUE HC Series nodes:

- Standalone GigaVUE-HC3 with SMT-HC3-C08 module installed.
- Standalone GigaVUE-HC1 nodes.
- Any GigaVUE HC Series node operating in a cluster with one of these node types.

NOTE: This section refers to any of these nodes as **GigaSMART-enabled** – they are all capable of using GigaSMART operations.

You can use both GigaVUE-FM and the CLI to create GigaSMART operations combining the GigaSMART applications, and then use them with other map rule criteria, and apply them in map rules on any network port in the node or cluster.

Quick Glance- How to Configure a GigaSMART Application

This section provides you a quick glance on document sections that can get you started on configuring the GigaSMART Applications.

GigaSMART Application	Refer to
Application Intelligence	
Application Intelligence	Create an Application Intelligence Session in Physical Environment
Application Filtering Intelligence	Create Application Filtering Intelligence for Physical Environment

GigaSMART Application	Refer to
Application Metadata Intelligence	Create Application Metadata Intelligence for Physical Environment
Application Session Filtering	Define ASF Session
Subscriber Intelligence	
GTP Correlation	• GTP Correlation Configuration Examples • GigaSMART GTP Whitelisting and GTP Flow Sampling Examples • Display Flow Ops Reports • GTP Engine Grouping Configuration Examples • GTP Stateful Session Recovery
SIP/RTP Correlation	• Configure SIP/RTP Correlation Engine
FlowVUE	• Configure FlowVUE o FlowVUE Configuration Examples
5G CUPS	• 5G Load Balancing • Configure CPN-UPN Communication Solution using Ansible • About Flow Sampling Rules and Maps • Monitoring CUPS Solution
GigaSMART TLS/SSL Decryption for Inline and Out-of-Band Tools	• Gigamon Inline Solutions
GigaSMART TLS/SSL Decryption for Out-of-Band Tools	• GigaSMART Passive TLS/SSL Decryption
Traffic Intelligence	
GigaSMART Adaptive Packet Filtering (APF)	• Implement APF Through the UI o Adaptive Packet Filtering Examples
Advanced Load Balancing	• Stateful Load Balancing • Stateless Load Balancing • Enhanced Load Balancing
De-duplication	• De-duplication Configuration Steps o Configure GigaSMART Parameters for Packet De-duplication o Example – GigaSMART De-duplication
Flow Masking	• GigaSMART Encapsulated Traffic Performance Enhancement
Header Stripping	• GigaSMART Header Stripping • Generic Header Stripping
Masking	• GigaSMART Masking

GigaSMART Application	Refer to
NetFlow Generation	• Configure Netflow Generation
Slicing	• Create Advanced Flow Slicing Profile • GigaSMART Packet Slicing
Tunneling	• Custom Tunnel Decapsulation Configuration Example • ERSPAN Tunnel Header Removal ○ ERSPAN Type III Tunnel Header Removal • GigaSMART IP Encapsulation/Decapsulation (GigaSMART Tunnel) • Configure L2GRE Tunnel Encapsulation and Decapsulation • GigaSMART VXLAN Tunnel Encapsulation • GigaSMART VXLAN Tunnel Decapsulation • Configuration

Application Intelligence Solutions

Designed to meet the needs of NetOps, SecOps and DevOps teams, these GigaSMART operations provide complete network and application visibility to reduce Mean Time to Detect (MTTD) and Mean Time to Respond (MTTR) and better secure and manage their hybrid cloud infrastructure.

GigaVUE-FMApplication Intelligence streamlines network telemetry by:

- Identifying and filtering applications and protocols in North-South and lateral network traffic.
- Generating metadata from applications and protocols involves creating detailed information about the transmitted data, including generating NetFlow/IPFIX for effective monitoring of standard information elements.

Additionally, Gigamon's Deep Observability Pipeline can enhance Security Information and Event Management (SIEM) and Observability tool deployments by :

- Adding deeper context to logs by providing full-stack visibility (L2-L7) for troubleshooting performance and security issues. This solution remains transparent to users, preventing easy tampering or falsification by threat actors.
- Providing visibility into traffic spanning unmanaged devices.

To know more about essential features and capabilities of Gigamon’s Application Intelligence Solution refer to ***GigaVUE Application Intelligence Solutions Guide***.

Subscriber Intelligence

Designed to meet the needs of service providers, these GigaSMART operations provide complete network and subscriber visibility to increase the efficiency of network monitoring, security, and customer experience tools.

NOTE: GigaSMART mobility is not supported in GigaVUE-HCT devices.

Subscriber Intelligence	
GTP Correlation	▪ GigaSMART GTP and CUPS Correlation ▪ GigaSMART GTP and CUPS Correlation ▪ GigaSMART GTP Whitelisting and GTP Flow Sampling ▪ Display Flow Ops Reports ▪ GTP Overlap Flow Sampling Maps ▪ GTP Scaling ▪ GTP Stateful Session Recovery
SIP/RTP Correlation	▪ GigaSMART SIP/RTP Correlation
FlowVUE	▪ GigaSMART IP FlowVUE
5G CUPS	• 5G Correlation • 5G Flow Sampling and Filtering • Monitoring of Subscriber Intelligence Solutions • User Plane Node Traffic Monitoring • Configure CPN-UPN Communication Solution using Ansible

Installation and Configuration of Subscriber Intelligence Solution using Ansible

The Gigamon Ansible Module consists of various playbooks that converts high level user intent in YAML format into JSON format.

The GigaVUE-FM uses the JSON format and translates the inputs into various individual components like GigaSMART GSgroups, GSOPs, FlowMaps, etc that are configured on the Physical or Virtual devices.

The Gigamon Ansible Module exposes playbooks that allows the configuration and maintenance of the Subscriber Intelligence solutions and the GigaVUE-FM GUI allows you to visualise, monitor and troubleshoot Subscriber Intelligence Solutions.

Refer to the following sections for configuring the Subscriber Intelligence solution:

- System Requirements
- Installation and Configuration of Gigamon Ansible Module
- Rules and Notes

The Gigamon Ansible Module allows you to configure the following:

S.No	Steps	Refer to..
1.	CUPS Solution	Configure CPN-UPN Communication Solution using Ansible
2.	Non-CUPS Solution	Configuration of Non-CUPS using Ansible

System Requirements

Ensure that the following environment is available before installing 'gigamon-ansible:

- **Python version:** 2.7.15 or greater
- **Operating System**: Linux
- **Ansible version**: 2.9.4 or greater
- **Python Packages**
  - **requests** - Install this using *pip install requests*
  - **ruamel.yaml** -Install this using *ruamel.yaml*
  - **jsonschema** - Install this using *pip install jsonschema*
  - **netaddr** - Install this using *pip install netaddr*

Installation and Configuration of Gigamon Ansible Module

Gigamon-Ansible module can be installed as follows:

Package	Operating System	Commands to install the Package
RPM package	CentOS	`sudo yum install <packageName>.rpm`
Deb package	Ubuntu	`sudo apt install <packageName>.deb`

The package is extracted under the path ***/usr/local/share/gigamon***.

For setting the Pythonpath, add the following to *~/.baschrc* and source it:

```
export INSTALL_DIR=/usr/local/share/gigamon
export PYTHONPATH=$PYTHONPATH:$INSTALL_DIR
```

Deployment Report

Everytime a Subscriber Intelligence solution is deployed/updated/deleted, a deployment report is generated in the path declared in the **ansible_inputs.json** file.

Check Mode

The checkmode feature allows you to find the difference between the configuration that is already present in the GigaVUE-FM against the payload that you are trying to apply without applying the new payload in GigaVUE-FM. This allows you to quickly find the components that gets affected if the new configuration is applied.

The check mode can be enabled by adding ***--check*** to the command that triggers the playbook.

```
/usr/bin/ansible-playbook --check -e '@~/cupsSolution/ansible_inputs.json' -i
~/cupsSolution/cups_inventory /usr/local/share/gigamon-ansible/playbooks/cups/deploy_cups.yml
```

The output of **--check** is a Deployment Report that contains the difference in configuration as shown:

Deployment Report Checkmode

```
Deployment_Report_Checkmode
deploymentPayloadDiff:
  updated:
  - path: //trafficPolicies/LTE/whitelisting/flowMaps/wlMapInternetToEEA/rules/rule_1/interface
    old_value: S11
    new_value: s11
  - path: //trafficPolicies/LTE/whitelisting/flowMaps/wlMapGeoProbe/rules/rule_1/apn
    old_value: apn.airtel.com
    new_value: apn.vodafone.com
  removed: []
  created: []
deploymentRequest: EDIT
deploymentResponse: CHECK MODE
timeStamp: 09-Mar-2020::12:12:1583781162
```

Reapplying Golden Payload

It is also possible to restore the configuration from the generated Golden Payload file using one of the following two commands:

- **/usr/bin/ansible-playbook -e '@~/mobilitySolution/ansible_inputs.json' -e 'applyGP=True' -i ~/mobilitySolution/mobility_inventory /usr/local/share/gigamon-ansible/playbooks/mobility_solution/deploy_mobility_solution.yml**
- **/usr/bin/ansible-playbook -e '@~/mobilitySolution/ansible_inputs.json' -e 'applyGP=True' --ask-vault-pass -i ~/mobilitySolution/mobility_inventory /usr/local/share/gigamon-ansible/playbooks/mobility_solution/deploy_mobility_solution.yml**

It is also possible to enable checkmode to find the difference between the configuration on GigaVUE-FM and the configuration of the Golden Payload. The command to do is:

- **/usr/bin/ansible-playbook --check -e '@~/mobilitySolution/ansible_inputs.json' -e 'applyGP=True' -i ~/mobilitySolution/mobility_inventory /usr/local/share/gigamon-ansible/playbooks/mobility_solution/deploy_mobility_solution.yml**
- **/usr/bin/ansible-playbook --check -e '@~/mobilitySolution/ansible_inputs.json' -e 'applyGP=True' --ask-vault-pass -i ~/mobilitySolution/mobility_inventory /usr/local/share/gigamon-ansible/playbooks/mobility_solution/deploy_mobility_solution.ymlS**

NOTE: It searches the golden payload file to reapply the configuration in the default location if the file path is not defined in the **ansible_inputs.json** file.

Rules and Notes

You must ensure the following rules and notes while deploying Subscriber Intelligence Solution:

- GigaVUE-FM is reachable from the server on which Ansible is executed.
- All the required licenses are installed on the Gigamon devices.
- The required permission for RBAC is available.
- Configurations that are not handled by the Mobility playbook are prec-configured.
- Migration of Subscriber Intelligence Solution from GigaVUE-FM 5.9 to any other version is impossible.

Configuration of Non-CUPS using Ansible

To configure a Non-CUPS solution, perform the following steps:

S.No	Steps	Refer to..
1.	Creating Inventory Directory	Creating Inventory Directory
2.	Creating ***fmInfo.yml***	Creating fmInfo.yml
3.	Creating ***ansible_inputs.json***	Creating ansible_inputs.json
4.	Creating Non-CUPS inventory file	Creating Non-CUPS inventory file
5.	Creating **host_vars** directory	Creating host_vars directory
6.	Creating **host_vars** files	Creating host_vars files

Creating Inventory Directory

Create an *Inventory Directory* to store all the Non-CUPS related configuration files.

```
username@fmreg26:~$ mkdir non-CupsSolution
username@fmreg26:~$ ls -l
drwxr-xr-x 2 ddaniel fmtaf 4096 May 11 11:59 non-CupsSolution
```

Creating ***fmInfo.yml***

Create ***fmInfo.yml*** file inside the Inventory Directory that contains the information such as ip-address, username and password.

To create the file, refer to Schema.

Creating ***ansible_inputs.json***

Create '***ansible_inputs.json***' file inside the Inventory Directory.

- ```
  gigamon@fmreg26:~/non-cupsSolution$ ls -l
  -rw-r--r-- 1 gigamon fmtaf 355 May 11 12:24 ansible_inputs.json
  -rw-r--r-- 1 gigamon fmtaf 172 May 11 12:10 fmInfo.yml
  ```

  File name: ansible_inputs.json

  To create a file, refer to Schema.
  ```

Creating Non-CUPS inventory file

Create the ***mobility_inventory*** file inside Inventory Directory.

```
gigamon@fmreg26:~/nonCupsSolution$ touch mobility_inventory

gigamon@fmreg26:~/nonCupsSolution$ ls -l
-rw-r--r-- 1 gigamon fmtaf 355 May 11 12:24 ansible_inputs.json
-rw-r--r-- 1 gigamon fmtaf 396 May 11 14:22 mobility_inventory
-rw-r--r-- 1 gigamon fmtaf 172 May 11 12:10 fmInfo.yml
```

The file contains the details of the following groups and provide the inputs to the groups as shown in the following table:

S.No	Groups-Input
Note: —You can provide the input or leave the field empty if you don't want to use the playbook to configure the following groups.	
1.	**Ports**—Name of the Cluster or standalone device IP that contains the ports that need to be configured.
2.	**IPInterfaceSolution**—Name of the Cluster or standalone device IP on which the IPInterfacesolution needs to be configured.
3.	**Tool Groups**—Name of the Cluster or standalone device IP on which the ToolGroup needs to be configured.
4.	**GigaStreams**—Name of the Cluster or standalone device IP on which the Gigastreams needs to be configured.
5.	**GTPWhitelist**—Name of the Cluster or standalone device IP on which the GTPWhitelist Data Base needs to be configured.
6.	**Policies**—Name of the Global policy or policies.
7.	**GTP**— Name of the GTP node or nodes.
6.	**CPN**—Name of the CPN or CPNs.
7.	**UPN**—Name of the UPN or UPNs.
8.	**SAM**—Name of SAMs Exporter node or nodes.
9.	**Sites**—Name of the site or sites and the names of the CPN/UPN participating in the site or sites.
10.	**CUPS**—Name of the file containing information of the solution level RBAC Tags.

File name: mobility_inventory (Single GigaVUE-FM instance)

```
[IPInterfaceSolution]
```

```
[ToolGroups]
cluster-two
cluster-one

[Gigastreams]
cluster-two
cluster-one

[GTPWhitelist]
cluster-two
cluster-one

[Ports]
cluster-two
cluster-one

[Policies]
5g_policy_1

[GTP]
gtp_CorrelationNode1

[CPN]

[UPN]

[SAM]

[Sites]
UK cpn_list='["cpnUkLTE"]' upn_list='[]' sam_list='[]'
Dallas cpn_list='[]' upn_list='["upnDallas"]' sam_list='[]'

[MobilitySolution]
mobilitySolution1 cups_5g_global_policy=5g_policy_1 sites='["UK", "Dallas"]'
```

File name: cups_inventory (Multiple GigaVUE-FM instances)

```
[IPInterfaceSolution]

[ToolGroups]
cluster-two fm_ip=192.168.36.2
cluster-one fm_ip=192.168.36.3

[Gigastreams]
cluster-two fm_ip=192.168.36.2
cluster-one fm_ip=192.168.36.3

[GTPWhitelist]
cluster-two fm_ip=192.168.36.2
cluster-one fm_ip=192.168.36.3
```

```
[Ports]
cluster-two fm_ip=192.168.36.2
cluster-one fm_ip=192.168.36.3

[Policies]
5g_policy_1

[GTP]
gtp_CorrelationNode1

[CPN]

[UPN]

[SAM]

[Sites]
UK cpn_list='["cpnUkLTE"]' upn_list='[]' sam_list='[]'
Dallas cpn_list='[]' upn_list='["upnDallas"]' sam_list='[]'

[MobilitySolution]
mobilitySolution1 cups_5g_global_policy=5g_policy_1 sites='["UK"]' fm_ip=192.168.36.2
mobilitySolution2 cups_5g_global_policy=5g_policy_1 sites='["Dallas"]' fm_ip=192.168.36.3
```

Creating **host_vars** directory

Create **host_vars** directory inside the Inventory Directory.

```
gigamon@fmreg26:~/nonCupsSolution$ mkdir host_vars

gigamon@fmreg26:~/nonCupsSolution$ ls -l
-rw-r--r-- 1 gigamon fmtaf 355 May 11 12:24 ansible_inputs.json
-rw-r--r-- 1 gigamon fmtaf 396 May 11 14:22 cups_inventory
-rw-r--r-- 1 gigamon fmtaf 172 May 11 12:10 fmInfo.yml
drwxr-xr-x 2 gigamon fmtaf 4096 May 11 14:48 host_vars
```

Creating **host_vars** files

Every unique element under each group in the ***mobility_inventory*** file needs to have a file, with the same name as the element, inside **host_vars** directory. This file has properties of the groups that it belongs to.

Below are the templates of various host_vars files.

Prerequisite

```
---
validate_certs: false
Ports:
- port:
  - 1/1/x1
  - 1/1/x2
  adminStatus: enable
  type: network

GTPWhitelist:
- alias: gtp1
  imsi: 310260564627811,310260564627812
  state: present
- alias: gtp2
  inputFile: './whitelistKeys/TenIMSIs_Valid.txt'
  state: present

Gigastreams:
- alias: toolGS_C11
  ports:
  - 4/1/x1
  - 4/1/x2
  type: hybrid
  state: present
- alias: toolGS_C12
  ports:
  - 4/1/x3..x4
  type: hybrid
  state: present


ToolGroups:
- alias: pgGrp_C11
  ports:
  - 2/1/x1
  smartLb: false
  type: tool
  state: present
- alias: pgGrp_C12
  ports:
  - 2/1/x2
  smartLb: false
  type: tool
  state: present
```

Site

For information about **Site**, refer to Schema.

cpNode

For information about **cpNode**, refer to Schema.

upNode

For information about **upNode**, refer to Schema.

5GPolicy

For information about **5GPolicy**, refer to Schema.

LTEPolicy

For information about **LTEPolicy**, refer to Schema.

Deployment of Non-CUPS Solution

To deploy the Non-CUPS solution, follow these steps:

Set up of additional variable for Single GigaVUE-FM instance

For a Single GigaVUE-FM instance deployment, you must set an additional environment variable as follows.

```
export ANSIBLE_FM_IP=192.168.36.2
```

It searches the login details of the GigaVUE-FM IP in ***fmInfo.yml*** file.

Execute the Playbook

- You can execute the playbook and deploy the CUPS solution using the following command:

  /usr/bin/ansible-playbook -e '@~/cupsSolution/ansible_inputs.json' -i ~/cupsSolution/cups_inventory /usr/local/share/gigamon-ansible/playbooks/cups/deploy_cups.yml

- If the **fmInfo** file is encrypted, use the following command to execute and deploy CUPS solution:

  /usr/bin/ansible-playbook -e '@~/cupsSolution/ansible_inputs.json' --ask-vault-pass -i ~/cupsSolution/cups_inventory /usr/local/share/gigamon-ansible/playbooks/cups/deploy_cups.yml

> **NOTE:** The multiple YML files created inside the **host_vars** are concatenated, converted into JSON format and sent to GigaVUE-FM.

GigaSMART GTP and CUPS Correlation

Required License: GTP Filtering and CUPS Correlation

Supported Devices: GigaVUE-HC3 Gen 2

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

The GigaSMART GTP application correlates traffic based on mobile subscriber IDs in the packet data networks of service providers. It provides a mechanism to filter and forward session traffic for subscribers to tools. GTP and CUPS correlation assists mobile carriers in debugging and analyzing GTP traffic in their 3G/4G networks.

NOTE: For Generation 3 cards, you can use Flow Sampling with sampling rate as 100% to achieve the same behavior as flow filtering.

- GPRS Tunneling Protocol (GTP) is an IP/UDP-based protocol that carries mobile data across service provider networks. The protocol is used in General Packet Radio Service (GPRS) networks such as: Global System for Mobile Communications (GSM), Universal Mobile Telecommunications System (UMTS), and Long Term Evolution (LTE). The protocol encapsulates user data that passes through the core network and carries subscriber-specific signaling traffic.
- GTP includes both control plane (GTP-c) and user-data plane (GTP-u) traffic. To gain an accurate view into the subscriber's session, GTP tunnels are used to correlate the subscriber-specific control plane and user-data plane traffic. A GTP session is the minimum unit of GTP and CUPS correlation consisting of one control and multiple user tunnels. All GTP traffic belonging to the same session is forwarded to the same tool port.
- Using GTP and CUPS correlation, you can filter, replicate, and forward specific subscriber sessions to specific tools by correlating the subscriber IDs that are exchanged as part of the control sessions to the corresponding tunnel IDs (TEIDs) that are part of the user-data plane traffic.

GTP and CUPS correlation provides the following:

- stateful filtering based on subscriber IDs (IMSI, IMEI, and MSISDN)
- stateful filtering based on GTP version or EPC interface
- stateful correlation of GTP-c with GTP-u traffic
- correlation of subscriber ID with corresponding tunnel ID
- forwarding of the subscriber-specific control and user-data plane traffic to a tool or group of tools
- It supports 12 million GTP subscriber sessions for GigaVUE-HC3 nodes.

- option to combine with GigaSMART Load Balancing to load balance GTP traffic to a set of tool ports. For information on GTP load balancing, refer to stateful load balancing in the section GigaSMART Load Balancing. For examples of GTP load balancing, refer to GTP Correlation Configuration Examples. Starting in software version 4.6, GTP load balancing in a cluster is supported for GTP flow filtering. For an example of GTP load balancing in a cluster, refer to GigaSMART GTP Whitelisting and GTP Flow Sampling.

Starting in software version 4.5, a GigaSMART group (gsgroup) associated with GTP applications can have multiple GigaSMART engine port members (**e** ports), up to four, forming an engine group. Refer to GTP Scaling.

Licensing Requirements

- For GTP flow sample with the rule percentage (0 or 100), the **GTP_MAX** license is required.
- For GTP flow sample with the rule percentage in all ranges (between 0 to 100) or for GTP Whitelist, the **GTP_MAX** and **FVUE** licenses are required
- For GTP flow filtering, the **GTP_MAX** license is rquired. The GTP flow filtering is supported only on Gen 2 devices.

Filtering on Subscriber IDs and Version

GTP stateful filtering supports filtering of GTP sessions based on the following subscriber IDs:

Component	Description
imsi	The International Mobile Subscriber Identity (IMSI) is a number that identifies a subscriber of a cellular network. It is a unique identification associated with all cellular networks. An IMSI is usually a 15 digit number, associated with GSM, UMTS, and LTE network mobile phone users.
imei	The International Mobile Station Equipment Identity (IMEI) is a number, usually unique, that identifies 3rd Generation Partnership Project (3GPP), for example, GPRS, LTE, as well as Integrated Digital Enhanced Network (iDEN) mobile phones, and some satellite phones. The IMEI identifies the device, but has no permanent relationship to the subscriber. Instead, the subscriber is identified by transmission of an IMSI number, stored on a SIM card.
msisdn	The Mobile Station International Subscriber Directory Number (MSISDN) is a unique number that identifies subscribers in a GSM or UMTS mobile network. This numbering plan is defined in the ITU-T recommendation E.164. The maximum length of an MSISDN is 15 digits.

In addition to filtering on subscriber IDs, you can optionally filter on GTP version (v1 or v2) or Evolved Packet Core (EPC) interface. Filtering on the EPC interface allows traffic to be segmented for a given interface.

The supported interfaces for EPC filtering are as follows:

- Gn/Gp
- S11U
- S11/S1-U
- S5/S8
- S10
- S2B

NOTE: When GigaSMART receives S5 and S11 GTP-C control traffic that share the same control TEID, and an S5 is received very close in time to the S5 Create Session Response, GigaSMART may retain only one of the S5 or S11 endpoints in the mobility statistics. The other endpoint might not be visible in whitelist map statistics or Flow-Ops reports.

When filtering on EPC interface, you do not also need to specify version, as the version is implied.

To create maps using GTP, specify a **Second Level Flow Sample** map and select **GTP** for the rule. When adding a map rule, you can specify the following:

- subscriber IDs (IMSI, IMEI, or MSISDN)
- number of digits. The maximum number of digits for the IMSI or MSISDN value is 15. The maximum number of digits for the IMEI value is 16. To specify the prefix for IMSI, IMEI, or MSISDN, you can use a wild card character or a digit string followed by a wild card character.
- map comment to label the purpose of a rule or the type of traffic covered by a rule
- GTP version 1 or version 2 (refer to Figure 1GTP Version) or EPC interface Gn/GP, S5/S8, or S10, S2B, or S11/S1U (refer to Figure 2GTP EPC Interface)

NOTE: In a map, version and EPC interface cannot be specified in the same flowrule, but they can be specified in different flowrules.

NOTE: The maximum number of GTP flowrules is 32 per map.

For examples of filtering on GTP version, refer to GigaSMART GTP and CUPS Correlation.

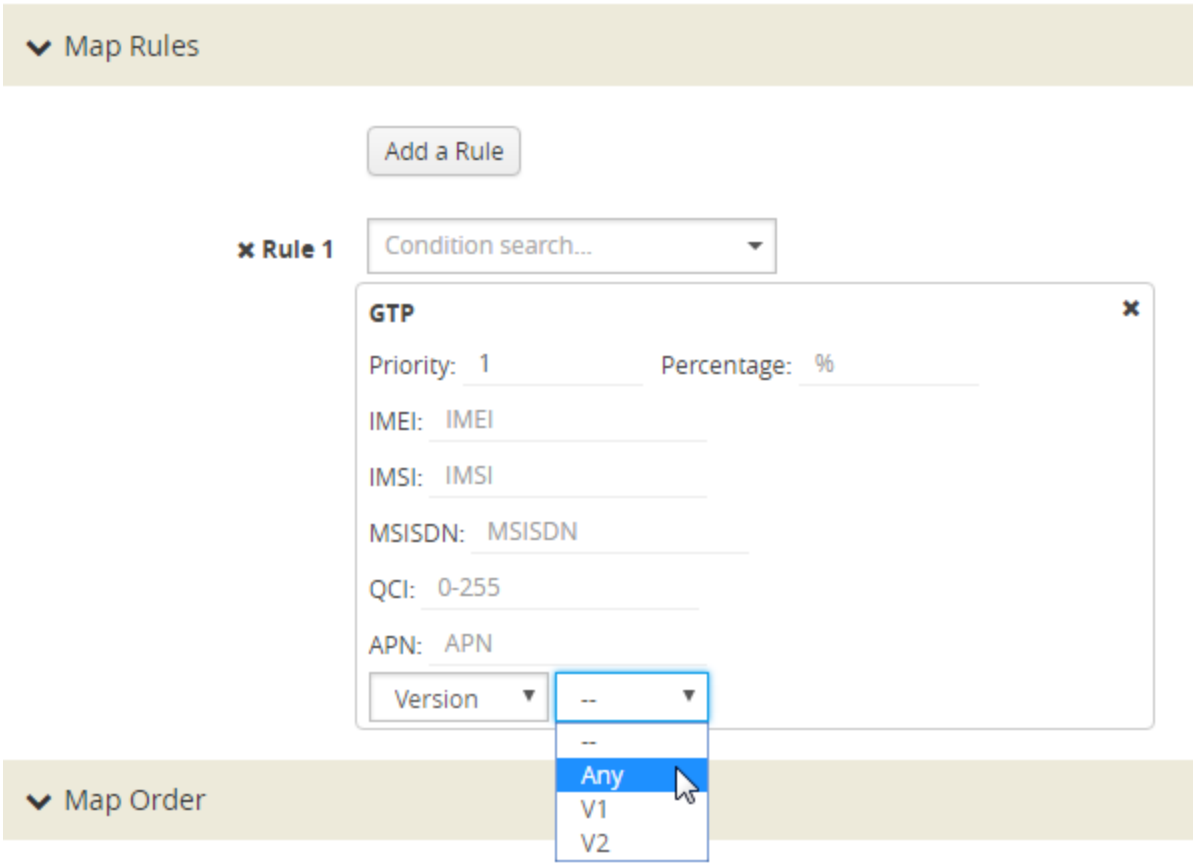


Figure 1 *GTP Version*

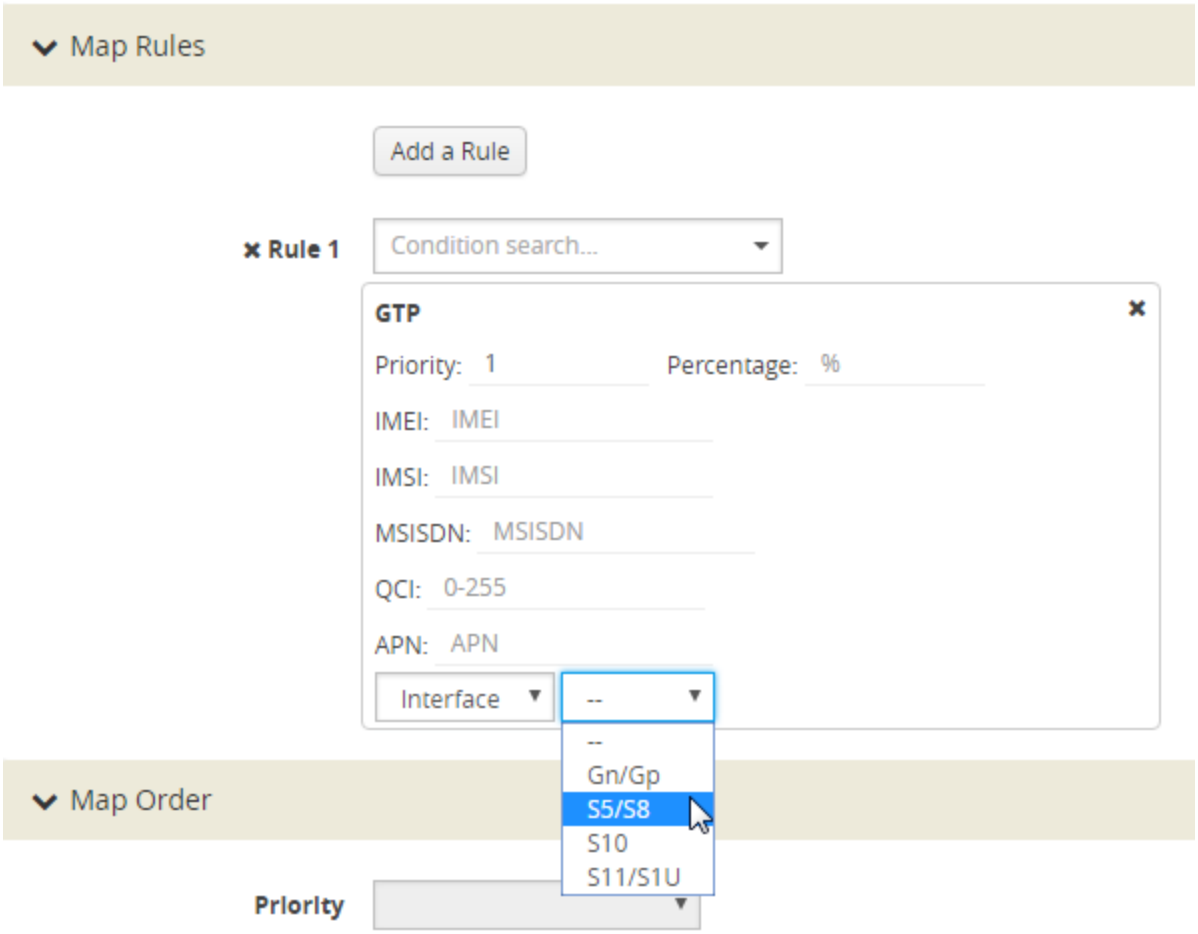
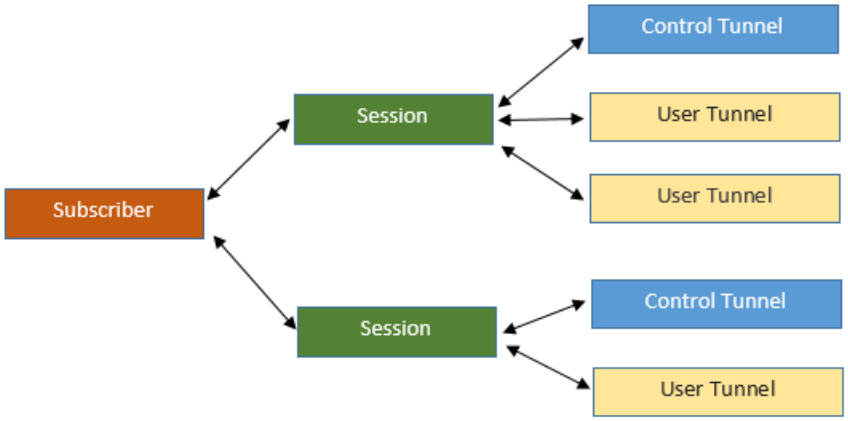


Figure 2 *GTP EPC Interface*

Session Correlation

Each GTP session has one control tunnel and one or more user tunnels. All the tunnels are correlated together into a session. Packets belonging to the same session will be forwarded to the same tool port. Refer to the following figure.



In a second level map, the following can be specified:

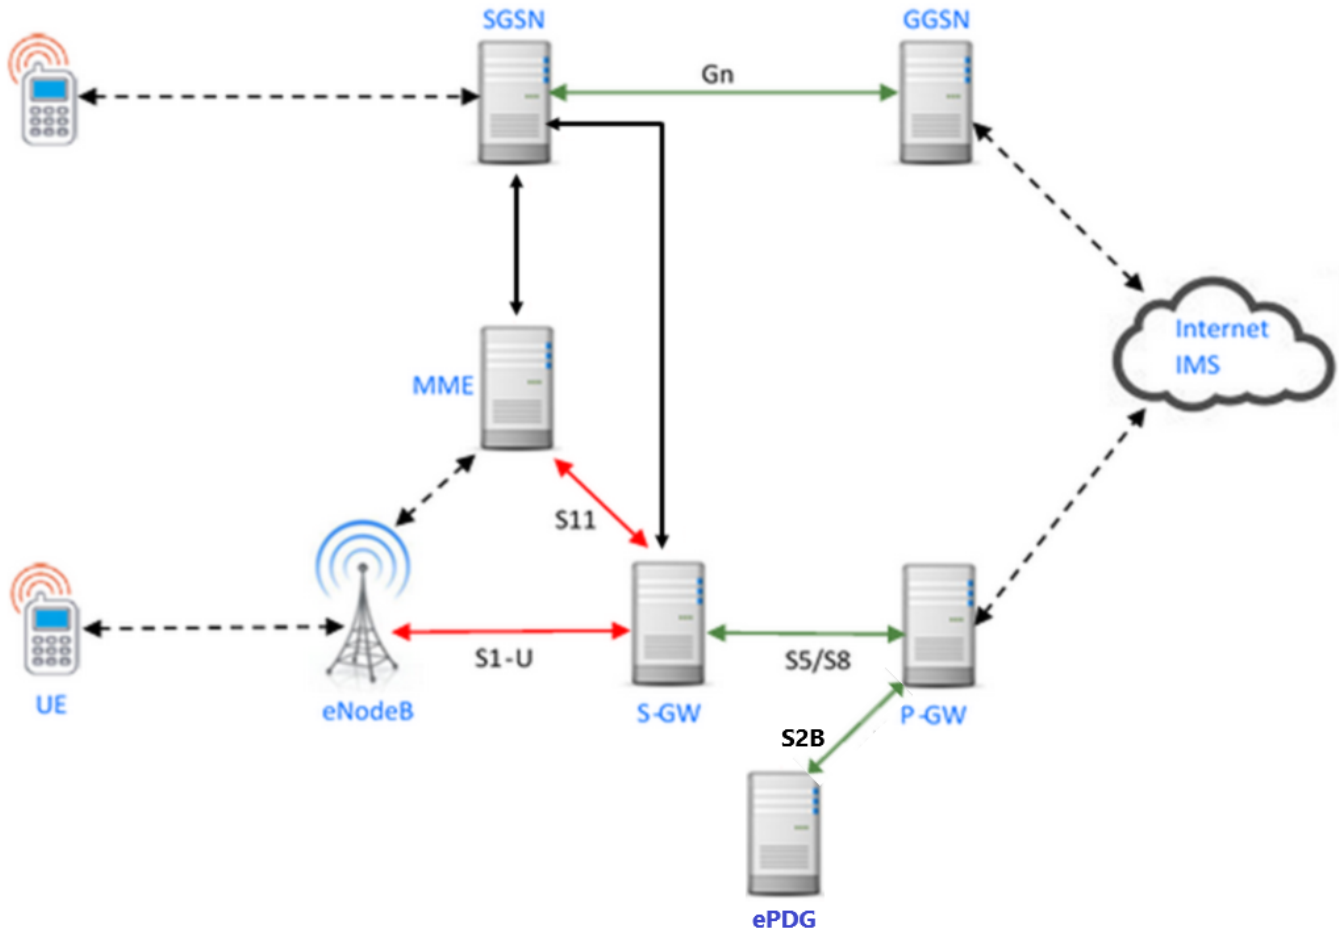
- one tool port—packets from one subscriber (same subscriber ID), from one or more GTP sessions, will be forwarded to the same tool port.
- multiple tool ports—packets from one subscriber (same subscriber ID), from multiple GTP sessions will be correlated and forwarded to same tool port. Using load balancing, GTP traffic that matches the same map but belongs to different subscribers can be load balanced to multiple tool ports.

Supported Interfaces

GTP is used at multiple interfaces by multiple devices in the core network. GTP stateful correlation is implemented for the following interfaces:

- Gn/Gp (for GPRS). The Gn interface is between SGSN-GGSN only.
- S5/S8 (for LTE)
- S1-U, S11U and S11 (for LTE)
- S10 (for S1-based mobility) Refer to Conditional S10 Support
- S2B

Support for interfaces for both GPRS and LTE networks includes the handovers between the different networks. Refer to the following figure.



For LTE networks, the following GTP traffic will be correlated to the specific mobile subscriber and routed to the same tool port:

- GTP-c traffic on the S11 interface between MME and S-GW
- GTP-u traffic on the S11u interface between MME and S-GW
- GTP-u traffic on the S1u interface between eNodeB and S-GW
- GTP-c traffic on the S10 interface between MMEs
- GTP-c traffic on the S5/S8 interface between S-GW and P-GW
- GTP-u traffic on the S5u interface between S-GW and P-GW
- GTP-c traffic on the S2b interface between P-GW and ePDG
- GTP-u traffic on the S2b-U interface between P-GW and ePDG

In order to correlate GTP-c and GTP-u traffic running on different interfaces, you must tap into the correct interfaces, as follows:

- Gn/Gp—one interface runs both GTP-c and GTP-u
- S5/S8—one interface runs both GTP-c and GTP-u

- S1u, S11u, and S11—these three interfaces have to be tapped at the same time to get both GTP-c and GTP-u to perform the correct correlation.
- S2b-C and S2b-U interfaces needs to be tapped to get GTP-c and GTP-u traffic for correlation.

For examples of filtering on GTP interface types Gn/Gp, S5/S8, S1 and S11, refer to GigaSMART GTP and CUPS Correlation.

Conditional S10 Support

The following table outlines support for the S10 interface. In the table, No means not supported, Conditional means there is limited support of the S10 interface.

S10	Support
Forward Relocation Request/Resp	Conditional; IMSI must be present in Forward Relocation Request
Forward Relocation Complete Notification/Ack	Conditional; IMSI must be present in Forward Relocation Request for Forward Relocation Complete Notification/Ack to be supported
Context Request/Response and Ac	Conditional; IMSI must be present
Identification Request/Response	No
Forward Access Context Notification/Ack	No
Relocation Cancel Request/Response	No
Configuration Transfer Tunnel	No

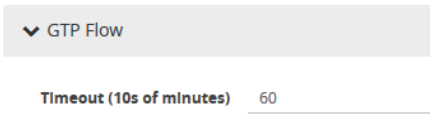
GTP Session Timeout

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

In prior software versions, the complete GTP session timeout was eight hours. Starting with software version 4.2, the GTP session timeout is configurable, with eight hours as the default.

To configure the GTP session timeout, do the following:

1. From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART Groups**.
2. Click **New** to create a new GigaSMART Group or **Edit** to modify an existing one.
3. Under GigaSMART Parameters, go to GTP Flow.
4. Enter the timeout in the **Timeout** field. The following figure shows an example where the timeout value is set to 60.



5. Click **Save**.

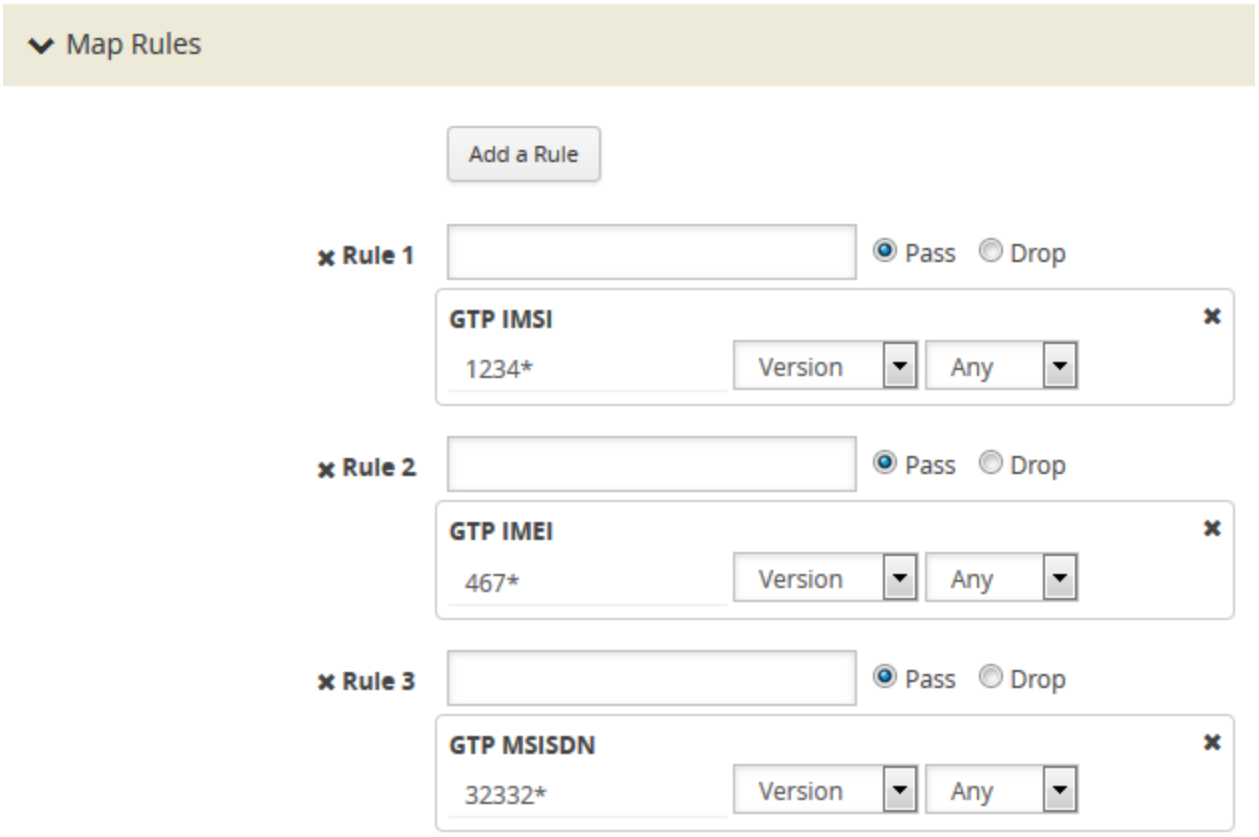
The GTP session timeout disconnects a GTP session if it has been inactive for the timeout value. The timeout can be configured as an integer from 1 to 6000, in increments of 10 minutes. The default value is 48, which is 480 minutes, which is 8 hours.

Priorities for Flow Rules and Maps

One virtual port can have multiple maps, and for each map, you can add multiple flow rules with different filtering attributes. The priorities for flow rules are as follows:

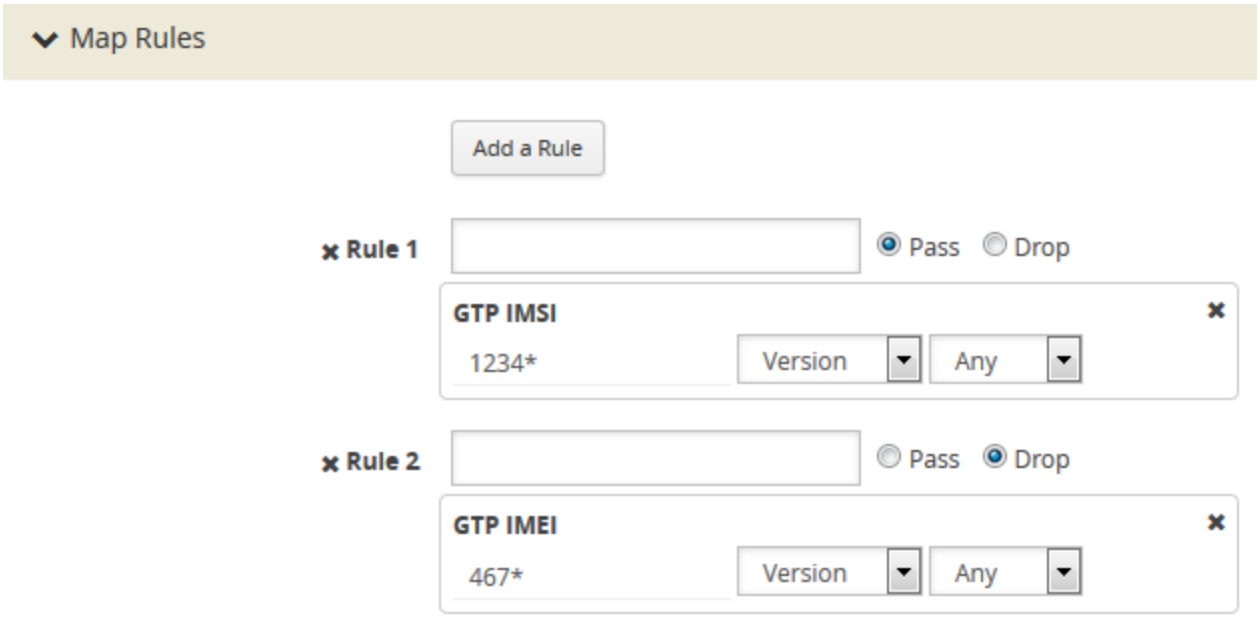
- a rule with a drop action has a higher priority than a rule with a pass action
- for the same pass or drop action, the priorities are IMSI, IMEI, or MSISDN in descending order

In a GTP session, if one IMSI, IMEI, or MSISDN rule is matched, the map is matched. For example, if any one of the following matches any rule shown in the following figure, map1 (which is a Second Level Flow Filter map) is matched:



In addition, in one map, all drop rules are matched first and all pass rules are matched next.

For example, in a GTP session, if an IMSI matches the first rule in map1 and an IMEI matches the second rule in map1, because the drop rule has higher priority, the packet will be dropped:



If multiple maps are matched, the map with the highest priority will be considered for further processing. For example, the Figure 3Rule in Map1 shows the rule for map1 while the rule Figure 82NetFlow Generation Gigamon Solution shows the rule for map2. In a GTP session, if an IMSI matches the first rule in map1 and an IMEI matches the first rule in map2, because map1 has higher priority, the packet will be passed.

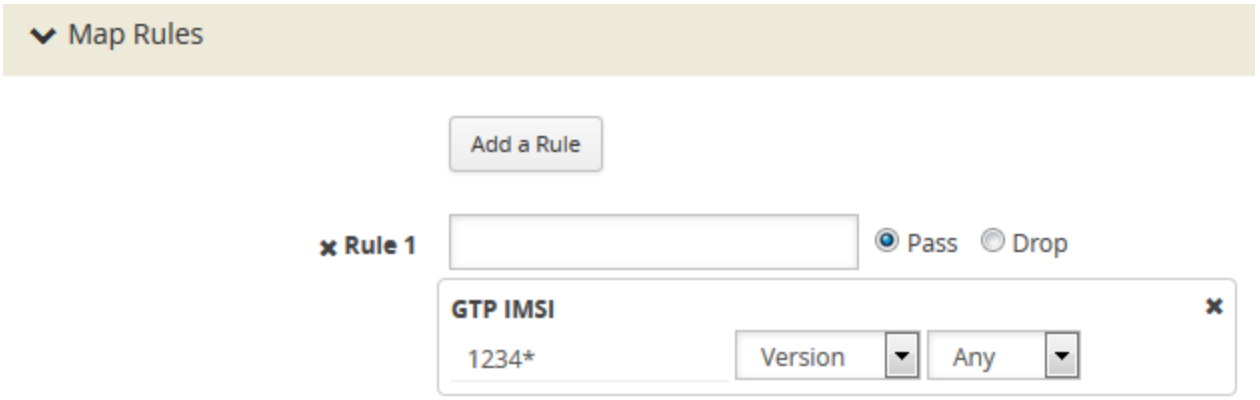


Figure 3 *Rule in Map1*



Figure 4 *Rule in Map2*

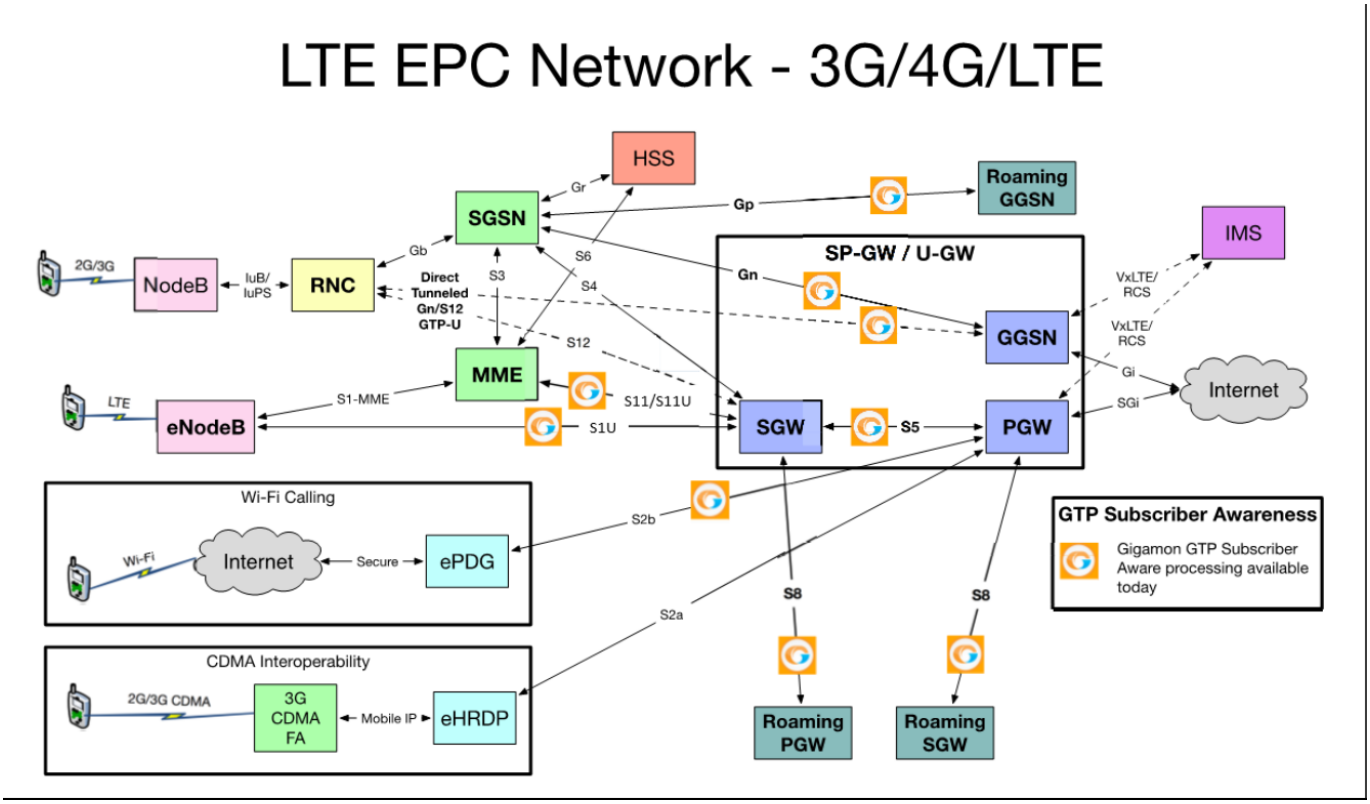
GTP Correlation Configuration Examples

The following sections provide examples of GTP correlation and GTP load balancing. Refer to the following examples:

- Example 1: Identifying High-Value and/or Roaming Subscribers Based on IMSIs
- Example 2: Identifying GTP Versions
- Example 3: Same Subscriber, Filter on Different Versions
- Example 4: Same Subscriber, Filter on Different Interfaces
- Example 5: EPC Filtering in Scenario1
- Example 6: EPC Filtering in Scenario2

Example 1: Identifying High-Value and/or Roaming Subscribers Based on IMSIs

Use GTP correlation to identify high value subscribers based on an IMSI or group of IMSIs. GTP correlation keeps track of the IMSIs that you are interested in monitoring. It correlates them to the corresponding data/user-plane sessions for the subscriber and/or group of subscribers. Filtering on subscriber ID (IMSI) limits the amount of traffic that is sent to monitoring tools.



In Example 1, filter rules are configured to identify and forward all the traffic related to subscribers identified by an IMSI prefix. All traffic specific to the filtered IMSIs 22222222222223*, including GTP-c and GTP-u, is forwarded to a monitoring tool. A shared collector is configured to which traffic not matching the filters is sent.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

Task	Description	UI Steps
1	Configure one network and two tool type of ports.	1. Select Ports > Ports > All Ports. 2. Click Quick Port Editor. 3. In the Quick View Editor set one port to Network and two ports to Tool. 4. Select Enable on each port. 5. Click OK. 6. Close the Quick Port Editor.

Task	Description	UI Steps
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the left navigation pane, go to **System > GigaSMART >**GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type a name for the GigaSMART Group in the Alias field. **4.** Click in the Port List field and select an engine port. **5.** Under GTP Flow set the Timeout. The default is value is 48, which is 480 minutes. **6.** Click Save.
3	Configure the GigaSMART operation and assign it to the GigaSMART group to enable GTP correlation.	**1.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations (GSOP > GigaSMART Operation**. **2.** Type a name for the GigaSMART Group in the Alias field. **3.** Click in the GigaSMART Operations (GSOP) field and select Flow Filtering. **4.** Click OK.
4	Configure a virtual port and assign it to the same GigaSMART group.	**1.** From the left navigation pane, go to **System > GigaSMART >Virtual Ports.** **2.** Click New. **3.** Type a name for the virtual port in the Alias field. **4.** Click in the GigaSMART Groups field, and select the GigaSMART Group created in Task 2. **5.** Click OK.
5	Create a first level map that directs GTP traffic from physical network port/s to the virtual port you created in the previous step. **Note**: In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. • Enter an alias • Select First Level for Type • Select By Rule for Subtype • Select a network port for the Source • Select the virtual port configured in Task 4 for the Destination **4.** Create Rule 1. **a.** Click Add a Rule **b.** Select Pass. **c.** Select Bi Directional. **d.** Select Port Source

Task	Description	UI Steps
		e. Set the source to 2123 **6.** Create Rule 2. **a.** Click Add Rule **b.** Select Pass. **c.** Select Bi Directional **d.** Select Port Source **e.** Set the source to 2152 **6.** Create Rule 3. **a.** Click Add Rule **b.** Select Pass. **c.** Select IPv4 Fragmentation **d.** Set Value to allFragNoFirst **5.** Click Save.
6	Create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, matches IMSIs specified by the flow rule, and sends matching traffic to physical tool ports.	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. ▪ Enter an alias ▪ Select Second Level for Type ▪ Select Flow Filter for Subtype ▪ Select the virtual port configured in Task for the Source ▪ Select the tool port configured in Task 1 for the Destination ▪ Select the GigaSMART Operation created in Task 3 from the GSOP list. **4.** Create a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select GTP IMSI. **d.** Enter 22222222222223* in the IMSI field. **5.** Click Save.
7	Add a shared collector for any unmatched data and send it to the second tool port.	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. ▪ Enter an alias ▪ Select Second Level for Type ▪ Select Collector for Subtype ▪ Select a the virtual port configured in Task for the Source

Task	Description	UI Steps
		▪ Select the second tool port configured in Task 1 for the Destination **4.** Click Save.
8	Display the configuration for Example 1.	To display the configuration for the GigaSMART Group: **1.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups**. **2.** Click on the alias for the GigaSMART Group to display the Quick View. To display the configuration for the GigaSMART Operation: **1.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations (GSOP) > GigaSMART Operation.** **2.** Click on the alias for the GigaSMART Operation to display the Quick View. To display the configuration for the maps: **1.** Select Maps > Maps > Maps. **2.** Click on a map alias to display the Quick View for the map.

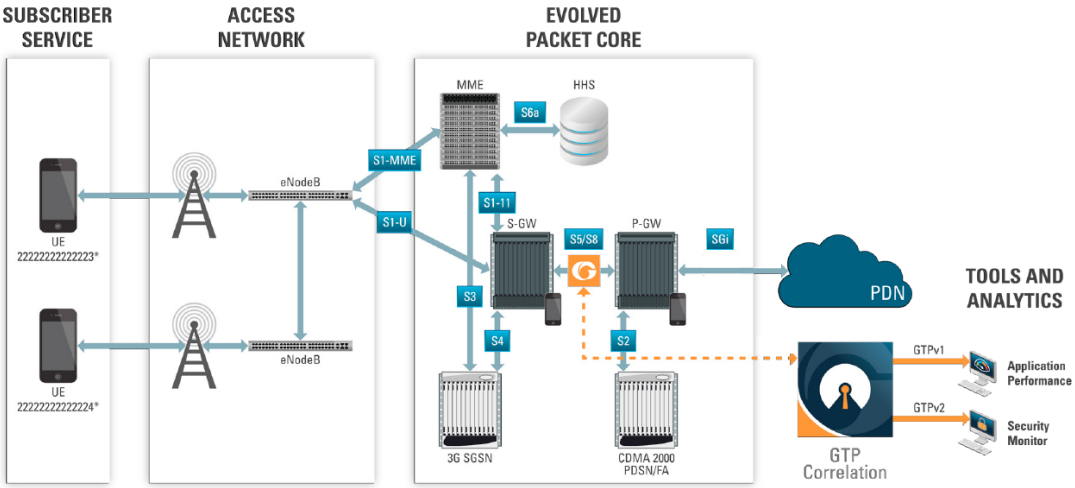
Display GTP Correlation Flow Ops Report Statistics

To display GTP correlation statistics associated with the GigaSMART group, select **GigaSMART > GigaSMART Operations > Statistics**.

Refer to *Flow Ops Report Statistics Definitions for GTP* on page 635 for descriptions of these statistics.

Example 2: Identifying GTP Versions

As part of GTP correlation, GigaVUE nodes also provide the flexibility to identify GTPv1 and GTPv2 messages. GTP version information is typically exchanged only as part of the control sessions. By correlating the control and user-plane sessions, GigaVUE nodes can identify, filter, and forward all sessions specific to a GTPv1 or v2 to one or more monitoring/analytic tools.



In Example 2, EMEI traffic is distributed based on GTP versions as follows:

- Filter and forward GTPv1 to a tool port
- Filter and forward GTPv2 to another tool port

Task	Description	UI Steps
1	Configure one network and two tool type of ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** In the Quick View Editor set one port to Network and two ports to Tool. **4.** Select Enable on each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART Groups.** **2.** Click New. **3.** Type a name for the GigaSMART Group in the Alias field. **4.** Click in the Port List field and select an engine port. **5.** Under GTP Flow set the Timeout. The default is value is 48, which is 480 minutes. **6.** Click Save.
3	Configure the GigaSMART operation and assign it to the GigaSMART group to enable GTP correlation.	**1.** From the device view, select GigaSMART > GigaSMART Operations > GigaSMART Operation. **2.** Type a name for the GigaSMART Group in the Alias field.

Task	Description	UI Steps
		3. Click in the GigaSMART Operations (GSOP) field and select Flow Filtering. **4.** Click Save.
4	Configure a virtual port and assign it to the same GigaSMART group.	**1.** From the left navigation pane, go to **System > GigaSMART >Virtual Ports.** **2.** Click New. **3.** Type a name for the virtual port in the Alias field. **4.** Click in the GigaSMART Groups field, and select the GigaSMART Group created in Task 2. **5.** Click Save.
5	Create a first level map that directs GTP traffic from physical network ports to the virtual port you created in the previous task. **NOTE:** In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. o Enter an alias o Select First Level for Type o Select By Rule for Subtype o Select a network port for the Source o Select the virtual port configured int Task 4 for the Destination **4.** Create Rule 1. **a.** Click Add a Rule **b.** Select Pass. **c.** Select Bi Directional. **d.** Select Port Source **e.** Set the source to 2123 **6.** Create Rule 2. **a.** Click Add Rule **b.** Select Pass. **c.** Select Bi Directional **d.** Select Port Source **e.** Set the source to 2153 **6.** Create Rule 3. **a.** Click Add Rule **b.** Select Pass. **c.** Select IPv4 Fragmentation **d.** Set Value to allFragNoFirst **5.** Click Save.
6	Create a second level map that takes	**1.** Select Maps > Maps > Maps

Task	Description	UI Steps
	traffic from the virtual port, applies the GigaSMART operation, matches IMEIs specified by the flow rule, and sends matching traffic to a tool port.	**2.** Click New. **3.** Configure the map. o Enter an alias o Select Second Level for Type o Select Flow Filter for Subtype o Select a the virtual port configured in Task for the Source o Select the a tool port configured in Task 1 for the Destination o Select the GigaSMART Operation created in Task 3 from the GSOP list. **4.** Create a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select GTP IMSI. **d.** Enter * in the IMSEI field and set Version to V1 **5.** Click Save.
7	Create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, matches IMEIs specified by the flow rule, and sends matching traffic to another tool port.	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. o Enter an alias o Select Second Level for Type o Select Flow Filter for Subtype o Select a the virtual port configured in Task for the Source o Select the second tool port configured in Task 1 for the Destination o Select the GigaSMART Operation created in Task 3 from the GSOP list. **4.** Create a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select GTP IMSEI. **d.** Enter * in the IMSI field and set Version to V2 **5.** Click Save.
8	Display the configuration for Example 2.	To display the configuration for the GigaSMART Group: **1.** From the device view, select GigaSMART > GigaSMART Groups > Maps. **2.** Click on the alias for the GigaSMART Group to display the Quick View.

Task	Description	UI Steps
		To display the configuration for the GigaSMART Group: **1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations **2.** Click on the alias for the GigaSMART Operation to display the Quick View. To display the configuration for the maps: **1.** Select Maps > Maps > Maps. **2.** Click on the map alias to display the Quick View for the map.

Example 3: Same Subscriber, Filter on Different Versions

In this example, traffic from the same subscriber is forwarded to two different load balancing groups based on version. GTP version 1 traffic is sent to one load balancing group and GTP version 2 traffic is sent to another load balancing group.

Task	Description	UI Steps
1	Configure one network and multiple tool type of ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** In the Quick View Editor set one port to Network and multiple ports to Tool. For example, set 1/2/g1 as a Network port and ports 1/2/g5 through 1/2/g9 as Tool ports. **4.** Select Enable on each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART Groups.** **2.** Click New. **3.** Type a name for the GigaSMART Group in the Alias field. **4.** Click in the Port List field and select an engine port. **5.** Click Save.
3	Configure the GigaSMART operation and assign it to the GigaSMART group.	**1.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations > GigaSMART**

Task	Description	UI Steps
		Operations. **2.** Type a name for the GigaSMART Group in the Alias field. **3.** Click in the GigaSMART Operations field and select Flow Filtering. **4.** Click in the GigaSMART Operations filed and select Load Balancing. **5.** Configure the Load Balancing as follows:. o Select Stateful o Select GTP for Type o Selecting Hashing o Select IMSI **6.** Click Save.
4	Configure a virtual port and assign it to the GigaSMART group.	**1.** From the left navigation pane, go to **System > GigaSMART >Virtual Ports.** **2.** Click New. **3.** Type a name for the virtual port in the Alias field. **4.** Click in the GigaSMART Groups field, and select the GigaSMART Group created in Task 2. **5.** Click Save.
5	Create two port groups (one for version 1 traffic and one for version 2 traffic) and enable load balancing on the port groups.	**1.** Select Ports > Port Groups > All Port Groups. **2.** Click New. **3.** Type a name for the GigaSMART Group in the Alias field. For example, pglbv1. **4.** Select SMART Load Balancing to enable load balancing **5.** Click in the Ports field and select half the tool ports configured in Task 1. **6.** Repeat steps 2 through 4, creating a second port group with the other ports configured in Task 1
6	Create an ingress (first level) map.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. o Enter an alias. For example, map1_1. o Select First Level for Type o Select By Rule for Subtype o Select a network port for the Source o Select the virtual port configured in Task 4 for the Destination **4.** Click Add a Rule to create a rule.

Task	Description	UI Steps
		o Select pass o Select MAC Destination o Enter a MAC address. For example, **00:a0:d1:e1:02:01** o Enter a MAC mask. For example, 0000.0000.0000 **5.** Click Save.
7	Create a second level map for version1.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. o Enter an alias. For example, map2_1. o Select Second Level for Type o Select Flow Filter for Subtype o Select the virtual port configured in Task 4 for the Source o Select the first port group for Destination o Select the GigaSMART Operation configured in Task 3 from the GSOP list. **4.** Click Add a Rule to create a rule. o Select Pass o Select GTP IMSI o Enter * in the IMSI field o Select **V1** for Version **5.** Click Save.
8	Create another second level map for version2.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. o Enter an alias. For example, map2_2. o Select Second Level for Type o Select Flow Filter for Subtype o Select the virtual port configured in Task 4 for the Source o Select the second port group for Destination o Select the GigaSMART Operation configured in Task 3 from the GSOP list. **4.** Click Add a Rule to create a rule. o Select Pass o Select GTP IMSI o Enter * in the IMSI field o Select **V2** for Version **5.** Click Save.

Example 4: Same Subscriber, Filter on Different Interfaces

In this example, traffic from the same subscriber is forwarded to two different load balancing groups based on interface. In this example, VLANs 1601 and 1602 are from S5/S8 interface and VLANs 1611 and 1612 are from S11/S1-U interface. The first level maps split the VLAN traffic to different virtual ports. The second level maps send the traffic to different load balancing groups.

Task	Description	UI Steps
1	Configure one network and multiple tool type of ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** In the Quick View Editor set one port to Network and multiple ports to Tool. For example, set 1/2/g1 as a Network port and ports 1/2/g5 through 1/2/g9 as Tool ports. **4.** Select Enable on each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type a name for the GigaSMART Group in the Alias field. **4.** Click in the Port List field and select an engine port. **5.** Click Save.
3	Configure the GigaSMART operation and assign it to the GigaSMART group.	**1.** From the device view, select GigaSMART > GigaSMART Operations > GigaSMART Operations. **2.** Type a name for the GigaSMART Group in the Alias field. **3.** Click in the GigaSMART Operations field and select Flow Filtering. **4.** Click in the GigaSMART Operations filed and select Load Balancing. **5.** Configure Load Balancing as follows: ▪ Select Stateful ▪ Select GTP for Type ▪ Selecting Hashing ▪ Select IMSI **6.** Click Save.

Task	Description	UI Steps
4	Configure virtual ports and associate them with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual Ports. **2.** Click New. **3.** Type a name for the virtual port in the Alias field. **4.** Click in the GigaSMART Groups field, and select the GigaSMART Group created in Task 2. **5.** Click Save. **6.** Repeat steps 1 through 5 to create a second virtual port.
5	Repeat task 4 to create another virtual port.	
6	Create two port groups (one for version 1 traffic and one for version 2 traffic) and enable load balancing on the port groups.	**1.** Select Ports > Port Groups > All Port Groups. **2.** Click New. **3.** Type a name for the GigaSMART Group in the Alias field. For example, pglbv1. **4.** Select SMART Load Balancing to enable load balancing **5.** Click in the Ports field and select half the tool ports configured in Task 1. **6.** Repeat steps 2 through 4, creating a second port group with the other ports configured in Task 1
7	Create a first level map for vport1.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter an alias. For example, map1_1. ▪ Select First Level for Type ▪ Select By Rule for Subtype ▪ Select a network port for the Source ▪ Select the first virtual port configured in Task 4 for the Destination **4.** Click Add a Rule to create a rule. ▪ Select Pass ▪ Select VLAN ▪ Enter 1601 for the Min value ▪ Enter 1602 for the Max value **5.** Click Save.
8	Create another first level map for	**1.** Select Maps > Maps > Maps.

Task	Description	UI Steps
	vport2.	**2.** Click New. **3.** Configure the map. ▪ Enter an alias. For example, map1_2. ▪ Select First Level for Type ▪ Select By Rule for Subtype ▪ Select a network port for the Source ▪ Select the second virtual port configured in Task 5 for the Destination **4.** Click Add a Rule to create a rule. ▪ Select Pass ▪ Select VLAN ▪ Enter 1611 for the Min value ▪ Enter 1611 for the Max value **5.** Click Save.
9	Create a second level map using vport1.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter an alias. For example, map2_1. ▪ Select Second Level for Type ▪ Select Flow Filter for Subtype ▪ Select the virtual port configured in Task 4 for the Source ▪ Select the first port group for Destination ▪ Select the GigaSMART Operation configured in Task 3 from the GSOP list. **4.** Click Add a Rule to create a rule. ▪ Select Pass ▪ Select GTP IMSI ▪ Enter * in the IMSI field **5.** Click Save.
10	Create another second level map using vport2.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter an alias. For example, map2_2. ▪ Select Second Level for Type

Task	Description	UI Steps
		▪ Select Flow Filter for Subtype ▪ Select the second virtual port configured in Task 5 for the Source ▪ Select the second port group for Destination ▪ Select the GigaSMART Operation configured in Task 3 from the GSOP list. **4.** Click Add a Rule to create a rule. ▪ Select Pass ▪ Select GTP IMSI ▪ Enter * in the IMSI field **5.** Click Save.

Example 5: EPC Filtering in Scenario1

In this example, traffic for all subscribers on interfaces S11/S1-U and Gn/Gp is sent to the same load balancing group. All other traffic is dropped.

Task	Description	UI Steps
1	Configure one network and two tool type of ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** In the Quick View Editor set one port to Network and two ports to Tool. For example, set 1/2/g1 as a Network port and ports 1/2/g5 and 1/2/g6 as Tool ports. **4.** Select Enable on each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type a name for the GigaSMART Group in the Alias field. **4.** Click in the Port List field and select an engine port. **5.** Click Save.

Task	Description	UI Steps
3	Configure the GigaSMART operation and assign it to the GigaSMART group.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **2.** Type a name for the GigaSMART Group in the Alias field. **3.** Click in the GigaSMART Operations field and select Flow Filtering. **4.** Click in the GigaSMART Operations (GSOP) field and select Load Balancing. **5.** Configure the Load Balancing as follows. ▪ Select Stateful ▪ Select GTP for Type ▪ Selecting Hashing ▪ Select IMSI **6.** Click Save.
4	Configure a virtual port and assign it to the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual Ports. **2.** Click New. **3.** Type a name for the virtual port in the Alias field. **4.** Click in the GigaSMART Groups field, and select the GigaSMART Group created in Task 2. **5.** Click Save.
5	Create a port group enable load balancing on the port group.	**1.** Select Ports > Port Groups > All Port Groups. **2.** Click New. **3.** Type a name for the GigaSMART Group in the Alias field. For example, pglbv1. **4.** Select SMART Load Balancing to enable load balancing **5.** Click Save.
6	Create an ingress (first level) map. **Note**: In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. ▪ Enter an alias ▪ Select First Level for Type ▪ Select By Rule for Subtype ▪ Select a network port for the Source ▪ Select the virtual port configured int Task 4 for the Destination

Task	Description	UI Steps
		4. Create Rule 1. **a.** Click Add a Rule **b.** Select Pass. **c.** Select Bi Directional. **d.** Select Port Source **e.** Set the source to 2123 **6.** Create Rule 2. **a.** Click Add Rule **b.** Select Pass. **c.** Select Bi Directional **d.** Select Port Source **e.** Set the source to 2152 **6.** Click Save.
7	Create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, matches IMSIs specified by the flow rules, and sends matching traffic to physical tool ports.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter an alias. ▪ Select Second Level for Type ▪ Select Flow Filter for Subtype ▪ Select the virtual port configured in Task 4 for the Source ▪ Select the port group for Destination ▪ Select the GigaSMART Operation configured in Task 3 from the GSOP list. **4.** Create the rule 1. **a.** Click Add a Rule **b.** Select Pass **c.** Select GTP IMSI **d.** Enter * in the IMSI field **e.** Select Interface and set it to Gg/Gp **6.** Create rule 2. **a.** Click Add a Rule **b.** Select Pass **c.** Select GTP IMSI **d.** Enter * in the IMSI field **e.** Select Interface and set it to S11/S1U **6.** Click Save.

Example 6: EPC Filtering in Scenario2

In this example, traffic for all subscribers from all interfaces except S5/S8 is sent to the same load balancing group. Traffic from the S5/S8 interface is dropped.

Step	Description	UI Steps
1	Configure one network and two tool type of ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** In the Quick View Editor set one port to Network and two ports to Tool. For example, set 1/2/g1 as a Network port and ports 1/2/g5 and 1/2/g6 as Tool ports. **4.** Select Enable on each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type a name for the GigaSMART Group in the Alias field. **4.** Click in the Port List field and select an engine port. **5.** Click Save.
3	Configure the GigaSMART operation and assign it to the GigaSMART group.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **2.** Type a name for the GigaSMART Group in the Alias field. **3.** Click in the GigaSMART Operations (GSOP) field and select Flow Filtering. **4.** Click in the GigaSMART Operations (GSOP) field and select Load Balancing. **5.** Configure Load Balancing as follows: ▪ Select Stateful ▪ Select GTP for Type ▪ Selecting Hashing ▪ Select IMSI **6.** Click Save.
4	Configure a virtual port and assign it to the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual Ports.

Step	Description	UI Steps
		2. Click New. **3.** Type a name for the virtual port in the Alias field. **4.** Click in the GigaSMART Groups field, and select the GigaSMART Group created in Task 2. **5.** Click Save.
5	Create a port group and enable load balancing on the port group.	**1.** Select Ports > Port Groups > All Port Groups. **2.** Click New. **3.** Type a name for the GigaSMART Group in the Alias field. For example, pglbv1. **4.** Select SMART Load Balancing to enable load balancing **5.** Click Save.
6	Create an ingress (first level) map. **Note**: In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. ▪ Enter an alias ▪ Select First Level for Type ▪ Select By Rule for Subtype ▪ Select a network port for the Source ▪ Select the virtual port configured int Task 4 for the Destination **4.** Create Rule 1. **a.** Click Add a Rule **b.** Select Pass. **c.** Select Bi Directional. **d.** Select Port Source **e.** Set the source to 2123 **6.** Create Rule 2. **a.** Click Add Rule **b.** Select Pass. **c.** Select Bi Directional **d.** Select Port Source **e.** Set the source to 2152 **6.** Click Save.
7	Create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, matches IMSIs specified by the flow rules, and sends	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map.

Step	Description	UI Steps
	matching traffic to physical tool ports.	▪ Enter an alias. ▪ Select Second Level for Type ▪ Select Flow Filter for Subtype ▪ Select the virtual port configured in Task 4 for the Source ▪ Select the port group for Destination ▪ Select the GigaSMART Operation configured in Task 3 from the GSOP list. **4.** Create the rule 1. **a.** Click Add a Rule **b.** Select drop **c.** Select GTP IMSI **d.** Enter * in the IMSI field **e.** Select Interface and set it to S5/S8 **6.** Create rule 2. **a.** Click Add a Rule **b.** Select Pass **c.** Select GTP IMSI **d.** Enter * in the IMSI field **5.** Click Save.

GigaSMART Rotational Sampling

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

Rotational Sampling samples sessions by International Mobile Subscriber ID (IMSI), using a rotating or sliding IMSI range, that rotates automatically on a configurable interval, without disrupting the treatment of existing sessions. It rotates through the entire pool of subscribers over time, such that all subscribers can be sampled in.

The IMSI range size is configurable, the length of time for each period or window is configurable, and as is the length of time for each period or window. Additionally, the amount it shifts or slides every period is also configurable.

This capability supports both network security and performance monitoring. Sampling should account for all active subscribers and be deterministic, so it is always evident whether a subscriber is included.

In Rotational Sampling, the pool of IMSIs is updated at regular intervals, ensuring that all subscribers have an opportunity to participate in the sampling.

Rotational Sampling is now enhanced to support two configuration options –

1. Entire Session Sampling for Mobile Subscribers
2. Dynamic Session Sampling for Fixed Wireless Access Subscribers.

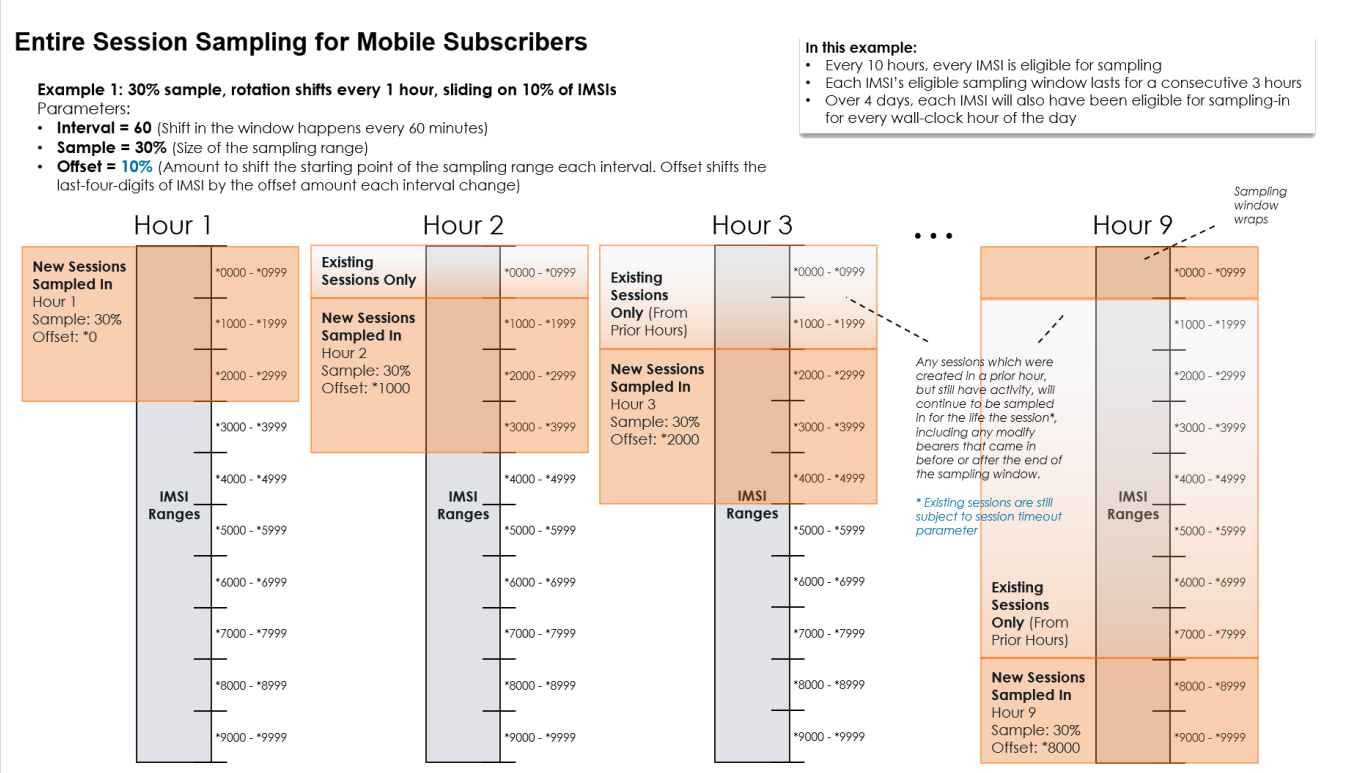
Entire session sampling for Mobile subscribers

The rotations do not disrupt treatment of existing sessions. Entire sessions are sampled in or out based upon the occurrence of the **"Create Session Requests in GTP-C"** or **"Session Establishment Requests in PFCP"**. Once a session is sampled in, it will remain sampled in for the life of the session, even if the session's IMSI later rotates out of the sampling range.

If the subscriber does not have any create session activity during their rotation window, then nothing will be forwarded for that subscriber in that round. The following parameters are considered while configuring an Entire Session Sampling for Mobile Subscribers:

- **Interval**: Defines how often the sampling window shifts. For example, an interval of 60 means the eligible IMSI range changes every 60 minutes.
- **Sample**: Specifies the size of the IMSI range (in percentage) that is selected for session sampling during each interval. For example, a sample of 30% means 30% of IMSIs are eligible within the active window.
- **Offset**: Determines how much the starting point of the IMSI sampling range shifts at every interval. The offset spreads IMSIs more evenly across intervals. For example, with an offset of 10%, the window moves forward by 10% of the IMSI range each time the interval changes.

The following diagram illustrates an example of an Entire Session Sampling for Mobile Subscribers:



Dynamic Session Sampling for Fixed Wireless Access Subscribers

Dynamic Session Sampling enables the system to periodically re-evaluate which subscriber sessions are actively being sampled. This feature is specifically designed for Fixed Wireless Access subscribers, as their sessions often last for extended periods and could otherwise remain sampled indefinitely.

Sampling Window

The subscriber IMSIs are divided into sampling windows, which rotate at configurable intervals (for example, every 12 hours). Only sessions with IMSIs in the current window are eligible for sampling. For Home Internet FWA subscribers, a 12-hour rotation is recommended to maintain balance and coverage. This approach prevents the same IMSIs from being repeatedly sampled and ensures broader data diversity.

Enabling Dynamic Sampling

Dynamic Sampling is controlled through a map rule flag called *Periodic Recalculation*.

- **When enabled:** It re-evaluates all ongoing sessions at each rotation. If a session's IMSI no longer qualifies for the current window, that session is excluded until its IMSI re-enters a future window.
- **When disabled**: Once a session is accepted, it stays sampled until the session ends, following the traditional "whole-session" sampling behavior.

This flexibility allows operators to choose between continuous fairness (enabled) or stable, session-based sampling (disabled).

Session Handling

At the end of every sampling interval, the system rotates the active IMSI window. A map lookup is performed for ongoing sessions whenever control or user packets are received. Only sessions that match the current IMSI window and the *Periodic Recalculation* setting continue to be sampled, while the rest are excluded until they fall into a later window. This ensures that only sessions aligned with the latest sampling policy remain active.

Offset Configuration

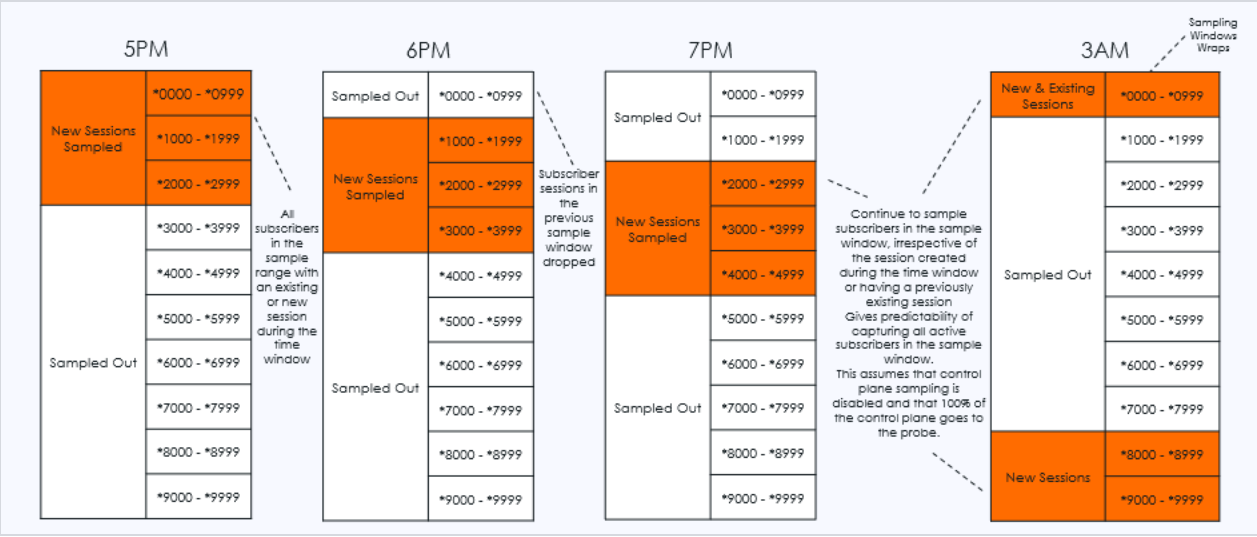
It is recommended to configure the offset value less than the sampling percentage to ensure even distribution of IMSI with maximum correlation. When the offset is less than the sampling percentage, each IMSI is sampled multiple times across different intervals, preventing missed or over-sampled entries. This overlap creates a rolling window effect that smooths out the sampling distribution and helps avoid gaps in the data.

Conversely, if the offset is equal to or greater than the sampling percentage, some IMSIs may be skipped or sampled too infrequently.

For example, with a 25% sampling percentage and a 25% offset, each IMSI would only be sampled once every four intervals, leading to potential data gaps.

NOTE: It is advisable to disable GTP-control sampling during the configuration of Dynamic Session Sampling.

The following diagram illustrates an example of Dynamic Session Sampling:



Supported Platforms Compatibility

The Rotational Sampling feature is supports Entire session sampling for Mobile Subscribers in Non-CUPS LTE (Non-CUPS GTP application) and UPN (Non-Overlap and Overlap). The table below outlines the supported compatibility.

Platforms	Non-CUPS		UPN	
	Non - Overlap	Overlap	Non - Overlap	Overlap
GigaVUE-HC3 Gen 3	✓	✓	✓	✓
GigaVUE-HC3 Gen 2	✓	×	×	×
GigaVUE-HC1-Plus Front	✓	✓	×	×
GigaVUE-HC1-Plus Rear	✓	✓	×	×

The Rotational Sampling feature supports Dynamic session sampling for Fixed Wireless Access Subscribers in Non-CUPS LTE (Non-CUPS GTP application) and UPN (Non-Overlap and Overlap). The table below outlines the supported compatibility.

Platforms	Non-CUPS		UPN	
	Non - Overlap	Overlap	Non - Overlap	Overlap
GigaVUE-HC3 Gen 3	✓	✓	✓	✓
GigaVUE-HC1-Plus Front	✓	✓	×	×
GigaVUE-HC1-Plus Rear	✓	✓	×	×

Configuring Rotational Sampling in GigaVUE-FM

GigaVUE-FM allows you to configure Rotational Sampling in the second level map of GTP Flow Sample and GTP Flow Overlap.

NOTE: The same configuration applies to UPN, with *3GPP node role* for the GigaSMART group as **user** and **mode** as *standalone*.

To configure Rotational Sampling in second level maps, do the following:

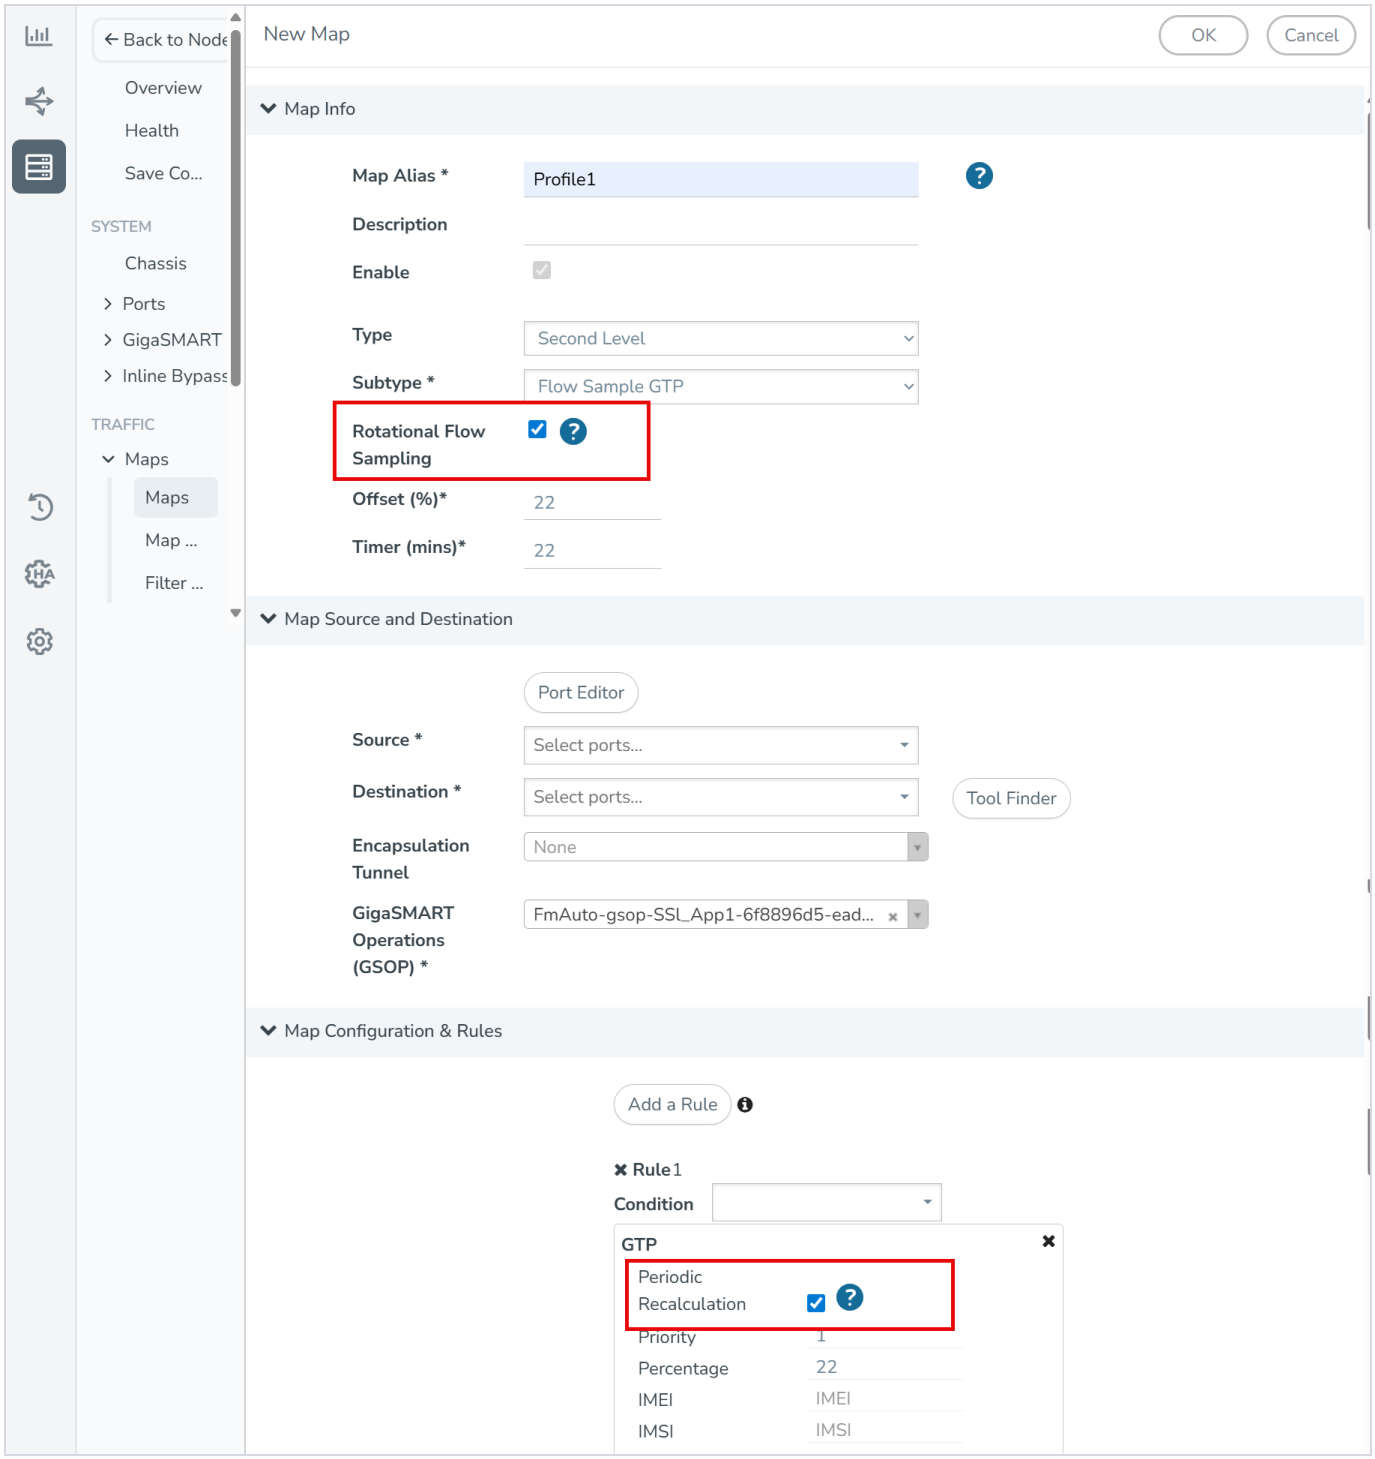
3. On the left navigation pane, click on  .
4. Go to **Physical** > **Nodes** > **Cluster ID**.
5. Go to **Traffic** > **Maps**.
6. Click **New**.

7. Enter the Map information:
   a. Enter an alias for the map.
   b. Enter a description about the map.
   c. Select the Type as **Second level**.
   d. Select the Subtype as any one of the following to configure Rotational Sampling:
      - Flow Sample GTP
      - Flow Sample Overlap
   e. Enable the **Rotational Flow Sampling** check-box.

   NOTE: In GigaVUE-FM, enabling Rotational flow Sampling means it selects the flow sampling type (fstype) as "rotational." Disabling the check-box means the flow sampling type is "default." The "fstype" is only applicable for 'secondLevel/flowSample' and 'secondLevel/flowSampleOverlap' map types.

   f. Enter the percentage of sessions to be sampled in the given interval in the **Offset** range.
   g. Enter the **Timer** range in minutes. The value ranges from 15 to 45000 minutes.
8. Refer to Create a new mapsection to complete the configuration for rotational sampling.
9. To enable Dynamic Session sampling , under *Map Configuration and Rules* , select the **Periodic Recalculation** flag.

   NOTE: In GigaVUE-FM , you can edit only the *Periodic Recalculation* flag in existing map rules.



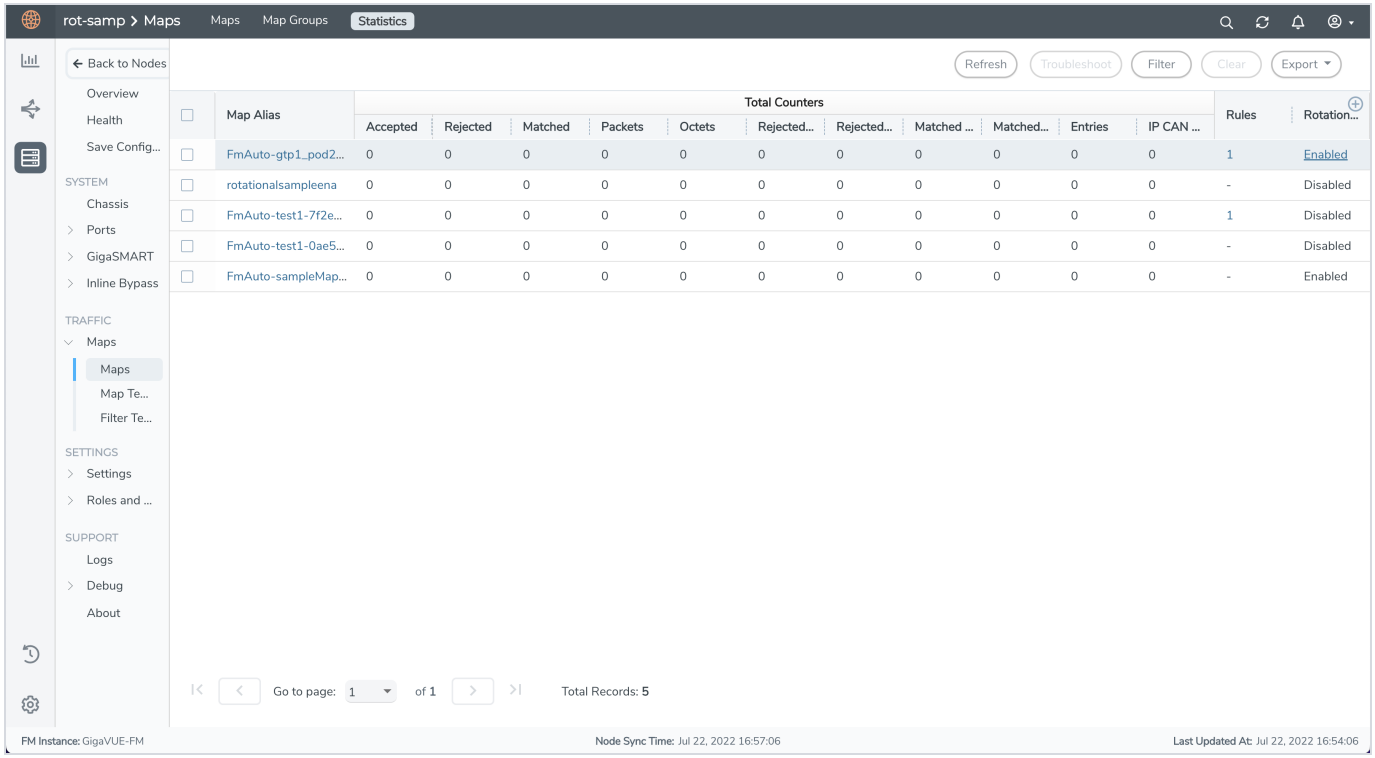
You can also configure Rotational Sampling through Ansible. Refer to *Installation and Configuration of Subscriber Intelligence Solution using Ansible* to know more.

You can also configure rotational sampling through GigaVUE-OS CLI. Refer to the **map** and **show map fs-stats all commands** in the GigaVUE-OS CLI Reference Guide to know more.

Viewing the Configuration of Rotational Sampling for non-CUPS LTE and UPN

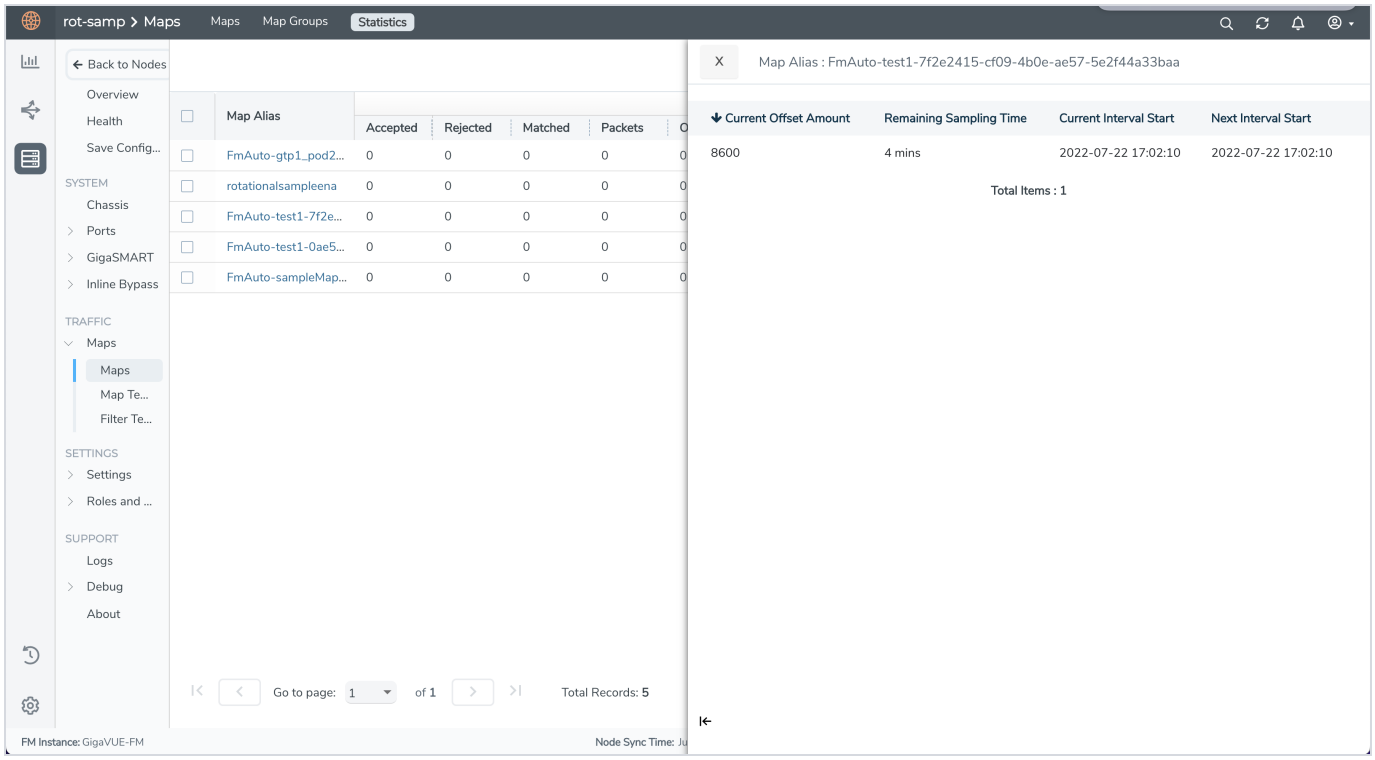
After configuring rotational sampling, GigaVUE-FM allows you to view the configuration for non-CUPS LTE and UPN. To view the configuration in GigaVUE-FM, follow these steps:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**.
2. In the **Map Statistics** page you can view the column **Rotational Sampling** which can be either **Enabled** or **Disabled**. If the **Rotational Sampling** is **Enabled**, it would provide a quick link view.



The quick view contains the following columns:

- Current Offset Amount
- Remaining Sampling Time
- Current Interval Start
- Next Interval Start



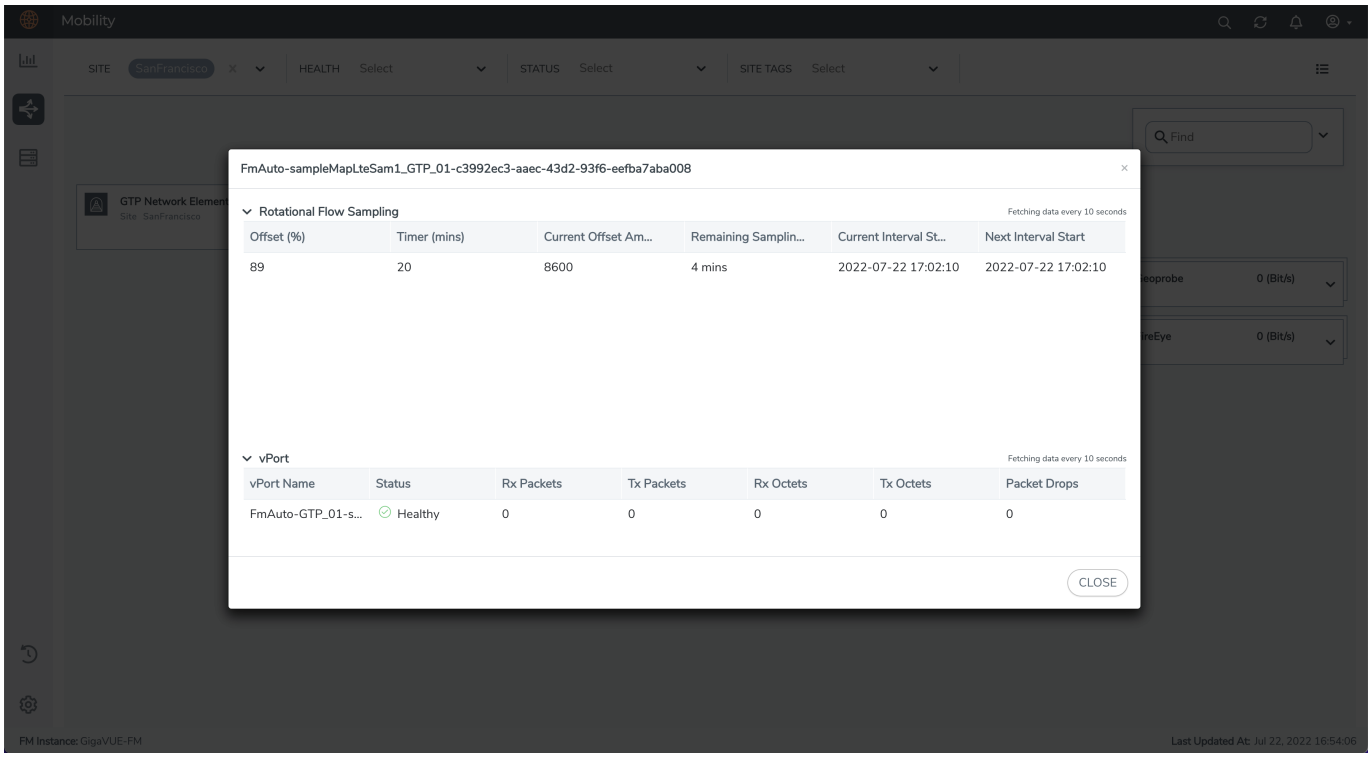
The Rotational Flow Sampling is also supported in the legacy maps pages in GigaVUE-FM. You can create and edit the "fstype" properties from the legacy Maps pages in GigaVUE-FM GUI.

View a Rotational Sampling configuration in Mobility Solution

If you have configured Rotational Sampling through Ansible, then to view the configuration in GigaVUE-FM, do the following:

1. On the left navigation pane, click on  select **Physical> Orchestrated Flows > Mobility,** and then select the required cluster or node ID.
2. Select the required site from the **Site** drop-down list box.
3. From the **Rules to Destination Tools** block, select the drop-down, and click on the ellipses named as **Details**.
4. A Details-Rules to Geoprobe page appear. In the left-navigation page, click on the **Second Level Maps**.
5. Click on the **Show Details** for the required **Fabric Map Name**.
6. A page with details of Flow Sampling appears. In this page, you can view the details of rotational sampling such as **Offset**, **Timer**, **CurrentOffset Amount**, **RemainingSampling**.

NOTE: When configuring rotational sampling in UPN, the *3GPP node role* for the GigaSMART group is **user** and **mode** is *standalone*.



GigaSMART 4G RAN Correlation

Required License:

A Radio Access Network (RAN) is part of a telecommunications system that connects individual radio devices to the core mobile network through radio connections.

RAN uses GTP correlation. It exposes RAN data fields from control plane (S11), for example:

- E-UTRAN Cell Global Identifier (ECGI)
- Tracking Area Identity (TAI)
- Tracking Area Code (TAC)

The above identifiers are extracted from the User location Information (ULI) and sends them to flow sampling and whitelisting.

The RAN correlation for 4G CUPS flow sampling filters require the ECGI, TAC and PLMN-ID values.

For 4G RAN based whitelisting requires ECGI includes the PLMN Id and the E-UTRAN Cell ID.

RAN data fields are available for filtering. GigaVUE-FM enables you to configure RAN based correlation for both 4G CUPS and 5G CUPS in the Flow Sample and Whitelist Maps. In CUPS solution, you can configure the RAN Correlation properties while creating and editing the properties in Flow Sample and Whitelist maps in Global policies for both CPN and UPN.

Rules and Notes

Keep in mind the following rules and notes when you work with the RAN Correlation feature:

- RAN correlation is not supported for the 3G.
- You must configure separate flow sampling map for 3G sampling parameters and 4G RAN sampling parameters to support 3G correlation and 4G RAN correlation in the same GSgroup.

GigaSMART SIP/RTP Correlation

Required Licenses: SIP/RTP Correlation (and FlowVUE for session-aware load balancing for RTP)

Supported Devices : GigaVUE-HC3 Gen 2 and GigaVUE-HC1 Plus.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

Session Initiation Protocol (SIP) is the dominant method to initiate, maintain, modify, and terminate voice calls in service provider and enterprise networks. Real-time Transport Protocol (RTP) is used to manage the real-time transmission of voice payload across the same networks. Visibility into a subscriber's voice traffic requires the ability to understand the subscriber attributes and stateful information contained within SIP to correlate subscriber-specific RTP traffic so that monitoring tools can achieve an accurate view of the subscriber's traffic on the network.

The GigaSMART SIP/RTP correlation application correlates the subscriber-specific attributes and the endpoint identifiers of the RTP streams where the session is carried,

as well as other SIP-related attributes that are exchanged as part of the control sessions. Use SIP/RTP correlation to leverage a subscriber-aware monitoring policy on Gigamon's Deep Observability Pipeline and to optimize current tool infrastructure investments by providing only relevant data to tools while increasing visibility into subscriber traffic. This helps improve QoE and performance. Carriers gain access to the subscriber's traffic by reliably correlating and passing all the identified subscriber's control and data sessions to the analytics/monitoring probes and/or billing subsystems for an accurate view of the subscriber's sessions.

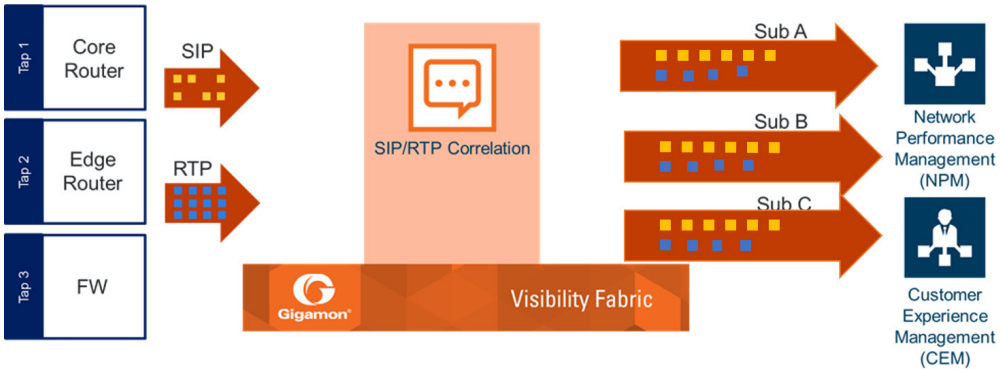


Figure 5 *SIP/RTP Correlation*

SIP is a signaling protocol for VoIP and VoLTE call initiation. It is implemented with RTP to control the payload. GigaSMART SIP/RTP correlation provides customers visibility into VoIP, VoLTE, SMSoLTE, and RCS traffic, and allows them to filter and forward traffic for a subscriber to the tools.

All SIP traffic is sent to all tool ports, as follows:

- all SIP packets sent to all ports within a load balancing port group.
- all SIP packets sent to all ports within a GigaStream.
- all SIP packets sent to all tool ports belonging to maps.

RTP traffic will be sampled and sent to the maps with the rules that match. RTP non-correlated traffic will be sent to the collector.

SIP/RTP correlation can be used for both enterprise and service providers, where ever

there is SIP/RTP traffic, such as in wire-line communications, wireless communications, and packet cable networks. This includes enterprise, IP Multimedia Subsystem (IMS), or fixed network implementations of SIP, as well as any media controlled by SIP, such as voice, text, or streaming media.

In addition, SIP/RTP correlation correlates SIP signaling and RTP payload for all sessions selected by a SIP User Agent (UA), a caller ID in a forward list, with flow sampling from 0 to 100%.

CallerID Tracking

The CallerID is tied to the Call-ID and remains constant for the duration of the session:

CallerID (A) <-----(Call-ID)-----> CalleeID (B)

CallerID tracking is based on the initial caller (A) and does not change until the Call-ID changes. Even if the callee, (B), sends a new INVITE during the SIP session, if that INVITE uses the same Call-ID, then the CallerID information in the SIP session display will still identify the initial caller, (A); it will not switch to reflect that (B) is now the caller.

Support for SIP, RTP, and RTCP

SIP/RTP correlation handles all SIP signaling and RTP/RTCP traffic, including RTCP control traffic, as well as that belonging to a call session.

The following is supported:

- Non-tunneled SIP, RTP, and RTCP correlation (all IMS interfaces)
- Tunneled SIP, RTP, and RTCP correlation (SIP/RTP/RTCP in GTP-U, through the GTP tunnel)

SIP/RTP Correlation Engine

When a packet containing SIP, RTP, or RTCP traffic is received, the SIP/RTP correlation engine looks up the session in the session table for load balancing ports and sampling maps or whitelist map. All SIP/RTP traffic with port or load balancing port group is forwarded based on the session table. The correlation engine load balancing keeps track of both the SIP session and the associated multiple RTP channels.

Each session identifies both sides of media streams (RTP) associated with the session. The SIP session has an aging timer that is configurable.

When a session matches one of the configured flow sampling rules, it is either accepted for sampling or rejected. If it is accepted, all packets belonging to that session are sent to the tool port. Otherwise, all packets belonging to the session are dropped.

Configure SIP/RTP Correlation Engine

To configure SIP/RTP Load balancing, follow these steps:

1. From the left navigation pane, go to **System > GigaSMART >GigaSMART Group** and then click **New**.
2. Specify an Alias in the **Alias** field.
3. Select Load Balance in **TCP Application Parameters**.
4. Specify the parameters. The following table explains the parameters for configuring the load balance:

TCP Application Parameters	Options
Application	• Broadcast - All the TCP handshake packets are broadcasted to the tool ports in all the maps. • Drop - All the unknown application packets of 5G are sent to the collector.
Load Balance	• Enable - Select Enable to perform any of the following: o SIP and RTP packets matching the flowsampling or whitelist maps are loadbalanced across the tool ports within a port-group. The SIP and RTP packets are loadbalanced to the same tool ports. o SIP and RTP packets matching the flowsampling or whitelist maps are sent to the tool port. o SIP and RTP packets matching the flowsampling or whitelist maps are sent to a tool GigaStream. The SIP and RTP packets may not be loadbalanced to the same tool ports within the GigaStream. • Disable - Select disable to perform any of the following: o all SIP packets sent to all ports within a load balancing port group o all SIP packets sent to all ports within a GigaStream o all SIP packets sent to all tool ports belonging to maps
TCP Control	• Broadcast - All the tcp handshake packets will be broadcasted to the tool ports in all the maps.. • Drop - All the TCP handshake packets will be sent to collector if configured

RTP traffic will be sampled and sent to the maps with the rules that match. RTP non-correlated traffic will be sent to the collector.

Only one SIP interface type is supported per engine, for example, S5. There is no mixing of interface types, such as S5 GTP-U with SGi.

SIP Whitelist

The SIP whitelist contains caller IDs, callee ID, the range for caller IDs, the range for callee IDs and the IP address. Each forward list entry in a file is a SIP caller ID, callee ID, caller ID range, callee ID range or IP address. The forward list can contain all types of entries.

forward list entries can be alphabetic, IP address, and numeric. For each entry, specify up to 64 alphanumeric characters. Some special characters are also supported.

You can manually add one entry at a time to a whitelist file, or you can upload files in.txt format. You must provide the whitelist caller ID range in numeric. You can also provide multipile entries to the forward list by ID range configuration. Each whitelist file can have up to 20,000 entries. One or more whitelist files can be fetched from a local directory or remote URL using HTTP or SCP.

On GigaVUE-HC3 nodes, the forward list database supports 1 million entries.

Multiple forward list databases can reside on a GigaVUE node, but only one forward list is applied to a GigaSMART group at a time.

Only one whitelist map is supported for a GigaSMART group. The GigaSMART operation does not have any rules for forward listing.

RTP Flow Sampling

FlowVUE is used for session-aware (stateful) load balancing and forward listing with sampling. Only RTP traffic will be sampled. There is no sampling of SIP traffic.

Up to five flow sample maps per GigaSMART group are supported. Each flow sample map can have 20 rules. Use rules to filter on caller ID. The rules support both alphabetic and numeric characters, up to 64 characters. Some special characters are also supported, such as wildcard characters.

Sampling is based on caller ID only (the from field).

Support for Sessions

The number of supported SIP and RTP sessions are as follows:

- GigaVUE-HC3—1 million SIP sessions and 4 million concurrent RTP sessions

Each SIP session can handle two RTP streams in both directions (bidirectional).

The number of supported TCP sessions are as follows:

- GigaVUE-HC3—2 million sessions

Support for IPv4 and IPv6

SIP is a text-based protocol, which is supported over UDP and TCP. The size of the

SIP message can vary greatly, so fragmentation and segmentation are common and are supported for tunneled SIP and non-tunneled SIP (IMS).

IPv4 and IPv6 are supported as follows:

- UDP Fragmentation—in-order packets, out-of-order packets
- TCP Segmentation—in-order packets, out-of-order packets

The following is not supported:

- GTP tunneled packets where the inner IP is fragmented
- IMS packets where the outer IP is fragmented

		UDP				TCP		
		Outer Frag	Inner Frag	In order	Out of order	Segmentation	In order	Out of order
GTP	IPv4	✓	×	✓	✓	✓	✓	✓
		✓	×	✓	✓	✓	✓	✓
	IPv6	✓	×	✓	✓	✓	✓	✓
		✓	×	✓	✓	✓	✓	✓
IMS	IPv4	N/A	✓	✓	✓	✓	✓	✓
		N/A	✓	✓	✓	✓	✓	✓
	IPv6	N/A	✓	✓	✓	✓	✓	✓
		N/A	✓	✓	✓	✓	✓	✓

Figure 6 *SIP/RTP UDP/TCP Support*

Support for Content Masking

SIP Common Presence and Instant Messaging (CPIM) content masking is supported, but only when the SIP transport is UDP.

The SIP method, MESSAGE, carried over UDP, might contain user-friendly, readable text messages. Use masking to replace these messages with x's, so they cannot be read.

NOTE: SIP/RTP correlation cannot mask text messages with a content type other than message/CPIM", such as plain text.

Behaviors of Some SIP Methods

The following are behaviors for some particular SIP methods:

- The SIP method, REGISTER, might not contain a user part. When there is no user part, it will be treated as a parse error.
- The SIP method, OPTIONS, (and response messages) might not contain a user part. When there is no user part, it will be treated as a parse error.

NOTE: SIP TCP packets with parse errors are not sent to collector. SIP TCP packets will be sent to the tool and incremented as parse errors in the session table stats.

SIP Whitelisting in a Cluster

The forward list (all whitelist files) reside on the leader of the cluster. The member nodes receive a copy of the forward list from the leader. Updates to the forward list are synchronized from the leader to the non-leaders. If a member node leaves the cluster and rejoins, its forward list will be resynchronized.

Use the cluster leader preference command to specify the highest preference for the leader, the second highest preference for the standby node, and lower preferences for the normal nodes in the cluster.

If there are GigaVUE TA Series nodes in the cluster, they will not receive a copy of the forward list.

Support for NAT

GigaSMART SIP Correlation engine correlates NAT/PAT enabled SIP/RTP packets. The correlation engine compares the top most "Via" address and the contact address of the device with the Layer 3 network address to find whether the device is configured with NAT.

To enable the NAT support, from the device view:

1. From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > SIP > SIP NAT**.
2. Enable the **SIP NAT** check box.

Not Supported by SIP/RTP Correlation

The following list is not currently supported by SIP/RTP correlation:

- encryption
- filtering based on Codecs
- SRVCC
- roaming
- SIP-I/SIP-T (supports SIP/RTP correlation based on SDP processsing even jf SIP carries SIP-I/SIP-T messages)
- forwarding of emergency calls to specific tools
- engine grouping. Only one engine is used for SIP/RTP correlation.

NOTE: SIP/RTP correlation and GTP correlation are not supported on the same GigaSMART engine port.

Display SIP/RTP Reports

To display SIP report, do the following:

1. From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups> Report**.
2. Select Group Type: **Flow SIP.**
3. From the device view, select GigaSMART Group: **gsg1**
4. Specify **Caller ID Pattern.**
5. Select Any. This return any pattern specified in the Caller ID Pattern field.

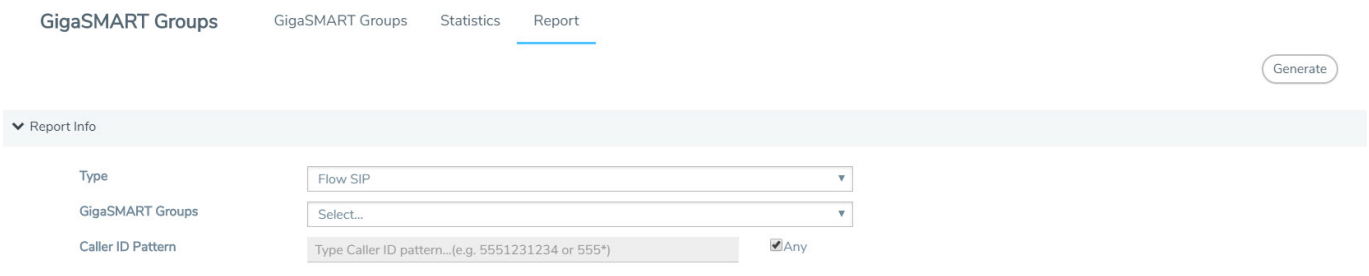


Figure 7 *Generate SIP Report*

6. Click the Generate button. The SIP Messages Report displays.



Figure 8 *SIP Report Page*

Display SIP Map Statistics

Map Statistics displays the counts sessions that matched a particular map.

For SIP Flow Sample maps, the counters show how many sessions matched the Caller-ID rule and were either accepted or rejected based on the sampling percentage.

For SIP Whitelist maps, the counter shows how many total entries are in the Whitelist and how many sessions matched those entries. .

To display SIP Map Statistics, do the following:

- On the left navigation pane, go to **Inventory** from **Traffic** select **Maps> Statistics.** The Statistics page displays a count of the rules that actually matched in a map.

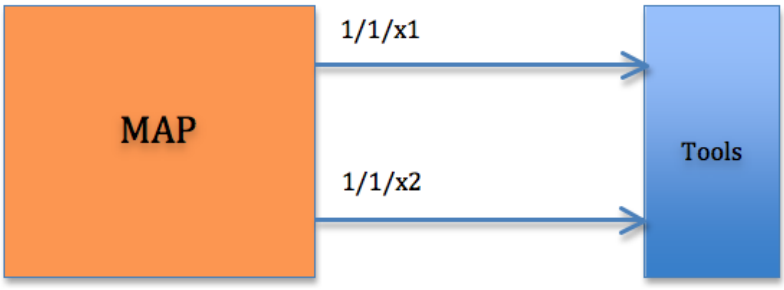


Figure 9 *SIP Map Statistics*

SIP/RTP Support for Tool Throttling

GigaSMART supports sampling/scaling based on fixed percentages, which remain in effect at all times, regardless of the tool port utilization. However, tool utilization may not be as efficient as not using fixed percentage for every use case. Staring with GigaSMART version 5.4 support for throttling sessions based on the traffic (pps) reaching a tool port is available. This feature helps avoid packets being drops during peak times by allowing users to adjust throttling start and stop levels.

The illustration below is a configuration and Intra Flow for Tool Port Throttling.



Admission Control

Each SIP session comes with RTP streams and each of the RTP stream uses a specific codec for information transfer. We can use this codec information to our advantage and do predictive analysis on how much pps would be generated for a given SIP session.

Based on the outcome of the pps for a given RTP stream codec, admission control module will check this value against the cumulative packet throughput on the destination tool-port to decide if the session will be Accepted or Rejected.

Example: Tool port 1/1/x1 is configured with a threshold of 3k pps.

Time	Port	Session	Codec (pps)	Cumulative pps	Throttle pps	Accepted
t0	1/1/x1	0		0	3000	
t1	1/1/x1	1	500	500	3000	Yes
t2	1/1/x1	2	1500	2000	3000	Yes
t3	1/1/x1	3	800	2800	3000	Yes
t4	1/1/x1	4	500	3300	3000	No

In software version 5.4 Tool port throttling applies only to SIP sessions for audio and only load balanced ports are supported in tool port throttling. Use case where there are tapping multiple interfaces using multiple engines, one SIP session can be throttled in one engine and not in another.

SIP/RTP Examples

This section provides information about:

- SIP/RTP Load Balancing Example
- SIP/RTP Minimum Configuration Example

SIP/RTP Load Balancing Example

This is a load balancing configuration example of SIP/RTP.

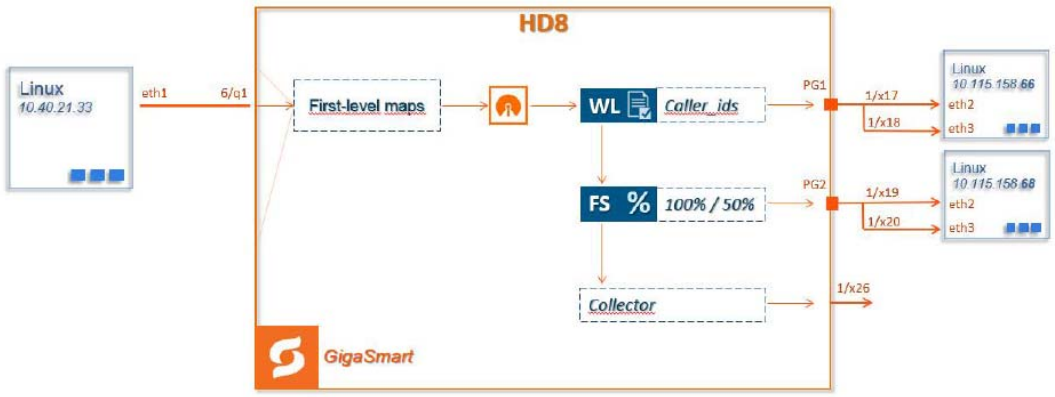


Figure 10 *SIP/RTP UDP/TCP Support*

SIP/RTP Minimum Configuration Example

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

Configure ports

To configure a GigaSMART group and associate it with a GigaSMART engine port do the following.

1. To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. **GigaSMART Groups.**
2. Click **New**.
3. Type an alias in the Alias field and enter **an engine port** in the Port List field.

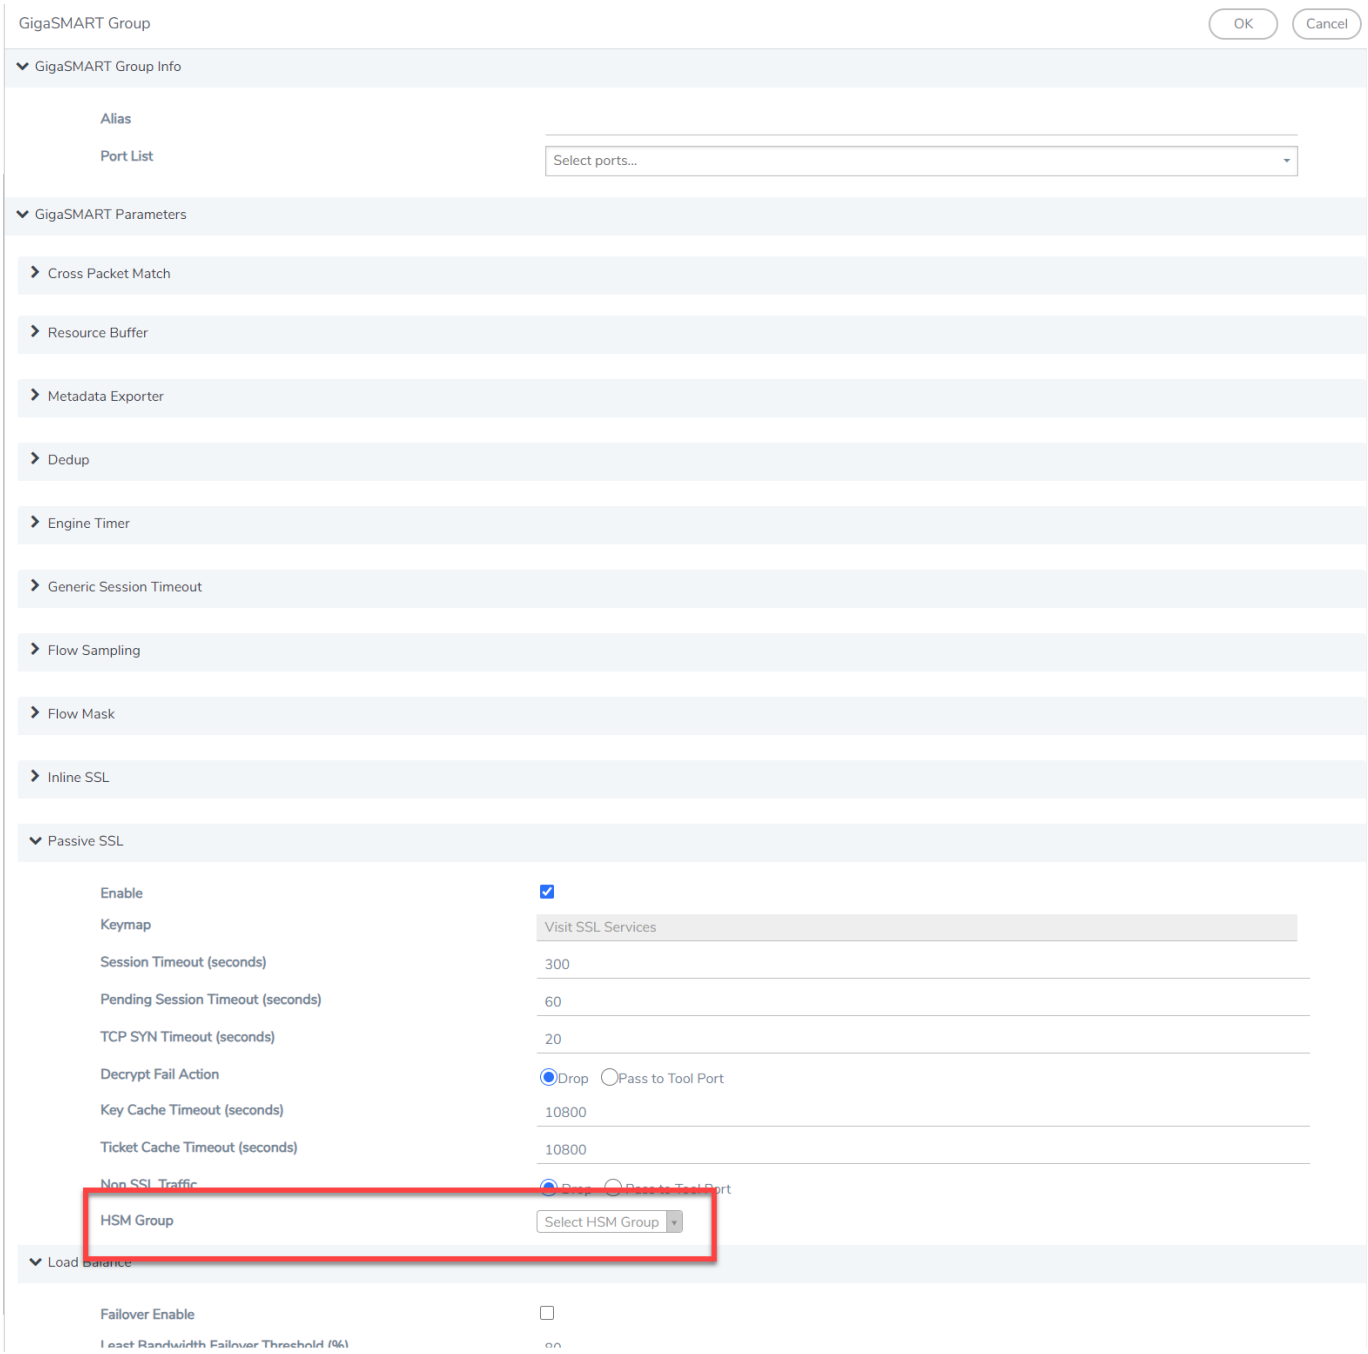


Figure 11 *GigaSMART Group Port Info*

4. Click **Save**.
5. Scroll drown the page to select **SIP Port** parameters.
6. Type parameters for **SIP Port and RTP Port.**
7. Enter parameters for **SIP Session Timeout** and the **SIP TCP Idle Timeout**.

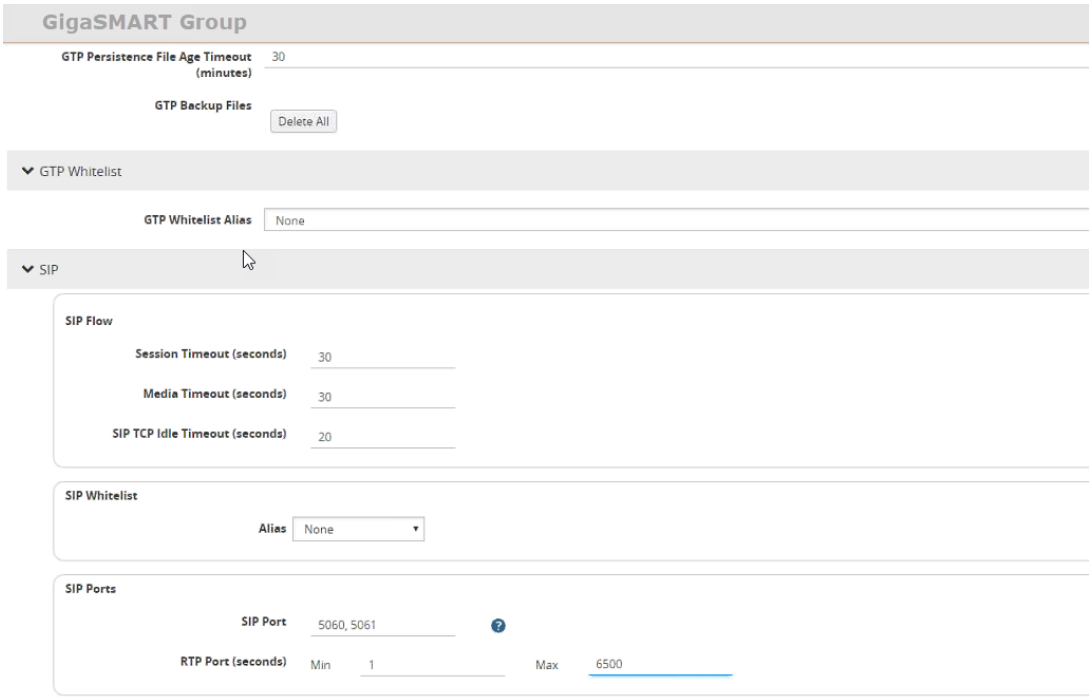


Figure 12 *SIP Port parameters*

8. Click **OK**.

Create Virtual Ports

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

To create virtual ports, do the following:

1. From the device view, select **GigaSMART > Virtual Ports**.
2. Click **New**.

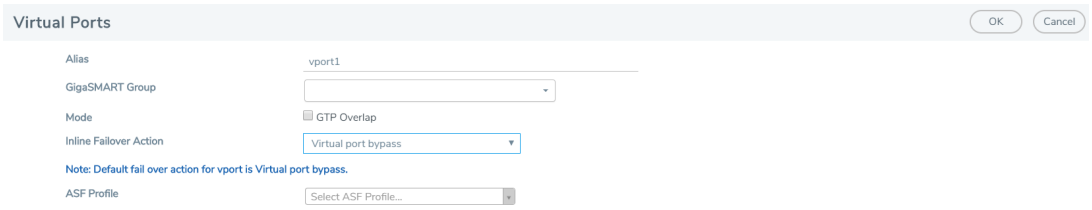


Figure 13 *Virtual Ports*

3. Enter an **alias** in the Alias field to identify the virtual port.
4. In the GigaSMART Groups field, select the GigaSMART Group configured in Step 1: of Configure a GigaSMART Group.

5. Click **Save**.

Configure GigaSMART Operation

Define a GigaSMART operation to enable SIP Flow Sampling. If combining Flow Sampling with Load Balancing GSOPs, make sure that you select both operations when creating the GigaSMART Operation.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

To configure the GigaSMART Operation, do the following:

1. From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations (GSOP).**
2. Click **New**. On the GigaSMART Operations page, do the following:

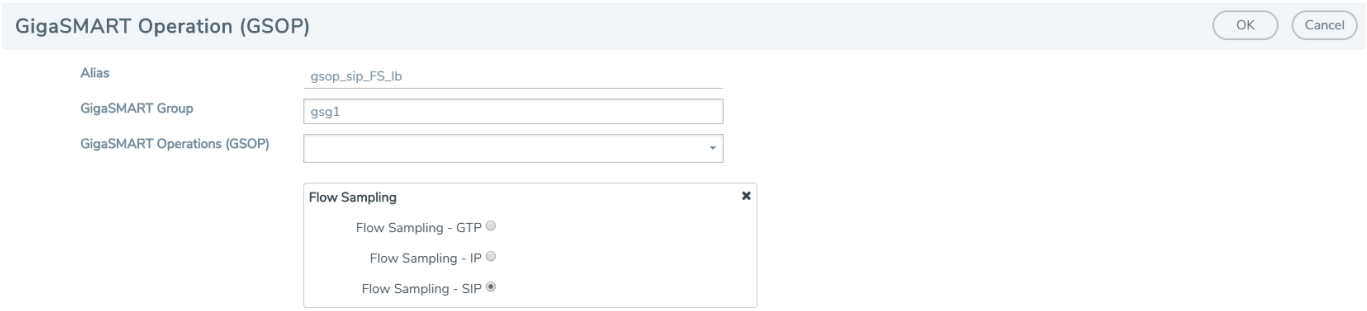


Figure 14 *GigaSMART Operations page*

3. In the Alias field, enter an **alias** to help identify this GSOP.
4. In the **GigaSMART Groups field**, select the **GigaSMARTgroup** configured in Step 1: Configure a GigaSMART Group.
5. In the GigaSMART Operations (GSOP) field, select **Flow Samplingfor SIP**.
6. Using the GSOP drop down list, select **Load Balancing** as the next GSOP operation.

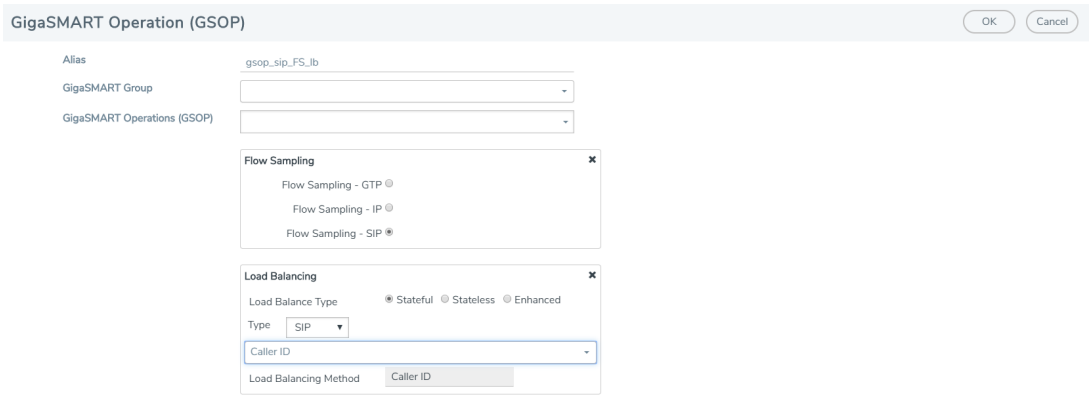


Figure 15 GigaSMART Operations - Load Balancing

Options:

- Stateful
- Stateless
- Enhanced

7. Select **Stateful**.
8. For Type, select **SIP** as the stateful application within a group of GigaSMART operations.
9. Select **Caller ID** as the Load Balancing Method.
10. Click **OK**.

Create first level map

1. On the left navigation pane, go to **Inventory** from **Traffic** select **Maps > Maps > Maps**.
2. Click **New**.
3. Type **map-level1** in the Alias field.
4. Select **First Level** for Type and **By Rule** for Subtype.
5. Select **port 1/1/x1** for the Source.

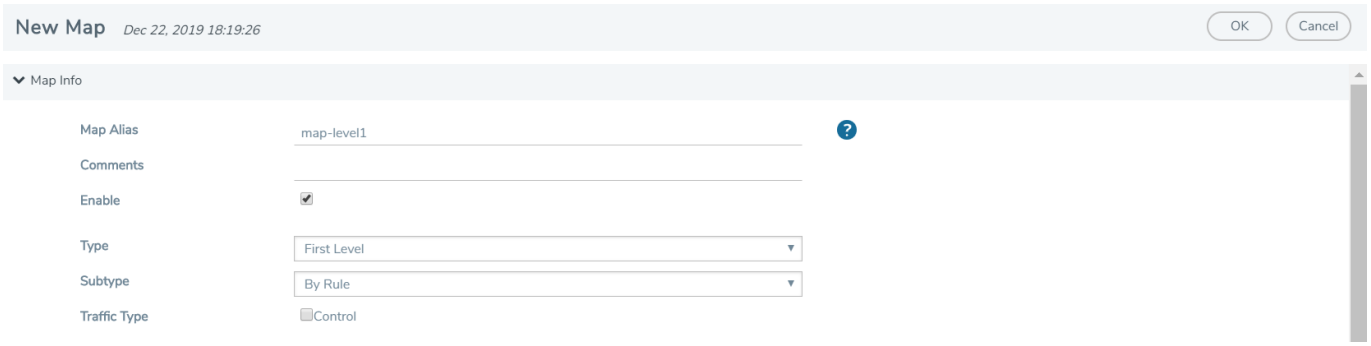


Figure 16 *Create New Map*

6. Select **virtual port vport1** for the Destination.
7. Click **Adda Rule** to add Rule 1
8. Click **Save**.

Create second level map for SIP Flow Sampling.

1. On the left navigation pane, go to **Inventory** from **Traffic** select **Maps > Maps > Maps.**
2. Click **New**.
3. Type **an alias** in the Alias field.
4. Select Second Level for Type and **Flow Sample SIP** for Subtype.
5. Select virtual port **vport1** for the Source.
6. Select **port group** for the Destination.
7. Select **the group** from the GSOP list.
8. Click **Add a Rule**.
9. Select **SIP** for the condition.
10. Select any of the following options:
    - Caller ID-For example, enter **408*** for Caller ID.
    - Caller ID Range
    - Callee ID
    - Callee ID Range
    - Caller/Callee ID Range

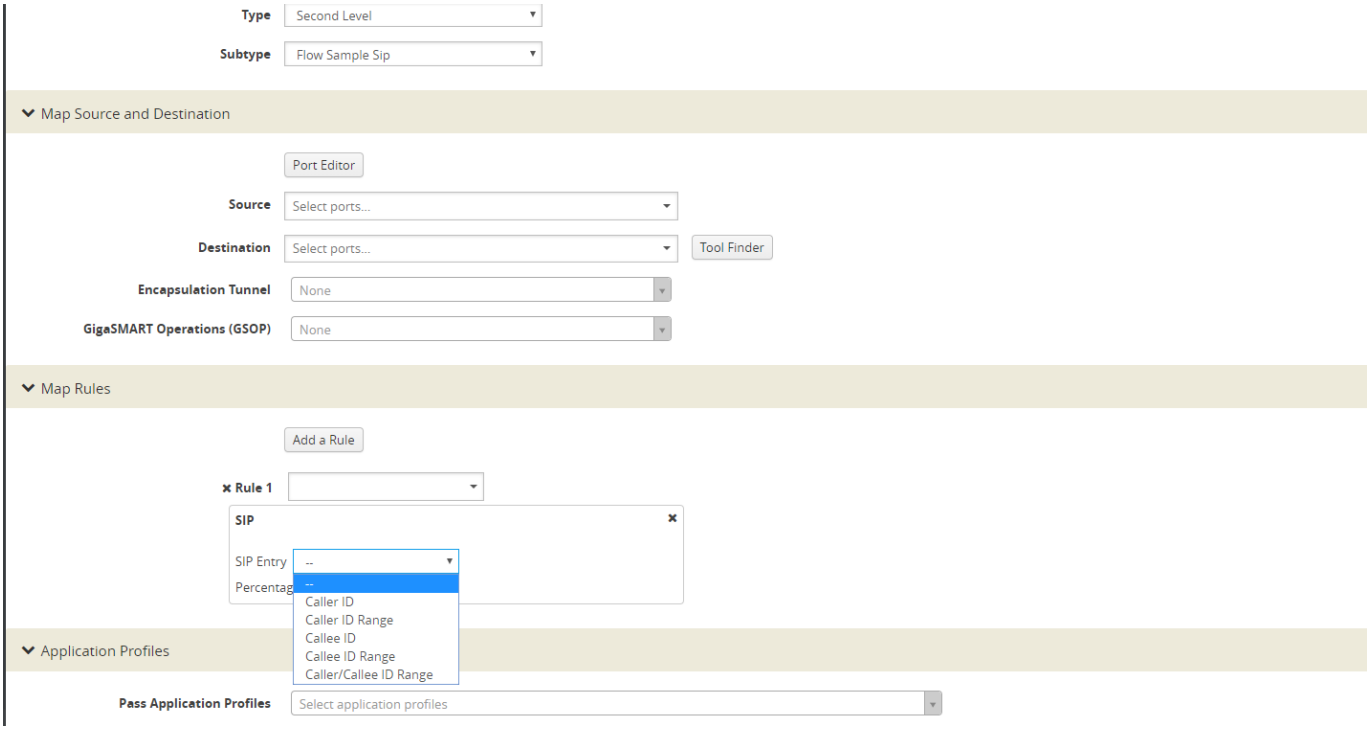


Figure 17 *Create Second Level Map*

11. Enter a **percentage** for amount the traffic you want to be affected by SIP flow sampling.
12. Click **Add a Rule**.
13. Select **SIP** for the condition
14. Enter **6501234*** for Caller ID.
15. Enter a **percentage** for amount the traffic you want to be affected by SIP flow sampling.
16. Click **OK**.

Create the SIP whitelist

1. From the left navigation pane, go to **System > GigaSMART > Whitelist.**
2. Click **New**.
3. Type an **alias** in the Alias field.
4. From the GigaSMART Groups drop-down list, select a .the GigaSMART group.
5. Load whitelist files from a specified location to populate the SIP whitelist.
   a. On the SIP Whitelist page, select **Bulk Upload**.
   b. Select **Bulk Entry Operation** for **Upload Type**
   c. Select Upload from URL from the Bulk Upload Type list.

d. Enter the **URL** in the **Enter Remote URL field**.

6. Click **OK**.

Associate the GigaSMART group to the SIP whitelist.

1. From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART Groups.**
2. Select the **GigaSMART Group** you previously created and click **Edit**.
3. From the **GigaSMART Groups drop-down list,** select a GigaSMART group previously created.
4. Under **SIP Whitelist**, click on the **SIP Whitelist Alias field** and select the **alias** previously created from the available list.
5. Click **OK**.

Configure GigaSMART operation for SIP whitelisting

1. From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations > GigaSMART Operation.**
2. Click **New**.
3. Type an **alias** in the Alias field.
4. Select the GigaSMART group created in task 1.
5. From the GigaSMART Operations (GSOP) drop-down list, select the following:
   - **SIP Whitelist** and select **Enabled.**
   - **Load Balancing**.
6. For Load Balancing, do the following:
   a. Choose: **Stateful**
   b. For Type select: **SIP**

7. Click **OK**.

Configure first level map

1. On the left navigation pane, go to **Inventory** from **Traffic** select **Maps > Maps > Maps**.
2. Click **New**.
3. Configure the map:
   a. Enter an **Alias**
   b. Type: **First Level,** Sub Type: **By Rule**

c. Source: **1/1/g2**

d. Destination: **vport**

4. Click **Add a Rule.**
5. Click **Save**.

Create another second level map for SIP flow whitelist

1. Select **Maps > Maps > Maps.**
2. Click **New**.
3. Configure the map:
   a. Alias: **alias name**
   b. Type: **Second Level**, Sub Type: **Flow Whitelist SIP**
   c. Source: **vport1**
   d. Destination: **1/2x23**
   e. Select from the **GSOP** list.
4. Click **Add a Rule**.
5. Select SIP. In SIP type, you can choose any of the following options:
   - Caller ID
   - Callee ID
   - Caller/Callee ID
   - Source IP
   - Destination IP
   - Source/Destination IP
   - All
6. Click **OK**.

GigaSMART IP FlowVUE

Required License: FlowVUE

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

The GigaSMART IP FlowVUE application enables existing tools to connect to the latest high-speed pipes by providing a representative view of traffic for diagnostic coverage. It allows organizations to sample traffic based on outer IP address or inner IP address across GTP tunnels. It intelligently reduces the amount of traffic while maintaining the integrity of the traffic by forwarding all the traffic flows associated with the sampled IP to the monitoring and analytic tools.

This feature can be used for troubleshooting, resolving tool's over-subscription, and ensuring optimal utilization of tool licenses.

GigaSMART IP FlowVUE supports the following:

- flow-aware sampling of IPs to filter and forward all flows sourced from a sampled set of IPs.
- random sampling on IPs and IP ranges, and at configurable sampling rates.
- Sampling GTP traffic based on inner IP header.
- User-configurable timeouts to detect and replace inactive IPs.

FlowVUE operations can be assigned to GigaSMART groups consisting of a single GigaSMART engine. Refer to Groups of GigaSMART Engine Ports for details.

Configure FlowVUE

The GigaSMART parameters for configuring FlowVUE are as follows:

GigaSMART Field	Definition
Type	Specifies whether inner or outer IP addresses are used for FlowVUE sampling as follows: • Device IP—Specifies a sample subset of traffic based on outer IP address. • Device IP in GTP—Specifies a sample subset of traffic based on inner IP address encapsulated in the GTP tunnels.
IPV4 Ranges /IPV6 Ranges	Specifies the range of IPV4/ IPV6 addresses. **Note**: The maximum IPV4 and IPv6 ranges that are configurable are 64 .
Rate (%)	Specifies the rate for random sampling of the traffic. The values range from 5 to 95%.
Timeout (minutes)	Specifies after how much time a flow in a sampled IP range is declared idle and is no longer sampled. The values range from 1 to 60 minutes.

FlowVUE Configuration Examples

Refer to the following section for examples:

- Sample a subset of IP Traffic
- Sample a subset of Subscribers' IP traffic encapsulated in GTP.
- GigaSMART IP FlowVUE

Display FlowVUE Statistics

To display FlowVUE statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics**.

Refer to FlowVUE Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

FlowVUE_Examples

This section contains:

- Sample a subset of IP Traffic
- Sample a subset of Subscribers' IP traffic encapsulated in GTP.
- Sample a subset of Subscribers' web traffic encapsulated in GTP

Sample a subset of IP Traffic

The following example samples on traffic (IP) where 10% of the traffic is forwarded.

Task	Description	UI Steps
1	Configure one network and two tool type of ports.	1. Go to **System > Ports > Ports > All Ports**. 2. Click Quick Port Editor. 3. Configure one network port and two tool ports. For example, select Network for port 1/1/x3. Select Tool for port 1/1/x4 and port 1/1/x1. 4. Select Enable for each port. 5. Click OK. 6. Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine	1. From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART**

Task	Description	UI Steps
	port and configure sampling parameters	**Groups**. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Click int he Port List field and select an engine port. For example. 1/1/e1. **5.** Configure the Flow Sampling parameters under GigaSMART Parameters. o Select Device IP under Flow Sampling type field. o Enter 1.1.1.0/255.255.255.0 in the IP Ranges field. o Enter 10 for Rate to set the flow sampling rate to 10 percent. o Enter Timeout between 1-60. **6.** Click Save.
3	Configure the GigaSMART operation and assign it to the GigaSMART group.	**1.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations (GSOP) > GigaSMART Operations.** **2.** Click New. **3.** Type gsfvue in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list. **5.** Select Flow Sampling from the GigaSMART Operations (GSOP) list. **6.** Select Flow Sampling - IP **7.** Click Save.
4	Create a Regular map.	**1.** On the left navigation pane, go to **Inventory** from **Traffic** and select **Maps > Maps > Maps**. **2.** Click New. **3.** Configure the map. • Enter to tool in the Alias field. • Select Regular for Type. • Select By Rule for Subtype. • Select the network port for the Source. • Select a tool port for the Destination. • Select gsfvue from the GSOP list. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Port Destination. **d.** Enter the port value. **5.** Click Save.

Sample a subset of Subscribers' IP traffic encapsulated in GTP.

The following example illustrates sampling and forwarding the 10% of subscriber's traffic (IP address encapsulated in GTP-U traffic) to the security and monitoring tools.

Task	Description	UI Steps
1	Configure one network and two tool type of ports.	**1.** Go to **System > Ports > Ports > All Ports**. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3. Select Tool for port 1/1/x4 and port 1/1/x1. **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port and configure sampling parameters	**1.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART Groups**. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Click int he Port List field and select an engine port. For example. 1/1/e1. **5.** Configure the Flow Sampling parameters under GigaSMART Parameters. o Select Device IP in GTP under Flow Sampling type field. o Enter 1.1.1.0/255.255.255.0 in the IP Ranges field. o Enter 10 for Rate to set the flow sampling rate to 10 percent. o Enter Timeout between 1-60. **6.** Click Save.

Task	Description	UI Steps
3	Configure the GigaSMART operation and assign it to the GigaSMART group.	**1.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations (GSOP) > GigaSMART Operations.** **2.** Click New. **3.** Type gsfvue in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list. **5.** Select Flow Sampling from the GigaSMART Operations (GSOP) list. **6.** Select Flow Sampling - IP **7.** Click Save.
4	Create a Regular map.	**1.** On the left navigation pane, go to **Inventory** from **Traffic** and select **Maps > Maps > Maps**. **2.** Click New. **3.** Configure the map. • Enter to tool in the Alias field. • Select Regular for Type. • Select By Rule for Subtype. • Select the network port for the Source. • Select a tool port for the Destination. • Select gsfvue from the GSOP list. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Port Destination. **d.** Enter 2152 as the port value. **Note**: In the rule, 2152 is GTP-U traffic port value. **5.** Click Save.

Sample a subset of Subscribers' web traffic encapsulated in GTP

By combining FlowVUE with other GigaSMART APF, the traffic can be further reduced by filtering on specific Layer 4 application ports.

The following example samples on a subset of IP address encapsulated in GTP and forwards only the HTTP traffic related to the sampled set of IPs.

Task	Description	UI Steps
1	Configure one network and two tool type of ports.	**1.** Go to **System > Ports > Ports > All Ports.** **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3. Select Tool for port 1/1/x4 and port 1/1/x1. **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART Groups**. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Click int he Port List field and select an engine port. For example. 1/1/e1. **5.** Configure the Flow Sampling parameters under GigaSMART Parameters. o Select Device IP in GTP. o Enter 1.1.1.0/255.255.255.0 in the IP Ranges field. o Enter 10 for Rate to set the flow sampling rate to 10 percent. o Enter timeout between 1-60. **6.** Click Save.
3	Configure the GigaSMART operation and assign it to the GigaSMART group. Also, configure APF.	**1.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations (GSOP) > GigaSMART Operation.** **2.** Click New. **3.** Type gsfvue_apf in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list. **5.** Select APF from the GigaSMART Operations (GSOP) list **6.** Enable APF. **7.** Select Flow Sampling from the GigaSMART Operations (GSOP) list. **8.** Select Flow Sampling - IP **9.** Click Save.

Task	Description	UI Steps
4	Configure virtual port and associate it with the GigaSMART group.	**1.** From the left navigation pane, go to **System > GigaSMART >Virtual ports.** **2.** Click New. **3.** Enter vp1 in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map and direct traffic to the virtual port.	**1.** On the left navigation pane, go to **Inventory** from **Traffic** select **Maps > Maps > Maps.** **2.** Click New. **3.** Configure the map. o Enter to_vp in the Alias field. o Select First Level for Type. o Select By Rule for Subtype. o Select the network port for the Source. o Select the virtual port vp1 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Port Source. **d.** Enter the port value. **Note**: In the rule, 2152 is GTP-U traffic port value. **5.** Click Save.
6	Create a second level map and use the APF GigaSMART operation. APF performs filtering according to the gsrules, sending only matching traffic to the tool port.	**1.** On the left navigation pane, go to **Inventory** from **Traffic** select **Maps > Maps > Maps.** **2.** Click New. **3.** Configure the map. o Enter map1 in the Alias field. o Select Second Level for Type. o Select By Rule for Subtype. o Select the virtual port vp1 for the Source. o Select the a tool port for the Destination. o Select gsfvue_apf from the GSOP list. **4.** Add a gsrule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Port Destination. **d.** Enter 80 for the port value. **e.** Select 2 for Position. **6.** Add a gsrule 2.

Task	Description	UI Steps
		a. Click Add a Rule. **b.** Select Pass. **c.** Select Port Source. **d.** Enter 80 for the port value. **e.** Select 2 for Position. **6.** Click Save.

GigaSMART GTP Whitelisting and GTP Flow Sampling

Required Licenses: GTP Filtering & Correlation and FlowVUE

Supported Devices : GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3 ,GigaVUE-HC1-Plus.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

Use GTP forward listing and GTP flow sampling to provide subsets of GTP correlated flows to tools. GTP forward listing selects specific subscribers based on IMSI, while GTP flow sampling uses map rules to select subscribers. Starting in software version 4.8, GigaSMART supports GTP overlap mapping, which combines both forward listing and flow sampling maps as part of a map group. Refer to GTP Overlap Flow Sampling Maps

Starting in software version 4.5, a GigaSMART group (gsgroup) associated with GTP applications can have multiple GigaSMART engine port members (**e** ports), up to four, forming an engine group. Refer to GTP Scaling.

Refer to the following sections:

- GTP Whitelisting
- GTP Flow Sampling
- GTP Subscriber Aware Random Sampling
- Display GTP Flow Ops Report Statistics

GTP Whitelisting

GTP forward listing selects specific subscribers based on IMSI. The forward list contains up to 2,000,000 subscriber IMSIs. For subscribers in the forward list, 100% of their traffic is always sent to a specified tool port.

For example, when a subscriber session comes in, GTP forward listing checks the IMSI of the subscriber. If the incoming IMSI or RAN matches an IMSI or RAN in the forward list, the session is sent to the tool port or load balancing group specified in the forward list map.

Starting in software version 4.7, GTP forward listing is supported in a cluster. Refer to GTP Whitelisting in a Cluster for more information.

Create Forward List

Subscriber IMSIs are added to a forward list that can contain up to 2,000,000 subscriber IMSIs. You can create multiple forward list database per GigaSMART group but the maximum number of whitlelist entries allowed are 2,000,000 IMSIs per GigaSMART group. You can have a maximum of 20 active forward list database in a GigaSMART group. You can also perform dynamic addition of a forward list database after deploying a solution. You can delete the forward list database, only after removing the GigaSMART group. Refer to Delete Forward List.

For the sequences of steps to create a forward list with the UI, refer to the configuration example for forward listing in GigaSMART GTP Whitelisting and GTP Flow Sampling.

Entries in the forward list can be added one at a time or whitelist files containing multiple IMSIs can be created and downloaded. Entries are added by using the GTP Whitelist page by selecting **GigaSMART > GTP Whitelist**. The GTP Whitelist page shows alias for the currently configured GTP Whitelists, the IMSI count for each Forward list and the GigaSMART Group associated with the GTP Whitelist. The GTP Whitelist is associated with the GigaSMART group by specifying its alias in the **GTP Whitelist Alias** field in GigaSMART Group configuration page and then clicking **New**.

An individual IMSI is added by selecting Individual **Entry Operation** and specifying the IMSI in the **Individual IMSI Entry** field.

The IMSIs in whitelist files must be distinct entries, with one IMSI on each line of a file and a maximum of 500,000 entries in each file. This means that 4 files of 500,000 entries will be needed to populate the forward list to its capacity. Wildcards are not supported in whitelist files.

Whitelist files must have a filename with a .txt suffix. Use the GTP Whitelist page to fetch the entries from a whitelist file at a specified location, using one of the following formats, which are specified in the **Enter Remote URL** field when **Bulk Entry Operation** is selected and the **Bulk Upload Type** is **Upload from URL**:

- http://IPaddress/path/filename.txt
- scp://username:password@IPaddress:/path/filename.txt
- tftp://IPaddress/path/filename.txt

To fetch a whitelist file from a local location, select **File Upload** for **Bulk Upload Type** and use the **Browse** button to select the file.

To update an existing forward list, download the whitelist file, add the forward list entry and then re-upload the file. This will not modify or remove the previous entries added in the file.

When a whitelist file is downloaded, the entries are compared to the forward list on the node. There may be new entries in the file that might already be part of the existing forward list. GigaSMART will add the new, non-duplicate entries to the forward list, without rejecting the entire file.

If the current number of entries in the forward list plus the new entries in the whitelist file is greater than the forward list capacity of 2,000,000 IMSIs, the **Append** operation will fail and the new entry or the entries from the new whitelist file will not be added.

GTP forward listing does not use map rules like GTP flow sampling does. The forward list is associated with a GigaSMART group, GigaSMART operation, and second level maps, called whitelist maps.

For the sequences of steps to create a forward list with the UI, refer to the configuration example for forward listing in GigaSMART GTP Whitelisting and GTP Flow Sampling.

Configure Whitelist Maps

The whitelist maps are configured per GigaSMART group. Each forward list map, associated with the same vport, uses the same underlying forward list.

Up to ten (10) whitelist maps are supported. Multiple whitelist maps provide a granular selection of tool ports for forward listing. Using multiple maps, traffic can be segregated and sent to multiple destinations. Forward list map rules allow you to select the subset of IMSIs sent to a particular tool.

Each forward list map can contain up to four rules. The rules specify the type of traffic to be forward listed by that map. Within any single map, the rules are evaluated in order. The rules in the first map have a higher priority than the rules in the second, third, and subsequent maps.

The rules will specify either an Evolved Packet Core (EPC) interface type or a GTP version as the attribute to match. An Access Point Name (APN) and can also be specified in a rule of a Second Level Flow Whitelist map, either by itself, or preceding the EPC interface type or in combination with the GTP version.

For APN, you must specify a pattern (a name) to match. Use APN to direct the traffic that matches the pattern to a specific tool.

GTP version and EPC interface are mutually exclusive. A mix of versions and interface types across whitelist maps, associated with the same vport, is not supported. For example, you can configure two whitelist maps with one map specifying a rule for version 1 and another map specifying a rule for version 2, OR four whitelist maps with each map specifying a rule for each interface type (Gn, S11, S5, and S10). For more information on interfaces, refer to Supported Interfaces.

An APN pattern is for example, three.co.uk. Wildcard prefixes and suffixes are supported, for example, *mobile.com or *ims*. The pattern can be specified in up to 100 case-insensitive alphanumeric characters and can include the following special characters: period (.), hyphen (-), and wildcard (*). A standalone wildcard (*) is not allowed for APN.

You must specify a pattern required for the forward list DataBase (DB) lookup in Type. The following three types of values are supported for the DB lookup:

- imsi/supi — only IMSI or SUPI value used for the DB lookup.
- ran — only RAN value used for the DB lookup.
- all — both RAN and IMSI/SUPI value used for the DB lookup.

By default, SUPI or IMSI is value is used for the DB lookup, if no type is configured.

You can configure a maximum of 10 forward list aliases in a single forward list map. The Database lookup happens only in the configured forward list alias based on the configured DB type.

When there is only DB type is configured and no forward list alias is configured, then the first forward list DB configured in the gsparams is used for the DB lookup.

When there is no DB type and no forward list alias are configured, then the lookup happens in all the forward list DB configured in the gsparams.

Each new subscriber session will be evaluated by the whitelist maps in the order of priority, which, by default, is the order in which the maps were created.

When a subscriber session comes in, GTP forward listing will check the IMSI of the subscriber. If the IMSI is present in the forward list, the rules in the first forward list map is evaluated to qualify the match further. Otherwise, the packet is evaluated against the rules in the subsequent whitelist maps for a possible match.

For example, with one forward list map having a rule specifying GTP version 1 and another forward list map having a rule specifying GTP version 2, when a subscriber session comes in, GTP forward listing will check the IMSI of the subscriber. if the IMSI is present in the forward list and if there is a match to version 1, the session (100% of subscriber packets) will be forwarded to the tool port, GigaStream, or load balancing group specified in the forward list map. If there is not a match to version 1, the next map is evaluated. If there is a match to version 2 in the next map, the session will be forwarded to the tool port, GigaStream, or load balancing group specified in the second forward list map.

NOTE: Both maps can specify the same destination.

Rules can be added to, or deleted from, a forward list map. Use the **Add a Rule** button to add a new forward list rule (a pass rule). Click **x** to delete a rule. A rule in a forward list map cannot be edited. To edit a rule, first delete it, then recreate it.

The default map configuration in which neither GTP version, EPC interface, or APN is specified in the map, continues to be supported. If the incoming IMSI matches an IMSI in the forward list, the session will be sent to the tool port, GigaStream, or load balancing group specified in the forward list map.

Whitelist maps cannot contain any other rules such as GigaSMART rules (gsrule), flow filtering rules (flowrule), or flow sampling rules (flowsample).

GTP whitelist-based forwarding is performed prior to GTP flow sampling (rule-based flow sampling) and GTP flow filtering.

NOTE: For GTP second level maps, a maximum of fifteen maps can be attached to a vport. For example, for the same vport you can have five whitelist maps and ten flow sampling maps, or ten whitelist maps, four flow sampling maps, and one flow filtering map. In addition, you can have a collector map, which is not counted.

For the steps to create a whitelist map with the UI, refer to the configuration example for forward listing in GigaSMART GTP Whitelisting and GTP Flow Sampling.

Change Priority of Whitelist Maps

Use the **Priority** field in the map to change the priority of whitelist maps.

Delete Whitelist Maps

When a forward list map is deleted, the priority of the remaining whitelist maps will be re-prioritized. For example, if the first forward list map is deleted, the second forward list map will increase in priority.

For the deleted forward list map, the traffic associated with the rules in the map will be reevaluated and then passed to subsequent maps.

When a forward list map is re-prioritized, the existing sessions will be reevaluated according to the new priority of the map. The traffic associated with the rules in the map will be reevaluated and then passed to subsequent maps.

When the last forward list map is deleted, the traffic associated with the rules in the map will also be reevaluated before being passed to subsequent maps. But the traffic associated with the rules in maps that were not matched, will not be reevaluated because that traffic was already passed to subsequent maps.

Apply Forward List

When a single forward list entry is added, forward listing is applied for new as well as existing subscribers. When a new whitelist file is fetched, forward listing is applied only for new subscribers.

Forward listed traffic is then sent to the port or load balancing group specified in the whitelist map.

Delete Entry from Forward List

Entries in the forward list can be deleted one at a time. Each entry is a single IMSI.

When a forward list entry is deleted, the session associated with the forward list entry stays active and traffic is still sent to the whitelist map. The forward list session will not be reevaluated or passed to subsequent maps.

To delete a single entry from the forward list, select **Individual Entry Operation**, set **Remove** as the **Operation Type,** and enter the IMSI in the **Individual IMSI Entry** field.

Delete Multiple Entries from Forward List

Multiple IMSIs can be deleted from the forward list. Specify the IMSIs to be deleted in a whitelist file, which can contain up to 20,000 IMSIs.

Whitelist files must have a filename with a .txt suffix. To remove multiple entries from the forward list, select **Bulk Entry Operation** and set **Remove** as the **Operation Type**.

Delete Forward List

The entire forward list can be deleted using one of the following options:

- Delete the forward list by deleting all the IMSI entries. With this option, you do not have to delete the forward list map, GigaSMART operation, or disassociate the GigaSMART group from the forward list. To delete all the IMSI entries, select **Delete All**.
- Destroy the forward list. With this option, you must first delete the forward list map, GigaSMART operation, and disassociate the GigaSMART group from the forward list before deleting the forward list.
- Alternatively ,select the forward list and click on edit option then select the **'Clear'** radio button, to remove existing forward list.

Destroy Forward List

To destroy a forward list, use the following sequence:

Task	UI Steps
Delete the forward list map	**1.** On the left navigation pane, go to **Inventory** from **Traffic** and select **Maps > Maps > Map**. **2.** Select the forward list map. **3.** Click Delete.
Delete the GigaSMART Operation	**1.** From the left navigation pane, go to **System > GigaSMART > GigaSMART Operation(s)** > **GigaSMART Operation**. **2.** Select the GigaSMART Operation. **3.** Click Delete.
Disassociate the GigaSMART group from the forward list. (You do not need to delete the GigaSMART group.)	**1.** From the left navigation pane, go to **System > GigaSMART > GigaSMART Groups > GigaSMART Groups.**

Task	UI Steps
	2. Select the GigaSMART group. 3. Click Edit. 4. Under GigaSMART Parameters, go the GTP Whitelist and set GTP Whitelist Alias to None.
Destroy (delete) the forward list.	1. From the left navigation pane, go to **System > GigaSMART > GigaSMART GTP Whitelist**. 2. Select the GTP Whitelist. 3. Click Delete. 4. Alternatively ,select the forward list and click on edit option then select the **'Clear'** radio button, to remove existing forward list.

GTP Whitelisting in a Cluster

The forward list(all whitelist files) must reside on the leader of the cluster. The member nodes receive a copy of the forward list from the leader. Updates to the forward list are synchronized from the leader to the member nodes. If a member node leaves the cluster and rejoins, its forward list will be resynchronized.

If there are GigaVUE TA Series nodes in the cluster, they will not receive a copy of the forward list.

GTP Flow Sampling

GTP flow sampling samples a configured percentage of GTP sessions. GTP flow sampling uses map rules to select subscribers and then forward a percentage of the packets to tool ports.

Pass rules are defined in flow sampling maps. Each rule contains some combination of IMSI, IMEI, and MSISDN numbers or patterns, Evolved Packet Core (EPC) interface type, GTP version, Access Point Name (APN), or QoS Class Identifier (QCI), as well as a percentage to sample. The flow is sampled to see if it matches a rule. The percentage of the subscriber sessions matching each rule are selected.

Map rules specify the type of traffic to be flow sampled by that map. For each new session, map rules are evaluated in top-down order of decreasing priority. If there is a match, the indicated percentage of the subscriber session is either accepted or rejected. If accepted,

the traffic is sent to the tool port or load balancing group specified in the map. If rejected, the traffic is dropped. If there is not a match to a rule, the traffic is passed to subsequent maps.

Starting in software version 4.6, GTP load balancing in a cluster is supported for GTP flow sampling. For an example of GTP load balancing in a cluster, refer to GigaSMART GTP Whitelisting and GTP Flow Sampling Examples.

About Flow Sampling Rules and Maps

Flow sampling rules are configured in maps called flow sampling maps. Up to ten (10) flow sampling maps per GigaSMART group are supported. Each flow sampling map supports up to 20 flow sampling rules, for a maximum of 200 rules per GigaSMART group.

GTP flow sampling (rule-based flow sampling) is performed after GTP forward list-based forwarding but before GTP flow filtering. So, flow sampling maps have a priority lower than whitelist maps and higher than flow filtering maps.

NOTE: For GTP second level maps, a maximum of fifteen maps can be attached to a vport. For example, for the same vport you can have one forward list map and ten flow sampling maps, or ten forward list map, four flow sampling maps, and one flow filtering map. In addition, you can have a collector map, which is not counted.

In the flow sampling maps, the rules in the first map have a higher priority than the rules in the second, third, and subsequent maps. Within any single map, rules are evaluated in order.

Rules can be added to, deleted from, or inserted into a flow sampling map when the subtype selected for a **Second Level** map is **Flow Sample**. Suffix wildcarding, such as IMSI 100*, is supported in the flow sampling map rules.

Use the **Add a Rule** button in the Maps page to add a new flow sampling rule (a pass rule). Specify IMSI, IMEI, or MSISDN subscriber IDs, as well as the percentage of the flow to be sampled. The percentage is a range from 1 to 100%. Use 0% to drop sampled data.

A rule can specify other packet attributes, such as an EPC interface type or GTP version. An APN pattern can also be specified in a rule, either by itself or preceding the EPC interface or GTP version. A QCI value can be specified, but only in combination with an APN pattern.

EPC interface and GTP version are mutually exclusive. They can be specified in a flow sampling rule, but not both in a single rule. The supported interface types for filtering are: Gn/Gp, S11/S1-U, S5/S8, S10, or S2B. The supported versions for filtering are 1 or 2. For example, you can send version 1 traffic to one tool port and version 2 traffic to another tool port. For more information on interfaces, refer to Supported Interfaces.

For APN, specify a pattern (a name) to match, for example, three.co.uk. Wildcard prefixes and suffixes are supported, for example, *mobile.com or *ims*. The pattern can be specified in up to 100 case-insensitive alphanumeric characters and can include the following special characters: period (.), hyphen (-), and wildcard (*).

QCI is a mechanism used in Long Term Evolution (LF TE) networks to ensure bearer traffic is allocated to the appropriate Quality of Service (QoS). For QCI, specify a value from 0 to 255. Wildcard prefixes and suffixes are not supported.

Use APN and QCI to send traffic that matches a certain APN pattern or that belongs to a certain bearer with a certain QCI to specified tool ports, based on the sampling percentage.

Click the **x** next to a rule to delete a specific rule. Rules are identified by a priority ID, which indicates the order of rules in a flow sampling map. For example, if a map has 12 pass flow sampling rules, there will be 12 priority IDs.

When creating Flow Sampling rules on the Maps page, the first rule created has the highest priority and the priority of subsequent rules is in the order that they are added. To change the priority of a Flow Sampling rule in a new map, do the following:

1. **Save** the rule.
2. Select the map and click **Edit**.
3. Enter a priority in the **Priority** field of each rule to order the rules in the map. (For details about map priority, refer to Map Priority)

NOTE: A flow sampling map can contain only flowsampling rules. A flow sampling map cannot contain other GigaSMART rules (gsrule) or flow filtering rules (flowrule).

For configuration examples for flow sampling, refer toGigaSMART GTP Whitelisting and GTP Flow Sampling.

Add Rule to Flow Sampling Map

Flow sampling is applied for new subscribers. When a new rule is added to the rules in a flow sampling map, traffic will be sent to the port or load balancing group specified in the map.G

Delete Rule from a Flow Sampling Map

When a rule is deleted from a flow sampling map, the session associated with the rule stays active. The traffic associated with the rule will not be reevaluated by subsequent maps.

Change Priority of Flow Sampling Maps

Use the **Priority** field in the GTP map rule to set the priority of flow sampling maps.

> **NOTE:** Editing the map priority in non-overlap will reset the active sessions. Total sessions will remain unaffected and continue incrementing with incoming user traffic.

Delete Flow Sampling Map

When a flow sampling map is deleted, the priority of the remaining flow sampling maps will be re-prioritized. For example, if the first flow sampling map is deleted, the second flow sampling map will increase in priority.

For the deleted flow sampling map, the traffic associated with the rules in the map will be reevaluated and then passed to subsequent maps.

When a flow sampling map is re-prioritized, the existing sessions will be reevaluated according to the new priority of the map. The traffic associated with the rules in the map will be reevaluated and then passed to subsequent maps.

When the last flow sampling map is deleted, the traffic associated with the rules in the map will also be reevaluated before being passed to subsequent maps. But the traffic associated with the rules in maps that were not matched, will not be reevaluated because that traffic was already passed to subsequent maps.

Flow-Ops Report Limitation for Multiple Flow Sampling Maps

The flow-ops report displays the flow sampling rule ID for sessions that have been accepted or rejected by the flow sampling map.

However, since rule IDs are not unique across maps, when there are multiple flow sampling maps, the flow-ops report is unable to identify the exact rule that the session matched. For example, with multiple flow sampling maps, each map can have a rule ID of 1. The rule ID will be identified in the flow-ops report, but not the map associated with it.

GTP Flow Sampling Percentage

The sampling Percentage field in a map for GTP flow sampling, represents the percentage of subscribers that will be sampled (not the sessions).

The GTP correlation engine tracks all of the subscribers and all of their sessions that it sees on the network. In this example, for those subscribers with an IMSI starting with the value 46*, the GTP correlation engine keeps a list of them and randomly selects 80% of those subscribers and sets them to be in the sample, which means that a tool port (or load balanced group) will see 100% of the packets for 100% of the sessions for those randomly selected 80% of subscribers.

For the other 20% of subscribers, the GTP correlation engine continuously tracks those subscribers through the network, but does not send any packets to the tool port (or load balanced group).

Refer to the GTP flow sampling configuration examples in GigaSMART GTP Whitelisting and GTP Flow Sampling.

Drop Unmatched Traffic

When a session matches one of the configured flow sampling rules, it is either accepted for sampling or rejected.

If it is accepted, all packets belonging to that GTP session are sent to the tool port or ports specified in the flow sampling maps. If a subscriber is in the sample, then both the control plane packets and the user-data plane packets are sent to the tools.

If it is rejected, all packets belonging to the session are dropped. If the subscriber is not in the sample, then neither the control plane packets nor the user-data plane packets are sent to the tools.

Control plane (GTP-c) and user-data plane (GTP-u) traffic are treated the same. For a matching session, all the control plane and user-data plane traffic will be accepted. Otherwise, all the control plane and user-data plane traffic will be rejected and dropped. Instead, to enable or disable GTP control plane traffic sampling, refer to Enable or Disable GTP Control Plane Traffic Sampling.

Enable or Disable GTP Control Plane Traffic Sampling

GTP control plane (GTP-c) traffic is typically a small percentage of total GTP traffic, but it contains useful information for analytics. Therefore, it is not always expedient to drop control plane traffic for sampled sessions.

Subscriber traffic by IMSI can be sampled such that network traffic for a subset of mobile subscribers can be selected to be sent to network monitoring tools. In some cases, network monitoring tools will want to see GTP control plane and GTP user plane traffic for a

percentage of the subscribers. In other cases, network monitoring tools will want to see all of the GTP control plane traffic, but see only the GTP user plane traffic for the sampled percentage of subscribers.

Starting in software version 4.5, all control plane traffic for all subscribers will be sent to tools if GTP control plane traffic sampling is disabled. When disabled, 100% of the control traffic that matches any of the flow sampling rules will be sent to the tool ports specified in the flow sampling maps. Control traffic for both accepted and rejected sessions will be sent to the tool ports.

When GTP control plane traffic sampling is enabled, GTP-c packets will be sampled and only the indicated percentage of the control traffic that matches any of the flow sampling rules will be sent to the tool ports specified in the flow sampling maps, as described in GTP Flow Sampling Percentage.

The default is enable.

To disable sampling of GTP-c traffic, which enables 100% of control plane traffic, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**. Under GigaSMART Parameters, go to **GTP Sampling** and make sure that **GTP Control Sampling** is not selected.

To enable sampling of GTP-c traffic, which enables 100% of control plane traffic, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**. Under GigaSMART Parameters, go to **GTP Sampling** and make sure that **GTP Control Sampling** is selected. This setting applies to all the flow sampling maps for a GigaSMART group.

GTP Subscriber Aware Random Sampling

GTP Subscriber Aware Random Sampling allows to randomly sample all the subscriber's IMSI on a rotational basis. Based on the configured sampling percentage, the selected sessions are either sampled in or out. The correlation engine takes the configurable interval as an input to rotate the random selection of each of the subscriber's sessions.

The configurable interval is a minimum of 12 hours and a maximum of 48 hours. Each GigaSMART node must be synchronized with an NTP/PTP server, as UTC time is involved in the random selection of the subscriber's sessions.

NOTE: This feature is effective for a new subscriber's sessions after enabling the random sampling.

The Map rules in the GTP random sampling are similar to GTP Flow Sampling. For more information refer to GigaSMART GTP Whitelisting and GTP Flow Sampling.

To enable GTP Random Sampling do the following:

1. From the left navigation pane, go to **System > GigaSMART > GigaSMART Groups**.
2. Select a GigaSMART Group and click **Edit**.
3. Under GigaSMART Parameters, go to **GTP** and select **GTP Random Sampling** check box.
4. Enter the time in **Rotation Interval in multiples of 12 hours**.
5. Click **OK**.

Display GTP Flow Ops Report Statistics

To display GTP statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics**.

Refer to *Flow Ops Report Statistics Definitions for GTP* on page 635 for descriptions of these statistics.

GigaSMART GTP Whitelisting and GTP Flow Sampling Examples

Refer to the following examples:

- Example 1: GTP Whitelisting
- Example 2: GTP Whitelisting with Multiple Maps
- Example 3: GTP Flow Sampling
- *Example 4: GTP Whitelisting, GTP Flow Sampling, and Load Balancing*
- Example 5: GTP Flow Sampling with Multiple Maps
- Example 6: APN for GTP Whitlelisting, GTP Flow Sampling
- Example 7: APN for FTP Whitelisting, APN and QCI for GTP Flow Sampling

Example 1: GTP Whitelisting

Example 1 is a GTP whitelisting configuration example. Traffic from network ports go to the three first level maps (GTP-Control, GTP-User, and Fragments-Not_First) and then to the virtual port (vport1). If there is a match to an IMSI in the whitelist (MyIMSIs), it is forwarded to a port.

Task	Description	UI Steps
1.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	• From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups**. • Click New. • Type an alias in the Alias field and enter an engine port in the Port List field. • Click Save.
2.	Create a virtual port.	• From the left navigation pane, go to **System > GigaSMART >Virtual Ports**. • Click New. • Type an alias in the Alias field and enter an engine port in the Port List field. • From the GigaSMART Groups drop-down list, select the GigaSMART group created in task 1. • Click Save.
3.	Create the GTP whitelist.	• From the left navigation pane, go to **System > GigaSMART >GTP Whitelist.** • Click New. • Type an alias in the Alias field. You can also create mulitiple whitelist aliases per gsgroup during the creation of solution. • From the GigaSMART Groups drop-down list, select the GigaSMART group created in Task 1. • Go to Task 4.
4.	Fetch whitelist files from a specified location to populate the GTP whitelist.	• On the GTP Whitelist page, select Bulk Upload. • Select Bulk Entry Operation for IMSI Upload Type • Select Upload from URL from the Bulk Upload Type list. • Enter the URL in the Enter Remote URL field. For example, http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx • Click Save.
5	Associate the GigaSMART group to the GTP whitelist.	• From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups.** • Select the GigaSMART Group created in Task 1 and click Edit. • Type an alias in the Alias field. You can also associate mulitiple whitelist aliases per gsgroup during the creation of solution. • From the GigaSMART Groups drop-down list, select the GigaSMART group created in task 1. • Under GTP Whitelist, click on the GTP Whitelist Alias field and select the alias from Task 3. • Click Save.
6.	Configure the GigaSMART operation for GTP whitelisting.	• From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations > GigaSMART Operation.** • Click New. • Type an alias in the Alias field. For example, GTP-Whitelist. • Select the GigaSMART group created in task 1.

Task	Description	UI Steps

		• From the GigaSMART Operations (GSOP) drop-down list, select the following:
		  ◦ GTP Whitelist and select Enabled.
		  ◦ Load Balancing.
		• For Load Balancing, do the following:
		  ◦ Choose Stateful
		  ◦ For Type select GTP
		• Choose Hashing for the metric and select IMSI
		• Click Save.

7.	Configure three first level maps.
	NOTE: In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.
		• Configure the first map as follows:
		  ◦ Alias: GTP-Control
		  ◦ Type and subtype: First Level By Rule
		  ◦ Source: network port or ports
		  ◦ Destination: virtual port created in Task 2.
		  ◦ Rule: Pass, Bi Directional, Port Destination 2123
		  ◦ Map Permissions: Select current user's group for Owner
		  ◦ Save the map
		• Configure the second map as follows:
		  ◦ Alias: GTP-User
		  ◦ Type and subtype: First Level By Rule
		  ◦ Source: Same network port or ports as first map.
		  ◦ Destination: virtual port created in Task 2.
		  ◦ Rule: Pass, Bi Directional, Port Destination 2152
		  ◦ Map Permissions: Select current user's group for Owner
		  ◦ Save the map
		• Configure the third map as follows:
		  ◦ Alias: Fragments-Not-First
		  ◦ Type and subtype: First Level By Rule
		  ◦ Source: Same network port or ports as first map
		  ◦ Destination: virtual port created in Task 2
		  ◦ Rule: Pass, IPv4 Fragmentation and select allFragNoFirst
		  ◦ Map Permissions: Select current user's group for Owner.
		  ◦ Save the map

8.	Configure a second level map for GTP whitelisting, the whitelist map. If there is a match to an IMSI in the whitelist (MyIMSIs), it is forwarded to a port.	• Configure the second level map as follows:
		  ◦ Alias: GTP-Whitelist
		  ◦ Type and subtype: Second Level By Rule
		  ◦ Source: virtual port created in Task 2
		  ◦ Destination: select a tool port
		  ◦ GSOP: GigaSMART Operation created in Task 6
		  ◦ Map Permissions: Select current user's group for Owner

Task	Description	UI Steps
		• Click Save.

Example 2: GTP Whitelisting with Multiple Maps

Example 2 is a GTP whitelisting configuration example that includes multiple GTP whitelisting maps, which provide a more granular selection of tool ports.

Traffic from network ports go to the three first level maps (GTP-Control, GTP-User, and Fragments-Not-First) and then to the virtual port (vport1). Two whitelist maps are configured. The first map specifies a rule for version 1 traffic. The second map specifies a rule for version 2 traffic.

Task	Description	UI Steps
1.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	• From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART Groups.** • Click New. • Type an gsg1 in the Alias field and enter an engine port in the Port List field, for example 10/7/e1. • Click Save.
2.	Create a virtual port.	• From the left navigation pane, go to **System > GigaSMART >Virtual Ports.** • Click New. • Type vport1 in the Alias field. • From the GigaSMART Groups drop-down list, select the GigaSMART group created in task 1. • Click Save.
3.	Create the GTP whitelist.	• From the left navigation pane, go to **System > GigaSMART >GTP Whitelist.** • Click New. • Type an MyIMSIs in the Alias field. • From the GigaSMART Groups drop-down list, select the GigaSMART group created in Task 1. • Go to Task 4.
4.	Fetch whitelist files from a specified location to populate the GTP whitelist.	• On the GTP Whitelist page, select Bulk Upload. • Select Bulk Entry Operation for IMSI Upload Type • Select Upload from URL from the Bulk Upload Type list. • Enter the URL in the Enter Remote URL field. For example, http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx

Task	Description	UI Steps
		• Click Save.
5.	Associate the GigaSMART group to the GTP whitelist.	• From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART Groups.** • Select the GigaSMART Group created in Task 1 and click Edit. • From the GigaSMART Groups drop-down list, select the GigaSMART group created in task 1. • Under GTP Whitelist, click on the GTP Whitelist Alias field and select the alias from Task 3. • Click Save.
5.	Configure the GigaSMART operation for GTP whitelisting.	• From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations > GigaSMART Operation**. • Click New. • Type gtp-whitelist in the Alias field. • Select the GigaSMART group created in task 1. • From the GigaSMART Operations (GSOP) drop-down list, select the following: o GTP Whitelist and select Enabled. o Load Balancing. • For Load Balancing, do the following: o Choose Stateful o For Type select GTP o Choose Hashing for the metric and select IMSI o Click Save.
6.	Configure three first level maps. **NOTE:** In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	Configure the first map. **a.** On the left navigation pane, go to **Inventory** from **Traffic** and select **Maps > Maps > Maps.** **b.** Click New. **c.** Configure the map: ▪ Alias: GTP-Control ▪ Type: First Level, Sub Type: By Rule ▪ Source: 8/1/x40, 8/1/x6 ▪ Destination: vport1 **d.** Click Add a Rule. ▪ Select Pass and Bi Directional ▪ Select Port Destination for the rule ▪ Set port value to 2123 **e.** Click Save. • Configure the second map.

Task	Description	UI Steps
		a. Click New. **b.** Configure the map: ▪ Alias: GTP-User ▪ Type: First Level, Sub Type: By Rule ▪ Source: 8/1/x40, 8/1/x6 ▪ Destination: vport1 **c.** Click Add a Rule. ▪ Select Pass and Bi Directional ▪ Select Port Destination for the rule ▪ Set port value to 2152 **d.** Click Save. • Configure the second map. **a.** Click New. **b.** Configure the map: ▪ Alias: Fragment-Not-First ▪ Type: First Level, Sub Type: By Rule ▪ Source: 8/1/x40, 8/1/x6 ▪ Destination: vport1 **c.** Click Add a Rule. ▪ Select Pass ▪ Select Port IPv4 Fragmentation for the rule ▪ Select allFragNoFirst for Value **d.** Click Save.
7.	Configure one second level map for GTP whitelisting, the first whitelist map. If there is a match to version 1 and if the IMSI is present in the whitelist (MyIMSIs), it is forwarded to the specified port.	**a.** On the left navigation pane, go to **Inventory** and from **Traffic** select **Maps > Maps > Map**s. **b.** Click New. **c.** Configure the map: ▪ Alias: GTP-Whitelist_v1 ▪ Type: Second Level, Sub Type: Flow Whitelist ▪ Source: vport1 ▪ Destination: 1/2x23 ▪ Select gtp-whitelist from the GSOP list. **d.** Click Add a Rule. ▪ Select GTP ▪ Set Version to V1 **e.** Click Save.
8.	Configure another second level map for GTP whitelisting, the second whitelist map. If there is	**a.** On the left navigation pane, go to **Inventory** and from **Traffic** select **Maps > Maps > Map**s.

Task	Description	UI Steps
	a match to version 2 and if the IMSI is present in the whitelist (MyIMSIs), it is forwarded to the specified port.	**b.** Click New. **c.** Configure the map: ▪ Alias: GTP-Whitelist_v2 ▪ Type: Second Level, Sub Type: Flow Whitelist ▪ Source: vport1 ▪ Destination: 1/2x24 ▪ Select gtp-whitelist from the GSOP list. **d.** Click Add a Rule. ▪ Select GTP ▪ Set Version to V2 **e.** Click Save.

Example 3: GTP Flow Sampling

Example 3 is a GTP flow sampling configuration example. Traffic from network ports go to the three first level maps (GTP-Control, GTP-User, and Fragments-Not_First) and then to the virtual port (vport1). The traffic flow is sampled based on the rules in one flow sampling map (GTP-Sample-01). The flow sampling rules specify IMSI, IMEI, and MSISDN numbers, as well as the percentage to sample. Packets are then accepted or rejected. Accepted packets are forwarded to a port. Rejected packets are dropped. Packets that do not match a rule will be passed to subsequent maps.

Task	Description	UI Steps
1.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**a.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART Groups**. **b.** Click New. **c.** Type an alias in the Alias field and enter an engine port in the Port List field. **d.** Click Save.
2.	Create a virtual port.	**a.** From the left navigation pane, go to **System > GigaSMART >Virtual Ports**. **b.** In the Alias field, type an alias for this virtual port. **c.** Type an alias in the Alias field and enter an engine port in the Port List field. **d.** From the GigaSMART Groups drop-down list, select the GigaSMART group created in task 1.

Task	Description	UI Steps
		e. Click Save.
3.	Configure three first level maps. **NOTE:** In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**a.** Configure the first map as follows: ▪ Alias: GTP-Control ▪ Type and subtype: First Level By Rule ▪ Source: network port or ports ▪ Destination: virtual port created in Task 2. ▪ Rule: Pass, Bi Directional, Port Destination 2123 ▪ Map Permissions: Select current user's group for Owner ▪ Save the map **b.** Configure the second map as follows: ▪ Alias: GTP-User ▪ Type and subtype: First Level By Rule ▪ Source: Same network port or ports as first map. ▪ Destination: virtual port created in Task 2. ▪ Rule: Pass, Bi Directional, Port Destination 2152 ▪ Map Permissions: Select current user's group for Owner ▪ Save the map **c.** Configure the third map as follows: ▪ Alias: Fragments-Not-First ▪ Type and subtype: First Level By Rule ▪ Source: Same network port or ports as first map ▪ Destination: virtual port created in Task 2 ▪ Rule: Pass, IPv4 Fragmentation and select allFragNoFirst ▪ Map Permissions: Select current user's group for Owner **d.** Save the map
4.	Configure the GigaSMART operation for GTP flow sampling.	**a.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations (GSOP) > GigaSMART Operation**s. **b.** Click New. **c.** Type an alias in the Alias field. For example, GTP-Whitelist. **d.** Select the GigaSMART group created in task 1. **e.** From the GigaSMART Operations (GSOP) drop-down list, select the following: ▪ GTP Whitelist and select Enabled. ▪ Load Balancing.

Task	Description	UI Steps
		f. For Load Balancing, do the following: **a.** Choose Stateful **b.** For Type select GTP **c.** Choose Hashing for the metric and select IMSI **d.** Click Save.
5.	Configure a second level map for GTP flow sampling, the flow sampling map. The traffic flow is sampled based on the rules in this map.	**a.** On the left navigation pane, go to **Inventory** and from **Traffic** select **Maps > Maps > Map**s. **b.** Click New. **c.** Configure the map. ▪ Type GTP-Sample-01 in the Alias field ▪ Select Second Level for Type ▪ Select Flow Sample for Subtype. ▪ Select the virtual port configured in Task 2 for the Source ▪ Select a tool port for the Destination ▪ Select the GigaSMART Operation configured in Task for from the GSOP list **d.** Use the Add a Rule button to create the following flow sampling rules: ▪ Percentage to 50, IMEI 01416800* ▪ Percentage to 80, IMSI 46* ▪ Percentage to 25, MSISDN 1509* ▪ Percentage to 15, IMSI 01400* ▪ Percentage to 20, IMSI, 31*, MSISDN 1909* **e.** Click Save.

Example 4: GTP Whitelisting, GTP Flow Sampling, and Load Balancing

Example 4 combines the GTP whitelisting configuration from Example 1 with the GTP flow sampling configuration from Example 3, and adds GigaSMART load balancing.

In Example 4, traffic from network ports go to the three first level maps (GTP-Control, GTP-User, and Fragments-Not-First) and then to the virtual port (vport1). If there is a match to an IMSI in the whitelist (MyIMSIs), it is forwarded to the port group (PG-Whitelist) for load balancing.

NOTE: In Example 4, the tool ports in the port group are on the same node as the GigaSMART group and GigaSMART operation.

If there is not a match to an IMSI in the whitelist, the traffic flow is sampled based on the rules in the flow sampling map (GTP-Sample-01). The flow sampling rules specify IMSI, IMEI, and MSISDN numbers, as well as the percentage to sample. Packets are then accepted or rejected. Accepted packets are forwarded to the port group (PG-Sample) for load balancing. Rejected packets are dropped. Packets that do not match a rule will be passed to subsequent maps.

Task	Description	UI Steps
1.	Create port groups and specify the tool ports and enable load balancing.	**a.** Go to **System > Ports >** select **Ports > Port Groups > All Port Groups.** **b.** Click New. **c.** Type PG-Whitelist in the Alias field. **d.** Select SMART Load Balancing **e.** Click in the Ports field and select the tool ports for the port group. **f.** Click Save. **g.** Repeat steps 2 through 6, to create a port group with the alias PF-Sample.
2.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**a.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART Groups.** **b.** Click New. **c.** Type an alias in the Alias field and enter an engine port in the Port List field. **d.** Click Save.
3.	Create a virtual port.	**a.** From the left navigation pane, go to **System > GigaSMART >Virtual Ports.** **b.** In the Alias field, type an alias for this virtual port. **c.** Type an alias in the Alias field and enter an engine port in the Port List field. **d.** From the GigaSMART Groups drop-down list, select the GigaSMART group created in task 1. **e.** Click Save.
4.	Configure three first level maps. **NOTE:** In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**a.** Configure the first map as follows: ▪ Alias: GTP-Control ▪ Type and subtype: First Level By Rule ▪ Source: network port or ports ▪ Destination: virtual port created in Task 2. ▪ Rule: Pass, Bi Directional, Port Destination 2123 ▪ Map Permissions: Select current user's group for Owner

Task	Description	UI Steps
		b. Save the map **c.** Configure the second map as follows: ▪ Alias: GTP-User ▪ Type and subtype: First Level By Rule ▪ Source: Same network port or ports as first map. ▪ Destination: virtual port created in Task 2. ▪ Rule: Pass, Bi Directional, Port Destination 2152 ▪ Map Permissions: Select current user's group for Owner ▪ Save the map **d.** Configure the third map as follows: ▪ Alias: Fragments-Not-First ▪ Type and subtype: First Level By Rule ▪ Source: Same network port or ports as first map ▪ Destination: virtual port created in Task 2 ▪ Rule: Pass, IPv4 Fragmentation and select allFragNoFirst ▪ Map Permissions: Select current user's group for Owner **e.** Save the map
5.	Create the GTP whitelist.	**a.** From the left navigation pane, go to **System > GigaSMART >GTP Whitelist.** **b.** Click New. **c.** Type an Alias for the Whitelist in the Alias field. For example, MyIMSIs
6.	Fetch whitelist files from a specified location to populate the GTP whitelist.	**a.** On the GTP Whitelist page, select Bulk Upload. **b.** Select Bulk Entry Operation for IMSI Upload Type **c.** Select Upload from URL from the Bulk Upload Type list. **d.** Enter the URL in the Enter Remote URL field. For example, http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx **e.** Click Save.
7.	(Optional) Add a single IMSI to the GTP whitelist.	**a.** On the GTP Whitelist page, select Individual Entry Operation. **b.** Select Append for Operation Type **c.** Enter the IMSI entry in the Individual IMSI Entry field. **d.** Click Save.
8.	Associate the GigaSMART group to the GTP whitelist.	**a.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART Groups.**

Task	Description	UI Steps
		b. Click New. **c.** Type an alias in the Alias field. **d.** Under GTP Whitelist, click on the GTP Whitelist Alias field and select the alias from Task 5. **e.** Click Save.
9.	Configure the GigaSMART operation for GTP whitelisting.	**a.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations (GSOP)> GigaSMART Operation.** **b.** Click New. **c.** Select the GigaSMART Group created in Task 8 from the GigaSMART Groups list. **d.** Type an alias in the Alias field. For example, gtp-whitelist. **e.** Select GTP Whitelist from the GigaSMART Operations (GSOP) list. **f.** Select Load Balancing from the GigaSMART Operations (GSOP) list. **g.** Configure Load Balancing as follows: ▪ Select Stateful ▪ Set Type to GTP ▪ Select Hashing ▪ Select IMSI **h.** Click Save.
10.	Configure a second level map for GTP whitelisting, the whitelist map. If there is a match to an IMSI in the whitelist (MyIMSIs), it is forwarded to a load balancing port group.	**a.** On the left navigation pane, go to **Inventory** and from **Traffic** select **Maps > Maps > Maps**. **b.** Click New **c.** Configure the map. ▪ Type an name in the Alias field. For example GTP-Whitelist. ▪ Select Second Level for Type ▪ Select By Rule for Subtype ▪ Select the GigaSMART Operation configured in Task 9 from the GigaSMART Operations (GSOP) list. ▪ Select the virtual port configured in Task 3 for Source ▪ Select PG-Whitelist for Destination **d.** Click Save.
11.	Configure the GigaSMART operation for GTP flow sampling.	**e.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations (GSOP) > GigaSMART Operation.** **f.** Click New.

Task	Description	UI Steps
		g. Select the GigaSMART Group created in Task 8 from the GigaSMART Groups list. **h.** Type an alias in the Alias field. For example, gtp-flowsample. **i.** Select Flow Sampling from the GigaSMART Operations (GSOP) list. **j.** Select Flow Sampling-GTP. **k.** Select Load Balancing from the GigaSMART Operations (GSOP) list. ■ Select Stateful ■ Set Type to GTP ■ Select Hashing ■ Select IMSI **l.** Click Save.
12.	Configure a second level map for GTP flow sampling, the flow sampling map. If there is not a match to an IMSI in the whitelist, the traffic flow is sampled based on the rules in this map. Accepted packets are forwarded to a load balancing port group.	**a.** On the left navigation pane, go to **Inventory** and from **Traffic** select **Maps > Maps > Maps**. **b.** Click New **c.** Configure the map. ■ Type an name in the Alias field. For example GTP-Sample-01. ■ Select Second Level for Type ■ Select Flow Sample for Subtype ■ Select the GigaSMART operation for flow sampling configured in Task 11 from the GSOP list. ■ Select the virtual port configured in Task 3 for Source ■ Select PG-Sample for Destination **d.** Create the following flow sample rules: ■ Percentage 50, IMEI 01416800*, IMSI 31* ■ Percentage 80, IMSI 46* ■ Percentage 25, MSISDN 1509* ■ Percentage 15, IMEI 01400*, imsi 31* ■ Percentage 20, IMSI 31*, MSISDN 1909* **e.** Click Save.

Example 5: GTP Flow Sampling with Multiple Maps

Example 5 includes multiple GTP flow sampling maps, which provide a more granular selection of tool ports for flow sampling.

In Example 5, traffic from network ports go to the three first level maps (GTP-Control, GTP-User, and Fragments-Not_First) and then to the virtual port (vport1). If there is a match to an IMSI in the whitelist (VoLTE_1MM), it is forwarded to the port-group (PG-Whitelist-1) for load balancing.

NOTE: In Example 5, the tool ports in the port group are on the same node as the GigaSMART group and GigaSMART operation.

If there is not a match to an IMSI in the whitelist, the traffic flow is sampled based on the rules in the four flow sampling maps (GTP-Sample-1 to GTP-Sample-4).

The flow sampling rules in each map specify IMSI, IMEI, and MSISDN numbers, as well as the percentage to sample. Packets are then accepted or rejected. Accepted packets are forwarded to the port-group (PG-Sample-1 to PG-Sample-4) for load balancing. Rejected packets are dropped. Packets that do not match a rule will be passed to subsequent maps, in this example, to a shared collector.

Task	Description	UI Steps
1.	Create port groups, specifying the tool ports and enabling load balancing.	**a.** Go to **System > Ports >** select **Ports > Port Groups > All Port Groups.** **b.** Click New. **c.** Type PG-Sample-1 in the Alias field. **d.** Select SMART Load Balancing **e.** Click in the Ports field and select the tool ports for the port group. **f.** Click Save. **g.** Repeat steps 2 through 6, to create a port groups with the aliases
2.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**a.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART Groups**. **b.** Click New. **c.** Type an alias in the Alias field and enter an engine port in the Port List field. **d.** Click Save.
3.	Create a virtual port.	**a.** From the left navigation pane, go to **System > GigaSMART >Virtual Ports**. **b.** Type vport1 in the Alias field. **c.** Select the GigaSMART Groups created in Task 2 from the GigaSMART Group list. **d.** Click Save.

Task	Description	UI Steps
4.	Configure three first level maps. **Note**: In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**a.** Configure the first map as follows: ▪ Alias: GTP-Control ▪ Type and subtype: First Level By Rule ▪ Source: network ports (for example, 10/1/x5, 10/3/x1,10/6/q1) ▪ Destination: virtual port created in Task 2. ▪ Rule: Pass, Bi Directional, Port Destination 2123 ▪ Save the map **b.** Configure the second map as follows: ▪ Alias: GTP-User ▪ Type and subtype: First Level By Rule ▪ Source: Same network ports as first map. ▪ Destination: virtual port created in Task 2. ▪ Rule: Pass, Bi Directional, Port Destination 2152 ▪ Save the map **c.** Configure the third map as follows: ▪ Alias: Fragments-Not-First ▪ Type and subtype: First Level By Rule ▪ Source: Same network ports as first map ▪ Destination: virtual port created in Task 2 ▪ Rule: Pass, IPv4 Fragmentation and select allFragNoFirst **d.** Save the map
5.	Create the GTP whitelist.	**a.** From the left navigation pane, go to **System > GigaSMART >GTP Whitelist.** **b.** Click New. **c.** Enter VoLTE_1MM in the Alias field. **d.** Go to Task 6.
6.	Fetch whitelist files from a specified location to populate the GTP whitelist.	**a.** On the GTP Whitelist page, select Bulk Upload. **b.** Select Bulk Entry Operation for IMSI Upload Type **c.** Select Upload from URL from the Bulk Upload Type list. **d.** Enter the URL in the Enter Remote URL field. For example, http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx **e.** Click Save.
7.	(Optional) Add a single IMSI to the GTP whitelist.	**a.** On the GTP Whitelist page, select Individual Entry Operation. **b.** Select Append for Operation Type **c.** Enter the IMSI entry in the Individual IMSI Entry field.

Task	Description	UI Steps
		d. Click Save.
8.	Associate the GigaSMART group to the GTP whitelist.	**a.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART Groups** **b.** Click New. **c.** Type gsg1 in the Alias field. **d.** Under GTP Whitelist, click on the GTP Whitelist Alias field and select VolTE_1MM. **e.** Click Save.
9.	Configure the GigaSMART operation for GTP whitelisting.	**a.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations (GSOP) > GigaSMART Operation.** **b.** Click New. **c.** Select the GigaSMART Group gsg1 created in Task 8 from the GigaSMART Groups list. **d.** Enter gtp-whitelistl in the Alias field. **e.** Select GTP Whitelist from the GigaSMART Operations (GSOP) list **f.** Select Load Balancing from the GigaSMART Operations (GSOP) list. **g.** Configure Load Balancing as follows: ▪ Select Stateful ▪ Set Type to GTP ▪ Select Hashing ▪ Select IMSI **h.** Click Save.
10.	Configure a second level map for GTP whitelisting, the whitelist map. If there is a match to an IMSI in the whitelist (VoLTE_1MM), it is forwarded to a load balancing port group.	**a.** On the left navigation pane, go to **Inventory** and from **Traffic** select **Maps > Maps > Maps**. **b.** Click New **c.** Configure the map. ▪ Enter GTP-Whitelist in the Alias field. ▪ Select Second Level for Type ▪ Select By Rule for Subtype ▪ Select gtp-whitelist from the GSOP list. ▪ Select the virtual port vport1 configured in Task 3 for Source ▪ Select port group PG-Whitelist-2 for Destination **d.** Click Save.
11.	Configure the GigaSMART	**a.** From the left navigation pane, go to **System >**

Task	Description	UI Steps
	operation for GTP flow sampling.	**GigaSMART >GigaSMART Operations (GSOP) > GigaSMART Operation.** **b.** Click New. **c.** Select the GigaSMART Group created in Task 8 from the GigaSMART Groups list. **d.** Enter gtp-flowsample-1 in the Alias field. **e.** Select Flow Sampling from the GigaSMART Operations (GSOP) list and then select the Flow Sampling-GTP option. **f.** Select Load Balancing from the GigaSMART Operations (GSOP) list. **g.** Configure Load Balancing as follows: ▪ Select Stateful ▪ Set Type to GTP ▪ Select Hashing ▪ Select IMSI **h.** Click Save.
12.	Configure a second level map for GTP flow sampling, the first flow sampling map. This map has 12 rules. Traffic flow is sampled based on the rules in this map. Accepted packets are forwarded to a load balancing port group.	**a.** On the left navigation pane, go to **Inventory** and from **Traffic** select **Maps > Maps > Maps**. **b.** Click New. **c.** Configure the map. ▪ Type GTP-Sample-1 in the Alias field ▪ Select Second Level for Type ▪ Select Flow Sample for Subtype. ▪ Select the virtual port vport1 configured in Task 3 for the Source ▪ Select a port grroup PG-Sampl-1 for the Destination ▪ Select the GigaSMART Operation configured in Task for from the GSOP list **d.** Use the Add a Rule button to create the following flow sampling rules: ▪ Percentage 75, IMSI 3182609833*, IMEI 35609506* ▪ Percentage 10, IMSI 3182609834*, IMEI 3560950* ▪ Percentage 20, IMSI 31826098350*, IMEI 356095* ▪ Percentage 20, IMSI 31826098351*, IMEI 35609* ▪ Percentage 20, IMSI 31826098352*, IMEI 3560* ▪ Percentage 20, IMSI 31826098353*, IMEI 356* ▪ Percentage 20, IMSI 31826098354*, IMEI 35* ▪ Percentage 20, IMSI 31826098355*, IMEI 31* ▪ Percentage 20, IMSI 31826098356*, IMEI 356095*

Task	Description	UI Steps
		▪ Percentage 20, IMSI 31826098356*, IMEI 356095* ▪ Percentage 20, IMSI 31826098357*, IMEI 3560* ▪ Percentage 20, IMSI 31826098358*, IMEI 35* ▪ Percentage 20, IMSI 31826098359*, IMEI 356095* **e.** Click Save.
13.	Configure a second level map for GTP flow sampling, the second flow sampling map. This map has 12 rules. Traffic flow is sampled based on the rules in this map. Accepted packets are forwarded to a load balancing port group.	**a.** On the left navigation pane, go to **Inventory** and from **Traffic** select **Maps > Maps > Maps**. **b.** Click New. **c.** Configure the map. ▪ Type GTP-Sample-2 in the Alias field ▪ Select Second Level for Type ▪ Select Flow Sample for Subtype. ▪ Select the virtual port vport1 configured in Task 2 for the Source ▪ Select a tool port group PG-Sample-2 for the Destination ▪ Select flow-sample-1 configured in Task 11 for from the GSOP list **d.** Use the Add a Rule button to create the following flow sampling rules: ▪ Percentage 30, IMSI 3182609836*, IMEI 35609506* ▪ Percentage 5, IMSI 3182609837*, IMEI 356095062* ▪ Percentage 50, IMSI 31826098380*, IMEI 356095062* ▪ Percentage 50, IMSI 3182609838l*, IMEI 35609506* ▪ Percentage 50, IMSI 31826098382*, IMEI 3560950* ▪ Percentage 50, IMSI 31826098383*, IMEI 356095* ▪ Percentage 50, IMSI 31826098384*, IMEI 35* ▪ Percentage 50, IMSI 31826098385*, IMEI 356* ▪ Percentage 50, IMSI 31826098386*, IMEI 3560* ▪ Percentage 50, IMSI 31826098387*, IMEI 35609* ▪ Percentage 50, IMSI 31826098388*, IMEI 356095* ▪ Percentage 50, IMSI 31826098389*, IMEI 3560950* **e.** Click Save.
14.	Configure a second level map for GTP flow sampling, the third flow sampling map. This map has 5 rules. Traffic flow is sampled based on the rules in this map. Accepted packets are	**a.** On the left navigation pane, go to **Inventory** and from **Traffic** select **Maps > Maps > Maps**, **b.** Click New. **c.** Configure the map. ▪ Type GTP-Sample-3 in the Alias field ▪ Select Second Level for Type

Task	Description	UI Steps
	forwarded to a load balancing port group.	▪ Select Flow Sample for Subtype ▪ Select the virtual port vport1 configured in Task 3 for the Source ▪ Select a port group PG-Sample-3 port for the Destination ▪ Select flow-sample-1 configured in Task 11 for from the GSOP list **d.** Use the Add a Rule button to create the following flow sampling rules: ▪ Percentage 10, IMSI 31826098390*, IMEI 35609506* ▪ Percentage 10, IMSI 31826098391*, IMEI 35609506* ▪ Percentage 10, IMSI 31826098392*, IMEI 35609506* ▪ Percentage 10, IMSI 31826098393*, IMEI 35609506* ▪ Percentage 10, IMSI 31826098394*, IMEI 35609506* **e.** Click Save.
15.	Configure a second level map for GTP flow sampling, the fourth flow sampling map. This map has one rule. Traffic flow is sampled based on the rules in this map. Accepted packets are forwarded to a load balancing port group.	**a.** On the left navigation pane, go to **Inventory** and from **Traffic** select **Maps > Maps > Maps**. **b.** Click New. **c.** Configure the map. ▪ Type GTP-Sample-4 in the Alias field ▪ Select Second Level for Type ▪ Select Flow Sample for Subtype ▪ Select the virtual port vport1 configured in Task 3 for the Source ▪ Select a tool port for the Destination ▪ Select flow-sample-1 configured in Task 11 for from the GSOP list **d.** Use the Add a Rule button to create the following flow sampling rule: ▪ Percentage 10, IMSI 31826098429*, IMEI 35609506* **e.** Click Save.
16.	Configure a collector map for any packets that do not match other rules.	**a.** On the left navigation pane, go to **Inventory** and from **Traffic** select **Maps > Maps > Maps**. **b.** Click New. **c.** Configure the map. ▪ Type GTP-Collector in the Alias field ▪ Select Second Level for Type ▪ Select Collector for Subtype ▪ Select the virtual port vport1 configured in Task 3 for the Source **d.** Click Save.

Example 6: APN for GTP Whitelisting, GTP Flow Sampling

Example 7 specifies APN patterns for GTP whitelisting and GTP flow sampling.

In Example 7, traffic from network ports go to the two first level maps (gtp_to_v1_c and gtp_to_v1_u) and then to the virtual port (v1).

In the whitelist map, if there is a match to the APN pattern and if the IMSI is present in the whitelist (IMSI), packets are forwarded to a tool port.

If there is not a match to an IMSI in the whitelist, the traffic is flow sampled based on the APN pattern in the flow sampling map. Accepted packets are forwarded to the same tool port as specified in the whitelist map.

Any unmatched traffic goes to a shared collector that sends it to a different tool port.

Task	Description	UI Steps
1.	Configure a network port and two tool ports and enable them.	**a.** Go to **System > Ports > Ports > All Ports.** **b.** Click Quick Port Editor. **c.** Configure a network port. Port 22/3/x3 in this example. **d.** Configure two tool ports. Port 22/4/x18 and 22/4/x19 in this example. **e.** Admin enable the ports by selecting Enable for each port. **f.** Click OK.
2.	Configure a GigaSMART group and associate it with two GigaSMART engine port.	**a.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART Groups.** **b.** Click New. **c.** Type an gsg2 in the Alias field. **d.** In the Port List field, select the engine ports, which are 22/2/e1 and 22/2/e2 in this example **e.** Click Save.
3.	Create a virtual port.	**a.** From the left navigation pane, go to **System > GigaSMART >Virtual Ports.** **b.** Type v1 in the Alias field. **c.** Select the GigaSMART Group created in Task 2 from the GigaSMART Group list. **d.** Click Save.

Task	Description	UI Steps
4.	Configure two first level maps, one for control traffic and one for user traffic.	**a.** Configure the first map as follows: ▪ Alias: gtp_to_v1_c ▪ Type and subtype: First Level By Rule ▪ Source: 22/3/x3 ▪ Destination: virtual port created in Task 2. ▪ Rule 1: Pass, Bi Directional, Port Destination 2123 ▪ Rule 2: Pass, Bi Directional, Port Destination 2122 **b.** Save the map **c.** Configure the second map as follows: ▪ Alias: gtp_to_v1_u ▪ Type and subtype: First Level By Rule ▪ Source: 22/3/x3. ▪ Destination: virtual port created in Task 2. ▪ Rule 1: Pass, Bi Directional, Port Destination 2152 ▪ Rule 1: Pass, Bi Directional, IPv4 Fragmentation, Value: allFragNoFirst. **d.** Save the map
5.	Create the GTP whitelist.	**a.** From the left navigation pane, go to **System > GigaSMART >GTP Whitelist.** **b.** Click New. **c.** Enter gtp-whitelist in the Alias field **d.** Go to Task 6.
6.	Fetch whitelist files from a specified location to populate the GTP whitelist.	**a.** On the GTP Whitelist page, select Bulk Upload. **b.** Select Bulk Entry Operation for IMSI Upload Type **c.** Select Upload from URL from the Bulk Upload Type list. **d.** Select Append. **e.** Enter the URL in the Enter Remote URL field. For example, http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx **f.** Click Save.
7.	(Optional) Add a single IMSI to the GTP whitelist.	**a.** On the GTP Whitelist page, select Individual Entry Operation. **b.** Select Append for Operation Type **c.** Enter the IMSI entry in the Individual IMSI Entry field. **d.** Click Save.
8.	Associate the GigaSMART group to the GTP whitelist.	**a.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART Groups**

Task	Description	UI Steps
		b. Select GS Group gsg2 created in Task 2 and click Edit **c.** Under GTP Whitelist, click on the GTP Whitelist Alias field and select. gtp-whitelist **d.** Click Save.
9.	Configure the GigaSMART operation for GTP whitelisting.	**a.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations (GSOP) > GigaSMART Operation.** **b.** Click New. **c.** Select the GigaSMART Group gsg2 created in Task 2 and associated with the GTP whitelist in Step 8. **d.** Enter gtp-correlat_gsp_wl in the Alias field. **e.** Select GTP Whitelist from the GigaSMART Operations (GSOP) list **f.** Select Load Balancing from the GigaSMART Operations (GSOP) list. **g.** Configure Load Balancing as follows: ▪ Select Stateful ▪ Set Type to GTP ▪ Select Hashing ▪ Select IMSI **h.** Click Save.
10.	Configure a second level map for GTP whitelisting, the whitelist map. If there is a match to the APN pattern and if IMSI is present in the whitelist (IMSI), it is forwarded to a tool port.	**a.** On the left navigation pane, go to **Inventory** and from **Traffic** select **Maps > Maps > Maps**. **b.** Click New **c.** Configure the map. ▪ Enter GTP-Whitelist in the Alias field. ▪ Select Second Level for Type ▪ Select Flow Whitelist for Subtype ▪ Select gtp-correlate_gsg_wl from the GSOP list. ▪ Select the virtual port v1 configured in Task 3 for Source ▪ Select 22/4/x18 for Destination ▪ Rule 1: GTP, APN: mobile.com **d.** Click Save.
11.	Configure the GigaSMART operation for GTP flow sampling.	**a.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations (GSOP) > GigaSMART Operation.** **b.** Click New. **c.** Select the GigaSMART Group gsg2 created in Task 2 and associated with the GTP whitelist in Step 8.

Task	Description	UI Steps
		d. Enter gtp-correlat_gsp_fs in the Alias field. **e.** Select GTP Whitelist from the GigaSMART Operations (GSOP) list **f.** Select Load Balancing from the GigaSMART Operations (GSOP) list. **g.** Configure Load Balancing as follows: ▪ Select Stateful ▪ Set Type to GTP ▪ Select Hashing ▪ Select IMSI **h.** Click Save
12.	Configure a second level map for GTP flow sampling, the flow sampling map. If there is not a match to an IMSI in the whitelist, the traffic flow is sampled based on the APN pattern in this map. Accepted packets are forwarded to the same tool port as specified in the whitlelist map	**a.** On the left navigation pane, go to **Inventory** and from **Traffic** select **Maps > Maps > Maps**. **b.** Click New. **c.** Configure the map. ▪ Type from_vp_fs1 in the Alias field ▪ Select Second Level for Type ▪ Select Flow Sample for Subtype. ▪ Select the virtual port v1 configured in Task 3 for the Source ▪ Select a 22/4/x18 for the Destination ▪ Select the GigaSMART Operation gtp-correlate_ gsg_fs ▪ Rule 1: GTP, Percentage: 100, APN: imsi* **d.** Click Save.
13.	Add a shared collector for any unmatched traffic from the virtual port and send it to a different tool port.	**a.** On the left navigation pane, go to **Inventory** and from **Traffic** select **Maps > Maps > Maps**. **b.** Click New. **c.** Configure the map. ▪ Type from_vp_scoll in the Alias field ▪ Select Second Level for Type ▪ Select Collector for Subtype ▪ Select the virtual port v1 configured in Task 3 for the Source **d.** Click Save.

Example 7: APN for FTP Whitelisting, APN and QCI for GTP Flow Sampling

Example 6 specified APN patterns for GTP whitelisting and GTP flow sampling. It also specifies QCI for GTP flow sampling.

In Example 7, traffic from network ports go to the two first level maps (gtp_to_v1_c and gtp_to_v1_u) and then to the virtual port (v1).

In the whitelist map, if there is a match to the APN pattern and if the IMSI is present in the whitelist (IMSI), packets are forwarded to a tool port.

If there is not a match to an IMSI in the whitelist, the traffic is flow sampled based on the APN pattern and QCI value in the flow sampling map. Accepted packets are forwarded to the same tool port as specified in the whitelist map. Only 50% of traffic with QCI 5 is sent to the tool port.

Any unmatched traffic goes to a shared collector that sends it to a different tool port.

Task	Description	UI Steps
1.	Configure a network port and two tool ports and enable them.	**a.** Go to **System > Ports > Ports > All Ports.** **b.** Click Quick Port Editor. **c.** Select a port (for example, 22/2/x3) and set Type to Network. **d.** Select a port (for example, 22/2/x18) and set Type to Tool **e.** Select a second port (for example, 22/2/x19) and set Type to Tool. **f.** Select Enable for Admin on the network and two tool ports. **g.** Click OK.
2.	Configure a GigaSMART group and associate it with two GigaSMART engine ports	**a.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART Groups.** **b.** Click New. **c.** Type gsg2 in the Alias field. **d.** Click in the Port List field and select two engine ports. For example, 22/2/e1 and 22/2/e2 **e.** Click Save.
3.	Create a virtual port.	**a.** From the left navigation pane, go to **System >**

Task	Description	UI Steps
		GigaSMART >Virtual Ports. **b.** Type v1 in the Alias field. **c.** Click in the GigaSMART Group field and select the GigaSMART Group created in Task 2. **d.** Click Save.
4.	Configure two first level maps, on for control traffic and one for user traffic	**a.** Configure the first map as follows: ▪ Alias: gtp_to_v1_c ▪ Type and Subtype: First Level By Rule ▪ Enable Control Traffic. ▪ Source: 22/2/3/x3 (network port created in Task 1) ▪ Destination: v1 (virtual port created in Task 3) ▪ Rule 1: Pass, Bi Directional, Port Destination 2123 ▪ Rule 2: Pass, Bi Directional, Port Destination 2122 **b.** Save the map **c.** Configure the second map as follows: ▪ Alias: gtp_to_v1_u ▪ Type and subtype: First Level By Rule ▪ Source: 22/2/3/x3 (network port created in Task 1) ▪ Destination: v1 (virtual port created in Task 3) ▪ Rule 1: Pass, Bi Directional, Port Destination 2152 ▪ Rule 2: Pass, Bi Directional, IPv4Fragmentation allFragNoFirst **d.** Save the map
5.	Associate the GigaSMART group to the active GTP Whitelist	**a.** From the device view, select GigaSMART > GigaSMART Groups. **b.** Select the GigaSMART Group created in Task 1 and click Edit. **c.** Locate the GTP Whitelist param, and enter the alias of whitelist in the GTP Whitelist Alias field. For example, IMSI. **d.** Save the GigaSMART Group.
6.	Configure the GigaSMART operation for GTP whitelisting.	**a.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations (GSOP) > GigaSMART Operation.** **b.** Click New. **c.** Select the GigaSMART Group gsg1 created in Task 2 from the GigaSMART Groups list. **d.** Enter gtp-correlate_gsp_wl in the Alias field. **e.** Select GTP Whitelist from the GigaSMART Operations (GSOP) list

Task	Description	UI Steps
		f. Select Load Balancing from the GigaSMART Operations (GSOP) list. **g.** Configure Load Balancing as follows: ▪ Select Stateful ▪ Set Type to GTP ▪ Select Hashing ▪ Select IMSI **h.** Click Save.
7.	Configure a second level map for GTP whitelisting, the whitelist map. If there is a match to the APN pattern and if the IMSI is present in the whitelist (IMSI), packets are forwarded to a tool port.	**a.** On the left navigation pane, go to **Inventory** and from **Traffic** select **Maps > Maps > Maps**. **b.** Click New **c.** Configure the map. ▪ Alias: GTP-whitelist ▪ Type an Subtype: Second Level Flow Whitelist ▪ Source: v1 (virtual port created in Task 3) ▪ Destination: 22/4/x18 ▪ GSOP: gtp-corelate_gsg_wl ▪ Select gtp-whitelist from the GSOP list. ▪ Rule: GTP, APN: mobile.com **d.** Click Save.
8.	Configure the GigaSMART operation for GTP flow sampling.	**a.** From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations (GSOP) > GigaSMART Operation.** **b.** Click New. **c.** Select the GigaSMART Group created in Task 2 from the GigaSMART Groups list. **d.** Enter gtp-corelate_gsg_fs in the Alias field. **e.** Select Flow Sampling from the GigaSMART Operations (GSOP) list and then select the Flow Sampling-GTP option. **f.** Select Load Balancing from the GigaSMART Operations (GSOP) list. **g.** Configure Load Balancing as follows: ▪ Select Stateful ▪ Set Type to GTP ▪ Select Hashing ▪ Select IMEI **h.** Click Save.
9.	Configure a second level map	**a.** On the left navigation pane, go to **Inventory** and

Task	Description	UI Steps
	for GTP flow sampling, the flow sampling map. If there is not a match to an IMSI in the whitelist, the traffic flow is sampled based on the APN pattern in this map. Accepted packets are forwarded to the same tool port as specified in the whitelist map.	from **Traffic** select **Maps > Maps > Maps**. **b.** Click New. **c.** Configure the map. ▪ Alias: from_vp_fs1 ▪ Type and Subtype: Second Level Flow Sample ▪ Source: vp1 ▪ Destination: 22/4/x18 ▪ GSOP: gtp-corelate_gsg_fs ▪ Rule 1: GTP, APN: *imsi*, QCI: 5, Percentage: 50 ▪ Rule: 2: GTP, IMSI: ims*, Percentage 100 **d.** Click Save.
10.	Add a shared collector for any unmatched traffic from the virtual port and send it to a different tool port.	**a.** On the left navigation pane, go to **Inventory** and from **Traffic** select **Maps > Maps > Maps**. **b.** Click New. **c.** Configure the map: ▪ Alias: from_vp_scoll ▪ Type and Subtype: Regular Collector ▪ Source: v1 ▪ Destination: 22/4/x19 **d.** Click Save.

4G/5G Traffic Monitoring using UPN

In 4G/5G traffic monitoring using User Processing Node (UPN), the UPN processes the subscriber information that is extracted from the PFCP packets and correlates with the GTP-u traffic without the CPN. When the traffic contains the user's field information such as IMSI, IMEI, you can use the 4G/5G traffic monitoring using UPN.

The PFCP allows optional Informational Elements (IE) that contain SUPI/IMSI, PEI/IMEI, and GPSI/MSISDN information of the subscriber during PFCP session establishment request.

NOTE: The amount of time to restore from backup can take up to 11 minutes.

The UPN performs the following activities in the Standalone mode:

- Processes the PFCP session establishment request and extracts the SUPI/IMSI, PEI/IMEI, GPSI/MSISDN User IP, and TEID for both endpoints. The information in the PFCP traffic is used to populate the UPN's session table.
- Correlates GTP-U traffic based on the subscriber information from PFCP. The GTP-U look-up is correlated based on the IP and TEID.

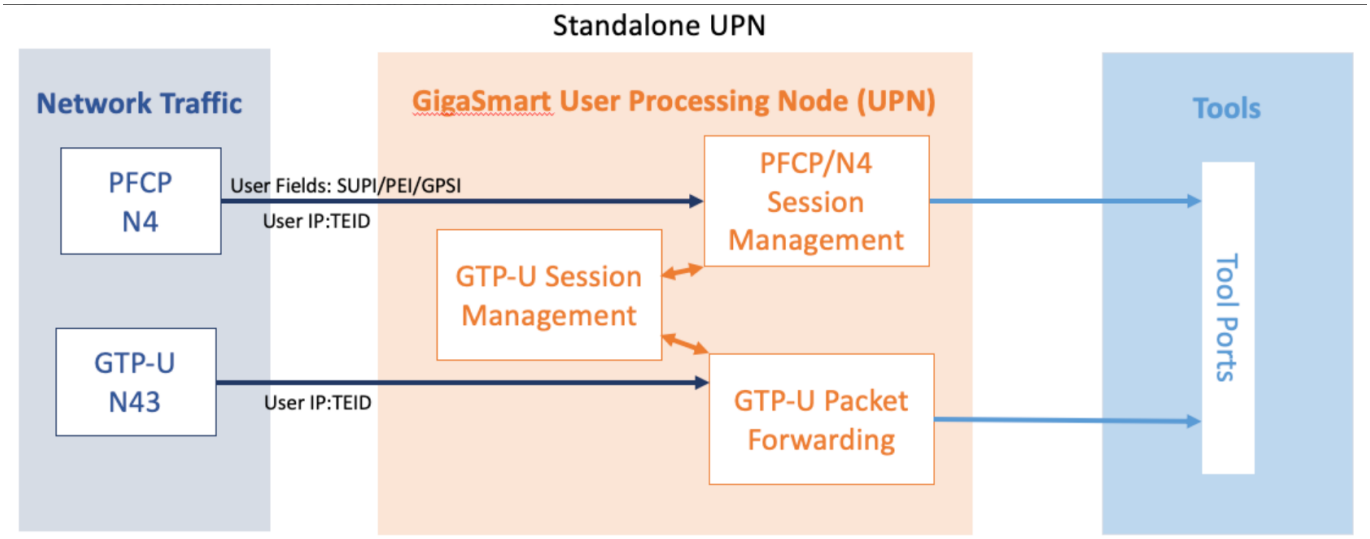
What is PFCP?

PFCP (Packet Forwarding Control Protocol) is a 3GPP Protocol that is communicated on the Sx/N4 interface between the Control Plane (CP) elements and User Plane (UP) elements. The CP element programs the UP element with policies on how to forward packets. PFCP packets convey information in the form of Information Elements (IE). PFCP allows an optional IE called User Fields that contains IMSI/SUPI, IMEI/PEI, MSISDN/GPSI information of the Subscriber during PFCP Session Establishment Request.

The PFCP (Packet Forwarding Control Protocol) is a 3GPP Protocol that is communicated on Sxa, Sxb, N4 interface between the Control Plane (CP) elements and User Plane (UP) elements.

The CP element programs the UP element with policies on how to forward packets. The CP function controls packet processing in the UP function by establishing, modifying, and deleting the PFCP session.

The following diagrams explain the functioning of UPN in Stand-alone mode:



Rules and Notes

- A Stand-alone UPN traffic monitoring session can be created only when there is atleast SUPI in the user fields.
- You cannot perform RAN correlation or Network Slicing Correlation in Stand alone UPN traffic monitoring
- UPN does not receive SFFP messages in the User field parse mode.

Configure 4G/5G Traffic Monitoring using UPN

The following are the points to remember while configuring 4G/5G traffic monitoring using UPN:

- You must enable the stand-alone mode in UPN while creating the solution in Ansible.
- Once the stand-alone mode is enabled, UPN cannot connect with the CPN, and you cannot change the mode. To change the mode, you need to delete the UPN from the solution and
- add a new one.

PFCP Messages

Session Related Messages

These messages are exchanged between Control Plane and User Plane function at the subscriber level. These messages are load balanced based on the **IMSI** of the subscriber.

- Session establishment request/response
- Session modification request/response
- Session deletion request/response
- Session report request/response

Node Related Messages

The PFCP node related message packets are broadcast to all the ports which are part of the port-group in FS/WL maps. These messages are not specific to any subscriber. The load balancing based on **IMSI** is not possible for these packets. They continue to broadcast as per the existing behavior.

- Heartbeat request/response
- PFCP PFD management request/response
- PFCP association setup request/response
- PFCP association update request/response
- PFCP association release request/response
- Version Not supported response

- PFCP Node report request/response
- PFCP session set deletion request/response

PFCP Load Balancing

The PFCP packets received are load balanced based on the **IMSI** present in the "IE Type: UserID".

Once the PFCP packet is received, the mobility application will parse the packet and create or update or delete the corresponding sessions or tunnels in the session table. The IMSI present in "IE Type: UserId" is stored in the session table and used for load balancing. Once sessions or tunnels are created, map lookup is performed for the PFCP packets. If there is a match in the IMSI, the packet is sent to only one port from the configured **Load Balancing** application based on the **IMSI** hashing. The hash value string from IMSI string present in the "IE Type: UserId" is added to the packet before sending to the **Load Balancing** application. The **Load Balancing** application will do **Load Balancing** logic based on this hash string and choose any one port to send the packet out.

Licensing Requirement:There is no need of a separate license for PFCP Load balancing. The existing license for Standalone UPN is sufficient for this feature.

Single Engine: The PFCP packet is sent to only one GigaSMARTS engine. Based on the map lookup, the packet will be forwarded to the **Load Balancing** application and sent to only one port based on the IMSI hashing.

Grouped Engine: ThePFCP packet is sent to all the GigaSMART engines configured as part of the engine group configuration. The PFCP packets are processed as per the existing design, but the packets are sent out only through one port in the **LoadBalancing** application in the leader engine**.**

Map Configuration

The 4G and 5G traffic passes through a single map in standalone mode. The user can configure a 4G map through which both 4G and 5G traffic passes. You do not need to configure separate maps for 4G and 5G traffic.

> **NOTE:** The single map configuration is applicable only for standalone UPN. For legacy GTP and 5G user need to configure 4G and 5G maps separately

Required License for UPN Standalone Single Map type: 5G.

- For 5G and GTP flow sample with rule percentage (0 or 100), the **5GC** and **GTP_MAX** licenses are required.
- For 5G and GTP flow sample with rule percentage in all ranges (between 0 to 100) or for GTP Whitelist, the **5GC**, **GTP_MAX**, and **FVUE** licenses are required.

Stateful Session Recovery

It is crucial to back up the information from the correlated session table to avoid discrepancies during engine reboots and connection failures. After a successful reboot, the information saved as part of the persistence file is populated back into the mobility database.

Mobility applications require stateful session recovery to persist their data between node reboots and connection failures. This feature aims to extend the support available in LTE/5G mobility applications to Standalone UPN functionality. The session table is backed up every 10 minutes. You cannot change the backup time interval.

File Age Timeout— Specifies the time the backup file is valid, in minutes. After this time expires, the backup file is stale. The default is 30 minutes.

<10 ~ 1440> - 10 minutes is the minimum and 1440 minutes is the maximum value for file-age-timeout.

Restart Age Timeout—Specifies the time interval following a reboot for aging-out sessions, in minutes. This is a shorter interval than that specified using the gtp-flow timeout. The gtp-flow timeout disconnects a GTP session if it has been inactive for the timeout value, which has a default of 8 hours. The restart-age-timeout default is 30 minutes.

<10 ~ 1440> - 10 minutes is the minimum and 1440 minutes is the maximum value for restart-age-timeout

GTP Persistence Interval—Specifies the time interval between backups, in minutes. The default is 10 minutes.

<10 ~ 1440> - 10 minutes is the minimum and 1440 minutes is the maximum value for restart-age-timeout.



5G Stateful Session Recovery

Required License: 5G Correlation

5G stateful session recovery provides session persistence for GigaSMART 5G applications, including 5G flow filtering, 5G forward listing, and 5G flow sampling. In this method of recovery, 5G sessions are backed up periodically so that they can be recovered faster after a GigaSMART line card reboot or a node reboot.

5G stateful session recovery requires additional memory for storing backups. GigaVUE-HC3 has the required memory.

Using 5G stateful session recovery, the 5G session tables in the GigaSMART line card memory are periodically backed up to the control card memory on the node and stored.

You should configure an interval for how often the backups occur, such as every 10 minutes. If 5G stateful session recovery is enabled and the GigaSMART line card is rebooted, the 5G session tables are restored automatically following the reboot.

The last stored backup file is downloaded from the control card to the GigaSMART line card using FTP. The session table is repopulated from the last stored backup file to each GigaSMART engine, up to 8 engines. Packet count statistics for sessions are saved and restored.

Depending on the size of the session table, the amount of time to restore from the backup might take as much as 3 minutes. During that interval, traffic is blocked to the virtual port on the GigaSMART line card. Once the session table is read and populated, traffic is allowed.

Depending on the interval between backups, there could be differences between the stored state and the current state of the system, for example, map configuration could change, or sessions could be added, modified, or deleted.

Load balancing information is not persisted, so after a session table is repopulated, a session that was once sent to one load-balanced port may be sent to a different load-balanced port after the reboot. However, for SUPI-based load balancing, the traffic is sent to the same port as it was before the reboot.

5G stateful session recovery works in a cluster environment; however, the cluster leader must remain the same.

Configure 5G Stateful Session Recovery

To enable 5G stateful session recovery, as well as to configure timers, do the following:

1. On the left navigation pane, click , and then select **Physical>Nodes**.

2. From the device view, select **GigaSMART > GigaSMART groups**.

3. Click on the alias of the **GigaSMART group**.

4. Select **GTP Persistence** in the **GTP Persistence** fields under GigaSMART Parameters. The timers are pre-configured with default values.

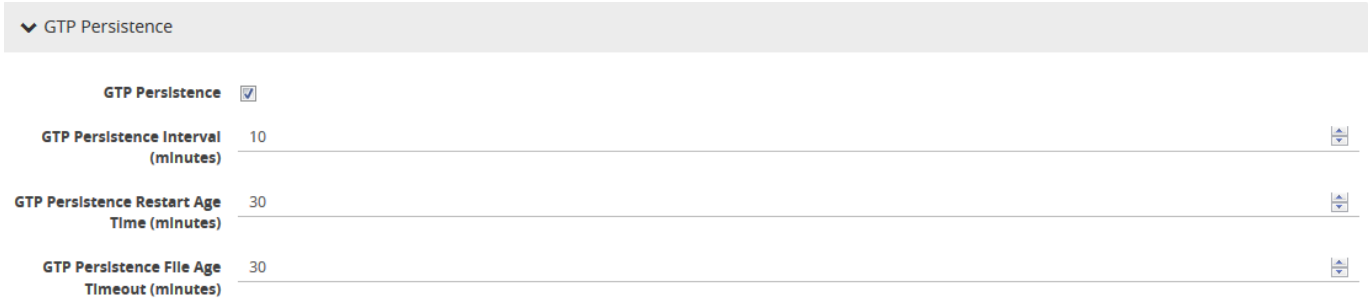


Figure 18 *GTP Persistence GigaSMART Parameters*

Use the **System** widget on the Overview page to determine the amount of memory. The size of memory is 24Gb in an upgraded system.

View Backup and Restore Information

To view the System information, select **Overview** from the Navigation pane. The amount of free and used memory is displayed in the **Memory** field.

To view backup and restore information for GTP Persistence:

1. Select **GigaSMART > GigaSMART Groups > GigaSMART Group**.

2. Click on the alias of the GigaSMART group.

A Quick View appears for the selected GigaSMART group.

3. Scroll down to GTP Persistence. In Figure 26-63, GigaSMART Group gsgrp-1_4_e1 is selected and the Quick View is displayed.

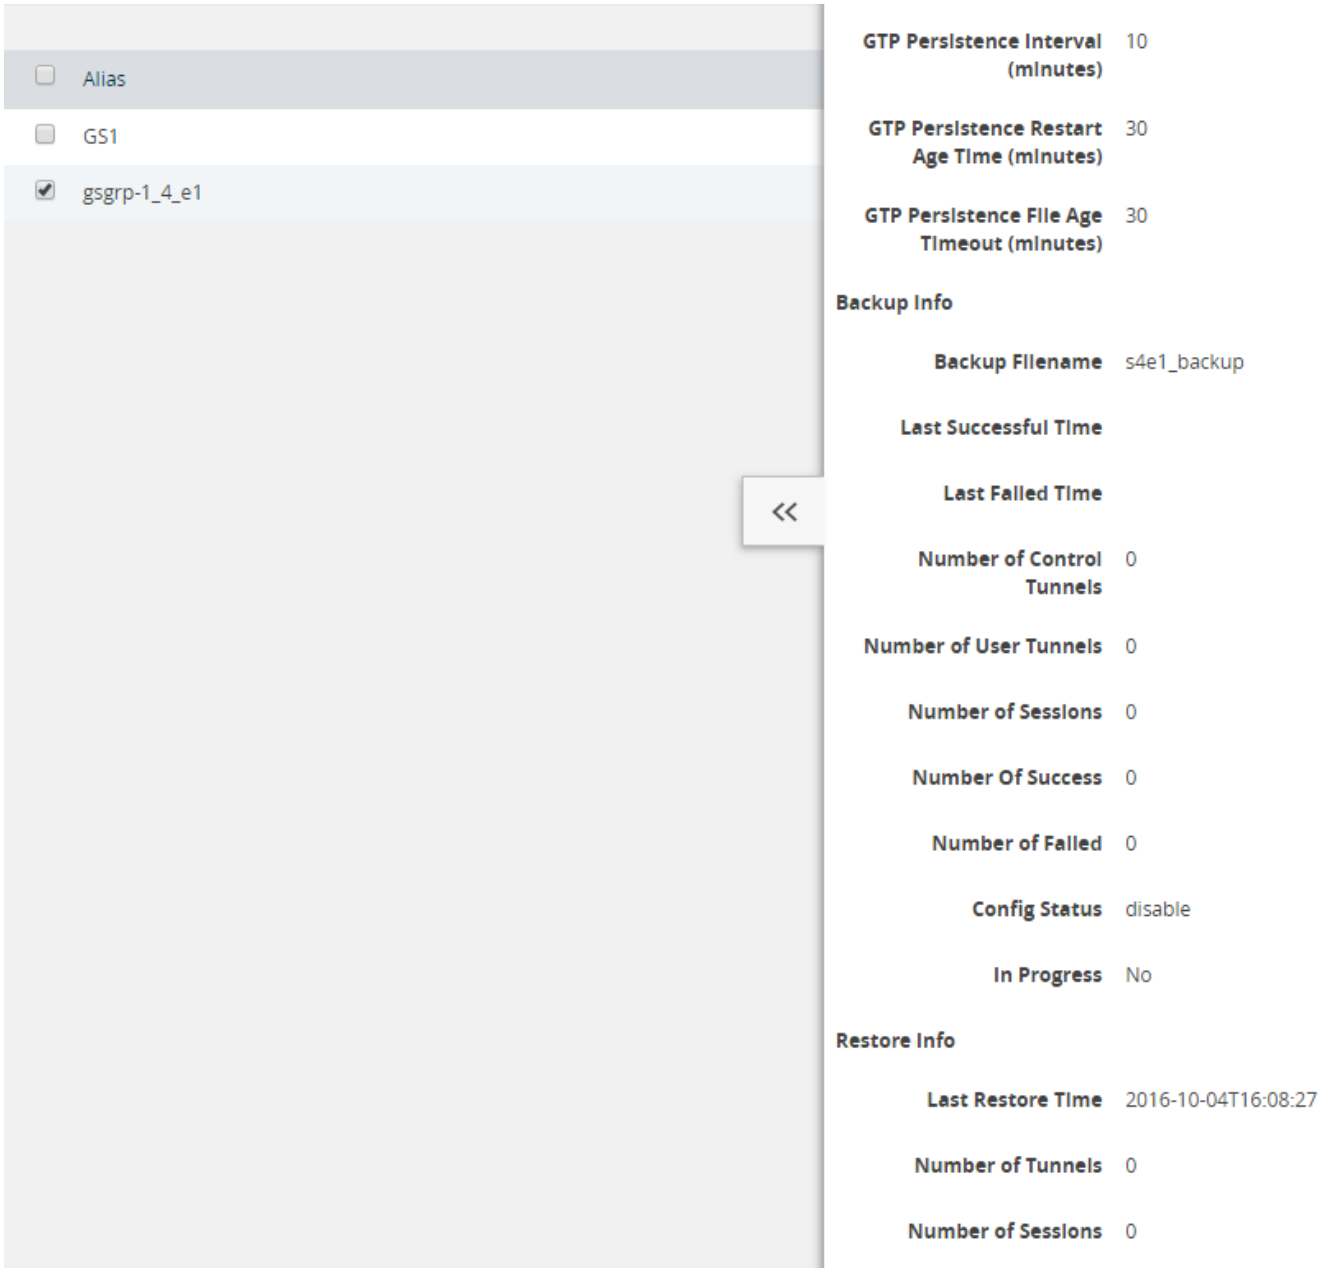


Figure 2 GTP Persistence Information

Configure GTP Overlap Mapping

The configuration of GTP forward listing and GTP flow sampling maps that are part of the GTP overlap flow sampling map group follow the same configuration considerations discussed previously in GigaSMART GTP forward listing and GTP Flow Sampling. As is the case with regular non-overlap GTP mapping, GTP forward listing selects specific subscribers based on IMSI, whereas GTP flow sampling uses map rules to select subscribers and then forward a percentage of the packets to tool ports.

Configuration Considerations

This section details certain configuration considerations that apply only to the configuration of GTP forward listing and flow sampling maps for GTP overlap flow sampling maps.

About Standalone UPN Flow Sampling Map Mode and Port Groups

A second level type map specifying Standalone UPN flow sampling map mode must be selected to configure Standalone UPN forward listing and flow sampling maps.

To configure a Standalone UPN whilelisting map in overlap flow sampling map mode, select **Type** as

Second Level and **Subtype** as **Flow Whitelist Overlap** in a map.

To configure a GTP flow sampling map in GTP overlap flow sampling map mode, select **Type**

as **Second Level** and **Subtype** as **Flow Sample Overlap** in a map .

You can configure one GTP forward listing map and one GTP flow sampling map pair that contain traffic policies corresponding to one destination port group. The load balanced port groups can contain a single port, a port range, or a GigaStream. Note that, starting in software version 4.8, port groups used in GTP overlapping maps support GigaStream.

The maximum number of port groups per single GTP overlap flow sampling map group is six.

For more information about port groups, refer to Port Groups.

Maximum Number of Port Group Members

The Maximum number of port group members can be determined as follows:

Standalone Configuration

In a standalone configuration, use the following sequence :

1. Determine the number of members per port group and add 1.

2. Multiply each port group result by each other.

3. The total multiplication should not exceed 4096.

For instance, assume the following configuration in a GTP overlap mapping group:

o Port Group 1—2 load-balanced GigaStream

o Port Group 2—3 load-balanced GigaStream

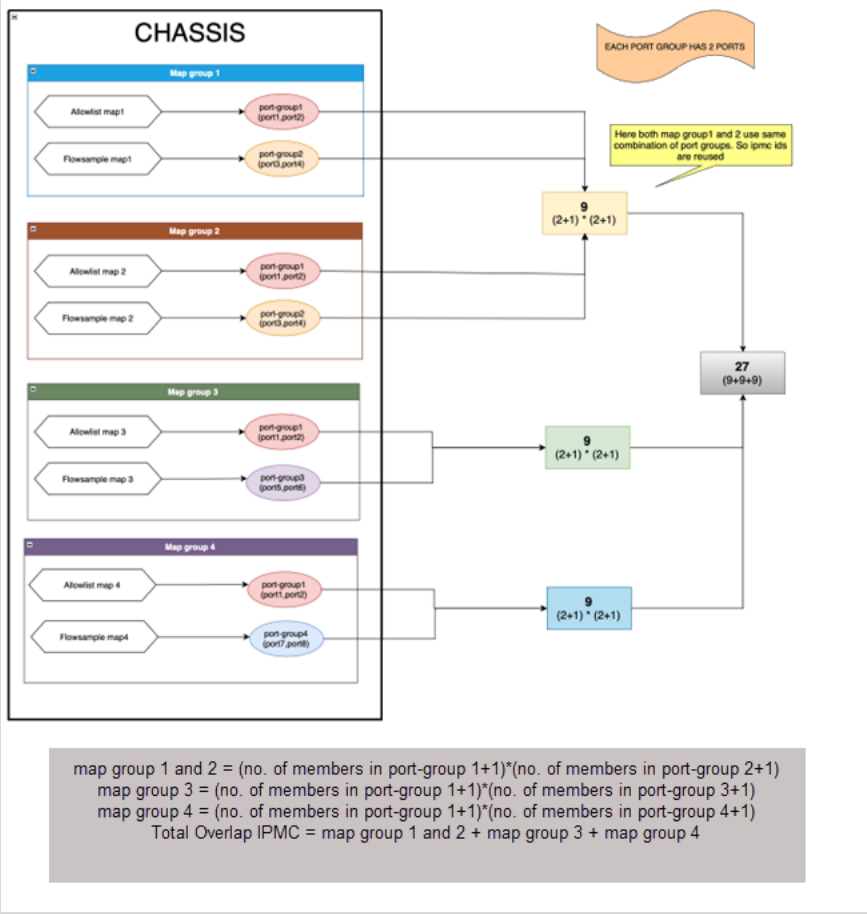
o Port Group 3—1 load-balanced tool port

o Port Group 4—1 load-balanced GigaStream

o Port Group 5—4 load-balanced tool ports. The total number becomes:(2+1)*(3+1)*(1+1)* (1+1)* (4+1) = 240

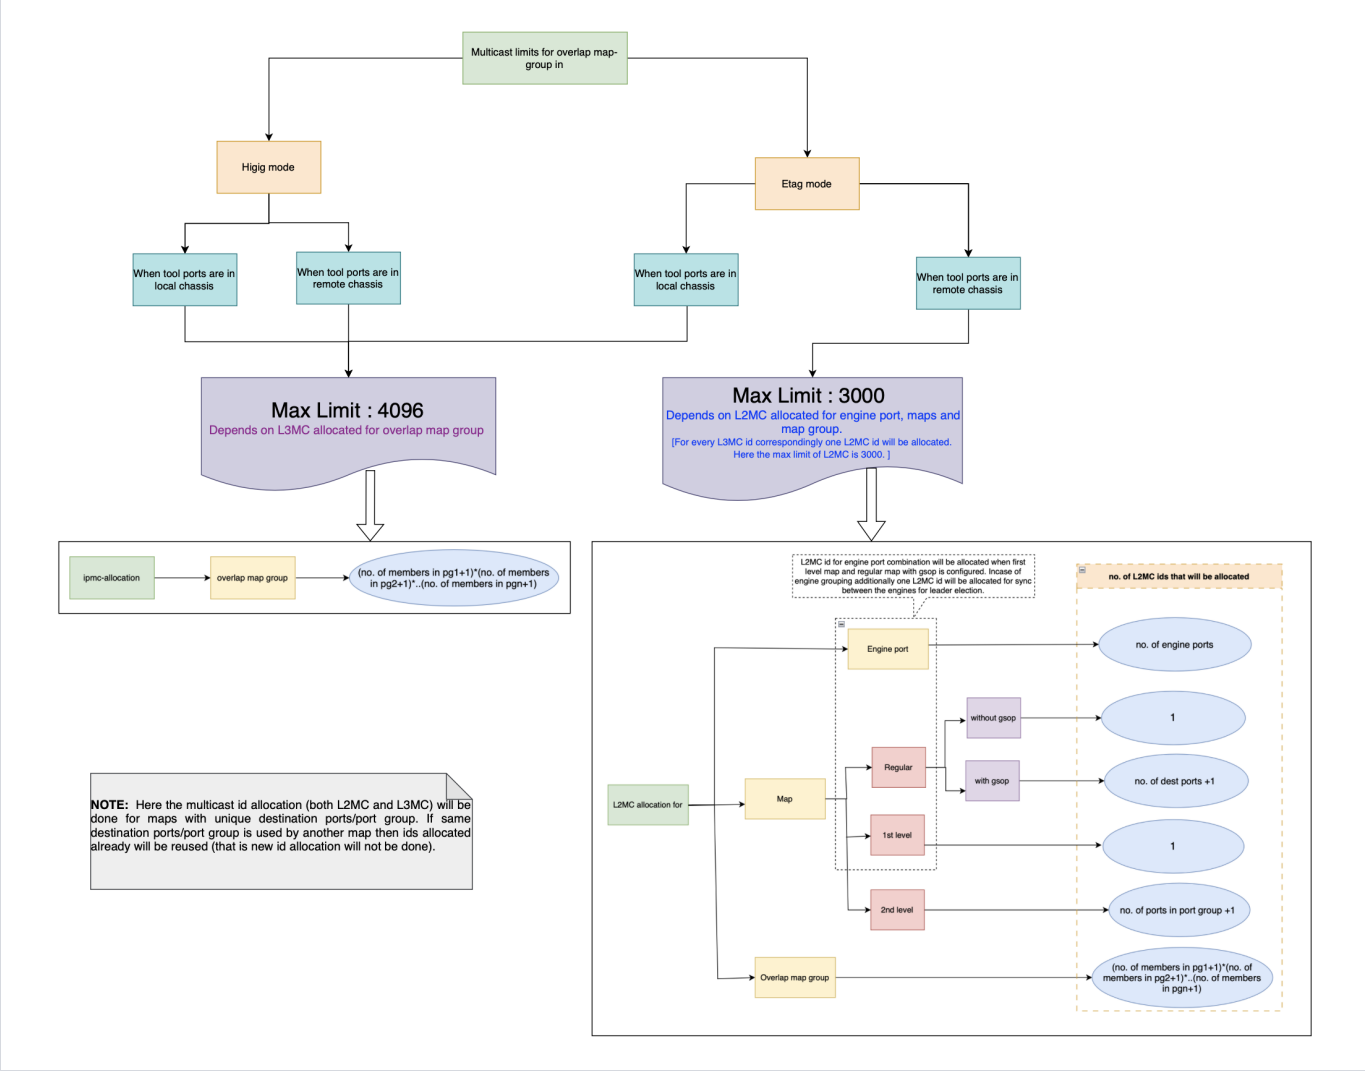
The tool configuration shown is accepted since this does not exceed the maximum number of multicast IDs (4096).

If Port Groups are utilized across multiple map groups in a single cluster, refer to the representation below on how the multicast IDs are calculated.



Cluster Configuration

In a cluster configuration, Multicast IDs can defer based on Hi-Gig mode or E-tag mode. The below is representation of the same.



Standalone UPN Flow Sampling Map Priority

Since a packet matches multiple maps independently the concept of second level map priority does not apply to Standalone UPN flow sampling maps. A standalone UPN flow sampling map pair consists of one Standalone UPN forward listing map and one Standalone UPN flow sampling map having the same destination port group. Within a Standalone UPN flow sampling map pair the forward listing map rules will be applied before the flow sampling map rules.

Virtual Port Configuration in Standalone UPN Mode

In Standalone UPN map configuration, the virtual port sending traffic to all the port groups need to be configured in Standalone UPN mode.

To configure the virtual port with Standalone UPN mode, select Standalone UPN when configuring the virtual port.

Overlap Support

Currently, a GTP packet can be destined to at most one destination tool. You can now receive multiple copies of a GTP packet that are sent to various tools (One copy per tool set).

Standalone UPN

In Standalone User Plane Node (UPN) traffic monitoring, the UPN processes the subscriber information that is extracted from the Packet Forwarding Control Protocol (PFCP) packets and correlates with the GTP-U traffic without the CPN. The PFCP allows optional Informational Elements (IE) that contain SUPI/IMSI, PEI/IMEI, and GPSI/MSISDN information of the subscriber during PFCP session establishment request.

NOTE: The amount of time to restore from backup can take up to 11 minutes.

The UPN performs the following activities in the Standalone mode:

- Processes the PFCP session establishment request and extracts the SUPI/IMSI, PEI/IMEI, GPSI/MSISDN User IP, and TEID for both endpoints. The information in the PFCP traffic is used to populate the UPN's session table.
- Correlates GTP-U traffic based on the subscriber information from PFCP. The GTP-U look-up is correlated based on the IP and TEID.

PFCP (Packet Forwarding Control Protocol)

The PFCP (Packet Forwarding Control Protocol) is a 3GPP Protocol that is communicated on Sxa, Sxb, N4 interface between the Control Plane (CP) elements and User Plane (UP) elements.

The CP element programs the UP element with policies on how to forward packets. The CP function controls packet processing in the UP function by establishing, modifying, and deleting the PFCP session.

PFCP Messages

Session Related Messages

These messages are exchanged between Control Plane and User Plane function at the subscriber level. These messages are load balanced based on the **IMSI** of the subscriber.

- Session establishment request/response
- Session modification request/response
- Session deletion request/response
- Session report request/response

Node Related Messages

The PFCP node related message packets are broadcast to all the ports which are part of the port-group in FS/WL maps. These messages are not specific to any subscriber. The load balancing based on **IMSI** is not possible for these packets. They continue to broadcast as per the existing behavior.

- Heartbeat request/response
- PFCP PFD management request/response
- PFCP association setup request/response
- PFCP association update request/response
- PFCP association release request/response
- Version Not supported response
- PFCP Node report request/response
- PFCP session set deletion request/response

PFCP Load Balancing

The PFCP packets received are load balanced based on the **IMSI** present in the "IE Type: UserID".

Once the PFCP packet is received, the mobility application will parse the packet and create or update or delete the corresponding sessions or tunnels in the session table. The IMSI present in "IE Type: UserId" is stored in the session table and used for load balancing. Once sessions or tunnels are created, map lookup is performed for the PFCP packets. If there is a match in the IMSI, the packet is sent to only one port from the configured **Load Balancing** application based on the **IMSI** hashing. The hash value string from IMSI string present in the

"IE Type: UserId" is added to the packet before sending to the **Load Balancing** application. The **Load Balancing** application will do **Load Balancing** logic based on this hash string and choose any one port to send the packet out.

Licensing Requirement:There is no need of a separate license for PFCP Load Balancing. The existing license for Standalone UPN is sufficient for this feature.

Single Engine: The PFCP packet is sent to only one GigaSMARTS engine. Based on the map lookup, the packet will be forwarded to the **Load Balancing** application and sent to only one port based on the IMSI hashing.

Grouped Engine: The PFCP packet is sent to all the GigaSMART engines configured as part of the engine group configuration. The PFCP packets are processed as per the existing design, but the packets are sent out only through one port in the **Load Balancing** application in the leader engine**.**

Map Configuration

The 4G and 5G traffic passes through a single map in standalone mode. The user can configure a 4G map through which both 4G and 5G traffic passes. You do not need to configure separate maps for 4G and 5G traffic.

> **NOTE:** The single map configuration is applicable only for standalone UPN. For legacy GTP and 5G user need to configure 4G and 5G maps separately

Required License for UPN Standalone Single Map type: 4G and 5G.

Stateful Session Recovery

It is crucial to back up the information from the correlated session table to avoid discrepancies during engine reboots and connection failures. After a successful reboot, the information saved as part of the persistence file is populated back into the mobility database.

Mobility applications require stateful session recovery to persist their data between node reboots and connection failures. This feature aims to extend the support available in LTE/5G mobility applications to Standalone UPN functionality. The session table is backed up every 10 minutes. You cannot change the backup time interval.

File Age Timeout— Specifies the time the backup file is valid, in minutes. After this time expires, the backup file is stale. The default is 30 minutes.

<10 ~ 1440> - 10 minutes is the minimum and 1440 minutes is the maximum value for file-age-timeout.

Restart Age Timeout—Specifies the time interval following a reboot for aging-out sessions, in minutes. This is a shorter interval than that specified using the gtp-flow timeout. The gtp-flow timeout disconnects a GTP session if it has been inactive for the timeout value, which has a default of 8 hours. The restart-age-timeout default is 30 minutes.

<10 ~ 1440> - 10 minutes is the minimum and 1440 minutes is the maximum value for restart-age-timeout

GTP Persistence Interval—Specifies the time interval between backups, in minutes. The default is 10 minutes.

<10 ~ 1440> - 10 minutes is the minimum and 1440 minutes is the maximum value for restart-age-timeout.



5G Stateful Session Recovery

Required License: 5G Correlation

5G stateful session recovery provides session persistence for GigaSMART 5G applications, including 5G flow filtering, 5G forward listing, and 5G flow sampling. In this method of recovery, 5G sessions are backed up periodically so that they can be recovered faster after a GigaSMART line card reboot or a node reboot.

5G stateful session recovery requires additional memory for storing backups. GigaVUE-HC3 has the required memory.

Using 5G stateful session recovery, the 5G session tables in the GigaSMART line card memory are periodically backed up to the control card memory on the node and stored.

You should configure an interval for how often the backups occur, such as every 10 minutes. If 5G stateful session recovery is enabled and the GigaSMART line card is rebooted, the 5G session tables are restored automatically following the reboot.

The last stored backup file is downloaded from the control card to the GigaSMART line card using FTP. The session table is repopulated from the last stored backup file to each GigaSMART engine, up to 8 engines. Packet count statistics for sessions are saved and restored.

Depending on the size of the session table, the amount of time to restore from the backup might take as much as 3 minutes. During that interval, traffic is blocked to the virtual port on the GigaSMART line card. Once the session table is read and populated, traffic is allowed.

Depending on the interval between backups, there could be differences between the stored state and the current state of the system, for example, map configuration could change, or sessions could be added, modified, or deleted.

Load Balancing information is not persisted, so after a session table is repopulated, a session that was once sent to one Load Balanced port may be sent to a different Load Balanced port after the reboot. However, for SUPI-based Load Balancing, the traffic is sent to the same port as it was before the reboot.

5G stateful session recovery works in a cluster environment; however, the cluster leader must remain the same.

Configure 5G Stateful Session Recovery

To enable 5G stateful session recovery, as well as to configure timers, do the following:

1. On the left navigation pane, click , and then select **Physical>Nodes**.

2. From the device view, select **GigaSMART > GigaSMART groups**.

3. Click on the alias of the **GigaSMART group**.

4. Select **GTP Persistence** in the **GTP Persistence** fields under GigaSMART Parameters. The timers are pre-configured with default values.



Figure 19 *GTP Persistence GigaSMART Parameters*

Use the **System** widget on the Overview page to determine the amount of memory. The size of memory is 24Gb in an upgraded system.

View Backup and Restore Information

To view the System information, select **Overview** from the Navigation pane. The amount of free and used memory is displayed in the **Memory** field.

To view backup and restore information for GTP Persistence:

1. Select **GigaSMART > GigaSMART Groups > GigaSMART Group**.

2. Click on the alias of the GigaSMART group.

A Quick View appears for the selected GigaSMART group.

3. Scroll down to GTP Persistence. In Figure 26-63, GigaSMART Group gsgrp-1_4_e1 is selected and the Quick View is displayed.

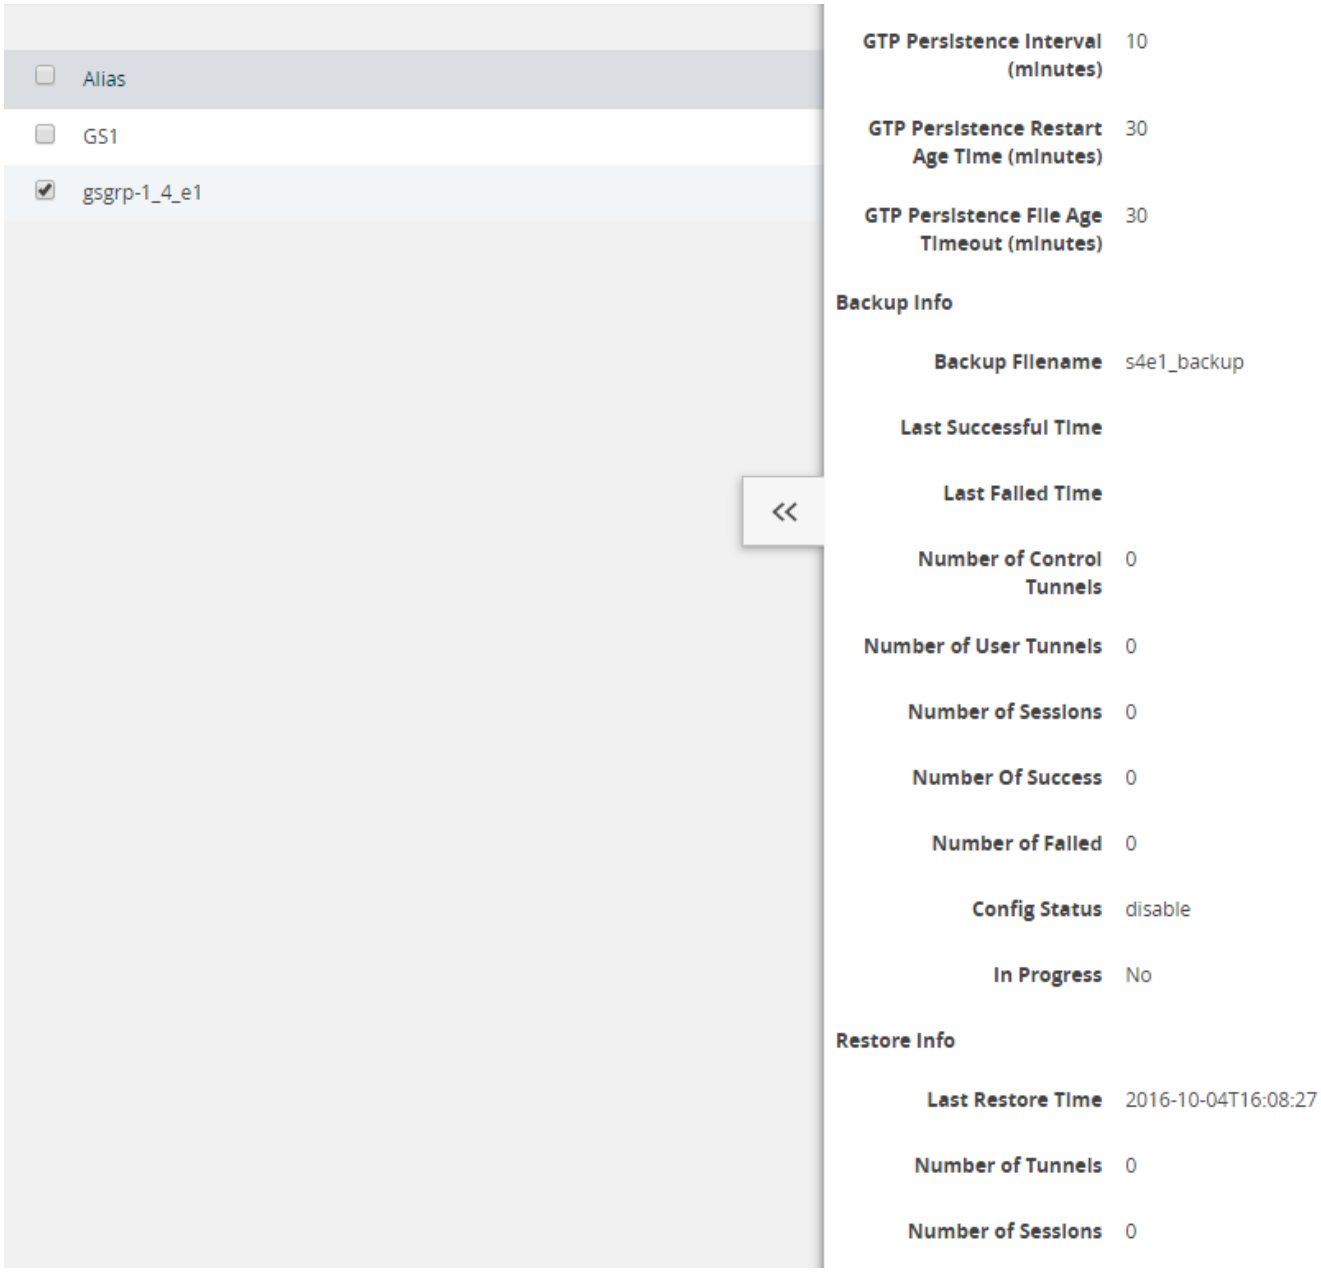


Figure 2 GTP Persistence Information

Configure GTP Overlap Mapping

The configuration of GTP forward listing and GTP flow sampling maps that are part of the GTP overlap flow sampling map group follow the same configuration considerations discussed previously in GigaSMART GTP forward listing and GTP Flow Sampling. As is the case with regular non-overlap GTP mapping, GTP forward listing selects specific subscribers based on IMSI, whereas GTP flow sampling uses map rules to select subscribers and then forward a percentage of the packets to tool ports.

Configuration Considerations

This section details certain configuration considerations that apply only to the configuration of GTP forward listing and flow sampling maps for GTP overlap flow sampling maps.

About Standalone UPN Flow Sampling Map Mode and Port Groups

A second level type map specifying Standalone UPN flow sampling map mode must be selected to configure Standalone UPN forward listing and flow sampling maps.

To configure a Standalone UPN whilelisting map in overlap flow sampling map mode, select **Type** as

Second Level and **Subtype** as **Flow Whitelist Overlap** in a map.

To configure a GTP flow sampling map in GTP overlap flow sampling map mode, select **Type**

as **Second Level** and **Subtype** as **Flow Sample Overlap** in a map .

You can configure one GTP forward listing map and one GTP flow sampling map pair that contain traffic policies corresponding to one destination port group. The load balanced port groups can contain a single port, a port range, or a GigaStream. Note that, starting in software version 4.8, port groups used in GTP overlapping maps support GigaStream.

The maximum number of port groups per single GTP overlap flow sampling map group is six.

For more information about port groups, refer to Port Groups.

Maximum Number of Port Group Members

The Maximum number of port group members can be determined as follows:

Standalone Configuration

In a standalone configuration, use the following sequence :

1. Determine the number of members per port group and add 1.

2. Multiply each port group result by each other.

3. The total multiplication should not exceed 4096.

For instance, assume the following configuration in a GTP overlap mapping group:

o Port Group 1—2 load-balanced GigaStream

o Port Group 2—3 load-balanced GigaStream

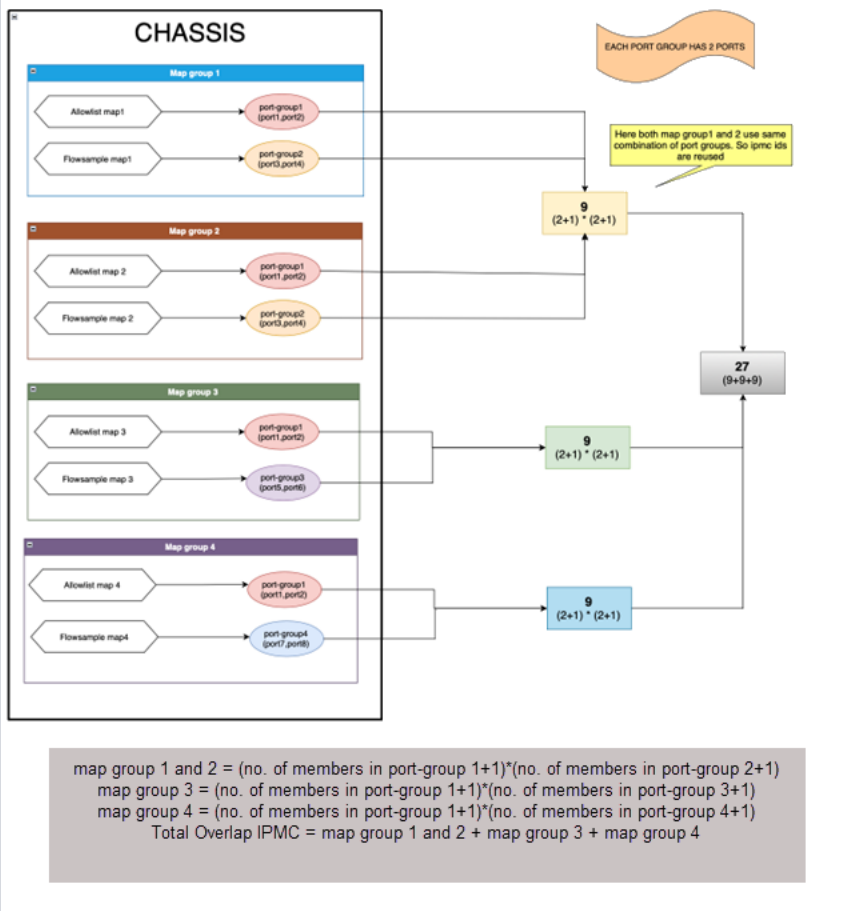
o Port Group 3—1 load-balanced tool port

o Port Group 4—1 load-balanced GigaStream

o Port Group 5—4 load-balanced tool ports. The total number becomes:(2+1)*(3+1)*(1+1)* (1+1)* (4+1) = 240

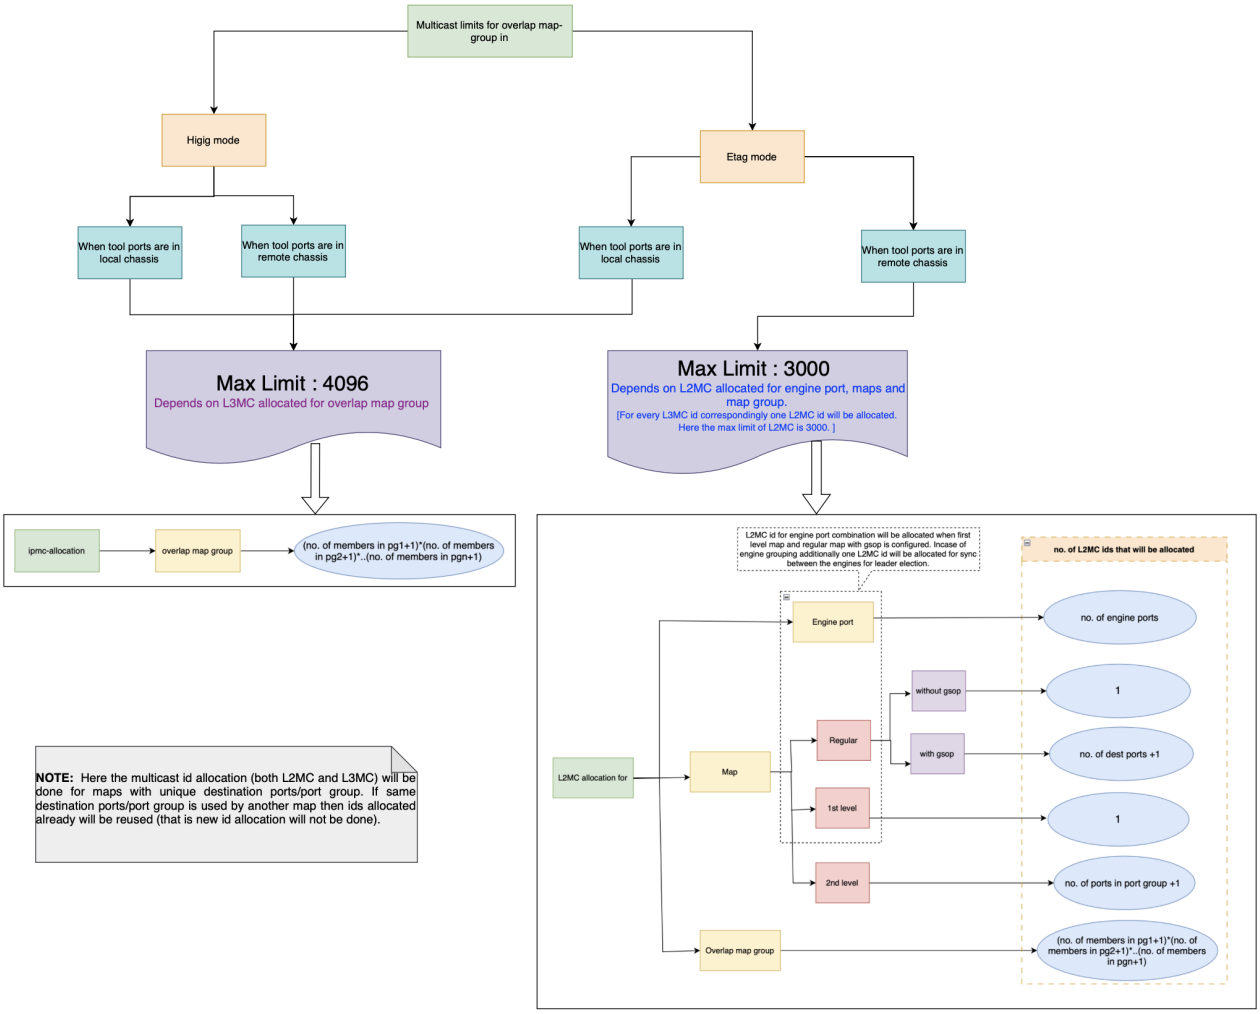
The tool configuration shown is accepted since this does not exceed the maximum number of multicast IDs (4096).

If Port Groups are utilized across multiple map groups in a single cluster, refer to the representation below on how the multicast IDs are calculated.



Cluster Configuration

In a cluster configuration, Multicast IDs can defer based on Hi-Gig mode or E-tag mode. The below is representation of the same.



Standalone UPN Flow Sampling Map Priority

Since a packet matches multiple maps independently the concept of second level map priority does not apply to Standalone UPN flow sampling maps. A standalone UPN flow sampling map pair consists of one Standalone UPN forward listing map and one Standalone UPN flow sampling map having the same destination port group. Within a Standalone UPN flow sampling map pair the forward listing map rules will be applied before the flow sampling map rules.

Virtual Port Configuration in Standalone UPN Mode

In Standalone UPN map configuration, the virtual port sending traffic to all the port groups need to be configured in Standalone UPN mode.

To configure the virtual port with Standalone UPN mode, select Standalone UPN when configuring the virtual port.

Overlap Support

Currently, a GTP packet can be destined to at most one destination tool. You can now receive multiple copies of a GTP packet that are sent to various tools (One copy per tool set).

GigaSMART Rotational Sampling Support

You can now configure GigaSMART Rotational Sampling in a standalone UPN, with *3GPP node role* for the GigaSMART group as **user** and **mode** as *standalone*. For more details refer to GigaSMART Rotational Sampling.

Interface Filtering and APN/DNN Filtering

Interface Filtering

In interface-based filtering, the traffic is filtered based on the interface from where it originates. When a new session is created, the interface information will be extracted and checked against the flow sample or whitelist rules. If there is a match with the flow sample or whitelist rules, the traffic belonging to this session will be forwarded to the corresponding tool ports.

Supported Platforms:

- GigaVUE-HC3 Gen 2
- GigaVUE-HC3 Gen 3

Note: You can configure interface-based filtering only through CLI and not through GigaVUE-FM.

You can configure only the interfaces **S1U, S5S8U, N3,** and **N9** for interface based filtering. It is not recommended to configure other interfaces such as **Gn, S5, S1** as they are not applicable for standalone UPN interface-based filtering.

The PFCP packets hit the configured interface rule irrespective of the interfaces from where they are originating.

For interface-based filtering (**S1U, S5S8U, N3,** and **N9**), you must configure the node role as **Standalone UPN** only. The CLI shows the following error, when you configure other node roles.

"% S1U, S5S8U, N3 and N9 interfaces can be configured only if the 3gpp-node-role is stand-alone"

The interface-based filtering supports the filtering of only GTP-U packets.

The following **Standalone UPN flow ops** report displays the **interface** information in the session table:

Joy (config) # sh gsgroup flow-ops-report alias gsg-g3 type flow-filtering gtp-imsi-pattern IMSIVALUE000000

==========
==========

Tunnel[Ver] Interface IP:Tunnel-ID => IP:Tunnel-ID IMSI WL FS ID LB port Pkts Timestamp

IMEI MSISDN

EBI:LBI[QCI] APN

==========
==========

CTRL[1] **S1U** 10.116.22.6:0xb289ad10 => 10.116.22.76:0x135e0620 IMSIVALUE000000 _ _ 3 45189089276

IMEIVALUE0000000 MSISDNVALUE00000

USER 5 10.116.22.44:0x00338cec => 10.116.22.79:0x535e0625 wap.mnc000.mcc000.+ N A 1 _

CTRL[2] **S5/S8-U** 10.254.156.136:0x0fb67d86 => 10.254.165.199:0x5dd70620 IMSIVALUE000000 _ _ 17 57994796676

IMEIVALUE0000000 MSISDNVALUE00000

APN (Access Point Name)/DNN (Data Network Name) Filtering

The Access Point Name (APN) is the name of a gateway between a 4G or 5G mobile network and another computer network, mostly the public Internet. A mobile device making a data connection must be configured with an APN.

In **APN/DNN** filtering, the traffic is filtered based on the **APN** string matching.

Note: You can configure APN/DNN filtering only through CLI.

When a new session is created, the **APN** pattern will be extracted and checked against the flow sample or whitelist rules, if there is a match with the **APN** pattern, the traffic belonging to this session will be forwarded to corresponding tool ports. In case of 5G traffic, the **DNN** information will be processed under the **APN** identifier.

The pattern match can be supported as an independent filtering or can be combined with the other filtering parameters such as the IMSI/IMEI/MSISDN.

The **APN/DNN** filtering supports the filtering of both the PFCP and GTP-U packets.

The following **Standalone UPN flow ops** report displays the **APN/DNN** information in the session table:

Joy (config) # sh gsgroup flow-ops-report alias gsg-g3 type flow-filtering gtp-imsi-pattern IMSIVALUE000000

==========
==========

Tunnel[Ver] Interface IP:Tunnel-ID => IP:Tunnel-ID IMSI WL FS ID LB port Pkts Timestamp

IMEI MSISDN

EBI:LBI[QCI] **APN**

==========
==========

IMSIVALUE000000 _ 0 0

IMEIVALUE0000000 MSISDNVALUE0000

CTRL 10.241.15.8 : 0x21480840 => 10.241.15.36 : 0x158a2740

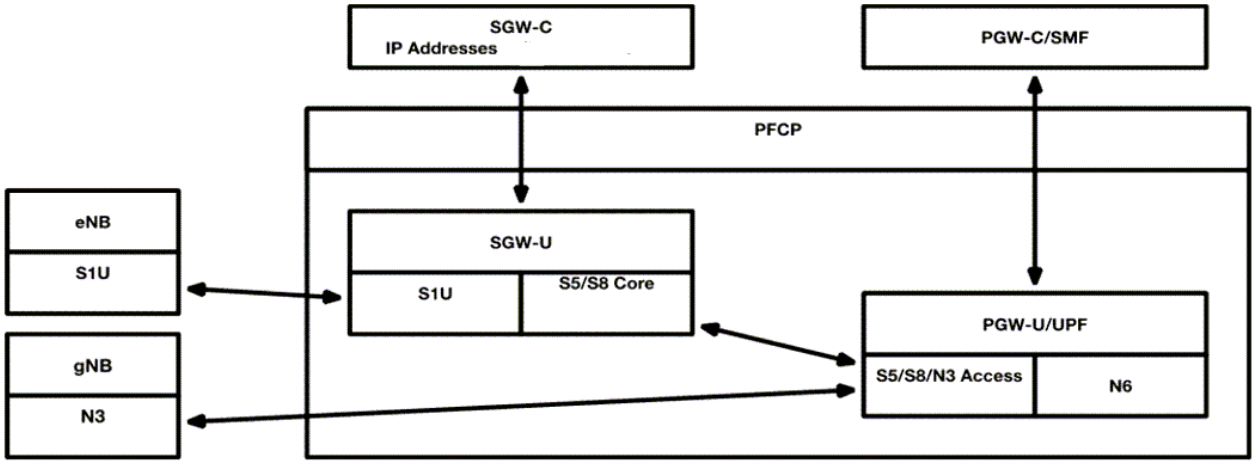
USER(S5S8U) 10.226.162.249:0x04f82ea1 => 0.0.0.0:0x00000000 ims.mnc000.mcc000.gp+ N A 2

USER(S5S8U) 0.0.0.0:0x00000000 => 10.241.15.41:0x558a2745 ims.mnc000.mcc000.gp+ N A 2

==========

Custom Interface Selection

When the interface IE is not available in the PFCP packets, you can explicitly provide the IP addresses of the network nodes such as SGW-C and SGW-U for identifying and filtering the traffic based on the interface (S1U, S5S8-U and N3). To configure customer interface selection, refer the Configure Custom Interface Selection topic.



Supported Platforms:

- UPN in GigaVUE-HC3 Gen 2 GigaSMART Card
- UPN in GigaVUE-HC3 Gen 3 GigaSMART Card

NOTE: This feature can be configured only through the GigaVUE-OS CLI and not through the GigaVUE-FM.

A new gsparam is introduced in the GigaVUE-OS CLI, which enables you to choose between the custom interface filtering option and the default 3GPP interface type IE based filtering and populate an IP address profile. As part of the profile configuration, you must enter the range of IP addresses associated with SGW-C and access IP addresses associated with S5S8-U in the GigaVUE-OS CLI.

Configuration of custom interface selection through CLI

- This feature is supported on gsgroup 3gpp-node-role user mode stand-alone.
- The custom mode must be enabled in the upn-interface-select gsparams and an IP profile must be attached.
- The newly added gsparam will take 3GPP Interface filtering as the default option.
- When upgrading to software release version GigaVUE 6.5.00, the N4/N3 user traffic will not be correlated and will be discarded since the QFI value will not be preserved through persistence records.
- Multiple IP profiles can be created, but only one IP profile can be associated per gsgroup.
- Interface qualifiers that are configured in second-level map rules are not applied to PFCP packets. Instead, PFCP packets will be matched against other configured rule qualifiers. This allows the tool to learn about all PFCP sessions. This approach facilitates session handovers between N3 and S5/S8 by processing PFCP independently of the interface qualifiers.
  For example, with the flow sample rule "add imsi 1234* interface N3 percentage 100," packets are handled as follows:
    - PFCP Packet 1 (IMSI 123411111111111 interface N3) matches the rule and is sent to the tool.
    - PFCP Packet 2 (IMSI 123422222222222 interface S5/S8u) also matches and is sent to the tool.
    - PFCP Packet 3 (IMSI 123433333333333 interface S1U) matches and is sent to the tool.
    - PFCP Packet 4 (IMSI 111144444444444 interface N3) does not match and is discarded.

  Similarly, for the rule "add interface N3 percentage 100," the packets are handled as follows
    - PFCP Packet 1 (interface S1U) matches the rule and are sent to the tool.
    - PFCP Packet 2 (interface N3) matches the rule and are sent to the tool.
    - PFCP Packet 3 (interface S5/S8u) matches the rule and are sent to the tool.

  This demonstrates that the interface criteria is not considered while forwarding the packets.

Limitations:

- The IP profile configuration and upn-interface-select gsparams is not supported through the GigaVUE-FM.
- The Flow-ops report will not display the interface specific counters for Standalone UPN (S1U, S5S8-U,N3).
- The UPN interface IP list does not support IPv6 addresses. Only IPv4 addresses are supported.
- Once the IP profile is created in the GigaVUE-OS CLI, you cannot modify the IP address or the IP address range specified in the IP profile. You can only create a new profile with a different IP address or IP address range and then associate the IP profile to the gsgroup.
-

GigaSMART 5G CUPS

This chapter describes about the GigaSMART 5G CUPS solutions and its operations.

Refer to the following sections for details:

- Overview
- 5G Correlation
- User Plane Node Traffic Monitoring
- Tool Traffic in Multi-Site Scenarios
- Configure CPN-UPN Communication Solution using Ansible

Overview

Required License : 5G-Correlation and CUPS

Supported Devices: GigaVUE-HC3 Gen 2, and GigaVUE-HC3 Gen 3.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

The GigaSMART 5G CUPS traffic visibility fabric solution provides network visibility across the control and user planes for the 5G stand-alone packet core network. A 5G stand-alone packet core network comprises the following network functions:

- Authentication Server Function (AUSF)
- Access and Mobility Management Function (AMF)
- Data Network (DN), e.g. operator services, Internet access or third party services

- Structured Data Storage network function (SDSF)
- Unstructured Data Storage network function (UDSF)
- Network Exposure Function (NEF)
- NF Repository Function (NRF)
- Policy Control function (PCF)
- Session Management Function (SMF)
- Unified Data Management (UDM)
- User plane Function (UPF)
- Application Function (AF)
- User Equipment (UE)
- (Radio) Access Network ((R)AN)

The following diagram shows the 5G Stand Alone Packet Core Network, in which only the N11, N3, and N4 interfaces are supported:

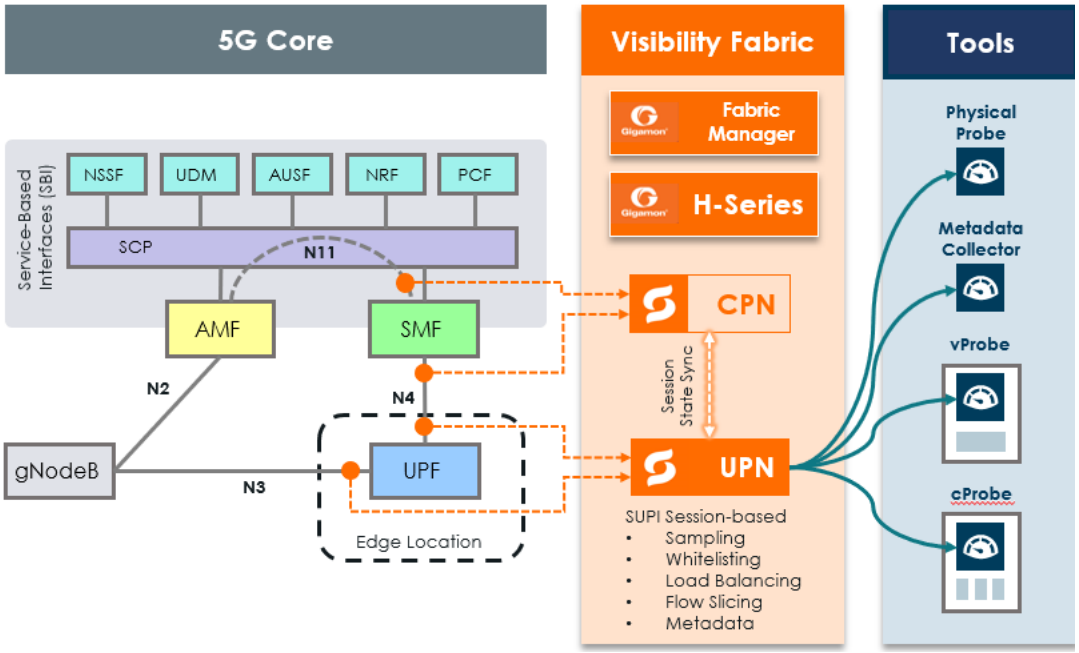


Figure 20 *5G Stand Alone Packet Core Network*

The GigaSMART 5G CUPS traffic visibility fabric solution has the following three functional areas:

Access	Accesses the traffic through physically or virtually tapping specific interfaces connecting various network element functions.
Correlation	Subscriber-aware session traffic correlation to enable sending both the control plane portion and user plane portion of a given session to the same tool.
Forwarding	Traffic forwarding through a combination of subscriber-aware white-listing, flow sampling, flow filtering, and load-balancing.

The following diagram explains the LTE and 5G CUPS traffic visibility solution:

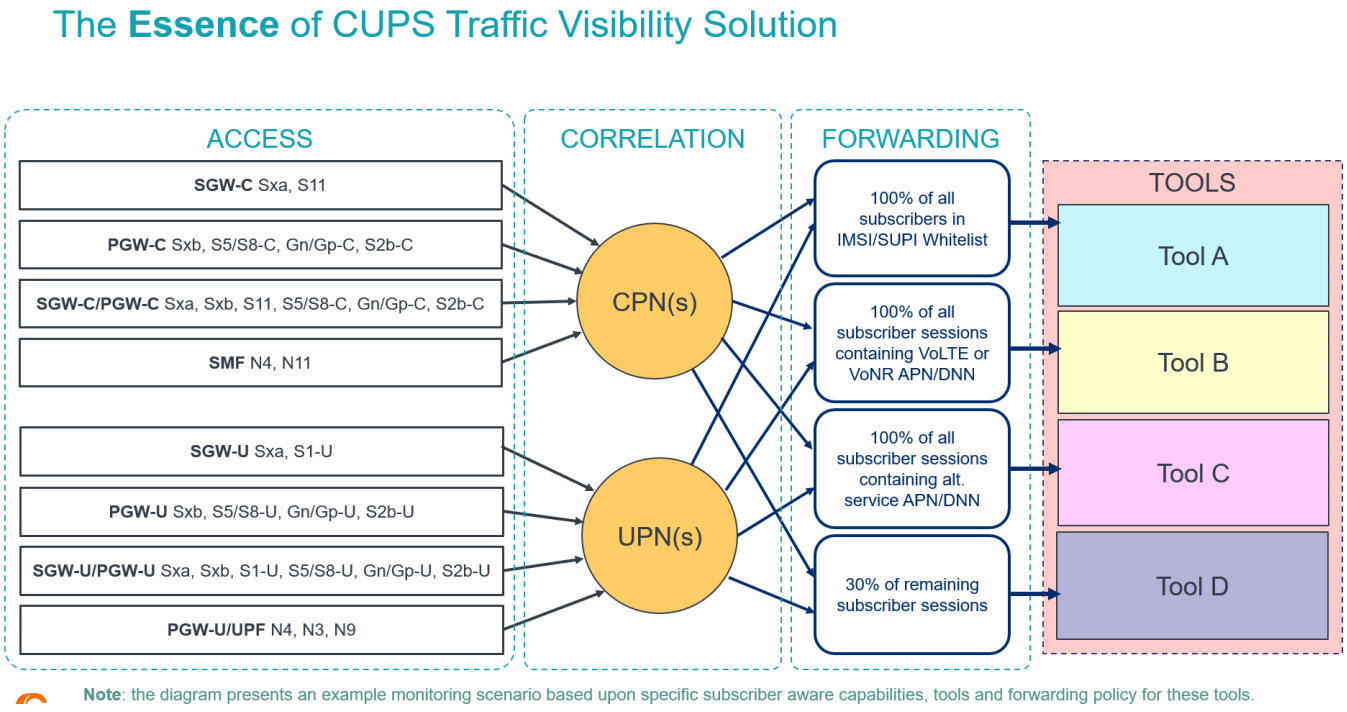


Figure 21 *CUPS Traffic Visibility Solution*

In the *Figure 2 : CUPS Traffic Visibility Solution*, specific interfaces of the user plane control functions such as SMF, SGW-C, PGW-C and the user plane functions such as UBF, SGW-U, PGW-U are accessed for getting a copy of the traffic that crosses each interface, and therefore collection of traffic from those access points becomes the basis for traffic monitoring and analysis.

Once the traffic is accessed, it is correlated in order to combine the control plane and user plane traffic for a given subscriber-related session. Then the traffic is forwarded through a

combination of white-listing, flow sampling, flow-filtering, and/or load-balancing to traffic monitoring/analysis tools in a load balanced way so that the tools can be used in the most efficient way.

The CUPS solution managed by GigaVUE-FM supports the following mobility networks scenarios based on visibility and analytics fabrics:

- 5G stand-alone mobility networks
- 3G/4G LTE mobility networks
- Mix 3G/4G LTE and 5G mobility networks

Supported Platforms

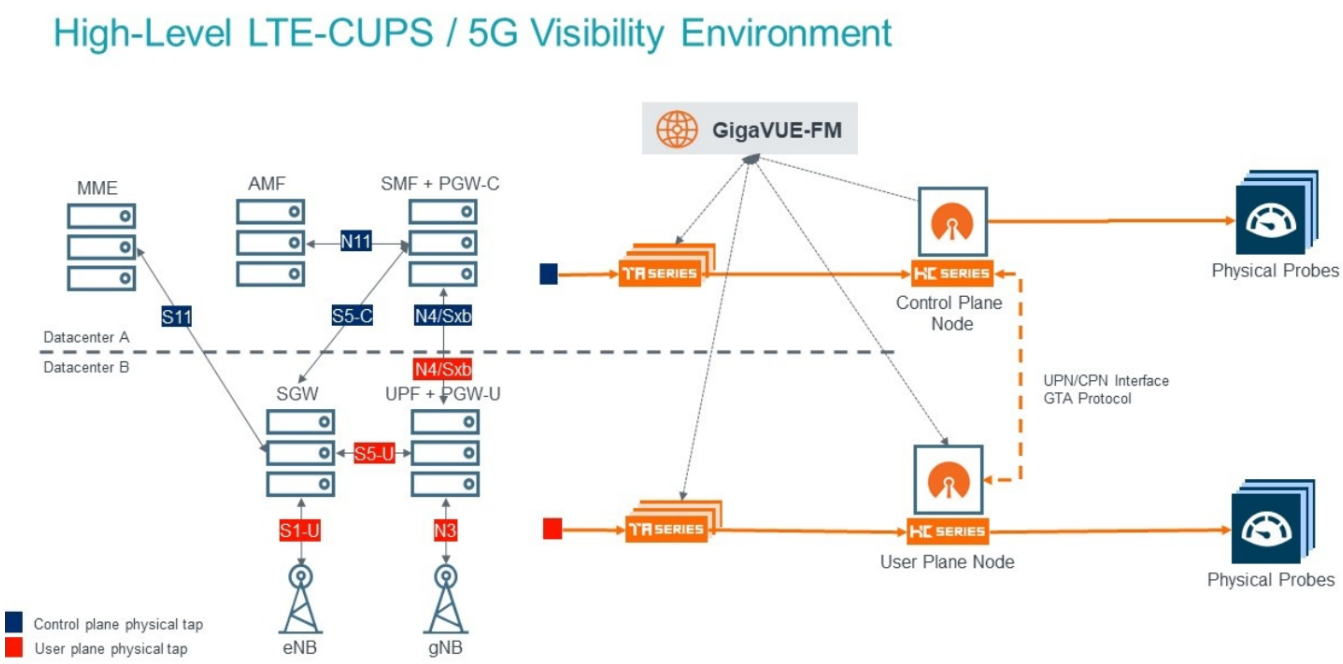
GigaSMART 5G CUPS traffic visibility fabric solution is supported in the following platform:

- GigaVUE-HC3

5G Correlation

The 5G Correlation feature correlates the 5G Control and User packets to deliver it to different tool ports based on the filtering policies configured. The control and user packets are processed in Control Plane Network (CPN) and User Plane Network (UPN) located in same or different locations, and then the packets are sent to the tools. The following diagram shows the High-Level LTE / 5G CUPS Visibility Environment.

NOTE: It is recommended that you unconfigure and reload the GigaSMART card when configuring a particular GS group with 5G and switching back to 4G or vice versa.



5G Load Balancing

If the node role is CPN, the load balancing of the packets depend upon the parameters configured for TCP in GigaVUE-FM.

The UPN node load balances the correlated packets across all the tool ports based on the configured hash key value such as SUPI, PEI, GSPI.

Configure 5G Load Balancing

To configure 5G Load balancing, follow these steps:

1. From the left navigation pane, go to **System > GigaSMART >GigaSMART Operations (GSOP)**, and then click **New**.
2. Specify an Alias in the **Alias** field.
3. Click in the **GigaSMART Group** field and select a GigaSMART group.
4. Click in the GigaSMART Operations (GSOP) field and select Load Balance in **TCP Application Parameters**.
5. Specify the paramaters. The following table explains the parameters for configuring the load balancing in CPN:

TCP Application Parameters	Options
Application	• Broadcast - All the TCP handshake packets are broadcasted to the tool ports in all the maps. • Drop - All the unknown application packets of 5G are sent to the collector. • Enhanced -
Load Balance	• Enable - All the 5G control packets will be load balanced across the port-groups configured in the flowsampling and whitelisting maps. • Disable - All the 5G control packets will be broadcasted.
TCP Control	• Broadcast - All the tcp handshake packets will be broadcasted to the tool ports in all the maps.. • Drop - All the TCP handshake packets will be sent to collector if configured

5G RAN Correlation

A Radio Access Network (RAN) is part of a telecommunications system that connects individual radio devices to the core mobile network through radio connections.

RAN uses GTP and / or 5G correlation. It exposes RAN data fields from control plane (S11), for example:

- New Radio Cell Global Identifier (NCGI)
- Tracking Area Identity (TAI)
- Tracking Area Code (TAC)

The Tracking Area Identity (TAI) and New Radio Cell Global Identifier (NCGI) are extracted from the User location Information (ULI) and sends them to flow sampling, filtering and forward listing.

The RAN correlation for 5G CUPS flow sampling uses the NCGI, TAC and PLMN-ID values.

For 5G RAN based forward listing requires NCGI includes the PLMN Id and the New Radio Cell Global Identifier (NCGI).

RAN data fields are available for correlation, flow sampling and forward listing. It does not support classic flow filtering.

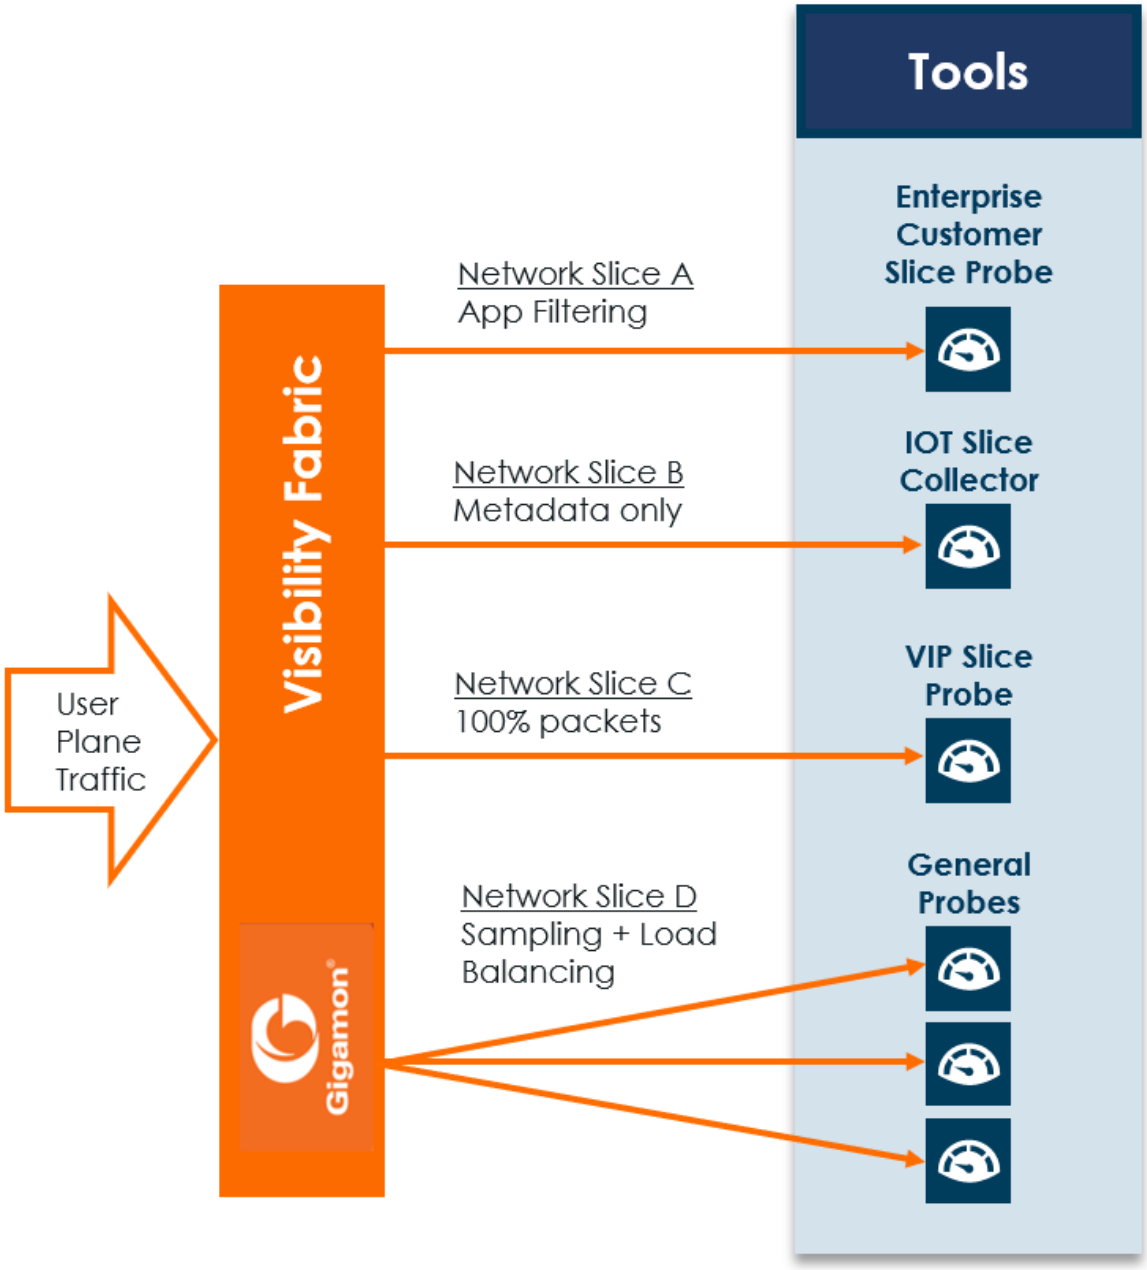
Configure 5G RAN Correlation

To configure 5G RAN Correlation, refer to Configure CPN-UPN Communication Solution using Ansible

Refer Deploying RAN Correlation for Forwarding Subscriber's Traffic Based on Geolocation for more detailed information.

5G Network Slice Correlation

Network Slicing is a feature of 5G networks whereby slices of network resources are dedicated to a specific application. Gigamon enables to have differentiated levels of visibility for each network slice.



5G Network Slice Correlation uses 5G Correlation. It exposes Network Slice instance ID (NSI). NSI is available for filtering through flow sampling maps. 5G subscriber aware filter includes the Network Slice Selection Assistance Information (NSSAI) value for 5G correlation.

Configure 5G Network Slice Correlation

To configure 5G Network Slice Correlation, refer to Configure CPN-UPN Communication Solution using Ansible

5QI Correlation

In 5G traffic, 5QI values present in the control traffic and its associated QoS Flow Identifiers (QFI) value are correlated and the data traffic is sent to the tools. QoS Flow Identifiers (QFI) identifies each QoS (Quality of Service) flows.

5QI flow sampling are performed based on the QoS flows instead of using other subscriber filtering parameters such as SUPI, PEI, GPSI, DNN.

You can configure 5QI along with other filtering parameters such as SUPI, PEI, GPSI, DNN, NSSAI.

Rules and Notes

Keep in mind the following rules and notes when you work with the 5QI sampling:

- 5QI does not support whitelisting.

5G Flow Sampling and Filtering

Required License : FlowVUE license is required for Flow Sampling

5G flow sampling samples a configured percentage of 5G sessions. 5G flow sampling uses map rules to select subscribers and then forward a percentage of the packets to tool ports.

Pass rules are defined in flow sampling maps. Each rule contains some combination of SUPI, PEI, and GPSI numbers or patterns as a percentage to sample. The flow is sampled to see if it matches a rule. The percentage of the subscriber sessions matching each rule are selected.

Map rules specify the type of traffic to be flow sampled by that map. For each new session, map rules are evaluated in top-down order of decreasing priority. If there is a match, the indicated percentage of the subscriber session is either accepted or rejected. If accepted,

the traffic is sent to the tool port or Load Balancing group specified in the map. If rejected, the traffic is dropped. If there is not a match to a rule, the traffic is passed to subsequent maps.

About Flow Sampling Rules and Maps

Flow sampling rules are configured in maps called flow sampling maps. Up to ten (10) flow sampling maps per GigaSMART group are supported. Each flow sampling map supports up to 100 flow sampling rules, for a maximum of 1000 (10*100) rules per GigaSMART group.

5G flow sampling (rule-based flow sampling) is performed after 5G whitelist-based forwarding. So, flow sampling maps have a priority lower than forward list maps and higher than flow filtering maps.

NOTE: For 5G second level maps, a maximum of fifteen maps can be attached to a vport. For example, for the same vport you can have one forward list map and ten flow sampling maps, or ten forward list map, four flow sampling maps, and one flow filtering map. In addition, you can have a collector map, which is not counted.

In the flow sampling maps, the rules in the first map have a higher priority than the rules in the second, third, and subsequent maps. Within any single map, rules are evaluated in order.

Rules can be added to, deleted from, or inserted into a flow sampling map when the subtype selected for a Second Level map is Flow Sample. Suffix wildcarding, such as SUPI 100*, is supported in the flow sampling map rules.

Use the Add a Rule button in the Maps page to add a new flow sampling rule (a pass rule). Specify SUPI, PEI, or GPSI subscriber IDs, as well as the percentage of the flow to be sampled. The percentage is a range from 1 to 100%. Use 0% to drop sampled data.

A DNN pattern can be specified in a rule by itself.

For DNN, specify a pattern (a name) to match, for example, three.co.uk. Wildcard prefixes and suffixes are supported, for example, *mobile.com or *ims*. The pattern can be specified in up to 100 caseinsensitive alphanumeric characters and can include the following special characters: period (.), hyphen (-), and wildcard (*).

Use DNN to send traffic that matches a certain DNN pattern to specified tool ports, based on the sampling percentage.

When creating Flow Sampling rules on the Maps page, the first rule created has the highest priority and the priority of subsequent rules is in the order that they are added.

NOTE: For 5G CPN flow sample second-level maps of subtype **flowSample5g**, only percentage 100 is supported. Configuring any other rule is not supported. Also, configuring or updating the priority of **flowSample5g** pass rules by using API, GigaVUE-FM, or Ansible is not supported.

Add Rule to Flow Sampling Map

Flow sampling is applied for new subscribers. When a new rule is added to the rules in a flow sampling map, traffic will be sent to the port or load balancing group specified in the map.

Delete Rule from a Flow Sampling Map

When a rule is deleted from a flow sampling map, the session associated with the rule stays active. The traffic associated with the rule will not be reevaluated by subsequent maps.

Change Priority of Flow Sampling Maps

Priority is set as per the order defined in the policy YAML file within the type.

Delete Flow Sampling Map

When a flow sampling map is deleted, the priority of the remaining flow sampling maps will be reprioritized. For example, if the first flow sampling map is deleted, the second flow sampling map will increase in priority.

For the deleted flow sampling map, the traffic associated with the rules in the map will be reevaluated and then passed to subsequent maps. When a flow sampling map is re-prioritized, the existing sessions will be reevaluated according to the new priority of the map. The traffic associated with the rules in the map will be reevaluated and then passed to subsequent maps.

When the last flow sampling map is deleted, the traffic associated with the rules in the map will also be reevaluated before being passed to subsequent maps. But the traffic associated with the rules in maps that were not matched, will not be reevaluated because that traffic was already passed to subsequent maps.

Flow-Ops Report Limitation for Multiple Flow Sampling Maps

The flow-ops report displays the flow sampling rule ID for sessions that have been accepted or rejected by the flow sampling map.

However, since rule IDs are not unique across maps, when there are multiple flow sampling maps, the flow-ops report is unable to identify the exact rule that the session matched. For example, with multiple flow sampling maps, each map can have a rule ID of 1. The rule ID will be identified in the flow-ops report, but not the map associated with it.

5G Flow Sampling Percentage

The sampling Percentage field in a map for 5G flow sampling, represents the percentage of subscribers that will be sampled (not the sessions).

The 5G correlation engine tracks all the subscribers and all of their sessions that it sees on the network. In this example, for those subscribers with an SUPI starting with the value 46*, the 5G correlation engine keeps a list of them and randomly selects 80% of those subscribers and sets them to be in the sample, which means that a tool port (or load balanced group) will see 100% of the packets for 100% of the sessions for those randomly selected 80% of subscribers.

For the other 20% of subscribers, the 5G correlation engine continuously tracks those subscribers through the network but does not send any packets to the tool port (or load balanced group).

Drop Unmatched Traffic

When a session matches one of the configured flow sampling rules, it is either accepted for sampling or rejected.

If it is accepted, all packets belonging to that 5G session are sent to the tool port or ports specified in the flow sampling maps. If a subscriber is in the sample, then both the control plane packets, and the user-data plane packets are sent to the tools.

If it is rejected, all packets belonging to the session are dropped. If the subscriber is not in the sample, then neither the control plane packets, nor the user-data plane packets are sent to the tools.

Control plane and user-data plane traffic are treated the same. For a matching session, all the control plane and user-data plane traffic will be accepted. Otherwise, all the control plane and user-data plane traffic will be rejected and dropped.

GTP Overlap Flow Sampling for 4G and 5G

GTP Overlap Flow Sampling for 4G and 5G enables tools to receive their own copy of traffic, through independent evaluation of second level maps. In this overlap mode, if a packet matches a map rule, such as for tool A, that packet is still available to be matched again in subsequent maps, such as for tool B and tool C. The multiple copies of a GTP packet are sent to multiple destinations simultaneously. This applies to both GTP forward listing and GTP flow sampling even if sampling is not employed.

The destination could be a single tool port, a GigaStream or an IMSI based load balanced tool port/GigaStream for a particular tool set. The GigaStream sends traffic to various links, connected to an external load balancer that would be connected to multiple tools.

Configure GTP Overlap Mapping

The configuration of GTP whitelisting and GTP flow sampling maps that are part of the GTP overlap flow sampling map group follow the same configuration considerations discussed previously in GigaSMART GTP Whitelisting and GTP Flow Sampling. As is the case with regular non-overlap GTP mapping, GTP forward listing selects specific subscribers based on IMSI, whereas GTP flow sampling uses map rules to select subscribers and then forward a percentage of the packets to tool ports.

Configuration Considerations

This section details certain configuration considerations that apply only to the configuration of GTP forward listing and flow sampling maps for GTP overlap flow sampling maps.

About GTP Overlap Flow Sampling Map Mode and Port Groups

A second level type map specifying GTP overlap flow sampling map mode must be selected to configure GTP forward listing and flow sampling maps.

To configure a GTP whilelisting map in overlap flow sampling map mode, select **Type** as **Second Level** and **Subtype** as **Flow Whitlelist Overlap** in a map.

To configure a GTP flow sampling map in GTP overlap flow sampling map mode, select **Type** as **Second Level** and **Subtype** as **Flow Sample Overlap** in a map .

You can configure one GTP forward listing map and one GTP flow sampling map pair that contain traffic policies corresponding to one destination port group. The load balanced port groups can contain a single port, a port range, or a GigaStream. Note that port groups used in GTP overlapping maps support GigaStream.

The maximum number of port groups per single GTP overlap flow sampling map group is six.

For more information about port groups, refer to Port Groups.

Maximum Number of Port Group Members

The Maximum number of port group members can be determined as follows:

Standalone Configuration

In a standalone configuration, use the following sequence :

1. Determine the number of members per port group and add 1.

2. Multiply each port group result by each other.

3. The total multiplication should not exceed 4096.

For instance, assume the following configuration in a GTP overlap mapping group:

o Port Group 1—2 load-balanced GigaStream

o Port Group 2—3 load-balanced GigaStream

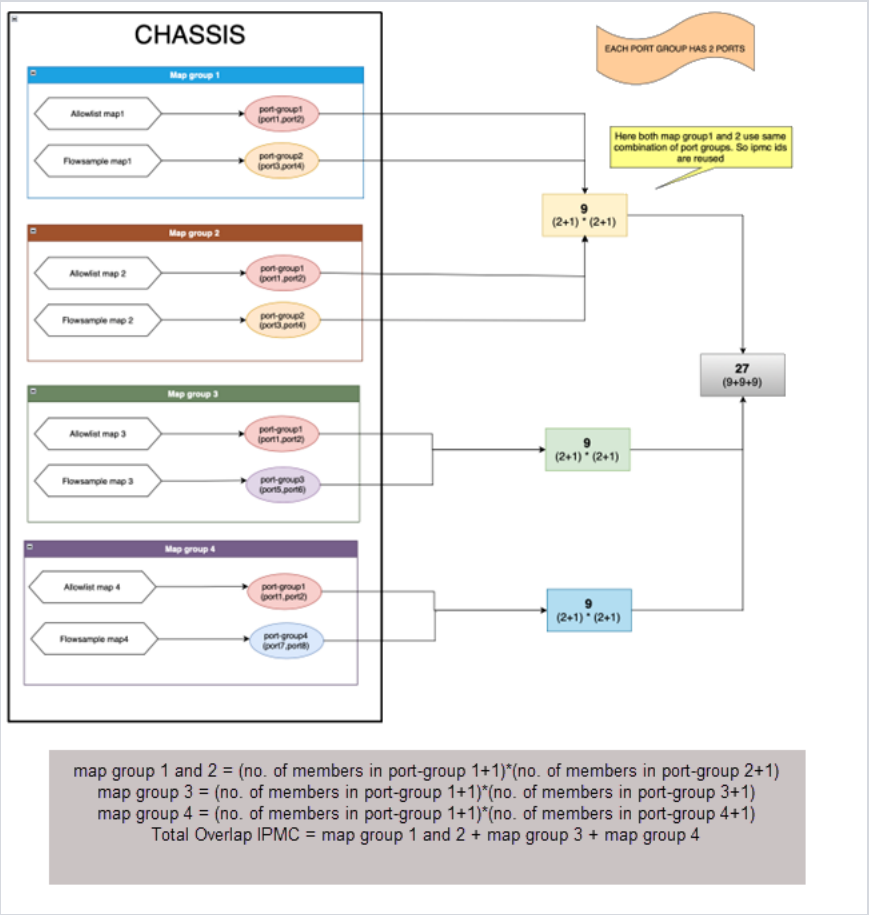
o Port Group 3—1 load-balanced tool port

o Port Group 4—1 load-balanced GigaStream

o Port Group 5—4 load-balanced tool ports. The total number becomes:(2+1)*(3+1)*(1+1)* (1+1)* (4+1) = 240

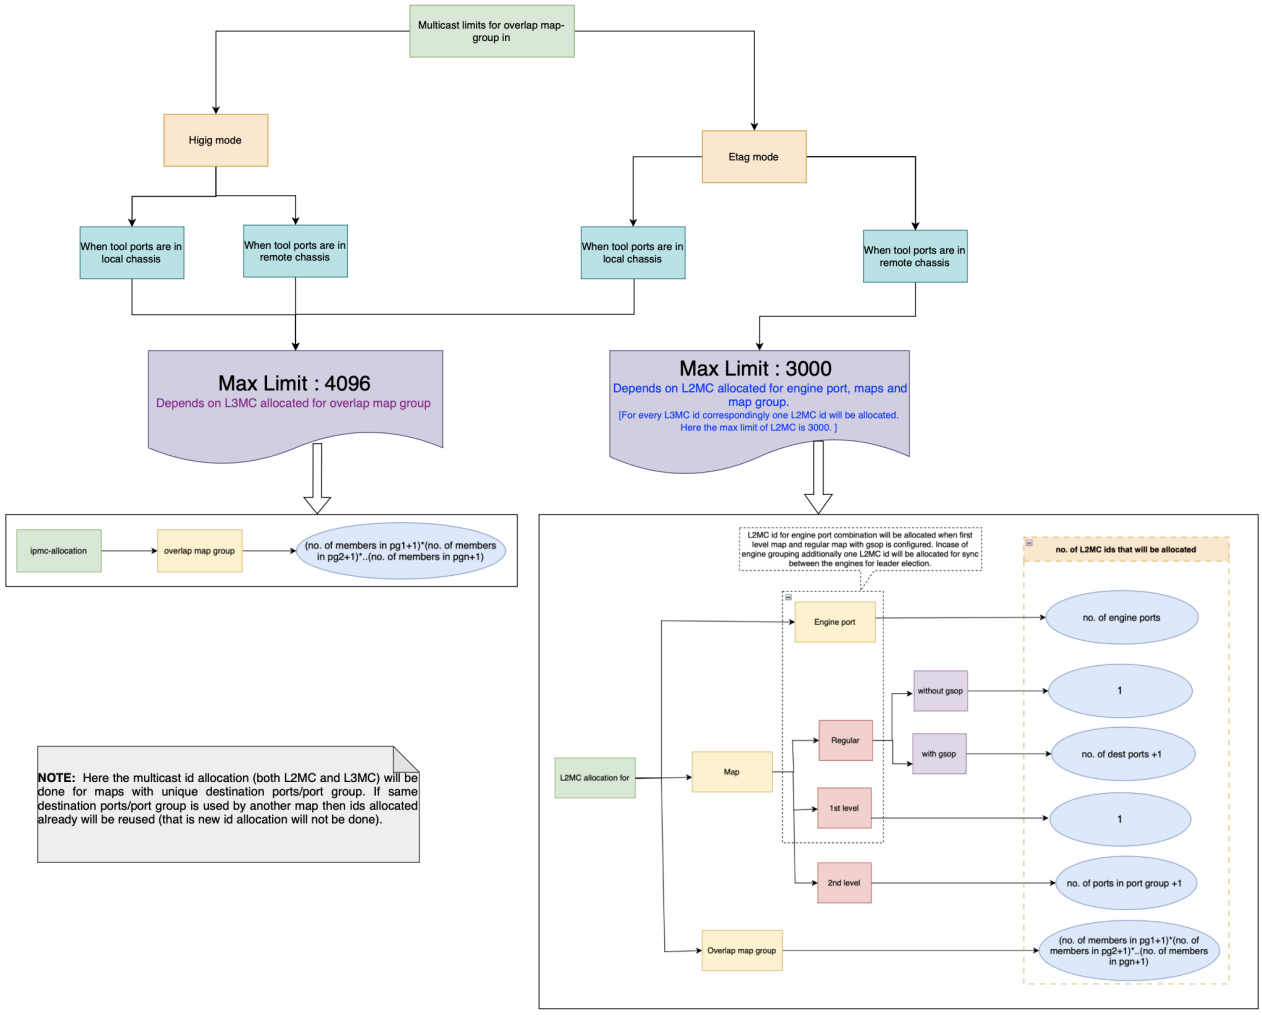
The tool configuration shown is accepted since this does not exceed the maximum number of multicast IDs (4096).

If Port Groups are utilized across multiple map groups in a single cluster, refer to the representation below on how the multicast IDs are calculated.



Cluster Configuration

In a cluster configuration, Multicast IDs can defer based on Hi-Gig mode or E-tag mode. The below is representation of the same.



GTP Overlap Flow Sampling Map Priority

A GTP overlap flow sampling map pair consists of one GTP forward listing map and one GTP flow sampling map having the same destination port group. Within a GTP overlap flow sampling map pair the forward listing map rules are applied before the flow sampling map rules.

Virtual Port Configuration in GTP Overlap Mode

In GTP Overlap map configuration, the virtual port sending traffic to all the port groups needs to be configured in GTP overlap mode.

To configure the virtual port with GTP overlap mode, select **GTP Overlap** when configuring the virtual port.

About Map Groups

To create a group of maps for GTP forward listing and GTP flow sampling, select **Maps > Maps > Map Groups**, and then click **New**. The maps for a map group are entered in the **Maps** field. All the maps in a map group receive traffic according to map rules, rather than map priority. Thus, multiple copies of a GTP packet can be sent to more than one tool.

The **Maps** field of the Map Group page groups the forward listing and flow sampling maps.

Keep in mind the following configuration considerations for map groups:

- A map group can be associated with only one GigaSMART group (gsgroup).
- All maps within a map group must be connected to the same vport.
- A map group can consist of only one GTP whitelisting map or only one GTP flow sampling map but it cannot contains two maps of the same type.
- Once a map group is created, it cannot be edited to change the type or subtype of the map. However, you can add and edit the map rules for a map while it is configured in a map group.

For more information about map groups, refer to Create Map Groups.

About Whitelist Maps

The GTP forward list is an IMSI list which is common to all whitelist maps. You can configure an optional rule within a whitelist map to specify a GTP version or interface-based policy.

Other than specifying a new second level type using **Type Second Level** and **Subtype Flow Whitelist Overlap** when creating the map, the configuration of GTP whitelist maps follows the same configuration guidelines as given in the section GTP Whitelisting.

A maximum of six whitelist maps sending traffic to six different port groups can be configured per GigaSMART group (gsgroup).

About Flow Sampling Maps

In GTP overlap flow sampling map mode, GTP flow sampling (rule-based flow sampling) is performed after GTP whitelist-based forwarding. Therefore, flow sampling maps have a lower priority than whitelist maps. Thus, within a GTP overlap map pair that consists of a single GTP forward list overlap map and a GTP flow sampling overlap map, the GTP whitelist map is of higher priority.

Within the flow sampling maps, the rules in the first map have a higher priority than the rules in the second, third, and subsequent maps. Within any single map, rules are evaluated in order.

A maximum of six flow sampling maps sending traffic to six different port groups can be configured per GigaSMART group (gsgroup).

Within each GTP forward listing and flow sampling pair, if there is not a match to an IMSI in the whitelist map, the traffic flow is sampled based on the rules in the flow sampling map. The flow sampling rules specify IMSI, IMEI, and MSISDN numbers, as well as the percentage to sample.

Within each map pair, packets are then accepted or rejected. Accepted packets are forwarded to the port groups for load balancing. Rejected packets are dropped.

Use the following steps to configure

Task	Description	UI Steps
1	Create GigaStreams that will be part of the port groups	a. Select **Ports** > **Port Groups** > **GigaStreams** b. Click **New**. c. Enter the name in the Alias field. d. In the Ports field, select port. e. Click Save. f. Configure a second GigaStream with the alias, select the required ports in the Ports field, and click **Save**.
2	Create port groups and specify the tool ports and assign GigaStreams to the port groups. The port groups will also be load balanced.	a. Select **Ports** > **Port Groups** > **All Port Groups**. b. Click **New**. c. Enter the name in the Alias field.

Task	Description	UI Steps
		d. Select Type GigaSMART Load Balancing. e. In the Ports field, select ports. f. In the GigaStream field, select the required GigaStream. g. Click **Save**. h. Configure a second Port Group. with the alias, select ports and in the Ports field, select GigaSMART Load Balancing, select in the GigaStream field, and then click **Save**.
3	Configure a GigaSMART group and associate it with a GigaSMART engine port.	a. From the device view, select GigaSMART > GigaSMART Groups. b. Click **New**. c. Enter the name in the Alias field. d. In the Port List field, select an engine port. e. Click **Save**.
4.	Create a virtual port. **Note**: You must enable GTP Overlap when configuring a virtual port for GTP overlap mapping.	a. From the device view, select GigaSMART > **Virtual Ports** > V**irtual Ports**. b. Click **New** c. Enter the name in the Alias field. d. In the GigaSMART Group field, select the group. e. Select GTP Overlap. f. Click **Save**.
5.	Create the GTP Whitelist	a. From the device view, select GigaSMART > **GTP Whitelist**. b. Click **New**. c. Enter Whitelist in the Alias field d. Go to Task 6.
6.	Fetch whitelist files from a specified location to populate the GTP forward list.	a. On the GTP Whitelist page, select Bulk Upload. b. Select Bulk Entry Operation for IMSI Upload Type.

Task	Description	UI Steps
		c. Select Upload from URL from the Bulk Upload Type list. d. Enter the URL in the Enter Remote URL field. e. Click **Save**.
7.	Associate the GigaSMART group to the GTP forward list.	a. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups b. Select GS Group GS31 created in Task 3 and click Edit c. Under GTP Whitelist, click on the GTP Whitelist Alias field and select Whitelist. d. Click **Save**.
8.	Configure the GigaSMART operation for GTP forward listing.	a. From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. b. Click **New**. c. Enter the name in the Alias field. d. Select the GigaSMART Group from the GigaSMART Groups list. e. Select GTP Whitelist from the GigaSMART Operations (GSOP) list f. Select Load Balancing from the GigaSMART Operations (GSOP) list. g. Configure Load Balancing as follows: • Select Stateful • Set Type to GTP • Select Hashing • Select IMSI h. Click **Save**.
9.	Configure the GigaSMART operation for GTP flow sampling.	a. From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. b. Click **New**. c. Enter the name in the Alias field.

Task	Description	UI Steps
		d. Select the GigaSMART Group from the GigaSMART Groups list. e. Select Flow Sampling-GTP. f. Select Load Balancing from the GigaSMART Operations (GSOP) list. g. Configure Load Balancing as follows: • Select Stateful • Set Type to GTP • Select Hashing • Select IMSI h. Click **Save**.
10.	Configure the first level maps. In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	a. Configure the first map as follows: • Alias. • Type and Subtype • Source • Destination • Rule 1 • Rule 2 • Click **Save**.
11.	Configure the first second level GTP overlap map for GTP forward listing If there is a match to an IMSI in the forward list for GTP version 1 traffic, it is then forwarded to load balancing port group.	a. Select **Maps** > **Maps** > **Maps**. b. Click **New**. c. Configure the map. • Alias • Type and Subtype • Source • Destination • GSOP • Rule 1. d. Click **Save**.
12.	Configure a second level map for GTP flow sampling, the flow sampling map. If there is not a match to an IMSI in the forward list, the traffic flow is sampled based on the rules in this map. Accepted packets are forwarded to load balancing port	a. Select **Maps** > **Maps** > **Maps**. b. Click **New**. c. Configure the map. • Alias

Task	Description	UI Steps
	group.	• Type and Subtype. • Source. • Destination • GSOP • Rule 1 d. Click **Save**.
13.	Configure the next second level GTP overlap map for GTP forward listing If there is a match to an IMSI in the forward list for GTP version 2 traffic, it is then forwarded to load balancing port group.	a. Select Maps > Maps > Maps. b. Click New. c. Configure the map. • Alias • Type and Subtype • Source • Destination • GSOP • Rule 1 d. Click **Save**.
14.	Configure the next second level map for GTP flow sampling. If there is not a match to an IMSI in the forward list as evaluated by the second level GTP whitelisting map *WLMAP2*, the traffic flow is sampled based on the rules in this map. Accepted packets are forwarded to load balancing port group.	a. Select **Maps** > **Maps** > **Maps**. b. Click **New**. c. Configure the map: • Alias • Type and Subtype: Sample Overlap • Source • Destination • Rule 1 d. Click **Save**.
15.	Configure a map group. Add the GTP forward listing and the two GTP flow sampling maps configured in previous steps.	a. Select **Maps** > **Map Groups**. b. Click **New**. c. Enter OverlapMap in the Alias field. d. In the Maps field, select the maps. e. Click Save.

Overlap Map Statistics

Overlap maps are displayed based on the following:

- If at least 1 flow-sample map accepts the packets among all the pairs of overlap maps, the Sample (Tx) counters in the GTP Interface stats will be incremented. If more than 1 pair of maps accepts the packets, the Sample (Tx) counters in the GTP Interface stats is incremented only once.
- If at least 1 whitelist map matches the packets among all the pairs of overlap maps, the Sample (Tx) counters in the GTP Interface stats will be incremented. If more than 1 pair of maps matches the packets, the Sample (Tx) counters in the GTP Interface stats is incremented only once.
- If there are no WL maps and all flow sample maps are no rule match, then Sample(Tx) and Sample Out counters in the GTP Interface stats is not incremented.

5G Stateful Session Recovery

Required License: 5G Correlation

5G stateful session recovery provides session persistence for GigaSMART 5G applications, including 5G flow filtering, 5G forward listing, and 5G flow sampling. In this method of recovery, 5G sessions are backed up periodically so that they can be recovered faster after a GigaSMART line card reboot or a node reboot.

5G stateful session recovery requires additional memory for storing backups. GigaVUE-HC3 has the required memory.

Using 5G stateful session recovery, the 5G session tables in the GigaSMART line card memory is periodically backed up to the control card memory on the node and stored.

You should configure an interval for how often the backups occur, such as every 10 minutes. If 5G stateful session recovery is enabled and the GigaSMART line card is rebooted, the 5G session tables is restored automatically following the reboot.

The last stored backup file is downloaded from the control card to the GigaSMART line card using FTP. The session table is repopulated from the last stored backup file to each GigaSMART engine, up to 8 engines. Packet count statistics for sessions are saved and restored.

Depending on the size of the session table, the amount of time to restore from the backup might take as much as 3 minutes. During that interval, traffic is blocked to the virtual port on the GigaSMART line card. Once the session table is read and populated, traffic is allowed.

Depending on the interval between backups, there could be differences between the stored state and the current state of the system, for example, map configuration could change, or sessions could be added, modified, or deleted.

Load balancing information is not persisted, so after a session table is repopulated, a session that was once sent to one load balanced port may be sent to a different load balanced port after the reboot. However, for SUPI-based load balancing, the traffic is sent to the same port as it was before the reboot.

5G stateful session recovery works in a cluster environment; however, the cluster leader must remain the same.

Configure 5G Stateful Session Recovery

To enable 5G stateful session recovery, as well as to configure timers, do the following:

1. On the left navigation pane, click  , and then select **Physical>Nodes**.
2. From the device view, select **GigaSMART** > **GigaSMART groups**.
3. Click on the alias of the **GigaSMART group**.
4. Select **GTP Persistence**in the **GTP Persistence** fields under GigaSMART Parameters as shown in the Figure 22GTP Persistence GigaSMART Parameters. The timers are pre-configured with default values.



Figure 22 *GTP Persistence GigaSMART Parameters*

Use the **System** widget on the Overview page to determine the amount of memory. The size of memory is 24Gb in an upgraded system.

View Backup and Restore information

To view the System information, select **Overview** from the Navigation pane. The amount of free and used memory is displayed in the **Memory** field.

To view backup and restore information for GTP Persistence:

1. Select **GigaSMART > GigaSMART Groups > GigaSMART Group**.
2. Click on the alias of the GigaSMART group.

A Quick View appears for the selected GigaSMART group.

3. Scroll down to GTP Persistence. In Figure 23GTP Persistence Information, GigaSMART Group gsgrp-1_4_e1 is selected and the Quick View is displayed.

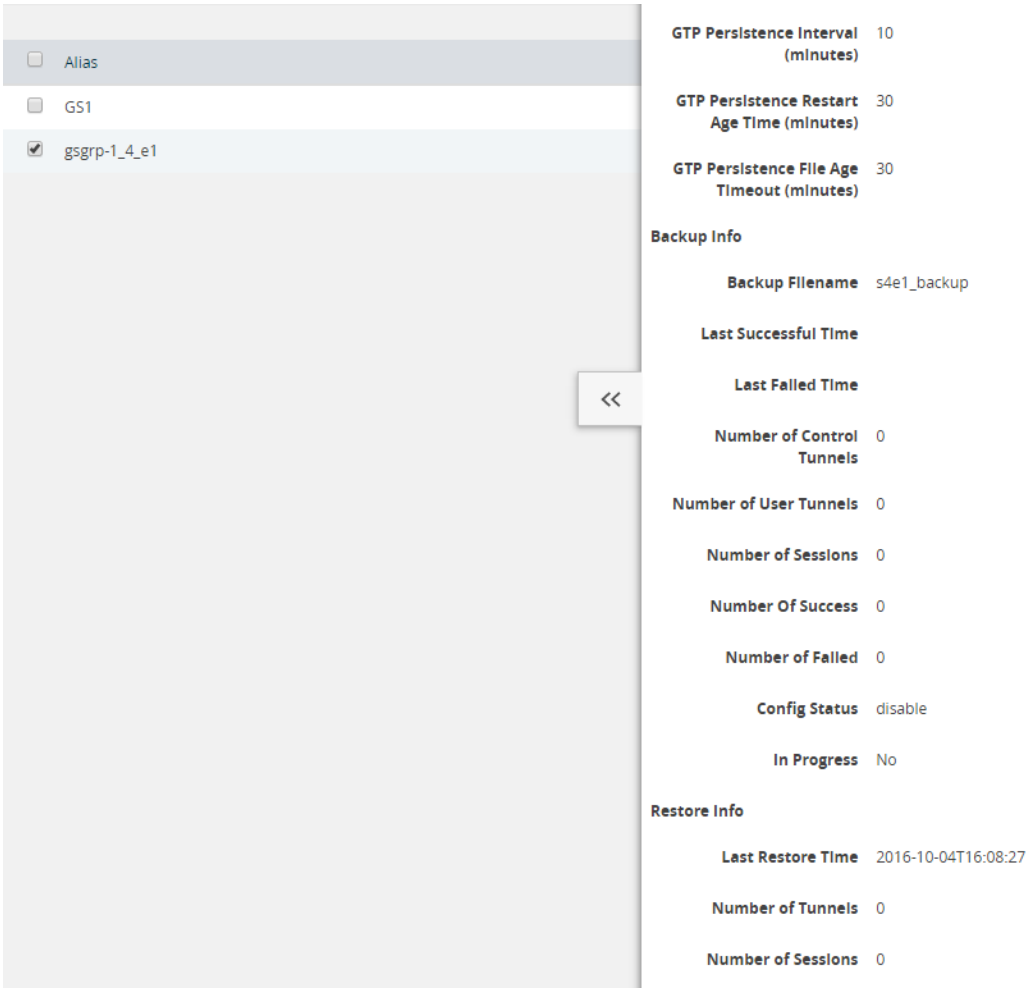


Figure 23 *GTP Persistence Information*

The following table describes persistence information.

Table 1: GigaSMART GTP Persistence Information

Name	Format
Backup Info	
Backup file name	The internal name of the backup file.
Last successful time	The time stamp of the last successful backup.

Name	Format
Last fail time	The time stamp of the last failed backup.
Number of control tunnels	The number of control tunnels backed up.
Number of user tunnels	The number of user tunnels backed up.
Number of sessions	The number of sessions backed up.
Number of success	The number of successful backups.
Number of failed	The number of failed backups.
Config Status	The status of a backup, which is either Enabled or Disabled.
In Progress	The progress, which is either Yes or No.
Restore Info	
Last restore time	The time stamp of the last restore.
Number of tunnels	The number of tunnels restored.
Number of sessions	The number of sessions restored.

Remove Backup files

To delete backup files, do the following:

1. Select the alias of GigaSMART Group.
2. Click **Edit**.
3. Scroll down to GTP Persistence (refer to Figure 24GTP Backup Files Delete)
4. Click **Delete All** under **GTP Backup Files**.

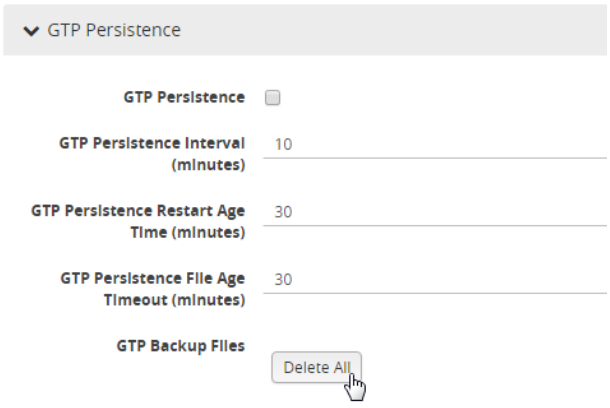


Figure 24 *GTP Backup Files Delete*

5G Whitelisting

Required License : 5G Whitelisting

5G Whitelisting selects specific subscribers based on SUPI. The whitelist contains up to 2,000,000 subscriber SUPIs in Gen2 GigaSMART. For subscribers in the whitelist, 100% of their traffic is always sent to a specified tool port.

For example, when a subscriber session comes in, 5G whitelisting checks the SUPI of the subscriber.

If the incoming SUPI matches an SUPI in the whitelist, all sessions associated with that SUPI are sent to the tool port or load balancing group specified in the whitelist map.

Configure Whitelist Maps

The whitelist maps are configured per GigaSMART group. Each whitelist map, associated with the same vport, uses the same underlying whitelist.

Up to ten (10) whitelist maps are supported. Multiple whitelist maps provide a granular selection of tool ports for whitelisting. Using multiple maps, traffic can be segregated and sent to multiple destinations. Whitelist map rules allow you to select the subset of SUPIs sent to a particular tool.

Each whitelist map can contain up to four rules. The rules specify the type of traffic to be whitelisted by that map. Within any single map, the rules are evaluated in order. The rules in the first map have a higher priority than the rules in the second, third, and subsequent maps.

The rules are specified based on the Data Network Name (DNN). A DNN can be specified in a rule of a Second Level Flow Whitelist map. 5G Whitelist map contains only DNN specific filters.

For DNN, you must specify a pattern (a name) to match. Use DNN to direct the traffic that matches the pattern to a specific tool.

A DNN pattern is for example, three.co.uk. Wildcard prefixes and suffixes are supported, for example, *mobile.com or *ims*. The pattern can be specified in up to 100 case-insensitive alphanumeric characters and can include the following special characters: period (.), hyphen (-), and wildcard (*). A standalone wildcard (*) is not allowed for DNN.

Each new subscriber session will be evaluated by the whitelist maps in the order of priority, which, by default, is the order in which the maps were created.

When a subscriber session comes in, 5G whitelisting will check the SUPI of the subscriber. If the SUPI is present in the whitelist, the rules in the first whitelist map is evaluated to qualify the match further. Otherwise, the packet is evaluated against the rules in the subsequent whitelist maps for a possible match.

NOTE: Both maps can specify the same destination.

Rules can be added to, or deleted from, a whitelist map. Use the Add a Rule button to add a new whitelist rule (a pass rule). Click x to delete a rule. A rule in a whitelist map cannot be edited. To edit a rule, first delete it, then recreate it.

The default map configuration DNN specified in the map, continues to be supported. If the incoming SUPI matches an SUPI in the whitelist, the session will be sent to the tool port, GigaStream, or Load Balancing group specified in the whitelist map. Whitelist maps cannot contain any other rules such as GigaSMART rules (gsrule), flow filtering rules (flowrule), or flow sampling rules (flowsample).

5G whitelist-based forwarding is performed prior to 5G flow sampling (rule-based flow sampling) and 5G flow filtering.

NOTE: For 5G second level maps, a maximum of fifteen maps can be attached to a vport. For example, for the same vport you can have five whitelist maps and ten flow sampling maps, or ten whitelist maps, and five flow sampling maps. In addition, you can have a collector map, which is not counted.

Whitelist maps cannot contain any other rules such as GigaSMART rules (gsrule), flow filtering rules (flowrule), or flow sampling rules (flowsample).

Change Priority of Whitelist Maps

Priority is set as per the order defined in the policy YAML file within the type.

Delete Whitelist Maps

When a whitelist map is deleted, the priority of the remaining whitelist maps are re-prioritized.

For example, if the first whitelist map is deleted, the second whitelist map increases in priority.

For the deleted whitelist map, the traffic associated with the rules in the map is reevaluated and then passed to subsequent maps.

When a whitelist map is re-prioritized, the existing sessions are reevaluated according to the new priority of the map. The traffic associated with the rules in the map are reevaluated and then passed to subsequent maps.

When the last whitelist map is deleted, the traffic associated with the rules in the map is also reevaluated before being passed to subsequent maps. But the traffic associated with the rules in maps that were not matched, are not reevaluated because that traffic was already passed to subsequent maps.

When a single whitelist entry is added, whitelisting is applied for new as well as existing subscribers.

When a new whitelist file is fetched, whitelisting is applied only for new subscribers.

Whitelisted traffic is then sent to the port or load balancing group specified in the whitelist map.

Entries in the whitelist can be deleted one at a time. Each entry is a single SUPI.

When a whitelist entry is deleted, the session associated with the whitelist entry stays active and traffic is still sent to the whitelist map. The whitelist session is not reevaluated or passed to subsequent maps.

To delete a single entry from the whitelist, select Individual Entry Operation, refer to Delete GTP Whitelist Maps in GigaSMART GTP Whitelisting and GTP Flow Sampling

To perform the following, refer to GigaSMART GTP Whitelisting and GTP Flow Sampling:

- Create Whitelist
- Apply Whitelist
- Delete Entry from Whitelist
- Delete Multiple Entries from Whitelist
- Delete Whitelist
- Destroy Whitelist

User Plane Node Traffic Monitoring

In User Plane Node (UPN) traffic monitoring, the UPN processes the subscriber information that is extracted from the Packet Forwarding Control Protocol (PFCP) packets and correlates with the GTP-u traffic without the CPN.

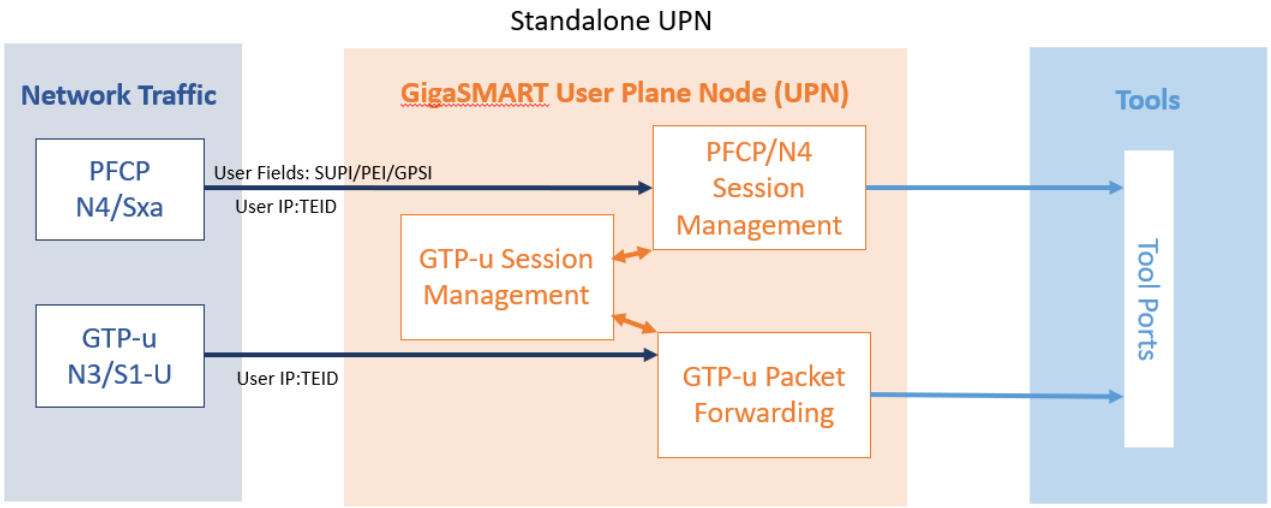
PFCP (Packet Forwarding Control Protocol) is a 3GPP Protocol that is communicated on the Sxa, Sxb, N4 Interface between the Control Plane (CP) elements and User Plane (UP) elements. The CP element programs the UP element with policies on how to forward packets.

PFCP allows an optional Informational Elements (IE) that contains SUPI/IMSI, PEI/IMEI, GPSI/MSISDN information of the Subscriber during PFCP Session Establishment Request. If your network enables these optional information elements, then you can use the Stand-Alone User Plane Node traffic monitoring.

The UPN performs the following activities in the Standalone mode:

- Processes the PFCP session establishment request and extracts the SUPI/IMSI, PEI/IMEI, GPSI/MSISDN User IP and TEID for both end points. The information in the PFCP traffic is used to populate the UPN's session table.
- Correlates GTP-U traffic based on the subscriber information from PFCP. GTP-U look up is correlated based on the IP and TEID .
- When UPN standalone mode is used, forward lists are limited to IMSI/SUPI only, and flowsampling maps for filtering or sampling are limited to IMSI/SUPI and IMEI/PEI. You will not be able to use any other CP fields for filtering or sampling, including APN, DNN, QCI, 5QI, ECGI, NCGI, NSI.

The following diagram explains the functioning of UPN in Standalone mode:



Rules and Notes

- You can create a Stand-Alone UPN traffic monitoring session only when there is atleast SUPI/IMSI in the user fields.
- The UPN only supports PFCP traffic from Sxa for LTE, N4 for 5G. GTP and 5G maps are required to forward traffic.
- The UPN determines if PFCP is from the N4 interface if it has QFI information.

Configure Stand-Alone User Plane Node Traffic Monitoring

To configure Stand-Alone User Plane Node Traffic Monitoring, refer to Configure CPN-UPN Communication Solution using Ansible

NOTE: You must enable the stand-alone mode in UPN while creating the solution in Ansible. Once the stand-alone mode is enabled, UPN cannot connect with the CPN, and you cannot the change the mode. To change the mode, you need to delete the UPN from the solution and add a new one .

Control Plane Metadata

The mobility control plane metadata feature is used to support the metadata export in JSON format for each transaction occurring in the 3G/4G/5G network.

The mobility network core control traffic is tapped and mirrored to GTP/HTTP2 GigaSMART engines for correlation. The GTP/HTTP2 correlation engine generates metadata for a transaction. It can either be a GTP transaction for 3G/4G control traffic or an HTTP2 transaction for 5G control traffic. The GTP correlation application supports generating metadata for 3G and 4G network. The 5G CPN application supports generating the control plane metadata for 5G network.

The metadata of subscriber activities is extracted using GTP or 5G correlation. After extraction, the metadata is exported to the data collection center in JSON format.

The records are exported for the following transactions:

For 3G:

- Create PDP
- Update PDP
- Delete PDP

For LTE:

- Create Session
- Delete Session
- Create Bearer
- Modify Bearer
- Update Bearer
- Delete Bearer

For 5G

- Create SM Context
- Modify SM Context
- Release SM Context
- N1 N2 Transfer

By default, values such as IMSI/SUPI, IMEI/PEI, MSISDN/GPSI and APN/DNN are added from the Create record to Modify and Delete records.

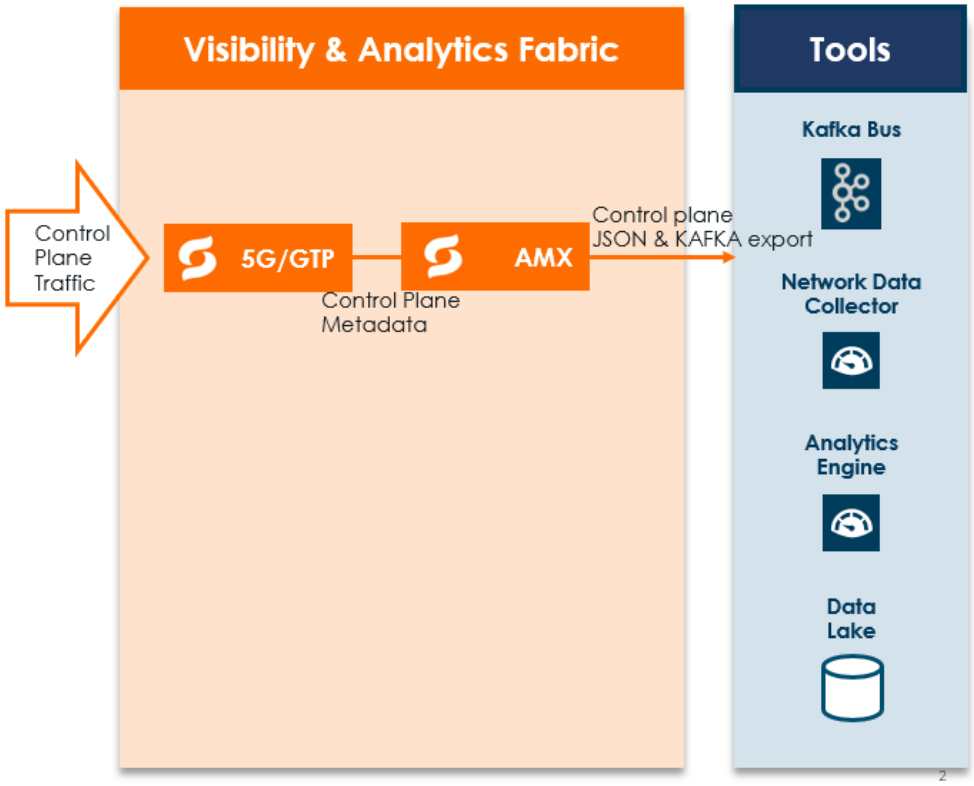
The control plane metadata can be exported in the following two formats:

- Flat
- Hierarchical

The mobility control metadata JSON is encoded from the session table information, current transaction and sent to the AMX application running on GigaVUE V Series Node which is connected at L4 protocol (UDP) level.

Once the transaction is complete in the session table and the metadata is created, it is scheduled for export to AMX. The AMX stores the received mobility control metadata in its internal database and is periodically sent to customer tools either by Kafka bus or as JSON file upload. For information on the export of 3G/4G/5G control plane metadata by AMX, refer Application Metadata Exporter.

NOTE: When the 5G subscriber reattaches through the createSMContext before the clean-up is done for the previous user tunnel with the same IMSI and PDU session id, the UE IP address and the Tunnel ID that are populated based on the old user tunnel information will be exported.



Configuration of 3G/4G Control Plane Metadata using Ansible

To configure 3G/4G control plane metadata, perform the following steps:

S.No	Steps	Refer to..
1.	Creating Inventory Directory	Creating Inventory Directory
2.	Creating ***fmInfo.yml***	Creating fmInfo.yml
3.	Creating ***ansible_inputs.json***	Creating ansible_inputs.json
4.	Creating 3G/4G inventory file	Creating 3G/4G control plane metadata inventory file
5.	Creating **host_vars** directory	Creating host_vars directory
6.	Creating **host_vars** files	Creating host_vars files

Creating Inventory Directory

Create an Inventory Directory to store all the 3G/4G control plane metadata related configuration files.

```
username@fmreg26:~$ mkdir CP-Metadata-Solution
username@fmreg26:~$ ls -l
drwxr-xr-x 2 ddaniel fmtaf 4096 May 11 11:59 CP-Metadata-Solution
```

Creating fmInfo.yml

Create fmInfo.yml file inside the Inventory Directory that contains the information such as ip-address, username and password. To create the file, refer to Schema.

Creating ansible_inputs.json

Create 'ansible_inputs.json' file inside the Inventory Directory.

gigamon@fmreg26:~/CP-Metadata-Solution $ ls -l

-rw-r--r-- 1 gigamon fmtaf 355 May 11 12:24 ansible_inputs.json

-rw-r--r-- 1 gigamon fmtaf 172 May 11 12:10 fmInfo.yml

File name: ansible_inputs.json To create a file, refer to Schema.

Creating 3G/4G control plane metadata inventory file

Create the **mobility_inventory** file inside Inventory Directory.

gigamon@fmreg26:~/CP-Metadata-Solution$ touch mobility_inventory

gigamon@fmreg26:~/CP-Metadata-Solution$ ls -l

-rw-r--r-- 1 gigamon fmtaf 355 May 11 12:24 ansible_inputs.json

-rw-r--r-- 1 gigamon fmtaf 396 May 11 14:22 mobility_inventory

-rw-r--r-- 1 gigamon fmtaf 172 May 11 12:10 fmInfo.yml

The file contains the details of the following groups and provide the inputs to the groups as shown in the following table:

S.No	Groups-Input
Note: —You can provide the input or leave the field empty if you don't want to use the playbook to configure the following groups.	
1.	**Ports**—Name of the Cluster or standalone device IP that contains the ports that need to be configured.
2.	**IPInterfaceSolution**—Name of the Cluster or standalone device IP on which the IPInterfacesolution needs to be configured.
3.	**Tool Groups**—Name of the Cluster or standalone device IP on which the ToolGroup needs to be configured.
4.	**GigaStreams**—Name of the Cluster or standalone device IP on which the Gigastreams needs to be configured.
5.	**GTPWhitelist**—Name of the Cluster or standalone device IP on which the GTPWhitelist database needs to be configured.
6.	**Policies**—Name of the Global policy or policies.
7.	**GTP**— Name of the GTP node or nodes.
6.	**CPN**—Name of the CPN or CPNs.
7.	**UPN**—Name of the UPN or UPNs.
8.	**SAM**—Name of SAMs Exporter node or nodes.
9.	**Sites**—Name of the site or sites and the names of the CPN/UPN participating in the site or sites.
10.	**CUPS**—Name of the file containing information of the solution level RBAC Tags.

File name: mobility_inventory (Single GigaVUE-FM instance)

```
[IPInterfaceSolution]
ip_interface_solution

[ToolGroups]
cluster-two
cluster-one

[Gigastreams]
cluster-two
cluster-one

[GTPWhitelist]
cluster-two
cluster-one

[Ports]
cluster-two
cluster-one

[Policies]
gtp_policy

[GTP]
gtp_1

[CPN]

[UPN]

[SAM]
sam1

[Sites]
UK gtp_list='["gtp1"]' upn_list='[]' sam_list='[sam1]'

[MobilitySolution]
mobilitySolution1 non_cups_lte_global_policy=gtp_policy sites='["UK"]'
```

File name: cups_inventory (Multiple GigaVUE-FM instances)

```
[IPInterfaceSolution]
ip_interface_solution_config
```

```
[ToolGroups]
cluster-two fm_ip=192.168.36.2
cluster-one fm_ip=192.168.36.3

[Gigastreams]
cluster-two fm_ip=192.168.36.2
cluster-one fm_ip=192.168.36.3

[GTPWhitelist]
cluster-two fm_ip=192.168.36.2
cluster-one fm_ip=192.168.36.3

[Ports]
cluster-two fm_ip=192.168.36.2
cluster-one fm_ip=192.168.36.3

[Policies]
gtp_policy

[GTP]
gtp_1

[CPN]

[UPN]

[SAM]
sam1

[Sites]
UK gtp_list='["gtp1"]' cpn_list='[]'upn_list='[]'
sam_list='["sam1"]'
Dallas gtp_list='["gtp1"]' cpn_list='[]'upn_list='[]'
sam_list='["sam1"]'

[MobilitySolution]
mobilitySolution1 non_cups_lte_global_policy=gtp_policy sites='["UK"]' fm_ip=192.168.36.2
mobilitySolution2 non_cups_lte_global_policy=gtp_policy sites='["Dallas"]' fm_ip=192.168.36.3
```

Creating host_vars directory

Create **host_vars** directory inside the Inventory Directory.

gigamon@fmreg26:~/CP-Metadata-Solution$ mkdir host_vars

gigamon@fmreg26:~/CP-Metadata-Solution$ ls -l

-rw-r--r-- 1 gigamon fmtaf 355 May 11 12:24 ansible_inputs.json

-rw-r--r-- 1 gigamon fmtaf 396 May 11 14:22 mobility_inventory

-rw-r--r-- 1 gigamon fmtaf 172 May 11 12:10 fmInfo.yml

drwxr-xr-x 2 gigamon fmtaf 4096 May 11 14:48 host_vars

Creating host_vars files

Every unique element under each group in the **mobility_inventory** file needs to have a file, with the same name as the element, inside **host_vars** directory. This file has properties of the groups that it belongs to.

Following are the templates of various host_vars files.

Prerequisite

```
validate_certs: false
Ports:
- port:
  - 1/1/x1
  - 1/1/x2
  adminStatus: enable
  type: network

GTPWhitelist:
- alias: gtp1
  imsi: 310260564627811,310260564627812
  state: present
- alias: gtp2
  inputFile: './whitelistKeys/TenIMSIs_Valid.txt'
  state: present

Gigastreams:
- alias: toolGS_C11
  ports:
  - 4/1/x1
  - 4/1/x2
  type: hybrid
  state: present
- alias: toolGS_C12
  ports:
  - 4/1/x3..x4
  type: hybrid
  state: present


ToolGroups:
- alias: pgGrp_C11
  ports:
  - 2/1/x1
  smartLb: false
  type: tool
  state: present
- alias: pgGrp_C12
  ports:
  - 2/1/x2
  smartLb: false
  type: tool
```

```
   state: present
```

Site

For information about Site, refer to Schema.

cpNode

For information about cpNode, refer to Schema.

upNode

For information about upNode, refer to Schema.

5GPolicy

For information about 5GPolicy, refer to Schema.

LTEPolicy

For information about LTEPolicy, refer to Schema.

Deployment of 3G/4G Control Plane Metadata Solution

To deploy the 3G/4G control plane metadata solution, follow these steps:

Set up of additional variable for Single GigaVUE-FM instance

For a single GigaVUE-FM instance deployment, you must set an additional environment variable as follows.

```
export ANSIBLE_FM_IP=192.168.36.2
```

It searches the login details of the GigaVUE-FM IP in ***fmInfo.yml*** file.

Execute the Playbook

You can execute the playbook and deploy the 3G/4G control plane metadata solution using the following command:

```
/usr/bin/ansible-playbook -e '@~/CP-Metadata-Solution/ansible_inputs.json' -i
~/CP-Metadata-Solution/mobility_inventory /usr/local/share/gigamon-
ansible/playbooks/mobility_solution/deploy_mobility_solution.yml
```

If the **fmInfo** file is encrypted, use the following command to execute and deploy 3G/4G control plane metadata solution:

```
/usr/bin/ansible-playbook -e '@~/CP-Metadata-Solution/ansible_inputs.json' --
ask-vault-pass -i ~/CP-Metadata-Solution/mobility_inventory
/usr/local/share/gigamon-ansible/playbooks/mobility_solution/deploy_mobility_
solution.yml
```

NOTE: The multiple YML files created inside the **host_vars** are concatenated, converted into JSON format and sent to GigaVUE-FM.

Configuration of 5G Control Plane Metadata using Ansible

To configure 5G control plane metadata, perform the following steps:

S.No	Steps	Refer to..
1.	Creating Inventory Directory	Creating Inventory Directory
2.	Creating ***fmInfo.yml***	Creating fmInfo.yml
3.	Creating ***ansible_inputs.json***	Creating ansible_inputs.json
4.	Creating 5G inventory file	Creating 5G control plane metadata inventory file
5.	Creating **host_vars** directory	Creating host_vars directory
6.	Creating **host_vars** files	Creating host_vars files

Creating Inventory Directory

Create an Inventory Directory to store all the 5G control plane metadata related configuration files.

```
username@fmreg26:~$ mkdir 5g-Metadata-Solution
username@fmreg26:~$ ls -l
drwxr-xr-x 2 ddaniel fmtaf 4096 May 11 11:59 5g-Metadata-Solution
```

Creating fmInfo.yml

Create fmInfo.yml file inside the Inventory Directory that contains the information such as ip-address, username and password. To create the file, refer to Schema.

Creating ansible_inputs.json

Create 'ansible_inputs.json' file inside the Inventory Directory.

gigamon@fmreg26:~/5G-Metadata-Solution $ ls -l

-rw-r--r-- 1 gigamon fmtaf 355 May 11 12:24 ansible_inputs.json

-rw-r--r-- 1 gigamon fmtaf 172 May 11 12:10 fmInfo.yml

File name: ansible_inputs.json To create a file, refer to Schema.

Creating 5G control plane metadata inventory file

Create the **mobility_inventory** file inside Inventory Directory.

gigamon@fmreg26:~/5g-Metadata-Solution$ touch mobility_inventory

gigamon@fmreg26:~/5g-Metadata-Solution$ ls -l

-rw-r--r-- 1 gigamon fmtaf 355 May 11 12:24 ansible_inputs.json

-rw-r--r-- 1 gigamon fmtaf 396 May 11 14:22 mobility_inventory

-rw-r--r-- 1 gigamon fmtaf 172 May 11 12:10 fmInfo.yml

The file contains the details of the following groups and provide the inputs to the groups as shown in the following table:

S.No	Groups-Input
Note: —You can provide the input or leave the field empty if you don't want to use the playbook to configure the following groups.	
1.	**Ports**—Name of the Cluster or standalone device IP that contains the ports that need to be configured.
2.	**IPInterfaceSolution**—Name of the Cluster or standalone device IP on which the IPInterfacesolution needs to be configured.
3.	**Tool Groups**—Name of the Cluster or standalone device IP on which the ToolGroup needs to be configured.
4.	**GigaStreams**—Name of the Cluster or standalone device IP on which the Gigastreams needs to be

S.No	Groups-Input
	configured.
5.	**GTPWhitelist**—Name of the Cluster or standalone device IP on which the GTPWhitelist database needs to be configured.
6.	**Policies**—Name of the Global policy or policies.
7.	**GTP**— Name of the GTP node or nodes.
6.	**CPN**—Name of the CPN or CPNs.
7.	**UPN**—Name of the UPN or UPNs.
8.	**SAM**—Name of SAMs Exporter node or nodes.
9.	**Sites**—Name of the site or sites and the names of the CPN/UPN participating in the site or sites.
10.	**CUPS**—Name of the file containing information of the solution level RBAC Tags.

File name: mobility_inventory (Single GigaVUE-FM instance)

```
[IPInterfaceSolution]
ip_interface_solution

[ToolGroups]
cluster-two
cluster-one

[Gigastreams]
cluster-two
cluster-one

[GTPWhitelist]
cluster-two
cluster-one

[Ports]
cluster-two
cluster-one

[Policies]
5g_policy_1

[GTP]


[CPN]
5g_cpn

[UPN]

[SAM]
sam1

[Sites]
UK cpn_list='["5g_cpn"]' upn_list='[]' sam_list='["sam1"]'

[MobilitySolution]
mobilitySolution1 cups_5g_global_policy=5g_policy_1 sites='["UK"]'
```

File name: cups_inventory (Multiple GigaVUE-FM instances)

```
[IPInterfaceSolution]
ip_interface_solution_config
```

```
[ToolGroups]
cluster-two fm_ip=192.168.36.2
cluster-one fm_ip=192.168.36.3

[Gigastreams]
cluster-two fm_ip=192.168.36.2
cluster-one fm_ip=192.168.36.3

[GTPWhitelist]
cluster-two fm_ip=192.168.36.2
cluster-one fm_ip=192.168.36.3

[Ports]
cluster-two fm_ip=192.168.36.2
cluster-one fm_ip=192.168.36.3

[Policies]
5g_policy_1

[GTP]

[CPN]
5g_cpn

[UPN]

[SAM]
sam1

[Sites]
UK cpn_list='["5g_cpn"]' upn_list='[]' sam_list='[sam1]'
Dallas cpn_list='[5g_cpn]' upn_list='[]' sam_list='[sam1]'

[MobilitySolution]
mobilitySolution1 cups_5g_global_policy=5g_policy_1 sites='["UK"]' fm_ip=192.168.36.2
mobilitySolution2 cups_5g_global_policy=5g_policy_1 sites='["Dallas"]' fm_ip=192.168.36.3
```

Creating host_vars directory

Create **host_vars** directory inside the Inventory Directory.

gigamon@fmreg26:~/5g-Metadata-Solution$ mkdir host_vars

gigamon@fmreg26:~/5g-Metadata-Solution$ ls -l

-rw-r--r-- 1 gigamon fmtaf 355 May 11 12:24 ansible_inputs.json

-rw-r--r-- 1 gigamon fmtaf 396 May 11 14:22 mobility_inventory

-rw-r--r-- 1 gigamon fmtaf 172 May 11 12:10 fmInfo.yml

drwxr-xr-x 2 gigamon fmtaf 4096 May 11 14:48 host_vars

Creating host_vars files

Every unique element under each group in the **mobility_inventory** file needs to have a file, with the same name as the element, inside **host_vars** directory. This file has properties of the groups that it belongs to.

Following are the templates of various host_vars files.

Prerequisite

```
validate_certs: false
Ports:
- port:
  - 1/1/x1
  - 1/1/x2
  adminStatus: enable
  type: network

GTPWhitelist:
- alias: gtp1
  imsi: 310260564627811,310260564627812
  state: present
- alias: gtp2
  inputFile: './whitelistKeys/TenIMSIs_Valid.txt'
  state: present

Gigastreams:
- alias: toolGS_C11
  ports:
  - 4/1/x1
  - 4/1/x2
  type: hybrid
  state: present
- alias: toolGS_C12
  ports:
  - 4/1/x3..x4
  type: hybrid
  state: present


ToolGroups:
- alias: pgGrp_C11
  ports:
  - 2/1/x1
  smartLb: false
  type: tool
  state: present
- alias: pgGrp_C12
  ports:
  - 2/1/x2
  smartLb: false
  type: tool
```

```
   state: present
```

Site

For information about Site, refer to Schema.

cpNode

For information about cpNode, refer to Schema.

upNode

For information about upNode, refer to Schema.

5GPolicy

For information about 5GPolicy, refer to Schema.

LTEPolicy

For information about LTEPolicy, refer to Schema.

Deployment of 3G/4G Control Plane Metadata Solution

To deploy the 3G/4G control plane metadata solution, follow these steps:

Set up of additional variable for Single GigaVUE-FM instance

For a single GigaVUE-FM instance deployment, you must set an additional environment variable as follows.

```
export ANSIBLE_FM_IP=192.168.36.2
```

It searches the login details of the GigaVUE-FM IP in ***fmInfo.yml*** file.

Execute the Playbook

You can execute the playbook and deploy the 5G control plane metadata solution using the following command:

```
/usr/bin/ansible-playbook -e '@~/5g-Metadata-Solution/ansible_inputs.json' -i
~/5g-Metadata-Solution/mobility_inventory /usr/local/share/gigamon-
ansible/playbooks/mobility_solution/deploy_mobility_solution.yml
```

If the **fmInfo** file is encrypted, use the following command to execute and deploy 5G control plane metadata solution:

```
/usr/bin/ansible-playbook -e '@~/5g-Metadata-Solution/ansible_inputs.json' --
ask-vault-pass -i ~/5g-Metadata-Solution/mobility_inventory
/usr/local/share/gigamon-ansible/playbooks/mobility_solution/deploy_mobility_
solution.yml
```

NOTE: The multiple YML files created inside the **host_vars** are concatenated, converted into JSON format and sent to GigaVUE-FM.

CPN-UPN Communication for Support of RAN and Network Slice Attributes

CPN-UPN Communication enables traffic management based on Radio Access Network (RAN) and Network Slice attributes in the Standalone User Plane Node (UPN). This feature allows you to enable flow sampling and forward listing based on the RAN attributes and flow sampling based on the Network Slice attributes. This is achieved by establishing communication between Control Plane Node and the Standalone UPN.

The CUPS architecture provides network visibility across the control and user planes for the 5G stand-alone packet core network. The UPN processes the GTP-u traffic using the subscriber correlation from the control plane PFCP protocols available at the user plane sites. UPN enriches the subscriber information extracted from the Gigamon PFCP message (GPFCP) packets received from CPN for enhanced flow sampling and forward listing capabilities.

By establishing a communication between CPN and UPN, the RAN and the Network Slice Attributes from the CPN are exchanged with UPN, which allows UPN to flow sample and forward list the traffic based on the RAN attributes and flow sample the traffic based on Network Slice attributes.

The CPN-UPN Communication solution includes the following components:

- CPN: can be 4G or 5G
- Standalone UPN

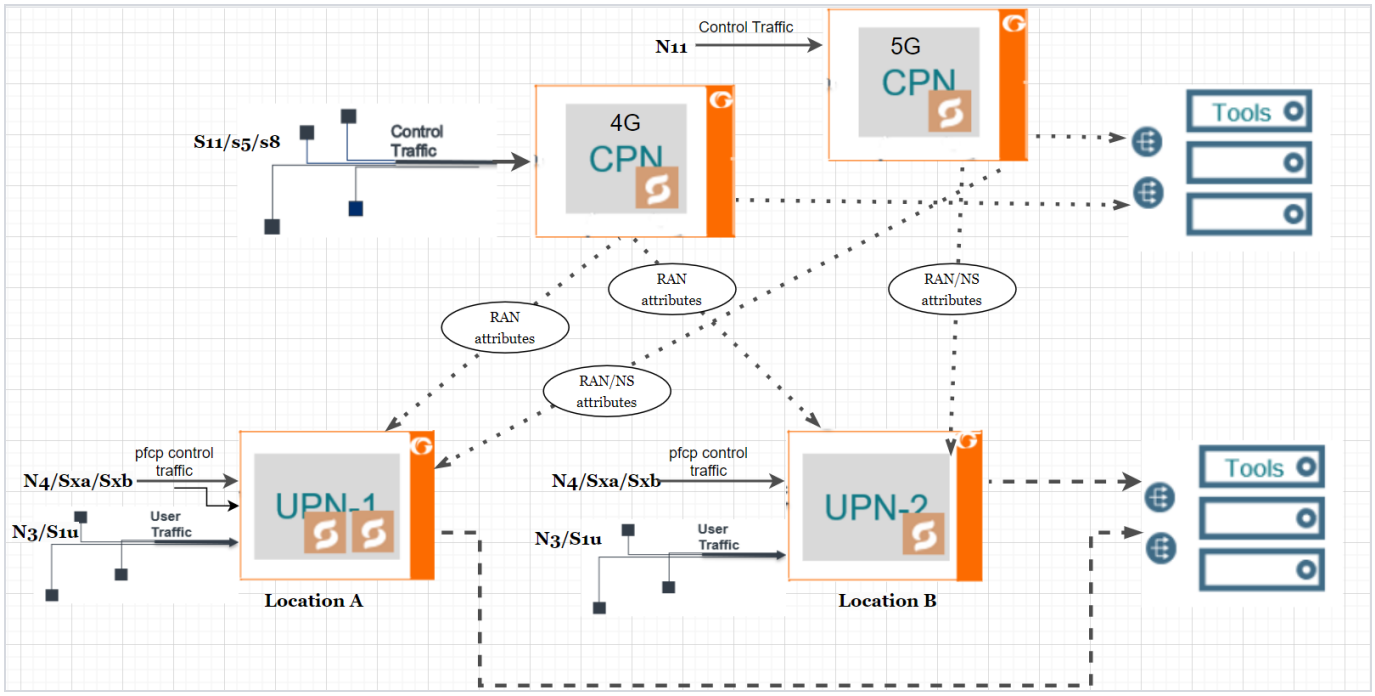
UPN supports the following parameters based on the attributes received from CPN:

RAN flow sampling and forward listing parameters

- E-UTRAN Cell Global Identifier (ECGI) (MCC+MNC+ECI)
- New Radio Cell Global Identifier (NCGI) (MCC+MNC+NCI)
- Tracking Area Code (TAC)
- Tracking Area Identity (TAI) (MCC+MNC+TAC)

Network Slice flow sampling parameters

- S-NSSAI with standardized SST (Slice/Service Type)
- S-NSSAI non-standard SST values with and without SD



The above topology explains how CPN and UPN communicate with each other.

In the above topology, the 4G CPN receives control traffic from the S11/S5/S8 interface, while the 5G CPN receives control traffic from the N11 interface. UPN-1 and UPN-2 receive PFCP traffic from the N4/Sxa/Sxb interface. Upon receiving PFCP traffic, UPN generates a session table that does not include the RAN and Network Slice Attributes.

The communication between CPN and UPN is facilitated using the GFPCP. GPFCP messages are never fragmented, and the message exchange between CPN and UPN is always less than 1500 bytes. Once the communication is established using GPFCP, the CPN pushes the RAN and Network Slice attributes to the UPN based on a request and appends the session table with the RAN and Network Slice attributes. In case of failures, UPN can also request the RAN and Network Slice attributes from CPN.

Once the UPN receives the user traffic from N3/S1u interface, based on the map rules configured in the UPN, the correlation takes place. With the RAN and Network Slice (NS) attributes from the CPN, the UPN now has the ability to flow sample and forward list the traffic based on the RAN and Network Slice Attributes. RAN and Network Slice information is not added to the packets sent out by UPN.

Keep in mind the following when configuring CPN-UPN Communication solution:

- UPN will back up and restore the sessions with the session attributes received from CPN.
- 4G CPN supports only RAN Attributes.
- The GPFCP messages are exchanged through UDP and are not encrypted.
- The UDP port used for the CPN-UPN communication must be within the ephemeral range (49152 to 65535).
- The GPFCP messages are not fragmented.

To know more about CPN-UPN communication configuration, communication statistics, map rules, upgrade, and rollback, refer to the following sections:

- CPN-UPN Communication Configuration
- View CUPS Communication Dashboard
- Map Rules for CPN-UPN Communication
- Upgrade Standalone UPN to CPN-UPN Communication Solution
- Rollback from CPN-UPN Communication Solution to Standalone UPN
- Quick Rollback from CPN-UPN Communication Solution to Standalone UPN

CPN-UPN Communication Configuration

CPN-UPN Communication can be configured in either of the following two ways. Refer to the following sections for step-by-step instructions:

- Configure CPN-UPN Communication using Ansible
- Configure CPN-UPN Communication using CLI

Configure CPN-UPN Communication using Ansible

To configure CPN-UPN Communication using Ansible:

Step No	Task	Refer to the following topics
1.	Configure CPN	GigaSMART 5G CUPS
2.	Configure UPN	Standalone UPN
3.	Install Gigamon Ansible Module	Installation and Configuration of Subscriber Intelligence Solution using Ansible
4.	Create Inventory Directory to store all the CPN-UPN Communication related configuration files	Create Inventory Directory
5.	Create ***fmInfo.yml*** for Ansible to access GigaVUE-FM	Create fmInfo.yml
6.	Create ***ansible_inputs.json***	Create ansible_inputs.json File
7	Create CPN-UPN Communication inventory file	Create an Inventory File for CPN-UPN Communication
8	Create **host_vars** directory	Create host_vars directory
9	Create **host_vars** files	Create host_vars files
10	Deploy CPN-UPN Communication solution in Ansible	Deploy CPN-UPN Communication Solution in Ansible
11	View the configured CPN-UPN Communication solution in GigaVUE-FM	View CPN-UPN Communication Solution in GigaVUE-FM
12	View Dashboard	View CUPS Communication Dashboard

Configure CPN-UPN Communication using CLI

Step No	Task	Refer to the following topics
1	Configure ports	ports
2	Configure GigaStream	gigastream
3	Create GPFCP profile	gpfcp profile
4	Configure UPN Interface profile	Configure Custom Interface Selection
5	Configure GigaSMART group	gsgroups
6	Configure IP interface	ip interface
7	Configure the exporter	apps exporter
8	Define the node role	gsparams
9	Create GigaSMART operations	gsop
10	Create a Vport	vport
11	Define first level map to send the traffic to Vport	map
12	Create a second level map with flow sample/ forward list rules	map
The following configuration example explains the entire configuration flow for 4G CPN (when one 4G CPN communicates with one UPN)		
CPN-UPN Communication configuration example		Configure CPN-UPN Communication solution using CLI

Configure CPN-UPN Communication Solution using Ansible

To configure CPN-UPN Communication using Ansible:

- Create Inventory Directory
- Create fmInfo.yml
- Create ansible_inputs.json File
- Create an Inventory File for CPN-UPN Communication
- Create host_vars directory
- Create host_vars files
- Create host_vars files
- Deploy CPN-UPN Communication Solution in Ansible
- View CPN-UPN Communication Solution in GigaVUE-FM

Create Inventory Directory

Create an *Inventory Directory* to store all the CPN-UPN Communication related configuration files.

```
username@fmreg26:~$ mkdir cupsSolution
```

Create ***fmInfo.yml***

Create ***fmInfo.yml*** file inside the inventory directory with the IP address, username, and password of GigaVUE-FM. This helps Ansible to fetch the details of GigaVUE-FM.

Create ansible_inputs.json File

Create a json file ***ansible_inputs.json*** inside the inventory directory that contains the path of the below mentioned files. This allows Ansible to identify and access the below mentioned files.

- ***fm_credential_file*** - Provide the path of the ***fmInfo.yml*** file.
- ***yaml_payload_path*** - Created automatically while running the cups playbook. Provide the path of the file that stores the payload sent to GigaVUE-FM.
- ***golden_payload_path*** - The payload of a successful cups solution deployment is saved in this file. Provide the file path that has the golden payload.
- ***deployment_report_path*** - Created automatically while running the cups playbook. This file stores the report of the deployment. Provide the path of the file.

Create an Inventory File for CPN-UPN Communication

Create ***mobility_inventory*** file inside the inventory directory.

Enter the details in the file as shown in the below example:

```
[IPInterfaceSolution]
<Name of the Cluster or standalone device IP that contains the ports that need to be configured.>

[ToolGroups]
<Name of the Cluster or standalone device IP on which the Tool Group needs to be configured.>
<Name of the Cluster or standalone device IP on which the Tool Group needs to be configured.>

[Gigastreams]
<Name of the Cluster or standalone device IP on which the Gigastreams needs to be configured.>

[GTPWhitelist]
<Name of the Cluster or standalone device IP on which the GTPWhitelist Data Base needs to be
configured.>

[Ports]
<Name of the Cluster or standalone device IP that contains the ports that need to be configured.>

[Policies]
<Name of the Global policy or policies.>

[CPN]
<Name of the CPN or CPNs.>

[UPN]
<Name of the CPN or CPNs.>

[Sites]
Site_Name1='["<Name of the CPN"]' upn_list='[]' sam_list='[]'
Site_Name2='[]' upn_list='["Name of the UPN"]' sam_list='[]'

[MobilitySolution]
<Name of the file CPN UPN Communication solution along with lts_policy file name, 5g_policy file
name, lte_policy file name and list of sites participating in the CPN UPN Communication solution.>
```

Create **host_vars** directory

Create **host_vars** directory inside the inventory directory to host the files that are used in the ***mobility_inventory*** file.

```
gigamon@fmreg26:~/cupsSolution$ mkdir host_vars
```

Create **host_vars** files

Each unique element like Ports, GigaStream, GTP forward list, and Tool Groups in the ***mobility_inventory*** file needs to have a file, with the same name as the element, inside **host_vars** directory.

For Example

- The element of name 'cpnUkLTE' under group 'CPN' has a file with name 'cpnUkLTE' inside host_vars directory. This file has the properties of the CPN.
- The element of name 'cluster-one' has a file with name 'cluster-one' inside host_vars directory. This has the properties of all the groups like Ports, IPInterface, ToolGroups, etc that the element is a member.

Below are a sample of a host_vars files.
Prerequisite:

```
---
validate_certs: false
Ports:
- port:
  - 1/1/x1
  - 1/1/x2
  adminStatus: enable
  type: network

GTPWhitelist:
- alias: gtp1
  imsi: 310260564627811,310260564627812
  state: present
- alias: gtp2
  inputFile: './whitelistKeys/TenIMSIs_Valid.txt'
  state: present

Gigastreams:
- alias: toolGS_C11
  ports:
  - 4/1/x1
  - 4/1/x2
  type: hybrid
  state: present
- alias: toolGS_C12
  ports:
```

```
  - 4/1/x3..x4
  type: hybrid
  state: present

ToolGroups:
- alias: pgGrp_C11
  ports:
  - 2/1/x1
  smartLb: false
  type: tool
  state: present
- alias: pgGrp_C12
  ports:
  - 2/1/x2
  smartLb: false
  type: tool
  state: present
```

Site

Mobility Solutions are generally spread across multiple geographically dispersed data centers called sites.

A site is a collection of the following:

- Network element functions.
- Traffic access points for interfaces of such network element functions.
- Visibility and Analytics Fabric (VAF) nodes.
- Traffic monitoring or analysis tool devices (called probes) that are locally connected without using any IP routed tunnels.

```
---
Site:
  # Name of the site
  alias: UK
  # 'skipDeployment' attribute is 'false' for the sites that are intended to be deployed in the
incremental deployment process
  skipDeployment: true
  # Tag values assigned to the site
  tags:
  - tagKey: Location
    tagValues:
    - UK
  # All the tools used in the site
  toolBindings:
  - alias: GeoProbe
    toolResourceType: GIGASTREAM
    toolClusterId: cluster-one
    toolResourceId: toolGS_C11
  - alias: EEA
    toolResourceType: PORTGROUP
```

```
    toolClusterId: cluster-one
    toolResourceId: pgGrp_C11
  # Network Ports. Fill as needed in the format 'clusterid:portId'
  networkPorts: []
  # Policies under 'siteOverrideOfPolicyArrangements' override global policies
  siteOverrideOfPolicyArrangements:
    forLTE:
      _file: /home/ddaniel/automationInventoryDirectory/host_vars/lte_policy_1.yml
    for5G: {}
  # Leave upNodes and cpNodes as empty
  upNodes: []
  cpNodes: []
```

cpNode

```
--
ProcessingNode:
  # Name of the control processing node
  alias: cpnUkLTE
  # Tags assigned to the processing node
  tags:
  - tagKey: Dept
    tagValues:
    - IT
    - Engg
  # Type of control node. Possible values for the nodeType are: 'PCPN_LTE', 'PUPN', 'PCPN_5G'
  nodeType: PCPN_5G
  # Location of the gigasmart engine port assigned to the processing node
  location:
    clusterId: cluster-one
    enginePorts:
    - 2/3/e1
  # IP interface that needs to be used by the processing node
  ipInterfaceAlias: dev3_IpIntUpn_1
  gtpControlSample: false
  gtpRandomSampling:
    enabled: false
    # min: 12, max: 48, multiples of 12hrs
    interval: 12
  numberOfLteSessions: 100000
  numberOf5gSessions: 100000
  # GS Group HTTP2 port list
  app5gHTTP2Ports:
  - 8080
  - 9000
  # Network Ports. Fill as needed in the format 'clusterid:portId'
  nodeOverrideNetworkPorts: []
  trafficSources:
  - networkFunctionName: cpn_pod1_SGW-C
    networkFunctionType: SGW-C
    tags:
    - tagKey: Dept
      tagValues:
      - IT
```

```
      - Engg
    networkFunctionInterfaces:
    - tunnelIdentifiers:
      - interfaceTunnelIdentifierType: IPADDRESS
        value: 255.255.255.0
        netMask: 198.51.100.42
      - interfaceTunnelIdentifierType: PORT
        value: 8805
      interfaceType: Sxa
      # Network Ports. Fill as needed in the format 'clusterid:portId'
      sourceOverrideNetworkPorts:
      - 192.168.65.8:8/1/x3
  #TCP loadbalancing properties applicable only for nodeType PCPN_5G
  appTcp:
    # Possible values for application are : 'broadcast', 'enhanced', 'drop'
    application: broadcast
    # Possible values for tcpControl are : 'broadcast', 'enhanced', 'drop'
    tcpControl: broadcast
    # To enable loadbalancing set value as true
    loadBalance: false
```

upNode

```
---
ProcessingNode:
  # Name of the control processing node
  alias: upnDallas
  # Tags assigned to the processing node
  tags:
  - tagKey: Dept
    tagValues:
    - IT
    - Engg
  # Type of user node. Possible values for the nodeType are: 'PCPN_LTE', 'PUPN', 'PCPN_5G'
  nodeType: PUPN
  # To make user node as standalone
  standAloneMode: true
  # Location of the gigasmart engine port assigned to the processing node
  location:
    clusterId: cluster-two
    enginePorts:
    - 6/3/e1
    - 6/3/e2
  # IP interface that needs to be used by the processing node
  ipInterfaceAlias: dev3_IpIntUpn_1
  gtpControlSample: false
  gtpRandomSampling:
    enabled: false
    # min: 12, max: 48, multiples of 12hrs
    interval: 12
  # Network Ports. Fill as needed in the format 'clusterid:portId'
  nodeOverrideNetworkPorts: []
  trafficSources:
  - networkFunctionName: upn_pod1_SGW-U
```

```
    networkFunctionType: SGW-U
    tags:
    - tagKey: Dept
      tagValues:
      - IT
      - Engg
    networkFunctionInterfaces:
    - tunnelIdentifiers:
      - interfaceTunnelIdentifierType: IPADDRESS
        mask: 255.255.255.0
        address: 198.58.100.45
      - interfaceTunnelIdentifierType: PORT
        value: 8805
      interfaceType: Sxa
      # Network Ports. Fill as needed in the format 'clusterid:portId'
      sourceOverrideNetworkPorts:
      - 192.168.65.9:9/1/x4
      - 192.168.65.9:9/1/x5
  - networkFunctionName: upn_pod2_UPF
    networkFunctionType: UPF
    tags:
    - tagKey: Dept
      tagValues:
      - IT
      - Engg
    networkFunctionInterfaces:
    - tunnelIdentifiers:
      - interfaceTunnelIdentifierValue:
          mask: 255.255.255.0
          address: 198.58.100.46
        interfaceTunnelIdentifierType: IPADDRESS
      - interfaceTunnelIdentifierValue:
          value: '2152'
        interfaceTunnelIdentifierType: PORT
      interfaceType: N11
      sourceOverrideNetworkPorts:
      - 192.168.65.9:9/1/x6
```

5GPolicy

```
---
5GPolicy:
  # gtpFlowTimeout is multiplied by 10 minutes to arrive at a timeout interval. (gtpFlowTimeout:
48 = 8 hours). Set this interval to match customer network's GTP session timeout for optimal
results
  gtpFlowTimeout: 48
  # gtpPersistence -- save state tables during reboot or box failure. Remove if not using
persistence
  gtpPersistence:
    # interval in minutes to save state table (min value is 10)
    interval: 10
    restartAgeTime: 30
    fileAgeTimeout: 30
  sampling:
```

```
    flowMaps:
    - alias : samplingMap2
      rules:
      - interface:
        dnn: internet.miracle
        pei: '*'
        supi: 46*
        gpsi:
        nas_5qi:
        tac:
        nci:
        plmndId:
        nsiid:
        # -- "controlPlanePercentage: specifies percentage of sampling at CPN (set to 100 for no
CPN sampling)
        controlPlanePercentage: 100
        # -- userPlanePercentage: specifies percentage of sampling at UPN (will be applied to all
UPNs -- use site override to set different rates per site)
        userPlanePercentage: 50
      tool: EEA
  whitelisting:
    whiteListAlias: gtp1
    flowMaps:
    - alias : whitelistMap2
      rules:
      - dnn: internet.miracle
        interface:
      - supi:
        ran:
      tool: EEA
  loadBalancing:
    # one of { flow5g }
    appType: flow5g
    # metric is the load balancing method -- one of { leastBw, leastPktRate, leastConn,
leastTotalTraffic, roundRobin, wtLeastBw, wtLeastPktRate, wtLeastConn, wtLeastTotalTraffic,
wtRoundRobin, flow5gKeyHash}
    metric: flow5gKeyHash
    # hashingKey -- ignored if metric is not of type 'flow5gKeyHash' -- one of { supi | pei| gpsi
}
    hashingKey: supi
```

LTEPolicy

```
---
LTEPolicy:
  # gtpFlowTimeout is multiplied by 10 minutes to arrive at a timeout interval. (gtpFlowTimeout:
48 = 8 hours). Set this interval to match customer network's GTP session timeout for optimal
results
  gtpFlowTimeout: 48
  # gtpPersistence -- save state tables during reboot or box failure. Remove if not using
persistence
  gtpPersistence:
    # interval in minutes to save state table (min value is 10)
    interval: 10
```

```
    restartAgeTime: 30
    fileAgeTimeout: 30
  sampling:
    flowMaps:
    - alias : samplingMap1
      rules:
      - interface:
        version: any
        apn: internet.miracle
        imei: 123*
        imsi:
        msisdn:
        qci:
        # -- controlPlanePercentage: specifies percentage of sampling at CPN (set to 100 for no
CPN sampling)
        controlPlanePercentage: 100
        # -- userPlanePercentage: specifies percentage of sampling at UPN (will be applied to all
UPNs -- use site override to set different rates per site)
        userPlanePercentage: 50
      tool: GeoProbe
  whitelisting:
    whiteListAlias: gtp1
    flowMaps:
    - alias : whitelistMap1
      rules:
      - version: v1
        interface:
        apn:
      tool: GeoProbe
  loadBalancing:
    # one of { gtp | aft | tunnel }
    appType: gtp
    # metric is the load balancing method -- one of # { leastBw, leastPktRate, leastConn,
leastTotalTraffic, roundRobin, wtLeastBw, wtLeastPktRate, wtLeastConn, wtLeastTotalTraffic,
wtRoundRobin, gtpKeyHash}
    metric: gtpKeyHash
    # hashingKey -- ignored if metric is not of type 'gtpKeyHash' -- one of { imsi | imei | msisdn
}
    hashingKey: imsi
```

Deploy CPN-UPN Communication Solution in Ansible

To deploy the CPN-UPN Communication solution, follow these steps:

1. Set an additional environment variable to search the login details of the GigaVUE-FM IP in ***fmInfo.yml*** file as follows.

   ```
   export ANSIBLE_FM_IP=<GigaVUE-FM IP Address>
   ```

2. Execute the playbook and deploy the CPN-UPN Communication solution using the following command:
   - If ***fmInfo.yml*** file is not encrypted, use the following command:

     ```
     /usr/bin/ansible-playbook -e '@~/cupsSolution/ansible_inputs.json' -i
     ~/cupsSolution/cups_inventory /usr/local/share/gigamon-ansible/playbooks/cups/deploy_
     cups.yml
     ```

   - If I***fmInfo.yml*** file is encrypted, use the following command:

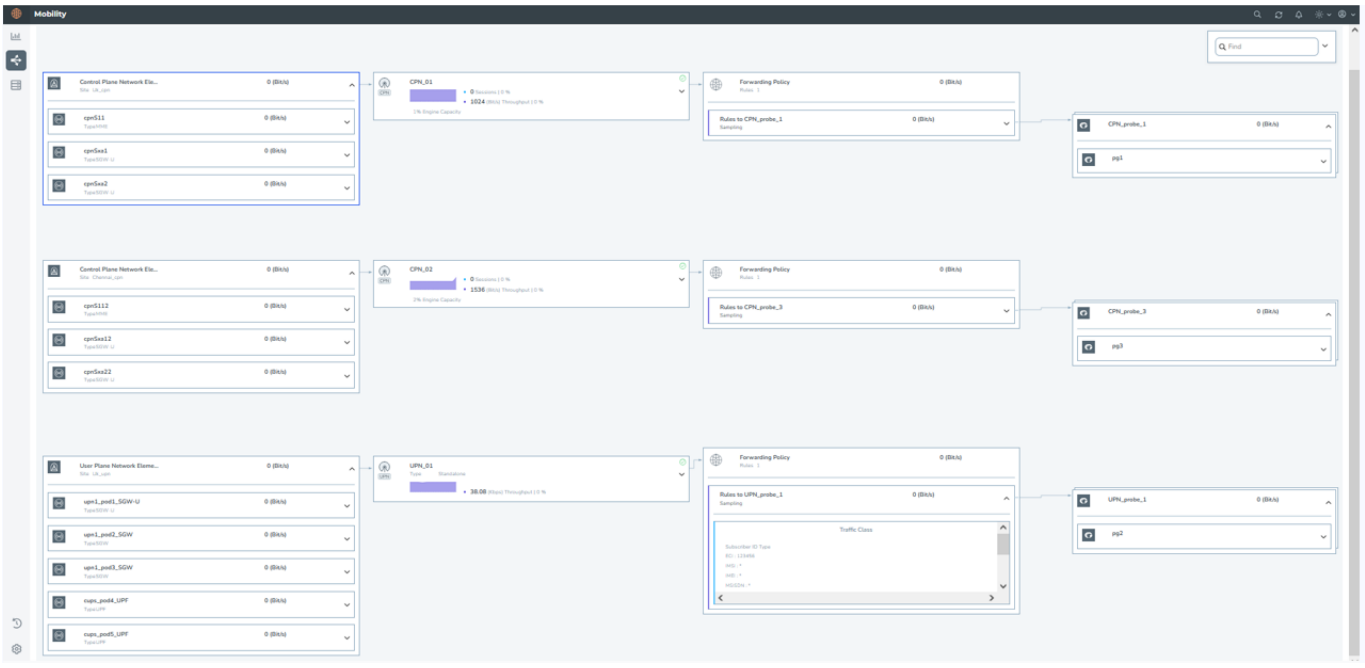
     ```
     /usr/bin/ansible-playbook -e '@~/cupsSolution/ansible_inputs.json' -i
     ~/cupsSolution/cups_inventory /usr/local/share/gigamon-ansible/playbooks/cups/deploy_
     cups.yml
     ```

View CPN-UPN Communication Solution in GigaVUE-FM

After deploying the CPN-UPN Communication in Ansible, you can view how the communication is established between CPN and UPN in GigaVUE-FM. Follow the steps given below:

1. In GigaVUE-FM, go to **Traffic > Physical > Orchestrated Flows**.
2. Click **Mobility**.
3. Click the **Alias**.

The following image shows the CPN-UPN communication solution:



After configuring CPN-UPN Communication solution, you can view the analytics and statistics in the CUPS Communication dashboard. Refer to View CUPS Communication Dashboard for more detailed information.

Remove CPN-UPN Communication

To remove the CPN-UPN Communication solution, follow these steps:

- Execute the playbook and use the following command:

```
/usr/local/bin/ansible-playbook -e '@~/<Path to Inventory directory>/ansible_
inputs.json' -i ~/<Path to Inventory directory>/mobility_inventory /<Path to Ansible
playbook>/delete_mobility_solution.yml
```

> **NOTE:** The multiple YML files created inside the **host_vars** are concatenated, converted into JSON format and sent to GigaVUE-FM.

Map Rules for CPN-UPN Communication

The following tables explain the behavior of PFCP and GTP-u packets on UPN node when Flow sample and Forward list maps are configured with different combination of attributes.

Flow Sample Map Rules for CPN-UPN Communication

Rules and Notes:

- For PFCP packets, RAN and Network Slice based flow sampling are not applicable as the PFCP packets do not contain any RAN or Network Slice parameters.
- For GTP-u packets, if RAN and Network Slice enrichment are not performed, the packets will trigger a RAN PULL to get RAN information from the CPN. During this process (~3 seconds), GTP-u packets will be evaluated against any configured rules. If no match is found, they will either be dropped or sent to the collector. After this (~3 seconds), RAN or Network Slice flow sampling will apply if RAN enrichment is successful. If enrichment fails, GTP-u packets will again be evaluated against matching rules, dropped, or sent to the collector.
- Below are the list of generic RAN, non RAN, and Network Slice parameters that can be configured for CPN-UPN communication:
  - **Generic filter parameters** - apn/imsi/imei/msisdn/version/interface
  - **RAN filter parameters** - eci/plmn-id/tac/tac-5g/nci
  - **Network Slice parameters** - nssid [SST or SST.SD] SD is optional

Generic Filter Parameters	RAN Filter Parameters	Network Slice Parameters	Flow sampling Status - PFCP Packets	Flow sampling Status - GTP-u packets
Yes	No	No	• PFCP packets matching the **Generic Filter** parameters configured are only flow sampled • RAN flow sampling is not applicable	• GTP-u packets matching the **Generic Filter** parameters configured are only flow sampled • RAN and Network Slice based flow sampling is not considered
No	Yes	No	• All PFCP packets are flow sampled • RAN flow sampling is not applicable	• GTP-u packets matching the **RAN Filter** parameters configured are only flow sampled • Network Slice based flow sampling is not considered
Yes	Yes	No	• PFCP packets matching the **Generic Filter** parameters configured are only flow sampled • RAN flow sampling is not applicable	• GTP-u packets matching the **RAN Filter and Generic Filter** parameters configured are only flow sampled • Network Slice based flow sampling is not considered
No	No	Yes	• All PFCP packets are flow sampled • RAN and Network Slice flow sampling is not applicable	• GTP-u packets matching the **Network Slice** parameters are only flow sampled • RAN based flow sampling is not considered

Generic Filter Parameters	RAN Filter Parameters	Network Slice Parameters	Flow sampling Status - PFCP Packets	Flow sampling Status - GTP-u packets
No	Yes	Yes	• All PFCP packets are flow sampled • RAN and Network Slice flow sampling is not applicable	• GTP-u packets matching the **RAN and Network Slice** parameters are only flow sampled
Yes	No	Yes	• PFCP packets matching the **Generic Filter** parameters configured are only flow sampled • RAN and Network Slice flow sampling is not applicable	• GTP-u packets matching the **Generic Filter and Network Slice** parameters are only flow sampled • RAN based flow sampling is not considered
Yes	Yes	Yes	• PFCP packets matching the **Generic Filter** parameters configured are only flow sampled • RAN and Network Slice flow sampling is not applicable	• GTP-u packets matching the **Generic Filter, RAN and Network Slice** parameters are only flow sampled

NOTE: In UPN, when the flow sampling map contains both RAN related and generic rules, the session count in the map statistics increments when both types of rules are matched. This behaviour is expected, as the rule lookup occurs twice: before and after enrichment.
The initial rule lookup takes place when user tunnels are created from the PFCP packet, incrementing the session count for generic matching rules. After enrichment, a RAN update from the CPN triggers a second rule lookup to match RAN rules. For 5G sessions, the session count is based on the QFI per PDU session after enrichment.

NOTE: In 5G CPN, for second-level maps of subtype **flowSample5g**, only percentage 100 is supported. Updating rule priority by using API, GigaVUE-FM, or Ansible is not supported.

Forward List Map Rules for CPN-UPN Communication

Rules and Notes:

- For PFCP packets filtered for matching, RAN forward list is Not Applicable as the PFCP packets do not contain any RAN parameters.
- For GTP-u packets filtered for matching, if RAN enrichment is not performed, the packets will trigger a RAN PULL to get RAN information from the CPN. During this process (~3 seconds), GTP-u packets will be evaluated against any configured rules. If no match is found, they will either be dropped or sent to the collector. After this (~3 seconds), RAN filtering will work if RAN enrichment is successful. If enrichment fails, GTP-u packets will again be evaluated against matching rules, dropped, or sent to the collector.

Database Contains	Forward list Type	Forward list Rule - APN	Forward list Rule - Interface	PFCP Packets Filtered for Matching	GTP-U Packets Filtered Matching
IMSI + RAN	IMSI	Yes	Yes	**IMSI + APN + Interface** • PFCP packets matching IMSI, APN, and Interface are only forward listed • RAN forward list is not applicable	**IMSI + APN + Interface** • GTP-u packets matching IMSI, APN, and Interface are only forward listed • RAN forward list is not considered
IMSI + RAN	IMSI	No	Yes	**IMSI + Interface** • PFCP packets matching IMSI and Interface are only forward listed • RAN forward list is not applicable	**IMSI + Interface** • GTP-u packets matching IMSI and Interface are only forward listed • RAN forward list is not considered
IMSI + RAN	IMSI	Yes	No	**IMSI + APN** • PFCP packets matching IMSI and APN are only forward listed • RAN forward list is not applicable	**IMSI + APN** • GTP-u packets matching IMSI and APN are only forward listed • RAN forward list is not considered
IMSI + RAN	IMSI	No	No	**IMSI**	**IMSI**

Database Contains	Forward list Type	Forward list Rule - APN	Forward list Rule - Interface	PFCP Packets Filtered for Matching	GTP-U Packets Filtered Matching
				• PFCP packets matching IMSI are forward listed • RAN forward list is not applicable	• GTP-u packets matching IMSI are only forward listed • RAN forward list is not considered
IMSI + RAN	RAN	Yes	Yes	**APN + Interface** • PFCP packets matching APN and Interface are only forward listed • RAN forward list is not applicable	**RAN + APN + Interface** • GTP-u packets matching the RAN, APN, and Interface are only forward listed
IMSI + RAN	RAN	No	Yes	**Interface** • PFCP packets matching the Interface are only forward listed • RAN forward list is not applicable	**RAN + Interface** • GTP-u packets matching the RAN and Interface are only forward listed
IMSI + RAN	RAN	Yes	No	**APN** • PFCP packets matching APN are only forward listed • RAN forward list is not applicable	**RAN + APN** • GTP-u packets matching RAN and APN are only forward listed
IMSI + RAN	RAN	No	No	**All PASS** • All PFCP packets are forward listed • RAN forward list is not applicable	**RAN** • GTP-u packets matching the RAN are only forward listed
IMSI + RAN	ALL	Yes	Yes	**IMSI + APN + Interface** • PFCP packets matching IMSI, APN, and Interface are only forward listed • RAN forward list is not applicable	**IMSI + RAN + APN + Interface** • GTP-u packets matching IMSI, RAN, APN, and Interface are only forward listed

Database Contains	Forward list Type	Forward list Rule - APN	Forward list Rule - Interface	PFCP Packets Filtered for Matching	GTP-U Packets Filtered Matching
IMSI + RAN	ALL	No	Yes	**IMSI + Interface** • PFCP packets matching the IMSI and Interface are only forward listed • RAN forward list is not applicable	**IMSI + RAN + Interface** • GTP-u packets matching IMSI, RAN, and Interface are only forward listed
IMSI + RAN	ALL	Yes	No	**IMSI + APN** • PFCP packets matching the IMSI and APN are only forward listed • RAN forward list is not applicable	**IMSI + RAN + APN** • GTP-u packets matching the IMSI, RAN, APN are only forward listed
IMSI + RAN	ALL	No	No	**IMSI** • PFCP packets matching the IMSI are only forward listed • RAN forward list is not applicable	**IMSI + RAN** • GTP-u packets matching IMSI and RAN are only forward listed

View CUPS Communication Dashboard

The CUPS (CPN-UPN) communication dashboard captures the statistics involved when the CPN communicates with the UPN to supplement the UPN with **RAN** and **Network Slice** attributes. Refer to CPN-UPN Communication for Support of RAN and Network Slice Attributes topic for the list of supported **RAN** and **Network Slice** attributes.

To access the dashboards:

1. Go to  -> **Analytics -> Dashboards.**
2. Navigate to **Mobility Dashboards -> CUPS Communication**.
3. The **CUPS Communication** consists of Packet Statistics, GPFCP Statistics, and Enrichment Statistics.
4. Choose the required statistics tab to view the visualizations listed in the tables below.
5. Select the required Site, CPN Alias and the UPN Alias to view the CPN UPN communication statistics of a specific site.
6. When the UPN GigaSMART engine goes down, an alarm is raised and displayed in the **UPN Engine Outage** section of the CUPS Communication dashboard with **Cluster ID**, **Engine ID** of the GigaSMART engine and the **Timestamp** of the engine outage.

- When there are zero values for all the KPIs in the **CUPS Communication** dashboard, it indicates that the UPN GigaSMART engine is down, and you can find the **Cluster ID** and the **Engine ID** of that UPN GigaSMART engine in the **UPN Engine Outage** visualization.
- The same statistics is replicated in all the UPN GigaSMART engines that are part of the GigaSMART engine group. In the **CUPS Communication Dashboard**, GigaVUE-FM displays only the statistics of the first UPN GigaSMART engine of the engine group. When you have zero values for all the KPIs in the **CUPS Communication** dashboard, it indicates that the first UPN GigaSMART engine is down, and you can find the **Cluster ID** and the **EngineID** of that UPN GigaSMART engine in the **UPN Engine Outage** visualization.

Click the  button to download the UPN GigaSMART engine outage details in .csv format.

The following tables list the various CUPS communication statistics visualizations:

Packet Statistics

The following table displays the visualization details of Packet Statistics.

Table 2: Packet Statistics Dashboard

Visualizations	Details
CPN Tx Packets	Displays the total number of packets transmitted from CPN to UPN.
UPN Rx Packets	Displays the total number of packets received at UPN from CPN.
CPN Tx Bits	Displays the total number of bits transmitted from CPN to UPN.
UPN Rx Bits	Displays the total number of bits received at UPN from CPN.
UPN Tx Packets	Displays the total number of packets transmitted from UPN to CPN.
CPN Rx Packets	Displays the total number of packets received at CPN from UPN.
UPN Tx Bits	Displays the total number of bits transmitted from UPN to CPN.
CPN Rx Bits	Displays the total number of bits received at CPN from UPN.
CPN Packets Drop	Displays the total number of packets dropped at CPN.
UPN Packets Drop	Displays the total number of packets dropped at UPN.

Packet Statistics

The following displays the visualization details of GPFCP Statistics.

Table 3: GPFCP Statistics Dashboard

Visualizations	Details
CPN GPFCP Session Lookup Fail	Counter to indicate the number of sessions not found for Enrichment update or Enrichment Pull at the CPN.
UPN GPFCP Session Lookup Fail	Counter to indicate the number of sessions not found for Enrichment update or Enrichment Pull at the UPN.
CPN GPFCP Ack Missing	Displays the GPFCP acknowledgment(s) missing at CPN. If the acknowledgment is missing for an Enrichment update, regenerate Enrichment update is accounted in this counter.
UPN GPFCP Ack Missing	Displays the GPFCP acknowledgment(s) missing at UPN. If the acknowledgment is missing for an Enrichment update, regenerate Enrichment update is accounted in this counter.
CPN GPFCP Tx HB	Displays the GPFCP heartbeats at CPN transferred.
CPN GPFCP Rx HB	Displays the GPFCP heartbeats at CPN received.
UPN GPFCP Rx HB	Displays the GPFCP heartbeats at UPN received.
UPN GPFCP Tx HB	Displays the GPFCP heartbeats at UPN transferred.

Enrichment Statistics

The following table displays the visualization details of Enrichment Statistics.

Table 4: Enrichment Statistics Dashboard

Visualizations	Details
CPN Tx Enrichment Update	Displays the number of Enrichment update messages transmitted from CPN to UPN.
UPN Rx Enrichment Update	Displays the number of Enrichment update messages received at UPN from CPN.
UPN Tx Enrichment Response Accept	Displays the number of acknowledgment messages transmitted from the UPN to the CPN for an Enrichment update stating that the update is accepted.
CPN RX Enrichment Response Accept	Displays the number of acknowledgment messages received at CPN from UPN for a n Enrichment update stating the update is accepted.
UPN Tx Enrichment Response Reject	Displays the number of acknowledgment messages transmitted from the UPN to the CPN for an Enrichment update stating that the update is rejected.
CPN Rx Enrichment Response Reject	Displays the number of acknowledgment messages received at CPN from the UPN for an Enrichment update stating that the update is rejected.

Visualizations	Details
UPN Tx Enrichment Pull Request	Displays the number of Enrichment Pull Request messages transmitted from UPN (UPN tries to get RAN and Network Slice attributes from the CPN).
CPN Rx Enrichment Pull Request	Displays the number of Enrichment Pull Request messages received at CPN (CPN receives the pull request from the UPN for the RAN and Network Slice attributes).
UPN Tx Enrichment Pull Request Retry	Displays the number of Enrichment Pull Request Retry messages transmitted from UPN. **Note**: When CPN does not respond to the Enrichment Pull Request from UPN, it sends the UPN Tx Enrichment Pull Request Retry message. The maximum number of retries that can be attempted is three.
CPN Enrichment Retry	Displays the number of Enrichment Retry messages transmitted from CPN (CPN retries to enrich the UPN for a REJECTED acknowledgment from UPN).

Upgrade Standalone UPN to CPN-UPN Communication Solution

To upgrade your existing Standalone UPN solution to CPN-UPN Communication solution, follow the steps given below:

Step No	Task	Refer the following topics
1.	Configure CPN **gsgroup**	gsgroups
2.	Configure **gpfcp** profiles for both CPN and UPN	gpfcp profile
3.	Configure **IP-Interface** for CPN-UPN communication	ip interface
4.	Add **exporters** to both CPN and UPN	apps exporter
5.	Add the configured **gpfcp** profiles to **gsparams**	gsparams
6.	Configure **gpfcp delay-interval** for CPN	gpfcp profile
7	Add **gsop** for the CPN	gsop
8	Add **vport** for the CPN	vport
9	Configure first level maps for the CPN	map

Step No	Task	Refer the following topics
10	Add first level maps for sending gpfcp messages to CPN and UPN engines	map
11	Configure second level maps for CPN	map
12	Configure or edit second level maps for UPN	map

Rollback from CPN-UPN Communication Solution to Standalone UPN

This section provides the steps to rollback from CPN-UPN Communication Solution to Standalone UPN and also clean up the configurations performed for the CPN-UPN Communication solution. If the CPN-UPN Communication fails due to some issues, you can roll back to the previous Standalone UPN Solution.

Follow the steps given below to roll back to Standalone UPN:

Step No	Task	Refer the following topics
1.	Remove CPN specific second level maps	map
2.	Remove **vports** configured for CPN	vport
3.	Delete the first level maps configured for **gpfcp** messages	map
4.	Delete other CPN specific first level maps	map
5.	Delete **gsops** attached to the CPN engine	gsop
6.	Delete **IP interfaces** for CPN UPN communication	ip interface
7	Delete the CPN **gsgroup**	gsgroups
8	Remove **gpfcp** profile from UPN **gsparams**	gsparams
9	Remove the **exporter** configurations for both the CPN and UPN	apps exporter
10	Delete all the configured **gpfcp** profiles	gpfcp profile
11	Remove **RAN** and **Network Slicing** rules from UPN specific second level maps	map rule

Quick Rollback from CPN-UPN Communication Solution to Standalone UPN

This section provides the steps to rollback from CPN-UPN Communication Solution to Standalone UPN without deleting the configurations created for CPN-UPN Communication Solution. The quick rollback will be handy while debugging issues in the CPN-UPN Communication Solution.

Follow the steps given below to perform quick roll back to Standalone UPN Solution:

1. Remove **gpfcp** profile from UPN **gsparams**. Refer to gsparams for more detailed instructions.
2. Remove **gpfcp** profile from CPN **gsparams**. Refer to gsparams for more detailed instructions.

Monitoring of Subscriber Intelligence Solutions

You can monitor the following solution in GigaVUE-FM after configuring the solution through Ansible:

- Monitoring CUPS Solution

Monitoring CUPS Solution

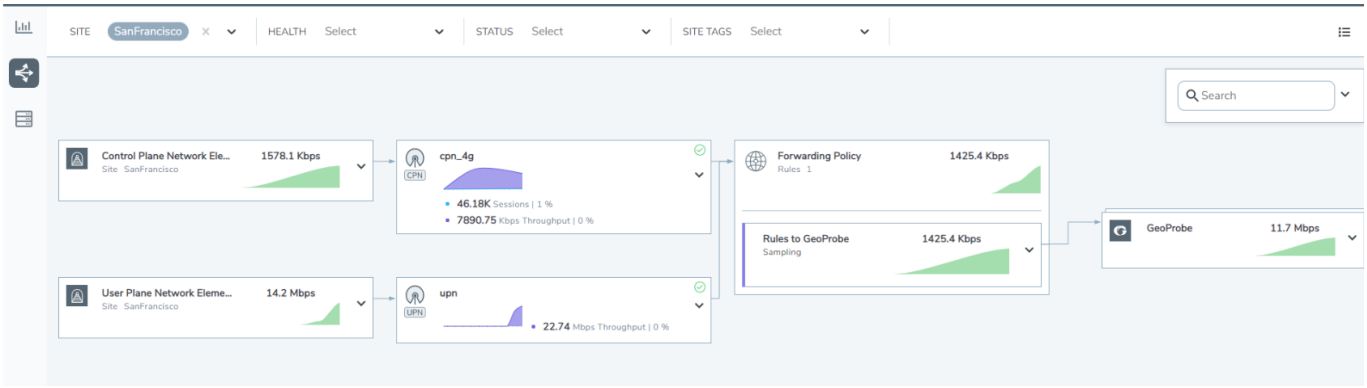
You can monitor the CUPS solution in GigaVUE-FM after configuring the solution through Ansible. You can monitor the following information from GigaVUE-FM:

- Interface
- CPN-UPN
- Forwarding Policies
- Tool

To view the CUPS solution, follow these steps:

1. Go to **Traffic** .
2. Go to **Physical** > **Orchestrate** > **Mobility**.
3. Select a Site from the **SITE** drop-down list. If required, you can select multiple sites.

The screen appears as shown in the figure:



GigaVUE-FM also helps you to view the statistics and health information of the following:

Components	Statistics and Health Information for...
CPN-UPN function	First level fabric maps
CPN-UPN node	GigaSMART Group, SFFP Profile, Exporter, Listener, and IP interface.
Forwarding Policies	GSOPS and Second level Map.

To view the statistics and health information, follow these steps:

1. Go to **Physical** > **Orchestrate** > **CUPS**.
2. Select a **Site** from the SITE drop-down list.
3. Select a component from the GUI.
4. Click **Troubleshoot** to view the statistics and health information.

Health Status of Solution

Health Status of a Solution depends on the health status of the CPN and UPN configured in the solution.

The health status of a solution is indicated by the following colors:

Color	Health Status of Solution CPN/UPN Node State
Green	Healthy -All the CPNs and UPNs health node status are green
Amber	Partially Healthy- None of the CPNs/UPNs nodes is red or at least a

Color	Health Status of Solution CPN/UPN Node State
	node is not functioning properly.
Red	Unhealthy- At least one None of the CPNs/UPNs node is not functioning properly.

CPN Health State

A single CPN's health is based on the health of following components configured in the CPN:

- Source Ports for CPN
- First Level Maps
- GigaSMART (vport, gsgroup and GSOP)
- Second Level Maps
- Destination Tools Ports
- CPN IP Interface

NOTE: Only active CPNs are considered in health computation.

UPN Health State

A single UPN's health is based on the health of following components configured in the UPN:

- Source Ports for UPN
- First Level Maps
- GigaSMART (vport, Gsgroup and gsop)
- Second Level Maps
- DestinationTools Ports
- UPN IpInterface

NOTE: Only Active UPNs are considered in health computation.

The following table explains the health status of a CPN/UPN in GigaVUE-FM based on the health status of the components configured in the CPN/UPN:

Color	UPN Health State
Green	Healthy - All the above components of the CPN/UPN are successfully configured.
Amber	Partially Healthy-None of the components or any of the CPN/UPN component is not successfully configured.
Red	Unhealthy- At least one or any of the configuration status of the CPN/UPN component is not successfully configured.

Monitor Session and Tunnel Utilization

This section describes how to monitor Session and Tunnel utilization of GigaSMART engines configured in a Subscriber Intelligence Solution.

Configure SNMP Traps for Subscriber Intelligence Solution

Real-time alerts for key resource and performance indicators within GTP engines can be accessed, such as:

- **Engine Approaching Tunnel/Session Limits** : Alerts are generated when tunnel or session usage nears capacity
- **GTP Persistence Restore Failure**: Notifications are sent if the persistence file is stale or corrupted.
- **Low Mempool Packet Buffers** : Alerts indicate when packet buffer availability is running low.
- **High CPU Utilization**: Notifications are triggered as CPU usage approaches a user-defined threshold.

These alerts enhance visibility into system health, allowing for more efficient troubleshooting and operational oversight.

These alerts are generated once the said resource parameter is enabled in the GigaSMART Group. To enable these configuration follow the below steps:

- From the left navigation pane, go to **Inventory > Physical > Nodes**.

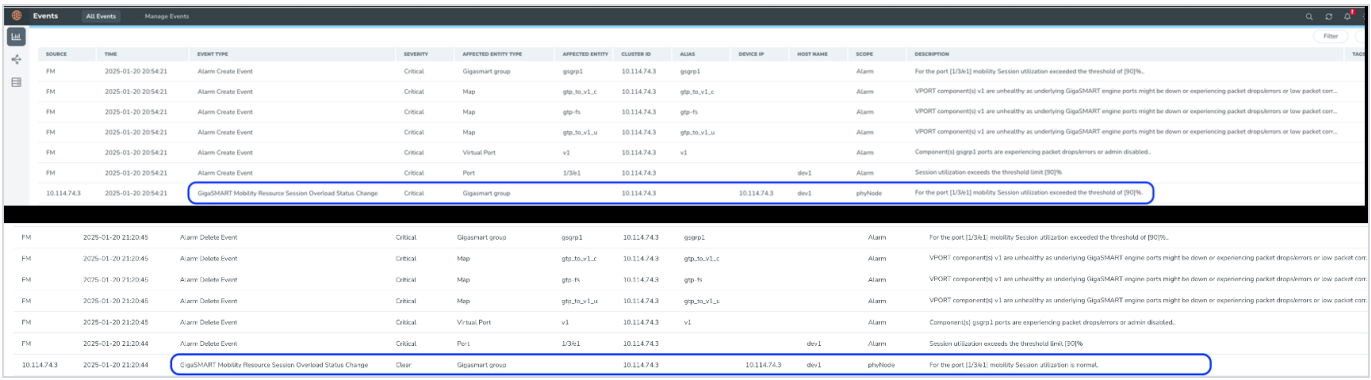
- Select the node, and then go to **GigaSMART > GigaSMART Groups> GigaSMART Parameters**.
- The GigaSMART Group parameters that need to be enabled are as follows:
  - **Enable Resource Packet Buffer** - This allows configuration of Packet Buffer Usage Threshold values. A clear alert is triggered when the packet buffer usage falls 20% below the high threshold value.
  - **Enable Tunnel Usage Threshold**- .This enables the configuration of tunnel usage thresholds. After enabling, configure the threshold values in the **Maximum Tunnel Usage Threshold** field within the range of 50-90%. The default value is 90%. If tunnel usage exceeds this configured value, an SNMP trap will be generated. A clear trap will be sent when the tunnel utilization drops to 10% below the high threshold value.
  - **Enable Session Usage Threshold**- This allows configuration of session usage thresholds. After enabling, configure the threshold values in the **Maximum Session Usage Threshold** field within the range of 50-90%. The default value is 90%. If session usage exceeds this configured value, an SNMP trap will be generated. A clear trap will be sent when session utilization drops to 10% below the high threshold value.
  - **GTP Persistence**- This allows configuration of GTP Persistence timer values. Refer to GTP Stateful Session Recovery to know more.
  - **Enable Resource CPU Utilization** -This allows the configuration of CPU utilization threshold values. A clear alert will be triggered when the CPU utilization drops below two-thirds of the high percentage value.

Additionally, these threshold values can also be configured through the GigaVUE-OS CLI refer to snmp-server and gsparams in GigaVUE-OS CLI Reference Guide

Once these values are configured in the GigaSMART Group, it is possible to enable the required event for SNMP notifications either on the device or on the GigaVUE-FM instance managing the device. For instructions about how to enable SNMP notification on a device, refer to Enable or Disable Events for SNMP Notifications in GigaVUE Administration Guide.

To view the SNMP trap ,its corresponding Object Identifier and Events refer to Mapping of SNMP Traps with GigaVUE-FM Events and Alarms in the GigaVUE Administration Guide.

The threshold values for sessions can also be visualized under the Analytics tab Refer to Analytics



Display Flow Ops Reports

GigaSMART provides support for Flow Ops reporting. The Flow Ops reports display session table statistics for the various features. Refer to GigaSMART Group Statistics Definitions for descriptions of these statistics.

To display the Flow Ops report perform the following steps:

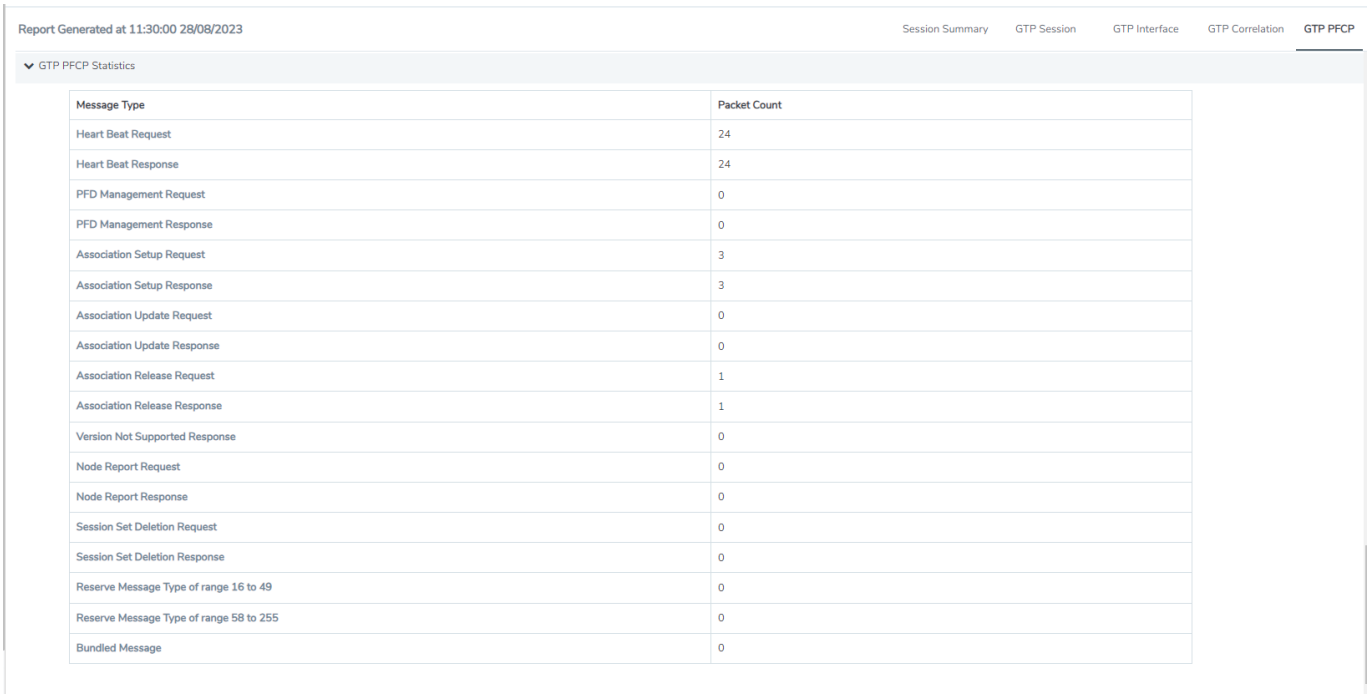
1. From the device view, select **GigaSMART > GigaSMART Groups > Report**.
2. Under **Report Info**, from the **Type** drop down list, select any one of the following report types to display the respective reports:
   - **Flow Filtering**- Flow Filtering reports.
   - **Flow Sampling** - Flow Sampling reports.
   - **Flow SIP** - Flow SIP reports.
   - **SSL Decryption** - SSL Decryption reports.
   - **Port Throttling** Port Session and Port Throttle reports.

     NOTE: For information about the fields that get displayed for each of the report type, refer to "Flow Ops Report - Field Reference"
3. Select a GigaSMART Group - for example: **grp1**.
4. Click **Generate**.

The corresponding Flow Ops report appears with the time stamp of when the report was generated. The field labels and the values selected (or entered) for generating the report appear as sticky info notes in the header pane. The following buttons are available in the header pane:

- **Generate New**: Use to generate a new report and to toggle between the report generated and the Reports page.
- **Export**: Use to export the report.
- **Upload:** Use to upload the report to the archive server.

NOTE: The IP based Flow Sampling reports cannot be exported or uploaded. To upload IP based flow sample report use GigaVUE-OS CLI commands:
show gsgroup flow-ops-report alias <> type flow-sampling any upload
show gsgroup flow-ops-report alias <> type flow-sampling device-ip-mask <netmask> upload
show gsgroup flow-ops-report alias <> type flow-sampling device-ip6-mask <upload>
Refer to **show**command in GigaVUE-OS CLI Reference Guide, and Definitions of GigaSMART Statistics



Message Type	Packet Count
Heart Beat Request	24
Heart Beat Response	24
PFD Management Request	0
PFD Management Response	0
Association Setup Request	3
Association Setup Response	3
Association Update Request	0
Association Update Response	0
Association Release Request	1
Association Release Response	1
Version Not Supported Response	0
Node Report Request	0
Node Report Response	0
Session Set Deletion Request	0
Session Set Deletion Response	0
Reserve Message Type of range 16 to 49	0
Reserve Message Type of range 58 to 255	0
Bundled Message	0

Flow Ops Report - Field Reference

Report Type	Fields	Description	Notes

Flow Filtering	• **Pattern**
	• **Any**
		• **Pattern**: Pattern for which the Flow Filtering report must be generated. The **Pattern** drop-down lists the following options.
		  • GTP IMSI
		  • GTP IMEI
		  • GTP MSISDN
		  Enter the required value for the selected option. The report is generated accordingly.
		• **Any-**Select **Any** for any of the pattens to be selected**.**
		  Note: The patterns and sub-pattern options are specific to device versions and may change accordingly.
			The Flow Filtering report appears as a single page with links to the various sections of the report. Click on the links to navigate to the following sections:
			• Session Summary
			• GTP Session Statistics
			• GTP Interface Statistics
			• GTP Correlation Statistics
			• GTP PFCP Statistics
			The report generation time is displayed on top of the report.
			Refer to the following sections for details:
			• Export Flow Filtering Reports
			• Generate Delta Reports
			Note: In the Hand Over scenario, the flow-ops report does not reflect the entire gtp-u packet count, since the older tunnels are marked for free before the sync mechanism starts between the leader and non-leader

Report Type	Fields	Description	Notes

		• **Duration**: Time at which the Flow Filtering report needs to be generated. The following options are available:
		  • **Now**: Report with current statistical data is generated. This is the default option.
		  • **Delta**: Report with statistical data for a specified time interval starting from the current time is generated. Select the required duration:
		  • 5 minutes
		  • 15 minutes
		  • 30 minutes
		  • 1 hour
		  • 4 hours
		  • 12 hours
		  • 1 day
		  • 2 days
		  • 7 days
			engines. The packets that are sent to the older tunnels are not taken into account in the flow-ops report.
			Note: When RAN attributes are not exchanged from CPN to UPN, the RAN metrics, such as ECI, TAC, and NSI_SST are displayed as zero in the flow-ops report. No values are displayed for other RAN metrics such as TMCC, TMNC, EMCC, EMNC, and NSI_SD.

Flow Sampling	• **Device IP Address/Mask** • **Any**	• **Device IP Address/Mask** - Device IP address/Mask. • **Any**- Select **Any** for any available device IP address. **Note**: Graphs will not be generated for Device IP Address /Mask.	Use the scroll bars and the pagination option at the bottom of the page to navigate to the various pages.

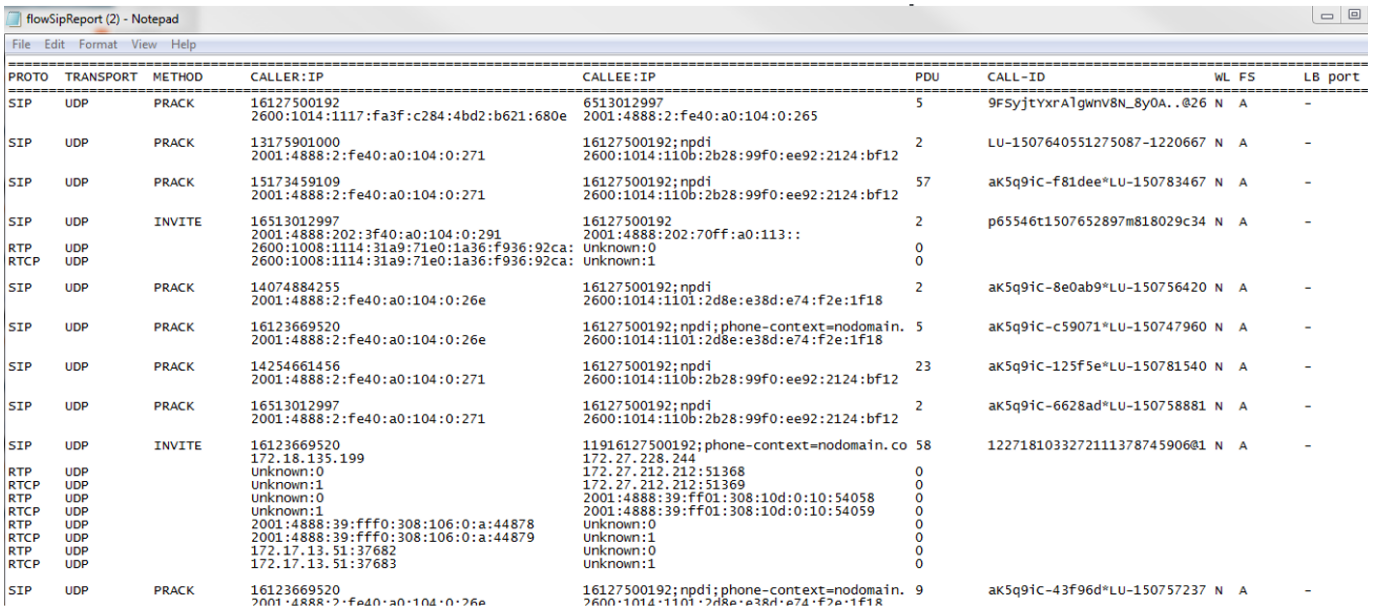
Report Type	Fields	Description	Notes
Flow SIP	• **Caller ID Pattern** • **Any**	• **Caller ID Pattern** - Caller ID pattern. • **Any**- Select **Any** for any available device IP address.	
SSL Decryption	• **Device IP Address/Mask** • **Any**	• **Device IP Address/Mask** - Device IP address/Mask. • **Any**- Select **Any** for any available device IP address.	

Export Flow SIP Session Reports

Support for exporting flow-ops session report is available for Flow SIP. To export Flow SIP session reports:

1. **From the device view,** select**GigaSMART > GigaSMART Groups > Report**.
2. Generate a Flow SIP report.
3. Click **Export** to download the Flow SIP report you just generated. A text file of the Flow SIP report is saved to your local directory.

> **NOTE:** The session table displays the first 1000 sessions only.



flowSipReport (2) - Notepad

File Edit Format View Help

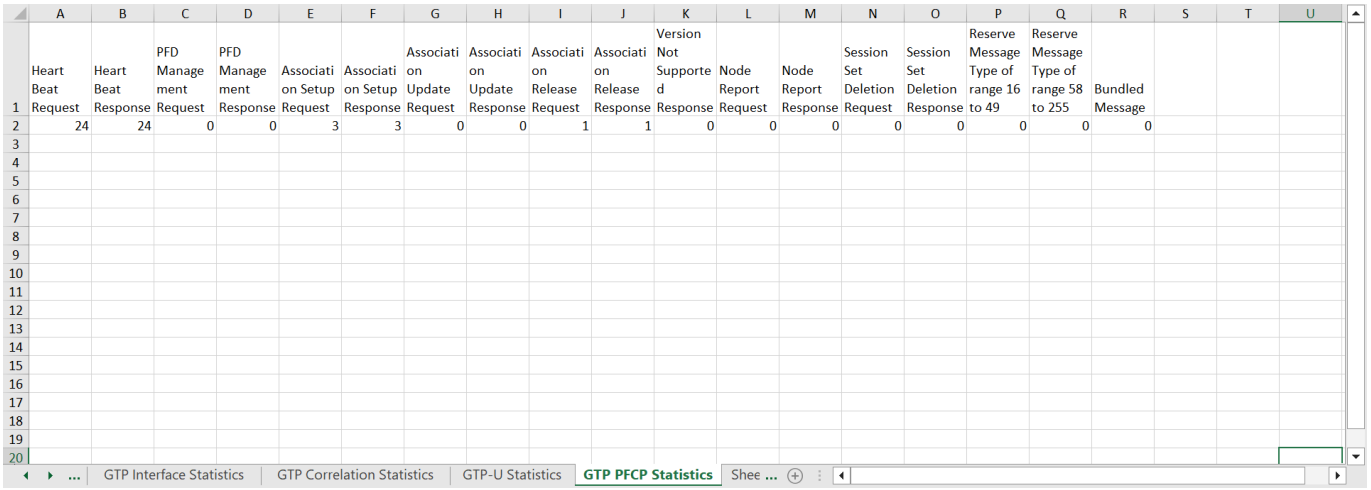
PROTO	TRANSPORT	METHOD	CALLER:IP	CALLEE:IP	PDU	CALL-ID	WL	FS	LB port
SIP	UDP	PRACK	16127500192 2600:1014:1117:fa3f:c284:4bd2:b621:680e	6513012997 2001:4888:2:fe40:a0:104:0:265	5	9FSyjtYxrAlgwnV8N_8yOA..@26	N	A	-
SIP	UDP	PRACK	13175901000 2001:4888:2:fe40:a0:104:0:271	16127500192;npdi 2600:1014:110b:2b28:99f0:ee92:2124:bf12	2	LU-1507640551275087-1220667	N	A	-
SIP	UDP	PRACK	15173459109 2001:4888:2:fe40:a0:104:0:271	16127500192;npdi 2600:1014:110b:2b28:99f0:ee92:2124:bf12	57	aK5q9iC-f81dee*LU-150783467	N	A	-
SIP	UDP	INVITE	16513012997 2001:4888:202:3f40:a0:104:0:291	16127500192 2001:4888:202:70ff:a0:113::	2	p65546t1507652897m818029c34	N	A	-
RTP	UDP		2600:1008:1114:31a9:71e0:1a36:f936:92ca:	Unknown:0	0				
RTCP	UDP		2600:1008:1114:31a9:71e0:1a36:f936:92ca:	Unknown:1	0				
SIP	UDP	PRACK	14074884255 2001:4888:2:fe40:a0:104:0:26e	16127500192;npdi 2600:1014:1101:2d8e:e38d:e74:f2e:1f18	2	aK5q9iC-8e0ab9*LU-150756420	N	A	-
SIP	UDP	PRACK	16123669520 2001:4888:2:fe40:a0:104:0:26e	16127500192;npdi;phone-context=nodomain. 2600:1014:1101:2d8e:e38d:e74:f2e:1f18	5	aK5q9iC-c59071*LU-150747960	N	A	-
SIP	UDP	PRACK	14254661456 2001:4888:2:fe40:a0:104:0:271	16127500192;npdi 2600:1014:110b:2b28:99f0:ee92:2124:bf12	23	aK5q9iC-125f5e*LU-150781540	N	A	-
SIP	UDP	PRACK	16513012997 2001:4888:2:fe40:a0:104:0:271	16127500192;npdi 2600:1014:110b:2b28:99f0:ee92:2124:bf12	2	aK5q9iC-6628ad*LU-150758881	N	A	-
SIP	UDP	INVITE	16123669520 172.18.135.199	11916127500192;phone-context=nodomain.co 172.27.228.244	58	122718103327211137874590 6@1	N	A	-
RTP	UDP		Unknown:0	172.27.212.212:51368	0				
RTCP	UDP		Unknown:1	172.27.212.212:51369	0				
RTP	UDP		Unknown:0	2001:4888:39:ff01:308:10d:0:10:54058	0				
RTCP	UDP		Unknown:1	2001:4888:39:ff01:308:10d:0:10:54059	0				
RTP	UDP		2001:4888:39:fff0:308:106:0:a:44878	Unknown:0	0				
RTCP	UDP		2001:4888:39:fff0:308:106:0:a:44879	Unknown:1	0				
RTP	UDP		172.17.13.51:37682	Unknown:0	0				
RTCP	UDP		172.17.13.51:37683	Unknown:1	0				
SIP	UDP	PRACK	16123669520 2001:4888:2:fe40:a0:104:0:26e	16127500192;npdi;phone-context=nodomain. 2600:1014:1101:2d8e:e38d:e74:f2e:1f18	9	aK5q9iC-43f96d*LU-150757237	N	A	-

Export Flow Filtering Reports

Support for exporting the flow filtering report is available.

To export Flow Filtering reports:

1. **From the device view,** select**GigaSMART > GigaSMART Groups > Report**.
2. Generate a Flow Filtering Report.
3. Click **Export** to download the Flow Filtering report you just generated. An excel sheet of the Flow Filtering report is saved to your local directory.



	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R
1	Heart Beat Request	Heart Beat Response	PFD Management Request	PFD Management Response	Association Setup Request	Association Setup Response	Association Update Request	Association Update Response	Association Release Request	Association Release Response	Version Not Supported Response	Node Report Request	Node Report Response	Session Set Deletion Request	Session Set Deletion Response	Reserve Message Type of range 16 to 49	Reserve Message Type of range 58 to 255	Bundled Message
2	24	24	0	0	3	3	0	0	1	1	0	0	0	0	0	0	0	0

GTP Interface Statistics | GTP Correlation Statistics | GTP-U Statistics | GTP PFCP Statistics | Shee ...

NOTE: The excel sheet contains separate tabs for each of the sections in the report.

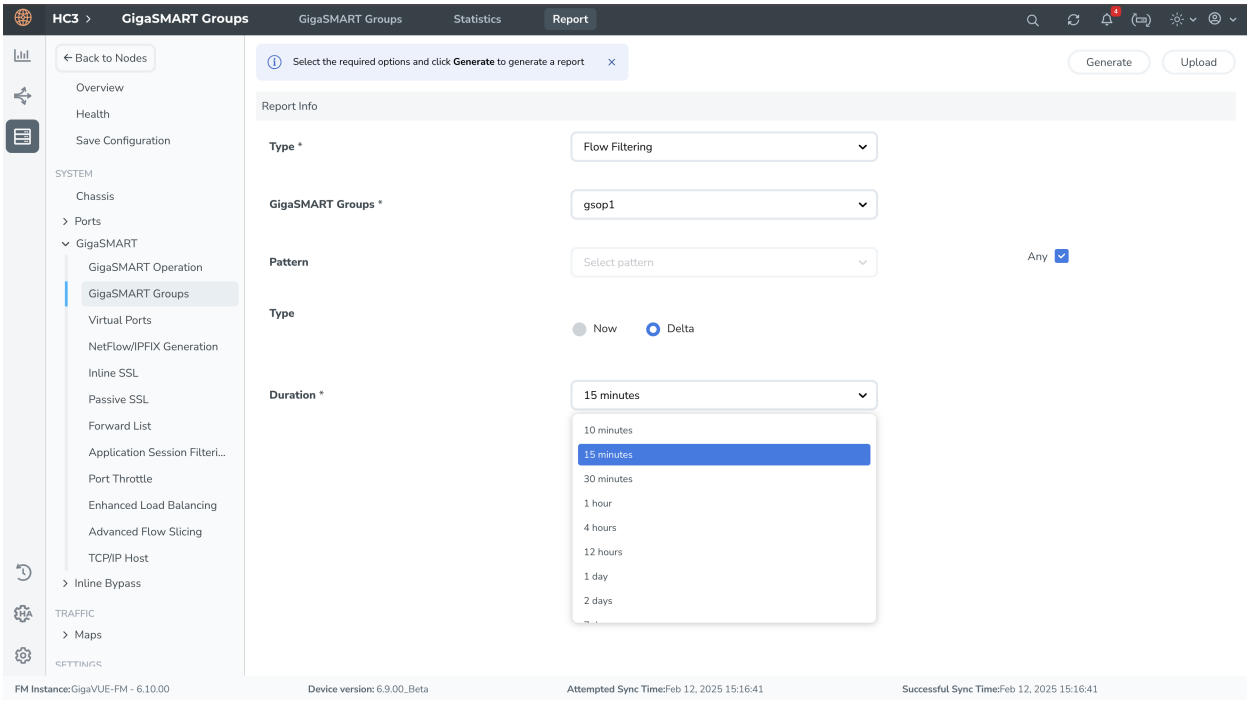
Generate Delta Reports

GigaVUE-FM provides support for generating delta reports that display statistical data for a specific time interval from the current time. Use delta reports to troubleshoot the flow filtering behavior across two different timestamps.

To generate a delta report:

1. From the device view, select **GigaSMART > GigaSMART Groups > Report**.
2. Select or enter the following details:
   - **Pattern**: Select the required pattern or choose any.
   - **Duration**: Select the required duration.

   Refer to the Flow Ops Report - Field Reference section for details.

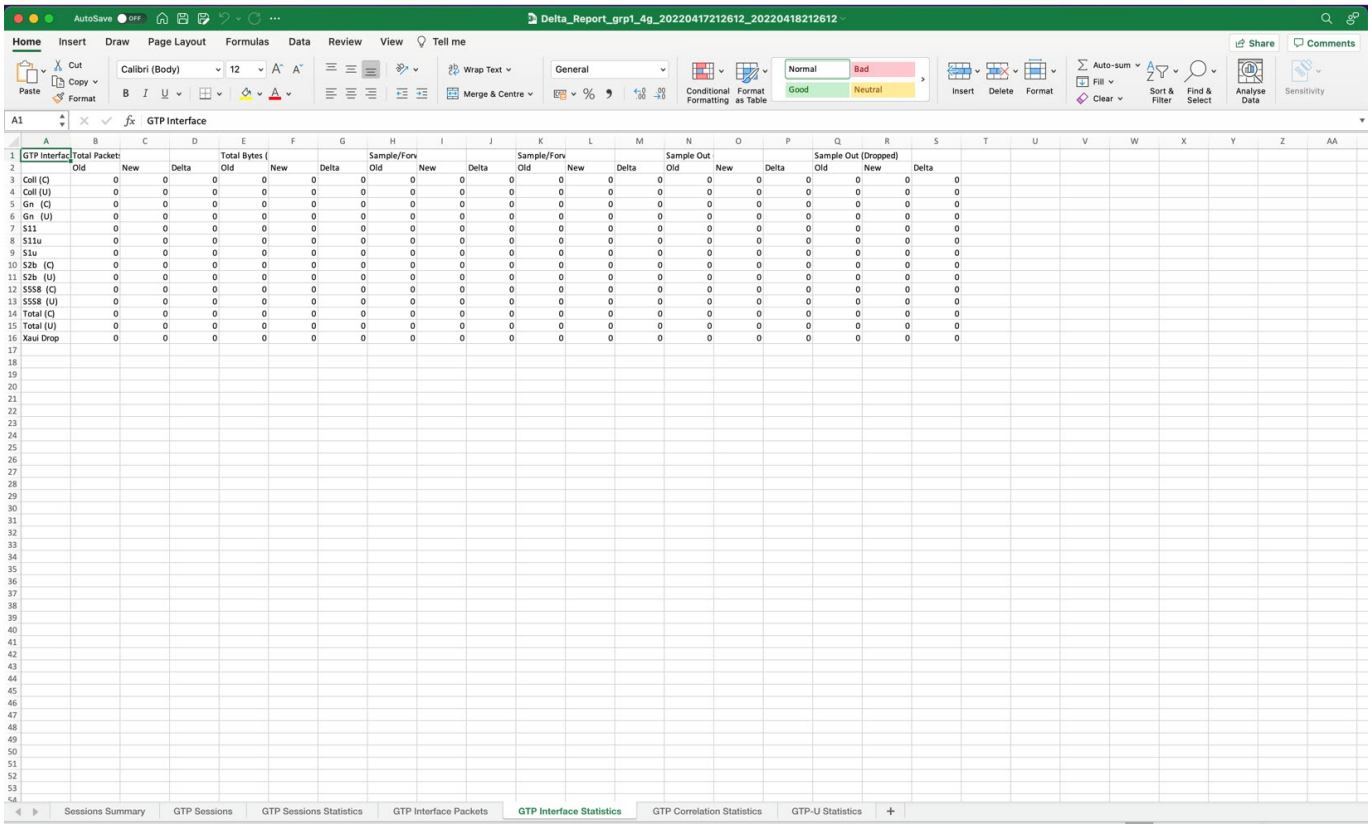


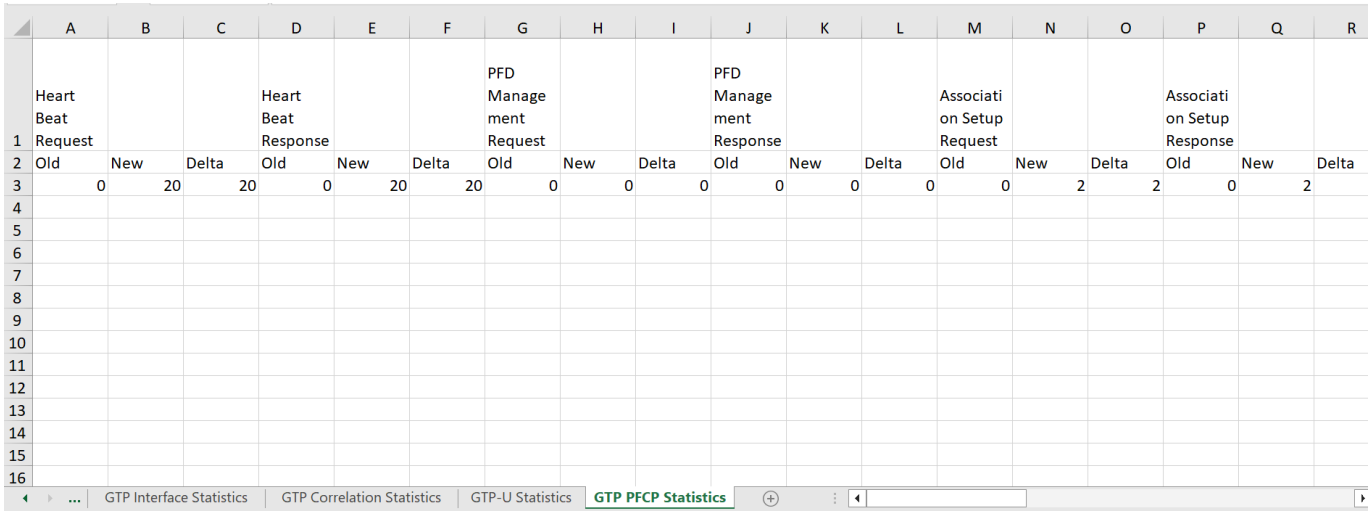
3. Click **Generate**. The delta report for the selected duration is displayed. The following details are displayed as part of the report:
   - Duration of the delta report in the page header
   - Detailed timestamps as part of the report header
   - Within the report, the following details are displayed:
     - **Old:** Earliest record available based on the duration selected. For example, if the duration is selected as 12 hours, the oldest data available for the past 12 hours from the current time is displayed.
     - **New**: Latest record available for the selected duration. The latest record displayed is based on the last stats collection in GigaVUE-FM and sometimes may not be the latest value on the device due to the stats collection interval.
     - **Delta**: Differential data calculated for the selected duration is displayed.

   NOTE: Negative values in the delta columns indicate that a device reset or maintenance has occurred between the selected duration. Refer to the Old and New columns for the statistical rate change.
4. Use the Export button to export the Flow Filtering report in .XLSX format. The format of the filename is as follows:

Delta_Report_<GigaSMART Group name>_<start time in YYYYMMDDHHmmss> _ <end time in YYYYMMDDHHmmss>

The following are the sample Flow Filtering delta reports:





GTP Overlap Flow Sampling Maps

Starting in software version 4.8, GTP overlap flow sampling maps combines GTP forward listing and GTP flow sampling maps into a GTP overlap flow sampling map group, which allows for selected traffic to be sent to multiple destinations simultaneously.

In this scenario, once traffic matches a map, it will be sent to the destination for that map. However, the matched traffic will also be evaluated by subsequent maps and, if a match occurs, it will be sent to each of the destinations pointed to by the subsequent maps.

Figure 26GTP Overlap Flow Sampling Map Groups illustrates regular non-overlap mapping where, once a traffic match is achieved in one map, all other maps are ignored.

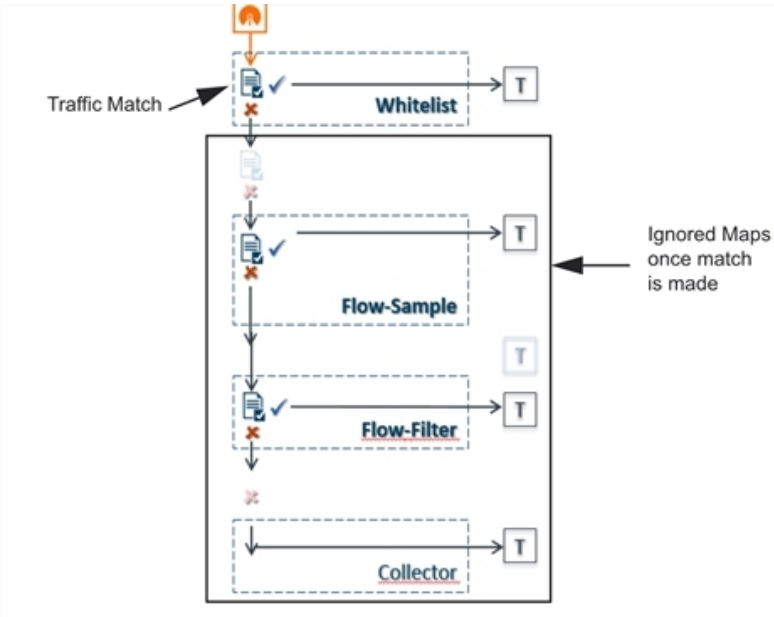


Figure 25 *Non-Overlap GTP Mapping Mode*

This contrasts with GTP overlap flow sampling maps. In Figure 25Non-Overlap GTP Mapping Mode matched traffic is sent to up to six GTP forward listing and flow sampling map pairs that in turn send accepted traffic to up to six load balanced port groups.

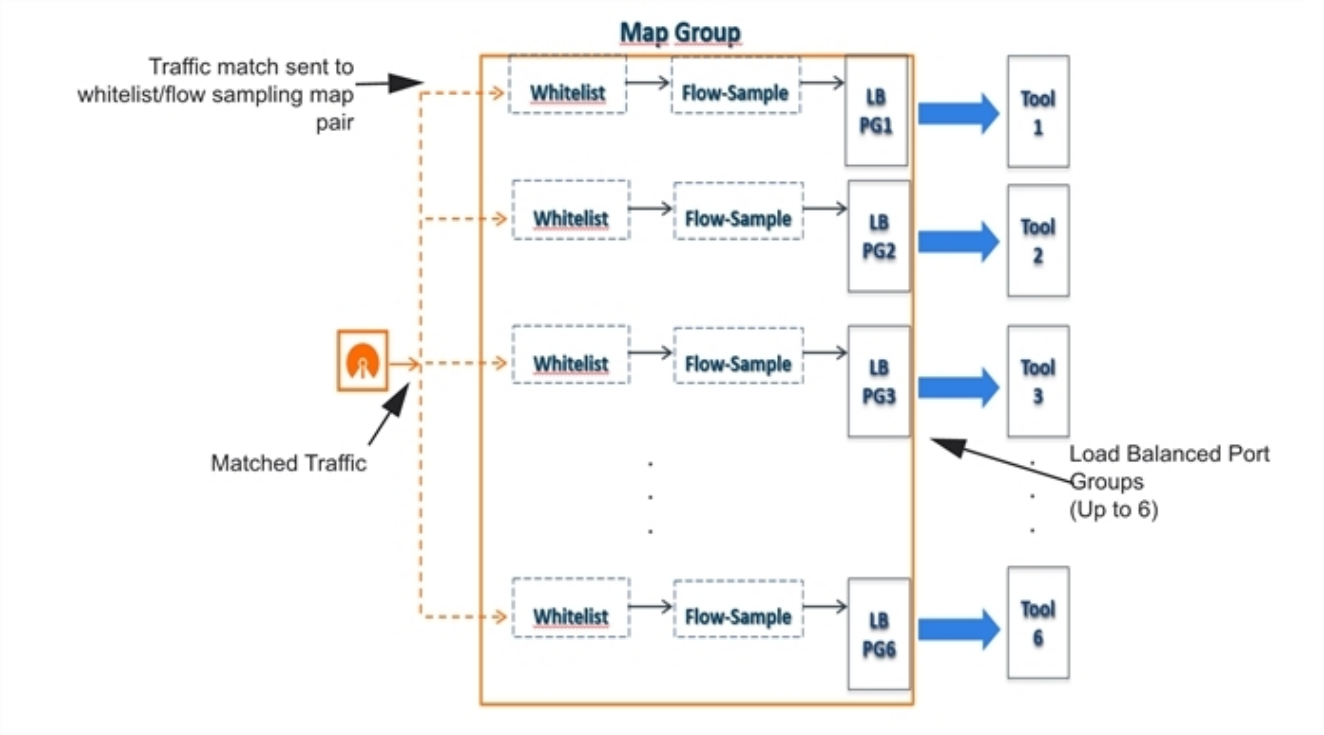


Figure 26 *GTP Overlap Flow Sampling Map Groups*

Configure GTP Overlap Mapping

The configuration of GTP forward listing and GTP flow sampling maps that are part of the GTP overlap flow sampling map group follow the same configuration considerations discussed previously in GigaSMART GTP Whitelisting and GTP Flow Sampling. As is the case with regular non-overlap GTP mapping, GTP forward listing selects specific subscribers based on IMSI, whereas GTP flow sampling uses map rules to select subscribers and then forward a percentage of the packets to tool ports.

Configuration Considerations

This section details certain configuration considerations that apply only to the configuration of GTP forward listing and flow sampling maps for GTP overlap flow sampling maps.

About GTP Overlap Flow Sampling Map Mode and Port Groups

A second level type map specifying GTP overlap flow sampling map mode must be selected to configure GTP forward listing and flow sampling maps.

To configure a GTP whilelisting map in overlap flow sampling map mode, select **Type** as **Second Level** and **Subtype** as **Flow Whitlelist Overlap** in a map.

To configure a GTP flow sampling map in GTP overlap flow sampling map mode, select **Type** as **Second Level** and **Subtype** as **Flow Sample Overlap** in a map .

You can configure one GTP forward listing map and one GTP flow sampling map pair that contain traffic policies corresponding to one destination port group. The load balanced port groups can contain a single port, a port range, or a GigaStream. Note that, starting in software version 4.8, port groups used in GTP overlapping maps support GigaStream.

The maximum number of port groups per single GTP overlap flow sampling map group is six.

For more information about port groups, refer to Port Groups.

Maximum Number of Port Group Members

The Maximum number of port group members can be determined as follows:

Standalone Configuration

In a standalone configuration, use the following sequence :

1. Determine the number of members per port group and add 1.

2. Multiply each port group result by each other.

3. The total multiplication should not exceed 4096.

For instance, assume the following configuration in a GTP overlap mapping group:

o Port Group 1—2 load-balanced GigaStream

o Port Group 2—3 load-balanced GigaStream

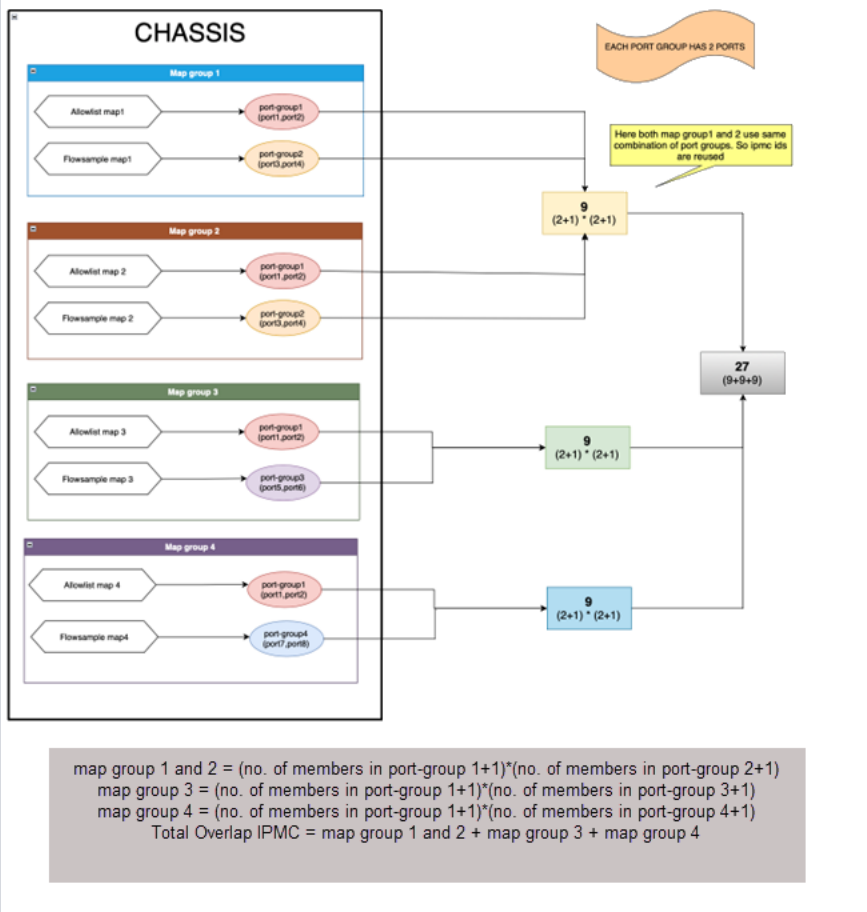
o Port Group 3—1 load-balanced tool port

o Port Group 4—1 load-balanced GigaStream

o Port Group 5—4 load-balanced tool ports. The total number becomes:(2+1)*(3+1)*(1+1)* (1+1)* (4+1) = 240

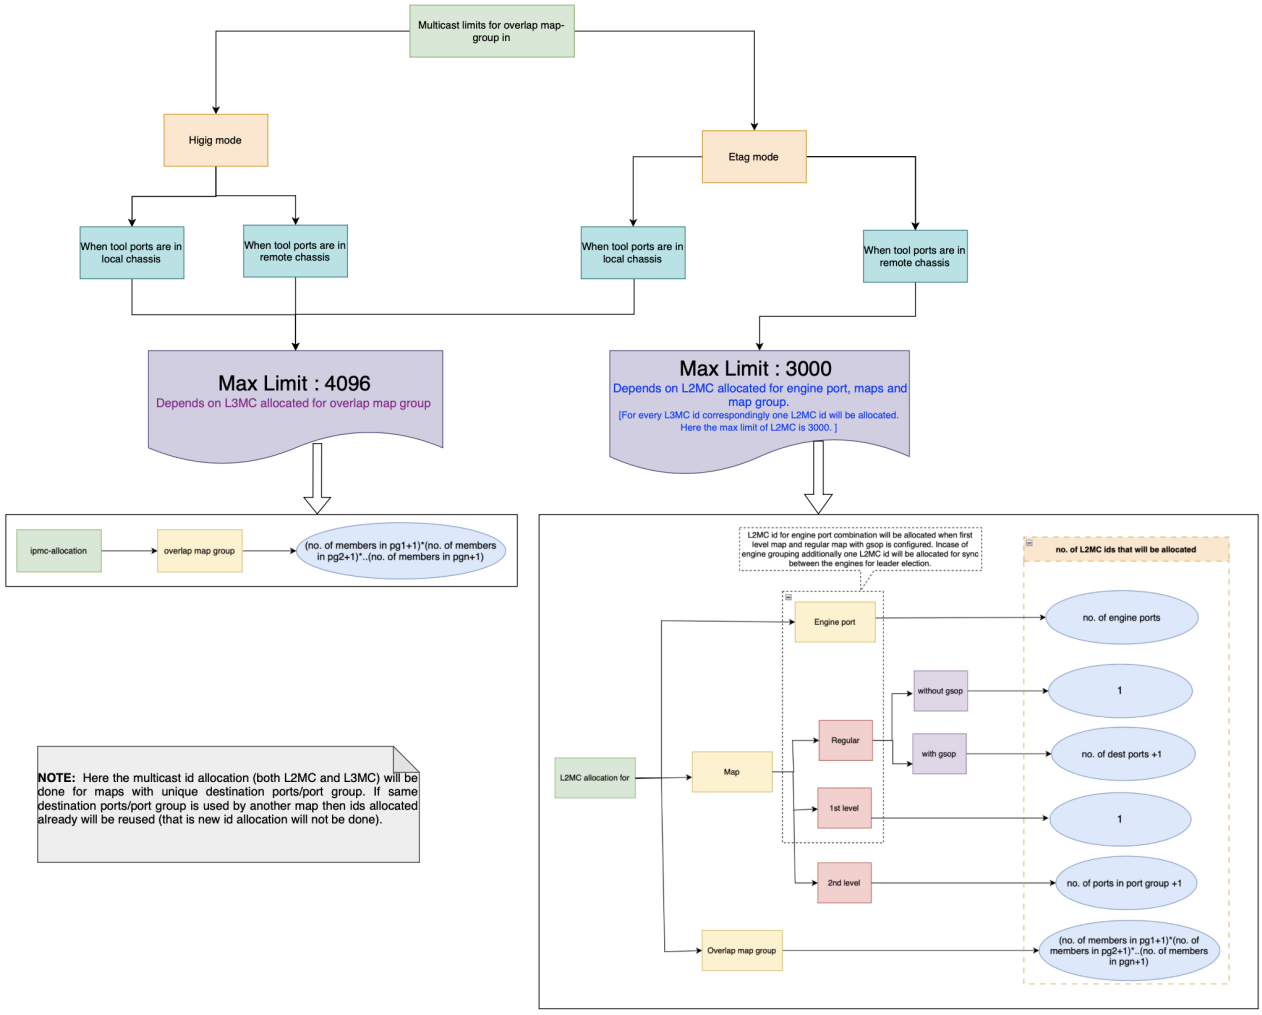
The tool configuration shown is accepted since this does not exceed the maximum number of multicast IDs (4096).

If Port Groups are utilized across multiple map groups in a single cluster, refer to the representation below on how the multicast IDs are calculated.



Cluster Configuration

In a cluster configuration, Multicast IDs can defer based on Hi-Gig mode or E-tag mode. The below is representation of the same.



NOTE: The maximum Layer 2 multicast limit in E-tag cluster mode is 3000. It is recommended to keep the configured multicast entries within this limit. Current resource usage can be checked using command **`show ipmc-resource`**. If you encounter issues due to limit exhaustion, delete device configurations that exceed the limit and upgrade to the latest image or version.
Prior to upgrading to the latest version/image, ensure that the maximum multicast limit is not exceeded.

GTP Overlap Flow Sampling Map Priority

Since a packet matches multiple maps independently the concept of second level map priority does not apply to GTP overlap flow sampling maps. A GTP overlap flow sampling map pair consists of one GTP forward listing map and one GTP flow sampling map having the same destination port group. Within a GTP overlap flow sampling map pair the forward listing map rules will be applied before the flow sampling map rules.

Virtual Port Configuration in GTP Overlap Mode

In GTP Overlap map configuration, the virtual port sending traffic to all the port groups needs to be configured in GTP overlap mode.

To configure the virtual port with GTP overlap mode, select **GTP Overlap** when configuring the virtual port.

About Map Groups

To create a group of maps for GTP forward listing and GTP flow sampling, select **Maps > Maps > Map Groups**, and then click **New**. The maps for a map group are entered in the **Maps** field. Refer to GTP Overlap Flow Sampling Maps. All the maps in a map group receive traffic according to map rules, rather than map priority. Thus, multiple copies of a GTP packet can be sent to more than one tool.

The **Maps** field of the Map Group page groups the forward listing and flow sampling maps. For example, assuming that two forward listing maps (**WLMAP1** and **WLMAP2**) and two flow sampling maps (**FSMAP1** and **FSMAP2**) have been configured in GTP overlap mode, the following example groups them all into the same map group called **map-group1**:

Keep in mind the following configuration considerations for map groups:

- A map group can be associated with only one GigaSMART group (gsgroup).
- All maps within a map group must be connected to the same vport.
- A map group can consist of only one GTP forward listing map or only one GTP flow sampling map but it cannot contains two maps of the same type.
- Once a map group is created, it cannot be edited to change the type or subtype of the map. However, you can add and edit the map rules for a map while it is configured in a map group.
- If multiple map groups are configured, the maps within each map group must point to the same port groups as the other map groups.

For more information about map groups, refer to Create Map Groups.

About Whitelist Maps

The GTP forward list is an IMSI list which is common to all forward list maps. You can configure an optional rule within a forward list map to specify a GTP version or interface-based policy.

Other than specifying a new second level type using **Type Second Level** and **Subtype Flow Whitelist Overlap** when creating the map, the configuration of GTP forward list maps follows the same configuration guidelines as given in the section GTP Whitelisting.

A maximum of six forward list maps sending traffic to six different port groups can be configured per GigaSMART group (gsgroup).

About Flow Sampling Maps

In GTP overlap flow sampling map mode, GTP flow sampling (rule-based flow sampling) is performed after GTP forward list-based forwarding. Therefore, flow sampling maps have a lower priority than forward list maps. Thus, within a GTP overlap map pair that consists of a single GTP forward list overlap map and a GTP flow sampling overlap map, the GTP forward list map is of higher priority.

Within the flow sampling maps, the rules in the first map have a higher priority than the rules in the second, third, and subsequent maps. Within any single map, rules are evaluated in order.

A maximum of six flow sampling maps sending traffic to six different port groups can be configured per GigaSMART group (gsgroup).

Overlap Map Statistics

Starting with version 5.4 overlap and non-overlap maps are available. Overlap maps are displayed based on the following:

- If at least 1 flow-sample map accepts the packets among all the pairs of overlap maps, the Sample (Tx) counters in the GTP Interface stats will be incremented. If more than 1 pair of maps accepts the packets, the Sample (Tx) counters in the GTP Interface stats is incremented only once.

- If at least 1 Forwar list map matches the packets among all the pairs of overlap maps, the Sample (Tx) counters in the GTP Interface stats will be incremented. If more than 1 pair of maps matches the packets, the Sample (Tx) counters in the GTP Interface stats is incremented only once.
- If there are no WL maps and all flow sample maps are no rule match, then Sample(Tx) and Sample Out counters in the GTP Interface stats is not incremented.

GTP Overlap Flow Sampling Maps Example

This section contains:

- Example 1: GTP Overlap Mode

Example 1: GTP Overlap Mode

Example 1 is a GTP overlap mapping mode example.

In Example 1, traffic from a single network port goes to a single first level map (mapLevel1-GTP) which directs GTP-Control, and GTP-User traffic to a virtual port (VP31). Traffic from VP31 is replicated to two GTP whitelisting maps (WLMAP1 and WLMAP2) and two GTP flow sampling maps (FSMAP1 and FSMAP2), which then forward accepted traffic to the final port-group destinations, pg1 and pg2, for load balancing (refer to Figure 27GTP Overlap Mode Example 1).

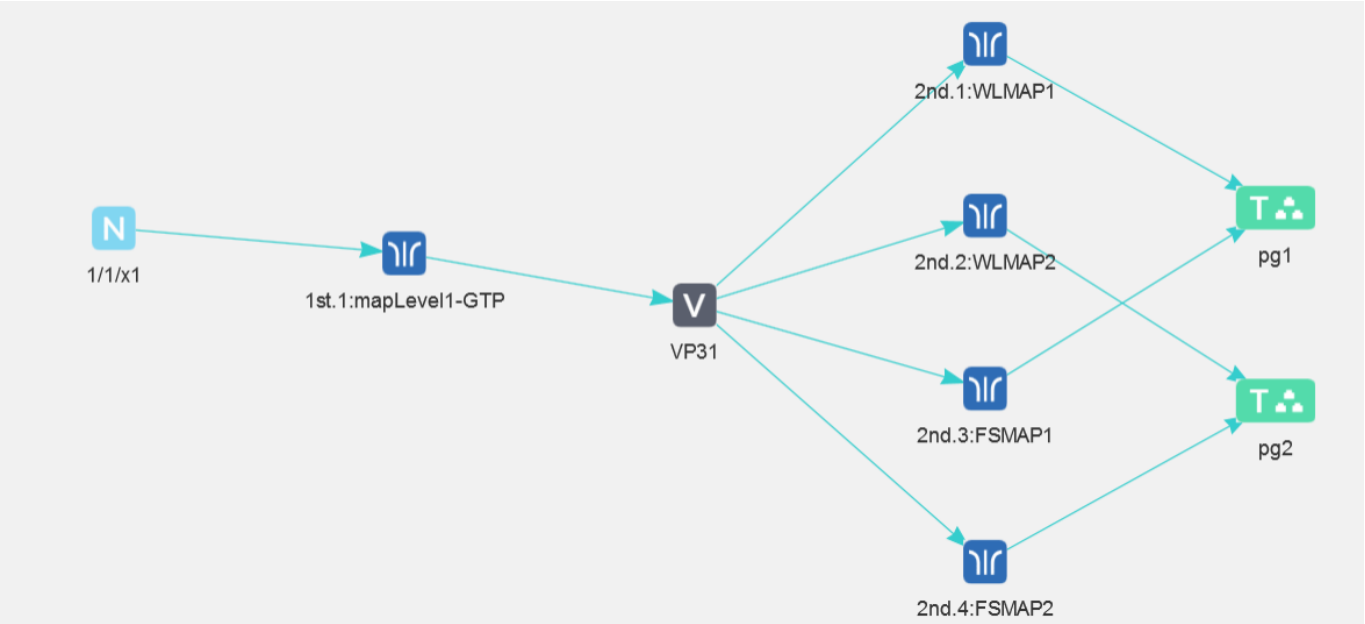


Figure 27 *GTP Overlap Mode Example 1*

NOTE: In Example 1, the tool ports and GigaStreams in the port group are on the same node as the GigaSMART group and GigaSMART operation.

Within each GTP whitelisting and flow sampling pair, if there is not a match to an IMSI in the whitelist map, the traffic flow is sampled based on the rules in the flow sampling map. The flow sampling rules specify IMSI, IMEI, and MSISDN numbers, as well as the percentage to sample.

Within each map pair, packets are then accepted or rejected. Accepted packets are forwarded to the port groups for load balancing. Rejected packets are dropped.

Use the following steps to configure example 1.

Task	Description	UI Steps
1	Create GigaStreams that will be part of the port groups	a. Select Ports > Port Groups > GigaStreams b. Click New. c. Enter gs1 in the Alias field. d. In the Ports field, select port 1/1x16 and 1/1/x17. e. Click Save. f. Configure a second GigaStream with the alias gs2, select ports1/1/x1 and 1/1/x2 in the Ports field, and click Save.
2	Create port groups and specify the tool ports and assign GigaStreams to the port groups. The port groups will also be load balanced.	a. Select Ports > Port Groups > All Port Groups. b. Click New. c. Enter pg1 in the Alias field. d. Select Type GigaSMART Load Balancing. e. In the Ports field, select ports 1/1/x6 and 1/1/x7. f. In the GigaStream field, select gs1 g. Click Save. h. Configure a second Port Group. with the alias pg2, select ports1/1/x18 and 1/1/x10 in the Ports field, select GigaSMART Load Balancing, select pg2 in the GigaStream field, and then click Save.
3	Configure a GigaSMART group and associate it with a GigaSMART engine port.	a. From the device view, select GigaSMART > GigaSMART Groups. b. Click New.

Task	Description	UI Steps
		c. Enter GS31 in the Alias field. d. In the Port List field, select an engine port. For example 1/3/e1. e. Click Save.
4.	Create a virtual port. **Note**: You must enable GTP Overlap when configuring a virtual port for GTP overlap mapping.	a. From the device view, select GigaSMART > Virtual Ports > Virtual Ports. b. Click New c. Enter VP31 in the Alias field. d. In the GigaSMART Group field, select GS31. e. Select GTP Overlap, f. Click Save.
5.	Create the GTP Whitelist	a. From the device view, select GigaSMART > GTP Whitelist. b. Click New. c. Enter Whitelist in the Alias field d. Go to Task 6.
6.	Fetch whitelist files from a specified location to populate the GTP whitelist.	a. On the GTP Whitelist page, select Bulk Upload. b. Select Bulk Entry Operation for IMSI Upload Type c. Select Upload from URL from the Bulk Upload Type list. d. Enter the URL in the Enter Remote URL field. For example, http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx e. Click Save.
7.	Associate the GigaSMART group to the GTP whitelist.	a. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups b. Select GS Group GS31 created in Task 3 and click Edit c. Under GTP Whitelist, click on the GTP Whitelist Alias field and select Whitelist. d. Click Save.
8.	Configure the GigaSMART operation for GTP whitelisting.	a. From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation.

Task	Description	UI Steps
		b. Click New. c. Enter gtp-overlapwhitelist1 in the Alias field. d. Select the GigaSMART Group GS31from the GigaSMART Groups list. e. Select GTP Whitelist from the GigaSMART Operations (GSOP) list f. Select Load Balancing from the GigaSMART Operations (GSOP) list. g. Configure Load Balancing as follows: • Select Stateful • Set Type to GTP • Select Hashing • Select IMSI h. Click Save.
9.	Configure the GigaSMART operation for GTP flow sampling.	a. From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. b. Click New. c. Enter gtp-overlapsampling1 in the Alias field. d. Select the GigaSMART Group GS31from the GigaSMART Groups list. e. Select Flow Sampling-GTP f. Select Load Balancing from the GigaSMART Operations (GSOP) list. g. Configure Load Balancing as follows: • Select Stateful • Set Type to GTP • Select Hashing • Select IMSI h. Click Save.
10.	Configure the first level maps. In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	a. Configure the first map as follows: • Alias: mapLevel1-GTP • Type and Subtype: First Level By Rule

Task	Description	UI Steps
		• Source: 1/1/x1 • Destination: VP31 • Rule 1: Pass, Bi Directional, Port Destination 2123 • Rule 2: Pass, Bi Directional, Port Destination 2152 • Click Save.
11.	Configure the first second level GTP overlap map for GTP whitelisting If there is a match to an IMSI in the whitelist for GTP version 1 traffic, it is then forwarded to load balancing port group *pg1*.	a. Select Maps > Maps > Maps. b. Click New c. Configure the map. • Alias: WLMAP1 • Type and Subtype: Second Level GTP Flow Whitelist Overlap • Source: VP3 • Destination: pg1 • GSOP: gtp-whitelist • Rule 1: GTP, APN: Version V1 d. Click Save.
12.	Configure a second level map for GTP flow sampling, the flow sampling map. If there is not a match to an IMSI in the whitelist, the traffic flow is sampled based on the rules in this map. Accepted packets are forwarded to load balancing port group *pg1*.	a. Select Maps > Maps > Maps. b. Click New c. Configure the map. • Alias: FSMAP1 • Type and Subtype: Second Level GTP Flow Sample Overlap • Source: VP3 • Destination: pg1 • GSOP: gtp-overlapsample1 • Rule 1: GTP, IMSI: 3102609834*, IMEI: 35609506*, Percentage: 20 d. Click Save.
13.	Configure the next second level GTP overlap map for GTP whitelisting If there is a match to an IMSI in the whitelist for GTP version 2 traffic, it is then forwarded to	a. Select Maps > Maps > Maps. b. Click New. c. Configure the map. • Alias: WLMAP2

Task	Description	UI Steps
	load balancing port group *pg2*.	• Type and Subtype: Second Level Flow Whitelist Overlap • Source: VP31 • Destination: pg2 • GSOP: gtp-whitelist • Rule 1: GTP, APN: Version V2 d. Click Save.
14.	Configure the next second level map for GTP flow sampling. If there is not a match to an IMSI in the whitelist as evaluated by the second level GTP whitelisting map *WLMAP2*, the traffic flow is sampled based on the rules in this map. Accepted packets are forwarded to load balancing port group *pg2*.	a. Select Maps > Maps > Maps. b. Click New. c. Configure the map: • Alias: FSMAP2 • Type and Subtype: Second Level GTP Flow Sample Overlap • Source: VP31 • Destination: pg2 • Rule 1: GTP, IMSI: 3102609835*, IMEI: 35609507*, Percentage: 20 d. Click Save.
15.	Configure a map group. Add the GTP whitelisting and the two GTP flow sampling maps configured in previous steps.	a. Select Maps > Map Groups. b. Click New. c. Enter OverlapMap in the Alias field. d. In the Maps field, select WLMAP1,WLMAP2,FSMAP1,FSMAP2. e. Click Save.

GTP Stateful Session Recovery

Required License: GTP Filtering & Correlation

GTP sessions can be backed up periodically so they can then be recovered faster after a GigaSMART line card reboot or a node reboot. GTP stateful session recovery provides session persistence for GigaSMART GTP applications, including GTP flow filtering, GTP forward listing, and GTP flow sampling.

GTP stateful session recovery requires additional memory for storing backups. GigaVUE-HC3 has the required memory.

Using GTP stateful session recovery, the GTP session tables in the GigaSMART line card memory will be periodically backed up to the control card memory on the node and stored.

You can configure an interval for how often the backups occur, such as every 10 minutes. If GTP stateful session recovery is enabled and the GigaSMART line card is rebooted, the GTP session tables will be restored automatically following the reboot.

The last stored backup file will be downloaded from the control card to the GigaSMART line card using FTP. The session table will be repopulated from the last stored backup file to each GigaSMART engine, up to 8 engines. Packet count statistics for sessions are saved and will also be restored.

Depending on the size of the session table, the amount of time to restore from the backup might take as much as 3 minutes. During that interval, traffic will be blocked to the virtual port on the GigaSMART line card. Once the session table is read and populated, traffic will be allowed.

Depending on the interval between backups, there could be differences between the stored state and the current state of the system, for example, map configuration could change, or sessions could be added, modified, or deleted.

Load balancing information is not persisted, so after a session table is repopulated, a session that was once sent to one load balanced port may be sent to a different load balanced port after the reboot. However, for IMSI-based load balancing, the traffic might be sent to the same port as it was before the reboot.

GTP stateful session recovery works in a cluster environment; however, the cluster leader must remain the same.

To enable **GTP Session Recovery**, the system must retain specific session information. This is where **GTP Persistence** comes in. It allows us to define how long session data is stored and how the application should handle session continuity across interruptions.

Configure GTP Persistence

To enable GTP persistence, as well as to configure timers, use the GTP Persistence fields under GigaSMART Parameters on the GigaSMART Group configuration page shown in Figure 28GTP Persistence GigaSMART Parameters and select **GTP Persistence**. The timers are preconfigured with default values.



Figure 28 *GTP Persistence GigaSMART Parameters*

Use the **System** widget on the Overview page to determine the amount of memory. The size of memory will be 24Gb in an upgraded system. To view the System information, select **Overview** from the Navigation pane. The amount of free and used memory is displayed in the **Memory** field.

To see backup and restore information for GTP Persistence, select **GigaSMART > GigaSMART Groups > GigaSMART Group**s, and then click on the alias of the GigaSMART group. A Quick View opens for the selected GigaSMART group. Scroll down do GTP Persistence. In Figure 29GTP Persistence Information, GigaSMART Group gsgrp-1_4_e1 is selected and the Quick View displayed.

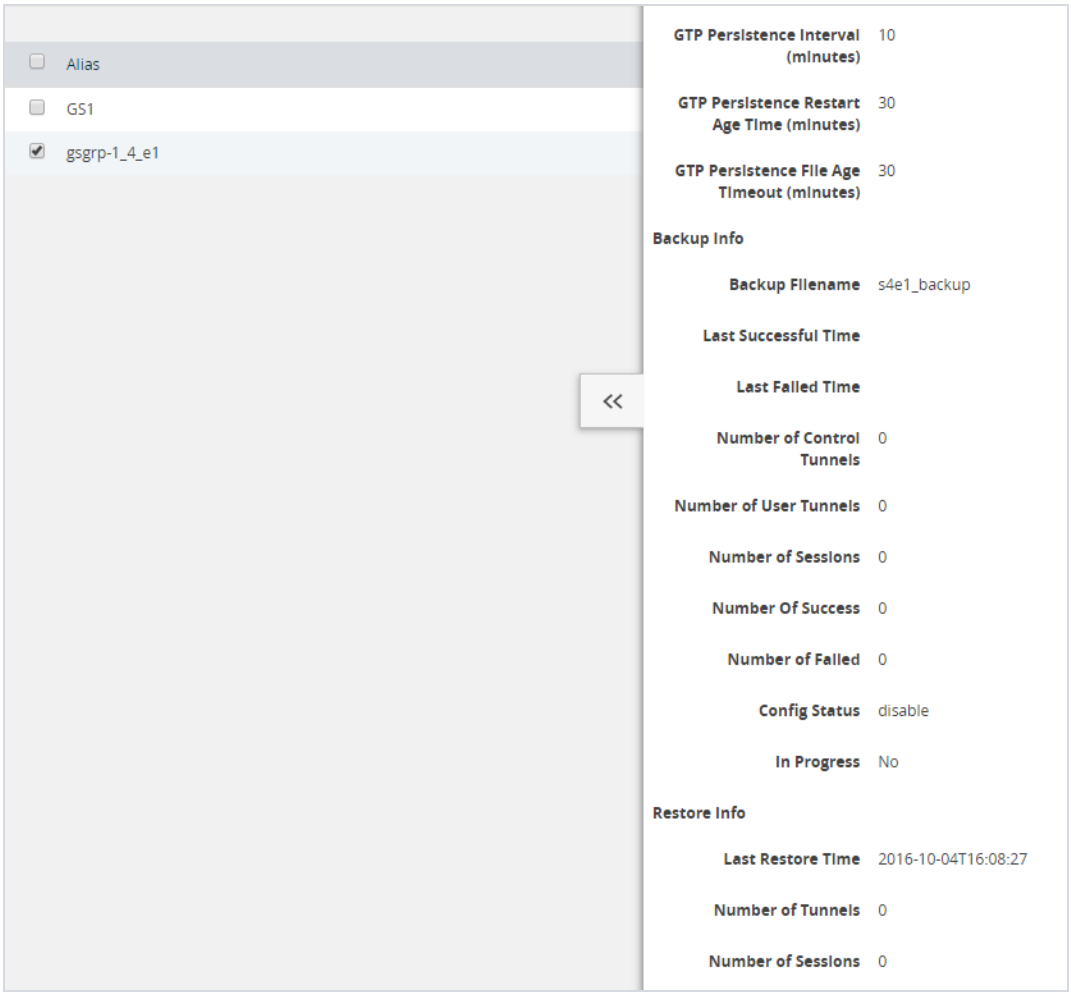


Figure 29 *GTP Persistence Information*

The following table describes persistence information.

Table 5: GigaSMART GTP Persistence Information

Name	Format
Backup Info	
Backup filename	The internal name of the backup file.
Last successful time	The timestamp of the last successful backup.
Last fail time	The timestamp of the last failed backup.
Number of control tunnels	The number of control tunnels backed up.
Number of user tunnels	The number of user tunnels backed up.
Number of sessions	The number of sessions backed up.
Number of success	The number of successful backups.
Number of failed	The number of failed backups.

Name	Format
Config Status	The status of a backup, which will be either Enabled or Disabled.
In Progress	The progress, which will be either Yes or No.
Restore Info	
Last restore time	The timestamp of the last restore.
Number of tunnels	The number of tunnels restored.
Number of sessions	The number of sessions restored.

To delete backup files, select the alias of GigaSMART Group and click **Edit**.Scroll down to GTP Persistence (refer to Figure 30GTP Backup Files Delete) and click **Delete All** under **GTP Backup Files**.

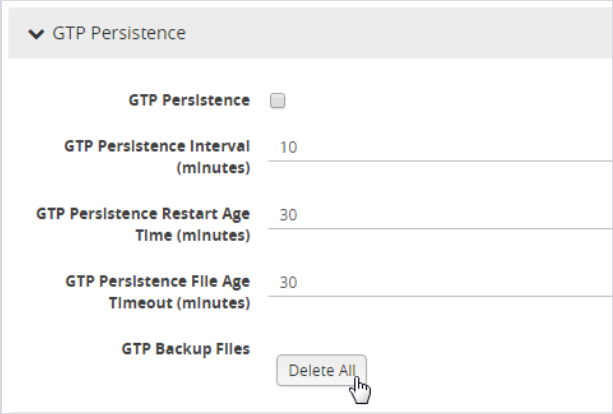


Figure 30 *GTP Backup Files Delete*

Refer to *GigaSMART GTP Stateful Session Recovery* in the GigaVUE-OS CLI Reference Guide.

GTP Scaling

GTP can be scaled as follows:

- GigaSMART Cards in GigaVUE-OS Devices
- GTP Engine Grouping

GigaSMART Cards in GigaVUE-OS Devices

Required License: GTP Filtering & Correlation

A total of four GigaSMART SMT-HC3-C08 line cards are supported on a single GigaVUE-HC3 node. This provides a total of eight GigaSMART engine ports, which increases the amount of GigaSMART processing available on the GigaVUE-HC3.

The increased number of GigaSMART line cards in the GigaVUE-HC3 can be used by the following GTP applications: GTP flow filtering, GTP flow sampling, and GTP whitelisting.

GTP Engine Grouping

Required License: GTP Filtering & Correlation

A GigaSMART group (gsgroup) associated with GTP applications can have multiple GigaSMART engine port members. You can combine up to eight engine ports to form an engine group. The engine group provides higher capacity to GTP applications by load balancing GTP user-data plane (GTP-u) traffic among the members of the group. Grouping multiple GigaSMART engine ports increases the effective throughput for GTP applications.

NOTE: A GigaSMART Engine Group supports supports 12 million GTP subscriber sessions for GigaVUE HC3 nodes.

GTP engine grouping is supported on GigaVUE-HC3 nodes.

GTP engine grouping can be used by the following GTP applications: GTP flow filtering, GTP flow sampling, and GTP whitelisting.

The following table lists GTP engine grouping support for GigaVUE nodes:

GigaVUE Node	Maximum Number of GigaSMART Line Cards per Node	Number of e ports per Line Card	Supported Number of e ports per GigaSMART Group	Location of e ports
GigaVUE-HC3	4	2	8	**e1** and **e2** on same module
GigaVUE-HC1	Not supported in this software version.			

Keep in mind the following recommendations and restrictions:

- Configure a GTP engine group on a single GigaVUE node.
- GTP engine grouping only supports IMSI hash-based load balancing.
- GTP engine grouping is limited to out-of-band cluster configurations in this software version.

Passing GTP Control Traffic

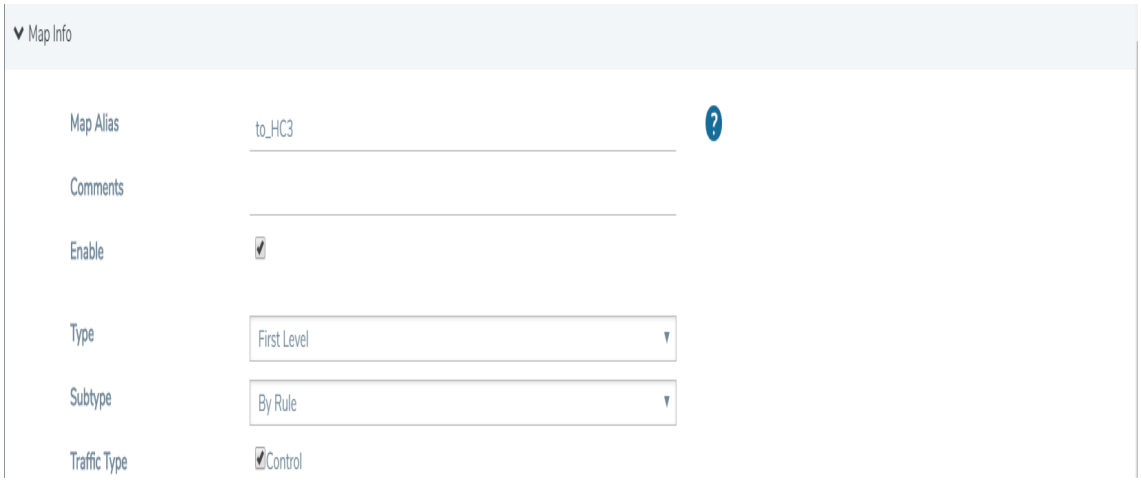
Enable the **ControlTraffic** checkbox when creating a First Level By Rule map, to allow GTP applications to pass GTP control traffic (GTP-c) to all GigaSMART engines in a GTP engine group. GTP-c traffic is sent to all members of the engine group in order to replicate the session tables.

> **NOTE:** In the map with the **ControlTraffic** enabled, only one vport is supported.

You can edit the map with **ControlTraffic** option enabled. For example, **Control Traffic** can be selected in an existing first level map, or it can be deleted from a first level map. Other editing, such as changing the **Source** or the **Destination** in the first level map is also allowed.

To set GTP Control Traffic:

1. From the device view, select **GigaSMART > Maps > Maps** and click **New**.
2. On the New Map page, select **First Level** and **By Rule** for the map type and subtype, respectively.
3. Enable **Control Traffic** as shown in the following figure.
4. Configure the map and click **Save** when done.



Upgrade from Earlier Release

When there is existing GTP configuration with one engine port per GigaSMART group in an earlier software version, an upgrade from that earlier software version to a higher release will succeed.

However you cannot convert that configuration to multiple engine ports per GigaSMART group. You must delete the configuration and reconfigure it, including the GigaSMART group, GigaSMART operation, virtual port, and maps. This is due to the need for separate maps for GTP control plane and GTP user plane traffic in higher releases.

Modify Engine Ports in GigaSMART Group

You can modify the engine ports in a GigaSMART group. For example, you can add an engine port to a GigaSMART group or remove an engine port from a GigaSMART group. After the change, reset all the GigaSMART line cards or modules that have engine ports configured in the GigaSMART group.

Modify vports in Map

You can modify the vport relating to the first level map with Control Traffic enabled. For example, you can change the vport configured in the map. After the change, reset all the GigaSMART line cards or modules that have engine ports configured in the vport.

Configure GTP Engine Grouping

Refer to the following examples:

- GTP Scaling
- GTP Scaling

Display Statistics

To display the GigaSMART Group statistics, select **GigaSMART > GigaSMART Groups > Statistics**.

Refer to GigaSMART Group Statistics Definitions for descriptions of the statistics.

GTP Engine Grouping Configuration Examples

To configure GTP Engine Grouping refer to the following examples:

- GTP Engine Grouping Configuration Example
- GTP Engine Grouping Configuration Complex Example

GTP Engine Grouping Configuration Example

This is an example of a GTP engine group consisting of two engine ports on a GigaVUE-HC1node. This example includes a GigaSMART operation for GTP flow filtering.

Task	Description	UI Steps
1	Configure ports as follows: o one network type of port. This will be used as the **Source** attribute in two first level maps in **Task 5** and **Task 6**. o one tool type of port for the **Destination** attribute in a second level flow filtering map in **Task 7**. o one tool type of port for the **Destination** attribute in a shared collector map in **Task 8**. Then administratively enable the ports.	**1.** Select Ports > Ports > All Ports **2.** Click Quick Port Editor **3.** Configure a network port and two tool ports. For example, select Network for port 22/3/c1 and select Tool for ports 22/3/c2 and 22/1/c3. **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with two GigaSMART engine port, to form the GTP engine group. The GigaSMART group will be used in **Task 5** and **Task 6**.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Type gsg2 in the Alias field. **3.** Click in the Port List field to add the two engine ports. **4.** Click Save.
3	For GTP flow filtering, configure a flow filtering GigaSMART operation and assign it to the GigaSMART group. The gsop will be used in the second level flow filtering map in **Task 7**.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **2.** Type gtp_gsg2 in the Alias field. **3.** Select gsg2 from the GigaSMART Groups list. **4.** Select Flow Filtering from the GigaSMART Operations list. (GSOP). **5.** Click Save.
4	Configure a virtual port and assign it to the same GigaSMART group. This virtual port will be used as the **Destination** in the first level maps in **Task 5** and **Task 6**, as the **Source** in the second level map in **Task 7**, and as the **Source** attribute in the shared collector map in **Task 8**.	**1.** From the device view, select GigaSMART > Virtual Ports. **2.** Click New. **3.** Type vp1 in the Alias field. **4.** Select gsg2 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map that directs GTP control traffic from the physical network port to the virtual port	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map.

Task	Description	UI Steps
	created in **Task 4**. **NOTE:** In the rule, 2123 is GTP-c traffic. This map, with the **Control Traffic** attribute, identifies the GTP-c control traffic needed for GTP engine grouping. **Note**: The order of configuration is important. Set **Control Traffic** before any map rules.	▪ Type gtp_to_vp1-c in the Alias field ▪ Select First Level for Type ▪ Select By Rule for Subtype ▪ Enable Control Traffic ▪ Select 22/3/c1 for Source ▪ Select vp1 for Destination **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass and Bi Directional **c.** Select Port Destination and specify 2123 **4.** Click Save.
6	Create another first level map that directs GTP user traffic from the physical network port to the virtual port created in **Task 5**. **Note**: In the rule, 2152 is GTP-u traffic. GTP-u traffic corresponding to the same GTP-c traffic will be sent to the same virtual port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type gtp_to_vp1 in the Alias field ▪ Select First Level for Type ▪ Select By Rule for Subtype ▪ Enable Control Traffic ▪ Select the network configured in Task 1 for Source ▪ Select vp1 for Destination **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass and Bi Directional **c.** Select Port Destination and specify 2123 **4.** Click Save.
7	Create a second level map for GTP flow filtering that takes traffic from the virtual port, applies the flow filtering GigaSMART operation, matches IMEIs and version specified by the flow rule, and sends matching traffic to a tool port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Type from_vp1 in the Alias field. **4.** Configure the map. ▪ Select Second Level for Type ▪ Select Flow Filter for Subtype ▪ Select vp1 for Source ▪ Select one of the tool ports configured in Task 1 for Destination. ▪ Select gtp_gsg2 for from the GSOP list. **5.** Add a Rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select GTP IMEI.

Task	Description	UI Steps
		d. Specify * in the IMEI field. **e.** Select V2 for Version. **6.** Click Save.
8	Add a shared collector for any unmatched traffic from the virtual port and send it to a different tool port than in **Task 7**.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type from_vp1_scoll in the Alias field. ▪ Select Second Level for Type ▪ Select Collector for Subtype ▪ Select vp1 for Source. ▪ Select the other tool port configured in Task 1 for Destination. **4.** Click Save.

GTP Engine Grouping Configuration Complex Example

This is a more complex example of GTP engine grouping than the previous example. This example has four engine ports on two GigaSMART line cards on the same GigaVUE-HC3 node. The GigaSMART line cards are in slots 1 and 3.

The GigaVUE-HC3 node is the cluster leader of a two-node out-of-band cluster. A GigaVUE-HC1 is the standby node in the cluster.

This example includes GigaSMART operations for GTP flow filtering with load balancing, GTP flow sampling with load balancing, and GTP whitelisting. The forward list must be associated with the GigaSMART group on the leader, the GigaVUE-HC3.

Task	Description	UI Steps
1	Configure ports on the GigaVUE-HC3 as follows: ▪ One network type of port. This will be used as the **Source** attribute in two first level maps in **Task 11** and **Task 12**. ▪ Twelve tool type of ports. There are four tool ports	**1.** On the GigaVUE-HC3, select Ports > Ports > All Ports **2.** Click Quick Port Editor. **3.** Configure one network port for **Task 11** and **Task 12**. For example, 23/2/c3. **4.** Configure 12 tool ports for three port groups in **Task 6**. For example, 23/3/x1..x4 for the first port group, 23/3/x9..x12 for the second tool group, and 23/3/x13..x16 for the third port group. **5.** Configure five tool ports for a GigaStream in **Task 2**. For example, 23/3/x20..x24.

Task	Description	UI Steps
	in each of three port groups used for load balancing. The port groups will be created in **Task 6**. ▪ Five tool type of ports for a GigaStream that will be created in **Task 2**. ▪ Two tool type of ports for another GigaStream that will be created in **Task 2**. Then administratively enable the ports.	**6.** Configure two tool ports for another GigaStream in **Task 2**. For example, 23/2/c1..c2 type tool. **7.** Select enable for each port. **8.** Click OK. **9.** Close the Quick Port Editor.
2	On the GigaVUE-HC3, configure one GigaStream using five tool ports. This will be used as the **Destination** attribute in the map in **Task 11**. Configure another GigaStream to be used in the stack link between the GigaVUE-HC3 and GigaVUE-HC1 that will be created in **Task 4**.	**1.** Select Ports > Port Groups > GigaStreams **2.** Configure a GigaStream with five tool ports. **a.** Click New. **b.** Type hc3-gs-1 in the Alias field. **c.** Select Tool GigaStream. **d.** Click in the Ports field and select the five tool ports configured in Task 1. **e.** Click Advanced Hash Settings and use the Default setting. **f.** Click Save. **7.** Configure another GigaStream with two tool ports. **a.** Click New. **b.** Type hc3-80g in the Alias field. **c.** Select Tool GigaStream. **d.** Click in the Ports field and select the two tool ports configured in Task 1. **e.** Click Advanced Hash Settings and use the Default setting. **f.** Click Save.
3	Configure ports on the GigaVUE-HC1 as follows: ▪ Two tool type of ports for a GigaStream that will be created in **Task 4**. ▪ One tool type of port that will be used as the **Destination** in a map in **Task 11**.	**1.** On the GigaVUE-HC1, select Ports > Ports > All Ports **2.** Click Quick Port Editor. **3.** Configure two tool ports for a GigaStream in **Task 4**. For example, 33/2/q1..q2 **4.** Configure one tool ports for the map in **Task 11**. For example, 33/3/x11. **5.** Configure four tool ports for a GigaStream in **Task 4**. For example, 33/2/x20..x24.

Task	Description	UI Steps
	▪ Four tool type of ports for a GigaStream that will be created in **Task 4**. Then administratively enable the ports.	**6.** Select Enable for each port. **7.** Click OK. **8.** Close the Quick Port Editor.
4	On the GigaVUE-HC1, configure a GigaStream using two tool ports. This will be used in the stack link created in **Task 5**. Configure another GigaStream using four tool ports. This will be used in the shared collector in **Task 17**.	**1.** Select Ports > Port Groups > GigaStreams **2.** Configure a GigaStream with five tool ports. **a.** Click New. **b.** Type hc3-gs-4 in the Alias field. **c.** Select Tool GigaStream. **d.** Click in the Ports field and select the five tool ports configured in Task 1. For example, 33/2/x20..x25. **e.** Click Advanced Hash Settings and use the Default setting. **f.** Click Save. **7.** Configure another GigaStream with two tool ports. **a.** Click New. **b.** Type hc3-80g in the Alias field. **c.** Select Tool GigaStream. **d.** Click in the Ports field and select the two tool ports configured in Task 1. For example, 33/2/q1..q2. **e.** Click Advanced Hash Settings and use the Default setting. **f.** Click Save.
5	Configure the stack link between the GigaVUE-HC1 and GigaVUE-HC3.	**1.** Select Ports > Port Groups > Stack Links **2.** Click New. **3.** Select Stack GigaStream. **4.** For First Member select hc3-80g. **5.** For Second Member select hc3-80g
6	Create three port groups and specify four tool ports each, for load balancing. Also, enable load balancing on each port group. The port groups, hc3-pg-1 and hc3-pg-2, will be used as the **Destination** in two second level flow sampling maps in **Task 14** and **Task 15**. The port group, hc3-q2x32-1-4, will be used as the **Destination** in a second level flow filtering map in **Task 16**	**1.** Select Ports > Port Groups > All Port Groups. **2.** Create the first port group. **a.** Click New. **b.** Type hc3-pg-1 in the Alias field. **c.** Select SMART Load Balancing. **d.** Click in the Ports field and select four tool ports. For example, 23/4/x9..x12. **e.** Click Save. **6.** Create the second port group. **a.** Click New.

Task	Description	UI Steps
		b. Type hc3-pg-1 in the Alias field. **c.** Select SMART Load Balancing. **d.** Click in the Ports field and select four tool ports. For example, 23/4/x13..x16 **e.** Click Save. **6.** Create the second port group. **a.** Click New. **b.** Type hc3-q2x32-1-4 in the Alias field. **c.** Select SMART Load Balancing. **d.** Click in the Ports field and select four tool ports. For example, 23/4/x13..x16 **e.** Click Save.
7	Configure a GigaSMART group and associate it with four GigaSMART engine ports, two in slot 1 and two in slot 3, to form the GTP engine group. The GigaSMART group will be used in **Task 8**, **Task 9**, and **Task 10**.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type hc3scale-4engines-slots1and3. **4.** Click in the Port List and select the engine ports. For example, 3/1/e1,23/1/e2,23/3/e1,23/3/e2. Go to **Task 8**.
8	Associate the GigaSMART group to an existing GTP forward list. **Note**: The forward list is only supported on the cluster leader, which is the GigaVUE-HC3 in this example.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Under GTP Whitelist, select the alias of an existing forward list. For example, gtp-whitelist. **3.** Click Save.
9	For GTP flow filtering, configure a flow filtering GigaSMART operation, specify load balancing, and assign the GigaSMART operation to the GigaSMART group. The hc3-scale-ff-lb gsop will be used in the second level flow filtering map in **Task 16**. For GTP flow sampling, configure a flow sampling GigaSMART operation, specify load balancing, and assign the GigaSMART operation to the GigaSMART group. The hc3-scale-fs-lb gsop will be used in the two second level flow sampling maps in **Task 15** and **Task 16**. For GTP forward listing, configure a	**1.** Configure a Flow Filtering operation. **a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Type hc3-scale-ff-lb in the Alias field. **d.** Select hc3scale-4engines-slots1and3 from the GigaSMART Groups list. **e.** Select Flow Filtering from the GigaSMART Operations (GSOP) list. **f.** Click Save. **7.** Configure a Flow Sampling operation. **a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation.

Task	Description	UI Steps
	whitelisting GigaSMART operation, and assign the GigaSMART operation to the GigaSMART group. (This GigaSMART operation is not load balanced.) The hc3-scale-wl gsop will be used in the second level forward listing map in **Task 13**.	**b.** Click New. **c.** Type hc3-scale-fs-lb in the Alias field. **d.** Select hc3scale-4engines-slots1and3 from the GigaSMART Groups list. **e.** Select Flow Sampling from the GigaSMART Operations (GSOP) list. **f.** Select Flow Sampling -GTP **g.** Click Save. **8.** Configure a Forward listing operation. **a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Type hc3-scale-wl in the Alias field. **d.** Select hc3scale-4engines-slots1and3 from the GigaSMART Groups list. **e.** Select GTP Whitelist from the GigaSMART Operations (GSOP) list. **f.** Select Enabled (default). **g.** Click Save.
10	Configure a virtual port and assign it to the same GigaSMART group. This virtual port will be used as the **Destination** in the first level maps in **Task 11** and **Task 12**, as the **Source** in the second level maps in **Task 13**, **Task 14**, **Task 15**, **Task 16**, and as the **Source** in the shared collector in **Task 17**.	**1.** From the device view, select GigaSMART > Virtual Ports **2.** Click New. **3.** Type vp-hc3scale-4engines-slots1and3 in the Alias field. **4.** Select hc3scale-4engines-slots1and3 from the GigaSMART Groups list. **5.** Click Save.
11	Create a first level map that directs GTP control traffic from the physical network port to the virtual port created in **Task 10**. **Note**: In the rule, 2123 is GTP-c traffic. This map, with the **Control Traffic** enabled, identifies the GTP-c control traffic needed for GTP engine grouping. In addition to the virtual port, traffic is also sent to a GigaStream and a tool port.	**1.** Select Maps > Maps > Maps. **2.** Configure the map. ▪ Type to_hc3_gtpc in the Alias field ▪ Select First Level for Type. ▪ Select By Rule for Subtype ▪ Enable Control Traffic ▪ Select a network port for the Source. For example, 23/3/q6. ▪ Select the GigaStream port hc3-gs-1, the virtual port p-hc3scale-4engines-slots1and3, and a tool port (for example,

Task	Description	UI Steps
		23/7/q6) for the Destination. **3.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass and Bi Directional. **c.** Select Port Destination for the rule. **d.** Enter 2123 for the port value. **5.** Click Save.
12	Create another first level map that directs GTP user traffic from the physical network port to the virtual port created in **Task 10**. **Note**: In the rule, 2152 is GTP-u traffic. GTP-u traffic corresponding to the same GTP-c traffic will be sent to the same virtual port.	**1.** Select Maps > Maps > Maps. **2.** Configure the map. ▪ Type to_hc3_gtpu_1 in the Alias field ▪ Select First Level for Type. ▪ Select By Rule for Subtype ▪ Select a network port for the Source. For example, 23/7/q6. ▪ Select the virtual port vp-hc3scale-4engines-slots1and3 for the Destination. **3.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass and Bi Directional. **c.** Select Port Destination for the rule. **d.** Enter 2152for the port value. **5.** Add a second rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Fragmentation for the rule. **d.** Enter 2152for the port value. **e.** Select allFragNoFirst for Value. **6.** Click Save.
13	Configure a second level map for GTP whitelisting, the forward list map, that takes traffic from the virtual port, applies the forward listing GigaSMART operation, and sends traffic to the remote GigaVUE-HC1 node through a GigaStream.	**1.** Select Maps > Maps > Maps. **2.** Configure the map. ▪ Type from_hc3_wl in the Alias field ▪ Select Second Level for Type. ▪ Select By Rule for Subtype ▪ Select the virtual port vp-hc3scale-4engines-slots1and3 for the Source. ▪ Select the GigaStream hc3-gs-1 for the Destination. ▪ Select hc3-scale-wl from the GSOP list. **3.** Click Save.

Task	Description	UI Steps
14	Configure a second level map for GTP flow sampling. This is the first of two flow sampling maps. This map filters for version 2. It takes traffic from the virtual port and applies the flow sampling GigaSMART operation. Traffic flow is sampled based on the flow sampling rule in this map. Accepted packets are forwarded to load balancing port group hc3-pg-2.	**1.** Select Maps > Maps > Maps. **2.** Configure the map. ▪ Type from_hc3_fs_v2 in the Alias field ▪ Select Second Level for Type. ▪ Select Flow Sample for Subtype ▪ Select the virtual port vp-hc3scale-4engines-slots1and3 for the Source. ▪ Select the port group hc3-pg-2 for the Destination. ▪ Select hc3-scale-fs-lb from the GSOP list. **3.** Add a rule. **a.** Click Add a Rule. **b.** Select GTP. **c.** Enter 60 for Percentage. **d.** Enter 5* in the IMSI field. **e.** Select V2 for Version **6.** Click Save.
15	Configure a second level map for GTP flow sampling. This is the second of two flow sampling maps. This map filters for version 1. It takes traffic from the virtual port and applies the flow sampling GigaSMART operation. Traffic flow is sampled based on the flow sampling rule in this map. Accepted packets are forwarded to load balancing port group hc3-pg-1.	**1.** Select Maps > Maps > Maps. **2.** Configure the map. ▪ Type from_hc3_fs_v1 in the Alias field ▪ Select Second Level for Type. ▪ Select Flow Sample for Subtype ▪ Select the virtual port vp-hc3scale-4engines-slots1and3 for the Source. ▪ Select the port group hc3-pg-1 for the Destination. ▪ Select hc3-scale-fs-lb from the GSOP list. **3.** Add a rule. **a.** Click Add a Rule. **b.** Select GTP. **c.** Enter 60 for Percentage. **d.** Enter 5* in the IMSI field. **e.** Select V1 for Version **6.** Click Save.
16	Create a second level map for GTP flow filtering that takes traffic from the virtual port, applies the flow filtering GigaSMART operation, matches IMSIs specified by the flow rule, and sends matching traffic to load balancing port group hc3-	**1.** Select Maps > Maps > Maps. **2.** Configure the map. ▪ Type from_hc3_ff in the Alias field ▪ Select Second Level for Type. ▪ Select Flow Filter for Subtype

Task	Description	UI Steps
	q2x32-1-4.	▪ Select the virtual port vp-hc3scale-4engines-slots1and3 for the Source. ▪ Select the port group hc3-pg-1 for the Destination. ▪ Select port group hc3-q2x32-1-4 from the GSOP list. **3.** Add a rule. **a.** Click Add a Rule. **b.** Select GTP IMSI **c.** Enter * in the IMSI field. **d.** Select Any for Version **5.** Click Save.
17	Add a shared collector for any unmatched traffic from the virtual port and send it to a GigaStream.	**1.** Select Maps > Maps > Maps. **2.** Configure the map. ▪ Type s_coll_hc3 in the Alias field ▪ Select Second Level for Type. ▪ Select Collector for Subtype ▪ Select the virtual port vp-hc3scale-4engines-slots1and3 for the Source. ▪ Select GigaStream hc3-gs-4 for the Destination. **3.** Click Save.

GigaSMART TLS/SSL Decryption for Inline Tools

A Gigamon inline solution refers to network configuration where a GigaVUE node is placed directly in the data path ("inline") between two network segments, allowing it to capture and analyze live network traffic while simultaneously directing traffic to security tools like firewalls or intrusion prevention systems (IPS) for inspection, all while maintaining network connectivity and minimizing disruption; essentially acting as a "traffic broker" for inline security tools with fail-safe bypass capabilities in case the security appliance fails.

To know more about essential features and capabilities of Gigamon's Inline Solution refer to Gigamon Inline Solutions

GigaSMART TLS/SSL Decryption for Out-of-Band Tools

Required License: TLS/SSL Decryption for Out-of-Band Tools

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

TLS/SSL decryption for out-of-band tools is described in GigaSMART Passive TLS/SSL Decryption. It is supported on all GigaVUE-HC1, GigaVUE-HC1-Plus, and GigaVUE-HC3 devices.

GigaSMART Passive TLS/SSL Decryption

GigaVUE H Series nodes support Secure Sockets Layer (SSL) decryption. TLS/SSL is a cryptographic protocol that adds security to (Transmission Control Protocol) TCP/IP communications such as Web browsing and email. The protocol allows the transmission of secure data between a server and client who both have the keys to decode the transmission and the certificates to verify trust between them. Passive TLS/SSL decryption delivers decrypted traffic to out-of-band tools that can then detect threats entering the network.

Passive TLS/SSL decryption is a pillar of the GigaSECURE Security Delivery Platform. For an overview of GigaSECURE, refer to GigaSECURE Security Delivery Platform.

On GigaVUE H Series nodes, GigaSMART line cards or modules perform the decryption of TLS/SSL traffic. Using GigaSMART for decryption offloads the decryption function from tools and offers improved tool performance by removing this computationally intensive task. GigaSMART provides a centralized decryption point. Decrypted TLS/SSL traffic can be sent from GigaSMART to inspection tools for further analysis, for example, to look at encrypted communications or to detect malware.

Before TLS/SSL traffic is decrypted, the de-duplication GigaSMART operation can be performed. Decrypted traffic from the GigaSMART line card or module can be filtered, aggregated, and replicated and then sent to one or more monitoring tools for analysis.

Passive TLS/SSL decryption is supported on the following GigaVUE H Series products with GigaSMART line cards or modules installed:

- GigaVUE-HC3
- GigaVUE-HC1
- GigaVUE-HC1-Plus
- GigaVUE-HCT

Use Passive TLS/SSL decryption on the GigaSMART line card or module with passive or offline traffic. Tap the traffic to and from a server and pass it to the GigaVUE H Series node with the GigaSMART line card or module.

Passive TLS/SSL decryption operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to Groups of GigaSMART Engine Ports for details.

For secure storage of private keys in Passive TLS/SSL decryption scenarios, integration with a Hardware Security Module (HSM) is supported. Thales Luna Network HSMs are compatible and can be used based on deployment requirements. Refer to Thales-Luna Network HSM for TLS/SSL Decryption for Out-of-Band Tools(Passive) for details.

Gigamon also offers inline TLS/SSL decryption, which inspects TLS/SSL encrypted traffic inline. Refer to Work With Inline TLS/SSL Decryption for details.

About Passive TLS/SSL Decryption

TLS/SSL encryption secures traffic between a client and a server, such as a Web server. TLS/SSL decryption uses keys to decode the traffic between the client and server.

SSL and Transport Layer Security (TLS) protocols consist of a set of messages exchanged between a client and server to set up and tear down the TLS/SSL connection between them. To set up the connection, the client and server use the Public Key Infrastructure (PKI) to exchange the bulk encryption keys needed for data transfer.

Figure 31Basic TLS/SSL Handshake shows the basic TLS/SSL handshake between a client and server to establish a session. The messages are unencrypted up to step 6 in Figure 31Basic TLS/SSL Handshake. The messages are encrypted after step 6, including the step 9 Finished message.

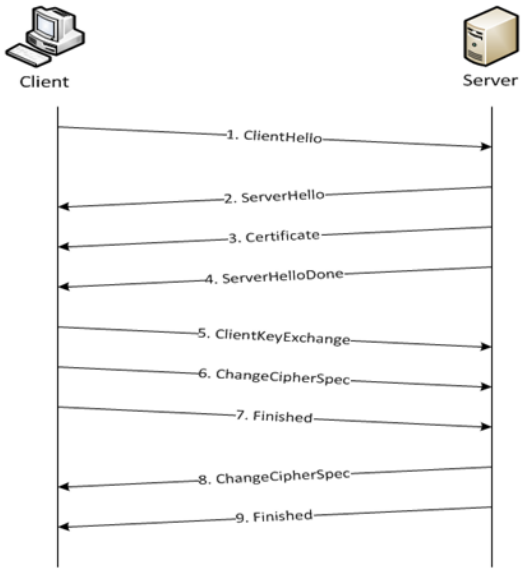


Figure 31 *Basic TLS/SSL Handshake*

Once a session has been established, the keys are saved so a session can be resumed efficiently later. Figure 32Resumed TLS/SSL Handshake shows the resumed TLS/SSL handshake, with fewer steps.

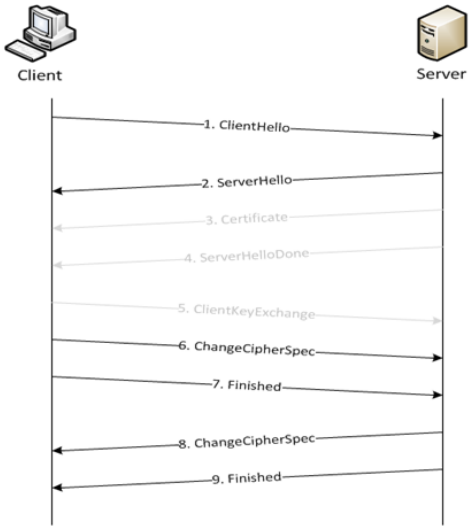


Figure 32 *Resumed TLS/SSL Handshake*

Passive TLS/SSL decryption can be deployed close to the server, as shown in Figure 33Inbound (Server Side).

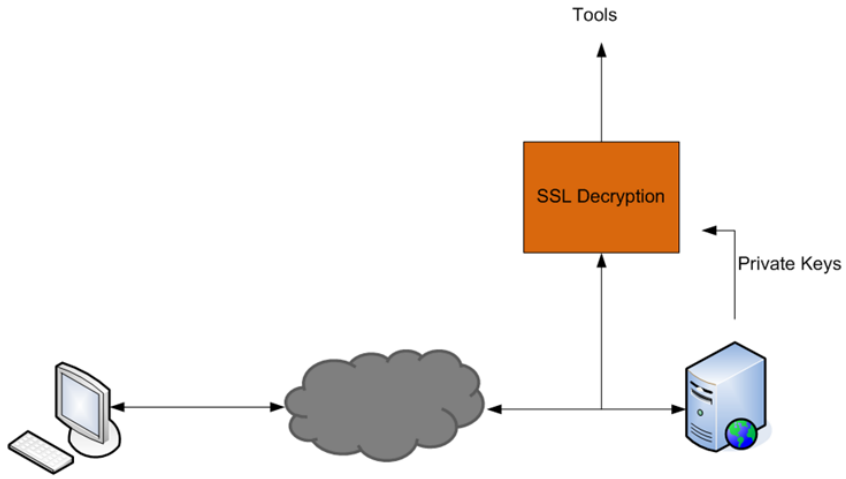


Figure 33 *Inbound (Server Side)*

Passive TLS/SSL decryption can also be deployed close to an TLS/SSL proxy, with the server in the Enterprise domain as shown in Figure 34Outbound (Client/Enterprise Side).

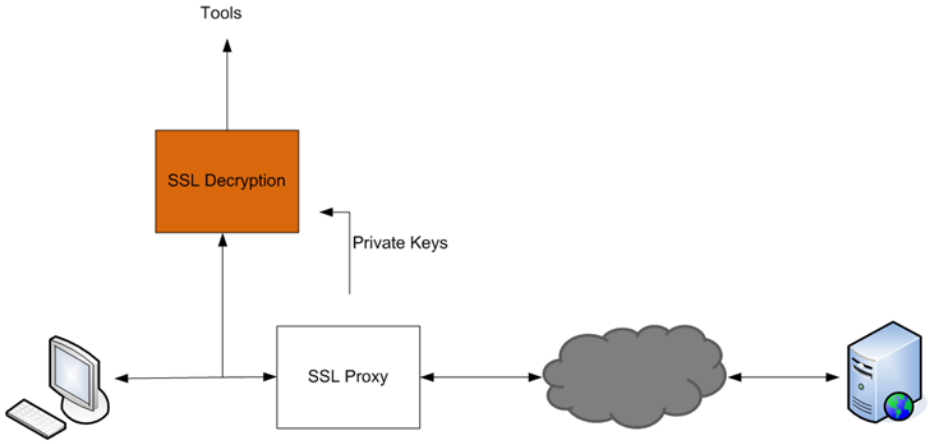


Figure 34 *Outbound (Client/Enterprise Side)*

In Figure 33Inbound (Server Side), tap traffic to the server and then send it for decryption. In Figure 34Outbound (Client/Enterprise Side), tap traffic to the proxy and then send it for decryption. You can have a deployment with either a server or a proxy, but not both.

The following sections describe Passive TLS/SSL decryption on GigaSMART:

- Supported Protocols, Algorithms, and Ciphers
- Limitations
- Create and Reset TLS/SSL Keychain Passwords
- Work with Keys and Services

Supported Protocols, Algorithms, and Ciphers

The supported protocols are as follows:

- SSL 3.0
- TLS 1.0
- TLS 1.1
- TLS 1.2

The supported authentication (Au) is as follows:

- RSA

The supported key exchange (Kx) is as follows:

- RSA

The supported encryption algorithms (Enc) are as follows:

- NULL
- RC4
- DES
- 3DES
- AES (including GCM mode)
- CAMELLIA
- SEED
- IDEA

The supported compression algorithm is as follows:

- NULL

The supported digest algorithms are as follows:

- MD5
- SHA1
- SHA2

The supported key sizes are 128, 256, 512, 1024, 2048, and 4096.

The supported TLS extensions are as follows:

- Extended Master Secret, RFC 7627
- Encrypt-then-MAC, RFC 7366

The supported ciphers are listed in Table 6: Supported Ciphers for Passive SSL decryption.

Table 6: Supported Ciphers for Passive SSL decryption

Cipher Name	Kx	Au	Enc	Bits	Mac
TLS_RSA_WITH_NULL_MD5	RSA	RSA	NULL	0	MD5
TLS_RSA_WITH_NULL_SHA	RSA	RSA	NULL	0	SHA
TLS_RSA_EXPORT_WITH_RC4_40_MD5	RSA_EXPORT	RSA_EXPORT	RC4_40	40	MD5
TLS_RSA_WITH_RC4_128_MD5	RSA	RSA	RC4_128	128	MD5
TLS_RSA_WITH_RC4_128_SHA	RSA	RSA	RC4_128	128	SHA
TLS_RSA_EXPORT_WITH_RC2_CBC_40_ MD5	RSA_EXPORT	RSA_EXPORT	RC2_CBC_40	40	MD5
TLS_RSA_WITH_IDEA_CBC_SHA	RSA	RSA	IDEA_CBC	128	SHA
TLS_RSA_EXPORT_WITH_DES40_CBC_SHA	RSA_EXPORT	RSA_EXPORT	DES40_CBC	40	SHA
TLS_RSA_WITH_DES_CBC_SHA	RSA	RSA	DES_CBC	56	SHA
TLS_RSA_WITH_3DES_EDE_CBC_SHA	RSA	RSA	3DES_EDE_CBC	168	SHA
TLS_RSA_WITH_AES_128_CBC_SHA	RSA	RSA	AES_128_CBC	128	SHA
TLS_RSA_WITH_AES_256_CBC_SHA	RSA	RSA	AES_256_CBC	256	SHA
TLS_RSA_WITH_CAMELLIA_128_CBC_SHA	RSA	RSA	CAMELLIA_128_CBC	128	SHA
TLS_RSA_WITH_CAMELLIA_256_CBC_SHA	RSA	RSA	CAMELLIA_256_CBC	256	SHA
TLS_RSA_WITH_SEED_CBC_SHA	RSA	RSA	SEED_CBC	128	SHA
TLS_RSA_WITH_NULL_SHA256	RSA	RSA	NULL	0	SHA256
TLS_RSA_WITH_AES_128_CBC_SHA256	RSA	RSA	AES_128_CBC	128	SHA256
TLS_RSA_WITH_AES_256_CBC_SHA256	RSA	RSA	AES_256_CBC	256	SHA256
TLS_RSA_WITH_AES_128_GCM_SHA256	RSA	RSA	AES_128_GCM	128	SHA256
TLS_RSA_WITH_AES_256_GCM_SHA384	RSA	RSA	AES_256_GCM	256	SHA384

All algorithms used for Passive TLS/SSL decryption are FIPS 140-3 compliant.

All key URLs must point to an RSA private key stored in the PEM or PKCS12 format, as follows:

- http://keyserver.domain.com/path/keyfile.pem
- https://keyserver.domain.com/path/keyfile.pem
- ftp://keyserver.domain.com/path/keyfile.pem
- tftp://keyserver.domain.com/path/keyfile.pem
- scp://username[:password]@keyserver.domain.com/path/keyfile.pem

The supported applications are as follows:

- HTTPS
- FTPS
- SMTP, IMAP, and POP3 with StartTLS

Limitations

The limitations of Passive TLS/SSL decryption are as follows:

- Only IPv4
- Only regular maps; no virtual ports (vports)
- Only combined with the de-duplication GigaSMART operation
- Only one private key per PKCS12 file
- Only server-side authentication
- Only the protocols and ciphers listed in Supported Protocols, Algorithms, and Ciphers.

> **NOTE:** If an TLS/SSL session cannot be decrypted due to having a non-supported protocol or cipher and if the GigaSMARTS Parameter **SSL Decryption** has **Decrypt Fail Action** is set to **Pass to Tool Port**, the packets will be forwarded to the tool without decryption . Non-supported ciphers and protocols include SSL 2.0, Diffie-Hellman (DHE keys), Ephemeral keys, Elliptic Curves Extension, compression, and 8K key size.

Licensing

The GigaSMART license for Passive TLS/SSL decryption is installed as any other license.

There are no limits to the number of Passive TLS/SSL decryption sessions or the number of users.

Create and Reset TLS/SSL Keychain Passwords

To perform the configuration in the following section, you must have an admin level access role.

Before uploading keys or configuring TLS/SSL, you must create an SSL keychain password. The password is used to encrypt the private keys that you upload to the node.

Keychain passwords are not saved on the node. Refer to Notes about Private Keys and Passwords.

> **NOTE:** When uploading TLS/SSL keys, make sure that you are not creating a duplicate key. Adding a duplicate key can cause errors.

To create an TLS/SSL keychain password, use the following steps:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**.
2. In the Physical Nodes page, select the node for which you want to create the Key Store password.
3. Go to **GigaSMART**>**Passive SSL** >**Key Store**.
4. Click **Actions**>**Keychain Password**.
5. Enter a password in the **Password** and **Confirm Password** fields.

   You can only configure a strong password. A strong password should include at least eight (8) or more characters (up to 64) and include the following:

   - one uppercase letter
   - one lowercase letter
   - one numerical character
   - one special character
6. Enable the **Auto login** check box to let GigaVUE-FM unlock the key store when the node reboots.
7. Click **Submit**.

After keys are installed on the node, you will be prompted to enter the password after any login as well as after a node reboot, for example:

If you are a user who does not have an admin level access role, when you enter the configure terminal mode, the following message is displayed:

Password required. Please contact administrator.

If you are a user with an admin level access role, but you enter an incorrect password, the following message is displayed:

Password does not match. Please reenter the password

If an TLS/SSL keychain password is lost, it can be reset, but all existing private keys will be revoked. When there are keys installed on the node, a warning is displayed before you are prompted for the new password.

Once you have a new password, you will have to upload the keys again.

Work with Keys and Services

This section describes working with private keys as well as services. Keys must be uploaded to the GigaVUE H Series node using a unique alias. Services must be defined for each server destination that needs decryption. There are two types of decryption. If the service is associated with a certificate then it is called service based decryption. If only certificate is added to the node without creating service then it is key certificate based decryption.

To perform the configuration in the following section, you must have an admin level access role.

Encrypted private keys are saved on the node. Refer to Notes about Private Keys and Passwords.

NOTE: When uploading TLS/SSL keys, make sure that you are not creating a duplicate key. Adding a duplicate key can cause errors.

Set Up Key Store Certificate Management for Passive TLS/SSL

A Key Store certificate can be setup to be auto- enabled, auto-deleted and auto-retained to a passive TLS/SSL service. The configuration can be done as follows:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**.
2. In the Physical Nodes page, select the node for which you want to configure Key Store settings.
3. Go to **GigaSMART**>**Passive SSL** >**Key Store**.
4. Click **Settings**. Configure the following settings:
    - **Auto Enable New Certificates** – This setting is not applicable for passive TLS/SSL service.
    - **Auto Delete Expired Certificates** – This option allows you to delete the expired certificates automatically. This setting will be triggered once a day at 12:00:00 UTC. Specify the number of days to retain an expired certificate in the **Number of days to retain expired certificates** field. The default value would be 30 days.
    - **Auto Delete Certificates with same entity** – This option allows you to automatically delete expired certificates that have a similar name and are associated with a passive TLS/SSL service. When you enable this option, you need to specify the maximum number of certificates to retain for the same entity in the

corresponding field. This means that if there are more certificates than the specified number, the oldest one will be deleted. This helps you manage your certificates and avoid cluttering your system with unnecessary or redundant certificates.

Upload TLS/SSL Private Keys

To upload an TLS/SSL private key, do the following:

1. From the device view, select **GigaSMART >Passive SSL > Key Store** to open the Key Store page.
2. Click **New**. The SSL Key page appears.
3. In the **Add Key** page, enter the following details:
   - For **Key Alias**, enter an alias for the SSL key.
   - For **Key Upload Type**, select **PEM**, **PKCS12**,**PKCS11,Luna-HSM**.
   - (optional) For **Passphrase**, enter a passphrase for the key.
   - Select a **Private Key** by pasting the copied key in PEM format or installing from URL or installing from local directory.
   - Select a **Certificate** by pasting the copied key in PEM format or installing from URL or installing from local directory.

   NOTE:
   - You can either use the server based private key or the key certificate based private key to decrypt a passive TLS/SSL traffic.
   - If the Thales-Luna Network HSM for TLS/SSL Decryption for Out-of-Band Tools(Passive) is enabled in TLS/SSL Service, you cannot use the server certificate to decrypt Passive TLS/SSL traffic.
4. Click **OK**.

If you choose to use Venafi, Inc. as your Electronic Key Management System (EKMS), ensure that you have the GigaVUE-FM host address and credentials to push the generated key into the selected nodes. For details on pushing keys from Venafi, see the GigaVUE-FM/Venafi Trust Protection Platform Integration Guide.

Delete TLS/SSL Key

To delete a particular SSL private key, select the key on the TLS/SSL Keys page, and then select Delete. To delete all TLS/SSL private keys, select multiple keys.

Display Key Store Details for Passive TLS/SSL

The key store certificates added would be displayed in the Key Store page.

To access the Key Store page:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**.
2. Select the node for which you want to view the key store certificate information.
3. From the left navigation pane, go to **System > GigaSMART > Passive SSL > Key Store**. The details about the key store certificates added for the selected node is displayed.

The following table describes the fields:

Component	Description
Key Alias	The alias name of the Key certificate.
Type	Defines whether the Key Store is a Certificate or a Private Key .
Health Status	The health indicator of a certificate used in traffic flow. The three major indicators with their respective color legend are as follows: Green - The key certificate is attached to a passive TLS/SSL service and the service is part of the passive TLS/SSL GSOP, which is used in a traffic map. This will also indicate a certificate which is being used as a signing CA in outbound deployment . Blue - The key certificate is not actively participating in any traffic flow. Red - The key certificate has expired .
Common name	A common name given to group the key based on domain.
Organization	Organization name that provided the key .
.**Organization Unit**	Organization unit name that provided the key.
Expiry Date	Date on which the key certificate would get expired.
Installed On	Date on which the key certificate was installed.
Description	Description or additional information about the key certificate.
Status	Status of the key certificate. The valid values are: Expiring—The key certificate is nearing the expiry date. Expired—The key certificate has expired.

You can control the key store certificates display by utilizing the filters provided.

Create TLS/SSL Service

After you have uploaded a private key, you can add a service. A service maps to a physical server, such as an HTTP server. One server can run multiple services. A service is a combination of an IP address and a server port number. Also, the key and the service must be tied together.

Prerequisites

Before creating a service, you must do the following:

- Upload a private key as described in Upload TLS/SSL Private Keys
- Create GigaSMART Group with TLS/SSL Decryption enabled.

To create a service, do the following:

1. From the device view, select **GigaSMART > Passive SSL > SSL Services**. The SSL Services page appears.
2. Click **New**.
3. On the Create SSL Service configuration page, do the following:
    - Enter an alias.
    - Click on the **Default Service** (inbuilt key will be used for decryption for any service) check box to allow GigaVUE-FM to enable default service.

NOTE: When opting for the Default service along with service based decryption, other than the particular service that is opted for service based, other connections will be decrypted using the key available on the default service.

    - Enter the information for the service: IP Address, Server Port.
    - Select the alias of SSL Key previously uploaded. For the steps, refer to Upload TLS/SSL Private Keys.
    - Select the GigaSMART Group with TLS/SSL decryption enabled to associate with this TLS/SSL service.
4. Click **Apply**.

Delete TLS/SSL Service

To delete a particular TLS/SSL service select the service on the TLS/SSL Services page, and then select Delete. To delete all TLS/SSL services, select multiple keys.

Notes about Private Keys and Passwords

Consider the following notes about private keys and passwords:

- Encrypted private keys are stored on the node. When a private key is uploaded, it is encrypted with a password before it is stored, therefore keys are password-protected. Keychain passwords are not stored on the node.
- Because only encrypted private keys are stored on the node and because the keychain password is not stored on the node, after any node reboot you will be prompted to enter the password. Until the password is entered, Passive TLS/SSL decryption is not working.
- Key content cannot be displayed.
- Keys that are synchronized across a cluster are encrypted.

ECODES for Troubleshooting Passive TLS/SSL Decryption

Use the following table of ECODE messages to assist with troubleshooting Passive TLS/SSL decryption:

ECODE	Description
81	TCP flow errors detected. Make sure you see the complete TCP flow. Use the de-duplication GigaSMART operation with Passive TLS/SSL decryption.
103	Session limit reached. The session table has been exhausted. If the session timeout (**session-timeout**) value is large, lower it.
104	Key/ticket cache limit reached. The allocated cache entries have been used up. If the timeout (**key-cache-timeout** or **ticket-cache-timeout**) value is large, lower it.
206	No server info. A flow has been received for which service-key mapping is not defined.
213	Packets for missed TCP handshake. Packets were received for TCP flows that do not exist. If the device was just started, this should trend down quickly.
218	Unknown TLS/SSL version. An TLS/SSL handshake processing error occurred. Use the de-duplication GigaSMART operation with Passive TLS/SSL decryption.
221	Unknown TLS/SSL version. An unsupported TLS/SSLv2 handshake was seen.
222	Protocol error. An unsupported protocol version was seen.
225	Unsupported cipher. The cipher suite cannot be decrypted.
226	Pre-master secret error. Check that the private key is correct and that the session is complete.
228	Generic decryption error. Usually indicates errors in the handshake. Check that you are getting the full session from both sides.
231	Invalid MAC. Likely indicates that invalid or truncated packets have been received.
232	Session not in cache. Indicates that you are trying to decrypt a restarted session where the

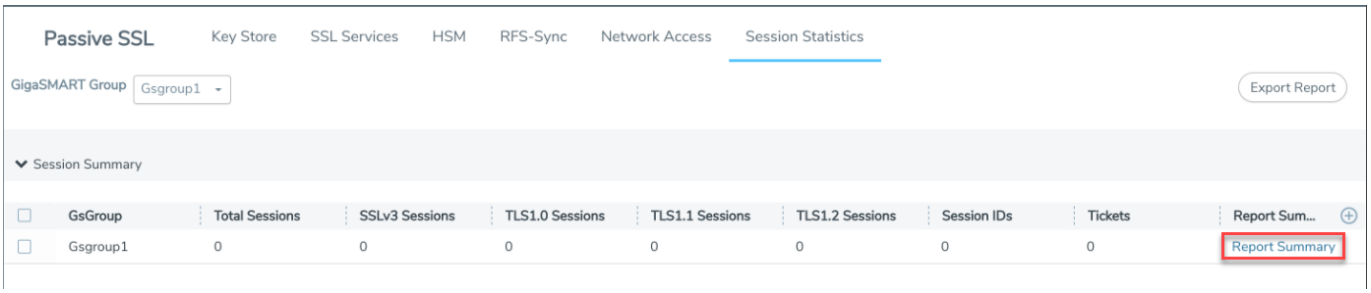
ECODE	Description
	original negotiation was not seen. These should trend down in time, but if they do not, increase the **key-cache-timeout** value.
237	Cannot decrypt ephemeral key based encryption. One of the Ephemeral/PFS ciphersuites, usually Diffie-Hellman Ephemeral, has been seen. These are not supported.
245	Ticket not in cache. This is usually not an error. Indicates that you are trying to decrypt a restarted session where the original negotiation was not seen. These should trend down in time, but if they do not, increase the **ticket-cache-timeout** value.

View Passive TLS/SSL Decryption Flow Ops Report

GigaSMART provides support for Flow Ops reporting. You can generate the Flow Ops report for Passive TLS/SSL Decryption, view the session summary and session statistics, and export the report to a remote server.

To view the session summary and session details of the Passive TLS/SSL Decryption Flow Ops report:

1. From the device view, go to **System > GigaSMART > Passive SSL > Session Statistics**.
2. From the **GigaSMART Group** drop-down list, select the required GigaSMART group. The Session Summary and Session Details appear. For descriptions of the session statistics, refer to Flow Ops Report Statistics for Passive TLS/SSL Decryption.
3. In the Session Summary table, click the **Report Summary** link to view the graphical representation of the trending data for the Passive SSL Decryption session summary as shown in the following figures:





You can also choose to export the session summary to a remote server. Select the required row in the Session Summary table, and then click **Export Report**. The Upload Flow Ops Report page appears. Enter the remote server path and password to access the server. Click **Upload**. The details are exported to the remote server.

View Certificate Statistics

To view certificate statistics, from the device view, go to **System > GigaSMART > Passive SSL > Certificate Statistics**. The page displays all the SSL server-certificates that are attached to the GigaSMART group.

View Service Statistics

To view service statistics, from the device view, go to **System > GigaSMART > Passive SSL > Service Statistics**. The page displays all the SSL services that are associated with the GigaSMART group.

View Error Statistics

To view error statistics from the device view, go to **System > GigaSMART > Passive SSL > Error Statistics**. The page displays all the SSL error code messages that were received for the device.

Configuring Passive TLS/SSL Decryption Examples - Command Line Reference

The following sections provide examples and commands (CLI) of TLS/SSL decryption. Refer to the following:

- Example 1: TLS/SSL Decryption with a Regular Map
- Example 2: TLS/SSL Decryption with De-duplication
- Other Usage Examples

Example 1: TLS/SSL Decryption with a Regular Map

In Example 1, a regular map is configured to use with the TLS/SSL decryption GigaSMART operation.

Step	Description	Command
1.	Upload a key and create a service. Refer to *Working with Keys and Services* on page 609.	**(config) # apps ssl key alias key1 download type private-key url https://keyserver.domain.com/path/keyfile.pem** **(config) # apps ssl service alias service1 server-ip 192.168.1.1 server-port 443**
2.	Configure a GigaSMART group..	**(config) # gsgroup alias gsgrp1 port-list 1/1/e1**
3.	Specify the GigaSMART group alias.	**(config) # gsparams gsgroup gsgrp1**
4.	Specify a failover action.	**(config gsparams gsgroup gsgrp1) # ssl-decrypt decrypt-fail-action drop**
5.	Configure session timeouts, in seconds.	**(config gsparams gsgroup gsgrp1) # ssl-decrypt pending-session-timeout 60** **(config gsparams gsgroup gsgrp1) # ssl-decrypt session-timeout 300** **(config gsparams gsgroup gsgrp1) # ssl-decrypt tcp-syn-timeout 20**
6.	Configure cache timeouts, in seconds.	**(config gsparams gsgroup gsgrp1) # ssl-decrypt key-cache-timeout 9000** **(config gsparams gsgroup gsgrp1) # ssl-decrypt ticket-cache-timeout 9000**
7.	Configure a key/service mapping that maps how a key is assigned to an IP address of a server.	**(config gsparams gsgroup gsgrp1) # ssl-decrypt key-map add service service1 key key1**
8.	Enable TLS/SSL decryption.	**(config gsparams gsgroup gsgrp1) # ssl-decrypt**

Step	Description	Command
		enable
9.	Exit the GigaSMART group configuration mode.	**(config gsparams gsgroup gsgrp1) # exit** **(config) #**
10.	Configure a GigaSMART operation for TLS/SSL decryption.	**(config) # gsop alias gdssl1 ssl-decrypt in-port any out-port auto port-list gsgrp1**

In the previous step, **gdssl1** is the alias for a GigaSMART operation, **in-port** specifies the destination port on which to listen, **out-port** specifies the destination port on which to send decrypted traffic, and **port-list** is set to the GigaSMART group alias previously configured. The **in-port** and **out-port** arguments can also be a port number between 1 and 65535.

Next, configure a traffic map, as follows:

Step	Description	Command
1.	Specify a map alias (m1) and specify the map type and subtype.	**(config) # map alias m1** **(config map alias m1) # type regular byRule**
2.	Specify the GigaSMART operation alias (gdssl1) as part of the map. This applies the associated GigaSMART functionality to packets matching a rule in the map.	**(config map alias m1) # use gsop gdssl1**
3.	Specify a map rule.	**(config map alias m1) # rule add pass ipver 4**
4.	Specify the destination for packets matching this map.	**(config map alias m1) # to 1/1/g2**
5.	Specify the source port(s) for this map.	**(config map alias m1) # from 1/1/g1**
6.	Exit the map prefix mode.	**(config map alias m1) # exit** **(config) #**
7.	Display the configuration.	**(config) # show gsop** **(config) # show map** **(config) # show gsparams**

Example 2: TLS/SSL Decryption with De-duplication

In Example 2, the configuration steps are the same except when you configure a GigaSMART operation you send the decrypted traffic to de-duplication for additional filtering, as follows:

```
(config) # gsop alias gdssl1 ssl-decrypt in-port any out-port auto dedup set port-list gsgrp1
```

Other Usage Examples

Two typical usage examples are as follows:

- Use map rules to filter on the IP address of the server and send everything to GigaSMART. Configure a GigaSMART operation to listen on the **in-port** used by the server. The GigaSMART will drop other traffic.
- Use map rules to filter on the IP address of the server and **in-port** and send specific port traffic to the GigaSMART. Configure a GigaSMART operation to listen on **in-port any**.

Thales-Luna Network HSM for TLS/SSL Decryption for Out-of-Band Tools(Passive)

Required License: Included with TLS/SSL Decryption for Out-of-Band Tools (Passive)

Hardware Security Modules (HSMs) are specialized systems that logically and physically safeguard cryptographic operations and cryptographic keys. HSMs protect sensitive data from being stolen by providing a highly secure operation structure. HSMs are comprehensive, self-contained solutions for cryptographic processing, key generation, and key storage. The hardware and firmware (i.e., software) required for these functions are automatically included in these appliances.

The application could be a web server or a database server, but, in the case of TLS/SSL decryption for out-of-band tools, the application is GigaSMART. The application interfaces with HSM to use the keys that are stored. There must be network connectivity between HSM and the application.

Keys are added to the HSM by an administrator. When an application's key is on HSM, the HSM creates an application key token. The key token is sent to the application. When the application wants to use a key, the application sends the token to HSM, which establishes a session with HSM to use the key. In this way, the use of keys by the application is secure because only key tokens are exchanged.

You can use Remote File System (RFS), a component in the HSM to store and manage encrypted keys. The RFS helps to automate the key distribution process. You can enable RFS on the GigaVUE-OS device using GigaVUE-FM so that the device can access the encrypted keys stored in RFS. You can synchronize RFS with GigaVUE-OS device to perform a bulk download of the encrypted keys.

Starting from software version 6.8, decryption for Out-of band Tools (Passive SSL/TLS Decryption) is also enhanced to include the Thales-Luna Network HSM support.

Thales-Luna HSM for TLS/SSL Decryption for Out-of-Band Tools—Rules and Notes

Keep in mind the following rules and notes before you configure and use HSM to store and manage keys:

- GigaSMART uses keys that are already stored on the HSM. There is no key generation.
- The key token that is uploaded to GigaSMART can only be in PKCS11 format.
- Only RSA keys (private keys) are supported.
- When TLS/SSL Decryption for Out-of-Band Tools is configured with HSM, only one map and GigaSMART operations will function. If you configure with more than one map will result in undefined behavior.
- The network connectivity between the HSM and GigaSMART must use a static IP address. Do not use DHCP because the IP address needs to remain the same.

  NOTE: If the GigaSMART® engine is configured using DHCP, the following issues may arise:
  1. Whenever a new DHCP IP is assigned to the GigaSMART® engine, the user must delete and re-create the ISSL App and deploy the solution.
  2. Additionally, the user needs to register the new DHCP IP with the HSM server for client use.

- Only IPv4 addresses are supported.
- Each GigaSMART card that interfaces with the HSM will use one license.
- Clustering is not supported.
- Increase the HSM timeout to 5000ms when using 4K size keys for decryption.
- When uploading RSA keys, validity check for protocol mismatch cannot be performed since the private keys are available on the HSM server.
- If a HSM Decryption deployment is modified follow the below steps:
  - Move the Inline Network traffic path to bypass mode.
  - Make the desired deployment change such as:
    - From non-HSM based decryption to Thales-Luna HSM based decryption.
    - From HSM based decryption to non-HSM based decryption.
  - Reboot the GigaSMART card.
  - Move the Inline Network out of bypass mode to ‘To inline Tool’ mode.

- There should be at least one active Luna HSM in the High Availability to ensure that the decryption is not interrupted.
- Do not attempt to reload the device, if a partition fails in the Luna HSM configuration and at least one active partition is present.

Configure HSM for TLS/SSL Decryption for Out-of-Band Tools

This section provides topics on how to configure and use HSM for TLS/SSL decryption for out-of-band tools:

Topics:

- Configure HSM Group
- Configure Set Key Handler
- Configure Passive TLS/SSL Network Access
- Use RFS to Manage Encrypted Keys
- Thales-Luna Network HSM for TLS/SSL Decryption for Out-of-Band Tools(Passive)
- Create a GigaSMART Operation (GSOP)
- Create TLS/SSL Keychain Password
- Upload TLS/SSL Private Keys
- Configure TLS/SSL Service
- Configure Maps

Configure HSM Group

To configure an HSM group ,add at least one HSM appliance by specifying an alias, a static IP address, and port number. Obtain the ESN and KNETI from your HSM administrator.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To add an HSM appliance, do the following:

1. From the device view, go to **GigaSMART**> **Passive SSL** > **HSM Groups**.
2. Click **New**.

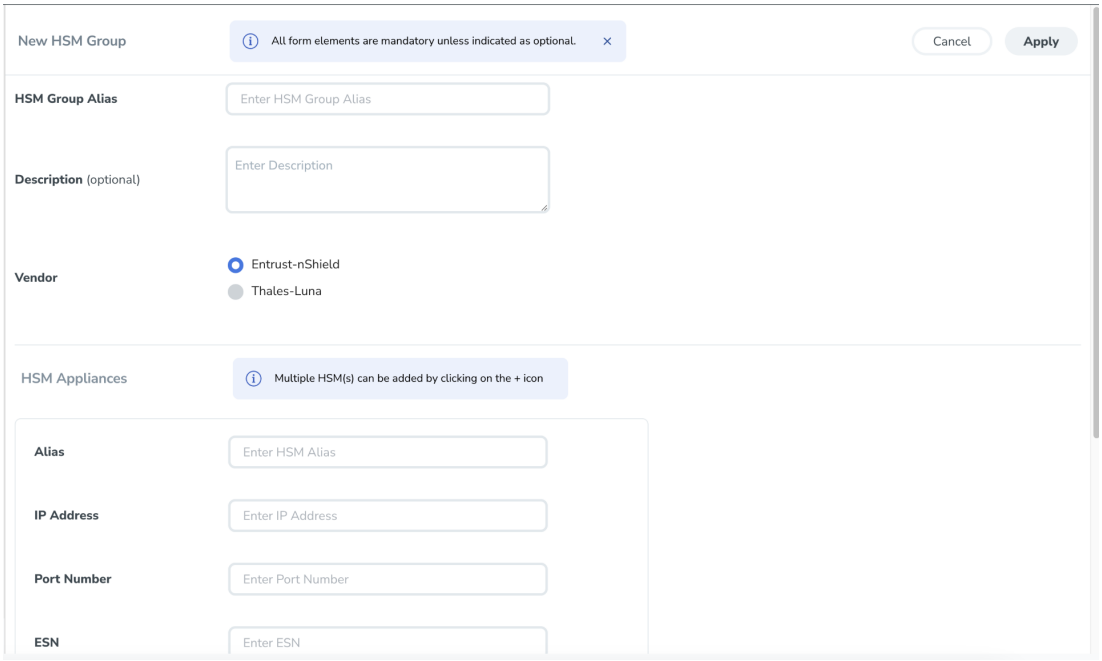


Figure 35 *Adding a New HSM Appliance*

3. Enter the details in the **HSM Group Alias** and **Description** fields under the **New HSM Group** section.
4. Select the required vendor type from the options **Thales-Luna**) to create the respective HSM Group.
5. Under the **HSM Appliances** section, in the **Alias** field, enter a name for the HSM appliance.
6. Enter a valid **IP address** and **Port Number**.

NOTE: For Thales-Luna HSM Group , the port group by default is 1792.

7. Configure the following fields for the vendor type **Thales-Luna**:
   a. Enter the valid username and password in the Server Username and Server Password fields.
   b. Enter the valid details in the Partition Label and Partition Password fields.

Click **Apply**.

NOTE: You cannot configure multiple HSM Groups for Passive SSL with HSM solution.

Configure Set Key Handler

1. From the device view, go to **GigaSMART > Passive SSL > HSM Groups**.

2. Select the **HSM appliance** you just created.
3. Click **Configure** from the **Actions** drop down.
4. Choose one of the following methods to install the key handler file:
   - **Install from URL**—Enter a valid directory path including the file name and enter the password to access the server.

     NOTE: SCP, SFTP, HTTP, FTP, and TFTP are the supported protocols from where you can select the key handler file.

   - **Install from Local Directory**—Browse and select the key handler file from your local directory.

NOTE: Ensure that the file name is "world"

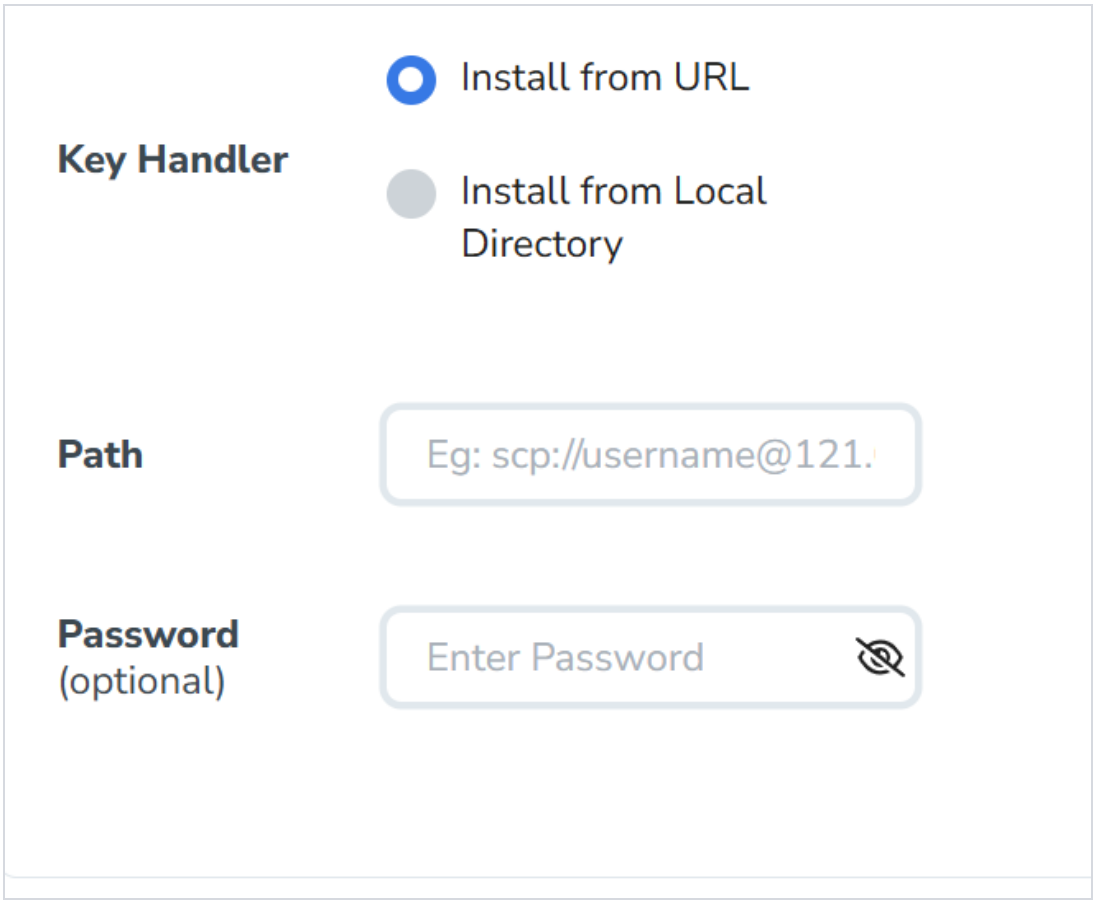


Figure 36 *HSM-Configure Key Handler*

5. Click **OK**.

NOTE: Configuring Key Handler is not applicable for a Thales-Luna HSM group.

Configure Passive TLS/SSL Network Access

Each GigaSMART card requires IP address configuration for network access. To configure IP address details:

1. From the device view, go to **GigaSMART > Passive SSL > Network Access.**
2. Select the **GigaSMART appliance.**
3. Click **Edit**.
4. Select **IP Address** from the **Network Access configuration** options.

   NOTE: For the Thales Luna HSM group, it is preferable to use a static IP address to prevent the Thales Luna registration from expiring.

   NOTE: If the IP address of the GigaSMART engine is changed, the GigaSMART engine needs a reboot to complete the HSM registration with the new IP address.
5. Enter **IP Address, Netmask, Gateway, DNS, MTU and VLAN** parameters.
6. Select the required management interface.
7. Under the **Ping Test** section, select the GigaSMART port and enter the **IP Address / Host Name** and the **Ping Test** parameters.

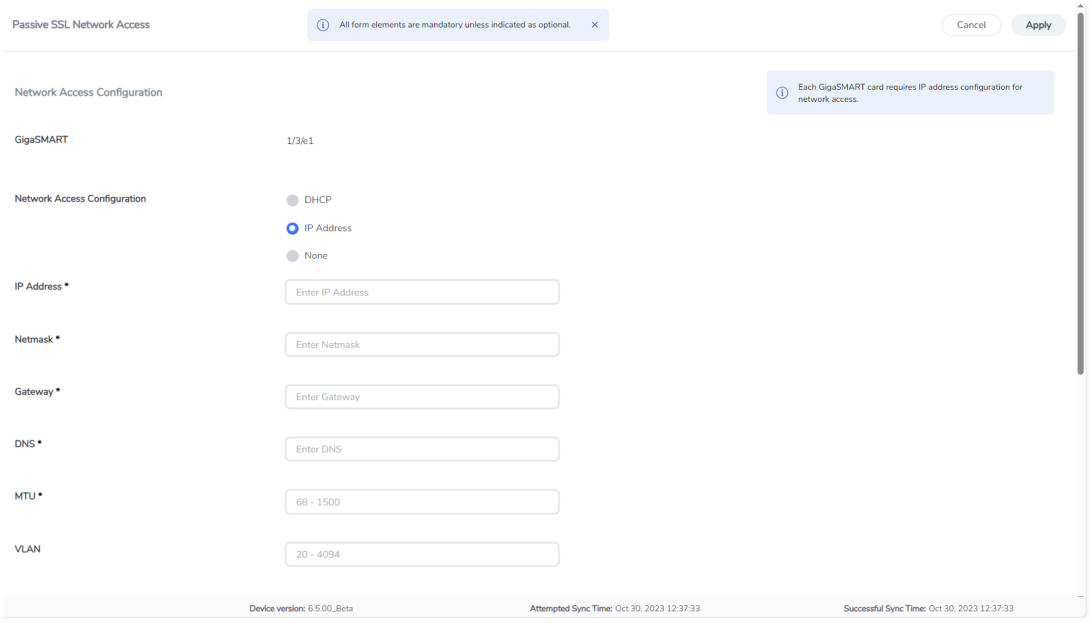


Figure 37 *Passive TLS/SSL Network Access - IP Configuration*

8. Click **Apply**.

Use RFS to Manage Encrypted Keys

Use Remote File System (RFS), a component in HSM to store and manage encrypted keys. The RFS helps to automate the key distribution process. You can enable RFS on the GigaVUE-OS device using GigaVUE-FM so that the device can access the encrypted keys stored in RFS. You can synchronize RFS with GigaVUE-OS device to perform a bulk download of the encrypted keys.

NOTE: Remote File System is not applicable for Thales-Luna HSM.

Refer to the following sections:

- Add RFS to GigaVUE-OS Device
- Map Encrypted Keys with Servers

Add RFS to GigaVUE-OS Device

To add and synchronize RFS to the GigaVUE-OS device:

1. From the device view, go to **GigaSMART > Passive SSL > RFS-Sync** .
2. Select your HSM Group.
3. From the RFS section, select the RFS got to **Actions>Sync** to automatically synchronize the RFS with the GigaVUE-OS device.
4. In the **Sync Period** field, enter the time interval for synchronization in hours.
5. Click **OK**
6.

The details of the RFS, such as the IP address, sync period, last sync time, next sync time, and the total number of keys stored and managed in the RFS appears in the RFS section. Click the **Show Details** link next to the **Total Keys** column to view the key token and key name mapping for the encrypted keys stored in RFS.

You can choose to modify the fields for the RFS you have added. Select the RFS click on **Actions > Edit** .

Click **Sync** to manually synchronize the RFS and the GigaVUE-OS device to fetch the latest encrypted keys from RFS at any point in time.

Map Encrypted Keys with Servers

A key name or a key token must be mapped to a server IP address. A total of up to 1000 key mappings is allowed per RFS. You can either manually map the keys with the servers or do a bulk key mapping using a text file. If a key mapping already exists, the new key mapping will be rejected. You must delete the existing key mapping to add the new mapping.

To map a key name or key token to a server IP address:

1. From the device view, go to **GigaSMART > Passive SSL > Key Mapping**.
2. In the Key Mapping section, click **Add**. The Add Key Mapping page appears.
3. Choose one of the following types to map the keys with the server IP address:
   - **Manual**—Manually map the key name or key token to the server IP address. You must keep adding the key mapping one at a time.
   - **From URL**—Create a text file with the key mappings and upload it to a server. Enter a valid directory path including the text file name and enter the password to access the server. It is recommended to use a secure protocol, such as SCP or HTTPS to access the URL.
   - **From Local Directory**—Create a text file with the key mappings and save the text file in your local directory. Browse and select the text file from your local directory.
4. Click **Submit**.

Configure a GigaSMART Group

Refer to the section GigaSMART Group for more details on GigaSMART Group.

To configure a GigaSMART group for passive TLS/SSL:

1. From the device view, go to **GigaSMART > GigaSMART Groups.**
2. Click **New**.
3. In the **Alias** field, enter a name for the GigaSMART group that you are creating for Passive TLS/SSL.
4. From the **Port List** drop-down list, select the required port you want to associate with this group.
5. Scroll down to the **GigaSMART Parameters > Passive SSL> HSM Group** section of the page, and then select the required HSM Group from the drop-down list.

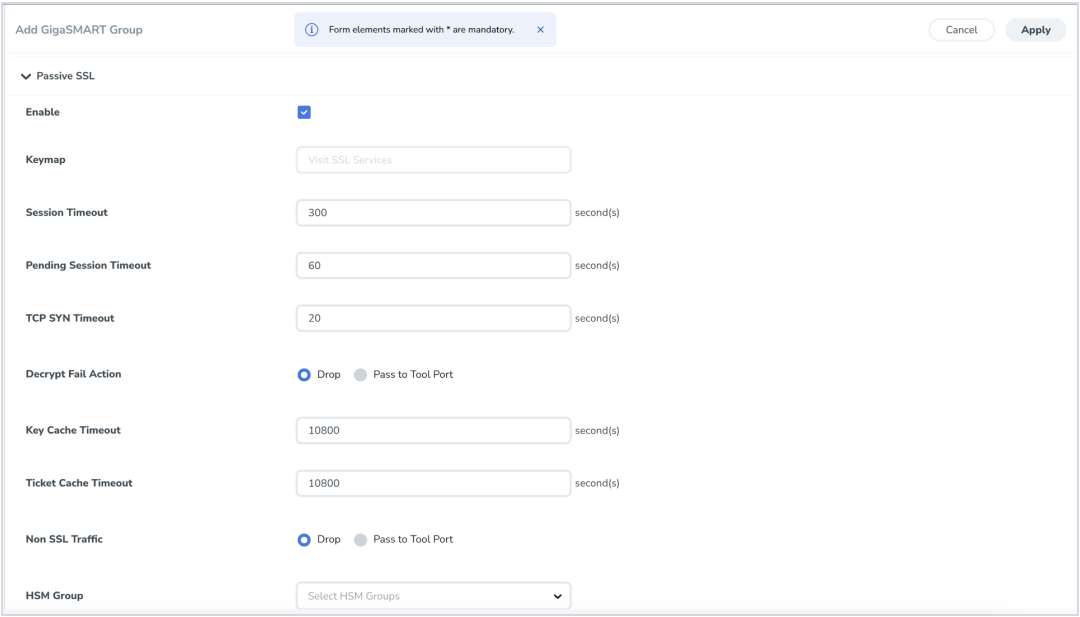


Figure 38 *GigaSMART Group Setup Page*

6. Click **OK**.

Create a GigaSMART Operation (GSOP)

Refer to the section GSOP for more details on GigaSMART Operation.

To create a GigaSMART operation with an TLS/SSL Decryption component:

1. From the device view, go to **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.
2. Click **New**.
3. In the **Alias** field, enter a name for the GigaSMART operation.
4. From the **GigaSMART Group** drop-down list, select the GigaSMART group that you have created for passive TLS/SSL.
5. From the **GigaSMART Operations (GSOP)** drop-down list, select **SSL Decryption**.

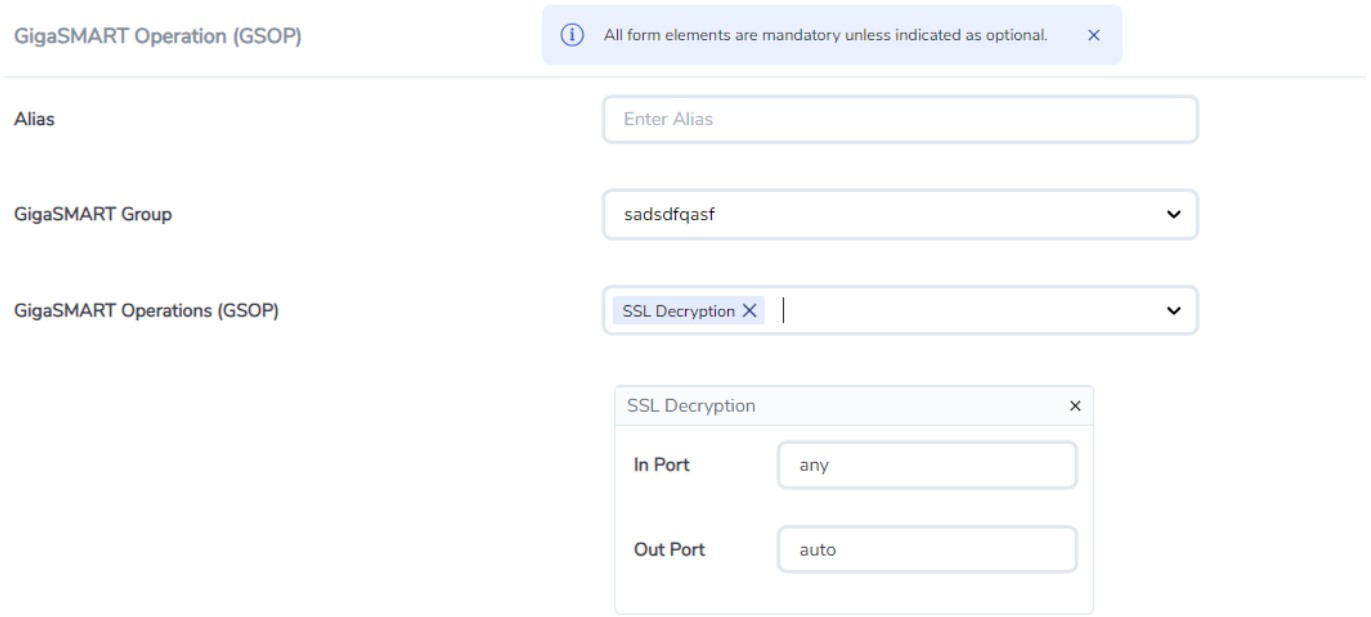


Figure 39 *GigaSMART Operations - Setup Page*

6. Click **Apply**.

Create TLS/SSL Keychain Password

Before uploading keys or configuring TLS/SSL, you must create an TLS/SSL keychain password. The password is used to encrypt the private keys that you upload to the node.

NOTE: When uploading TLS/SSL keys, make sure that you are not creating a duplicate key (adding same key with different key alias). Adding a duplicate key can cause errors.

To create an TLS/SSL keychain password:

1. From the device view, go to **GigaSMART > Passive SSL > Key Store.**
2. Click **Keychain Password** from the **Actions** drop down list.

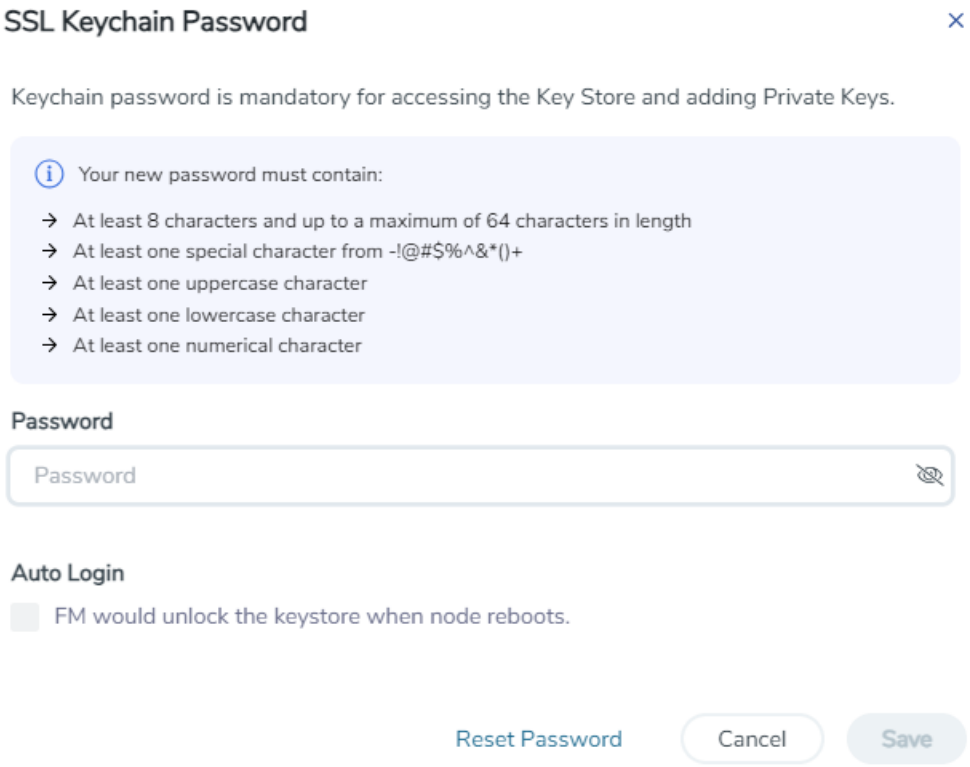


Figure 40 *SSL Keychain Password Setup Page*

3. In the **Password** field, enter a valid password. Ensure that the password meets the following specifications:
   - Password must be at least eight (8) characters in length.
   - Password must contain at least one:
     - uppercase character
     - lowercase character
     - numerical character
     - special character
4. Click on the **Auto Login** check box to allow GigaVUE to unlock the keystore when the node reboots.
5. Click **Save**.

Upload TLS/SSL Private Keys

To upload an TLS/SSL private key:

1. From the device view, go to **GigaSMART >Passive SSL > Key Store** to open the Key Store page.

2. Click **Add**. The SSL Key page appears.
3. In the **Alias** field, enter a name for the SSL key.
4. Select the **Key Upload Type** from the options **PEM, PKCS12, Entrust nShield**, **Luna-HSM**.
   a. Passphrase- SSH passphrases allows you to protect your private key from being used without the passphrase. Enter the passphrase created with the private key.
   b. Private key- Enter the Private Key using any of the following options:
      - Copy and Paste
      - Install from URL
      - Install from Local Directory
   c. Certificate- Enter the Certificate using any of the following options:
      - Copy and Paste
      - Install from URL
      - Install from Local Directory
   d. Key Label - The Key label that is imported from the HSM server. This field is applicable when you configure Thales-Luna HSM.

   NOTE: Ensure that Luna key labels match the correct certificate and decryption keys when using the default service in Passive SSL.

   e. Path and Password- Configure the file path and password if you select 'Install form URL', or else choose the file if you select 'Install form Local directory'. The supported protocols are **HTTP, HTTPS, FTP, TFTP, SCP, and SFTP**.
5. Click **Save**.

Configure TLS/SSL Service

After you have uploaded a private key, you can add a service. A service maps to a physical server, such as an HTTP server. One server can run multiple services. A service is a combination of an IP address and a server port number. Also, the key and the service must be tied together.

To create an TLS/SSL service:

1. From the device view, go to **GigaSMART > Passive SSL > SSL Services.**
2. Click **New**. The SSL Service page appears.

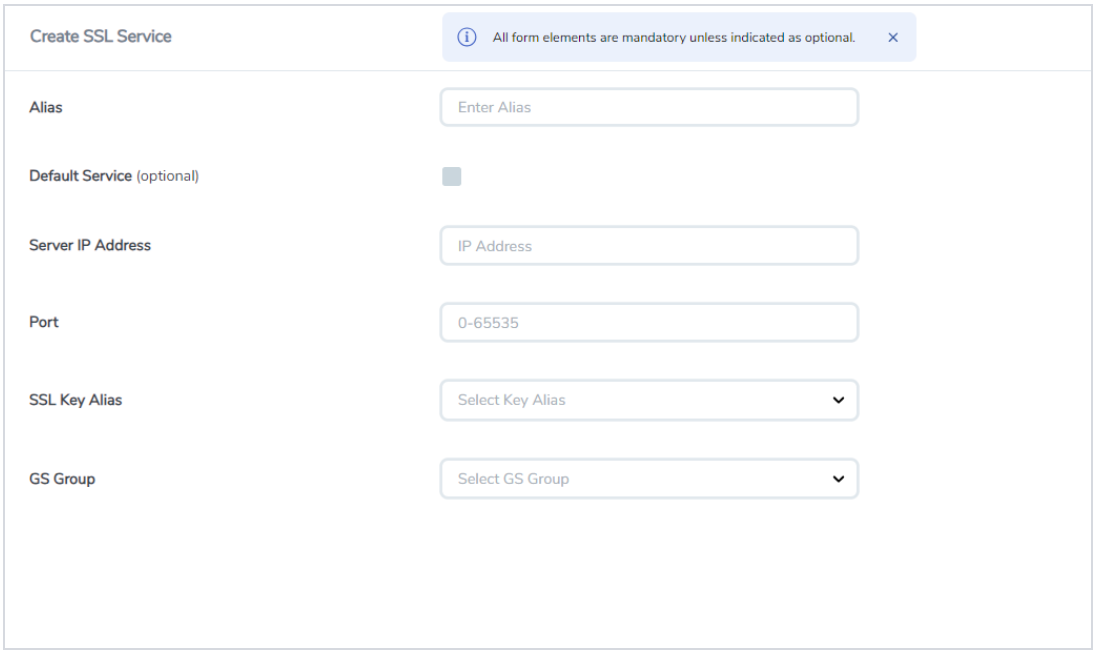


Figure 41 *SSL Service*

3. In the **Alias** field, enter a name for the SSL service.
4. Map the TLS/SSL service to a server IP address and a server port using one of the following methods:
   - Select the **Enabled** check box next to the **Default Service** field to dynamically map the server IP address and server port.

     NOTE: If you select the **Enabled** check box, the **Server IP Address** and **Server Port** fields are disabled.

   - In the **Server IP Address** and **Server Port** fields, enter an IP address and port to which you want to map the TLS/SSL service.
5. From the **SSL Key Alias** drop-down list, select the name of the SSL Key previously uploaded.
6. From the **GS Group** drop-down list, select the GigaSMART group with TLS/SSL decryption enabled to associate with this TLS/SSL service.
7. Click **Apply**.

Configure Maps

1. From the device view, go to **Traffic> Maps > Maps.**
2. Click **New**.
3. Configure the map.

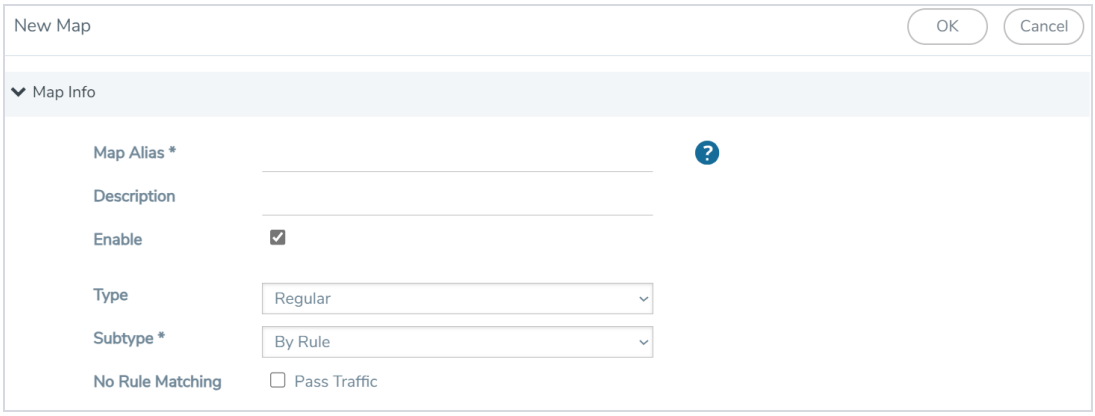


Figure 42 *Create New Map*

- Type **map11** in the Alias field.
- Select Regular for **Type**.
- Select **ByRule** for **Subtype**.
- Select the network port for the Source.
- Select **Tool port/Hybrid port** for Destination.

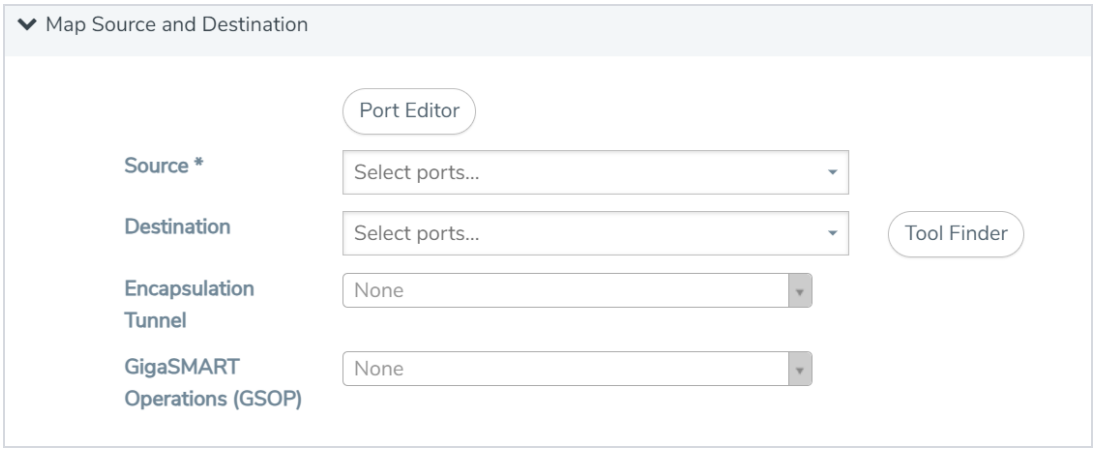


Figure 43 *Configure Map Details*

4. Add **Configuration & Rules**.

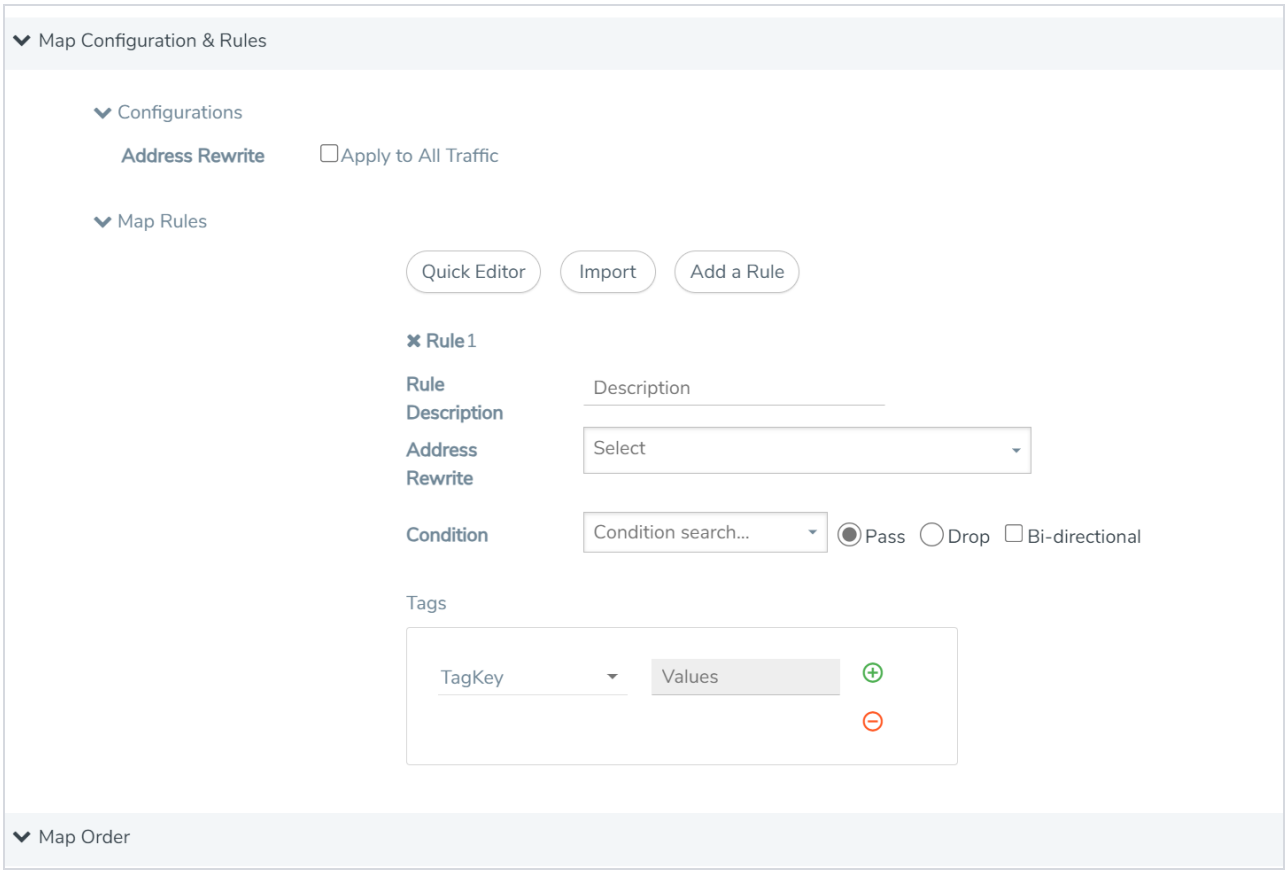


Figure 44 *Figure 20-123: Map Details - Create Rule*

a. Under Configurations, select the **Apply to All Traffic** check box to rewrite the address.

b. Select the address from the drop-down list.

c. Under Map Rules, Click **Add a Rule**.

d. Select **Pass** as condition.

e. Select **IPv4 Version** and set **Version to v4.**

5. Click **OK**.

To view the configured Thales-Luna HSM Group status, select the **Thales Luna profile> Action> Diagnostic.** You can view the statistics of the below components:

1. Ping Result
2. Verify
3. High Availability
4. Luna Key Label

Traffic Intelligence

Designed to serve a wide variety of network operations environments, these GigaSMART operations provide complete network visibility to increase the efficiency of network performance tools.

Traffic Intelligence	
Adaptive Packet Filtering	▪ GigaSMART Adaptive Packet Filtering (APF)
Advanced Load Balancing	▪ GigaSMART Load Balancing
De-duplication	▪ GigaSMART De-Duplication
Flow Masking	▪ GigaSMART Traffic Performance Enhancement
Header Stripping	▪ GigaSMART Header Addition ▪ GigaSMART Header Stripping
Masking	▪ GigaSMART Masking
NetFlow Generation	▪ GigaSMART NetFlow Generation
Slicing	▪ GigaSMART Advanced Flow Slicing ▪ GigaSMART Packet Slicing
PCAPng Application	• PCAPng Application
Tunneling	▪ GigaSMART Custom Tunnel Decapsulation ▪ GigaSMART ERSPAN Tunnel Decapsulation ▪ GigaSMART IP Encapsulation (GigaSMART Tunnel) ▪ GigaSMART IP Encapsulation/Decapsulation (GigaSMART Tunnel) ▪ GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation ▪ IP Fragmentation and Reassembly on L2GRE and GMIP Tunnels ▪ GigaSMART VXLAN Tunnel Decapsulation ▪ GigaSMART TCP tunnel ▪ GRE-In-UDP Tunnel Decapsulation

GigaSMART Adaptive Packet Filtering (APF)

Required License: Adaptive Packet Filtering

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus, and GigaVUE-HCT.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

Feature Overview

Adaptive Packet Filtering (APF) provides filtering on specific encapsulation protocol parameters. Additionally, it has the ability to look beyond the encapsulation protocol parameters into the original (encapsulated) data packet, to filter on source and destination IP or Layer 4 port numbers. APF offers the ability to look for content anywhere in the data packet and make intelligent filtering and forwarding decisions.

APF filters packet-by-packet, but does not have the concept of sessions. For Application Session Filtering (ASF) and packet buffering on ASF, refer to Application Session Filtering with Buffering. Adaptive Packet Filtering includes fragmentation awareness whereby all IP fragments associated with the filtered data packet are always forwarded allowing a complete view of the traffic stream for accurate analytics. This is applicable when APF is combined with ASF.

APF operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to Groups of GigaSMART Engine Ports for details.

In APF second level maps, a maximum of twenty-five (25) maps can be attached to a virtual port (vPort). Each map can contain up to 25 GigaSMART rules. The higher map limit applies only when the vPort is configured solely with APF (no other GigaSMART applications or operations are enabled); if any other GigaSMART application or operation is enabled on the vPort, the per vPort limit for APF second level maps remains 5.

NOTE: When you configure a large number of APF second level maps per vPort (for example, approaching the 25-map limit), additional APF processing can impact overall system performance.

Adaptive Packet Filtering (APF) goes deeper into packets to search for a condition, then filter and forward packets to tools, as follows:

- *Content-based Filtering*
- *Encapsulation Awareness*
- *Pattern Matching on Gen 2 GigaSMART modules*
- *Pattern Matching on Gen 3 GigaSMART modules*

Implement APF Through the UI

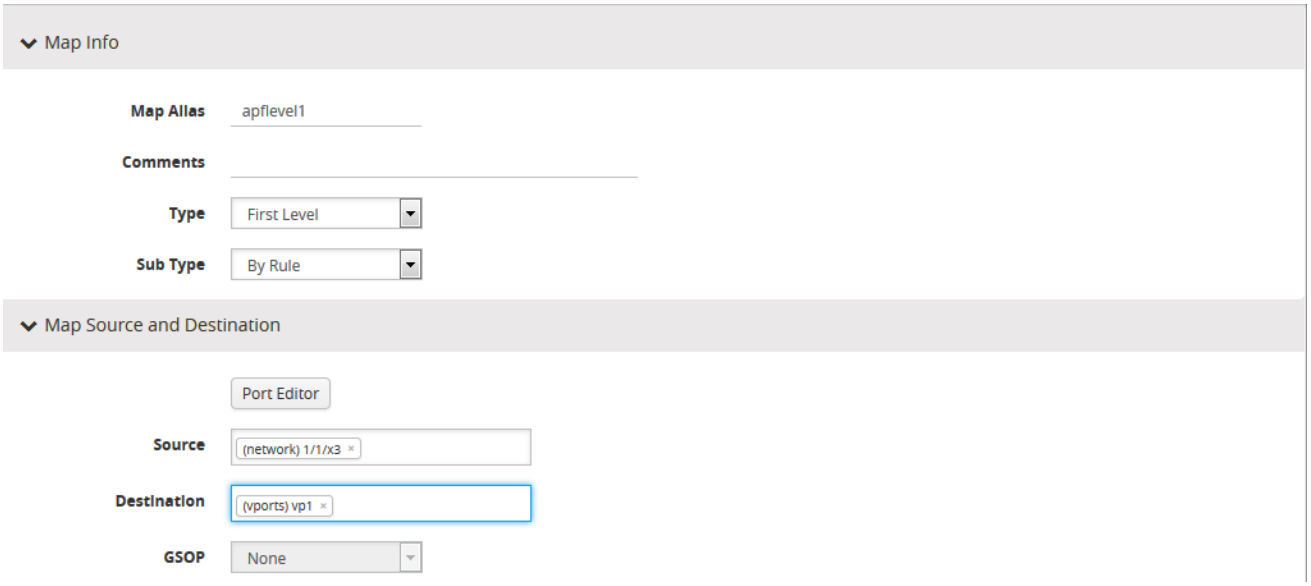
To create vports through the UI and implement APF, do the following:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**.
2. From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART Groups**, and then click **New**.
3. On the GigaSMART Group page, select an available engine ports in the Port List field to associate group with one of the available engine ports.
   You can associate the GigaSMART Group with one or multiple eports. For APF, no GigaSMART parameters are required unless combined with other gsops.
4. From the device view, select **GigaSMART > Virtual Ports**, and then click **New**.
5. On the Virtual Ports page, enter an alias and select the GigaSMART groups created in Step1, and then click **Save**.
6. To enable the APF operation, do the following:

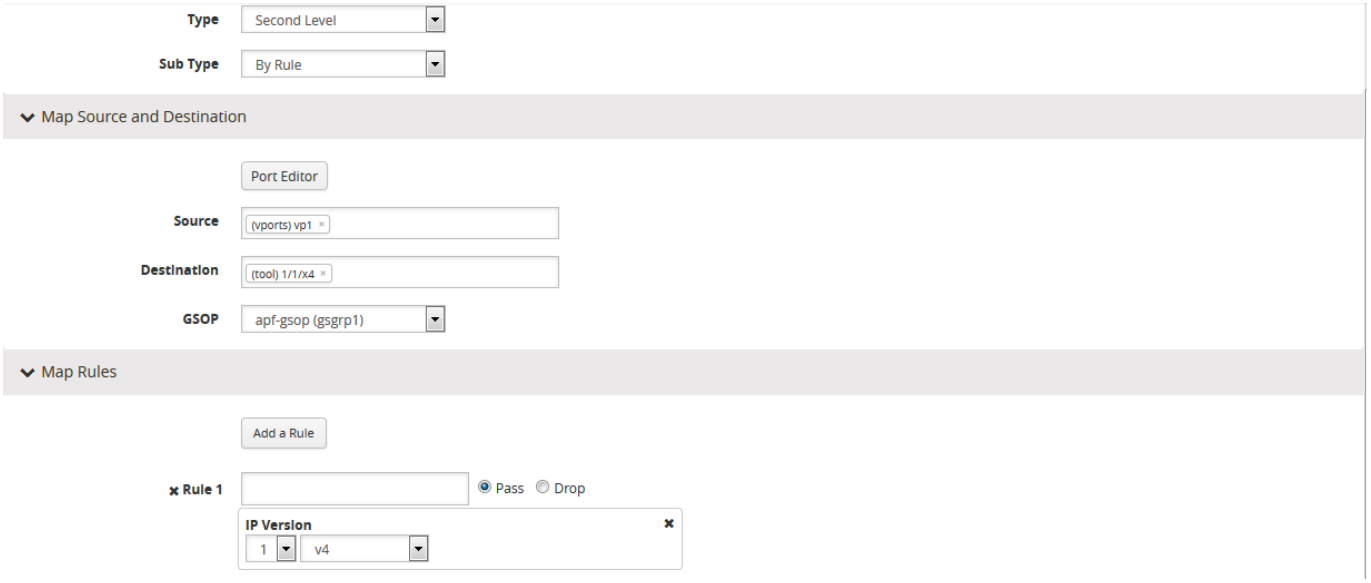
a. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.
b. On the GigaSMART Operations page, enter an alias in the **Alias** field
c. In the **GigaSMART Groups** drop-down list, select the GigaSMART group from step 1.
d. From the **GigaSMART Operations (GSOP)** list, select **Adaptive Packet Filtering** and select **Enabled**.
e. Click **Save**.

Once APF is enabled, maps can be created that the APF and the virtual port.

6. Create the first level map with virtual port created in step 2 as the destination and without applying a GigaSMART Operation.



7. Create a second level map with the APF GigaSMART operation, the virtual port as the source, and a rule. The following figure shows an example.



This completes the process to create an APF GigaSMART operation and corresponding rules. To learn more about the rules applicable for APF, see following sections.

Content-based Filtering

Content-based filtering is based on packet contents beyond Layer 2, 3, and 4 headers. The following four groups of attributes of rules in a map support content-based filtering.

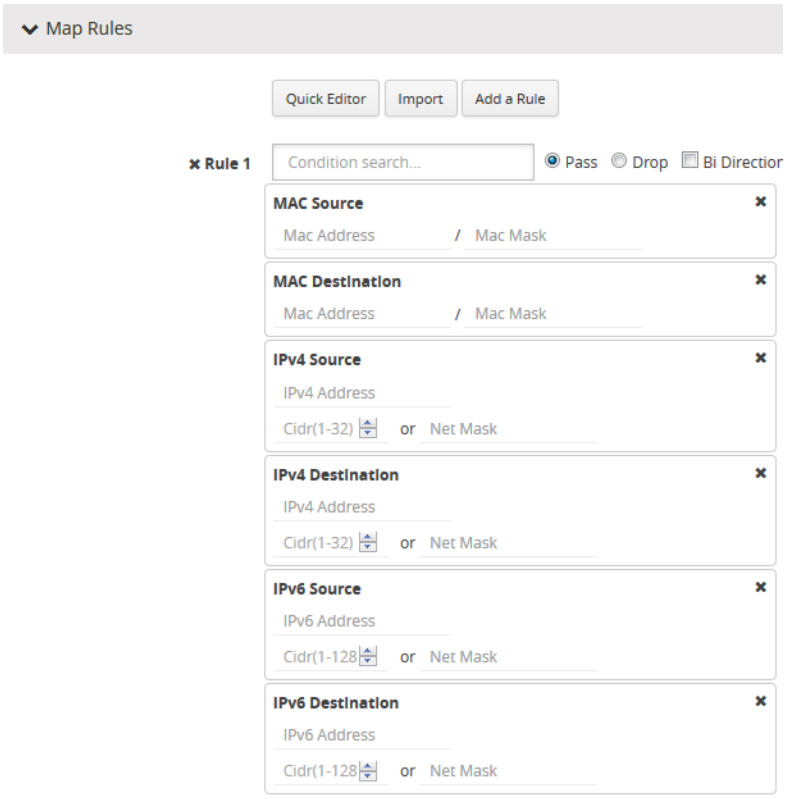
The first group of attributes has the following format:

```
<attribute> <address> <cidr>|<mask>
```

The first group of attributes that use this format are as follows:

- ipv4 src and dst
- ipv6 src and dst
- mac src and dst

The following figure shows the attributes as displayed in the UI.



The second group of attributes has the following format:

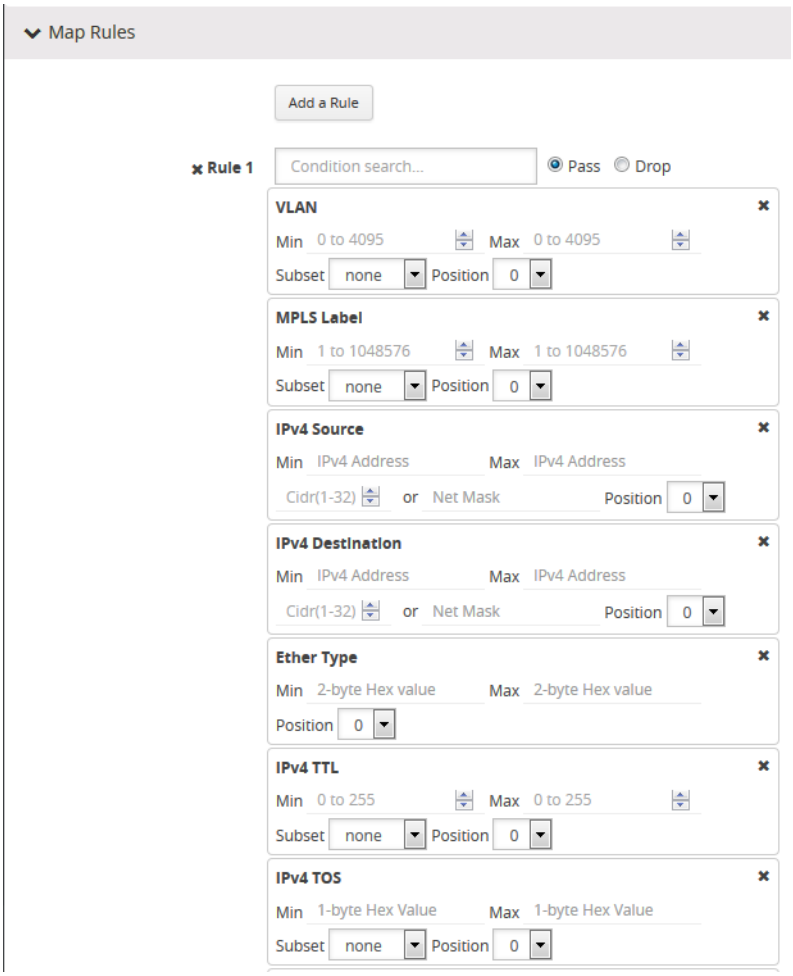
```
<attribute> min <value> max <value> subset <odd|even|none> pos
```

The second group of attributes that use this format are as follows:

- vlan id
- mpls label
- l4port src and dst
- ethertype
- ipv4 ttl, tosval, and protocol
- ipv6 flow-label

- vntag dvifid, svifid, and viflistid

The following figure shows the attributes as displayed in the UI.



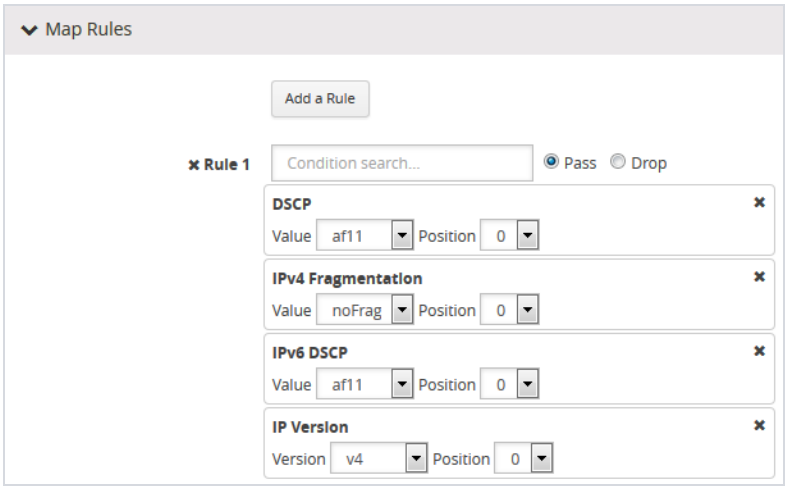
The third group of attributes has the following format: <value> <position>

```
<attribute> value <value> pos <0|1|..|n>
```

The third group of attributes that use this format are as follows:

- ipv4 dscp and frag
- ipv6 dscp
- ipver

The following figure show the attributes as displayed in the UI.



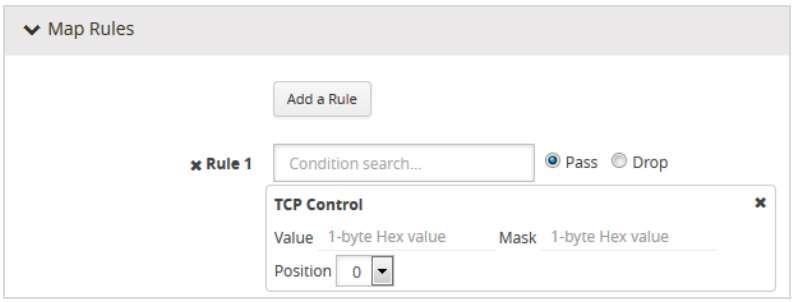
The fourth group of attributes has the following format:

```
<attribute> value <value> mask <mask> pos <0|1|2|3>
```

The fourth group of attribute that uses this format is as follows:

- tcp ctl

The following figure shows the attribute as displayed in the UI.



The maximum occurrences of each attribute supported are as follows:

Attribute	Maximum Occurrences
Attributes in IPv4 header	3
Attributes in IPv6 header	3
Attributes in MAC header	3
VLAN ID	4
MPLS label	4
Attributes in L4port	3
Ethertype	6

Attribute	Maximum Occurrences
Attributes in VNTag header	3
Attributes in TCP header	3
IP ver	3

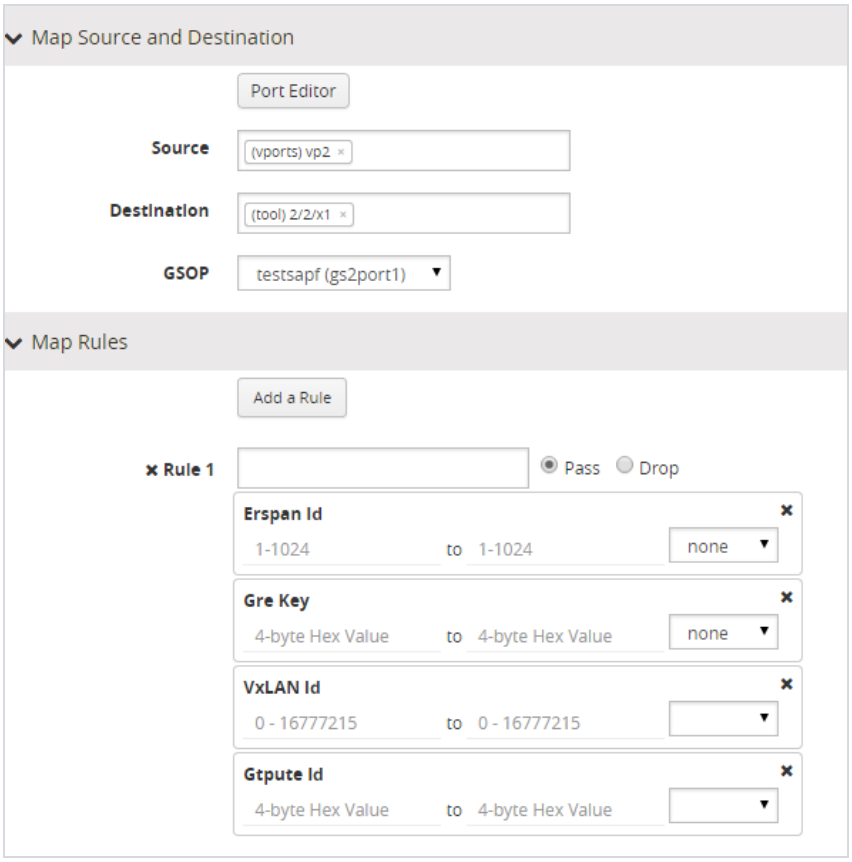
NOTE: In Generation 2 cards, for fragmented packets, only the head fragment gets processed and forwarded based on the filtering criteria while the subsequent fragments are dropped.

Encapsulation Awareness

Encapsulation awareness offers filtering across advanced encapsulation headers, including GTP tunnel ID, VXLAN ID, ERSPAN ID, and GRE key.

The following attributes of rules in a map support encapsulation awareness:

1. Enter a GTP tunnel identifier as a four-byte hex value, either a range or a single value.
2. Enter a VXLAN ID as a three-byte hex value, either a range or a single value.
3. Enter an ERSPAN ID as a decimal value from 1~1024, either a range or a single value using the corresponding arguments.
4. Enter a GRE key as a four-byte hex value, either a range or a single value.



Pattern Matching on Gen 2 GigaSMART modules

Use APF to create pattern matching filters in which the pattern is a particular sequence of data bytes at a variable or fixed offset from the start of a packet. Thus you can filter on any data patterns within a packet.

Pattern matching identifies content based on patterns in any part of the packet, including the payload. Patterns can be a static string at a user configured offset or a subset of Perl Compatible Regular Expression (PCRE) at a variable offset.

The Pattern Match attribute in a map rule supports pattern matching.

Multiple pattern matches are supported. A map can have multiple gsrules, each rule can have a pattern matching expression, and a single packet can match multiple rules.

The Pattern Match attribute in a map rule is shown in Figure 45Use Pattern Match Under Maps for Pattern Matching.

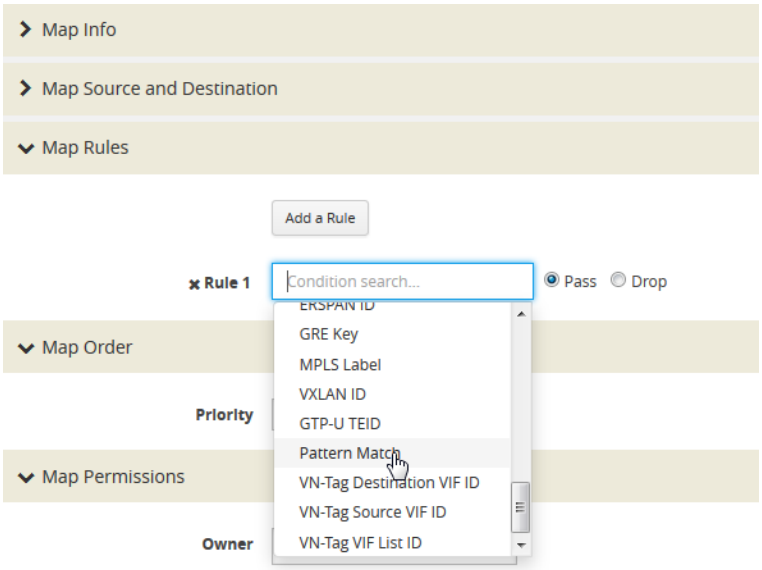


Figure 45 *Use Pattern Match Under Maps for Pattern Matching*

After selecting pattern matching for the rule, you can enter a Perl-compatible regular expression or a string to be used as a filter when pattern matching. For example to pass all packets including the string www.gigamon.com select **string** as type for the pattern match as shown in Figure 46Pattern Match with Type String.

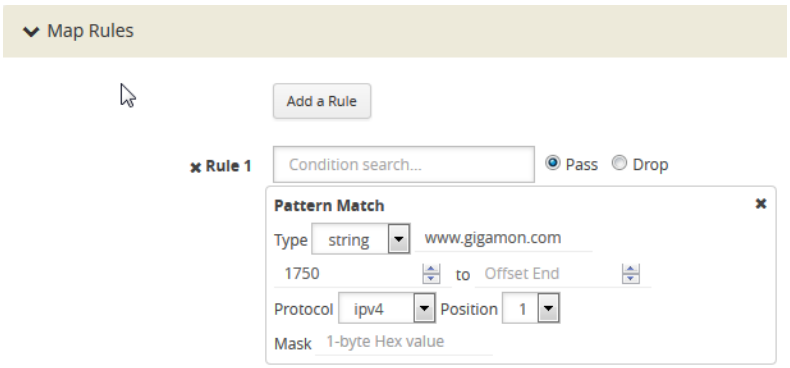


Figure 46 *Pattern Match with Type String*

To pass packets that match any phone number in the nnn-nnn-nnnn format, select **regex** for the pattern match type and enter the following regular expression in the value field: \d {3}-\d{3}-\d{4} as shown in Figure 47Pattern Match with Type RegEx.

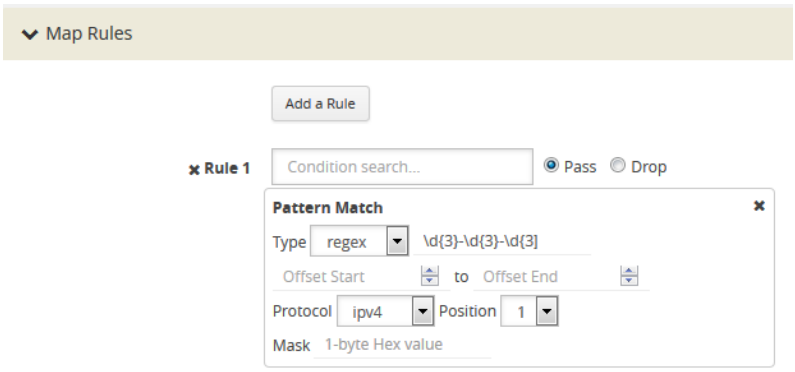


Figure 47 *Pattern Match with Type RegEx*

The offset is a value or range from 0 to 1750. The offset indicates where the pattern under search is located, specify, a value to indicate that the pattern has to start at that offset in the packet in order to be considered a match. Specify a range (beginning and ending) to indicate that the pattern can be anywhere in the packet in that range.

The optional protocol argument of the Pattern Match specifies that the matching will start after the protocol header specified in the command (IPv4, IPv6, TCP, or UDP). Pos 1 or 2 indicates the position. For example, position 2 indicates that matching is to start after the second protocol header. The offset and start and end values are also counted after the protocol header.

For example, to mask an SSL client hello packet pattern starting from the first position after the TCP header with an offset of 0 (located right after the TCP header), you define the pattern match rule as shown in Figure 48Pattern Match for SSL Client Hello Packet.

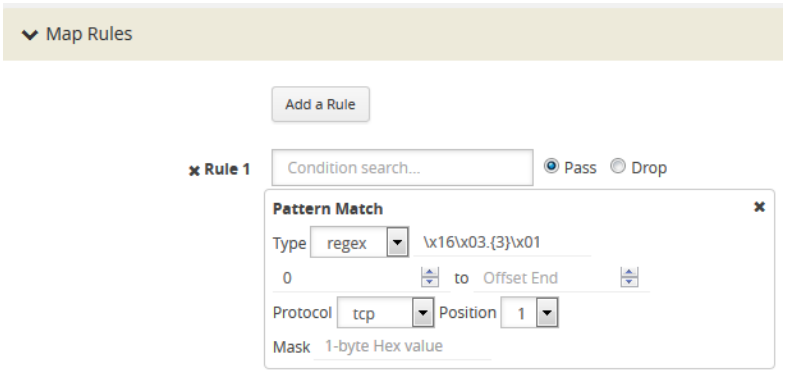


Figure 48 *Pattern Match for SSL Client Hello Packet*

Pattern Matching for IMSI (International Mobile Subscriber Identity)

An IMSI is a unique number usually fifteen digits, identifying a mobile subscriber. The IMSI field in GTP is 8 bytes and is encoded in Telephony Binary Coded Decimal (TBCD) format.

You can create a APF rule to match a specific IMSI. The APF RegEx gsrule needs to match the IMSI in TBCD format instead of the decimal format.

The procedure to convert the IMSI to TBCD hexadecimal format is as follows:

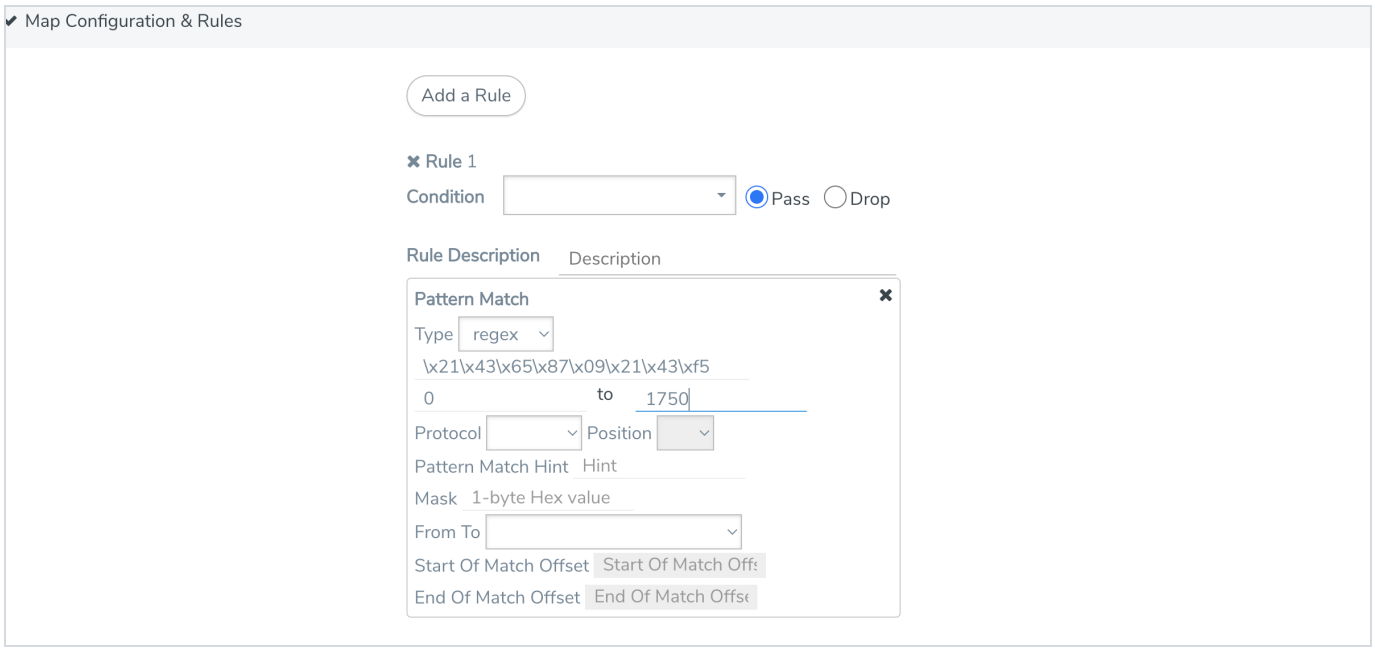
1. If the IMSI length is less than 16 digits, pad the remaining digits with F. For example, 15 digits IMSI 123456789012345 becomes 123456789012345F.

2. Group each pair of digits as 12, 34, 56, 78, 90, 12, 34, 5F.

3. Swap the most significant digit and the least significant digit of each pair.



The end result, 21436587092143F5, is the TBCD equivalent of the IMSI.



Masking with Pattern Matching

APF allows masking when there is a match through pattern matching. Use masking with pattern matching to mask out a specific portion of a packet due to security reasons or to hide sensitive information in packets.

Multiple pattern matches are supported in a map. If there is masking associated with a rule and a packet matches multiple rules, the masking action is enforced for all the matching rules in the map.

The mask specifies that the matched pattern in the gsrule will be masked with the pattern specified in the 1-byte masking pattern.

The pattern specified in the gsrule will be overwritten. The overwritten length is the length of the matching pattern specified in either a string or a RegEx pmatch. Use the 1-byte to overwrite the original pattern match pattern. If there are multiple matches in the packet, up to 10 matches will be masked.

For example, to find Social Security numbers in the format xxx-xx-xxxx, between offset 40 and 80 and replace them with zeros, create a map with a pass rule in a Second Level byRule map with the regular expression \d{3}-?\d{2}-?{4} and a mask with a 1-byte masking value of 0 as shown in Figure 49Map Rule with RegEx for Masking SSNs.

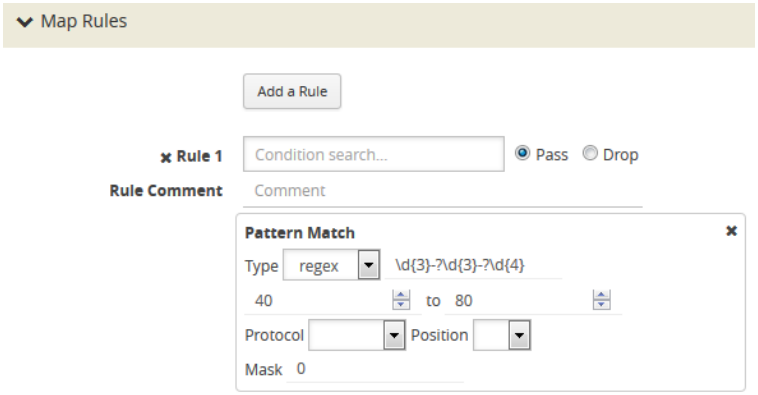


Figure 49 *Map Rule with RegEx for Masking SSNs*

Pattern Matching Hint

To optimize APF pattern matching performance in second level maps with gsrules, you can optionally use a pattern matching hint.

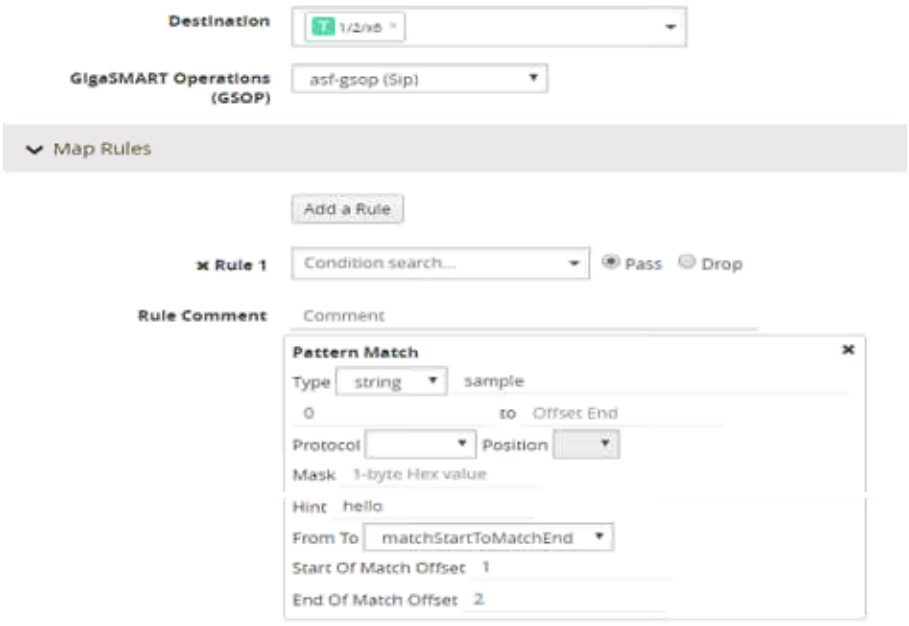


Figure 50 *Pattern Match with Hint*

The addition of the hint leads to two levels of filtering. First, the packet is subjected to a check for the simpler match comprising “gamon|GIM”. If a match is found, a second level check for a match in the complete RegEx, “a[gG]igamon|aGIMO\\s[a-f]\\d{4}”, is performed.

A hint must be selected so that all the packets that are expected to match the actual RegEx must have that string in them, otherwise the first level check will not be cleared. The hint in the example, “gamon|GIM”, was selected because a packet containing either “gamon” or “GIM” in it is a potential match to the actual RegEx.

Best Practices of Pattern Matching Hint

The pattern matching hint is optional and, to optimize performance, it should be specified for all gsrules in a map. In that map, its usage is all or none, meaning you cannot have a mix of gsrules with some having the pattern matching hint and others not. However, if there are two maps, one map can have gsrules that include the pattern matching hint, while the other map can have gsrules that do not.

The use of the pattern matching hint improves performance in complex RegEx patterns involving “lookbehind” and “lookahead” constructs of PCRE syntax. Using them in conjunction with maps with simple patterns, such as fixed length string, is not advisable as it might lead to performance degradation in some cases. Since the RegEx rule set is limitless, there are no specific rules in which the degradation happens. A best practice is to try out both options, with and without the pattern matching hint, to find out what works best.

The rule of thumb while constructing the pattern matching hint is to keep it as simple as possible. Also, it must be a subset of the configured RegEx pattern. First, try out a 3 to 6 character-wide hint. If that does not provide the necessary scale, you can make the hint wider and more specific to prevent false positives. A maximum length of 63 bytes is supported.

Cross-Packet Pattern Matching

Cross-packet pattern matching refers to a scenario where a match initiates in one packet and ends in a subsequent packet. Staring with Gigamon software release 5.4 this feature enhancement extends the support for GSOP cross packet pattern spanning two packets.

Cross packet matching applies to connection oriented exchanges only and available for 5-tuple flows. Cross packet matching scan will be performed on frames with the following header encapsulations:

- IPv4/TCP, IPv4/UDP
- IPv6/TCP, IPv6/UDP
- IPv4/IPv6/TCP, IPv4/IPV6/UDP
- IPV6/IPv4/TCP, IPV6/IPv4/UDC

Every packet of a flow is subjected to pattern matching scan starting with the inner most L4 payload section. For example, 5-tuple TCP session with nested TCP layer will position the scan starting from start of innermost TCP payload to the end of frame. Bi-directional flow maintains match context for each direction separately and this feature supports up to 1Million flows.

The figure below illustrates the Cross-packet pattern matching concept where the pattern search "**abcdef**" spans two packets.

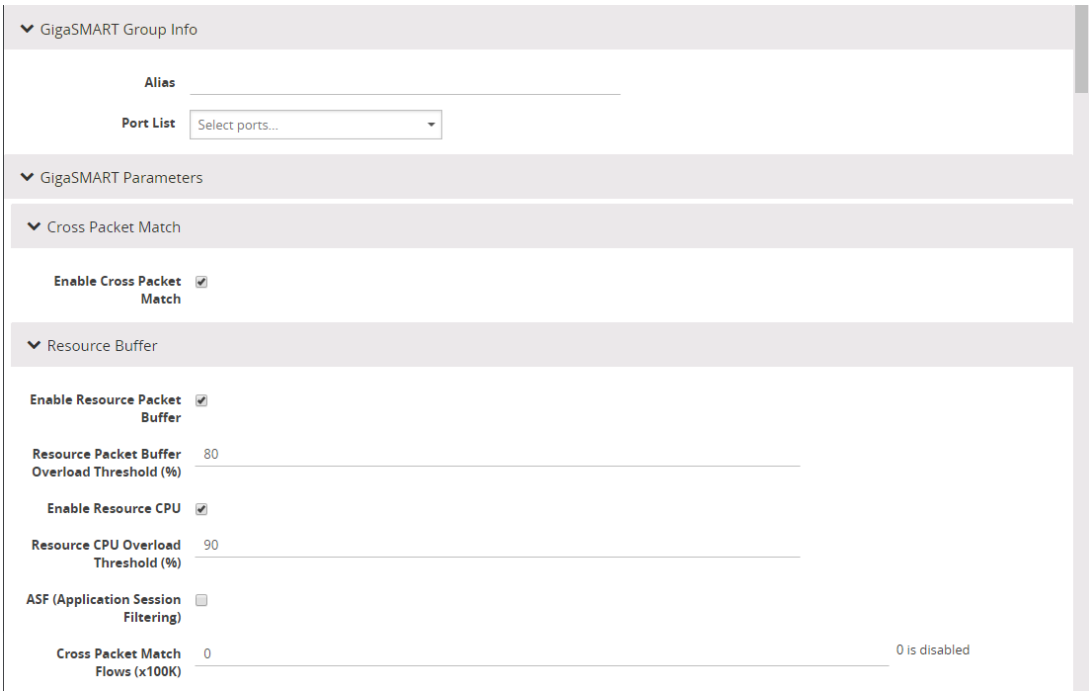


Enable/Disable Cross-packet Matching

You can enable or disable Cross-packet pattern matching from the GigaSMART GSOP operation.

1. Select a Physical Node.

2. **GigaSMART** > **GigaSMART > GigaSMART Groups**.
3. Click **New**. The GigaSMART Group parameter page displays.



4. Click the **Enable Cross-packet Match** check box to enable.
5. Enter a range from 1 to 10 for the **Cross Packet Match Flows** parameter. Each unit is 100K bi-directional flows.
6. Click **OK**.

NOTE: When disabling this functionality you will be notified that change will be effective only after chassis or GigaSMART card reboot.

Disable Cross Packet Matching

1. Repeat Steps 1 through 3 from the “Enabling Cross Packet Matching” task.
2. Uncheck the **Enable Cross-packet Match** check box to disable this functionality**.**
3. Click **OK**.

View Cross-packet Matching

1. Select a Physical Node.
2. **GigaSMART** > **GigaSMART > GigaSMART Groups**.

3. Select a Group.
4. Click **Edit**. The GigaSMART Group parameters including cross pattern match details pane displays.

Pattern Matching on Gen 3 GigaSMART modules

In Gen 3 GigaSMART module, regular expression (RegEx) pattern matching is introduced to extend Adaptive Packet Filtering (APF) capabilities. This feature enables deeper analysis of packet contents, enhancing security and traffic management functionalities by going beyond basic header filtering.

Gen 3 GigaSMART rules enable advanced traffic filtering with support for up to 25 GigaSMART rules per second level map and up to 25 APF second level maps per vPort. RegEx Pattern Matching can be combined with other filtering attributes in a single rule. Both pass and drop actions are supported with RegEx.

Pattern matching operates within the first 1750 bytes of a packet and begins from protocol-specific locations. For Layer 3 protocols (IPv4 and IPv6) and Layer 4 protocols (TCP and UDP) the pattern matching starts after the last protocol header. For all other protocols it begins at the start of the packet. If Layer 3 and Layer 4 inner headers are present in the packets, the pattern matching starts after the last inner protocol header. The Character limit is 127 for configuring the RegEx pattern. RegEx pattern can be either regular expression or string. Once a RegEx rule is configured, it cannot be modified. You must delete the rule and reconfigure it as a new rule.

It supports ASF for session-based filter and fragment packet forwarding. ASF can buffer up to 20 packets for pattern matching before forwarding traffic to the tools.

Multiple pattern matches are supported. It is recommended to configure multiple pattern in different GS rule within the same map instead of configuring multiple patterns in a single GS rule. To achieve optimal performance, each rule can be configured with a different and unique pattern, especially when multiple filtering rules exist. Each rule can have a pattern matching expression, and a single packet can match multiple rules. For Example, when you configure regex pattern “testN” in one rule and another regex pattern “Next” in another rule, then payload “testNext” can match both rules.

The differences between the Gen 2 and Gen 3 map configurations for Pattern Match RegEx are outlined in the "Gen 2 Map and Gen 3 Map Configuration for Pattern Match RegEx" section in the GigaVUE-OS CLI Reference Guide.

The Pattern Match attribute in a map rule is shown in Figure 45Use Pattern Match Under Maps for Pattern Matching.

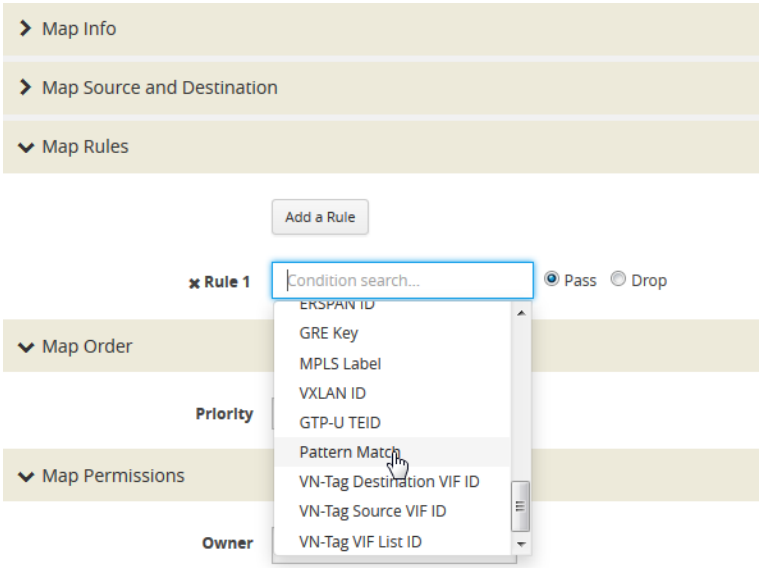


Figure 51 *Use Pattern Match Under Maps for Pattern Matching*

After selecting pattern matching for the rule, you can enter a Perl-compatible regular expression or a string to be used as a filter when pattern matching. For example to pass all packets including the string www.gigamon.com select **regex** as type for the pattern match as shown in Figure 46Pattern Match with Type String.

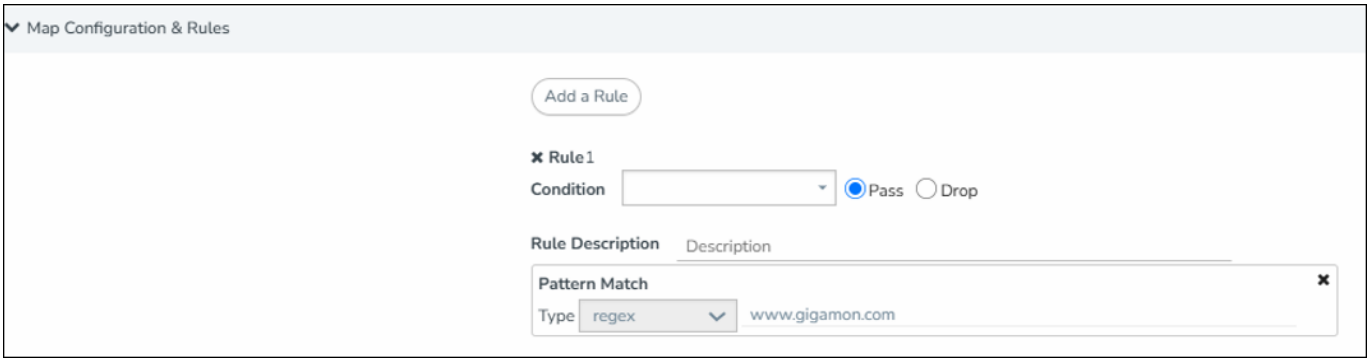


Figure 52 *Pattern Match with Type RegEx*

To pass packets that match any phone number in the nnn-nnn-nnnn format, select **regex** for the pattern match type and enter the following regular expression in the value field: \d {3}-\d{3}-\d{4} as shown in Figure 47Pattern Match with Type RegEx.

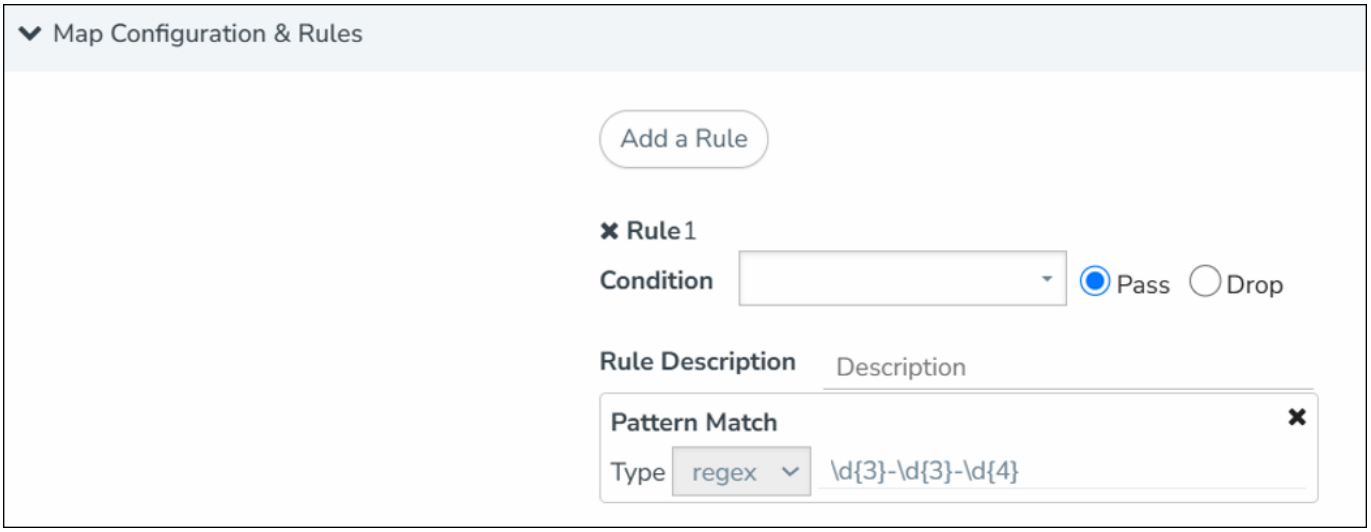


Figure 53 *Pattern Match with Type RegEx*

Regex Masking on Gen 3 GigaSMART modules

On Gen 3 GigaSMART modules, APF supports regex pattern matching with masking to mask sensitive data in the packet payload before traffic is sent to tools. Typical use cases include masking personally identifiable information (PII) such as Social Security numbers, credit card numbers, account IDs, or user IDs to help support privacy and compliance requirements.

Masking is configured on APF GigaSMART rules that use Pattern Match with the RegEx type. When a packet matches the regex, APF overwrites the matched region or a sub-range relative to the start or end of the match , with a 1 byte mask value. This allows you to:

- Mask an entire matched value (for example, a full SSN or credit card number).
- Mask only the value part of a key-value pair (for example, the bytes following user_id=) while leaving the key visible.

When ASF is configured together with APF regex-based masking, the system must mask only the first matching packet in a session. All subsequent packets in the same session must be sent to the tool without masking (if applicable).

Configuring a Regex Masking on Gen 3 GigaSMART modules

To configure Regex Masking on Gen3 GigaSMART modules :

1. From the left navigation pane, go to **Inventory > Physical > Nodes**.
2. Under **Traffic > Maps > Maps.**

3. Create or (edit) a second-level map that uses an APF GSOP on a Gen 3 GigaSMART module. Refer to Implement APF Through the UI
4. Under the **Map Configuration & Rules**
   a. Set Action to **Pass**.
   b. Set Condition to **Pattern Match**.
   c. Set Type to **RegEx**.
5. In the Pattern Match area, configure:
   a. Value: a regex that matches the sensitive data.
   b. Range: an appropriate offset or begin to end range where the pattern can appear (for example, 0..1750).
6. In the **Masking** section (visible only for RegEx on Gen 3 nodes), fill in:
   a. Mask byte: A 1-byte mask value (for example, 0x2a to mask with *).
   b. From – Choose where masking should start relative to the regex match and set an offset:
   - Start of match: offset 0–1749
   - End of match: offset 0–1748

   c. To – Choose how far masking should extend and set the allowed range:
   - End of match: offset –1749 to +1749
   - End of packet: masks from the From position to the end of the packet.
   - Length:
     - 1–1750 bytes when From is Start of match
     - 1–1749 bytes when From is End of match.

The masking range is defined relative to the regex match—for example, from the start to the end of the match, or from the end of the match for a fixed length, so that you can either mask the entire matched value or only the value portion of a key–value pair (for example, the bytes after user_id=) while leaving the key visible.

> **NOTE:** Masking is applied only for pass rules. Drop rules can still use regex to decide whether to drop a packet, but they do not perform masking.

On Gen 3, 25 APF second-level maps per vPort are supported only when all of those maps use APF (with or without regex masking). If any other GigaSMART application is configured on that vPort (alone or in combination with APF), a maximum of five second-level maps is supported for that vPort.

Unsupported PCRE RegEx patterns

The following PCRE regex constructs are not supported by the Gen 3 APF regex engine. These limitations apply to both regex-based pattern matching and regex-based masking on Gen 3 GigaSMART modules.

- Continuation escape
- Collating elements
- Equivalence classes
- Line endings escape
- Character class subtraction
- Character class symmetric difference
- Character class complement
- Possessive quantifiers
- Backreferences
- Line comments
- Branch reset
- Lookahead
- Lookbehind
- Non-backtracking expressions
- Recursion
- Conditional expressions
- Subroutines
- Callouts
- Backtracking control verbs

Feature Parity with Gen 2 GigaSMART module and Gen 3 GigaSMART module

Pattern Matching Features	Gen 2	Gen 3
Pattern Matching	✓	✓
Masking with Pattern Matching	✓	✓
Cross-Packet Pattern Matching	✓	✗
Pattern Matching Hint	✓	✗

NOTE: On Gen 3, masking with pattern matching is supported only for APF gsrules that use Pattern Match with the RegEx type.

Map Statistics

Go to **Map > Statistics** to display counts of the rules that actually matched in a map. A single packet can match one or more rules. For example, if a single packet matches multiple rules in an APF map, all matching rules will have that packet counted against them and the overall map status pass counter will show 1.

Display APF Statistics

Refer to APF Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

Adaptive Packet Filtering Examples

The following are APF examples:

- Identify Social Security Numbers in User-Level Transactions
- Mask Social Security Numbers
- Filter on Fiber Channel over Ethernet (FCOE) Traffic
- Multi-Encapsulation Filtering
- Filter on Subscriber Device IP (User-Endpoint IP or UE-IP)

- Filter on Inner Layer 2-4 Parameters for Unrecognized Headers
- GTP Tunnel ID-Based Filtering
- ERSPAN Tunneling
- Distribute Traffic Based on Inner IP Addresses and Inner TCP Port Values
- MPLS Label Based Filtering
- Combine APF with GigaSMART Operations
- Conditional Header Stripping
- Facilitate Overlapping Rules

Identify Social Security Numbers in User-Level Transactions

The following example looks for packets containing Social Security Numbers in an incoming traffic stream using pattern matching. Once a match is detected, the packets are forwarded to a monitoring tool for additional analysis.

Task	Description	UI Steps
1	Configure one network and two tool ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3. Select Tool for port 1/1/x4 and port 1/1/x1. **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsgrp1 in the Alias field. **4.** Click Save.
3	Configure the GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **2.** Click New. **3.** Type gsfil in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Select APF from the GigaSMART Operations (GSOP) list. **6.** Select Enable. **7.** Click Save.
4	Create a virtual port.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp1 in the Alias field.

Task	Description	UI Steps
		4. Select gsgrp1 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map to forward traffic from network port 1/1/x3 to virtual port vp1.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the port 1/1/x3 for the Source. ▪ Select the virtual port vp1 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IP Version. **d.** Select v4 for Version. **5.** Click Save.
6	Create a second level map to forward traffic from the virtual port vp1 to GigaSMART with pattern matching.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map2 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x1 for the Destination. **4.** To add a rule for Gen 2 GigaSMART modules. Follow the below steps. ▪ Click Add a Rule. ▪ Select Pass. ▪ Select Pattern Matching. ▪ Select regex for Type and enter the value \d{3}-?\d{2}-?\d{4}. ▪ Set the Offset Start to 40. ▪ Set the Offset End to 8. **5.** To add a rule for Gen 3 GigaSMART modules. Follow the below steps. ▪ Click Add a Rule. ▪ Select Pass. ▪ Select Pattern Matching. ▪ Select regex for Type and enter the value \d{3}-?\d

Task	Description	UI Steps
		{2}-?\d{4}. **6.** Click Save.

Mask Social Security Numbers

In the following pattern matching example, IPv4 packets contain Social Security Numbers (SSNs) in the format xxx-xx-xxxx. If the SSNs are between offset 40 and 80, they will be replaced with zeros. This example is not applicable to Gen 3 GigaSMART module.

Task	Description	UI Steps
1	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsgrp1 in the Alias field. **4.** Click Save.
2	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter gsTraffic in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Click Save.

Task	Description	UI Steps
3	Create a first level map to direct traffic from network port 1/1/x1 to virtual port gsTraffic.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the port 1/1/x3 for the Source. ▪ Select the virtual port gsTraffic for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IP Version. **d.** Select v4 for Version. **5.** Click Save.
4	Configure the GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations. **2.** Click New. **3.** Select gsgrp1 from the GigaSMART Groups list. **4.** Select Adaptive Packet Filtering from the GigaSMART Operations (GSOP) list. **5.** Select Enable. **6.** Click Save.
5	Create a second level map to direct traffic from the virtual port gsTraffic to GigaSMART.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map2 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x6 for the Destination. ▪ Select gsop1 from the GSOP list. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Pattern Matching. **d.** Select regex for Type and enter the value \d{3}-?\d{2}-?\d{4} **e.** Set the Offset Start to 40. **f.** Set the Offset End to 80

Task	Description	UI Steps
		g. Enter 0 for Mask. 8. Click Save.

Mask Passwords and Account IDs (Gen 3 Regex Masking)

In this example, Adaptive Packet Filtering (APF) on Gen 3 GigaSMART modules uses regex pattern matching with masking to protect sensitive fields—such as passwords or alphanumeric account IDs—embedded in application payloads. The regex defines what to match, and the From/To masking window controls which portion of the match is overwritten in the packet before it is sent to tools.

Example 1 – Mask the entire password value

Many applications send credentials in key–value form, such as password=P@ssw0rd!123, in URLs, form bodies, or headers. In this use case, APF regex masking hides the entire password value, while leaving the password= key visible for troubleshooting.

- Regex pattern
  - password=[^&\s]{5,64}
  - Matches strings that start with password= followed by 5–64 characters that are neither a space nor &.
    - Example matches:
    - password=P@ssw0rd!123
    - password=abc123XYZ!
- Masking window
  - From: start-of-match 9
    - password= is 9 characters.
    - Skips the key and starts masking at the first character of the password value.
  - To: end-of-match 0
    - Masks from that point to the end of the regex match, so the full password value is overwritten.
- Effect on payload
  - Original: password=P@ssw0rd!123
  - Masked (mask byte 0x2a, *): password=**********

For example, you can configure a GigaSMART rule similar to:

```
gsrule add pass pmatch mask 0x2a from start-of-match 9 to end-of-match 0 RegEx
"password=[^&\s]{5,64}" 0..1750
```

Example 2 – Partially mask account IDs, keeping the last four characters

Some environments require partially masking identifiers, allowing operators to see only the tail of an ID (for example, for audit or troubleshooting) while hiding most of the value. In this use case, APF regex masking hides the leading portion of an account ID and preserves the last four characters.

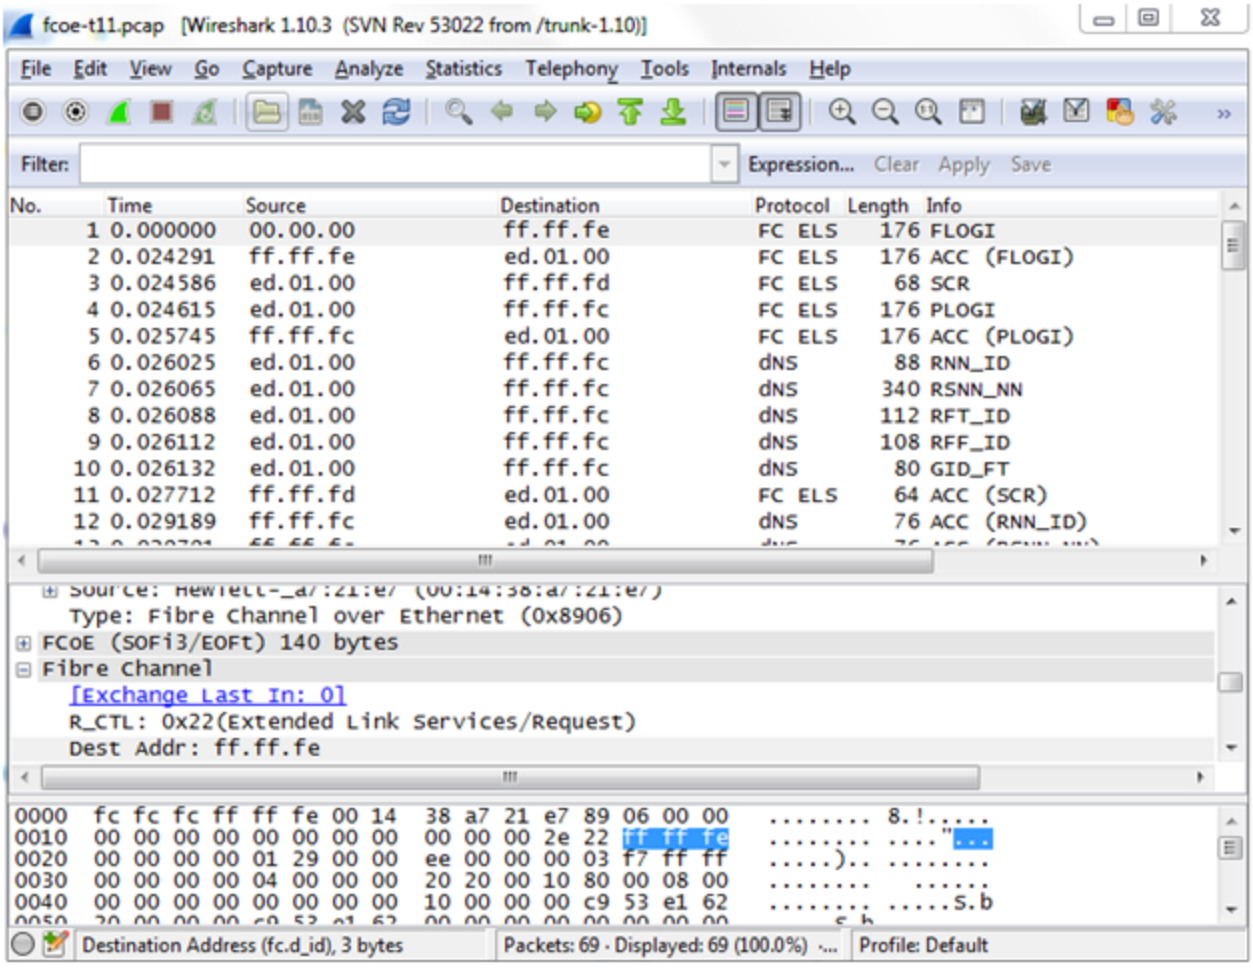
- Regex pattern for accountId=[A-Za-z0-9]{8,20}
  - Matches strings that start with accountId= followed by 8–20 alphanumeric characters. Example match: accountId=AB12CD34EF56
- Masking window
  - From: start-of-match 10
    - accountId= is 10 characters.
    - Skips the key and starts masking at the first character of the account ID value.
  - To: end-of-match -4
    - Stops masking 4 bytes before the end of the match, leaving the last four characters visible.
- Effect on payload
  - Original: accountId=AB12CD34EF56
  - Masked (mask byte 0x2b): accountId=********EF56

A corresponding GigaSMART rule can look like:

```
gsrule add pass pmatch mask 0x2b from start-of-match 10 to end-of-match-4 RegEx
"accountId=[A-Za-z0-9]{8,20}" 0..1750
```

Filter on Fiber Channel over Ethernet (FCOE) Traffic

The flexibility offered by regular expression-based filters can be used as an infrastructure to classify traffic streams with protocol headers that are typically unsupported on traditional TAP/SPAN aggregation devices. In this example, regular expression-based filters are used for filtering on the source address in a Fiber Channel header.



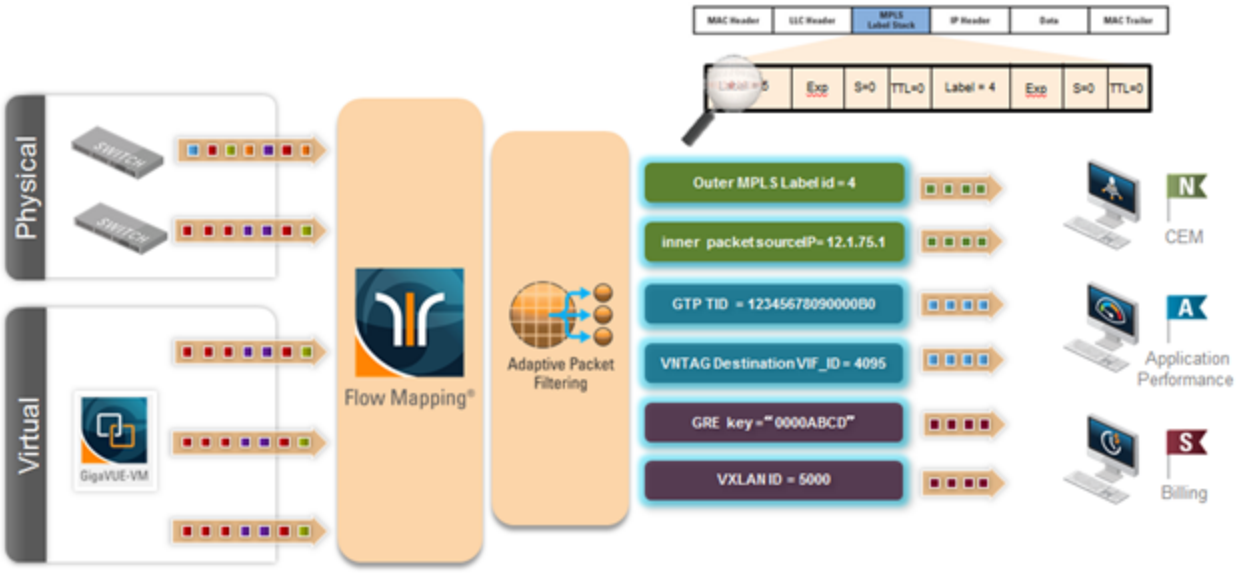
Task	Description	UI Steps
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3. Select Tool for port 1/1/x4 and port 1/1/x1. **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsgrp1 in the Alias field. **4.** Click Save.
3	Configure the GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation.

Task	Description	UI Steps
		2. Click New. **3.** Select gsfil from the GigaSMART Groups list. **4.** Select Adaptive Packet Filtering from the GigaSMART Operations list. **5.** Select Enable. **6.** Click Save.
4	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter gsTraffic in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Click Save.

Task	Description	UI Steps
5	Create a first level map to forward FCOE traffic to the virtual port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter to_vp in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the port 1/1/x3 for the Source. ▪ Select the virtual port vp1 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select EtherType **d.** Enter 8906 in the Value field. **5.** Click Save.
6	Create a second level map to filter on regular expression, using a string match to the destination address in the FCOE packet.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map2 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x1 for the Destination. ▪ Select gsfil from the GSOP list. **4.** To add a rule for Gen 2 GigaSMART modules. Follow the below steps. ▪ Click Add a Rule. ▪ Select Pass. ▪ Select Pattern Match. ▪ Select string for Type and enter txff\xff\xfe. ▪ Set the Offset Start to 0. ▪ Set the Offset End to 29. **5.** To add a rule for Gen 3 GigaSMART modules. Follow the below steps. ▪ Click Add a Rule. ▪ Select Pass. ▪ Select Pattern Match. ▪ Select regex for Type and enter txff\xff\xfe. **6.** Click Save.

Multi-Encapsulation Filtering

In order to complement the mobility brought about by the virtualized server infrastructure, network virtualization overlays like VXLAN, VNTag, NVGRE are being designed and implemented in Data Centers and Enterprise environment. Across Service Provider environments, huge volumes of traffic are being tunneled over GTP. Until now, the Gigamon Deep Observability Pipeline provided the option of stripping out these headers, thus providing visibility to monitoring tools that do not understand these overlays and encapsulation protocol. With APF, this capability is further enhanced where operators now have the option of making forwarding decisions based on the encapsulation and inner packet contents.



With encapsulation awareness enabled by APF, operators have multiple options to act on the packet including the flexibility to:

- Filter on encapsulation header parameters, Layer 2 – 4 parameters in the outer or inner headers (up to 5 layers of encapsulation) in any combination. For example:
  - Forward traffic specific to a subset of VXLAN IDs to one or more monitoring tools.
  - Distribute traffic based on MPLS label values across one or more monitoring tools.
- In combination with Header Stripping:
  - Implement “conditional” header-stripping, based on encapsulation header parameters or inner/outer packet contents, as follows:

    – Forward a subset of traffic “as-is” to monitoring tools that need these encapsulations for analysis.

– Alternatively, strip out the outer headers/encapsulations and distribute traffic to monitoring tools that do not require these outer headers for analysis.

- Since APF is implemented as a second level map, operators can also implement overlapping rules where:
  - A copy of the traffic can be distributed across a group of monitoring tools.
  - A refined subset from the same incoming stream is distributed across a different set of tools.

Filter on Subscriber Device IP (User-Endpoint IP or UE-IP)

Encapsulation awareness enabled by APF allows mobile operators to filter on Layer 2 – 4 header parameters found in an encapsulated packet.

This allows operators to filter and forward traffic specific to a mobile subscriber device or a group of subscriber devices, identified by their IP address (User-Endpoint IP) to one or more monitoring tools.

In this example, we are:

- Identifying and forwarding traffic from / to a UE-IP of 1.1.1.1 to a monitoring tool connected to 1/1/x1
- Identifying and forwarding traffic from / to a UE-IP of 1.1.1.2 to a different monitoring tool connected to tool port 1/1/x4

In many cases, the GTP control sessions are low-volume and are useful in providing some level of visibility in to the quality of experience of the subscribers. To this end, operators prefer to replicate the control sessions across all the monitoring tools, while filtering and forwarding a subset of the user-plane sessions to a subset of monitoring tools. The following example also illustrates configuration commands, leveraging the patented flow-mapping technology to replicate the GTP control sessions across all the monitoring tools involved in the traffic analysis.

Task	Description	UI Steps
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3. Select Tool for port 1/1/x4 and port 1/1/x1. **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and	**1.** From the device view, select GigaSMART > GigaSMART

Task	Description	UI Steps
	associate it with a GigaSMART engine port.	Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Select engine port 1/1/e1 in the Port List field. **5.** Click Save.
3	Configure the GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations. **2.** Click New. **3.** Type gsfil in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list. **5.** Select Adaptive Packet Filtering from the GigaSMART Operations (GSOP) list. **6.** Select Enable. **7.** Click Save.
4	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map to forward GTP-u traffic to the virtual port. **NOTE:** In the rule, 2152 is GTP-u traffic.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter to_vp in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the port 1/1/x3 for the Source. ▪ Select the virtual port vp1 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Port Source **d.** Enter 2152 for the port value. **5.** Click Save.
6	Create a first level map to forward GTP-c traffic to the tools. **Note**: In the rule, 2123 is GTP-c traffic.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type to_tool in the Alias field. ▪ Select Regular for Type. ▪ Select By Rule for Subtype.

Task	Description	UI Steps
		▪ Select the port 1/1/x3 for the Source. ▪ Select port 1/1/x1 and port 1/1/x4 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Port Source **d.** Enter 2123 for the port value. **5.** Click Save.
7	Create a second level map to filter on source and destination IP (bi-directional).	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x1 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Source. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Add a Rule 2. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Click Save.
8	Create another second level map to filter on source and destination IP (bi-directional).	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source.

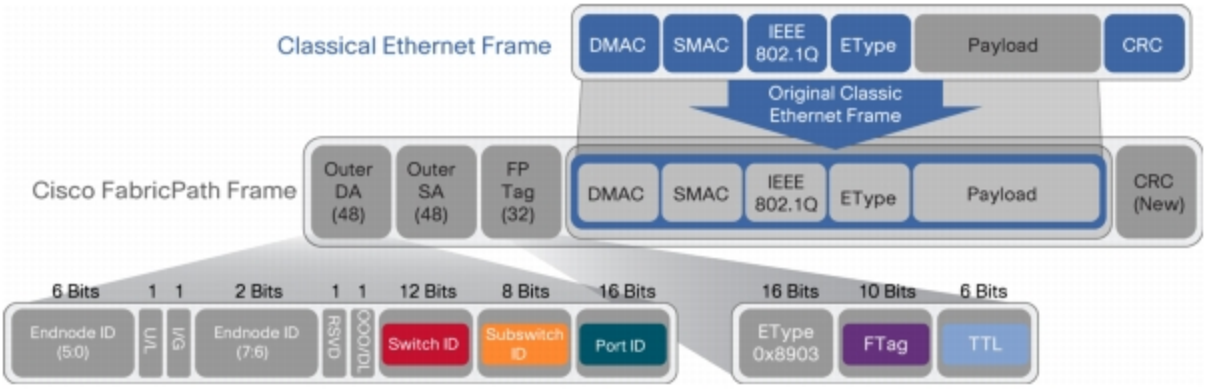
Task	Description	UI Steps
		▪ Select the tool port 1/1/x4 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Source. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Add a Rule 2. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Click Save.

Filter on Inner Layer 2-4 Parameters for Unrecognized Headers

The flexibility of encapsulation awareness enables filtering on encapsulated contents even if APF does not recognize the outer encapsulation header. The following example illustrates a packet encapsulated in Fabric Path headers. Fabric Path headers (as shown in the figure) are mac-in-mac headers that are currently not recognized by APF. However operators can still filter and forward traffic flows based on Layer 2 – 4 parameters found in the encapsulated packets.

In this example, we are:

- Identifying and forwarding traffic from/to ip 1.1.1.1 in the inner / original packet to monitoring tool connected to tool port 1/1/x1
- Identifying and forwarding traffic from/to ip 1.1.1.2 in the inner / original packet to monitoring tool connected to tool port 1/1/x4



Task	Description	UI Steps
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3. Select Tool for port 1/1/x4 and port 1/1/x1. **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Select engine port 1/1/e1 in the Port List field. **5.** Click Save.
3	Configure the GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **2.** Click New. **3.** Type **4.** Select gsfil from the GigaSMART Groups list. **5.** Select Adaptive Packet Filtering from the GigaSMART Operations (GSOP) list. **6.** Select Enable. **7.** Click Save.
4	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map to forward fabric path packets to the virtual	**1.** Select Maps > Maps > Maps.

Task	Description	UI Steps
	port.	**2.** Click New. **3.** Configure the map. ▪ Enter to_vp in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the port 1/1/x3 for the Source. ▪ Select the virtual port vp1 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select EtherType **d.** Enter 8903 in the Value field. **5.** Click Save.
6	Create a second level map to filter on source and destination IP (bi-directional).	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x1 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Source. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 1. **7.** Add a Rule 2. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 1. **7.** Click Save.
7	Create another second level map to filter on source and destination IP (bi-	**1.** Select Maps > Maps > Maps. **2.** Click New.

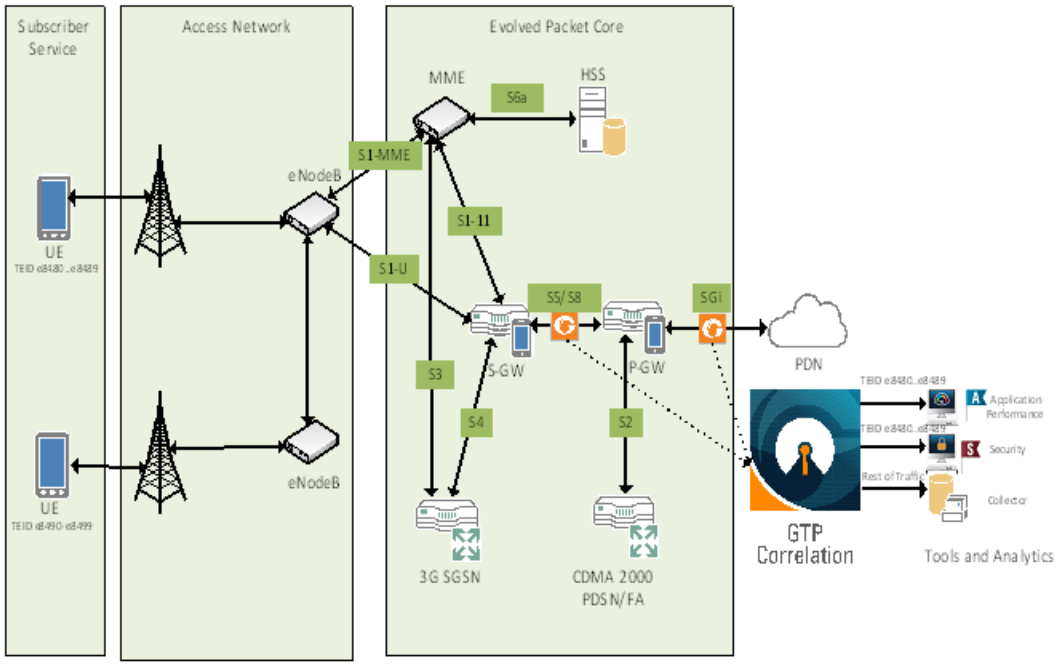
Task	Description	UI Steps
	directional).	**3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x4 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Source. **d.** Enter 1.1.1.2 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 1. **7.** Add a Rule 2. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination. **d.** Enter 1.1.1.2 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 1. **7.** Click Save.

GTP Tunnel ID-Based Filtering

The following example demonstrates filtering and forwarding traffic based on tunnel IDs included as part of the GTP user-plane messages. It also illustrates the concept of a shared collector to which traffic not matching any of the configured filters can be optionally sent. GTP control sessions are forwarded to all the monitoring tools leveraging the power of flow mapping by filtering on Layer-4 UDP port 2123.

For GTP-u:

- Filter and forward teid ranges 0x001e8480..0x001e8489 to a monitoring tool
- Filter and forward teid ranges 0x001e8490..0x001e8499 to another monitoring tool
- Forward the rest of the traffic to a shared collector



Task	Description	UI Steps
1	Configure one network and three tool type of ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x9. Select Tool for the port s 1/1/x13, 1/1/x14, and 1/1/x15. **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Select engine port 1/1/e1 in the Port List field. **5.** Click Save.
3	Configure the GigaSMART operation and assign it to the GigaSMART group. Packets processed by this operation are evaluated using Adaptive Packet Filtering (APF) rules.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **2.** Click New. **3.** Type gsfil in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list. **5.** Select Adaptive Packet Filtering from the GigaSMART (GSOP) Operations list. **6.** Select Enable. **7.** Click Save.

Task	Description	UI Steps
4	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map that directs GTP-u traffic from physical network port/s to the virtual port created in the previous step. **NOTE:** In the rule, 2152 is GTP-u traffic.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter to_vp in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the port 1/3/x9 for the Source. ▪ Select the virtual port vp1 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select EtherType **d.** Enter 8903 in the Value field. **5.** Click Save.
6	Create a first level map that directs GTP-u traffic from physical network port/s to the tool ports. **NOTE:** In the rule, 2123 is GTP-c traffic.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter ctrl_to_tool in the Alias field. ▪ Select Regular for Type. ▪ Select By Rule for Subtype. ▪ Select the port 1/3/x9 for the Source. ▪ Select the port 1/3/x13 and port 1/3/x15 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Port Source. **d.** Enter **2123** for the port value. **5.** Click Save.
7	Create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, and matches tunnel IDs specified by the gsrule.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type m1 in the Alias field. ▪ Select Second Level for Type.

Task	Description	UI Steps
		▪ Select By Rule for Subtype. ▪ Select the port 1/3/x15 for the Source. ▪ Select the virtual port vp1 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select GTP-U TEID. **d.** Enter 0x001e8480 for Min and 0x001e8489 for Max. **e.** Set Subset to none. **6.** Click Save.
8	Create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, and matches tunnel IDs specified by the gsrule.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type m2 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the port 1/3/x15 for the Source. ▪ Select the virtual port vp1 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select GTP-U TEID. **d.** Enter 0x001e8490 for Min and 0x001e8499 for Max. **e.** Set Subset to none. **6.** Click Save.
9	Add a shared collector for any unmatched data and send it to the third tool port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type scoll in the Alias field. ▪ Select Second Level for Type. ▪ Select Collector for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the port 1/3/x14 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Click Save.

ERSPAN Tunneling

In this example, APF is used to filter packets based on ERSPAN ID. The ERSPAN header is not removed from the packet.

A second level map is configured in the example. A virtual port feeds traffic to the second level map. APF filters the packets and forwards those that match the filter criteria in the map.

Task	Description	UI Steps
1	Configure a tool type of port.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Select Tool for a port. For example, port 1/1/g1. **4.** Select Enable. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsgp2 in the Alias field. **4.** Select an engine port 1/3/e1 in the Port List field. For example, 1/3/e2 **5.** Click Save.
3	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp in the Alias field. **4.** Select gsgrp2 from the GigaSMART Groups list. **5.** Click Save.

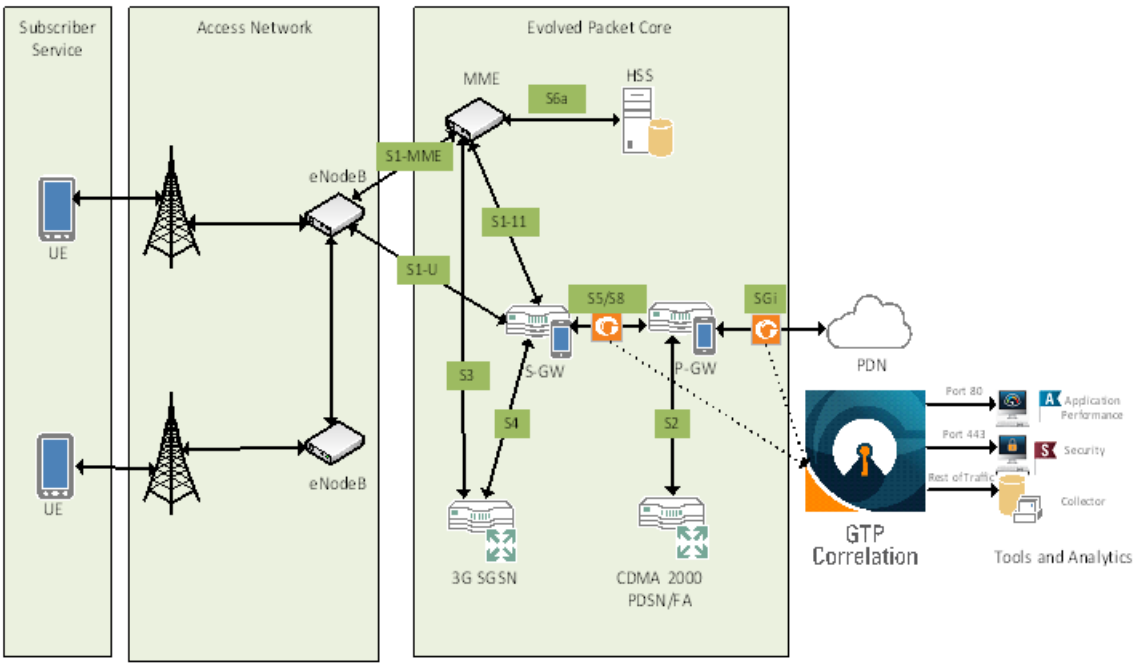
Task	Description	UI Steps
4	Configure the GigaSMART operation and assign it to the GigaSMART group.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations. **2.** Click New. **3.** Type er2 in the Alias field. **4.** Select gsgp2 from the GigaSMART Groups list. **5.** Select Adaptive Packet Filtering from the GigaSMART Operations list. **6.** Select Enable. **7.** Click Save.
5	Create a first level map.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type test1a in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the network port for the Source. For example, 1/1/g3. ▪ Select the virtual port vp for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select MAC Source **d.** Enter the address 0000.0000.0000 for Min and the address 0000.0000.0000 for Max. **5.** Click Save.
6	Create a second level map.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type test1b in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the network port for the Source. For example, 1/1/g3. ▪ Select the virtual port vp for the Destination. ▪ Select er2 from the GSOP list. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select MAC Source

Task	Description	UI Steps
		d. Enter the address 0000.0000.0000 for Min and the address 0000.0000.0000 for Max. 5. Click Save.

Distribute Traffic Based on Inner IP Addresses and Inner TCP Port Values

In the following example, traffic is distributed based on inner IP addresses and inner TCP port values as follows:

- Packets from VLAN 20 with GTP inner IP 65.128.7.21 and 98.43.132.70, inner TCP port 80 is forwarded to one tool port
- Packets from VLAN 20 with GTP inner IP 65.128.7.21 and 98.43.132.70, inner TCP port 443 is forwarded to a second tool port
- All packets not matching these rules is forwarded to a third tool port



Task	Description	UI Steps
1	Configure one network and three tool type of ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and three tool ports. For example, select Network for port 1/1/x1. Select Tool for the port s 1/1/x10, 1/1/x11, and 1/1/x12. **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsgrp1 in the Alias field. **4.** Select an engine port in the Port List field. For example, 1/1/e1. **5.** Click Save.
3	Configure the GigaSMART operation and assign it to the GigaSMART group. Packets processed by this operation are evaluated using Adaptive Packet Filtering (APF) rules.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **2.** Click New. **3.** Type g1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Select APF from the GigaSMART Operations (GSOP) list. **6.** Select Enable. **7.** Click Save.
4	Configure a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter gsTraffic in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map that directs traffic from the physical network port to the virtual port created in the previous step.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type map1 in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the network port for the Source. For example, 1/1/x1 ▪ Select the virtual port gsTraffic for the Destination. **4.** Add a rule with three conditions. **a.** Click Add a Rule. **b.** Select Pass.

Task	Description	UI Steps
		c. Select VLAN and enter 20 for Min. d. Select IPv4 Protocol and select UDP for Value. e. Select Port Destination and enter 2152 for the port value 6. Click Save.
6	Create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, matches the rules, and sends the traffic to one tool port.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. ▪ Type map2 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port gsTraffic for the Source. ▪ Select the port 1/1/x10 for the Destination. ▪ Select g1 from the GSOP list. 4. Add a rule with three conditions. a. Click Add a Rule. b. Select Pass. c. Select IPv4 Destination then enter 65.128.721 for the IP address and 255.255.255.255 for the Net Mask. Set position to 2. d. Select IPv4 Protocol and set the Potion to 2. e. Select Port Destination and enter 80 for the port value and select 2 for Position. 6. Add a rule with three conditions. a. Click Add a Rule. b. Select Pass. c. Select IPv4 Destination then enter 98.43.132.70 for the IP address and 255.255.255.255 for the Net Mask. Set Position to 2. d. Select IPv4 Protocol and set the Position to 2. e. Select Port Destination and enter 80 for the port value and select 2 for Position. 6. Click Save.
7	Create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, matches the rules, and sends the traffic to another tool port.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. ▪ Type map3 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port gsTraffic for the Source.

Task	Description	UI Steps
		▪ Select the port 1/1/x10 for the Destination. ▪ Select g1 from the GSOP list. **4.** Add a rule with three rule conditions. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination then enter 65.128.721 for the IP address and 255.255.255.255 for the Net Mask. Set Position to 2. **d.** Select IPv4 Protocol. Set Position to 2 **e.** Select Port Destination and enter 443 for the port value and select 2 for Position. **6.** Add another rule with three rule conditions. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination then enter 98.43.132.70 for the IP address and 255.255.255.255 for the Net Mask. Set position to 2. **d.** Select IPv4 Protocol. Set position to 2. **e.** Select Port Destination and enter 443 for the port value and set Position to 2. **6.** Click Save.
8	Add a shared collector for any unmatched data and send it to the third tool port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type mapclin the Alias field. ▪ Select Second Level for Type. ▪ Select Collector for Subtype. ▪ Select the virtual port gsTraffic for the Source. ▪ Select the port 1/1/x12 for the Destination. **4.** Click Save.

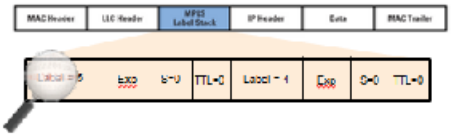
MPLS Label Based Filtering

Multiprotocol Label Switching (MPLS) is a mechanism in high-performance telecommunications networks that directs data from one network node to the next based on short path labels rather than long network addresses, avoiding complex lookups in a routing table. The labels identify virtual links (paths) between distant nodes rather than endpoints.

MPLS is a scalable, protocol-independent transport. In an MPLS network, data packets are assigned labels. Packet-forwarding decisions are made solely on the contents of this label, without the need to examine the packet itself. This allows one to create end-to-end circuits across any type of transport medium, using any protocol.

However in the context of Gigamon Deep Observability Pipeline nodes, traffic flows encapsulated in MPLS labels cannot be filtered and forwarded. With the wide-scale adoption of MPLS as a technology across enterprise and service provider environments, the ability to classify traffic flows based on MPLS labels would be a huge value add to granularly control the flow of traffic to the monitoring tools. APF can be leveraged to filter and forward traffic flows based on MPLS label values. MPLS can stack multiple labels to form tunnels within tunnels. The flexibility of APF facilitates traffic classifications across up to 5 levels of MPLS label stacks in addition to the capability to filter and forward based on Layer 2-4 parameters found in the encapsulated packet. The following example illustrates filtering and forwarding traffic based on MPLS labels, as follows:

- Filter and forward traffic flows specific to mpls label = 4 at the second level in the MPLS label stack to tool 1
- Filter and forward traffic flows specific to mpls label = 3 at the first level in the MPLS label stack to tool 2



Step	Description	Command
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3 and select Tool for the port s 1/1/x4 and 1/1/x1 **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Select an engine port in the Port List field. For example, 1/1/e1 **5.** Click Save.
3	Configure the GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART

Step	Description	Command
		Operations (GSOP) > Operations. **2.** Click New. **3.** Type gsfil in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list. **5.** Select APF from the GigaSMART Operations list. **6.** Select Enable. **7.** Click Save.
4	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp1 in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list. **5.** Click Save.

Step	Description	Command
5	Create a first level map to forward traffic to the virtual port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type to_vp in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the network port for the Source. For example, 1/1/x3 ▪ Select the virtual port vp1 for the Destination. **4.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IP Version and set Version to v4. **4.** Add Rule 2. **a.** Click Add a Rule. **b.** Select Pass and Bi Directional. **c.** Select MAC Source and enter 00:00:00:00:00:00 for the address. **d.** Set Version to v4. **5.** Click Save.
6	Create another second level map to filter on MPLS label.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the port 1/1/x1 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select MPLS Label. **d.** Set the value to 4 and the Position to 1 **5.** Click Save.
7	Create another second level map to filter on MPLS label.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type map2 in the Alias field.

Step	Description	Command
		▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the port 1/1/x4 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select MPLS Label. **d.** Set the value to 3 and the Position to 1 **5.** Click Save.

Combine APF with GigaSMART Operations

APF can also be combined with other GigaSMART functions including Header Stripping, Packet Slicing or Masking, De-duplication and FlowVUE. This provides network administrators and operators to perform a second layer of filtering in combination with the GigaSMART tool optimization and packet manipulation operations.

In the following example, operators can distribute traffic to monitoring tools based on decapsulated contents, more specifically, after Header stripping VXLAN:

- Identifying and forwarding traffic from/to ip 1.1.1.1 from the decapsulated packets to monitoring tool connected to tool port 1/1/x1
- Identifying and forwarding traffic from/to ip 1.1.1.2 in the decapsulated packets to monitoring tool connected to tool port 1/1/x4

NOTE: This can be applied to any protocol that is supported through header-stripping, for example:

- GTP, VXLAN, ISL, MPLS, MPLS+VLAN, VLAN, VN-Tag, fabric-path.
- This is also supported for Gigamon tunnel decapsulation.

Step	Description	Command
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3 and select Tool for the port s 1/1/x4 and 1/1/x1 **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Select an engine port in the Port List field. For example, 1/1/e1 **5.** Click Save.
3	Configure the GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations. **2.** Click New. **3.** Type gsfil_vxlanhs in the Alias field. **4.** Select gsg1 from the GS Groups list. **5.** Select Adaptive Packet Filtering from the GigaSMART Operations (GSOP) list and Enable. **6.** Select Strip Header from the GigaSMART Operations (GSOP) list and select VXLAN. **7.** Click Save.
4	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp1 in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map to forward VXLAN traffic to the virtual port. VXLAN accepts destination UDP ports 8472 and 4789. Starting in software version 4.5.01, VXLAN also accepts destination UDP port 48879.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type to_vp in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the network port for the Source. For example, 1/1/x3 ▪ Select the virtual port vp1 for the Destination. **4.** Add a rule.

Step	Description	Command
		a. Click Add a Rule. **b.** Select Pass. **c.** Select Port Source and set the port value to 8472. **4.** Click Save.
6	Create a second level map to filter on source and destination IP (bi-directional).	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x1 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Source. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Add a Rule 2. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Click Save.
8	Create another second level map to filter on source and destination IP (bi-directional).	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x1 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add Rule 1. **a.** Click Add a Rule.

Step	Description	Command
		b. Select Pass. **c.** Select IPv4 Source. **d.** Enter 1.1.1.2 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Add a Rule 2. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination. **d.** Enter 1.1.1.2 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Click Save.

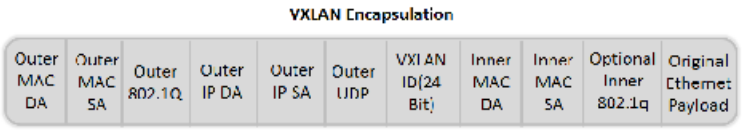
Conditional Header Stripping

Another use-case that can be addressed leveraging the flexibility of APF would be the capability to header strip packets based on specific contents found across the packet including the inner packet contents. Since the APF rules are enforced before any other GigaSMART operation, operators can filter based on encapsulation protocol values and /or encapsulated (original) packet contents and apply conditional Header Stripping operations.

The following example shows how an end-user can filter and strip out outer VXLAN headers for a subset of the traffic based on inner IP addresses, while sending the rest of the traffic “as-is” to monitoring tools that need the VXLAN headers for traffic analysis, as follows.

- Identifying and forwarding traffic from/to ip 1.1.1.1 in the inner / encapsulated packets to monitoring tool connected to tool port 1/1/x1 *after* Header Stripping VXLAN.
- Identifying and forwarding traffic from/to ip 1.1.1.2 in the inner / encapsulated packets to monitoring tool connected to tool port 1/1/x4 *without* stripping the VXLAN header.

NOTE: This can be applied to any GigaSMART operation. While this example shows filtering based on inner packet contents, conditional SMART operations can be applied by filtering on encapsulation headers as well.



NOTE: This can be applied to any protocol that is supported through Header Stripping. GTP, VXLAN, ISL, MPLS, MPLS+VLAN, VLAN, VN-Tag, and fabric-path are all supported, as is Gigamon tunnel decapsulation.

Step	Description	Command
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3 and select Tool for the port s 1/1/x4 and 1/1/x1 **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups (GSOP) > GigaSMART Groups. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Select an engine port in the Port List field. For example, 1/1/e1. **5.** Click Save.
3	Configure the GigaSMART operations.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations and create two GigaSMART Operations. **2.** Create the first operation. **a.** Click New. **b.** Type gsfil_vxlanhs in the Alias field. **c.** Select gsg1 from the GigaSMART Groups list. **d.** Select Adaptive Packet Filtering from the GigaSMART Operations list and Enable. **e.** Select Strip Header from the GigaSMART Operations list and select VXLAN. **f.** Click Save. **7.** Create second first operation. **a.** Click New. **b.** Type gsfil apf in the Alias field. **c.** Select gsg1 from the GigaSMART Groups list. **d.** Select Adaptive Packet Filtering from the GS Operations (GSOP) list and Enable. **e.** Click Save.
4	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp1 in the Alias field.

Step	Description	Command
		4. Select gsg1 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map to forward VXLAN traffic to the virtual port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type to_vp in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the network port for the Source. For example, 1/1/x3 ▪ Select the virtual port vp1 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Port Source and set the port value to 8472. **4.** Click Save.
6	Create a second level map to filter on source and destination IP (bi-directional), using first GigaSMART operation.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x1 for the Destination. ▪ Select **gsfil_vxlanhs** from the GSOP list. **4.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Source. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Add a Rule 2. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask

Step	Description	Command
		f. Set Position to 2. **7.** Click Save.
7	Create another second level map to filter on source and destination IP (bi-directional), using second GigaSMART operation.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x4 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Source. **d.** Enter 1.1.1.2 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Add a Rule 2. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination. **d.** Enter 1.1.1.2 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Click Save.

Facilitate Overlapping Rules

Because APF is implemented as a second level map operation, APF can also be leveraged for implementing basic overlapping rules. For the same incoming input stream, a copy of the traffic can be sent out to a group of monitoring tools while a refined subset of the traffic stream can be sent to a different set of monitoring tools. Typically overlapping rules would be implemented by combining APF with the patented Flow Mapping® technology.

Note that Role-Based Access control in the case of APF is applied at the gsgroup / **e** port.

In the following example, for the same input stream:

- HTTP traffic is identified and distributed to a monitoring tool connected to tool port 1/1/x1.
- At the same time, the same stream of HTTP packets are being sent out after slicing unwanted packet contents to a different monitoring tool connected to tool port 1/1/x4.

Step	Description	Command
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3 and select Tool for the port s 1/1/x4 and 1/1/x1 **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Select an engine port in the Port List field. For example, 1/1/e1. **5.** Click Save.
3	Configure the GigaSMART operations.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOPS) > Operations and create two GigaSMART Operations. **2.** Create the first operation. **a.** Click New. **b.** Type gsfil in the Alias field. **c.** Select gsg1 from the GigaSMART Groups list. **d.** Select APF from the GigaSMART Operations list and Enable. **e.** Click Save. **6.** Create second operation. **a.** Click New. **b.** Type gsfil_slice in the Alias field. **c.** Select gsg1 from the GigaSMART Groups list. **d.** Select APF from the GigaSMART Operations (GSOP) list and Enable. **e.** Select Slicing from the GigaSMART Operations (GSOP) list and select None. Set Offset to 150. **6.** Click Save.
4	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New.

Step	Description	Command
		3. Enter vp1 in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map to forward traffic to the virtual port. Port 1/1/x1 and vertual ort vp1 are sent destination port 80 traffic, which is HTTP.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type to_vp in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the network port 1/1/x3 for the Source. ▪ Select the virtual port vp1 and the tool port 1/1/x1 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Port Source and set the port value to 2152. **4.** Click Save.

Step	Description	Command
6	Create a second level map to filter on HTTP traffic and slice it.	**1.** Click New. **2.** Configure the map. ▪ Type map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the virtual tool port 1/1/x4 for the Destination. ▪ Select gsfil_slice form the GSOP list. **3.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IP Version **d.** Set Version to v4 **e.** Set Position to 1 **6.** Click Save.
7	Create another second level map for the rest of the traffic.	**1.** Click New. **2.** Configure the map. ▪ Type map2 in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the virtual port 1/1/x1 for the Destination. ▪ Select gsfil from the GSOP list. **3.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IP Version. **d.** Set Version to v4. **e.** Set Postilion to 4 **6.** Click Save.

In the following example, for the same traffic stream, TCP traffic is sent to one monitoring tool while forwarding a subset of TCP flows specific to HTTP to another monitoring tool connected to tool port 1/1/x4.

Step	Description	Command
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3 and select Tool for the port s 1/1/x4 and 1/1/x1 **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Select an engine port in the Port List field. For example, 1/1/e1. **5.** Click Save.
3	Configure the GigaSMART operations.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOPS) > GigaSMART Operation and create two GigaSMART Operations. **2.** Click New. **3.** Type gsfil in the Alias field. **4.** Select gsg1 from the GS Groups list. **5.** Select Adaptive Packet Filtering from the GS Operations list and Enable. **6.** Click Save.

Step	Description	Command
4	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp1 in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map to forward TCP traffic to the virtual port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type to_vp in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the network port 1/1/x3 for the Source. ▪ Select the virtual port vp1 and the tool port 1/1/x4 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Protocol and set the value to TCP. **4.** Click Save.
6	Create a second level map to filter on HTTP traffic.	**1.** Click New. **2.** Configure the map. ▪ Type map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the virtual tool port 1/1/x1 for the Destination. ▪ Select gsfil form the GSOP list. **3.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination. **d.** Set Position to 2 **e.** Set the port value to 80. **6.** Add Rule 2. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Source. **d.** Set Position to 2 **e.** Set the port value to 80.

Step	Description	Command
		6. Click Save.

GigaSMART Load Balancing

GigaSMART Enhanced Load Balancing require a separate license. Stateless Load Balancing is included with Base licenses. Stateful Load Balancing for GTP is included with the GTP Filtering & Correlation license.

Stateful Load Balancing for ASF is included with the Application Session Filtering (ASF) license. Stateful Load Balancing for tunnel is included with the Advanced Tunneling license (GigaVUE-HC3), and Tunneling license (GigaVUE-HC1)

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

Load balancing distributes GigaSMART outgoing traffic to multiple tool ports or multiple tunnel endpoint destinations. In this way, traffic processed by GigaSMART is shared.

Stateful load balancing distributes GigaSMART processed traffic to multiple tool ports or tunnel endpoints based on GigaSMART application-specific flow sessions. Stateless load balancing distributes GigaSMART processed traffic to multiple tool ports or tunnel endpoints based on hash values generated from predefined protocol fields in the packet.

Load balancing operations to tool ports can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to Groups of GigaSMART Engine Ports for details.

NOTE: You can combine the front and rear Gen 3 modules of GigaVUE-HC1-Plus. When combined, weighted load balancing is automatically applied for traffic processing.

The following sections describe the available load balancing schemes:

- GigaSMART Load Balancing
- Stateless Load Balancing
- Enhanced Load Balancing

Stateful Load Balancing

Stateful load balancing distributes GigaSMART processed traffic to multiple tool ports or tunnel endpoints based on GigaSMART application-specific flow sessions.

With stateful load balancing, packets belonging to the same flow session maintained by GigaSMART applications are forwarded to the same tool port or tunnel endpoint within a port group.

NOTE: When stateful tunnel load-balancing is configured in Generation 3 GigaSMART card (SMT-HC1-S), the load balancing is undefined until the end-points accessibility become stable.

Use the **GigaSMART Operations (GSOP)** page to configure load balancing. Specify one stateful application within a group of GigaSMART operations and specify a load balancing metric.

The following are the supported stateful applications:

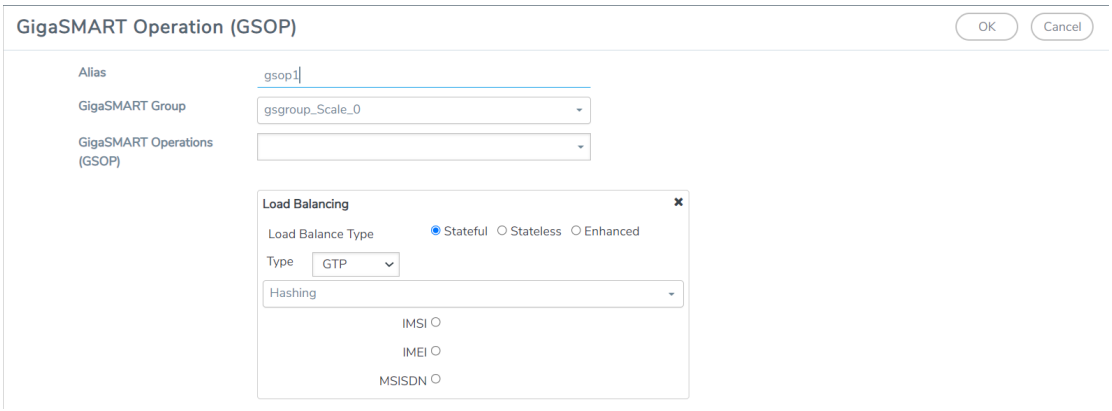
Application	Reference
GTP	GigaSMART GTP and CUPS Correlation
Application Session Filtering (ASF)	GigaSMART Application Session Filtering ASF and Buffer ASF
Tunnel	GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation
SIP	GigaSMART SIP/RTP Correlation

To select Stateful Load Balancing, do the following:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP)**, and then click **New**.
2. Specify an alias in the **Alias** field.
3. Click in the **GigaSMART Group** field and select a GigaSMART Group.
4. Click in the **GigaSMART Operations (GSOP)** field and select **Load Balancing** from the drop-down list.
5. Select **Stateful**.
6. For **Type**, select one stateful application within a group of GigaSMART operations.

   GTP, **ASF**, **Tunnel**, and **SIP** are the supported stateful applications.

7. Specify a load balancing metric. For example, **Hashing** as shown in the following figure.
8. Click Ok.



For details on stateful load balancing, refer to the following sections:

- Stateful Load Balancing Metrics
- Hashing Key Support
- Configure Stateful Load Balancing

Stateful Load Balancing Metrics

The load balancing metrics available for stateful load balancing are described in the following table.

For weighted metrics, such as Weighted Least Bandwidth, you can either define a weight for each port such as 5,10, 25, 50, or you can use link speed: 1 for 1Gb, 10 for 10Gb, 40 for 40Gb, 100 for 100Gb. Use the **Port Groups** configuration page to select weight and use the **Port** configuration page to select link speed.

NOTE: Only the traffic from the stateful application (for example, GTP, ASF, or tunnel) is used to perform load balancing. Other traffic in the map that does not match the application's filter rule is excluded.

Metric	Description
Least Bandwidth	A tool port is selected from a port group based on the least bits per second load to the port. To compensate for bursty traffic, the history of the last 10 second bandwidth is considered on the load balancing decision.

Metric	Description
	This metric is not supported for tunnel.
Weighted Least Bandwidth	A tool port is selected from a port group based on the least bits per second load to the port, as described under Least Bandwidth. This metric is not supported for tunnel. If this metric is selected, link speed is considered in addition to the bandwidth of the port. If this metric is not selected, the weight configured for each port in the port group is considered in addition to the bandwidth of the port.
Least Packet Rate	A tool port is selected from a port group based on Least Packet Rate. To compensate for bursty traffic, the history of the last 10 second packet count is considered on the load balancing decision.
Weighted Least Packet Rate	A tool port is selected from a port group based on Least Packet Rate, as described under Least Packet Rate. With Weighted Least Packet Rate, if a port has a higher weight, it will be sent more traffic. If this metric is selected, link speed is considered with packet rate. If this metric is not selected, the weight configured for each port in the port group is considered with packet rate.
Round Robin	A tool port is selected from a port group based on round robin.
Weighted Round Robin	A tool port is selected from a port group based on least packet rate, as described under Round Robin. If this metric is selected, link speed is considered with Round Robin. If this metric is not selected, the weight configured for each port in the port group is considered with Round Robin.
Least Connection	A tool port is selected from a port group based on the current Least Connection assigned to each tool port. The port with the least number of connections assigned is selected. **Note**: The meaning of connection is defined by the application.
Weighted Least Connection	A tool port is selected from a port group based on the current Least Connection assigned to each tool port, as described under Least Connection. If this metric is selected, link speed is considered with Least Connection. If this metric is not selected, the weight configured for each port in the port group is considered with Least Connection. **Note**: The meaning of connection is defined by the application.
Least Cumulative Traffic	A tool port is selected from a port group based on the least total bytes sent to each tool port. The port with the least number of connections assigned is selected. **Note**: The meaning of connection is defined by the application.
Weighted Least Cumulative Traffic	A tool port is selected from a port group based on the least total bytes sent to each tool port, as described under Least Cumulative Traffic.

Metric	Description
	If this metric is selected, link speed is considered with least cumulative traffic. If this metric is not selected, the weight configured for each port in the port group is considered with Least Cumulative Traffic.
Hashing	A tool port is selected from a port list based on hashing of data provided by the GSOP application, which is normally extracted from the packet. The values for hashing key are: IMSI GTP key (imsi), IMEI GTP key (imei), and MSISDN GTP key (msisdn). The hashing key only applies to the GTP stateful application. Refer to *Hashing Key Support* on page 672 for details.

Hashing Key Support

The following table describes the support for GTP hashing key.

GTP Key	Hashing	(Weighted) Least Bandwidth	(Weighted) Least Packet Rate	(Weighted) Least Round Robin	(Weighted) Least Connection	(Weighted) Least Cumulative Traffic
IMSI	Supported	Supported	Supported	Supported	Supported	Supported
IMEI	Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
MSISDN	Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported

Configure Stateful Load Balancing

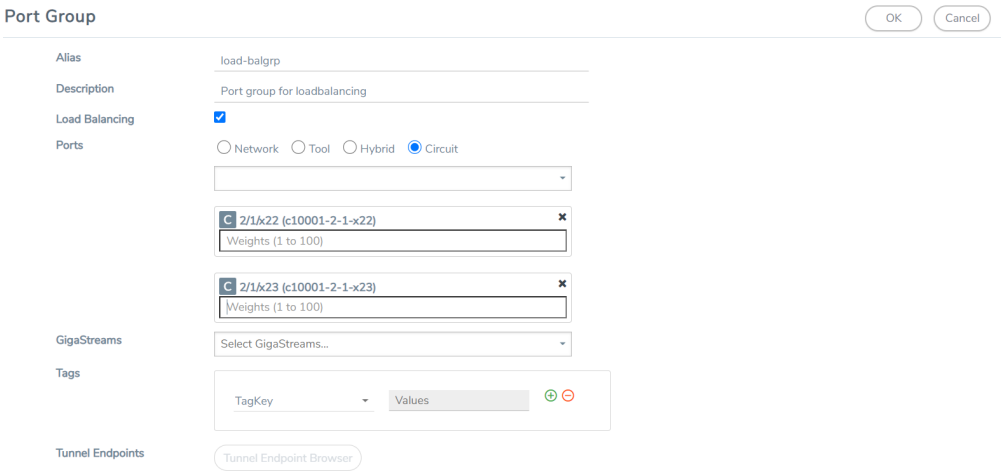
Use the Port Group page to specify the list of tool ports or tunnel endpoints for stateful load balancing and to enable load balancing in a port group.

To enable load balancing in a port group, do the following:

1. Select **Ports > Port Groups > All Port Groups**.
2. Click **New**.
3. In the Alias filed, enter an alias. For example, load-balgrp.
4. Select the port type.
5. If applicable, enable **Load Balancing.** Load Balancing is available only for specific port types—Tool, Hybrid, and Circuit.

6. Define the weights for each of the ports.

NOTE: Weight determines the traffic sent to a particular port. The weight of the individual ports must be less than 100. The combined value of the ports can be greater than 100, as the actual load balancing ratio is computed with individual values divided by the combined value



7. Click **OK** to save the configuration.

NOTES:

- Up to 50 load balancing port groups are supported.
- Ports within port groups can have different rates.

Examples

Refer to the following examples:

- Example 1: GigaSMART Stateful Load Balancing
- Example 2: GigaSMART Stateful Load Balancing

For an example of load balancing on L2GRE encapsulation tunnel, refer to GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation.

Example 1: GigaSMART Stateful Load Balancing

Example 1 configures stateful load balancing of GigaSMART GTP traffic among tool ports 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4 based on bandwidth with different weights for each port. The same subscriber (imsi) traffic will be forwarded to the same tool port. GTP-c packets are

replicated to all tool ports.

Task	Description	UI Steps
1.	Create a port group, specify the tool ports for load balancing, and weights for each tool port.	**a.** Select Ports > Port Groups. **b.** Click New **c.** Type portgrp1 in the Alias field. **d.** Enable Load Balancing **e.** Click in the Ports field to select the tool ports for the group. For example, 1/1/x6,1/1/x7,1/2/x3, and 1/2/x4. **f.** Specify the weights for each port as follows: weight 1/1/x6 5 weight 1/1/x7 10 weight 1/2/x3 20 weight 1/2/x4 10 **g.** Click Save.
2.	Create a GigaSMART group and specify a port and enable replicate GTP-c packets to all tool ports in the load balancing port group.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type gsgrp1 in the Alias field. **d.** Select the engine port. For example 1/3/e1. **e.** Under the Load Balance, select Replicate GTP-c. **f.** Click Save.
3.	Create a GSOP, including GTP application and load balancing metric.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Type gsop1 in the Alias field **d.** Select gsgrp1 from the GigaSMART Groups list. **e.** Select the GigaSMART Operation Flow Filtering **f.** Select the GigaSMART Operation Load Balancing and set the operation as follows: o Select Stateful. o Select GTP for Type. o Select Weighted Least Bandwidth for the metric.
4.	Create a virtual port and associate it with the GigaSMART group.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > Virtual Ports. **b.** Click New. **c.** Type vp1 in the Alias field.

Task	Description	UI Steps
		d. Select gsgrp1 from the GigaSMART Group list. **e.** Click OK.
5.	Create an ingress (first level) map. Note the following: ▪ You can specify only one port group as part of the map tool port in the **to** statement. ▪ You can define the same load balancing port group in multiple maps, however, the load balancing metrics defined in the GSOPs on those maps have to be the same. ▪ You cannot use a shared collector map for load balancing. **Note**: In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map11 in the Alias field. **d.** Select First Level for Type and By Rule for Subtype. **e.** Select port 1/1/x1 for the Source. **f.** Select virtual port vp1 for the Destination. **g.** Add Rule 1. o Click Add a Rule o Select Pass o Select Port Destination for the condition. o Set the port value to 2123. **h.** Add Rule 2. o Click Add a Rule o Select Pass o Select Port Destination for the condition. o Set the port value to 2125. **i.** Click Save.
6.	Create a second level map.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map22 in the Alias field. **d.** Select Second Level for Type and By Rule for Flow Filter. **e.** Select virtual port vp1 for the Source. **f.** Select port group portgrp1 for the Destination. **g.** Select gsop1 from the GSOP list. **h.** Click Add a Rule o Select Pass o Select GTP IMSI for the condition. o Enter 234567* for IMSI. o Select Any for Version. **i.** Click Save.

Example 2: GigaSMART Stateful Load Balancing

Example 2 configures stateful load balancing of GigaSMART GTP traffic among tool ports 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4 based on hashing of the imei value. The same device ID (imei) traffic will be forwarded to the same tool port. GTP-c packets are replicated to all tool ports.

Task	Description	Steps
1.	Create a port group and specify the tool ports for load balancing.	**a.** Select Ports > Port Groups. **b.** Click New **c.** Type portgrp1 in the Alias field. **d.** Enable Load Balancing **e.** Click in the Ports field to select the tool ports for the group. For example, 1/1/x6,1/1/x7,1/2/x3, and 1/2/x4. **f.** Click OK.
2.	Create a GigaSMART group and specify ports and enable replicate GTP-c packets to all tool ports in the load balancing port group.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type gsgrp1 in the Alias field. **d.** Select the engine port. For example 1/3/e1. **e.** Under the Load Balance, select Replicate GTP-c. **f.** Click OK.
3.	Create a GSOP, including GTP application and load balancing metric.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations. **b.** Click New. **c.** Type gsop1 in the Alias field **d.** Select gsgrp1 from the GigaSMART Groups list. **e.** Select the GigaSMART Operation Flow Filtering **f.** Select the GigaSMART Operation Load Balancing. **g.** Configure Load Balancing as follows: o Select Stateful. o Select GTP for Type. o Select Hashing for the metric. o Select IMEI

Task	Description	Steps
		h. Click OK.
4.	Create a virtual port and associate it with the GigaSMART group.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > Virtual Ports. **b.** Click New. **c.** Type vp1 in the Alias field. **d.** Select gsgrp1 from the GigaSMART Groups list. **e.** Click OK.
5.	Create an ingress (first level) map. Note the following: ▪ You can specify only one port group as part of the map tool port in the **to** statement. ▪ You can define the same load balancing port group in multiple maps, however, the load balancing metrics defined in the GSOPs on those maps have to be the same. ▪ You cannot use a shared collector map for load balancing. **Note**: In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map11 in the Alias field. **d.** Select First Level for Type and By Rule for Subtype. **e.** Select port 1/1/x1 for the Source. **f.** Select virtual port vp1 for the Destination. **g.** Add Rule 1. o Click Add a Rule o Select Pass o Select Port Destination for the condition. o Set the port value to 2123. **h.** Add Rule 2. o Click Add a Rule o Select Pass o Select Port Destination for the condition. o Set the port value to 2125. **i.** Click Save.
6.	Create a second level map.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map22 in the Alias field. **d.** Select Second Level for Type and By Rule for Flow Filter. **e.** Select virtual port vp1 for the Source. **f.** Select port group portgrp1 for the Destination. **g.** Select gsop1 from the GSOP list. **h.** Click Add a Rule o Select Pass o Select GTP IMSI for the condition. o Enter 234567* for IMSI. o Select Any for Version. **i.** Click Save.

Stateless Load Balancing

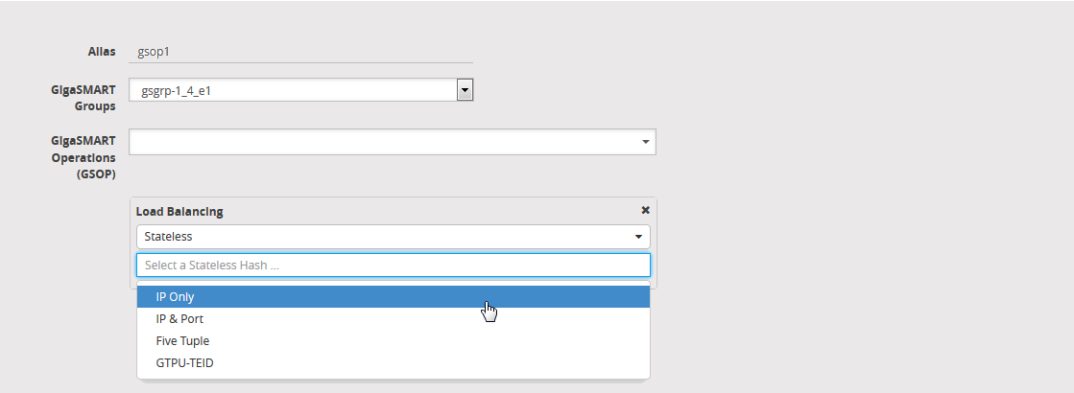
Stateless load balancing distributes GigaSMART processed traffic to multiple tools based on predefined protocol fields in the packet.

Unlike stateful load balancing, stateless load balancing can be configured together with most other GigaSMART operations or as a separate GigaSMART operation to provide more flexible traffic distribution options over what is available from a tool GigaStream. Packets processed by stateless load balancing are forwarded to a tool port within a port group.

Stateless load balancing supports packets with MPLS encapsulation and IEEE 802.1 Q-in-Q VLAN tags.

To select stateless load balancing, do the following:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) GigaSMART Operation**, and then click **New**.
2. Type an alias in the **Alias** field.
3. From the **GigaSMART Groups** drop-down list, select a GigaSMART group.
4. From the **GigaSMART Operations (GSOP)** drop-down list, select **Load Balancing**.
5. Configure Load Balancing:
   - Select **Stateless**
   - Specify a load balancing metric. For example, **IP Only** as shown in the following figure.
6. Click **Save**.



For details on stateless load balancing, refer to the following sections:

Stateless Load Balancing Metrics

The load balancing metrics available for stateless load balancing are described in the following table.

A tool port is selected from a port list based on hashing. The fields to be hashed are described in the table.

To specify a metric for stateless load balancing:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.
2. Click **New** to create a new GigaSMART Operation or **Edit** to modify an existing one.
3. Select **Load Balancing** from the **GigaSMART Operations (GSOP)** drop-down list and configure load balancing as follows.
   a. Select **Stateless**
   b. Select one of the hashing metrics: IP Only, IP & PORT, Five Tuple, or GTPU-TEID.
      Table 7: Stateless Load Balancing Metrics describes each of the metrics.
4. Click **Save**.

Table 7: Stateless Load Balancing Metrics

Metric	Description
IP Only	The source IP and destination IP addresses.
IP & Port	The source IP and destination IP addresses, and Layer 4 source port and destination port numbers.
Five Tuple	The source IP and destination IP addresses, source port and destination port numbers, and protocol field in the IP header.

Metric	Description
GTPU-TEID	The GTP-u tunnel identifier (ID). **Note**: There is no inner or outer field location for **GTPU-TEID**.
outer	The first occurrence of header or field. For example, **IP Only outer** is the first IP header in the packet, which could be IPv4 or IPv6.
inner	The second occurrence of header or field. For example, **ip-only inner** is the second IP header in the packet. The first IP header could be IPv4 or IPv6, as follows: ▪ IPv4-IPv4—**IP Only inner** is the IP addresses in the second IPv4 header ▪ IPv6-IPv4—**IP Only inner** is the IP addresses in the IPv4 header ▪ IPv4-IPv6—**IP-Only inner** is the IP addresses in the IPv6 header The supported IP encapsulation types are: IP-in-IP, VXLAN, GTP, GRE(including L2GRE), and ERSPAN.

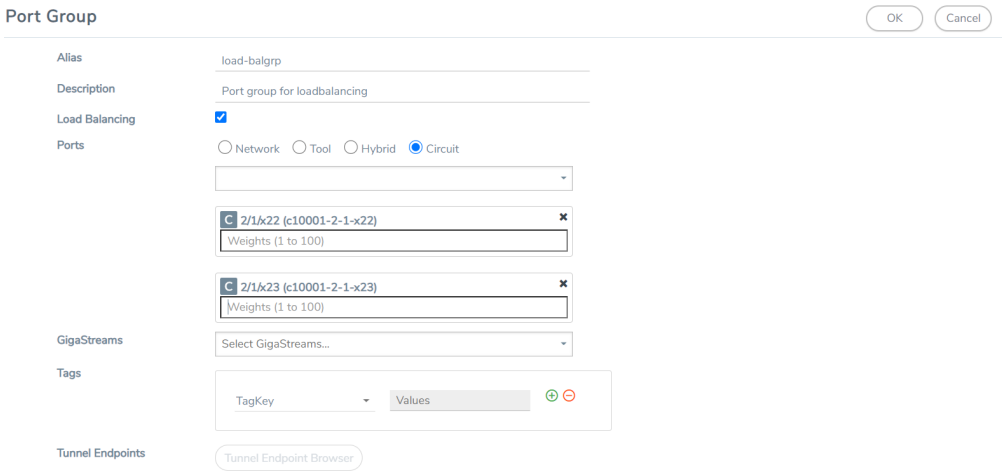
NOTE: On Gen3 GigaVUE HC Series, L2GRE traffic that uses an 8-byte GRE header (4-byte base header + 4-byte key) follows the same inner-header hashing behavior as other supported tunnels. When you select an inner stateless load balancing metric (for example, IP Only inner, IP & Port inner, or Five Tuple inner), L2GRE packets with and without a GRE key are hashed based on the inner IP and port fields. In earlier releases, L2GRE packets with a 4-byte GRE key were always hashed on the outer header, which could result in uneven distribution across GigaSMART CPU cores when many flows shared the same outer 5-tuple.

Configure Stateless Load Balancing

To configure stateless load balancing, specify the group of tool ports and enable load balancing in a port group.

1. Select **Ports > Port Groups > All Port Groups**.
2. Click **New**.
3. Type an alias in the **Alias** field. For example, load-balgrp.
4. Select the Port Type.
5. If applicable, enable **Load Balancing.** Load Balancing is available only for specific port types—Tool, Hybrid, and Circuit.
6. Define the weight for each of the ports used.

NOTE: Weight determines the traffic sent to a particular port. The weight of the individual ports must be less than 100. The combined value of the ports can be greater than 100, as the actual load balancing ratio is computed with individual values divided by the combined value



7. Click **OK** to save the configuration.

Notes:

- Up to 50 load balancing port groups are supported, with a maximum of 16 ports for each group.
- Ports within port groups must be on the same chassis.
- Ports within port groups can have different rates.

Examples

Refer to the following examples:

- Example 1: GigaSMART Stateless Load Balancing
- Example 2: GigaSMART Stateless Load Balancing
- Example 3: GigaSMART Stateless Load Balancing

For an example of load balancing on L2GRE encapsulation tunnel, refer to GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation.

Example 1: GigaSMART Stateless Load Balancing

Example 1 configures stateless load balancing of traffic among tool ports 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4 after slicing the packet to an offset of 70 bytes.

Task	Description	Steps
1.	Create a port group and specify the tool ports for load balancing.	**a.** Select Ports > Port Groups. **b.** Click New **c.** Type portgrp1 in the Alias field. **d.** Click in the Ports field to select the tool ports for the group. For example, 1/1/x6,1/1/x7,1/2/x3, and 1/2/x4. **e.** Enable Load Balancing **f.** Click OK.
2.	Configure a GigaSMART group and associate it with GigaSMART engine ports.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type gsgrp1 in the Alias field. **d.** Select the engine ports. For example 1/3/e1 and 1/3/e2. **e.** Click OK.
3.	Create a GSOP, with load balancing.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > Operations. **b.** Click New. **c.** Type lbiponlyouter in the Alias field **d.** Select gsgrp1 from the GigaSMART Groups list. **e.** Select the GigaSMART Operation Load Balancing. **f.** Configure Load Balancing as follows: o Select Stateless o Select IP Only Outer for the hash metric **g.** Click OK.
4.	Create a first level map. Note the following: ▪ You can specify only one port group as part of the map tool port in the **to** statement. ▪ You can define the same load	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map1 in the Alias field. **d.** Select First Level for Regular and By Rule for Subtype. **e.** Select port 1/1/x1 for the Source.

Task	Description	Steps
	balancing port group in multiple maps, however, the load balancing metrics defined in the GSOPs on those maps have to be the same. ▪ You cannot use a shared collector map for load balancing.	**f.** Select port group portgrp1 for the Destination. **g.** Select lbiponlyouter from the GSOP list. **h.** Click Add a Rule o Select Pass o Select IP Version for the condition. o Select v4 for Version. **i.** Click Save.

Example 2: GigaSMART Stateless Load Balancing

Example 2 configures stateless load balancing of GTP traffic among tool ports 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4. Data packets with the same GTP-u tunnel ID will be forwarded to the same tool port.

Task	Description	Steps
1.	Create a port group and specify the tool ports for load balancing.	**a.** Select Ports > Port Groups. **b.** Click New **c.** Type portgrp1 in the Alias field. **d.** Enable Load Balancing **e.** Click in the Ports field to select the tool ports for the group. For example, 1/1/x6,1/1/x7,1/2/x3, and 1/2/x4. **f.** Click OK.
2.	Configure a GigaSMART group and associate it with GigaSMART engine ports.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type gsgrp1 in the Alias field. **d.** Select the engine ports. For example 1/3/e1 and 1/3/e2. **e.** Click OK.
3.	Create a GSOP, including load balancing metric.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations. **b.** Click New.

Task	Description	Steps
		c. Type gsop1 in the Alias field **d.** Select gsgrp1 from the GigaSMART Groups list. **e.** Select the GigaSMART Operation Load Balancing. **f.** Configure Load Balancing as follows: o Select Stateless o Select GTPU-TEID for the hash metric **g.** Click OK.
4.	Create first level maps. Note the following: ▪ You can specify only one port group as part of the map tool port in the **to** statement. ▪ You can define the same load balancing port group in multiple maps, however, the load balancing metrics defined in the GSOPs on those maps have to be the same. ▪ You cannot use a shared collector map for load balancing.	Create the first map: **a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map1 in the Alias field. **d.** Select First Level for Regular and By Rule for Subtype. **e.** Select port 1/1/x1 for the Source. **f.** Select port group portgrp1 for the Destination. **g.** Click Add a Rule, and then select Pass **h.** Select IPv4 Protocol for the first condition. **i.** Select UDP for Value **j.** Select Port Destination for the second condition. **k.** Enter 2123 for the port value. **l.** Click OK. Create the second map: **a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map1 in the Alias field. **d.** Select First Level for Regular and By Rule for Subtype. **e.** Select port 1/1/x1 for the Source. **f.** Select port group portgrp1 for the Destination. **g.** Select gsop1 from the GSOP list. **h.** Click Add a Rule, and then select Pass **i.** Select IPv4 Protocol for the first condition. **j.** Select UDP for Value **k.** Select Port Destination for the second condition. **l.** Enter 2152 for the port value. **m.** Click OK.

Example 3: GigaSMART Stateless Load Balancing

Example 3 configures stateless load balancing of HTTP on GTP traffic among tool ports 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4. Data packets with the same inner IP will be forwarded to the same tool port.

Task	Description	Steps
1.	Create a port group and specify the tool ports for load balancing.	**a.** Select Ports > Port Groups. **b.** Click New **c.** Type portgrp1 in the Alias field. **d.** Enable Load Balancing **e.** Click in the Ports field to select the tool ports for the group. For example, 1/1/x6,1/1/x7,1/2/x3, and 1/2/x4. **f.** Click OK.
2.	Configure a GigaSMART group and associate it with GigaSMART engine ports.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type gsgrp1 in the Alias field. **d.** Select the engine ports. For example 1/3/e1 and 1/3/e2. **e.** Click OK.
3.	Create a GSOP, including load balancing metric.	**f.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **g.** Click New. **h.** Type gsop1 in the Alias field **i.** Select gsgrp1 from the GigaSMART Groups list. **j.** Select the GigaSMART Operation Load Balancing. o Select Stateless o Select IP Only Inner for the hash metric **k.** Click OK.
4.	Create first level maps. Note the following: o You can specify only one port group as part of the map tool port in the **to** statement.	Create the first map: **a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map1 in the Alias field.

Task	Description	Steps
	o You can define the same load balancing port group in multiple maps, however, the load balancing metrics defined in the GSOPs on those maps have to be the same. o You cannot use a shared collector map for load balancing.	**d.** Select Regular for Type and By Rule for Subtype. **e.** Select port 1/1/x1 for the Source. **f.** Select port group portgrp1 for the Destination. **g.** Click Add a Rule, and then select Pass **h.** Select IPv4 Protocol for the first condition. **i.** Select UDP for Value **j.** Select Port Destination for the second condition. **k.** Enter 2123 for the port value. **l.** Click Save. Create the second map: **a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map1 in the Alias field. **d.** Select First Level for Regular and By Rule for Subtype. **e.** Select port 1/1/x1 for the Source. **f.** Select port vp1 for the Destination. **g.** Select gsop1 from the GSOP list. **h.** Click Add a Rule, and then select Pass **i.** Select IPv4 Protocol for the first condition. **j.** Select UDP for Value **k.** Select Port Destination for the second condition. **l.** Enter 2152 for the port value. **m.** Click Save.
5.	Create the second level map.	Create the third map: **a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map22 in the Alias field. **d.** Select Second Level for Type and By Rule for Subtype. **e.** Select virtual vp1 for the Source. **f.** Select port group portgrp1 for the Destination. **g.** Select gsop1 from the GSOP list. **h.** Click Add a Rule, and then select Pass **i.** Select IPv4 Destination for the first condition.

Task	Description	Steps
		j. Enter 80 for the destination value. **k.** Select 2 for position. **l.** Click Save.

Enhanced Load Balancing

Required License: Enhanced Load Balancing

GigaSMART Enhanced Load Balancing supports evenly distributed traffic among multiple tool ports based on one or more user defined fields. When a tool port fails, the traffic is redistributed just for that tool port to other member tool ports. When the failed tool port recovers, the traffic that was redistributed is restored to the recovered tool port. Traffic across other member tool ports remain undisturbed during this process.

Traffic Handling and Load Balancing Distribution

- Non GTP traffic and (subsequent) fragmented packets to be load balanced to all tool ports based on outer IP.
- Rebalance traffic when the following occurs:
  - Tool port goes up or down
  - Tool port group membership changes
  - When a tool port fails redistribute the traffic just for that tool port and when the failed tool port recovers restore traffic just for that tool port

Enhanced Load Balancing Metrics

GigaSMART provides configuration support for a new enhanced load balancing (enhanced-lb) app. The enhanced-lb app allows users to define the fields used for load balancing.

Enhanced Load Balancing supports the following hashing metrics:

- inner IP address
- outer IP address
- inner L4 port

- outer L4 port
- GPRS Tunnel Endpoint Identifier (TEID)

Configure Enhanced Load Balancing

To configure enhanced load balancing, do the following:

1. Select >**Physical**.
2. Click on a node you want apply enhance load balancing.
3. Select **GigaSMART>Enhanced Load Balancing**. The enhance load balance screen displays.
4. Click **New**. The enhanced load balancing screen appears.
5. Enter an Alias.
6. Select the Hash Field Name. The following options are available:

   - IP

   - L4 Port

   - GTP-U TEID
7. Select Hash Field Location. Options:

   - inner

   - outer

   NOTE: Use the “+” or “-” icons to add and delete hash fields.

Hash Field Name	Hash Field Location
IP	**Inner**
IP	**Outer**
L4 Port	**Outer**
L4 Port	**Inner**
GTP-U TEID	**None**

You can configure up to five (5) different hash fields and location.

8. Click **OK**.

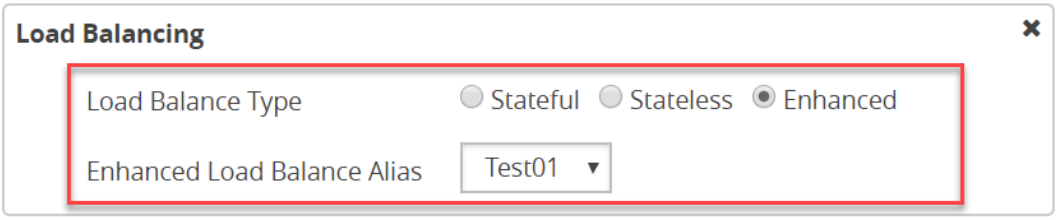
NOTE: To view details about the enhanced load balancing parameters, click the **Alias**. The detail/edit dialog box displays.

Configure GigaSMART Operation (GSOP)

1. Select > **Physical**.
2. Select **GigaSMART > GSOP**.
3. Click **New**.
4. Enter **Alias.**
5. Select a GSOP Group.
6. Select a GSOP type**: Load Balancing**.
7. Select GSOP type: Enhance load balancing. The Load Balancing dialog appears.



8. Select **Enhanced**.
9. Select an Enhanced Load Balance Alias from the drop-down. This is the enhanced load balance you previously created.



10. Click **OK**. The enhanced load balance GSOP is now available on the GSOP page.

GigaSMART De-Duplication

Required License: De-duplication

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

GigaSMART de-duplication detects duplicates of the following types:

- IPv4 packets

- IPv6 packets
- non-IP packets (including non-IPv4 and non-IPv6 packets)

Duplicates are packets in which the fields (including the headers and payload) are the same, with the exception of some fields such as Time-to-Live (TTL). For example, if two packets are identical except for TTL, they will be counted as duplicates.

Duplicate packets are common in network analysis environments where both the ingress and egress data paths are sent to a single output (for example, as a result of a SPAN operation on a switch). They can also appear when packets are gathered from multiple collection points along a path. GigaSMART de-duplication lets you eliminate these packets, only forwarding a packet once and thus reducing the processing load on your tools.

Feature Overview

There are two actions that can be specified for handling the duplicate packets detected:

- drop, which drops the duplicate packets
- count, which counts the duplicate packets, but does not drop them

A time interval can be configured within which an identical packet will be considered a duplicate. The greater the interval over which traffic can be checked for duplicates, the higher the accuracy of the de-duplication detection and subsequent elimination.

For example, if two of the same packets are seen in the specified time interval, the packets will be detected as duplicates. If one packet is seen in the time interval and another packet is seen in a later time interval, the packets will not be detected as duplicates.

For IPv4 and IPv6 packets, to determine if a packet is considered to be a duplicate, parts of the IP headers (Layer 3 and Layer 4), as well as part of the payload are compared.

For non-IP packets, a packet is considered to be a duplicate if it is identical.

Keep in mind the following when configuring GigaSMART de-duplication:

Feature	Description
Layer 2 Retransmissions Not Removed	Valid Layer 2 retransmissions are part of normal network behavior and are not removed by the de-duplication feature. Layer 2 retransmissions will show differences in the IP Window ID field.
Encapsulated Duplicates Not Removed	If the same packet is seen once with encapsulation (for example, GRE) and once without encapsulation, the GigaSMART will not detect and remove the duplicate.

Feature	Description
No NAT or PAT	Packets tapped on opposite sides of a NAT or PAT boundary will differ in the Network layer and will not be detected as duplicates.
MPLS and VLAN Tags	De-duplication properly parses VLAN and MPLS tags to get to the IP headers.
VN-Tag Packets	VN-Tag packets are treated as non-IP packets. User Header Stripping to strip VN-Tag to get to the IP headers for de-duplication. Refer to GigaSMART Header Stripping.
GigaSMART Engine Ports	De-duplication operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to Groups of GigaSMART Engine Ports for details.

De-duplication Configuration Steps

To configure de-duplication, use the following steps:

- Configure GigaSMART parameters on a specified GigaSMART group.
- Configure a GigaSMART operation.
- Configure a map that will use the de-duplication GigaSMART operation. This ties de-duplication to rules defined in a flow map, which applies the GigaSMART operation to specific traffic flows.

These steps are detailed in GigaSMART De-Duplication.

Configure GigaSMART Parameters for Packet De-duplication

Use the **de-dupe** section under GigaSMART Parameters on the GigaSMART Groups configuration page to configure options for GigaSMART de-duplication operations. The following table describes the configuring parameters for de-duplication on a specified GigaSMART group:

Parameter	Description
Action	Specifies whether duplicate packets are to be counted or dropped as follows: o Count– GigaSMART counts the duplicate packets, but does not drop them. o Drop– GigaSMART drops the duplicate packets. The default is drop.

Parameter	Description
IP Tclass **IP TOS** **TCP Sequence** **VLAN**	These options are useful when applying de-duplication operations to packets in a NAT environment. Different NAT implementations can change certain packet header fields (for example, the TCP sequence number). If you want to be able to detect duplicates without requiring that these fields match (ToS field, TCP sequence number, VLAN ID), you can disable the corresponding option. o **IP Tclass** – Ignore or include IPv6 traffic class. Use for IPv6. The default is include. o **IP TOS** – Ignore or include the IP ToS bits when detecting duplicates. Use for IPv4. The default is include. o **TCP Sequence** – Ignore or include the TCP Sequence number when detecting duplicates. The default is include. o **VLAN** – Ignore or include the VLAN ID when detecting duplicates. The default is ignore. Include means the field will be included when GigaSMART compares packets. Ignore means the field will be ignored when GigaSMART compares packets.
Timer <Value: 10-500000 µs>	Configures the time interval within which an identical packet will be considered a duplicate. The greater the interval over which traffic can be checked for duplicates, the higher the accuracy of the de-duplication detection and subsequent elimination. The default is 50,000µs. For example, if two same packets are seen in the specified time interval, the packets will be detected as duplicates. If one packet is seen in the time interval and another packet is seen in a later time interval, the packets will not be detected as duplicates. **NOTE:** Retransmissions are not counted as duplicates.

Example – GigaSMART De-duplication

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

This example shows the configuration steps for a de-duplication operation in which the GigaSMART application drops duplicate packets.

Task	Description	UI Steps
1.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**a.** From the device view, select GigaSMART > GigaSMART Groups. **b.** Click New. **c.** Type an alias in the Alias field. For example, gs2port1. **d.** Click in Port List field and select an engine port. For example, 2/1/e1. **e.** Go to Task 2.
2.	Configure parameters on the GigaSMART group.	**a.** Under the De-duplication section on the GigaSMART Group configuration page, set the parameters as the follows: Action: drop IP Tclass: Include IP TOS: Ignore TCP Sequence: Ignore Vlan: Ignore Timer (us) 500000 **b.** Click Save.
3.	Configure the GigaSMART operation for de-duplication and assign it to the GigaSMART group.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Type a name for the operation in the alias field. For example, testdedup. **d.** Select the GigaSMART Group create in task 1. **e.** Select De-duplication from the GigaSMART Operations (GSOP) list and select Enable. **f.** Click Save.
4.	Create a map.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type an alias in the Map Alias field that will help you identify this map. For example, testingdedup **d.** Select Regular and By Rule for the map type and subtype. **e.** Specify the network and tool ports in the Source and Destination fields. For example, 2/2/x4 and 2/2/x6 for Source and 2/2/x9 for Destination. **f.** From the GigaSMART Operation (GSOP) drop-down list, select the GigaSMART operation configured in Task 3. For example, testdedup in this example.

Task	Description	UI Steps
		g. Click Add a Rule under Map Rules and create the following rule: Select Pass, then select Bi Directional, and then select Port Source from the drop-down list and set the Min to 0 and Max to 443. **h.** Click Save.

Display De-duplication Statistics

To display the statistics for de-duplication in a cluster environment, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics**. The de-duplication statistics will be in the row labeled De-duplication in the GS Operations column.

Refer to De-duplication Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

GigaSMART Traffic Performance Enhancement

This section explains about :

- Flow Masking
- GigaSMART Traffic Performance Enhancement

Flow Masking

Flow Masking helps to identify flows, by selecting the portion of the packet for flow identification. The flow mask consists of an offset and a length, in bytes. Use the offset to specify the number of bytes from the beginning of the packet to the start of the mask within the packet. Use the length to specify the number of bytes, following the offset, to mask within the packet. The length identifies the traffic flow. Both the offset and the length are variable; however, the offset plus the length cannot be greater than 112 bytes. A default mask is provided with an offset of 14 bytes and a length of 28 bytes.

The GigaSMART encapsulated traffic performance is enhanced by using the feature Flow Masking.

For more information about encapsulated traffic performance, refer to GigaSMART Traffic Performance Enhancement.

GigaSMART Encapsulated Traffic Performance Enhancement

GigaSMART Traffic Performance Enhancement does not require a separate license.

The GigaSMART traffic performance enhancement provides a method to improve GigaSMART packet processing for MPLS traffic and other traffic having Layer 2 encapsulation, such as L2GRE or VNTag. This type of traffic has a header in the packet between the MAC address and the IP address. Figure 54MPLS Header Between MAC and IP Address in Packet shows the MPLS example.



Figure 54 *MPLS Header Between MAC and IP Address in Packet*

The GigaSMART processor is able to identify IP flows if there is only the MAC address and no other header in the packet before the IP address or if the only header before the IP address is VLAN. Without this enhancement, the GigaSMART processor cannot identify IP flows if there are MPLS headers or Layer 2 encapsulation other than VLAN before the IP address.

Most of theGigaSMART features require packets belonging to the same flow connection to be processed in the correct order. The processor uses the IP addresses and protocol to identify a traffic flow. When packets belonging to the same flow are identified, GigaSMART process these packets in a sequential manner. GigaSMART process all the encapsulated packets one at a time sequentially even though they belong to different traffic flows (different IP address pair).

Performance is impacted if the GigaSMART processor cannot use the IP source and destination (ipsrc, ipdst) to identify flows. This enhancement provides another method to identify flows. Using a flow mask, you select the portion of the packet for flow identification.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure the MPLS traffic performance enhancement, do the following:

1. From the device view, select **GigaSMART > GigaSMART Group > GigaSMART Groups**.
2. Click **New** to create a new GigaSMART group or **Edit** to modify an existing one.
3. Go to Flow Mask under GigaSMART Parameters and select **Enable**.
4. Enter the offset and length in the **Offset (bytes)** and **Length (bytes)** fields.

   If you do not enter any values in these fields, the default offset and length is used.
5. Click **Save**.

Refer to the following sections for examples:

- Flow Masking Example 1
- Flow Masking Example 2

Flow Masking Example 1

In Example 1 packets are expected to have two MPLS labels before the IP header, and no VLAN tag between the MAC and MPLS headers. IP addresses will be used to identify the flows.

The offset will be the sum of the following: 14 bytes for the Ethernet header+ 8 bytes for the MPLS headers +12 bytes offset from the beginning of the IP header = 34 bytes.

The length will be the sum of the following: 4 bytes for ipsrc + 4 bytes for ipdst = 8 bytes.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure Example 1, do the following:

1. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.
2. Click **New** to create a new GigaSMART Group or **Edit** to modify an exiting one.
3. Under GigaSMART Parameters, go to Flow Mask and select **Enable**.
4. Enter 34 in the **Offset (bytes)** field and 8 in the **Length (bytes)** field as shown in the following figure



5. Click **Save**.

Flow Masking Example 2

In Example 2, packets are expected to have one VLAN tag and two MPLS labels before the IP header. IP addresses will be used to identify the flows.

The offset will be the sum of the following: 14 bytes for the Ethernet header + 4 bytes for the VLAN tag + 8 bytes for the MPLS headers +12 bytes offset from the beginning of the IP header = 38 bytes.

The length will be the sum of the following: 4 bytes for ipsrc + 4 bytes for ipdst = 8 bytes.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure Example 2, do the following:

To configure Example 1, do the following:

1. From the device view, select **GigaSMART > GS Groups > GS Groups**.
2. Click **New** to create a new GigaSMART Group or **Edit** to modify an exiting one.
3. Click **Enable** under GS Params Flow Mask.
4. Enter 38 in the **Offset (bytes)** field and 8 in the **Length (bytes)** field as shown in the following figure.



5. Click Save.

GigaSMART Header Stripping

Required License: Header Stripping

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

GigaSMART operations with a **Strip Header** component can identify and remove headers from tagged packets or headers and trailers from tunneled (encapsulated) packets.

The following types of packets can be stripped:

- **Header Stripping** – Remove headers from ERSPAN, MPLS, MPLS+VLAN, VLAN, VN-Tag, VXLAN, GRE, tagged packets, Cisco FabricPath Headers, or FM6000 timestamps before they are sent to tool ports. This feature is handy when working with tools that either cannot recognize these headers or have to engage in additional processing to adjust for them. Figure 55Strip Header GigaSMART Operation Configured shows the GigaSMART Operations page for configuring Header Stripping.
- **Tunnel Stripping** – Remove both the header and trailer of ISL or GTP-encapsulated packets, preserving the packet within for analysis. This is handy when sending data to tools that cannot parse the ISL or GTP tunnel information and analyze the packets within. Figure 56Tunnel Stripping GigaSMART Operation Configured shows an example of the GS Operations page configured for tunnel stripping.

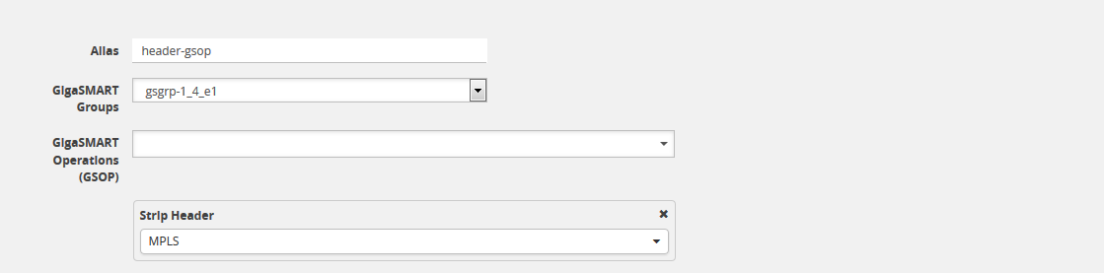


Figure 55 *Strip Header GigaSMART Operation Configured*

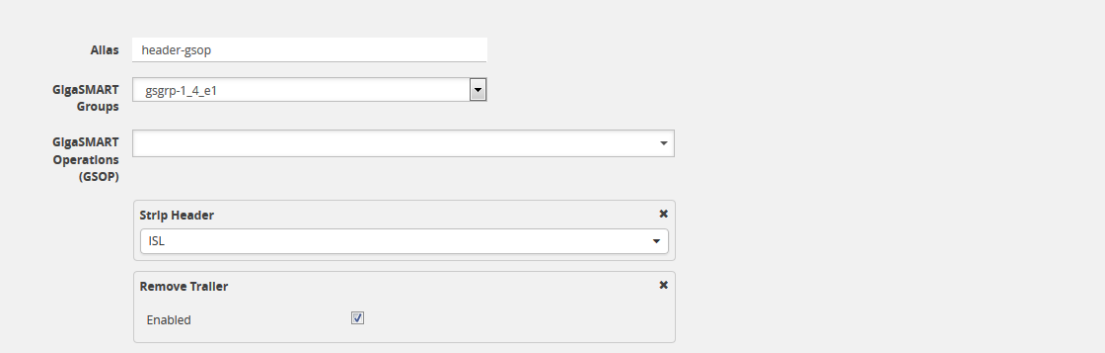


Figure 56 *Tunnel Stripping GigaSMART Operation Configured*

You can also use the **Strip Header** feature in tandem with the **Add VLAN** component to differentiate stripped packets from non-stripped packets. This is particularly useful when seeing stripped/non-stripped packets on common IP ranges (10.x.x.x; 192.168.x.x). Refer to the following table for more information.

On GigaVUE-TA400 both L3 MPLS outer Header Stripping and L2 MPLS outer Header Stripping of traffic carrying Pseudowire MPLS control word (PWMCW) or Pseudowire associated channel header (PWACH) are supported. L2 MPLS Header Stripping is not supported on traffic without PWMCW or PWACH. E

Keep in mind the following when configuring GigaSMART operations with a **Strip Header** component:

Summary	Description
ERSPAN Header Stripping	The ERSPAN header can be stripped. Specify an ERSPAN flow ID, from 0 to 1023. Use this option to strip ERSPAN Type II and Type III headers. A flow ID of zero is a wildcard value that matches all flow IDs.
Cisco FabricPath Header Stripping	The Cisco FabricPath headers can be stripped. The ability to decapsulate all packets with Cisco FabricPath headers; that is, all packets matching a destination switch ID and source switch ID. Also apply filters based on outer src/dst switch ID or ability to filter based on inner packet parameters with or without decapsulating the packet. The **Fabric Switch ID Source** and **Fabric Switch ID Destination** attributes are mandatory. Enter a value from 0 to 4095 (<0~(2^12-1)>) for a 12-bit switch ID. Enter 0 to strip all switch IDs.
FM6000 Timestamp Header Stripping	Packets entering GigaSMART from other devices may contain FM6000 timestamps. FM6000 is an Intel chip used for timestamping. The FM6000 timestamp can be stripped or it can be converted to UTC and appended to one of two Gigamon timestamping trailer formats. FM6000 has a hardware timestamp in the packet. For GigaSMART, the hardware timestamp needs to be translated into UTC time. An FM6000 device sends time mapping information in separate control packets called keyframes, which enable the UTC timestamp to be calculated. The calculated UTC timestamp can then be appended to the packets as a trailer. There are three timestamp formats: **None**, or **GigaSMART**, and **X12-TS** (for PRT-H00-X12TS). If the timestamp format is none, the FM6000 timestamp is stripped from the packet. If the timestamp format is **GigaSMART** or **X12-TS**, the FM6000 timestamp is stripped, converted to UTC, and a trailer containing the UTC timestamp is appended to the packets. The **GigaSMART** timestamp is added to the GigaSMART trailer. For the format of the GigaSMART trailer, refer to GigaSMART Trailer Reference. The **X12-TS** timestamp is added to the PRT-H00-X12TS trailer. For the format of the PRT-H00-X12TS trailer, refer to *GigaVUE-OS CLI Reference Guide*. **NOTES:** • FM6000 timestamp Header Stripping only supports an FM6000 timestamp that is present in the packet before the Frame Check Sequence (FCS). Packets containing an FM6000 timestamp that overwrites the existing FCS are discarded. Packets with a bad FCS are also discarded.

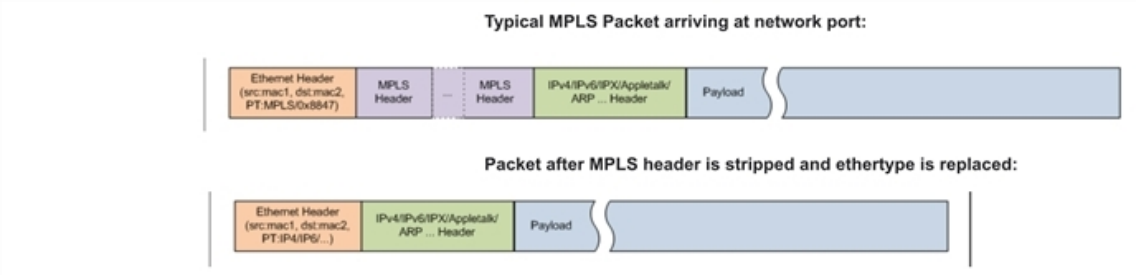
Summary	Description
	• Keyframes must only be sent to GigaSMART from one FM6000 device. If there are multiple FM6000 devices, their clocks must be synchronized using PTP. • The GigaVUE node maintains the keyframe database per GigaSMART operation (gsop). There is a one-to-one mapping of the keyframe database to each map associated with an FM6000 gsop. Ten (10) maps with an FM6000 gsop are supported per GigaSMART engine. • The keyframe rate on the FM6000 device is important. If GigaSMART does not see the keyframe, the time that is converted and appended to the header will be a default start time (1969). In certain conditions, the keyframe rate on the FM6000 device might need to be increased so that GigaSMART does not miss seeing this frame. • In H-Series Generation3 cards, FM6000 timestamping is not supported. For an FM6000 timestamping example, refer to Example – FM6000 Timestamping.
GRE Header Stripping	By specifying a GigaSMART Operation with a GSOP type of **Strip Header GRE**, the GigaSMART can strip GRE headers. It will automatically strip either Layer 3 or Layer 2 headers depending on the incoming packet. **Layer 3** – The GigaSMART can strip the outer IPv4/IPv6 delivery header and the GRE header to expose the encapsulated packet. Only IPv4 as the delivery protocol is supported in first and second generation cards. IPv4 and IPv6 address are supported in SMT-HC3-C08, SMT-HC1-S (third generation) cards. Any packet inside the GRE tunnel will be exposed, including IPv6 payloads. For an example, refer to Example – Stripping Layer 3 GRE IP Encapsulated Packets. **Layer 2** – The GigaSMART can strip GRE MPLS encapsulated and GRE Ethernet encapsulated packets, as follows: • GRE MPLS encapsulation – strip outer Ethernet header, outer IP header, GRE header, and MPLS header. • GRE Ethernet encapsulation (Transparent Ethernet Bridging) – strip outer Ethernet header, outer IP header, and GRE header. For an example, refer to Example – Stripping Layer 2 GRE Ethernet Encapsulated Packets. **Note**: The first-level **GRE Header Stripping** is supported on Gen 2 and Gen 3 devices.
Maximum MPLS Label Stack	The GigaSMART can strip MPLS headers up to a depth of seven labels.
Supported VLAN Types	The GigaSMART can strip both 802.1Q and Q-in-Q VLAN headers. Refer to How to Handle Q-in-Q Packets in Maps.
VXLAN Stripping	GigaSMART can strip VXLAN (Virtual eXtensible Local Area Network) headers. You can strip either matching VXLAN headers or all VXLAN headers. Select Strip Header from the GigaSMART Operation, select VXLAN for the protocol and use the following value in the Vxlan Id field: **0~(2^24-1)**. The VXLAN header is 8 bytes long with a 3-byte VXLAN Network Identifier (VNI) field. The VNI field is matched with the configured value and if it matches, the outer header (L2+IP+UDP+VXLAN) is stripped and the inner frame is sent to the tool. Specify a value of 0 to strip all outer VXLAN headers. **Note**: When processing packets with multiple encapsulation layers – for example, an ERSPAN-tunneled packet with a VXLAN tag – a VXLAN header-stripping

Summary	Description
	operation strips all the way to the end of the VXLAN layer instead of just the VXLAN tag.
ISL Tunnel Stripping	ISL tunnel stripping removes the 26-byte header and the 4-byte FCS trailer associated with Cisco ISL VLAN encapsulation. **IMPORTANT:** Make sure the packets processed by a GigaSMART operation with a **Strip Header ISL** component are all using ISL encapsulation. GigaSMART operations do not distinguish between packets using ISL and packets that do not – it strips the requisite bytes from all packets it processes.
GTP Tunnel Stripping	GTP tunnel stripping removes the header and trailer for GTP-u packets inside the GTP tunnel between the SGSN and GGSN interfaces in a 3G network, and between the eNodeB (eNb) and the SGW and between the SGW and the PGW in an LTE network. The SGSN and GGSN interfaces are also referred to as the Gn (or Gp) interface. The interface between eNb and SGW is referred to as S1U. The user plane interface between SGW and PGW is referred to as S5-U/S8-U. Both use GTPv1. Both GTPv1 and GTPv0 are supported for stripping. GTP-c control packets are not stripped. GTP[1] (also referred to as "GTP-Prime") is not supported for stripping.
Ethertype Replaced	After the VLAN/MPLS headers are stripped, the original ethertype carried in the Layer Two header is no longer valid. The GigaSMART replaces the ethertype field differently for MPLS and VLAN packets: **Ethertype Replacement for VLAN Packets** VLAN-tagged packets carry the original value for the ethertype field immediately after the VLAN tag. After the four-byte VLAN header is stripped, GigaSMART simply sets the ethertype field in the Layer 2 header to the value that was originally present in the packet past the VLAN tag. **Ethertype Replacement for MPLS Packets** Unlike VLAN-tagged packets, the Layer 3 protocol type is not carried in the packet for an MPLS packet – instead, it is applied by an egress router. To handle this, the GigaSMART examines the byte following the MPLS header to determine whether the packet is IPv4/IPv6 and takes the following actions: o **IPv4** – The GigaSMART replaces the ethertype from the MPLS packet with the IPv4 ethertype (0x0800) o **IPv6** – The GigaSMART replaces the ethertype from the MPLS packet with the IPv6 Ethertype (0x86DD). o **Non-IPv4/IPv6** – The GigaSMART passes the packet to destination tool ports without stripping the header. MPLS Header Stripping is only supported for IPv4/IPv6 packets.
CRCs Recalculated	The GigaVUE H Series node automatically recalculates and applies correct CRC checksums based on the new packet length after the header is stripped.
Viewing Statistics	From the device view, select **GigaSMART > GigaSMART Operations (GSOPS) > Statistics** to see statistics related to ongoing Header Stripping operations. Refer to View GigaSMART Statistics for more information.
Combine with Other	You can combine the **Strip Header** component with other GigaSMART components

Summary	Description
Components	in a single operation. Refer to How to Combine GigaSMART Operations for details on the combinations of GigaSMART operations. Refer to Order of GigaSMART Operations for information on the order in which components of a single GigaSMART operation are applied.
GigaSMART Engine Ports	Header stripping operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to Groups of GigaSMART Engine Ports for details.
Generic	Use Generic Header Stripping to remove any arbitrary header from a packet by specifying the offset and the length of the header. For information about Generic Header Stripping, refer to GigaSMART Header Stripping.

Header Stripping Illustrated

The following figures illustrate how the header-stripping operations work – you can see the original MPLS packet with its label stack intact, followed by a stripped packet with a recalculated CRC and a new ethertype field.



Example – Stripping MPLS Headers and Adding a VLAN ID

The example shown in Figure 57Strip Header GigaSMART Operation for Stripping MPLS Headers illustrates a simple GigaSMART operation named **HeaderStrip** configured to strip MPLS headers and add a VLAN tag of 200. The operation is assigned to the GigaSMART group with the alias of gsGrp1.

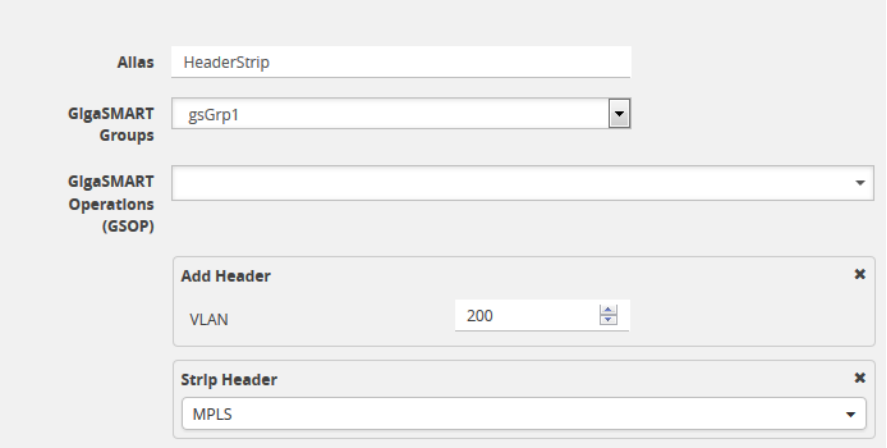


Figure 57 *Strip Header GigaSMART Operation for Stripping MPLS Headers*

The example shown in Figure 58Strip Header GigaSMART Operation for Stripping GTP Tunnel Informationillustrates a simple GigaSMART Header Stripping operation named **gtp-strip** configured to strip GTP tunnel information. The operation is performed by the GigaSMART group with the alias of GS1.

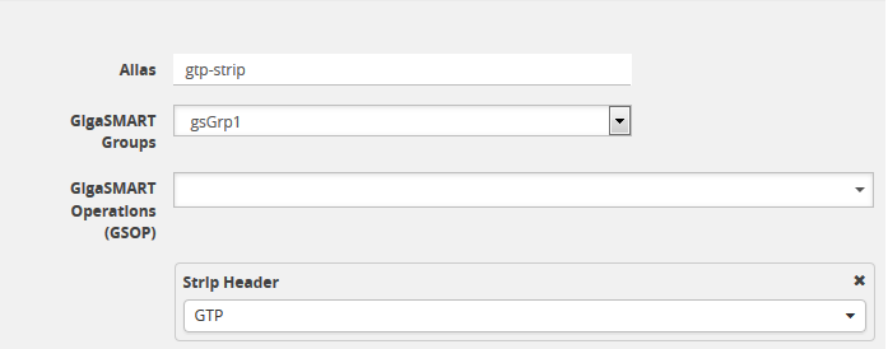
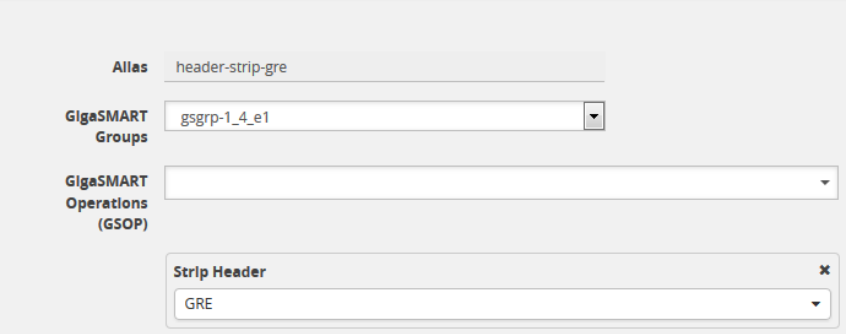


Figure 58 *Strip Header GigaSMART Operation for Stripping GTP Tunnel Information*

Example – Stripping Layer 3 GRE IP Encapsulated Packets

Use the configuration shown in the following figure to strip Layer 3 GRE IP encapsulated packets.

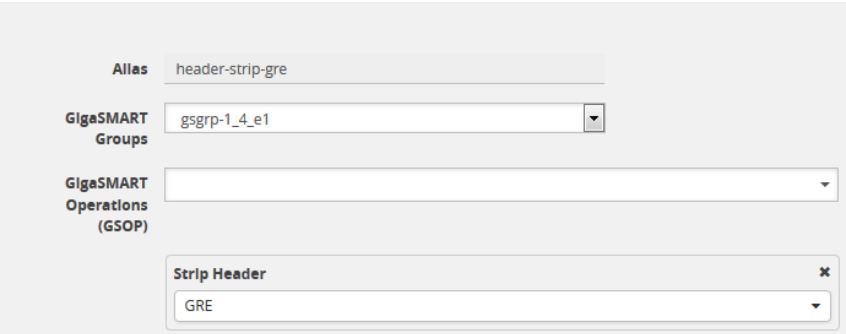


The following figure shows L3 GRE IP encapsulation before and after stripping.



Example – Stripping Layer 2 GRE Ethernet Encapsulated Packets

Use the configuration shown in the following figure to strip Layer 2 GRE Ethernet encapsulated (Transparent Ethernet Bridging) packets.



The following figure shows L2GRE Ethernet encapsulation before and after stripping.



Display Header Stripping Statistics

To display Header Stripping statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics**.

Refer to Header Stripping Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

Example – FM6000 Timestamping

Figure 59Strip Packets with FM6000 Timestamp shows an example GigaSMART Operation configured to strip packets containing the FM6000 timestamp:

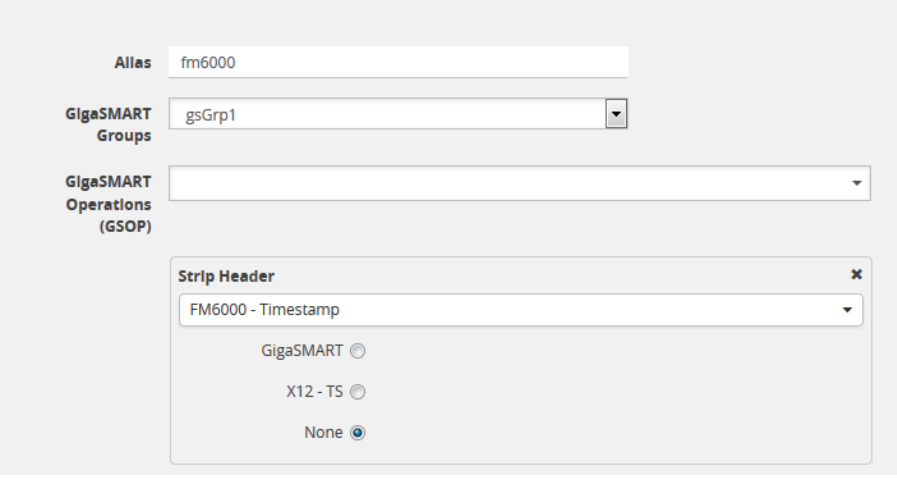


Figure 59 *Strip Packets with FM6000 Timestamp*

Figure 81GigaSMART Operation with a Static Offset is an example GigaSMART Operation configured to convert packets containing the FM6000 timestamp to UTC and append the UTC timestamp to the Gigamon trailer:

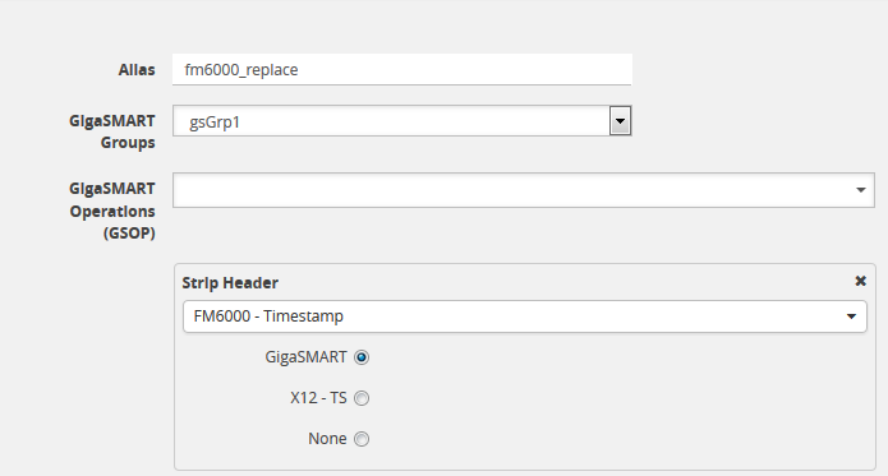


Figure 60 *Strip Packets and Append UTC Timestamp to the Gigamon Trailer*

Figure 61Map with Strip Header GigaSMART Operation is an example map using the strip header GigaSMART operation in Figure 60Strip Packets and Append UTC Timestamp to the Gigamon Trailer.

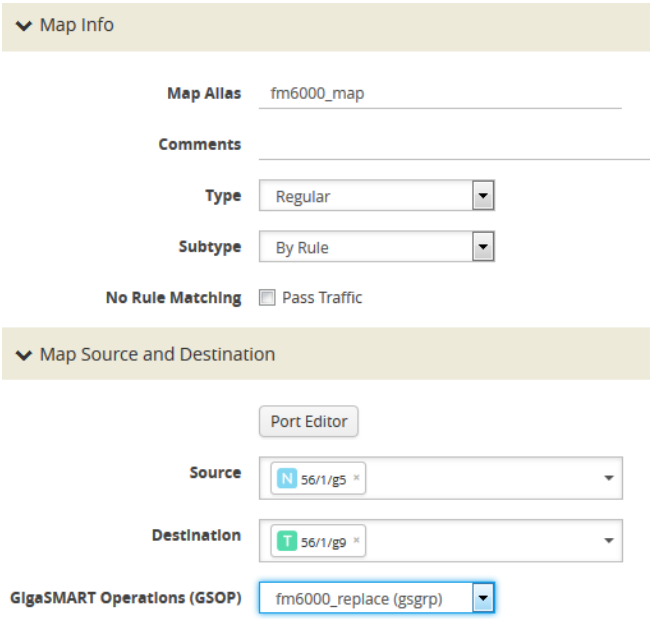


Figure 61 *Map with Strip Header GigaSMART Operation*

NOTE: There is one-to-one mapping between the GigaSMART Operation (gsop) and the map.

If there are multiple devices, each device can be configured with a different timestamp format. To configure this, use a different gsop and a different map for each device. For example, for packets arriving from FM6000 device1, configure a gsop for FM6000 device1 and associate it with map1. For packets arriving from FM6000 device2, configure a gsop for FM6000 device2 and associate it with map2.

All the maps can send all the packets to the same tool port.

Generic Header Stripping

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

GigaSMART operations use **Generic** Header Stripping to remove any arbitrary headers from a packet. The headers are stripped based on the offset and the length of the header.

To perform the generic Header Stripping operation:

1. Click **GigaSMART** on the left navigation pane. The GigaSMART Operation (GSOP) page is displayed.

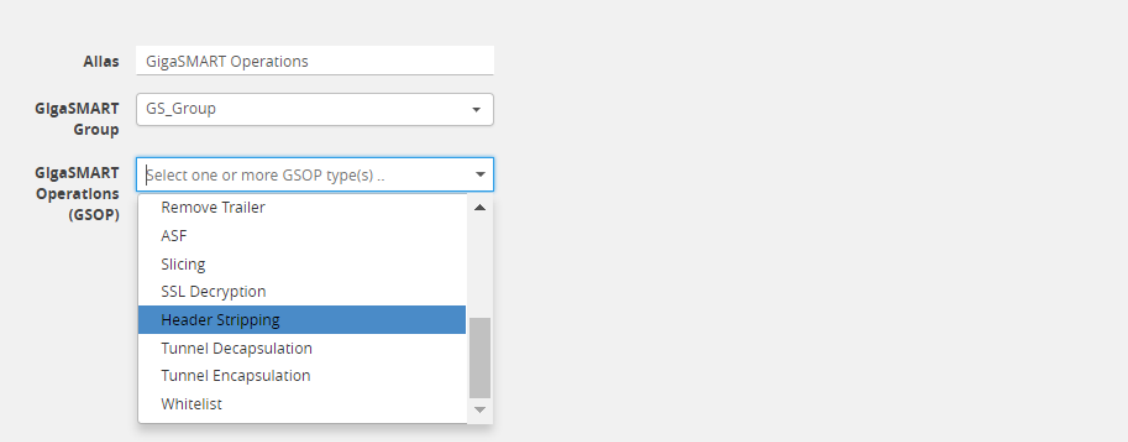


Figure 62 *GigaSMART Header Stripping*

2. In the **Alias** field, enter a name.
3. From the **GigaSMART group** drop-down list, select a GigaSMART group.
4. From the GigaSMART Operations (GSOP) drop-down list, select **Header Stripping**. A **Header Stripping** drop-down list is displayed.

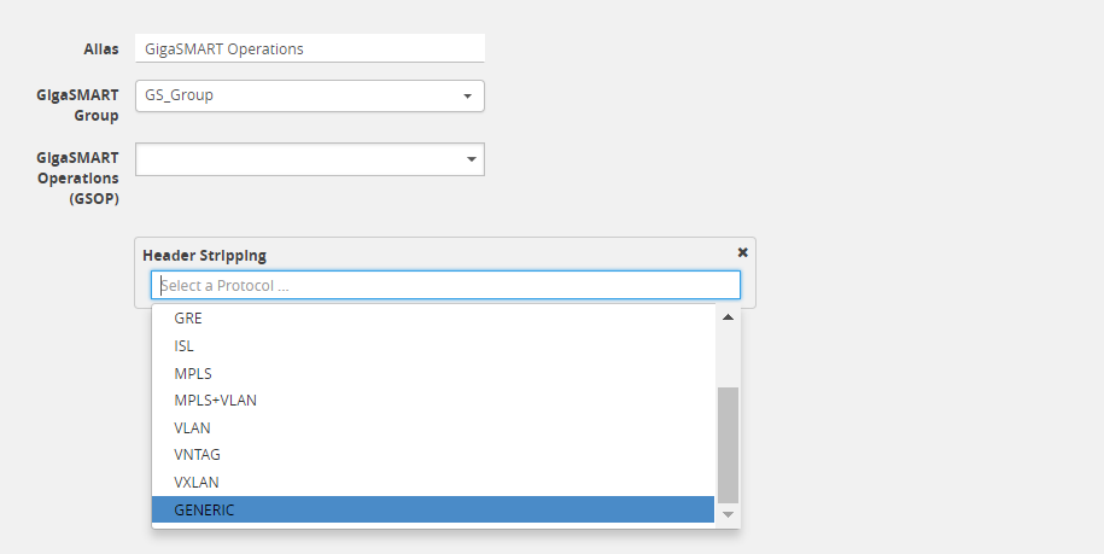


Figure 63 *GigaSMART Generic Header Stripping*

5. From the **Header Stripping** drop-down list, select **Generic**.

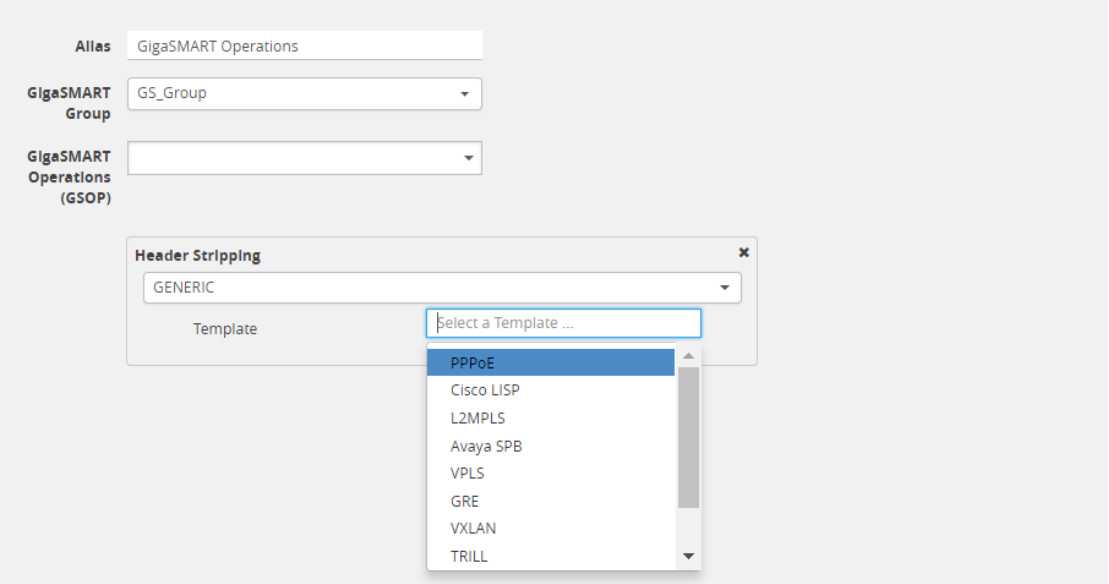


Figure 64 *Generic Header Stripping Template*

6. From the **Template** drop-down list, select one of the following options:

- Custom
- PPPoE
- Cisco LISP
- L2MPLS
- Avaya SPB
- VPLS
- GRE

- VXLAN
- TRILL
- Brocade VCS

Custom

A custom template lets you strip any arbitrary headers from a packet.

To strip any arbitrary header from a packet:

1. From the **Template** drop-down list, select **Custom.**

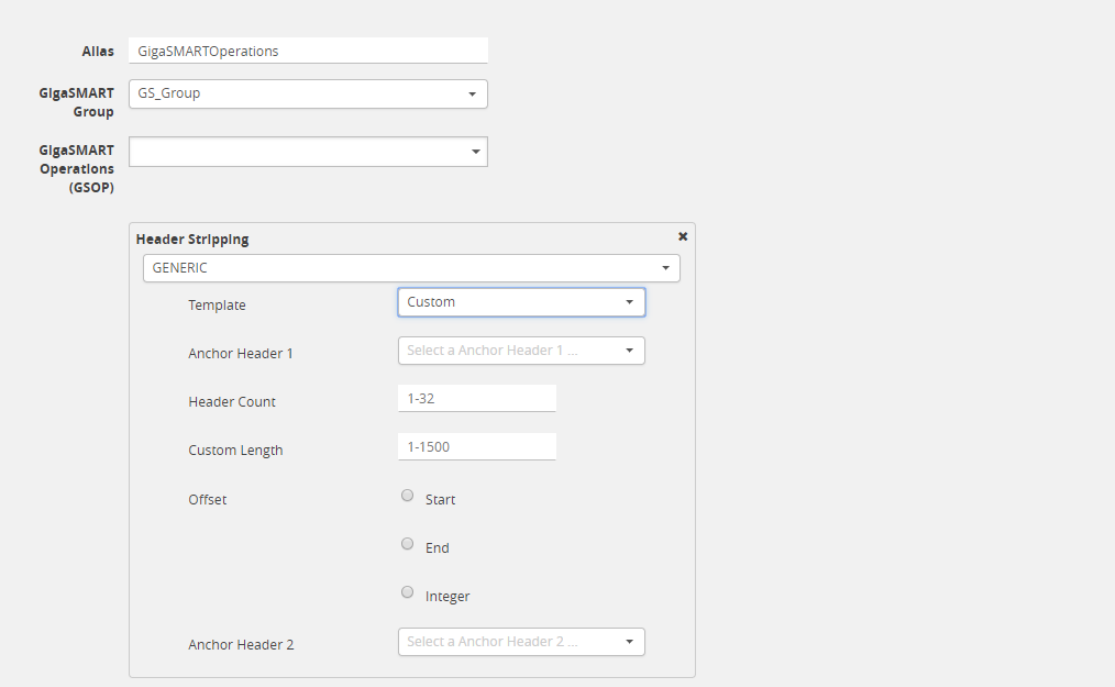


Figure 65 *Generic Header Stripping - Custom*

2. Select the following options to determine the headers to be stripped:

Component	Description
Anchor Header 1	Specifies the protocol from where GigaSMART should start stripping the header. The following values can be specified: o None o Eth o VLAN o MPLS

Component	Description
	o IPv4, and IPv6 The value **None** starts the Header Stripping operation from the start of the packet. If **None** is selected for Anchor Header 1, the Anchor Header 2 must also be set to **None.** The offset must not be set to **end**.
Offset	Specifies exactly from which end of Anchor Header 1 the stripping operation should start. You can specify the offset in terms of the following: **Start**—the Header Stripping operation starts from the left end of the Anchor Header 1. **End**—the Header Stripping operation starts from the right end of the Anchor Header 1. **Integer**—the Header Stripping operation starts from the specified integer offset of the Anchor Header 1. The integer value varies depending on the Anchor Header 1 specified.
Header Count	Specifies how many headers from the offset GigaSMART should remove. This is applicable when the packet headers are of known type. The known headers are as follows: • Ethernet • VLAN • MPLS • IPv4, and IPv6. It is important to note that, if **start** is specified for offset, the Anchor Header 1 is already counted for stripping. So, the value specified in the Header Count excludes the Anchor Header 1. If **None** is specified for Anchor Header 1, the Anchor Header 1 is not counted for stripping. So, the Header Count counts the Anchor Header 1 for stripping operation.
Custom Length	Specifies how many bytes of packet GigaSMART should remove. If the packet headers are unknown, the custom length of the unknown header can be specified to strip the packets. A combination of Header Count and Custom Length can also be used to strip the known and unknown headers. If Custom Length is specified, do not select **Any** for Anchor Header 2.
Anchor Header 2	Specifies the protocol that should become the next header after the stripping operation is complete. The following values can be specified: • None • Eth • VLAN • MPLS • IPv4 • IPv6 • TCP, UDP, and Any The value **Any** indicates that the next possible header can be any one of the options displayed for Anchor Header 2. The value **None** indicates that it is not necessary to mention the Anchor Header 2.

NOTE: Generic Header Stripping cannot strip unknown headers with variable length.

NOTE: On GigaVUE-TA400 both L3 MPLS outer Header Stripping and L2 MPLS outer Header Stripping of traffic carrying PWMCW (Pseudowire MPLS control word) or Pseudowire associated channel header (PWACH) are supported. L2 MPLS Header Stripping is not supported on traffic without PWMCW or PWACH.

PPPoE

The PPPoE template lets you strip the PPPoE encapsulated packets from the packet structure. Figure 66PPPoE Encapsulated Packets illustrates a red outline around the frame that needs to be striped.



Figure 66 *PPPoE Encapsulated Packets*

To strip the PPoE encapsulated packets:

1. From the **Template** drop-down list, select **PPPoE**. Refer to Figure 67Generic Header Stripping - PPPoE.

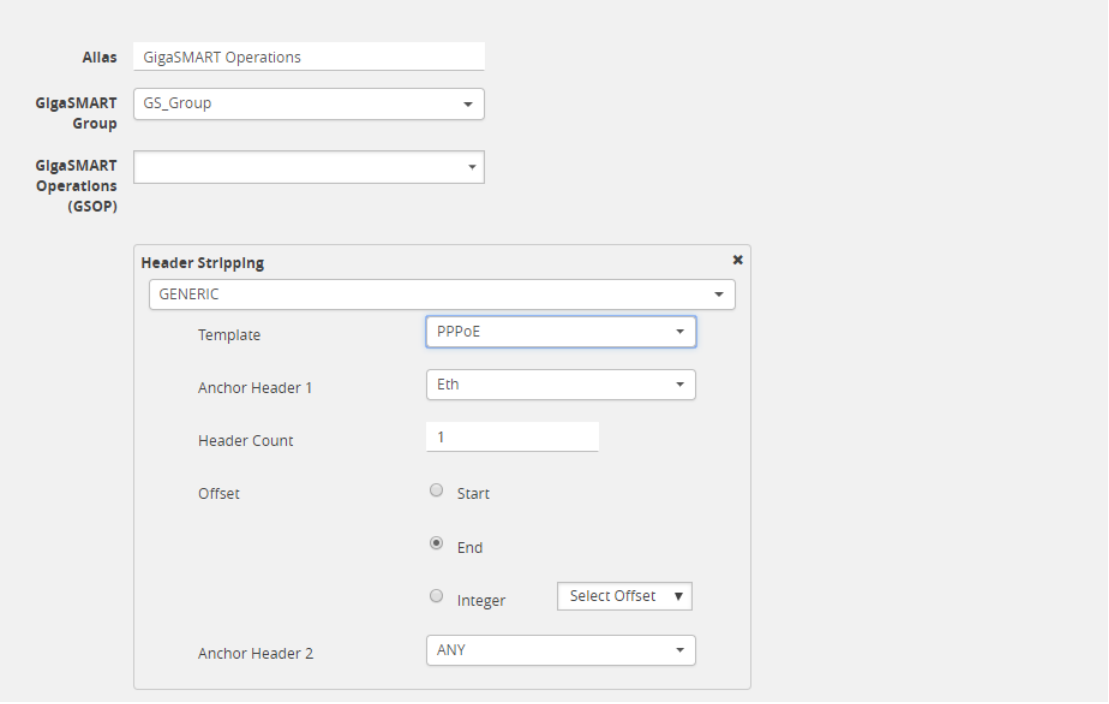


Figure 67 *Generic Header Stripping - PPPoE*

By default, the following values are selected:

Field	Value	Description
Anchor Header 1 Offset	Eth End	Starts the Header Stripping operation from the right end of the Ethernet header.
Header Count	1	Strips the header next to the Ethernet Header.
Anchor Header 2	Any	Updates a valid protocol as the Anchor Header 2 in the packet. In this case, any IPv4 or IPv6 protocol can become the Anchor Header 2.

2. Click **OK**. The Header Stripping operation is displayed in the GigaSMART Operations (GSOP) page.

Cisco LISP

Cisco LISP is used to carry original IP packets to support multi-homing. In this example, the IPv4 outer header, UDP header, and LISP header are stripped from the Cisco LISP header format. The LISP header is considered as an unknown header.

Figure 68Cisco LISP Encapsulated Packets illustrates a red outline around the frame that needs to be striped.



Figure 68 *Cisco LISP Encapsulated Packets*

To strip the Cisco LISP encapsulated packets:

1. From the **Template** drop-down list, select **Cisco LISP**.

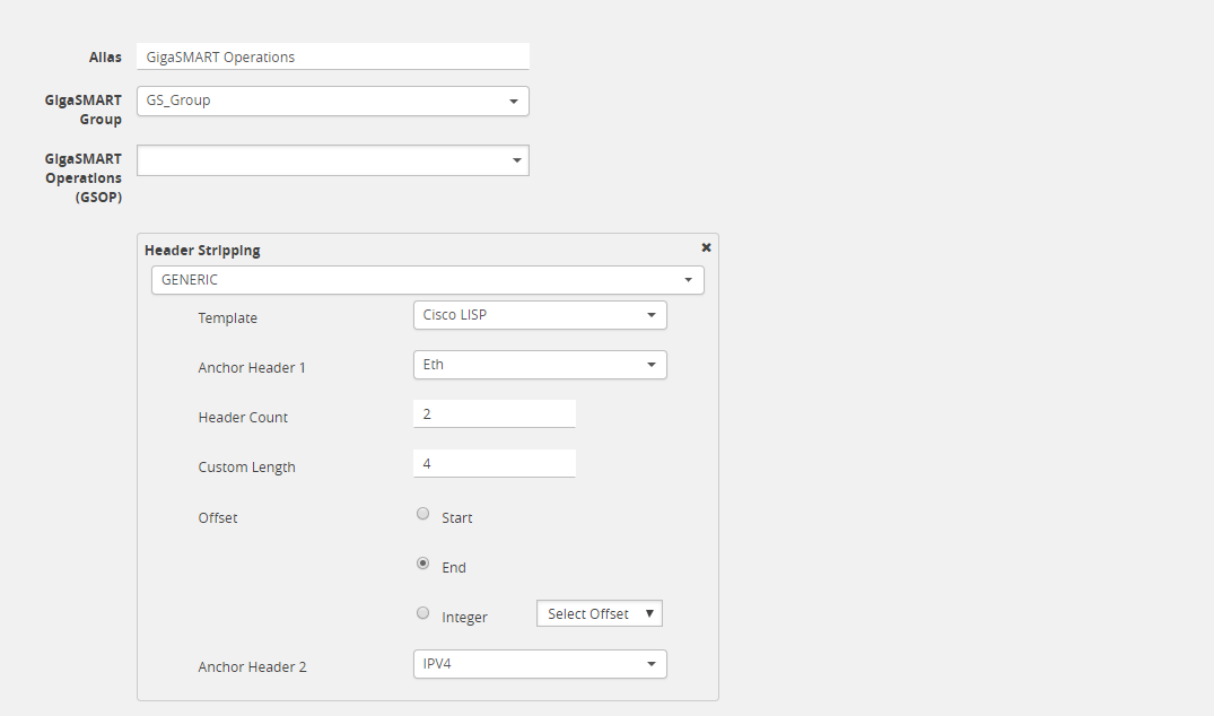


Figure 69 *Generic Header Stripping - Cisco LISP*

By default, the following values are selected:

Field	Value	Description
Anchor Header 1	Eth	Starts the Header Stripping operation from the right end of the Ethernet header.
Offset	End	
Header Count	2	Strips the next two headers— IPv4 Outer Header and UDP from the packet.
Custom Length	8	Strips 8 bytes of the unknown packet header. LISP is an unknown header.
Anchor Header 2	IPv4	Updates IPv4 protocol as the Anchor Header 2 in the packet.

2. Click **OK**. The Header Stripping operation is displayed in the GigaSMART Operations (GSOP) page.

L2 MPLS

The L2 MPLS packet, also known as VPLS, encapsulates Ethernet packets in the MPLS label stack. In this example, the outer Ethernet header and MPLS [PW Label] are stripped from the L2 MPLS encapsulated packets.

Figure 70L2 MPLS Encapsulated Packets illustrates a red outline around the frame that needs to be striped.



Figure 70 *L2 MPLS Encapsulated Packets*

To strip the outer MAC header from the L2 MPLS encapsulated packets:

1. From the **Template** drop-down list, select **L2MPLS**.

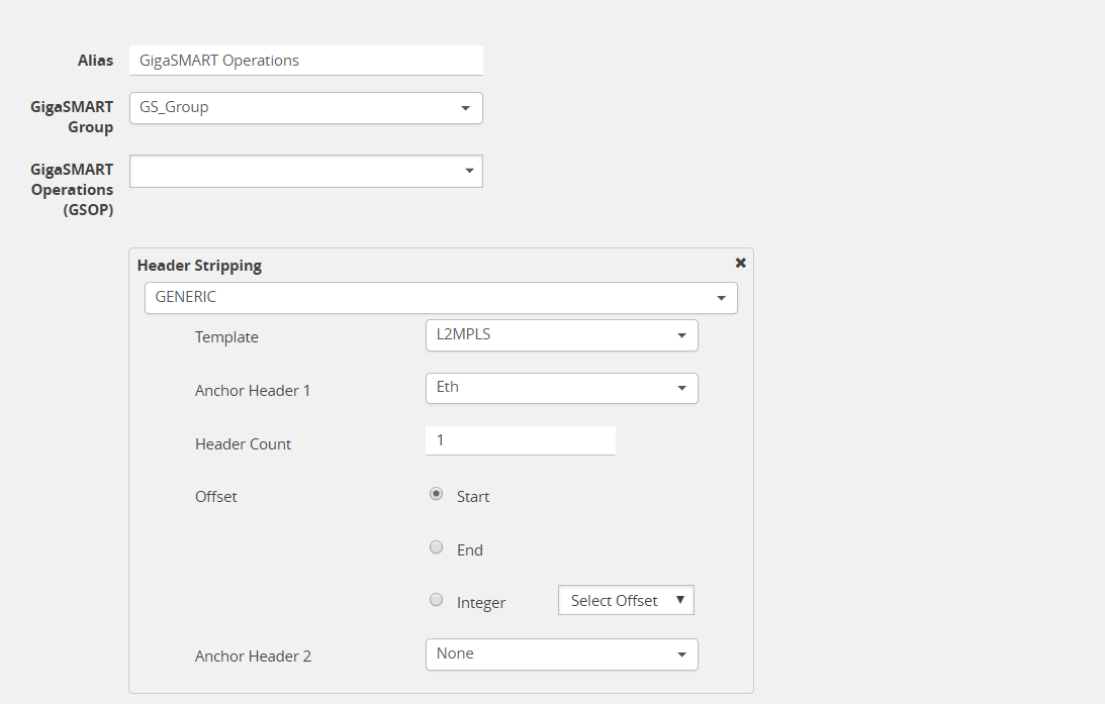


Figure 71 *Generic Header Stripping - L2 MPLS*

By default, the following values are selected:

Field	Value	Description
Anchor Header 1	None Start	Starts the Header Stripping operation from the start of the Ethernet header.

Field	Value	Description
Offset		
Header Count	2	Strips the first and the second header from the packet. The outer Ethernet header and MPLS [PW label] packet header are both removed. As Anchor Header 1 is set to none, the Header Count counts the first header for stripping.
Anchor Header 2	None	Signifies that there is no need to specify the Anchor Header 2. In this case, the IPv4 protocol forms the first header of the packet after the stripping operation is complete.

2. Click **OK**. The Header Stripping operation is displayed in the GigaSMART Operations (GSOP) page.

VXLAN

VXLAN encapsulates Ethernet packets in IP using VXLAN header. In this example, the outer Ethernet header, outer IP header, outer UDP header, and VXLAN Header are stripped from the VXLAN encapsulated packets.

Figure 72VXLAN Encapsulated Packets illustrates a red outline around the frame that needs to be striped.



Figure 72 *VXLAN Encapsulated Packets*

To strip the outer Ethernet frame from the VXLAN encapsulated packets:

1. From the **Template** drop-down list, select **VXLAN**.

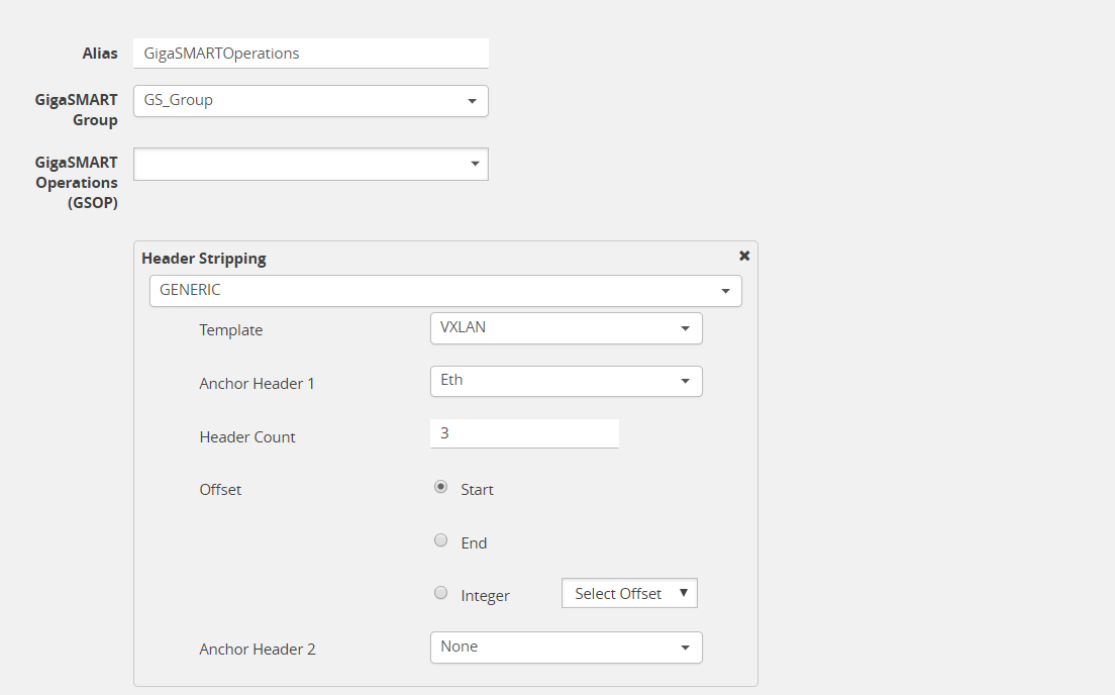


Figure 73 *Generic Header Stripping - VXLAN*

By default, the following values are selected:

Field	Value	Description
Anchor Header 1 **Offset**	Eth Start	Starts the Header Stripping operation from the start of the Ethernet header.
Header Count	3	Strips the next three headers—outer IP header, outer UDP header, and VXLAN header.
Anchor Header 2	None	Signifies that there is no need to specify the Anchor Header 2. In this case, the IPv4 protocol forms the first header of the packet. **Note**: When the Anchor Header 1 is set to **None**, the Anchor Header 2 must also be set to **None**.

2. Click **OK**. The Header Stripping operation is displayed in the GigaSMART Operations (GSOP) page.

TRILL

TRILL encapsulates Ethernet packets in Ethernet frame to provide L2 layer routing in data centers. In this example, consider TRILL as an unknown header. This TRILL frame is stripped with the inner Ethernet header from the encapsulated packets. The combined length of TRILL header (6 bytes) and inner Ethernet header (14 bytes) is 20 bytes.

Figure 74TRILL Encapsulated Packets illustrates a red outline around the frame that needs to be striped.



Figure 74 *TRILL Encapsulated Packets*

To strip TRILL from the encapsulated packets:

1. From the **Template** drop-down list, select **TRILL**.

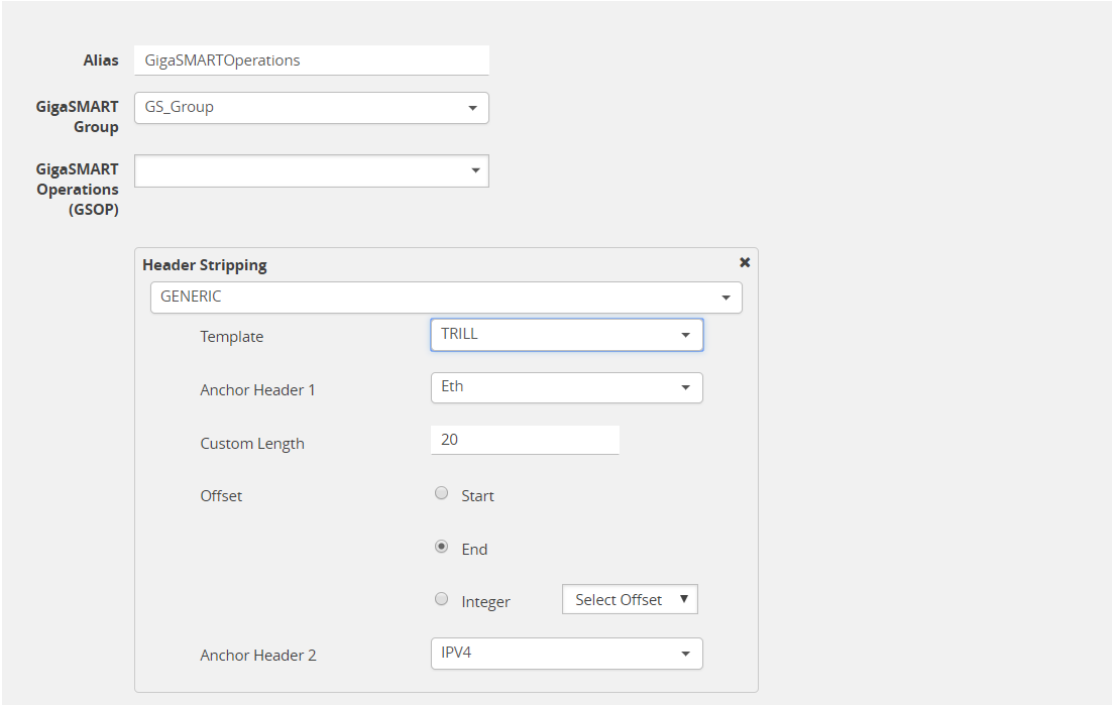


Figure 75 *Generic Header Stripping - TRILL*

By default, the following values are selected:

Field	Value	Description
Anchor Header	TRILL	Starts the Header Stripping operation from the right

Field	Value	Description
1 Offset	End	end of the outer Ethernet header.
Custom Length	20	Strips 20 bytes of unknown header from the packets. In this case, the TRILL and the inner Ethernet headers are stripped.
Anchor Header 2	IPv4	Updates IPv4 protocol as the second header in the packet.

2. Click **OK**. The Header Stripping operation is displayed in the GigaSMART Operations (GSOP) page.

Avaya SPB

Avaya SPB (802.1ah) fabric encapsulates Ethernet packets using MAC-In-MAC headers. In this example, the outer Ethernet header and ITAG are removed from the packet structure.

Figure 76Avaya SPB Encapsulated Packets illustrates a red outline around the frame that needs to be striped.

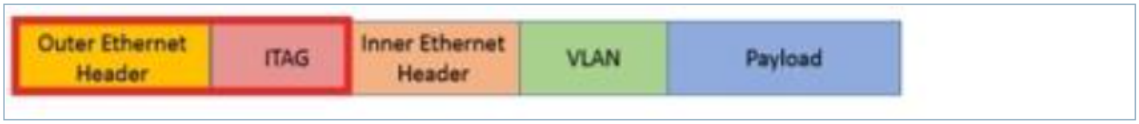


Figure 76 *Avaya SPB Encapsulated Packets*

To strip the outer Ethernet headers from the Avaya SPB encapsulated packets:

1. From the **Template** drop-down list, select **Avaya SPB**.

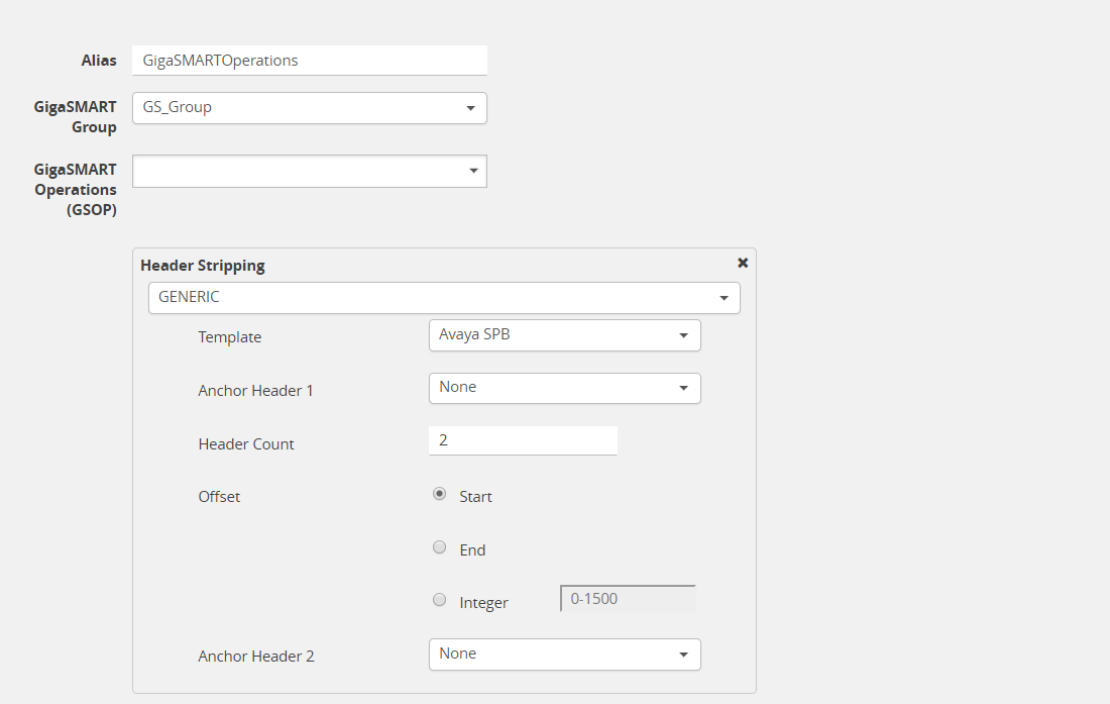


Figure 77 *Generic Header Stripping - Avaya SPB*

By default, the following values are selected:

Field	Value	Description
Anchor Header 1 Offset	None Start	Starts the Header Stripping operation from the left end of the outer Ethernet header.
Header Count	2	Strips the outer Ethernet and ITAG headers from the packet.
Anchor Header 2	None	Signifies that it is not necessary to specify the next header. The inner Ethernet header becomes the first header after the stripping operation is complete.

2. Click **OK**. The Header Stripping operation is displayed in the GigaSMART Operations (GSOP) page.

You can also strip the ITAG, inner Ethernet header, and VLAN from the packet structure.

Figure 78Avaya SPB Encapsulated Packets illustrates a red outline around the frame that needs to be striped.



Figure 78 *Avaya SPB Encapsulated Packets*

To strip the inner Ethernet headers from the encapsulated packets:

1. Select the following values to strip the inner Ethernet headers from the encapsulated packets:

Field	Value	Description
Anchor Header 1 Offset	Eth End	Starts the Header Stripping operation from the right end of the outer Ethernet header.
Header Count	3	Strips the ITAG, inner Ethernet, and VLAN headers from the packet.
Anchor Header 2	Any	Indicates that any valid protocol available after the Header Stripping operation can become the next header in the packet.

2. Click **OK**. The Header Stripping operation is displayed in the GigaSMART Operations (GSOP) page.

GigaSMART Header Addition

Required License: Header Stripping

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

GigaSMART operations with an **Add Header** selected can add VLAN tags to packets. This operation is useful in the following situations:

- Differentiating stripped packets from non-stripped packets on common IP ranges (for example,10.x.x.x; 192.168.x.x).
- Removing an arbitrary-length MPLS label stack and replacing it with a single, predictable, four-byte VLAN tag between the source address and ethertype field in the Layer 2 header. Many tools that are unable to parse the arbitrary length of an MPLS label stack can work with a predictable VLAN tag.

Keep in mind the following when configuring GigaSMART operations with an **Add Header** component:

Add VLAN Tag	You can combine **Strip Header** with **VLAN add** to help identify packets with stripped headers. This approach lets you remove an arbitrary-length MPLS label stack and replace it with a single, predictable, four-byte VLAN tag between the source address and ethertype field in the Layer 2 header. Many tools that are unable to parse the arbitrary length of an MPLS label stack can work with a predictable VLAN tag. **Packet Modifications for add_vlan** The **Add Header** operation makes the following modifications to a packet: o **TPID –** 0x8100 (802.1Q VLAN) or 0x88A8 and 0x9100 (Q-in-Q). The two-byte ethertype originally present in the Ethernet header is moved past the new VLAN header to identify the original Layer 3 header. o **CFI** – 0 o **Priority** – 0 o **VLAN ID** – User-provided value in the **VLAN** field of an **Add Header** GigaSMART Operation. Refer to How to Handle Q-in-Q Packets in Maps for TPID.
CRCs Recalculated	The GigaVUE H Series node automatically recalculates and applies correct CRC checksums based on the new packet length after the header is stripped.
Viewing Statistics	Use **GS Operations Statistics** page to see statistics related to ongoing GigaSMART operations. Refer to View GigaSMART Statistics for more information.
Combine with Other Components	You can combine the **Add Header** component with other GigaSMART components in a single operation. Refer to How to Combine GigaSMART Operations for details on the combinations of GigaSMART operations. Refer to Order of GigaSMART Operations for information on the order in which components of a single GigaSMART operation are applied.
GigaSMART Engine Ports	Header addition operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to Groups of GigaSMART Engine Ports for details.

GigaSMART Masking

Required License: Base

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus.

GigaSMART operations with **Masking** selected write over a specific portion of a packet with a specified one-byte pattern. Masking operations consist of an **offset**, **length**, and **pattern** as shown in Figure 79GigaSMART Operations Page with Masking Selected.

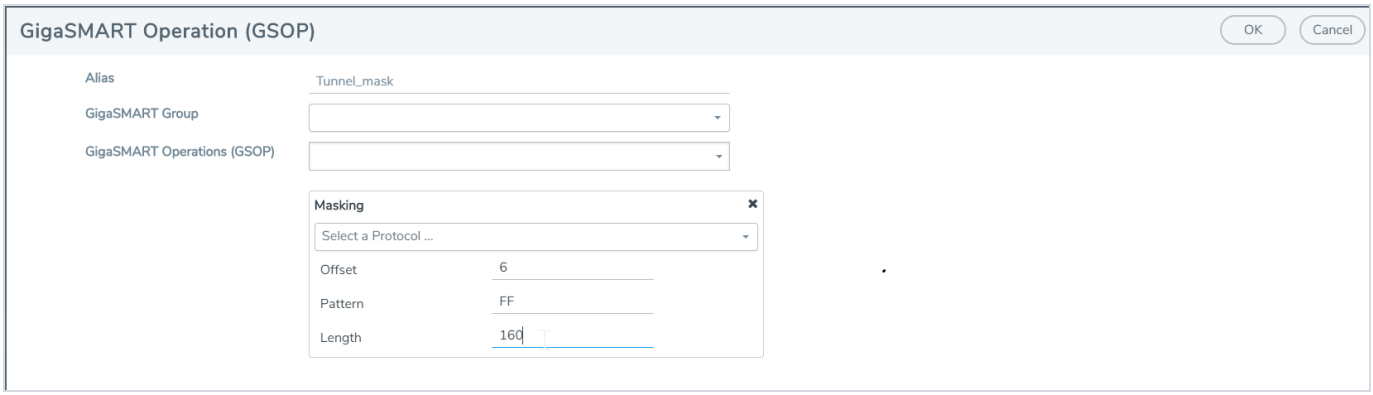


Figure 79 *GigaSMART Operations Page with Masking Selected*

The following table describes the fields.

Component	Description
Offset	Specifies **where** GigaSMART should start masking data with the supplied pattern. You can specify this in terms of either a static offset from the start of the packet or a relative offset from a particular protocol layer. This lets you automatically compensate for variable length headers, specifying a mask target in terms of a particular packet header.
Length	Specifies **how much** of the packet GigaSMART should mask. The specified one-byte pattern can be repeated to mask from 1-9600 bytes.
Pattern	Specifies **what** pattern GigaSMART should use to mask the specified portion of the packet. You can specify a one-byte hex pattern (for example, 0xFF).

Refer to the following when configuring GigaSMART operations with a **Masking** component:

Feature	Description
Protocol	The following are the protocols that you can select for from the protocol drop-down list: o **IPV4** – Mask starting a specified number of bytes after the IPv4 header. o **IPV6** – Mask starting a specified number of bytes after the IPv6 header. o **UDP** – Mask starting a specified number of bytes after the UDP header. o **TCP** – Mask starting a specified number of bytes after the TCP header. o **FTP**– Identify using TCP port 20. Mask payloads using offset from the TCP header. o **https** – Identify using TCP port 443. Mask payloads using offset from the TCP header. o **SSH** – Identify using TCP port 22. Mask payloads using offset from the TCP header. The GigaSMART-HC0 module can provide masking for GTP tunnels, provided the user payloads are unencrypted. Both GTPv1 and GTPv2 are supported – GTP' (GTP

Feature	Description
	prime) is not supported. Keep in mind that only GTP-u (user plane packets) are masked. Control plane packets (GTP-c) are left unmodified. o **GTP** – Mask starting a specified number of bytes after the outer GTP header. o **GTP-IPV4** – Mask starting a specified number of bytes after the IPv4 header inside the encapsulating GTP packet. o **GTP-UDP** – Mask starting a specified number of bytes after the UDP header inside the encapsulating GTP packet. o **GTP-TCP** – Mask starting a specified number of bytes after the TCP header inside the encapsulating GTP packet.
Masking Offset and Length	You can specify either a **relative** offset or a **static** offset for the masking pattern: • **Static** offsets begin masking a specific number of bytes from the start of the packet. Choose a static offset by setting **Protocol** to **None** and supplying an **Offset** from **<0~9000>** bytes. Zero (0) indicates the start of the Ethernet frame. • **Relative** offsets begin masking a specified number of bytes from the end of a particular header – IPv4, IPv6, and so on. Choose a relative offset by selecting any of the following values for the **protocol** argument and supplying an offset from the specified protocol header of **<1~9000>** bytes: **Note**: You can only mask one contiguous portion of a packet.
Recalculated CRC	GigaSMART automatically calculates a new Ethernet CRC based on the masked packet's length and data, and uses it to replace the existing CRC. This way, analysis tools do not report CRC errors for masked packets. **Note**: IP or UDP checksum is not recalculated if masking is done on the existing IP or UDP checksum.
GigaSMART Trailer	Masking operations can optionally include the GigaSMART Trailer. If you do elect to include the GigaSMART Trailer, it will include the original packet length. Refer to GigaSMART Trailer Reference for details.
In Combination with Slicing	Masking operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to Groups of GigaSMART Engine Ports for details.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

Examples – GigaSMART Masking

The example shown in Figure 80GigaSMART Masking Operation creates a GigaSMART masking operation named **Tunnel_mask**. This example starts masking six bytes after the end of the TCP layer in the GTP-encapsulated packet and continues for 150 bytes, writing over the existing data with an FF pattern.

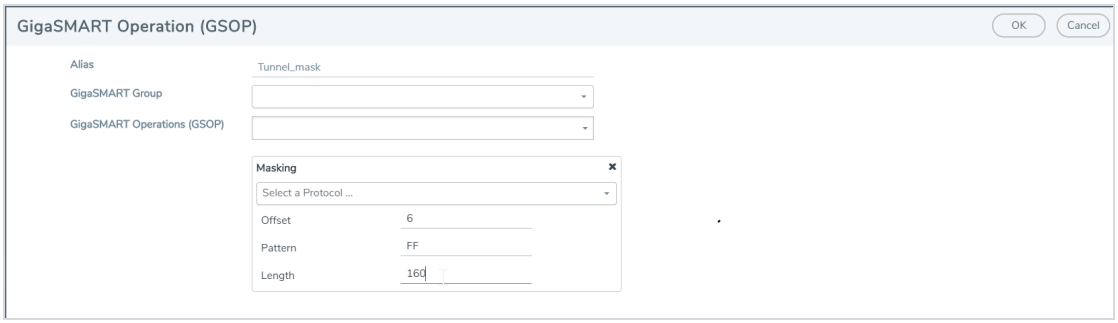


Figure 80 *GigaSMART Masking Operation*

This example shown in Figure 81GigaSMART Operation with a Static Offset creates a GigaSMART masking operation named **Mask_FIX**. This example uses a static masking offset of 148 bytes and continues for the next 81 bytes, writing over the existing data with an **FF** pattern. This GigaSMART operation is assigned to the GigaSMART group with the alias of gs2port2.

This example simulates how to mask a FIX (Financial Information eXchange) packet so that generic information is preserved at the start and end of the FIX data portion of the packet while private information within is masked. This example does not include the optional GigaSMART Trailer.

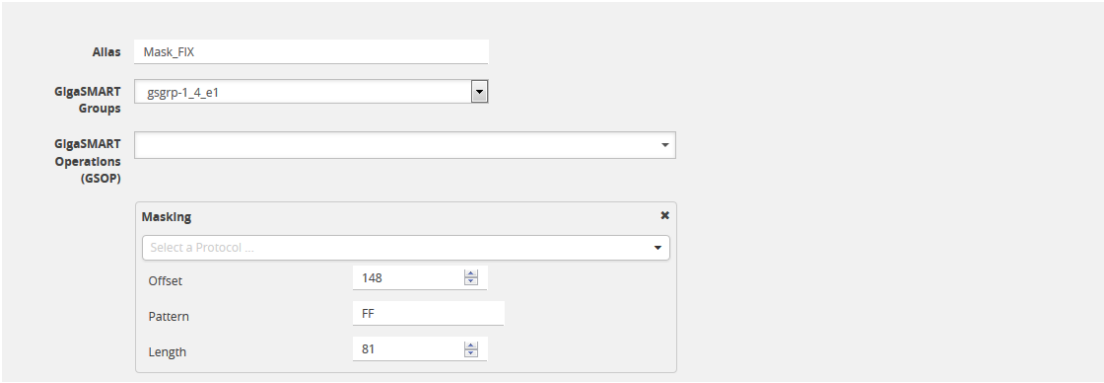


Figure 81 *GigaSMART Operation with a Static Offset*

Display Masking Statistics

To display masking statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics**. The statistics for masking will be in the row labeled Masking in the GS Operations column.

Refer to Masking Statistics Definitions for descriptions of the masking statistics as well as to GigaSMART Operations Statistics Definitions.

GigaSMART NetFlow Generation

Required License: NetFlow GenerationRequired License for NetFlow with Second Level Maps: Adaptive Packet Filtering (APF)

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC3 Gen 2.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

NetFlow Generation is a simple and effective way to increase visibility into traffic flows and usage patterns across systems. The flow-generated data can be used to build relationships and usage patterns between nodes on the network. Routers and switches that support NetFlow can collect IP traffic statistics to be exported as NetFlow records.

However, the processor and memory load of enabling NetFlow can cause service degradation and affect their ability to pass traffic without introducing latency and packet drops. Due to this processing overhead, sampled NetFlow is implemented in most of the high-end routers. Sampling in every “N” packets for NetFlow processing can severely limit the visibility needed to monitor flows.

The advanced capabilities of GigaSMART® technology can be leveraged to summarize and generate unsampled NetFlow statistics from incoming traffic streams. Offloading NetFlow Generation to an out-of-band solution like the Gigamon Deep Observability Pipeline completely eliminates the risk of using core production network resources in generating this data. Combined with the flexibility offered by Gigamon patented Flow Mapping® technology, operators can pick and choose from which flows to generate NetFlow statistics, while at the same time sending the original packets to other monitoring tools.

Support for NetFlow versions 5 and 9 and IP Information Export (IPFIX), as well as CEF, enables seamless integration with standards-based collectors. NetFlow records can also be exported to multiple collectors concurrently, providing a single flow source for business-critical management applications such as security, billing, and capacity planning. Exported flows can also be filtered so that collectors only receive the specific records relevant to them.

> **NOTE:** Legacy NetFlow supports only one NetFlow version (v5, v9 or IPFIX) record and NetFlow exporter format version per engine unless the exporter format is CEF.

Gigamon has also extended IPFIX to include URL information, providing insight into HTTP and SIP traffic. Other enterprise extensions for IPFIX are HTTP, DNS, and SSL certificates, which provide metadata that can be used for security analysis.

Additionally, Gigamon's Deep Observability Pipeline architecture is the first in the industry to summarize flow statistics as well as to provide the flexibility of aggregating, replicating, filtering, and forwarding raw traffic streams to monitoring tools for detailed troubleshooting and analytics.

The Gigamon Deep Observability Pipeline establishes a scalable framework to deliver pervasive flow-level visibility across enterprises, data centers, and service provider environments to accurately design, engineer, optimize, and manage their network infrastructure.

NOTE: NetFlow Generation exports records using IPv4. IPv6 is not supported.

GigaSMART operations with a NetFlow component can be assigned to multiple GigaSMART groups or GigaSMART groups consisting of multiple GigaSMART engine ports.

NetFlow/IPFIX Generation is a pillar of the GigaSECURE Security Delivery Platform.

NetFlow Generation is displayed in Figure 82NetFlow Generation Gigamon Solution.

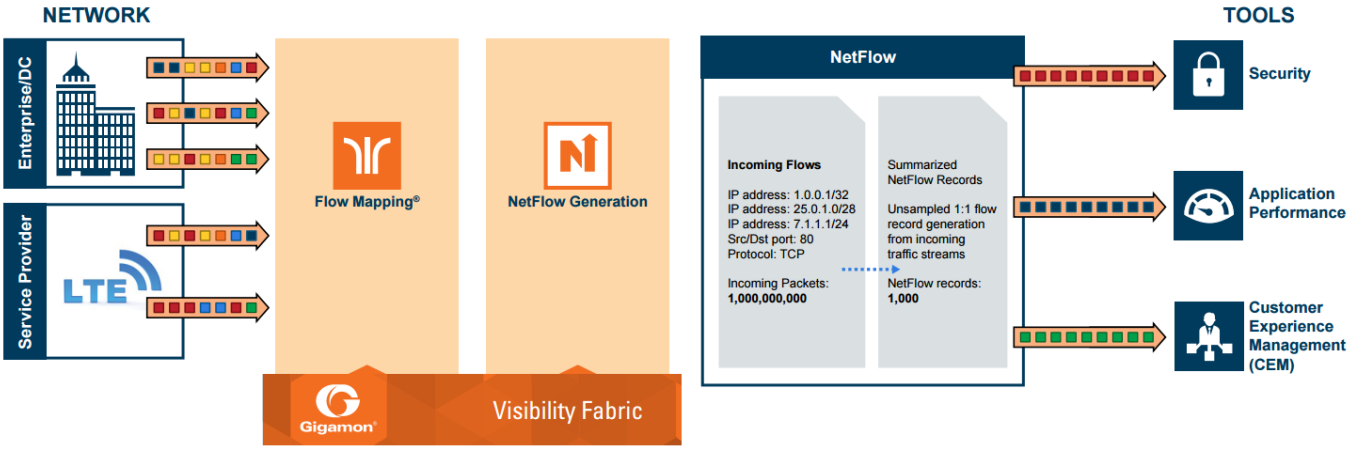


Figure 82 *NetFlow Generation Gigamon Solution*

In Figure 82NetFlow Generation Gigamon Solution, incoming packets from network(s) enter the Gigamon Deep Observability Pipeline and are directed by maps to NetFlow. NetFlow examines the incoming packets and converts the packets of choice into flows records. Specific flows are then forwarded to specific tools, such as Security, Application Performance, and Customer Experience Management (CEM) tools.

Active Timeout

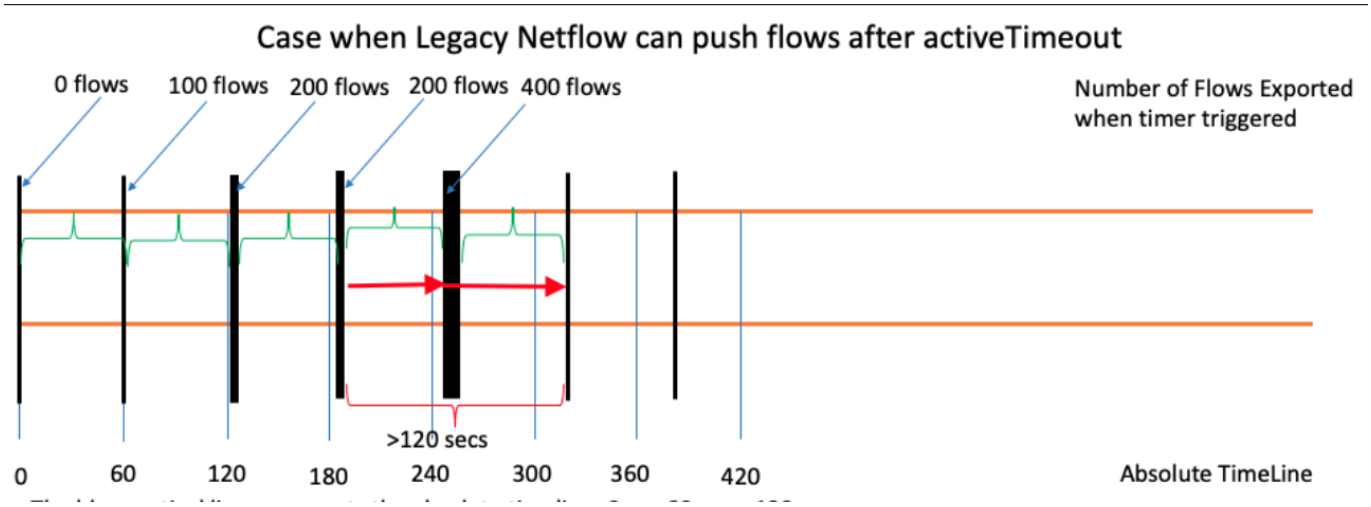
When a flow is active (GigaSMART engine receives packet) and sends packets for more than delta seconds, it hits an active timeout of delta seconds.

Inactive Timeout

When a flow stops sending packets for more than delta seconds, it hits inactive timeout of delta.

However, there can be some flows which are not exported at the end of active/inactive timeout.

The following diagram shows an example when Legacy NetFlow can push flows after active timeout.



In the above diagram, the blue vertical line represents the absolute time. The black vertical line represents the time at which GigaSMART engine starts exporting data after inactive time out. The green curly brackets represent the active timeout. The width of the black line represents the time taken by the exporter to push all the existing flows.

After a black line ends, it takes 60 seconds for another black line to start. This 60 second gap is represented by the green curly brackets. To ensure that each flow could be sent only once, there is only one vertical black line at any point of time.

In the above example, the red arrow represent a new flow in the network. When the flow starts and reaches the first black line, the GigaSMART engine calculates if the flow has reached the active time out. Here, since the flow has not reached the active Time out, the GigaSMART engine does not export the data for the flow.

The time consumed by the export process depends on the number of the flows that are being exported. The export process restarts after active timeout (the second black line). When the flow reaches the second black line, GigaSMART engine exports data for that flow.

NOTE: The difference between the time taken at which data is first exported for the flow and the flow start time is greater than the active timeout.

Hence, it is possible for the GigaSMART engine to rarely consume more time to export data of active flows.

NOTE: For inactive timeout also the Legacy NetFlow pushes the flows in same manner as explained in the example.

NetFlow Generation Components

NetFlow Generation collects IP traffic statistics on all interfaces where a NetFlow Monitor is enabled. It then gathers the statistics of the traffic flows and exports the NetFlow records to at least one NetFlow collector (typically a device that performs the actual traffic analysis based on the information from the NetFlow records).

Figure 83NetFlow Generation Components shows the NetFlow Generation components.

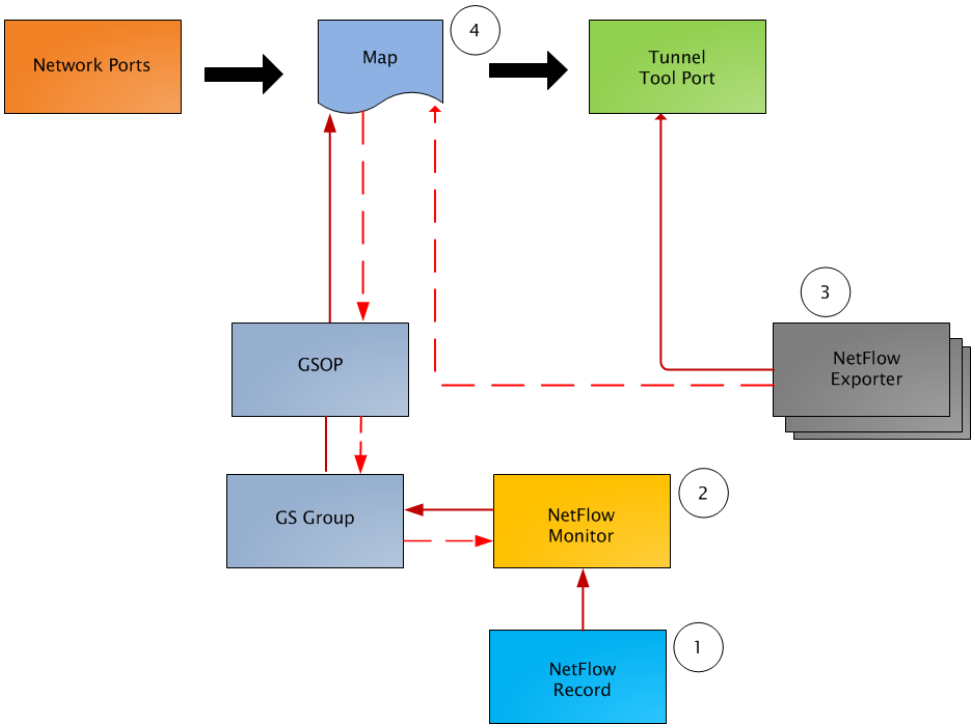


Figure 83 *NetFlow Generation Components*

GigaSMART NetFlow Generation illustrates the NetFlow Generation and how its components are associated. The NetFlow Generation associates its components in the following order:

1. One or more Records are associated to the Monitor.
2. The Monitor is associated to the GigaSMART group.
3. The Exporter is associated to the IP interface with tool port.

4. The map will eventually bind to the Exporter, Record, and Monitor.

NOTE: The dotted line from the map represents the interaction between the NetFlow Generation components.

Refer to *Example 1: NetFlow Generation Configuration* on page 631 for an example configuration of the following components.

Network Ports

NetFlow operates on the network flow. The incoming traffic on the network ports contains inputs such as, source and destination IP addresses, source and destination ports, interfaces, and so on. The network ports provide traffic to maps.

Map(s)

Traffic is received and acted upon according to maps. Maps determine what traffic is forwarded to NetFlow. Through map configuration, you add rules to filter the packets that need to go to NetFlow, and associate the map to the IP interface with tool port to specify where to send the filtered traffic.

Starting in software version 4.3.01, NetFlow supports both first level and second level maps. First level maps contain flow mapping rules to filter traffic that is needed by NetFlow and then send the filtered traffic to the IP interface with tool ports.

Second level maps are used for configuring filtering rules enabled through Adaptive Packet Filtering (APF). A virtual port is configured that directs traffic to the second level map. After the APF rules are applied, the filtered traffic that is needed by NetFlow is sent to the IP interface tool ports.

For examples of first level maps, refer to GigaSMART NetFlow Generation and GigaSMART NetFlow Generation.

For examples of first and second level maps, refer to GigaSMART NetFlow Generation and GigaSMART NetFlow Generation.

GigaSMART Group

The GigaSMART group specifies the GigaSMART engine to use, such as 8/1/e1 or 8/1/e2.

GSOP

The GigaSMART operation enables NetFlow. If a second level map is configured, the GigaSMART operation directs traffic to APF first, and then to NetFlow.

NetFlow Records

A NetFlow record contains key elements that specify what to match in the flow, such as all packets with the same source and destination port, or anything that comes in on a particular interface. A flow record also contains non-key elements that specify what information to collect for the flow, such as when the flow started or the number of bytes in the flow.

For NetFlow-v5, the fields in the flow record are fixed. For details, refer to *V5 Fixed Record Template* on page 657.

For NetFlow-v9 and IPFIX, you configure the fields, and thus create a record template. You specify how the fields are organized and in what order. The template is sent to the collector, so the collector knows what fields to expect in a NetFlow record. The template is sent periodically.

Starting in software version 4.6, multiple records are supported. An increased number of records allows more NetFlow data to be exported.

The maximum number of records is five. For all five records, each record must have the same match fields but differing collect fields. The same match fields will define the flows being considered. The different collect fields will define multiple templates sent to the NetFlow servers.

Starting in software version 5.1 for IPFIX and software version 5.2 for v5 and v9, a mix of IPv4 and IPv6 collect fields (IPv4 source/destination and IPv6 source/destination) are not supported in one record. Instead, create two records, one for IPv4 collect fields and one for IPv6 collect fields. When the IPv4/IPv6 collect fields are in separate records, an exporter will only send out records with non-blank elements.

NetFlow Monitors

Monitors store the NetFlow records associated with them. The configuration of a monitor includes the definition of the cache that specifies the records that you want to store, as well as timeouts associated with the cache. The cache can contain up to 4 million entries.

There can be a maximum of two monitors on a GigaSMART line card or module, one associated with each **e** port.

Starting in software version 4.6, up to five records can be added to the monitor. This results in the creation of five templates. For all five records, each record must have the same match fields but differing collect fields.

Sampled NetFlow Data

NetFlow data can be sampled. Sampling reduces the amount of ingress traffic sent to NetFlow for processing, which reduces the load on external collectors.

A NetFlow monitor can have multiple records with different sampling rates. The records are only updated with packets at the rate specified.

The following types of sampling are available: single-rate or multi-rate, as well as no sampling.

Sampling is enabled and disabled on the NetFlow monitor, across all flows. When sampling is enabled, you define the sampling rate by specifying a number for 1 in N, where N is the packet count.

For single-rate, the number can be from 10 to 16000. For multi-rate, the number can be from 1 to 16000. Single-rate applies to all records, whereas multi-rate applies to any record.

> **NOTE:** In a single-rate sampling type, all the NetFlow records are sampled in the same rate. In multi-rate sampling type, the sampling rate of the NetFlow records differ according to the settings defined in the individual records.

For example, if sampling is 1 in 1024, 1 packet in 1024 will be selected for NetFlow. The default is 1 in 1, which means no sampling.

> **NOTE:** The sampling mode in this release is deterministic. The selection of the packet is not random. Deterministic sampling means that if the rate is 1 in 1024, after 1023 packets, the 1024th packet is selected, while packets 1 to 1023 are ignored.

NetFlow Exporters

NetFlow records are sent to exporters. Each exporter is associated with one external collector. Records can be exported to both IPv4 and IPv6 destination. Either IPv4 or IPv6 destination address can be configured in an exporter. There can be up to six exporters that

send flow records to up to six external collectors. The six destinations are per GigaSMART engine.

The configuration of an exporter includes the IP address of a collector, the transport protocol and destination port, and the template refresh interval, which specifies the frequency of when the record template is sent to the collector.

Starting in software version 5.1, an option is added to assign different exporters to different records. Instead of records being sent to all exporters, you can add an exporter to a record, which defines the exporter to which the record is sent.

IP Interface with Tool Ports

NetFlow exporters are associated with IP interface, since exporters route both records and templates to collectors in the network.

NOTE: It is expected that the gateway specified in the IP interface configuration does Layer 3 routing. However, when the IP interface and the collector's IP address are in the same subnet, the following applies:

- Configure the IP interface's gateway IP address to the same as the collector's IP address.
- Configure the IP interface's subnet mask.
- The maximum number of exporters supported per GigaSMART group is six.

Enhancements to NetFlow

In addition to the NetFlow components, there are also the following enhancements:

- *Exporter Filtering*
- *Remote Interface IDs*
- *NetFlow Option Templates*
- *IPFIX Extension: HTTP Response Code*
- *IPFIX Extension: Packet URL*
- *IPFIX Extension: User Agent*
- *IPFIX Extension: Domain Name Service (DNS)*
- *IPFIX Extension: SSL Metadata*
- *SNMP Packet Support on IP Interfaces with Tool Ports*
- *NetFlow Format Support on Exporters*

Exporter Filtering

Not all collectors are interested in all kinds of packets. On each exporter, you can configure pass filters to filter the records transmitted to a collector. Thus, you can send a subset of records to a collector, such as the flow records for UDP packets or for packets coming in on a particular port.

Filtering is based on criteria, such as ports or IP addresses. For example, you can filter on different interfaces, such as single port (1/1/x1) or a contiguous range of ports (1/1/x1..x4). Note that you can only filter the criteria or a subset of the criteria that you configured for the match fields in the record.

NOTE: If no filters are configured, all records are sent to the collectors.

The exporter pass filters are as follows:

- Input interface
- IPv4 and IPv6 DiffServ Code Point (DSCP)
- IPv4 and IPv6 source address
- IPv4 and IPv6 destination address
- IPv4 protocol
- IPv4 Type of Service (TOS)
- IPv6 flow label
- L4 source and destination port
- MAC source and destination address
- VLAN ID

Take into account the following considerations:

- an exporter can have up to 5 filter rules
- each rule can have up to 4 attributes
- input interface can only be specified once per filter
- other attributes can be specified multiple time in a rule
- two rules cannot be identical

For an example of exporter filtering, refer to GigaSMART NetFlow Generation.

Remote Interface IDs

Interface ID, ingress as well as egress, can be configured as match and collect fields. Interface IDs can be local or remote. If you are interested in the interface ID on which a packet arrives, you need the port number of the node sending the packet. To get that information, you can use the LLDP/CDP discovery protocols that talk to neighbors to fetch the remote interface ID.

Discovery has to be either enabled or disabled on all the ports in a map. If discovery is enabled, the remote interface ID is sent in the NetFlow data record, as learned through LLDP/CDP.

To configure port discovery with NetFlow, enable discovery on the port or ports that are specified in the **Source** field of the associated map.

Note the following:

- You cannot modify discovery once the map is defined.
- Local port IDs are unique across a cluster. Remote IDs might not be unique. With port discovery enabled, there is a possibility of port ID collisions.

NOTE: If port discovery is not enabled, the local port ID is sent in the NetFlow data record.

When port discovery is enabled, the sending of the remote ID requires the collaboration of the end nodes. NetFlow expects an integer for port ID. If end nodes send an alphanumeric string, MAC address, or IP address (non-integers), that cannot be translated in an integer, NetFlow interprets them as either 0xFFFF or 0xFFFFFFFF.

When port discovery is enabled and ingress LLDP/CDP packets contain interface IDs that cannot be translated into an integer, use the collect field **interface input name** in the flow record definition. Using an interface name will send meaningful information about a network port to help identify the port to which the flow record refers.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To use the collect field interface input name, do the following:

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Records**.
2. Click **New** to create a new Record or select the recored and click Edit to change an existing record.

3. Click in the field **Non-Key Fields (Collect)** and select **Interface** from the list.
4. Select **Input** and then **Input Name**. Specify an input **Width** as shown in Figure 84Collect Field Interface input Name.
5. Click **Save**.

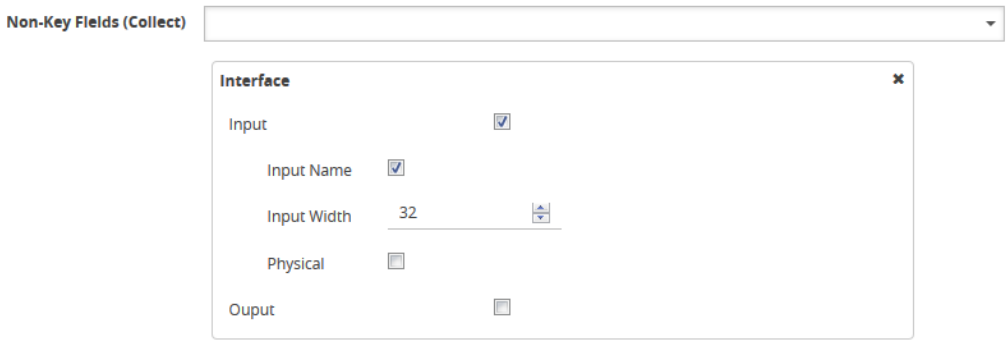


Figure 84 *Collect Field Interface input Name*

NetFlow Option Templates

For NetFlow-v9 and IPFIX, each exporter periodically sends option templates and option data records. There are two supported option templates, as follows:

- Interface ID to name mapping template and data record
- Exporter statistics template and data record

The option template for interface ID to name mapping contains an interface ID and name pair. Instead of a local port ID, the actual port number is available. For NetFlow-v9, the name field has a fixed length of 32 bytes. Names shorter than 32 bytes will be padded, while names longer than 32 bytes will be truncated. For IPFIX, the name field is of variable length.

NOTE: When port discovery is disabled for the port or ports specified in the **Source** field of the associated map, the interface option data record sends the interface ID to name mapping. But when discovery is enabled, interface option data records are not sent.

Each exporter sends out statistics, based on the standards. The exporter statistics option template includes information such as the exported flow record count, the exported message total count, and the exported octet total count.

By default, the transmission of option templates from the exporter is always enabled. The frequency of the transmission can be configured using the **Template Refresh interval** field in the NetFlow Exporter configuration page. To open the configuration page, select **GigaSMART > NetFlow / IPFIX Generation > Exporters** and click **New**.

IPFIX Extension: HTTP Response Code

For IPFIX only, use the collect field **Private PEN HTTP Response Code** in the flow record definition for capturing any packet with an HTTP response code embedded in it. This is a private information element extension, specific to Gigamon. The only valid private enterprise name (pen) is gigamon.

The HTTP response code information is captured from the packet and reported in the NetFlow record. The captured HTTP response code will be from the first packet that has HTTP/1 at the start of the HTTP header.

The field length in the flow record is a fixed length of 2 bytes.The range of response code values is from 100 to 599, as follows:

- 100-199 (informational)
- 200-299 (success related)
- 300-399 (redirection)
- 400-499 (client requests)
- 500-599 (server related)

If there is no HTTP response code in the flow, a zero value will be reported.

NOTE: In releases prior to software version 5.2, **HTTP Response Code** was directly under **Private PEN**. Starting in software version 5.2, there is a new **HTTP** section with **Response Code** under it. For backwards compatibility, both are supported.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure the NetFlow record for capturing any packet with an HTTP response code embedded in it, do the following:

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Records**.
2. Click **New**. Refer to Figure 85NetFlow Record Configuration for HTTP Response Code.

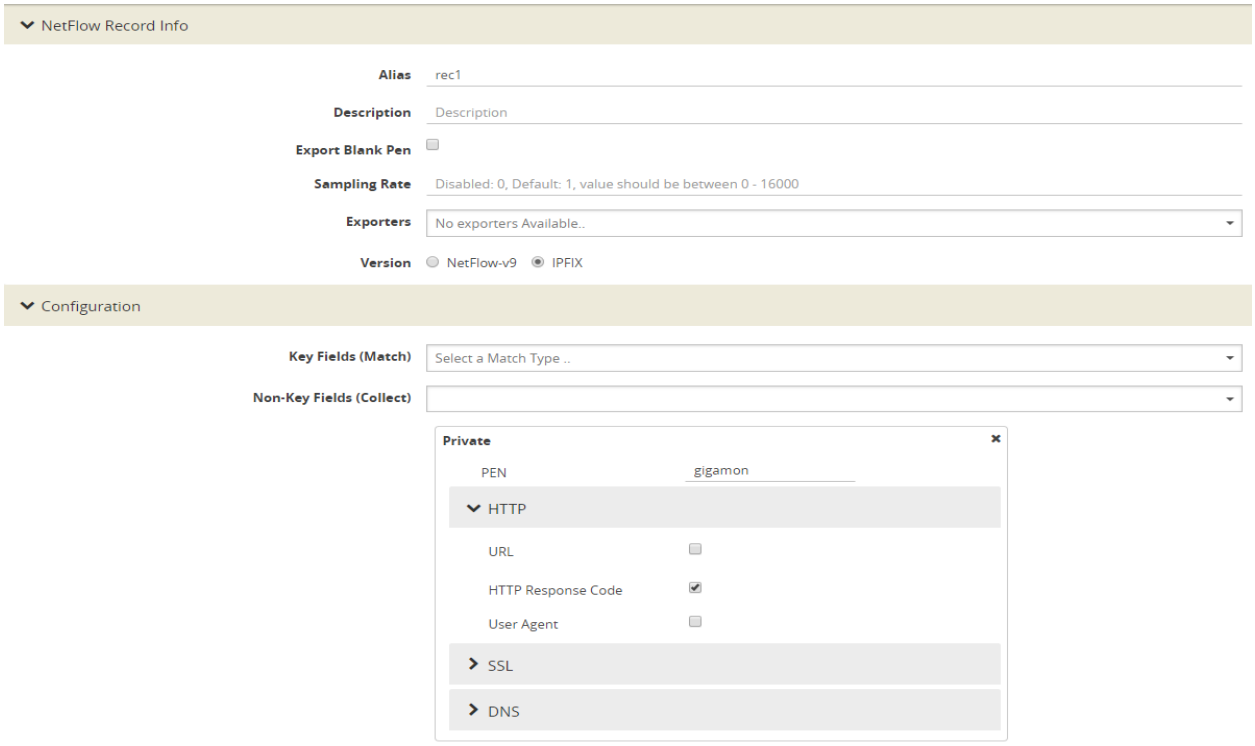


Figure 85 *NetFlow Record Configuration for HTTP Response Code*

3. For Version, select **IPFIX**.
4. Click the **Non-Key Fields (Collect)** drop-down list and select **Private**.
5. In the Private non-key field, do the following:
   - Set **PEN** to gigamon. (This is the default.)
   - Select **HTTP Response Code**.
6. Click **Save**.

IPFIX Extension: Packet URL

For IPFIX only, use the collect field **Private PEN HTTP URL** in the flow record definition for capturing any packet with a URL embedded in it. This is a private information element extension, specific to Gigamon. The only valid private enterprise name (pen) is gigamon.

The following URL information is captured from the packet and reported in the NetFlow record:

- HTTP: GET, POST, PUT, DELETE, and HEAD method types
- SIP: INVITE, ACK, BYE, REGISTER, OPTIONS, and CANCEL request types

The captured URL will be from the first packet that contains a URL. If there are additional URLs in subsequent packets in the flow, they will be ignored. If there is no URL in the flow, a zero length will be reported.

NOTE: The URL will always appear as the last element in a template, no matter the order in which the collect fields were configured.

NOTE: In releases prior to software version 5.2, **URL** was directly under **Private PEN**. Starting in software version 5.2, there is a new **HTTP** section with **URL** under it. For backwards compatibility, both are supported.

In Figure 85NetFlow Record Configuration for HTTP Response Code in the Private non-key field, select URL and enter an optional width.

IPFIX Extension: User Agent

For IPFIX only, use the collect field **Private PEN HTTP User Agent** in the flow record definition for capturing any packet with a user agent in the HTTP request header to gather information about clients user agents.

In general, the HTTP request is sent from the browser to the web application, so **User Agent** is filled in by the browser. As such, different browsers fill in this field with different values.

The maximum user agent length that is allowed in the data record is 250 bytes. The default is 150 bytes. Use the width parameter to specify a user agent length of up to 250 bytes.

In Figure 85NetFlow Record Configuration for HTTP Response Code in the Private non-key field, select User Agent and enter an optional width.

IPFIX Extension: Domain Name Service (DNS)

For IPFIX only, use the non-key or collect field **Private PEN DNS** in the flow record definition for capturing any packet with Domain Name Service (DNS) parameters embedded in it. This is a private information element, specific to Gigamon. The only valid private enterprise name (pen) is gigamon.

A domain name service translates host names into IP addresses. DNS has been exploited by attackers. Use this NetFlow enterprise element to gather metadata to help protect against security threats.

When certain DNS parameters are configured, the corresponding values for those collect parameters can be displayed in hexadecimal format or in text format. The following DNS parameters can display their values as text when the text version of that parameter is used:

- Additional Class, Additional Class Text
- Additional Type, Additional Type Text
- Authority Class, Authority Class Text
- Authority Type, Authority Type Text
- Query Class, Query Class Text
- Query Type, Query Type Text
- Response Class, Response Class Text
- Response IPv4 Address, Response IPv4 Address Text
- Response IPv6 Address, Response IPv6 Address Text
- Response Type. Response Type Text

For example, if the DNS **query-type** parameter collects a hexadecimal value of **0x1**, the **query-type-text** parameter collects the text string A, which refers to the IP address of the host.

The DNS parameters are captured from the packet and reported in the NetFlow record. Refer to Display Exporter Statistics and NetFlow Exporter Statistics Definitions.

Handle Blank Records for IPFIX

In the NetFlow record, the collect fields may contain one the following:

- Only private enterprise elements such as SSL, HTTP, or DNS
- Only non-private enterprise elements such as source IP address
- Both private and non-private elements

If all the collect fields contain only the private enterprise elements, and if during run-time, the records are blank or empty, they will not be added to NetFlow, however they will be counted in the exporter statistics as Empty Records Not Added.

If the collect fields contain both private and non-private enterprise elements, and if during run-time, the private enterprise elements are blank or empty, the records can be exported to the collector.

Configure DNS Record for IPFIX

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure the NetFlow record for capturing any packet with (DNS) parameters embedded in it, do the following:

1. Go to **GigaSMART > NetFlow / IPFIX Generation > Records**.
2. Click **New**. Refer to Figure 86NetFlow Record Configuration for DNS.

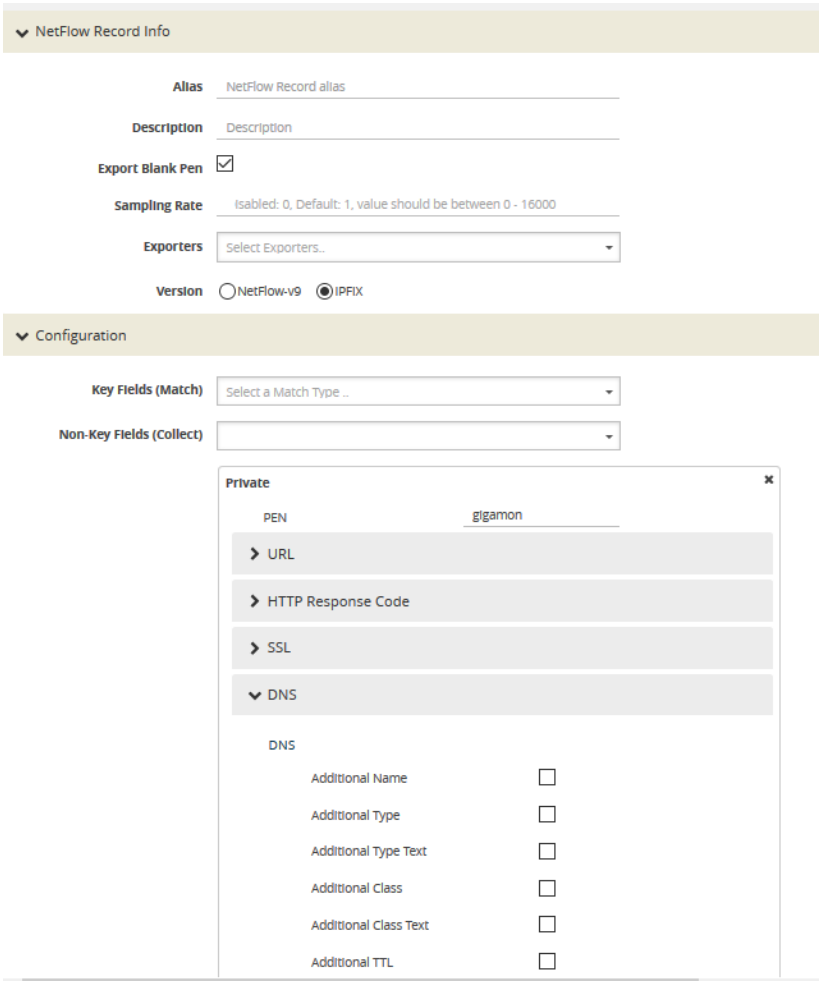


Figure 86 *NetFlow Record Configuration for DNS*

3. In the **Alias** field, enter a name.

4. To export the blank pen records, select **Export Blank Pen**.
5. For Version, select **IPFIX**.
6. Click the **Non-Key fields (Collect)** drop-down list and select **Private**.
7. In the **PEN** field, enter **gigamon**. (This is the default.)
8. Click **DNS** and select the parameters. The Number of Collects field is displayed for some DNS parameters. Refer to Table 8: DNS Parameters .
9. In the **Number of Collects** field, specify the number of instances of elements to collect from the DNS request. The value ranges from 1 to 10. The default value is 1.
10. Click **Save**.

Table 8: DNS Parameters

Method	For more information:
Additional Name	The domain name in the additional records section.
Additional Type	The additional type containing one of the RR type code.
Additional Type Text	The text string that corresponds to the hexadecimal value of the additional type containing one of the RR type code.
Additional Class	The additional class containing one of the RR class code.
Additional Class Text	The text string that corresponds to the hexadecimal value of the additional class containing one of the RR class code.
Additional TTL	The time-to-live (TTL), which is the time interval in seconds that the record is cached in the additional records section.
Additional RData	The content that describes the resource in the additional records section.
Additional RData Length	The length of the rdata field in the additional records section.
AN Count	The number of resource records in the answer section.
AR Count	The number of resource records in the additional records section.
Authority Name	The domain name in the authority section.

Method	For more information:
Authority Type	The authority type containing one of the RR type code.
Authority Type Text	The text string that corresponds to the hexadecimal value of the authority type containing one of the RR type code.
Authority Class	The authority class containing one of the RR class code.
Authority Class Text	The text string that corresponds to the hexadecimal value of the authority class containing one of the RR class code.
Authority TTL	The time-to-live (TTL), which is the time interval in seconds that the record is cached in the authority section.
Authority RData	The content that describes the resource in the authority section. The format of the rdata field varies according to the type and class of the resource record.
Authority RData Length	The length of the rdata field in the authority section.
Bits Count	The variable length of a bit map. The bit map must be a multiple of 8 bits long. For example: "/QR=1/AA=0/TC=0/RD=1/RA=1/AD=0/CD=0/Z=0", where /QR is the query (0) or a response (1), /AA is the authoritative answer, /TC is the truncation, /RD is the recursion desired, /RA is the recursion available, /AD is the authentic data, /CD is the checking disabled, and /Z is the reserved for future use.
Identifier	The identifier (Transaction ID) generated by the device that creates the DNS query and is copied by the server into the response so it can be used by that device to match that query to the corresponding reply received from the DNS server.
NS Count	The number of the name server (NS) resource records in the authority records section.
Op Code	The query type.
Qd Count	The number of entries in the question section.
Query Class	The query format containing one of the RR class codes.

Method	For more information:
Query Class Text	The text string that corresponds to the hexadecimal value of the query format containing one of the RR class codes.
Query Name	The domain name requested in the query. The maximum name length is 64 bytes. If the name is longer, it will be truncated.
Query Type	The query format containing one of the RR type codes.
Query Type Text	The text string that corresponds to the hexadecimal value of the query format containing one of the RR type codes.
Response Code	The type of the response.
Response Class	The response format containing one of the RR class codes.
Response Class Text	The text string that corresponds to the hexadecimal value of the response format containing one of the RR class codes.
Response Name	The domain name in the response. The maximum name length is 64 bytes. If the name is longer, it will be truncated.
Response Type	The query type specified in the response.
Response Type Text	The text string that corresponds to the hexadecimal value of the query type specified in the response.
Response RData Length	The length of the rdata field in the response data field.
Response RData	The content that describes the resource in the response data field. The format of the rdata field varies according to the type and class of the resource record.
Response-TTL	The time-to-live (TTL), which is the time interval in seconds that the record is cached.
Response iPv4 Address	The IPv4 address in the response if the response type host and class are Internet/IPv4.
Response iPv4 Address Text	The text string that corresponds to the hexadecimal value of the IPv4 address in the response if the response type host and class are Internet/IPv4. The format is dotted decimal.

Method	For more information:
Response IPv6 Address	The IPv6 address in the response if the response type host and class are Internet/IPv6.
Response iPv6 Address Text	The text string that corresponds to the hexadecimal value of the IPv6 address in the response if the response type host and class are Internet/IPv6. The format is dotted decimal.

IPFIX Extension: SSL Metadata

For IPFIX only, use the collect field **Private PEN ssl** in the flow record definition for capturing any packet with Secure Sockets Layer (SSL) or server metadata embedded in it, such as common name. This is a private information element extension, specific to Gigamon. The only valid private enterprise name (pen) is gigamon.

Examining the parameters associated with the SSL certificate or the SSL server provides visibility into SSL flows in the network and helps detect malicious activity. For example, checking the issuer might reveal an unknown self-signed certificate or a certificate signed by a questionable Certificate Authority (CA). Checking the certificate validity dates might reveal an expired certificate.

When NetFlow collects SSL certificate metadata, it makes use of the GigaSMART SSL application, described in GigaSMART Passive TLS/SSL Decryption. The data is routed to the SSL application first and then to NetFlow. If de-duplication is also enabled, the data is routed from de-duplication to SSL, and then to NetFlow. The SSL application does not decrypt the data.

NOTE: Only the NetFlow Generation license is needed for NetFlow to collect SSL certificate metadata.

When certain SSL certificate parameters are configured, the corresponding values for those collect parameters can be displayed in hexadecimal format or in text format. The following SSL certificate parameters can display their values as text when the text version of that parameter is used:

- Serial Number, Serial Number Text
- Signature Algorithm, Signature Algorithm Text
- Subject Algorithm, Subject Algorithm Text
- Valid Not After, Valid Not After Text
- Valid Not Before, Valid Not Before Text

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure a record for SSL Certificate or SSL Server, do the following:

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Records**.
2. Click **New**.
3. For Version, select **IPFIX**.
4. Click in the **Non-Key fields (Collect)** field.
5. Select **Private** from the drop-down list.
6. In the Private non-key field, do the following:

- Set **PEN** to gigamon.
- Under SSL, select the SSL Certificate and SSL Server parameters and specify a width value in bytes. For example, **Certificate Issuer Common Name** and **Certificate Subject Common Name** as shown in Figure 87NetFlow Record Configuration for SSL.

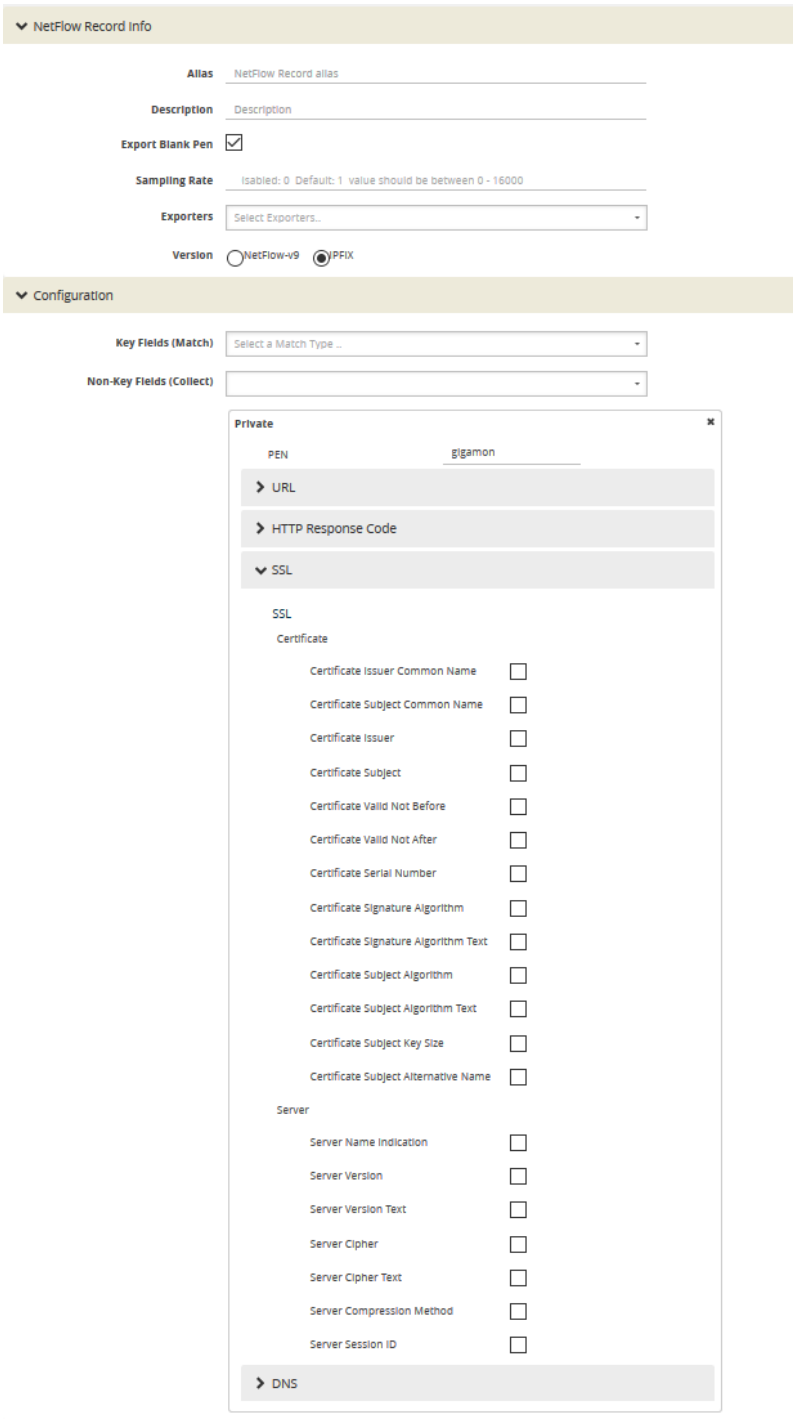


Figure 87 *NetFlow Record Configuration for SSL*

SSL Certificate Parameters

When certain SSL server parameters are configured, the corresponding values for those collect parameters can be displayed in hexadecimal format or in text format. The following SSL server parameters can display their values as text when the text version of that parameter is used:

- Cipher, Cipher Text
- Version, Version Text

For example, if the **ssl server cipher** parameter collects a hexadecimal value of **C027**, the **ssl server cipher-text** parameter collects the following text string:

TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256

The parameters supported for the SSL certificate are as follows:

- **Certificate Issuer**—the certificate issuer, which identifies the entity that has signed and issued the certificate. For example: "/C=US/ST=Arizona/L=Scottsdale/O=MyCo2.com, Inc./OU=http://certs.myco2.com/repository//CN=MyCo2 Secure Certificate Authority", where /C is the country name, /ST is the state or province, /L is the locality name, /O is the organization name, /OU is the organizational unit name, and /CN is the common name.
- **Certificate Issuer Common Name**—the certificate issuer common name, which is a subset of **Issuer**.
- **Certificate Subject**—the certificate subject, which identifies the entity associated with the public key stored in the subject public key. The **Certificate Subject** has the same fields as the **CertificateIssuer**.
- **Certificate Subject Common Name**—the certificate subject common name, which is a subset of **Subject**.
- **Certificate Subject Alternative Name**—the subject alternative name, which allows identities to be bound to the subject of the certificate. This parameter is useful to detect if the certificate claims to sign for anything else and to detect anomalies such as certificates that claim to sign for a wildcard (*). The first subject alternative name present in the certificate is collected.
- **Certificate Valid Not Before** and **Certificate Valid Not After**—the date on which the certificate validity period begins and ends. The certificate validity period is the time interval during which the CA warrants that it will maintain information about the status of the certificate. It is expressed in universal time. The format is YYMMDDHHMMSSZ, where Z is Zulu time (GMT).
- **Certificate Valid Not Before Text** and **Certificate Valid Not After Text**—the text string that corresponds to the hexadecimal value of the date on which the certificate validity period begins and ends. The format is MMM DD HH:MM:SS YYYY GMT.

- **Certificate Serial Number**—the unique number for each certificate issued by a given CA. The issuer name and serial number identify a unique certificate. This parameter is useful to detect any certificate changes or substitutions.
- **Certificate Serial Number Text**—the text string that corresponds to the hexadecimal value of the unique number for each certificate issued by a given CA.
- **Certificate Signature Algorithm**—the identifier for the cryptographic algorithm used by the CA to sign the certificate, defined in ASN.1 format. This parameter is useful to detect servers that are not compliant with an organization's cryptographic standards.
- **Certificate Signature Algorithm Text**—the text string that corresponds to the hexadecimal value of the identifier for the cryptographic algorithm used by the CA to sign the certificate, defined in ASN.1 format. This parameter is useful to detect servers that are not compliant with an organization's cryptographic standards.
- **Certificate Subject Algorithm**—the subject public key algorithm used, defined in ASN.1 format, such as RSA or DSA.
- **Certificate Subject Algorithm Text**—The text string that corresponds to the hexadecimal value of the subject public key algorithm used, defined in ASN.1 format, such as RSA or DSA.
- **Certificate Subject Key Size**—the subject public key size.

Optionally, on the **issuer**, **Certificate Issuer Common Name**, **Certificate Subject**, **Certificate Subject Common Name**, and **Certificate Subject Alternative Name** parameters, you can indicate the width of the field in bytes.

SSL Server Parameters

The parameters supported for the SSL server are as follows:

- **Server Name Indication**—the extension to the Transport Layer Security (TLS) protocol by which a client indicates the host name to which it is attempting to connect at the start of the handshaking process.
- **Server Version**—the version of SSL, including the major and minor version.
- **Server Version Text**—the text string that corresponds to the hexadecimal value of the identifier for the version of SSL, including the major and minor version.
- **Server Cipher**—the cipher that the server agreed to use for that session.
- **Server Cipher Text**—the text string that corresponds to the hexadecimal value of the identifier for the cipher that the server agreed to use for that session.
- **Server Compression Method**—the server compression method, which is typically not set (in other words, NULL). This parameter is useful to detect attacks that use compression.
- **Server Session ID**—the session identifier, generated by a server, which identifies a particular session. This parameter is useful to detect a session restart.

Optionally, on the **Server Name Indication** parameter, you can indicated the width of the field in bytes.

Restrict Ports for NetFlow SSL Sessions

SSL metadata is collected by sending all traffic to the SSL module. The SSL module accepts all IPv4 TCP packets and attempts to find SSL sessions. During the process of finding these sessions, the metadata required by NetFlow is extracted.

To improve the throughput of SSL metadata extraction for NetFlow, the TCP ports can be restricted. Reducing the TCP packets inspected by limiting the TCP ports inspected reduces the amount of packets sent to the SSL module.

Configure the monitor to scan specific ports for SSL. Options are available to scan all ports, a list of up to 10 ports, or well-known ports.

The following are the well-known ports:

- MAP_SSL_PORT 993
- POP3_SSL_PORT 995
- SMTP_SSL_PORT 465
- LDAP_SSL_PORT 636
- NNTP_SSL_PORT 563
- HTTP_SSL_PORT 443

Configure SSL Certificate and Server Parameters

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure the SSL certificate and server parameters to collect, do the following in the UI, for example:

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Records**.
2. Click **New**.
3. Enter an alias in the **Alias** field to identify this record. For example, ipfixrec.
4. Select **IPFIX**.
5. From the **Non-Key Field (Collect)** list, select **Private**.

6. Set **PEN** to gigamon. (This is the default.)
7. Under SSL, select any of the following:
    - Certificate issuer Common Name Width 30
    - Certificate Subject Common Name Width 40
    - Certificate issuer Width 150
    - Certificate Subject Width 120
    - Certificate Valid Not Before
    - Certificate Valid Not After
    - Certificate Serial Number
    - Certificate Signature Algorithm
    - Certificate Signature Algorithm Text
    - Certificate Subject Algorithm
    - Certificate Subject Algorithm Text
    - Certificate Subject Key Size
    - Certificate Subject Alternative Name
    - Server Name Indication Width 64
    - Server Version
    - Server Version Text
    - Server Cipher
    - Server Cipher Text
    - Server Compression Method
    - Server Session ID
8. Click **Save**.

Best Practices for Collecting SSL Metadata

When collecting SSL certificate metadata, the match conditions must be configured so that the NetFlow sessions match the SSL sessions. To do this, configure the following in the NetFlow record for IPv4 flows:

1. From the device view, select **GigaSMART >**NetFlow Record / IPFIX Generation.
2. Select the record ipfixrec and click Edit. (This is the record configured in the previous section Configure SSL Certificate and Server Parameters.)
3. From the **Key Fields (Match)** list, select **IPv4**.
4. Select **Protocol**.
5. Under **Source** select **Address**.
6. Under **Destination** select **Address**.

7. From the **Key Fields (Match)** list, select **Transport** and then select the following:
   - **Source Port**
   - **Destination Port**
8. Click **OK**.

Or, configure the following in the NetFlow record for IPv6 flows:

1. From the device view, select **GigaSMART >**NetFlow Record / IPFIX Generation.
2. Select the record ipfixrec and click Edit. (This is the record configured in the previous section Configure SSL Certificate and Server Parameters.)
3. From the **Key Fields (Match)** list, select **IPv6**.
4. Select **Protocol**.
5. Under **Source** select **Address**.
6. Under **Destination** select **Address**.
7. From the **Key Fields (Match)** list, select **Transport** and then select the following:
   - **Source Port**
   - **Destination Port**
8. Click **OK**.

Or, configure the following in the NetFlow record for a mix of IPv4 and IPv6 flows:

1. From the device view, select **GigaSMART >**NetFlow Record / IPFIX Generation.
2. Select the record ipfixrec and click Edit. (This is the record configured in the previous section Configure SSL Certificate and Server Parameters.)
3. From the **Key Fields (Match)** list, select **IPv4**.
4. Select **Protocol**.
5. Under **Source** select **Address**.
6. Under **Destination** select **Address**.
7. From the **Key Fields (Match)** list, select **IPv6**.
8. Select **Protocol**.
9. Under **Source** select **Address**.
10. Under **Destination** select **Address**.
11. From the **Key Fields (Match)** list, select **Transport** and then select the following:
    - **Source Port**
    - **Destination Port**
12. Click **OK**.

Refer to NetFlow Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

Notes and Considerations for SSL Certificate Metadata

Refer to the following notes and considerations for SSL certificate metadata:

- When the GigaSMART SSL application is gathering SSL certificate metadata for NetFlow, it is not able to decrypt SSL packets at the same time.
- Using the SSL application plus NetFlow to collect SSL certificate metadata consumes GigaSMART resources, resulting in less memory for other GigaSMART applications.
- To retrieve SSL certificate metadata, there must be a valid SSL session in which most of the packets that form the session are received, as follows:
  - For the SSL certificate parameters, all packets up to at least the certificate packet must be received.
  - For the SSL server parameters except **SeverName Indication**, all packets up to the server hello packet must be received.
  - For the **ServerName Indication** SSL server parameter, all packets up to the client hello packet must be received.

SNMP Packet Support on IP Interfaces with Tool Ports

SNMP packet support on IP interfaces with tool ports processes SNMP packets arriving on IP interfaces with tool ports and forwards them to GigaSMART so that external NetFlow collectors can integrate with GigaSMARTNetFlow Generation.

External collectors need to recognize GigaSMARTNetFlow Generation as a valid NetFlow interface. The validation can be manual or automatic, as follows:

- with manual validation, configuration on the collector side must provide the details of GigaSMARTNetFlow Generation
- with automatic validation, the recipient collector initiates an SNMP query requesting relevant information and GigaSMARTNetFlow Generation responds to it with the required information

SNMP packet support on IP interfaces with tool ports provides automatic validation. Starting in software version 4.5, GigaSMARTNetFlow Generation processes SNMP packets arriving on IP interfaces with tool ports.

NetFlow records are sent to exporters through IP interfaces with tool ports. Each NetFlow exporter is associated with one external collector. Up to six exporters can send flow records to up to six external collectors.

An IP interface with tool port can have multiple exporters. An external collector can listen to multiple exporters.

To listen to SNMP packets from external collectors, enable SNMP under the NetFlow exporter. The following are the steps to allow external collectors to send SNMP packets to the IP interface with tool port:

1. From the device view, select **GigaSMART > NetFlow > Exporters**.
2. Click **New** to create a new exporter or Edit to configure an existing exporter.
3. Select **SNMP** under the SNMP section.

These steps enables SNMP on the default port, which is port number 161.

NOTE: Only the default SNMP port is supported for packets arriving on the IP interface. If the incoming request uses a non-default SNMP port, they will be dropped at the IP interface.

To disable listening for SNMP packets by a specified NetFlow exporter, uncheck the SNMP check box.

By default, listening to SNMP packets from external collectors is disabled.

NetFlow Format Support on Exporters

NetFlow Exporters support versions IPFIX, v5, and v9. Starting in software version 5.3, the Common Event Format (CEF) version 23 is also supported. CEF is a standard format used by event collection/correlation Security Information and Event Management (SIEM) vendors. SIEMs such as Arcsight, Splunk, and QRadar accept CEF format. By supporting CEF, NetFlow metadata can integrate with and use a variety of SIEMs.

CEF is a logging format that uses the syslog message as a transport mechanism, meaning that the CEF message (header and payload) is included within the syslog message. The transport protocol that is supported is UDP and the default port number is 514.

Metadata that is generated by NetFlow can be exported in the supported formats to one or more collectors. Each exporter must have the same export type (v5, v9, IPFIX, or CEF). One CEF message is sent out per record per flow.

Also, starting in software version 5.3, IP fragmentation is supported. CEF does not allow a message to be split over multiple CEF payloads. Since CEF messages are verbose, they can be larger than the MTU.

To support CEF messages that exceed the MTU, a single IP datagram containing a CEF message will be broken up into multiple packets of smaller sizes. The reassembly of the datagram will occur at the receiving end (at the SIEMs).

For details on the CEF message format, refer to

CEF Message Format

An example of the CEF message format is as follows:

```
Fri Feb 23 02:25:37 2018 9/3/e1
CEF:23|Gigamon|metadata|5.3.00|4|metadatageneration|6| src=68.94.156.1
GigamonMdataDnsAdditionalType=41GigamonMdataDnsAdditionalTypeText=OPT
```

In the example CEF message, there is a syslog header, a CEF header, and an extension that contains the CEF payload. The fields are delimited with a vertical bar (|).

The syslog header contains the following:

- timestamp—Fri Feb 23 02:25:37 2018
- host name identifier—9/3/e1

NOTE: The host name identifier has the format <box ID>/<slot ID>/<engine ID>.For example, 9/3/e1 means 9 is the box ID, 3 is the slot ID, and e1 is the engine ID.

The CEF header contains the following:

- version—CEF:23
- device vendor—Gigamon
- device product—metadata
- device version—5.3.00
- signature identifier—4
- name—metadata generation
- severity—6

The CEF extension contains key-value pairs delimited with a space. In the example CEF message, the following is the CEF payload, in plaintext:

- src=68.94.156.1
- GigamonMdataDnsAdditionalType=41
- GigamonMdataDnsAdditionalTypeText=OPT

The CEF standard specifies key-value pairs. There are some predefined standardkeys, for example, src is a predefined key for source IP address.

For keys that are not predefined in the CEF standard, such as the NetFlow metadata elements in the CEF extension, there are custom-defined keys. Custom-defined keys have the following format:

- <VendorNameProductNameExplanatoryKeyName>

For example, GigamonMdataDnsAdditionalTypeText, is a custom-defined key that contains the following:

- VendorName—Gigamon
- ProductName—Mdata
- ExplanatoryKeyName—DnsAdditionalTypeText

Another example of the CEF format is the following SSL record:

```
Thu Mar 1 08:21:28 2018 1/1/e1 CEF:23|Gigamon|metadata|5.3.00|4|metadata
generation|6|GigamonMdataSslIssuerName=DigiCert SHA2 High Assurance S
dpt=54839 GigamonMdataSslValidNotBefore=3137303130363030303030305a
GigamonMdataSslSerialNo=0118ee3c2167b99e1b718c6eadb8fb4d00000000
GigamonMdataSslValidNotAfter=3230303131353132303030305a
GigamonMdataSslCertSigAlgo=2a864886f70d01010b
GigamonMdataSslCertSubAlgo=2a864886f70d010101
GigamonMdataSslCertSubKeySize=270 GigamonMdataSslServerVersion=771
GigamonMdataSslCertSubAltName=*.stickyadstv.com
GigamonMdataSslServerCompressionMethod=192 GigamonMdataSslServerCipher=49199
GigamonMdataSslServerVersionText=TLSv1.2 GigamonMdataSslServerSessionId=63
GigamonMdataSslIssuer=2f433d55532f4f3d44696769436572742049484132204869676820417373757261 6e636
69676963
6572742e636f6d2f434e3d44696769436572742053484132204869676820417373757261 6e636
52053657276 6572204341 GigamonMdataSslCertSubCommonName=*.stickyadstv.com
GigamonMdataSslSub=2f433d55532f53543d4e657720596f726b2f4c3d4e657720596f726b2f4f3
d4672656
```

5776865656c204d6564696120496e632f4f553d4672656577686565 6c2f434e3d2a2e737469636b796164737 4762e636f6d dst=10.50.22.59 src=38.106.34.118 spt=443

Display Exporter Statistics

To display exporter statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics** and open the GigaSMART Statistics Quick View to view the NetFlow Statistics.

Refer to NetFlow Exporter Statistics Definitions for descriptions of the statistics.

Display Monitor Statistics

To display exporter statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics** and open the GigaSMART Statistics Quick View to view the NetFlow Statistics.

Refer to NetFlow Monitor Statistics Definitions for descriptions of these statistics.

Display IP Interfaces Statistics

To display IP interfaces statistics, select **Ports > IP Interfaces > Statistics** and look for the IP interface ID in the statistics table.

Refer to IP Interfaces Statistics Definitions for descriptions of these statistics.

NetFlow Generation Configuration Modification and Removal

There may be instances where a NetFlow Generation configuration may require alteration by modifying a NetFlow Generation Monitor Configuration or a NetFlow Generation Record Configuration. It may further require that the configuration be removed entirely. In such instances, refer to the following.

Modify NetFlow Generation Monitor Configuration

This example shows the modification of a NetFlow Generation Monitor configuration.

1. Unlink the monitor from GigaSMART Parameters.
    a. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.
    b. Select the GigaSMART group to modify.
    c. Click **Edit**.
    d. Under NetFlow, select **None** in the **Monitor** field.
    e. Click **Save**.
2. Modify the monitor parameters.
    a. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Monitors**.
    b. Select the Monitor to modify.

c. Click **Edit**.

d. Under Config, modify the monitor parameters.

e. Select the record from the **Record(s)** list to re-add it to the monitor.

3. Re-add the monitor to GigaSMART Parameters for the changes to take affect.

a. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.

b. Select the GigaSMART group to modify.

c. Click **Edit**.

d. Under NetFlow, select the monitor in the **Monitor** field.

e. Click **Save**.

Modify NetFlow Generation Record Configuration

This example shows the modification of a NetFlow Generation Record configuration.

1. Unlink the monitor from gsparams.

a. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.

b. Select the GigaSMART group to modify.

c. Click **Edit**.

d. Under NetFlow, select **None** in the **Monitor** field.

e. Click **Save**.

2. Modify the record bound to the monitor.

a. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Records**.

b. Select the record to modify.

c. Click **Edit**.

d. Modify the record configuration.

3. Re-add the monitor to the GigaSMART Parameters for changes in record to take affect.

a. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.

b. Select the GigaSMART group to modify.

c. Click **Edit**.

d. Under NetFlow, select the monitor in the **Monitor** field.

e. Click **Save**.

Remove NetFlow Generation Configuration

Use the following steps to remove a NetFlow Generation Configuration:

1. Remove the NetFlow parameter from the GigaSMART Group.
   a. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.
   b. Select the GigaSMART Group.
   c. Click **Edit**.
   d. Under NetFlow, select **None** in the **Monitor** field.
2. Delete the Maps.
   a. Select **Maps > Maps > Maps**.
   b. Select Table View.
   c. Select the Maps.
   d. Click **Delete**.
3. Delete the IP interface.
   a. Select **Ports > IP Interfaces**.
   b. Select the port.
   c. Click **Delete**.
4. Delete the monitor, records, and exporter
   a. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Monitors**.
   b. Select the monitor, and then click **Delete**.
   c. Select **Records**
   d. Select the record, and then click **Delete**
   e. Select **Records**.
   f. Select the record, and then click **Delete**.

V5 Fixed Record Template

NetFlow v5 records have a template of fixed fields that cannot be edited. The template contains Match/Key and Collect/Non-Key elements. It has an alias of **predefined_netflow_v5_record**.

To display the template, select **GigaSMART > NetFlow / IPFIX Generation > Records** and click on **predefined_netflow_v5_record** to display the Record Quick View shown in Figure 88NetFlow Record predefined_netflow_v5_record.

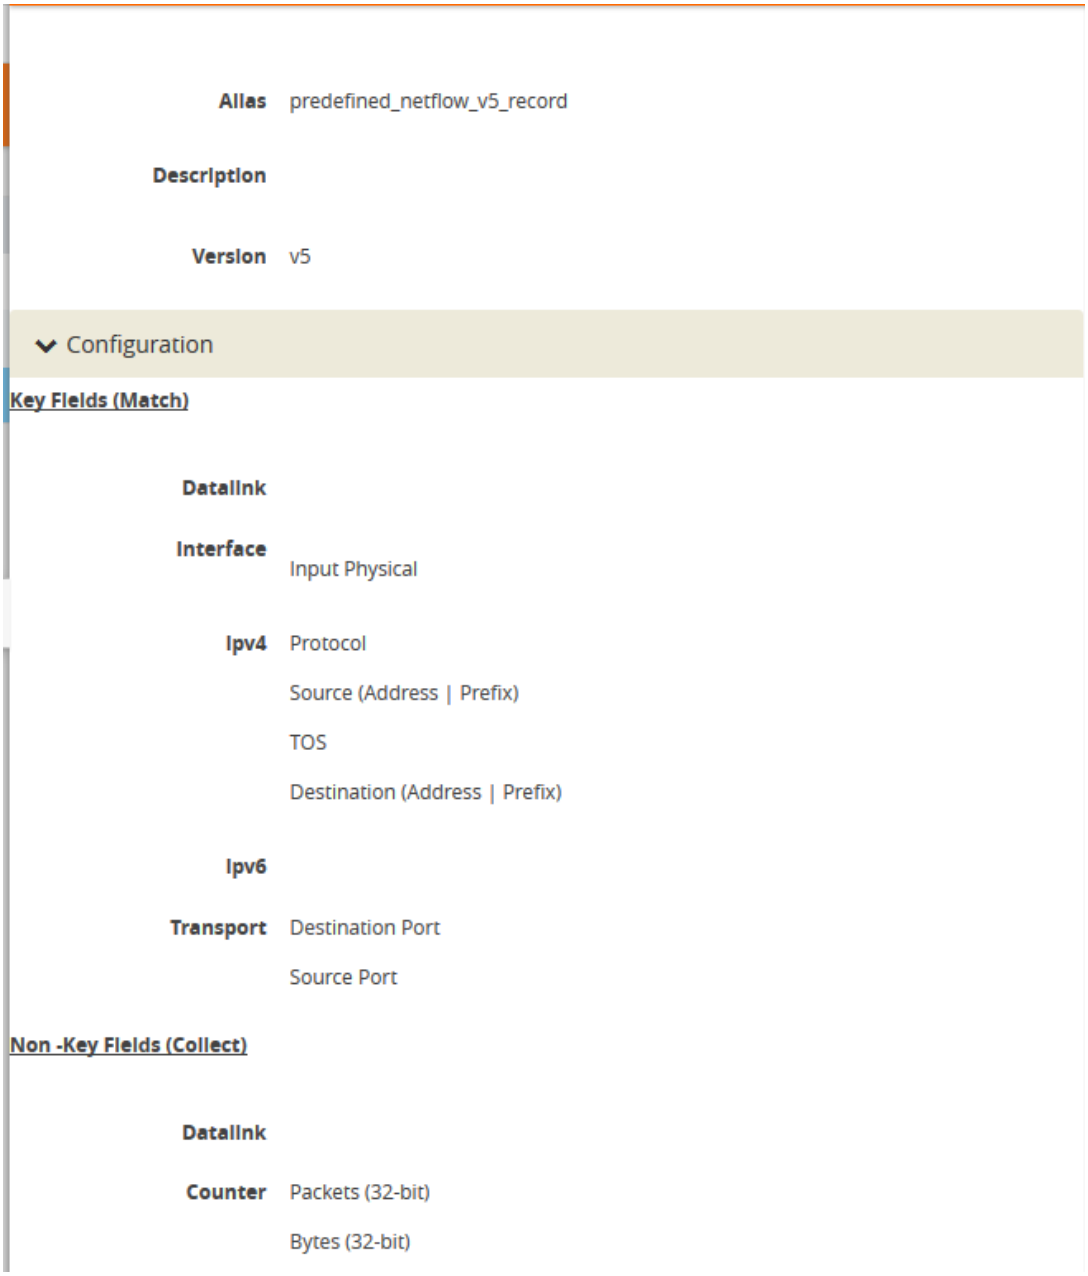


Figure 88 *NetFlow Record predefined_netflow_v5_record*

NetFlow Generation Match/Key and Collect/Non-Key Elements

NetFlow v9 and IPFIX records allow the user to configure Match/Key and Collect/Non-Key elements.

Match/Key Syntax

NetFlow v9 and IPFIX records allow the you to configure Match/Key elements.

NOTE: NetFlow v9 does not support Match/Key elements whose ID on the specified link is greater than 128. For additional information, refer to the following:

http://www.iana.org/assignments/ipfix/ipfix.xhtml

To configure the Match/Key elements, click in the Key Fields (Match) field in the NetFlow Record configuration page and select the match type.

The supported combinations of Match/Key elements are outlined in the following table:

Match Type	Parameters			Description
Data Link	**Source Mac**			Supported for v9 and IPFIX.
	Destination			Supported for v9 and IPFIX.
	VLAN			Supported for v9 and IPFIX.
Interface	**Input physical**	**Physical Width-2** **Physical Width-4**		Supported for v9 and IPFIX. for width, the only supported values are 2 or 4.
IPv4	**Destination**	**Address**		Configures the IPv4 destination address as a key field. Supported for v9 and IPFIX.
		Prefix	**<netmask | mask_ length>**	Configures a prefix for the IPv4 destination address as a key field. Supported for v9 and IPFIX.
	DSCP			Supported only for IPFIX.
	Fragmentation Flags			Supported only for IPFIX.
	Fragmentation ID			Supported for v9 and IPFIX.
	Fragmentation			Supported for v9 and IPFIX.

Match Type	Parameters			Description
	Offset			
	Header Length			Supported only for IPFIX.
	Option Map			Supported only for IPFIX.
	Precedence			Supported only for IPFIX.
	Protocol			Supported for v9 and IPFIX.
	Section	**Header Size**	**<size>**	Configures the number of bytes of raw data starting at the IPv4 header, to use as a key field. The range is from 1 to 128. Supported only for IPFIX.
		Payload Size	**<size>**	Configures the number of bytes of raw data starting at the IPv4 payload, to use as a key field. The range is from 1 to 128. Supported only for IPFIX.
	Source	**Address**		Supported for v9 and IPFIX.
		Prefix	**<netmask | mask_ length>**	Supported for v9 and IPFIX.
	TOS			Supported only for IPFIX.
	Total Length	**maximum | minimum**		Supported only for IPFIX.
	TTL			Supported only for IPFIX.
	Version			Supported for v9 and IPFIX.
IPv6	**Destination**	**Address**		Supported for v9 and IPFIX.
		Prefix	**<netmask | mask_ length>**	Supported only for IPFIX.
	DSCP			Supported only for IPFIX.
	Extension Map			Supported for v9 and IPFIX.
	Flow Label			Supported for v9 and IPFIX.
	Fragmentation Flags			Supported only for IPFIX.
	Fragmentation			Supported for v9 and IPFIX.

Match Type	Parameters			Description
	ID			
	Fragmentation Offset			Supported for v9 and IPFIX.
	Hop Limit			Supported only for IPFIX.
	Length	**Header**		Supported only for IPFIX.
		Payload		Supported only for IPFIX.
		Total		Supported only for IPFIX.
	Next Header			Supported only for IPFIX.
	payload-length			Supported only for IPFIX.
	Precedence			Supported only for IPFIX.
	Protocol			Supported for v9 and IPFIX.
	Section	**Header Size**	**<size>**	Supported only for IPFIX. The range is from 1 to 128.
		Payload Size	**<size>**	Supported only for IPFIX. The range is from 1 to 128.
	Source	**Address**		Supported for v9 and IPFIX.
		Prefix	**<netmask | mask_ length>**	Supported only for IPFIX.
	Traffic Class			Supported for v9 and IPFIX.
	Version			Supported for v9 and IPFIX.
Transport	**Destination Port**			Supported for v9 and IPFIX.
	ICMP	**IPv4**	**Code**	Supported only for IPFIX.
			Type	Supported only for IPFIX.
		IPv6	**Code**	Supported only for IPFIX.
			Type	Supported only for IPFIX.
	Source Port			Supported for v9 and IPFIX.
	TCP	**ACK Number**		Supported only for IPFIX.
		Destination Port		Supported only for IPFIX.

Match Type	Parameters			Description
		Flags <enable | disable>	**[ACK] | [CWR] | [ECE] | [FIN] | [PSH] | [RST] | [SYN] | [URG]**	Supported only for v9 and IPFIX.
		Header Length		Supported only for IPFIX.
		Sequence Number		Supported only for IPFIX.
		Source Port		Supported only for IPFIX.
		Urgent Pointer		Supported only for IPFIX.
		window-size		Supported only for IPFIX.
	UDP	**Destination Port**		Supported only for IPFIX.
		Message Length		Supported only for IPFIX.
		Source Port		Supported only for IPFIX.

Collect/Non-Key Syntax

NetFlow v9 and IPFIX records allow the user to configure Collect/Non-Key elements.

The number of Collect/Non-Key elements in a record can be up to 32. Each Collect/Non-Key element has a size. The accumulated size of the Collect/Non-Key elements in the record cannot exceed 1024 bytes. The supported Collect/Non-Key elements is determined either by the maximum number of elements in a record (32) or by the maximum size (1024 bytes), whichever is reached first.

NOTE: NetFlow v9 does not support Collect/Non-Key elements whose ID on the specified link is greater than 128. For additional information, refer to the following:

http://www.iana.org/assignments/ipfix/ipfix.xhtml

To configure the Collect/Non-Key elements, click in the **Non-Key Fields (Collect)** field in the NetFlow Record configuration page and select the match type.

The supported combinations of Collect/Non-Key elements are outlined in the following table:

Collect Type	Parameters	Size		Description
Counter	**Bytes**	**32** **64**		Supported for v9 and IPFIX.
	Packets	**32** **64**		Supported for v9 and IPFIX.
Datalink	**Source**			Supported for v9 and IPFIX.
	Mac Destination			Supported for v9 and IPFIX.
	VLAN			Supported for v9 and IPFIX.
Flow	**End Reason**			Supported only for IPFIX.
Interface	**Input Name**	**Input Width**	**[width]**	Supported for v9 and IPFIX. for width, the range is from 1 to 32.
	Physical	**Physical Width-2** **Physical Width-4**		Supported for v9 and IPFIX. For width, the only supported values are 2 or 4.
	Output	**Physical Width-2** **Physical Width-4**		Supported for v9 and IPFIX. For width, the only supported values are 2 or 4.
IPv4	**Destination**	**Address**		Configures the IPv4 destination address as a non-key field. Supported for v9 and IPFIX.
		Prefix	**<netmask | mask_ length>**	Supported for v9 and IPFIX.
	DSCP			Supported only for IPFIX.
	Fragmentation Flags			Supported only for IPFIX.
	Fragmentation ID			Supported for v9 and IPFIX.

Collect Type	Parameters	Size		Description
	Offset			Supported for v9 and IPFIX.
	Header Length			Supported only for IPFIX.
	Option Map			Supported only for IPFIX.
	Precedence			Supported only for IPFIX.
	Protocol			Supported for v9 and IPFIX.
	Section	**Header Size**	**<size>**	Configures the number of bytes of raw data starting at the IPv4 header, to use as a key field. The range is from 1 to 128. Supported for v9 and IPFIX.
		Payload Size	**<size>**	Configures the number of bytes of raw data starting at the IPv4 payload to use as a key field. The range is from 1 to 128. Supported for v9 and IPFIX.
	Source	**Address**		Supported for v9 and IPFIX.
		Prefix	**<netmask | mask_ length>**	Configures a prefix for the IPv4 destination address as a non-key field. Supported for v9 and IPFIX.
	TOS			Supported only for IPFIX.
	Total Length	**[maximum]**		Supported only for IPFIX.
		[minimum]		Supported only for IPFIX.
	TTL			Supported only for IPFIX.
	Version			Supported for v9 and IPFIX.
IPv6	**Destination**	**Address**		Supported for v9 and IPFIX.
		Prefix	**<netmask | mask_ length>**	Supported only for IPFIX.
	DSCP			Supported only for IPFIX.
	Extension Map			Supported for v9 and IPFIX.

Collect Type	Parameters	Size		Description
	Flow Label			Supported for v9 and IPFIX.
	Fragmentation Flags			Supported only for IPFIX.
	Fragmentation ID			Supported for v9 and IPFIX.
	Fragmentation Offset			Supported for v9 and IPFIX.
	Hop Limit	**[maximum]**		Supported only for IPFIX.
		[minimum]		Supported only for IPFIX.
	Length	**Header**		Supported for v9 and IPFIX.
		Payload		Supported only for IPFIX.
		Total	**[maximum]**	Supported only for IPFIX.
			[minimum]	Supported only for IPFIX.
	Next Header			Supported only for IPFIX.
	Precedence			Supported only for IPFIX.
	Protocol			Supported for v9 and IPFIX.
	Section	**Header Size**	**<size>**	Supported only for IPFIX. The range is from 1 to 128.
		Payload Size	**<size>**	Supported only for IPFIX. The range is from 1 to 128.
	Source	**Address**		Supported for v9 and IPFIX.
		Prefix	**<netmask | mask_ length>**	Configures a prefix for the IPv4 destination address as a non-key field. Supported only for IPFIX.
	Traffic Class			Supported for v9 and IPFIX.
	Version			Supported for v9 and IPFIX.
Private	**PEN <pen name>**	**DNS**	**<additional-class**	Supported only for IPFIX.

Collect Type	Parameters	Size		Description
			[number-of-collects <1-10>] | additional-class-text [number-of-collects <1-10>] |additional-name [number-of-collects <1-10>] | additional-rd-length [number-of-collects <1-10>] | additional-rdata [number-of-collects <1-10> | width <1-128>] | additional-ttl [number-of-collects <1-10>] | additional-type [number-of-collects <1-10>] | additional-type-text [number-of-collects <1-10>] | an-count | ar-count |authority-class [number-of-collects <1-10>] | authority-class-text [number-of-collects <1-10>] | authority-name [number-of-collects <1-10>] | authority-rd-length [number-of-collects <1-10>] | authority-rdata [number-of-collects <1-10> | width <1-128>] | authority-ttl [number-of-collects <1-10>] | authority-type	

Collect Type	Parameters	Size		Description
			s <1-10>] | authority-type-text [number-of-collects <1-10>] | bits | identifier | ns-count | op-code | qd-count | query-class [number-of-collects <1-10>] | query-class-text [number-of-collects <1-10>] | query-name [number-of-collects <1-10>] | query-type [number-of-collects <1-10>] |	
Private (continued)	**PEN <pen name>**	**DNS**	**query-type-text [number-of-collects <1-10>] | response-class [number-of-collects <1-10>] | response-class-text [number-of-collects <1-10>] | response-code | response-ipv4-addr [number-of-collects <1-10>] |response-ipv4-addr-text [number-of-collects <1-10>] | response-ipv6-addr [number-of-collects <1-10>] | response-ipv6-addr-text [number-of-collects <1-10>] |response-name [number-of-collects <1-10>] |**	Supported only for IPFIX.

Collect Type	Parameters	Size		Description
			response-rd-length [number-of-collects <1-10>] | response-rdata [number-of-collects <1-10> | width <1-128>] | response-ttl [number-of-collects <1-10>] | response-type [number-of-collects <1-10>] |response-type-text [number-of-collects <1-10>]>	
Private	**PEN <pen name>**	**HTTP**	**Response Code**	Supported only for IPFIX.
Private	**PEN <pen name>**	**HTTP**	**URL**	Supported only for IPFIX. For width, the range is from 1 to 250.
Private	**PEN <pen name>**	**HTTP**	**User Agent**	Supported only for IPFIX. For width, the range is from 1 to 250.
Private	**PEN <pen name>**	**SSL Certificate**	**<Issuer [width] | Issuer Common Name [width] | Serial Number | Serial Number Text | Signature Algorithm |Signature Algorithm Text | Subject [width] | Subject Algorithm | Subject Algorithm Text | Subject Alternative Name [width] | Subject Common Name | [width] | Subject Key Size | Valid Not After | Valid Not After Text | Valid Not Before | Valid Not Before Text>**	Supported only for IPFIX. For width of Issuer and Subject, the range is from 1 to 250. For width of Issuer Common Name, Subject Alternative Name, and Subject Common Name, the range is from 1 to 64.

Collect Type	Parameters	Size		Description
Private	**PEN <pen name>**	**SSL Server**	**<Cipher | Cipher Text | Compression Method | Name Indication [width] | Session ID | Version | Version Text>**	Supported only for IPFIX. For width, the range is from 1 to 64.
Private	**PEN <pen name>**	**URL**	**[width]**	Supported only for IPFIX. For width, the range is from 1 to 250.
timestamp	**Sys-uptime First**			Supported for v9 and IPFIX.
	Sys-uptime First Last			Supported for v9 and IPFIX.
transport	**Destination Port**			Supported for v9 and IPFIX.
	ICMP	**IPv4 Code**		Supported only for IPFIX.
		IPv4 Code Type		Supported only for IPFIX.
		ipv6 Code		Supported only for IPFIX.
		ipv6 Type		Supported only for IPFIX.
	Source Port			Supported for v9 and IPFIX.
	TCP Flags	**[ACK] | [CWR] | [ECE] | [FIN] | [PSH] | [RST] | [SYN] | [URG]**		Supported for v9 and IPFIX.
	TCP	**ACK Number**		Supported only for IPFIX.
		Destination Port		Supported only for IPFIX.
		Header Length		Supported only for IPFIX.
		Sequence Number		Supported only for IPFIX.
		Source Port		Supported only for IPFIX.
		Urgent Pointer		Supported only for IPFIX.
		Window Size		Supported only for IPFIX.

Collect Type	Parameters	Size		Description
	UDP	**Destination Port**		Supported only for IPFIX.
		Message Length		Supported only for IPFIX.
		Source Port		Supported only for IPFIX.

NetFlow Cacheless export using TCP protocol

Classic NetFlow operates in a cacheless mode and exports packets using TCP protocol. NetFlow monitor enables you to set the cache timeout to zero and makes the NetFlow to operate in a cacheless mode. Each packet received by the NetFlow is treated as a complete flow and the metadata extracted from the packet are immediately exported without caching the flow.

The two main advantages of a cacheless mode are as follows:

- The metadata are directly sent to the export buffer as they are not copied to cache.
- The cacheless mode provides more record space.

When NetFlow operates in a cacheless mode, the NetFlow sends metadata for every incoming packet, hence the number of flows are equal to the number of incoming packets.

For NetFlow to export packets in TCP, you should Configure Apps Exporter, which is also used by 3GPP CUPS and Tunneling.

Rules and Notes

- You must define at least 2 exporters, such as **apps exporter**, **apps netflow exporter**.
- You must provide same name for the exporters. The name must be followed by a number in the range from 1 to 9.
- You can add up to 64 apps exporters to one NetFlow Engine. Flows is load balanced across these exporters based on the inner IP address and the direction of the flow.
- You can export only IPv4 using TCP through the exporter.
- The maximum number of bytes exported for the inner payload data is 9600.
- If a record has any inner collects and the packet does not have any inner fields , the record will not be collected.

- In a cluster environment, when there is a NetFlow configuration, the backup-and-restore functionality does not work as expected. The restore operation does not push the configuration to the GigaSMART engine of non-master nodes. It is recommended to copy and paste the configuration backup file instead of running the “config text apply” operation.

For more information about the commands, refer to GigaVUE-OS CLI Reference Guide.

Configure Netflow Generation

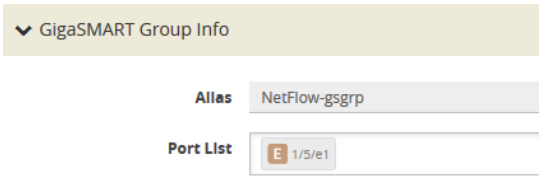
The following are the step for setting up a typical NetFlow Generation configuration with GigaVUE-FM:

- Step 1: Configure a GigaSMART Group
- Step 2: Configure the NetFlow Exporter
- Step 3: Configure an IP Interface
- Step 4: Configure the NetFlow Record
- Step 5: Configure the NetFlow Monitor
- Step 6: Add the NetFlow Monitor to GigaSMART Group
- Step 7: Configure the GigaSMART Operation
- Step 8: Configure Mapping Rules to Filter Packets

Step 1: Configure a GigaSMART Group

Configure a GigaSMART Group using the following steps. you will use this GigaSMART Group in Step 6: Add the NetFlow Monitor to GigaSMART Group, where you assign a NetFlow Monitor to the group.

1. From the device view, select **GigaSMART > GigaSMART Groups**.
2. Click **New** to create a new GigaSMART Group.
3. Enter an alias to help identify this GigaSMART group. For example, Netflow-gsgrp
4. Select an engine port (the **e** port references the GigaSMART line card or module) Your GigaSMART group should look similar to the example shown in the following figure.



5. Click **Save**.

Notes:

- The GigaSMART Group can contain multiple GigaSMART engine ports.
- Only one NetFlow Generation Monitor can be configured per GigaSMART Group.
- Once a GigaSMART group is created, the Alias and port List values are no longer modifiable. You will need to recreate the group to change these values.

Step 2: Configure the NetFlow Exporter

Configure one or more NetFlow Generation Exporters. There can be up to six NetFlow Generation Exporters for each NetFlow Generation Monitor.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure the NetFlow Exporter, do the following:

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Exporters**.
2. Click **New**. The NetFlow Exporters page appears.
3. On the NetFlow Exporter page, enter the information for the exporter. Table 9: NetFlow ExporterConfiguration Fields describes the fields.

   NOTE: The NetFlow version must be configured with the same version of the Exporter and the Record. If no version is specified, version 9 is the default.
4. Under the **Filters** section, click **Add a Rule** to create a filter for the exporter.
5. Click **Save**.

Table 9: NetFlow ExporterConfiguration Fields

Field	Description
Alias	The alias name for the NetFlow Exporter.
Description	An optional description of the NetFlow record.
Format	The format is either NetFlow or CEF.
Version	The version is either NetFlow-v9, NetFlow-v5, or IPFIX.
Template Refresh Interval	After each template-refresh-interval, the record template is sent to the collector. Also, the option template is sent.
SNMP	Enables SNMP packet support on IP interfaces associated with the NetFlow Exporter.
Transport Protocol	The UDP port of the collector. This value cannot be changed.

Field	Description
IP Version	IP Version of the destination IP. You can select IPv4 or IPv6. Default is set as v4.
Destination IP	The IP address of the NetFlow/IPFIX collector. Default is set as 0.0.0.0.
Destination Port	Port for the destination IP. Default is set as 2055.
DSCP	The DSCP priority of the packet. Default is set as 0.
TTL	The Time to Live of the packet. Default is set as 64.

Step 3: Configure an IP Interface

In this step, you identify the collector port and configure it as a tool port, where the NetFlow collector will be connected, and then configure an IP interface. The steps are as follows:

1. Select the port to use and configure it as a tool port.
    a. Select **Ports > Ports > All Ports**.
    b. Click the **Quick Port Editor** button to open the Quick Port Editor.
    c. In the Quick Port Editor select the port to use for the IP interface, provide an alias to help identify the port (for example, NetFlow_Tunnel_Port), select **Tool** for the port type, and select **Enable**.
    d. Click **OK**.
2. Select **Ports > IP Interfaces**.
3. Click **New**.
4. On the IP Interface page, do the following:
    a. In the **Alias** and **Description** fields, enter a name and description for the IP interface.
    b. From the **Port** drop-down list, select the tool port that you configured in Step 1.
    c. Select the type of IP interface as either **IPv4** or **IPv6**.
    d. Enter the **IP Address**, **IP Mask**, **Gateway** address, and **MTU** value.
    e. From the **GigaSMART Group** drop-down list, select the GigaSMART group you created in Step 1: Configure a GigaSMART Group.
    f. From the Exporters drop-down list, select the NetFlow exporter you created in Step 2: Configure the NetFlow Exporter.

Step 4: Configure the NetFlow Record

Configure a NetFlow Generation Record that has the following:

- **match** parameters that identify unique flows
- **collect** parameters that identify fields you want to collect for the unique flows

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure the NetFlow Record, do the following:

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Records**.
2. Click **New**. The NetFlow record page shown in Figure 89NetFlow Record Page displays.

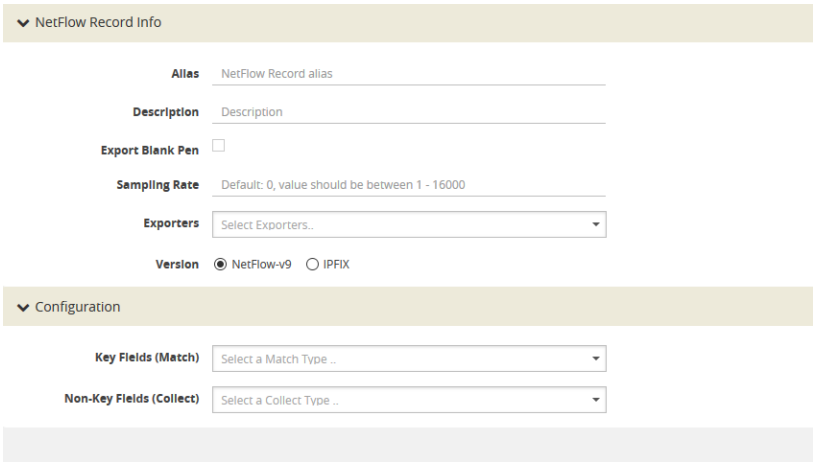


Figure 89 *NetFlow Record Page*

3. On the NetFlow Record page, do the following:

Table 10: NetFlow Record Configuration Fields

Field	Description
Alias	The alias name that would help identify the NetFlow record.
Description	An optional description of the NetFlow record.
Sampling Rate	The **Sampling Rate** is **multi-rate** only, and is specified as 1 in N, where N is the packet count. The packet count can be a number from 1 to 16000. Refer to Configure Netflow Generation. The **Sampling Rate** is disabled by default.
Exporter	Select the Exporter that you want from the **Exporters** menu.

Field	Description
Version	The version is either NetFlow-v9 or IPFIX. The NetFlow version must be configured with the same version of the Exporter and the Record. **NetFlow-v9** is the default. NetFlow-v9 and **IPFIX** let you configure Match/Key and Collect/Non-Key elements. Make sure that you configure the NetFlow version prior to configuring the match and collect parameters because the subsequent parameters depend on the NetFlow version configured.
Key Fields (Match)	The parameters that identify unique flows. The available Match/Key fields are based on the configured NetFlow version.
Non-Key Fields (Collect)	The parameters that identify what you want to collect for the unique flows. The number of Collect/Non-Key elements in a record can be up to 32. From the drop down select the following option: Exporter- You can collect IPv4 or IPv6 switch management interface address by including additional collects in their configuration. You can select any of the following options: **IPv4 address**—Adds new switch or router management IPv4 address. **IPv6 address**—Adds new switch or router management IPv6 address. For details about the match and collect parameters, refer to Configure Netflow Generation

Step 5: Configure the NetFlow Monitor

Configure a NetFlow Generation Monitor and associate the NetFlow Generation Record to the specified NetFlow Generation Monitor by doing the following:

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Monitors**.
2. Click **New**. The Monitors page displays.
3. On the Monitors page, do the following:
   a. Enter an **Alias** to identify the monitor.
   b. Enter a **Description** (optional).
   c. Configure the **Cache** parameters. Refer to Table 11: NetFlow Monitor Parameters.
   d. Configure the **Sampling** parameters. Refer to Table 11: NetFlow Monitor Parameters.
   e. Select the **Record** that you configured in Step 4: Configure the NetFlow Record.
4. Click **Save**.

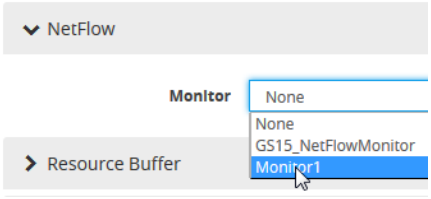
Table 11: NetFlow Monitor Parameters

Parameter	Description
Cache Type	Set as Normal.
Cache Timeout Active	Despite the flow being active, it is “flushed out” to the Exporter after this timeout, which is set in seconds.
Cache Timeout Inactive	Inactive flows are “flushed out” to the Exporter after this timeout, which is set in seconds.
Cache Timeout Event	Applies to the TCP flow. The flow is “flushed out” to the Exporter after detecting a FIN or RST.
Mode	Select the sampling mode that you want: • No sampling • Multi rate • Single rate
Single Sampling Rate	Refer to Configure Netflow Generation.

Step 6: Add the NetFlow Monitor to GigaSMART Group

Return to the GigaSMART Group configured in Step 1: Configure a GigaSMART Groupand set the NetFlow Monitor to the monitor created in Step 2: Configure the NetFlow Exporter.

1. From the device view, select **GigaSMART > GigaSMART Groups> GigaSMART Groups**.
2. Select the GigaSMART Group configured inStep 1: Configure a GigaSMART Group, and then click **Edit**.
3. Under GigaSMART Parameters, go to NetFlow. Click in the **Monitor** field and select the NetFlow monitor configured in Step 5: Configure the NetFlow Monitor as shown in the following figure.



4. Click **Save**.

Step 7: Configure the GigaSMART Operation

Define a GigaSMART operation to enable NetFlow Generation. If combining NetFlow with APF or De-duplication GSOPs, make sure that you select both operations when creating the GigaSMART Operation.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure the GigaSMART Operation, do the following:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP)**.
2. Click **New**. The GigaSMART Operations (GSOP) page displays. (Refer to Figure 90GigaSMART Operation (GSOP) Page.)
3. On the GigaSMART Operations page, do the following:
   a. In the **Alias** field, enter a alias to help identify this gsop.
   b. In the **GigaSMART Groups** field, select the gsop configured in Step 1: Configure a GigaSMART Group.

      aIn the **GigaSMART Operations (GSOP)** field, select **NetFlow**. The NetFlowGigaSMART Operation is enabled by default as shown in Figure 90GigaSMART Operation (GSOP) Page.

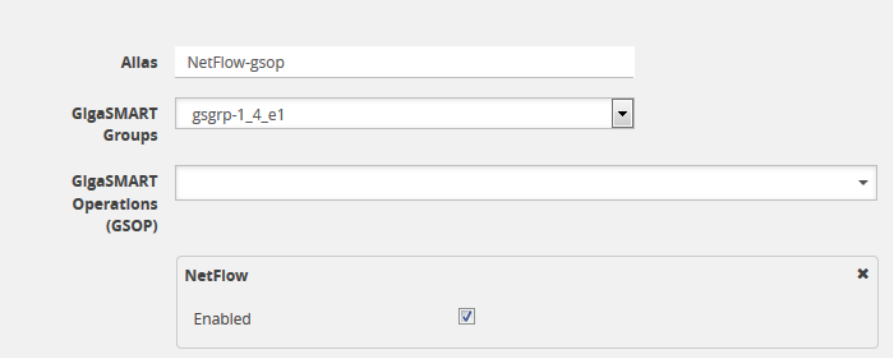


Figure 90 *GigaSMART Operation (GSOP) Page*

4. Click **Save**.

Step 8: Configure Mapping Rules to Filter Packets

To add Flow Mapping® rules to filter packets that are needed to run NetFlow, configure a map and associate the map to the IP interface with tool port.

For more detailed information about Flow Mapping®, refer to About Flow Mapping® and Manage Maps.

Notes:

- For a single NetFlowGigaSMART Operation, make sure that you create a Regular By Rule map. When combining with APF or De-duplication, use First Level or Single Level map types.
- Make sure that the other combining GigaSMART Operations are configured before creating maps using NetFlow.
- When combining NetFlow with APF or De-duplication, create virtual ports to use with the second level maps.
- The destination tool port must be the IP interface with tool port identified in Step 3: Configure an IP Interface

For second level maps, you will need to create virtual ports. To create virtual ports, do the following:

1. From the device view, select **GigaSMART > Virtual Ports**.
2. Click **New**. The Virtual Ports page displays.
3. Enter an alias in the **Alias** field to identify the virtual port.
4. In the **GigaSMART Groups** field, select the GigaSMART Group configured in Step 1: Configure a GigaSMART Group.
5. Click **Save**.

To configure mapping rules to filter packets, do the following:

1. Select **Maps > Maps > Maps**.
2. Click **New** to create a new map.
3. On the New Map page, do the following:
    a. Enter an alias in the **Alias** field and select the map **Type** and **Subtype**.
    b. Specify **Source** and **Destination** ports.
    c. In the **GigaSMART Operations (GSOP) field**, select the GigaSMART Operation configured in Step 7: Configure the GigaSMART Operation.
    d. Click **Add a Rule** to add the rules needed for the map.
4. Click **Save**.

GigaSMART Packet Slicing

Required License: Base

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

GigaSMART operations with a **Slicing** selected truncate packets after either a specified header/layer and offset (a **relative** offset) or at a specific offset. Slicing operations are typically configured to preserve specific packet header information, allowing effective network analysis without the overhead of storing full packet data.

Packets can have multiple variable-length headers, depending on where they are captured, the different devices that have attached their own headers along the way, and the protocols in use (for example, IPv4 versus IPv6). Because of this, slicing operations with a hard-coded offset will not typically provide consistent results.

To address this, the GigaSMART lets you configure packet slicing using **relative offsets** – a particular number of bytes after a specific packet header (IPv4, IPv6, UDP, and so on). The GigaSMART parses through Layer 4 (TCP/UDP) to identify the headers in use, slicing based on the variable offset identified for a particular header instead of a hard-coded number of bytes.

Keep in mind the following when configuring GigaSMART operations with a **Slicing** component:

Feature	Description
Protocol	The following are the protocols that you can select for from the protocol drop-down list: o **IPV4** – Slice starting a specified number of bytes after the IPv4 header. o **IPV6** – Slice starting a specified number of bytes after the IPv6 header. o **UDP** – Slice starting a specified number of bytes after the UDP header. o **TCP** – Slice starting a specified number of bytes after the TCP header. o **FTP** – Identify using TCP port 20 and slice payloads using offset from the TCP header. o **HTTPS** – Identify using TCP port 443. Slice encrypted payloads using offset from the TCP header. o **SSH** – Identify using TCP port 22. Slice encrypted payloads using offset from the TCP header. The GigaSMART can provide slicing for GTP tunnels, provided the user payloads are

Feature	Description
	unencrypted. Both GTPv1 and GTPv2 are supported – GTP' (GTP prime) is not supported. Keep in mind that only GTP-u (user plane packets) are sliced. Control plane packets (GTP-c) are left unmodified because of their importance for analysis. o **GTP** – Slice starting a specified number of bytes after the outer GTP header. o **GTP-IPV4** – Slice starting a specified number of bytes after the IPv4 header inside the encapsulating GTP packet. o **GTP-UDP** – Slice starting a specified number of bytes after the UDP header inside the encapsulating GTP packet. o **GTP-TCP** – Slice starting a specified number of bytes after the TCP header inside the encapsulating GTP packet.
Slicing Offsets	You can specify either a **relative** offset or a **static** offset for the start of the packet slice: ▪ **Static** offsets begin slicing a specific number of bytes from the start of the packet. Choose a static offset by setting **protocol** to **none** and supplying an **offset** from **<64~9000>** bytes. ▪ **Relative** offsets begin slicing a specified number of bytes from the end of a particular header – IPv4, IPv6, and so on. Choose a relative offset by selecting any of the values listed for the **protocol** argument, along with an offset of **<4~9000>** bytes from the end of the specified layer:
Recalculated CRC	GigaSMART packet slicing automatically calculates and appends a new four-byte Ethernet CRC based on the sliced packet's length and data and uses it to replace the existing CRC. This way, analysis tools do not report CRC errors for sliced packets. **Note**: The minimum relative offset is 4 bytes to allow the recalculated CRC to be added. The packet is sliced at the relative offset, and then the recalculated 4 bytes CRC is added to the sliced packet.
GigaSMART Trailer	Slicing operations can optionally include the GigaSMART Trailer. If you do elect to include the GigaSMART Trailer, it will include the original packet length before slicing. **Note**: Refer to How to Use GigaSMART Trailers for details on when the GigaSMART Trailer is required for a GigaSMART Operation as well as the information found in it.
In Combination with Masking	Slicing operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to Groups of GigaSMART Engine Ports for details.

Example – GigaSMART Packet Slicing

This example creates a GigaSMART slicing operation named **IPv6_Headers**. This operation truncates all packet data starting four bytes after the IPv6 header. The sliced packet would include the DLC, IPv6, and TCP headers, which are often enough for analysis needs.

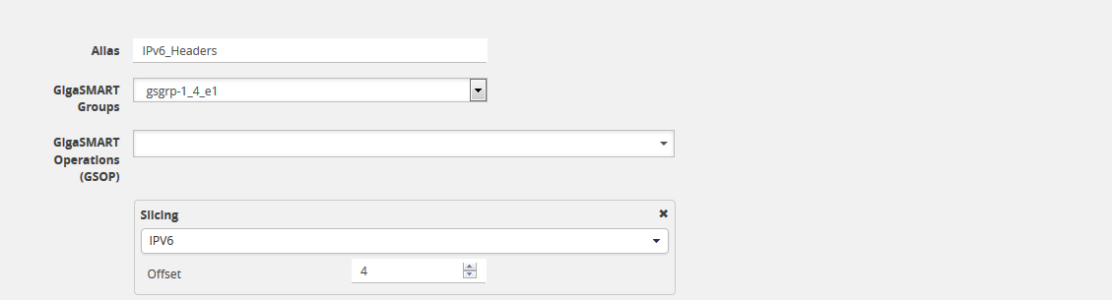


Figure 91 *GigaSMART Operations (GSOP) Page with Slicing Selected*

Display Slicing Statistics

To display slicing statistics, select **GigaSMART > GigaSMART Operations > Statistics**. The statistics for slicing will be in the row labeled Slicing in the GS Operations column.

Refer to Slicing Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

GigaSMART Advanced Flow Slicing

Required License: Advanced Flow Slicing

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus.

Advanced Flow Slicing (AFS) allows you to slice traffic on multiple network protocols, each with different offset. In Advanced Flow Slicing, you can configure multiple protocols within a profile with rules for slicing of a packet which in turn reduces the number of GigaSMART Operations (GSOP) and the maps to be configured. In this feature, the slicing of packets occurs after the packet count has reached a configured value thereby preventing the slicing of control information and other important data of the networking protocols.

Limitations

- Supports up to five enhanced profiles.
- Supports up to ten protocol rules for each profile.
- Slicing profile cannot be edited after attaching it to GSOP.
- Supports up to 80 million sessions per GigaSMART group.
- When enhanced slicing or AFS profile is configured with 7-tuple, the protocol rule must be configured with "flow-session inner" only.

NOTE: Only 5-tuple hashing is supported for Gen2. The 7-tuple hashing is supported for Gen3.

NOTE: It is recommended not to edit an Advanced Flow Slicing (AFS) or the enhanced-slicing profile once created. For any changes to a profile, the user must always delete and recreate with new configurations.

Create Advanced Flow Slicing Profile

To create an advanced flow slicing profile, follow these steps:

1. Go to **Physical > Physical Nodes**.
2. Click the required cluster ID.
3. On the left navigation pane, go to **System** > **GigaSMART> Advanced Flow Slicing**. The Advanced Flow Slicing page appears.
4. Click **New**.
   - Enter an allias.
   - Enter the Maximum Sessions. In a profile, you can configure from 4 to 80 million sessions.
   - Select the Protocol from the drop-down list.
   - Choose the protocol type as Inner or Outer.
   - Enter the value for Offset. The value specifies the number of bytes that should be sliced after the protocol header. The value ranges from 64 to 9000, if there is no protocol selected. The value ranges from 0 to 9000 when other protocol is selected.
   - Select the Flow Session from the drop-down list. For a profile with flow-session defined, a session is created when a packet is received and when there is no existing flow-session for that flow.

     Slicing or dropping starts on the next packet of a session after the number of packets reaches the specified value in the packet count.
   - Choose either to Slice or Drop the session. The default action is Slice.

     The slicing or drop occurs for the first matching rule in a profile.
   - Enter the timeout sessions from 10 to 300 seconds. The default value is 30 seconds.
   - Enter the Skip Packet Count value to slice or drop the packet after it reaches the given count value.

   Each profile must have atleast one protocol field. You can add a maximum of 10 rules in a profile. Click **+** to add new rules in a profile.
5. Click **OK**.

Slicing or drop cannot be performed in the following conditions:

- When the traffic does not match any rule in the profile.
- When a packet does not contain the configured inner or outer IP and l4-port in a defined flow-session.

NOTE: It is recommended not to edit an Advanced Flow Slicing (AFS) or the enhanced slicing profile once created. For any changes to a profile, you must always delete the existing profile and recreate the profile with the new configurations.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

Example-GigaSMART Advanced Slicing

This example shown in creates a GigaSMART Advanced slicing operation named **eslice-gtpu**. In this operation, slicing occurs at HTTPS (port 443) over GTPu-TCP and GTPu-UDP traffic after 10 packets for each TLS connection. This operation truncates all packet data starting 40 bytes after the TCP header. The slicing starts after TCP header gets encapsulated by GTP.

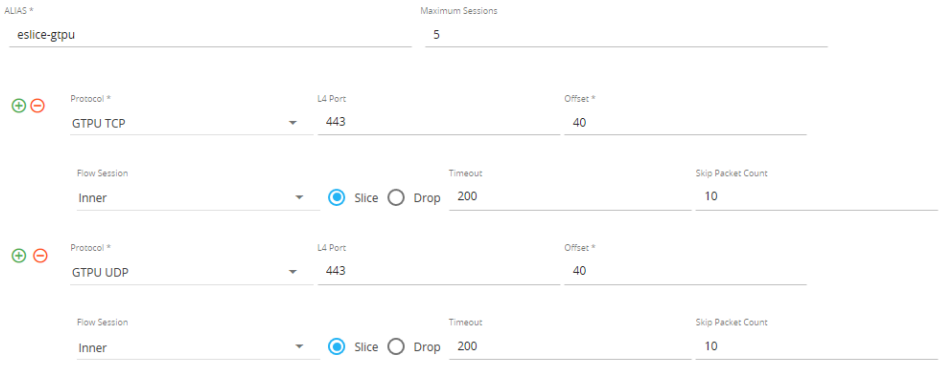


Figure 92 *Advanced Flow Slicing*

Display Slicing Statistics

To display Advanced Slicing statistics, select **System** > **GigaSMART> Advanced Flow Slicing**> **Statistics**. You can view the statistics of the Advanced Flow Slicing.

GigaSMART SSL Decryption

Refer to GigaSMART TLS/SSL Decryption for Out-of-Band Tools for more details on SSL decryption for inline and out-of-band tools.

PCAPng Application

The PCAPng application is a GigaSMART parser application that reads the various blocks in the received PCAPng files and validates the blocks to be sent to the destination application or to the tools. The PCAPng file contains the following blocks:

- Mandatory Blocks
  - Section Header Block (SHB)
- Optional Blocks
  - Interface Description Block (IDB)
  - Enhanced Packet Block (EPB)
  - Simple Packet Block
  - Name Resolution Block
  - Interface Statistics Block

The actual packets are present in the Enhanced Packet Block. The block data is parsed to find the start and end offset of the valid packets and the packet is sent out to the next application.

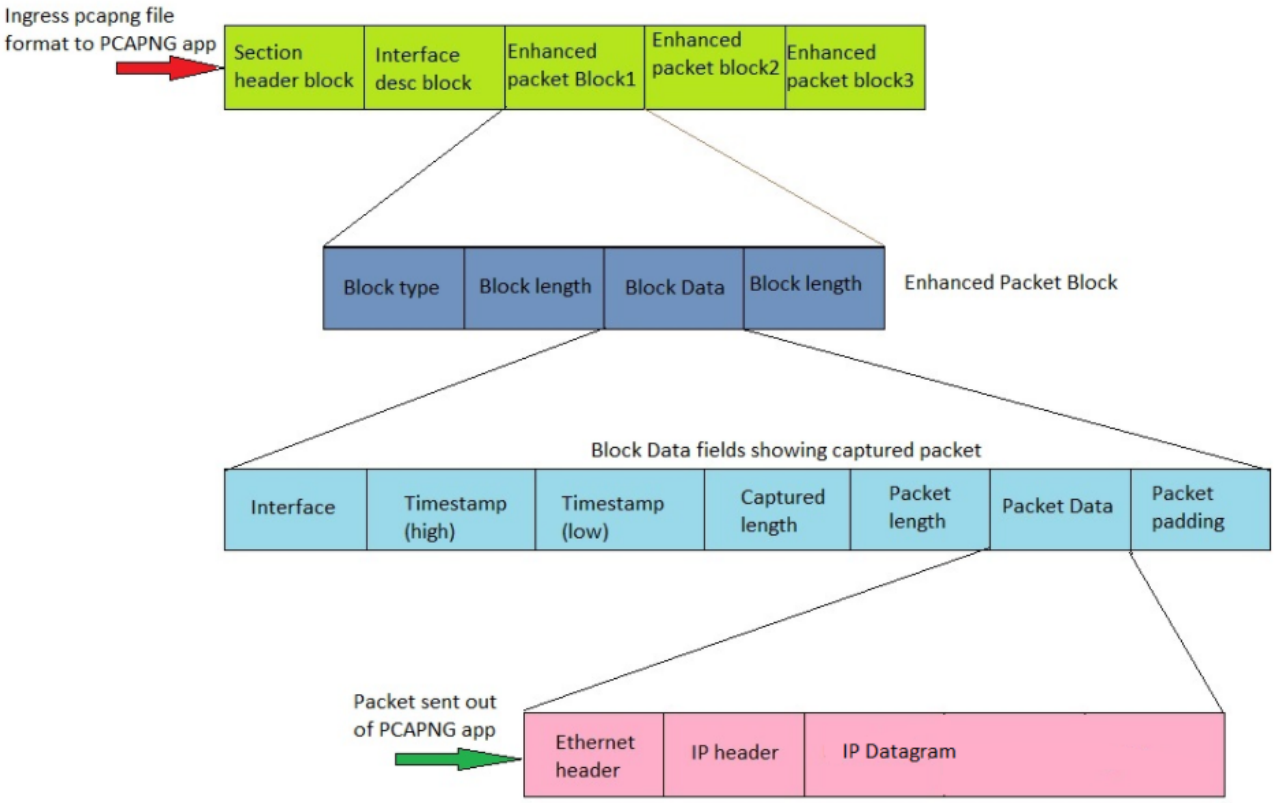
NOTE: Only one EPB in a PCAPng file is supported.

The PCAPng application processes the data depending on the packet type that contains a combination of the blocks mentioned above:

Block Combination	Process
SHB+IDB+EPB+data	Packets are parsed, validated, and the data packet is sent out.
SHB+IDB	Packets are dropped.
EPB+Data	Packets are parsed, validated, and the data packet is sent out.

The PCAPng application validates if the incoming data matches any of the above three formats in the same order, and processes the packets accordingly.

The following figure shows a sample PCAPng file format that contains one section header block:



Tunneling Operations

GigaSMART tunneling operations:

GigaSMART IP Encapsulation/Decapsulation (GigaSMART Tunnel)

Required License for IP Decapsulation: Base (GigaVUE-HC3), Tunneling (GigaVUE-HC1)Required License for IP Encapsulation: Advanced Tunneling (GigaVUE-HC3), Tunneling (GigaVUE-HC1)

Use GigaSMART encapsulation and decapsulation operations to send traffic arriving on one GigaSMART-enabled node over the Internet to a second GigaSMART-enabled node. There, the traffic is decapsulated and made available to local tool ports.

This feature is useful when instrumenting remote data centers – you can tunnel selected portions of the traffic from the remote GigaSMART-enabled node to tools in a central location. Traffic is encapsulated at the sending end of the tunnel and decapsulated at the receiving end.

GigaSMART drops packets when it receives packets with multiple encapsulation headers that are added due to unpredictable looping in the network.

IP fragmentation and reassembly are supported. Refer to IP Fragmentation and Reassembly on L2GRE and GMIP Tunnels.

The source of the GigaSMART tunnel can be any of the following:

- **GigaSMART-Enabled GigaVUE H Series Node**
  - Standalone GigaVUE-HC3 node with SMT-HC3-C08 modules installed.
  - Standalone GigaVUE-HC1 nodes.
  - Any GigaVUE HC Series node operating in a cluster with the previous node types.
- **GigaVUE V Series node or a GigaVUE-VM**

NOTE: You can also create GigaSMART operations that allow a GigaVUE H Series node to act as the receiving end of an ERSPAN tunnel for data mirrored over the Internet from Cisco equipment. However, this feature requires the Advanced Tunneling license; refer to GigaSMART ERSPAN Tunnel Decapsulation.

Display GMIP Tunnel Decapsulation Statistics

To display tunnel decapsulation statistics, select **GigaSMART > GigaSMART Operations > Statistics** and click a on the GS Operation in the table to open the Quick View for GS Operations Statistics.

Refer to Tunnel Decapsulation Statistics Definitions and GigaSMART Operations Statistics Definitions for descriptions of these statistics.

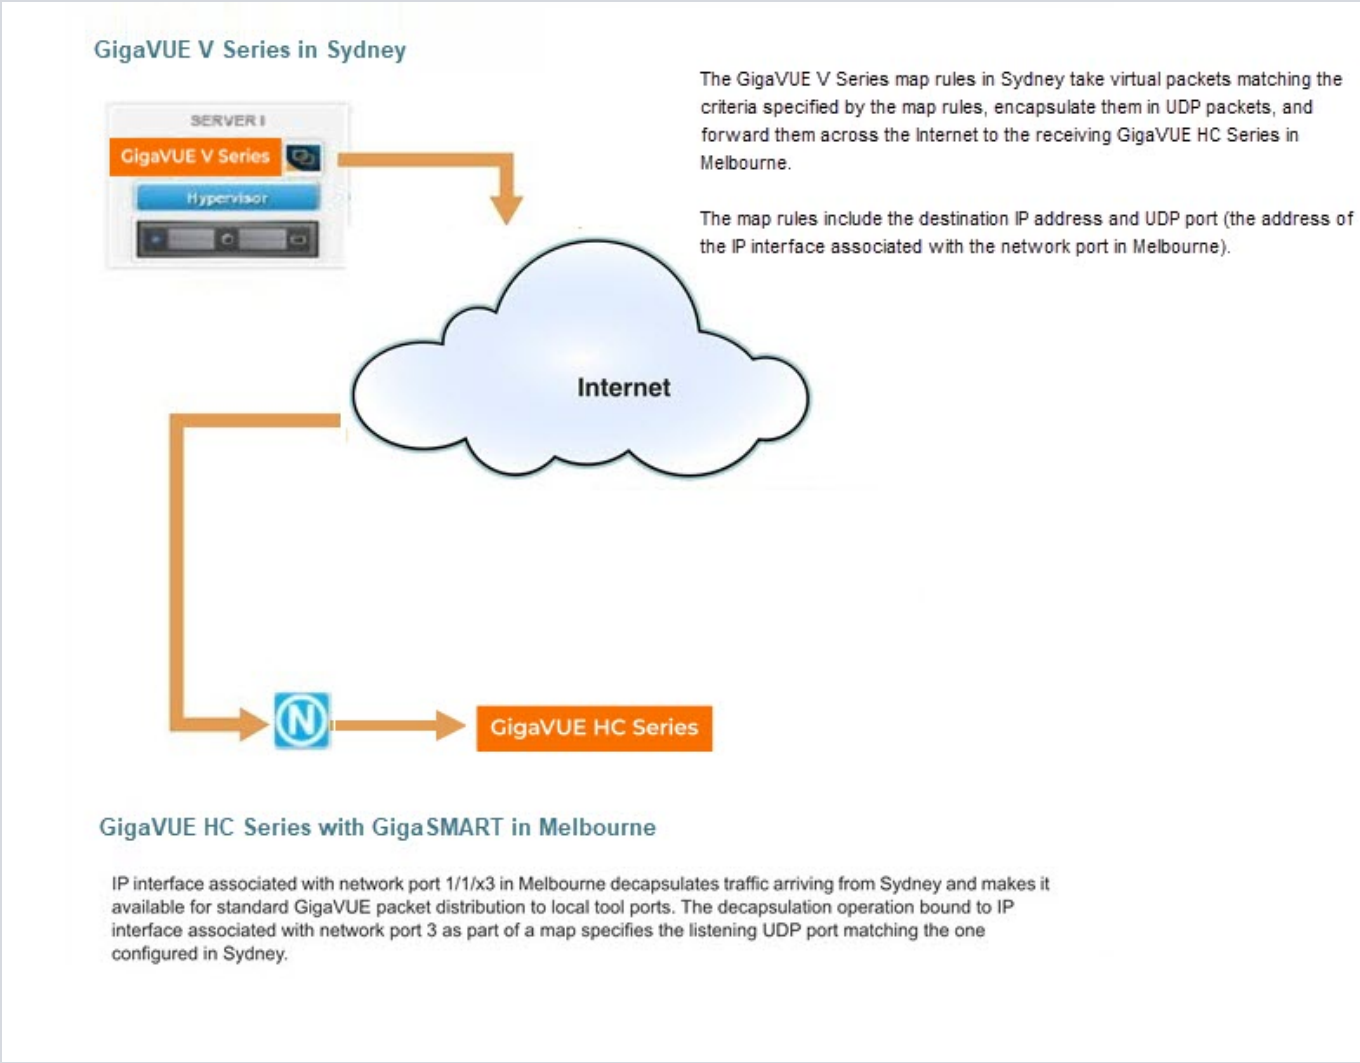
Display GMIP Tunnel Encapsulation Statistics

To display tunnel encapsulation statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics** and click on the GS Operation in the table to open the Quick View for GS Operations Statistics.

Refer to Tunnel Encapsulation Statistics Definitions and GigaSMART Operations Statistics Definitions for descriptions of these statistics.

Example: GigaSMART Encapsulation/Decapsulation (GigaVUE V Series)

The following figures demonstrate how to create a sample IP tunnel between a sending GigaVUE V Series node in Sydney and a receiving GigaVUE HC Series in Melbourne. First, the overall tunnel is summarized, followed by configuration descriptions for the sending and receiving ends.



Configure Sending End of Tunnel: GigaVUE V Series map rules in Sydney

A GigaVUE V Series node in this location is configured with map rules that will send data over the Internet to the IP interface associated with a network port on a GigaVUE HC Series.

Map rules are created in the GigaVUE-FM user interface – Step 2 in the Create Map wizard includes **Tunnel Traffic To** settings that specify where matching traffic should be sent:

Create "Tunnel Traffic To" Option	Setting
UDP IP	This is the destination IP address for the IP interface associated with network port on the GigaVUE H Series in Melbourne. We will set it to 10.150.68.222
UDP Source Port	This is the UDP source port from which tunneled packets will be sent. We will set this to 5000.
UDP Destination Port	This is the listening port on the receiving GigaVUE H Series IP interface associated with network port. We will set this to 10000.

Configure Receiving End of Tunnel: GigaVUE HC Series with GigaSMART in Melbourne

Now we need to configure the receiving end of the tunnel with an IP interface associated with network port. The GigaVUE H Series in this location will have an IP interface associated with network port configured on network port 1/1/3 with an IP address of 10.150.68.222 and a GigaSMART decapsulation operation that listens on UDP port 10000.

The following table summarizes the steps necessary to configure the receiving end of the tunnel using the UI:

Task	UI Steps
Start by designating port 1/1/x3 as an IP interface with network port, configuring its IP profile, and assigning its GigaSMART operations to a GigaSMART group. This command sets the IP address, subnet mask, default gateway, and MTU for the IP interface associated with a tool port on port 1/1/x3.	**1.** Select Ports > IP Interfaces. **2.** Click New. **3.** Configure the IP Interface: ▪ Alias: 1_1_x3 ▪ Port: 1/1/x3 ▪ IP Address: **10.150.68.222** ▪ IP Mask: 255.255.255.255 ▪ Gateway: **10.150.68.1** ▪ MTU: 9400 ▪ GigaSMART Group: GS2 **4.** Save
Now, create an IP decapsulation GigaSMART operation (**gmipdecap**) that will decapsulate traffic received on UDP port 10000. Recall that we configured the sending end of the tunnel to send	**1.** From the device view, select GigaSMART > GigaSMART Operations > GigaSMART Operation. **2.** Click New. **3.** Configure the GigaSMART Operation:

Task	UI Steps
to that UDP port. The operation has the alias **gv_ipdecap**. Note that this operation uses the same GigaSMART group (GS2) as the IP interface associated with network port we set up in the first step.	▪ Alias: **gv_ipdecap** ▪ GigaSMART Groups: GS2 ▪ GigaSMART Operations (GSOP): Tunnel Decapsulation **4.** Configure the Tunnel Encapsulation: ▪ GMIP ▪ GMIP Port: 10000 **5.** Save.
Once we have our IP decapsulation operation, we can include it as part of a map. ▪ Open the map configuration page to create a map named **decapper**. ▪ The **Source** field specifies the ingress ports for this map. ▪ The **GSOP** field applies the gv_ipdecap GigaSMART operation to all packets matching the rules in the map, decapsulating them from the tunnel. ▪ The **Destination** field specifies where matching packets will be sent (tool port 1/1/x11). ▪ The rule with **Pass** selected specifies that packets arriving on this port with an IP Source address of 10.10.10.10 /32 will be processed by the gv_ipdecap GSOP and sent to tool port 1/1/x11.	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. ▪ Alias: decapper ▪ Type: Regular ▪ Subtype: By Rule ▪ Source: 1/1/x3 ▪ Destination: 1/1/x11 ▪ GigaSMART Operation (GSOP): gv_ipdecap (GS2) **4.** Click Add Rule. ▪ Select Pass. ▪ Select IPv4 Source for Rule 1. ▪ Set the IPv4 Address to 10.10.10.10 ▪ Set the Net Mask to 255.255.255.255 **5.** Click Save.

Configure IP Interfaces for GigaSMART Tunnels

When you configure IP interfaces for GigaSMART tunnels:

- You can associate the IP interfaces with multiple GigaSMART groups that are created either in the same node or in another node that resides in the same cluster.
- You can associate multiple GigaSMART engines to a GigaSMART group.
- You can also associate NetFlow exporters to the IP interface.

About IP Interface Centralization

A tunnel that originates from a node in a cluster can terminate on a remote port in another cluster. Also, a tunnel can have multiple termination points. You can associate the IP interface with multiple GigaSMART groups that are created either in the same node or in another node that resides in the same cluster. Moreover, you can associate multiple GigaSMART engines to a GigaSMART group.

The following figure illustrates the tunnel centralization feature.

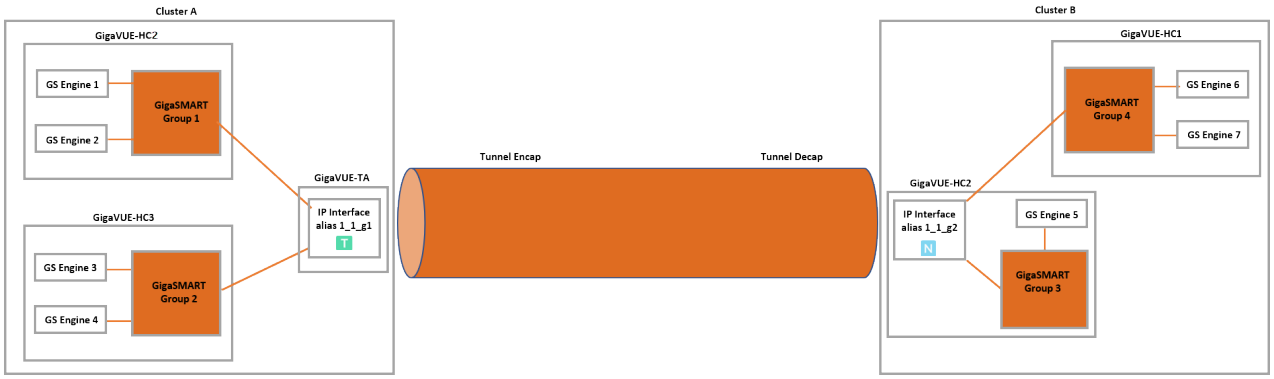


Figure 93 *Tunnel Centralization*

In this example, a GigaVUE-HC3, and GigaVUE-TA nodes reside in cluster A. In GigaSMART engines 1 and 2 are associated with the GigaSMART group 1. Similarly, in GigaVUE-HC3, GigaSMART engines 3 and 4 are associated with GigaSMART group 2. An IP interface with alias 1_1_g1 is configured with a tool port in GigaVUE-TA. Both the GigaSMART groups, 1 and 2 are associated with the IP interface. The IP interface 1_1_g1 is the originating point for the tunnel, where encapsulation happens.

Now, let us look at the termination point of the tunnel. The GigaVUE-HC1 a nodes reside in cluster B. In GigaVUE-HC1-Plus, the GigaSMART engine 5 is associated with GigaSMART group 3 and in GigaVUE-HC1, the GigaSMART engines 6 and 7 are associated with GigaSMART group 3. An IP interface with alias 1_1_g2 is configured with a network port in GigaVUE-HC1-Plus. Both the GigaSMART groups, 3 and 4 are associated with the IP interface. The IP interface 1_1_g2 is the termination point for the tunnel, where decapsulation happens. Thus the tunnel terminates on a remote port in another cluster.

NOTE: Do not attempt to configure with IP /Tunnel Interface and GigaVUE-FM Management Interface in the same subnet. As per routing functionality this is not a valid configuration.

For configuring IP Interface for GigaSMART tunnels, refer to Configure IP Interface.

Configure GigaSMART IP Encapsulation/Decapsulation

This section contains the following topics:

- Configure Both Ends of GigaSMART Tunnel
- Configure Sending End of Tunnel: GigaVUE-HC1 in Reno
- Configure Receiving End of Tunnel: GigaVUE-HC3 with GigaSMART in San Francisco

Configure Both Ends of GigaSMART Tunnel

Creating a GigaSMART tunnel requires configuration on both the sending and receiving ends:

Sending End of Tunnel	Receiving End of Tunnel
The sending end of a GigaSMART tunnel can be either a GigaVUE-VM deployment or a GigaSMART-enabled GigaVUE H Series node. **Sending Data from a GigaSMART-Enabled GigaVUE H Series Node** o Configure an IP interface with an IP address, subnet mask, default gateway, MTU setting and assign it to a GigaSMART group. o Create a GigaSMART operation with a tunnel-encap component. The encapsulation settings include the IP address and listening UDP port of the P interface that is associated with a network port on the destination GigaVUE H Series. o Bind the GigaSMART operation to one or more network ports as part of a map. The network ports must be mapped to the IP interface associated with a tool port. **Sending Data from GigaVUE-VM/GigaVUE-FM** When you provision a vMap for a GigaVUE-VM node in GigaVUE-FM, in addition to selecting the virtual traffic to be forwarded, you also specify the destination to which traffic should be tunneled with the following settings: o **UDP IP** – The IP address of the P interface that is associated with a network port on the receiving end of the tunnel. o **UDP Source Port** – The source port from which traffic will be sent to the receiving end of the GigaSMART tunnel. o **UDP Destination Port** – The listening UDP port at the destination end of the GigaSMART tunnel. **Sending Data from GigaVUE-2404/GigaSMART-6X** o Configure an IP interface with an IP address, subnet mask, default gateway, and MTU setting. Associate the IP interface with a tool port. o Create a GigaSMART operation with an encapsulation component. The encapsulation settings include the IP address and listening UDP port of the IP interface thst is associated with network port on the destination device. o Bind the GigaSMART operation to one or	o Configure an IP interface with an IP address, subnet mask, and default gateway. The IP address must match the destination IP address specified at the sending end of the tunnel. o Create a GigaSMART operation with a decapsulation component. The decapsulation settings include the same listening UDP port you specified as the destination port at the sending end of the tunnel. o Bind the GigaSMART operation to the IP interface that is associated with a network port as part of a map that distributes arriving traffic to local tool ports for analysis with local tools.

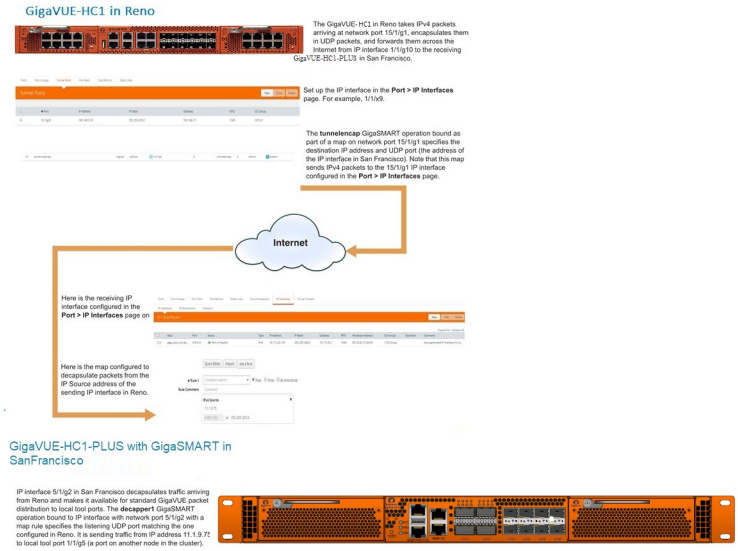
Sending End of Tunnel	Receiving End of Tunnel
more network ports as part of a map rule with at least one regular map rule criterion. The network ports must be mapped to the IP interface associated with a tool port.	

Keep in mind the following when configuring GigaSMART operations with encapsulation/decapsulation components:

Feature	Description
Viewing Statistics	Use the **show tunneled-port** commands to see statistics related to ongoing tunnel operations. Refer to View GigaSMART Statistics for more information.
Packet Order	Packer sequence is not preserved if the packets are reordered while traversing the Internet. The receiving GigaSMART delivers them in the same order received.
GMIP Header	The GMIP header is 46 bytes consisting of 14 Ethernet + 20 IP + 8 UDP + 4 tunnel version.
Tunnel Decap Type GMIP portdst	Use the GigaSMART Operations page to specify the UDP port on which the P interface that is associated with a network port on the receiving GigaVUE HC Series is listening. Use this option when decapsulating traffic from a either GigaSMART-enabled node or a GigaVUE-VM deployment. The setting must match the configuration of the portdst configured on the sending end of the tunnel.
GigaSMART Engine Ports	GigaSMART operations with a tunnel component can be assigned to GigaSMART groups consisting of multiple GigaSMART engine ports. Refer to Groups of GigaSMART Engine Ports for more information.

Example: GigaSMART Encapsulation/Decapsulation (GigaVUE-HC1 Node)

The following figures demonstrate how to create a sample IP tunnel between a sending GigaVUE-HC1 in Reno and a receiving GigaVUE H Series cluster in San Francisco. First, the overall tunnel is summarized, followed by configuration descriptions for the sending and receiving ends.



Configure Sending End of Tunnel: GigaVUE-HC1 in Reno

The GigaVUE-HC1 in this location has an IP interface configured on tool port 1/1/g1 with an IP address of 11.1.9.75. Maps to this port that use a tunnel encapsulation GigaSMART operation can send data over the Internet. The following table summarizes the commands necessary to configure the sending end of the tunnel in the GigaVUE-FM:

Task	UI Steps
Start by designating port 1/1/g1 as a tool port.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** In the Quick View Editor find port 1/1/g1. **4.** Set Type to Tool. **5.** Select Enable **6.** Click OK. **7.** Close the Quick Port Editor.
Use the **IP Interfaces** page to set up the network parameters for 1/1/g1. This page sets the IP address, subnet mask, default gateway, and MTU for the IP interface associated with a tool port on port 1/1/g1. Notice that the GigaSMART group in this example has the alias **gsport1**.	**1.** Select Ports > IP Interfaces. **2.** Click New. **3.** Configure the IP interface: ▪ Alias: 1_1_g1 ▪ Port: 1/1/g1 ▪ IP Address: **11.1.9.75** ▪ IP Mask: 255.255.255.0 ▪ Gateway: 11.1.9.1 ▪ MTU: 9400

Task	UI Steps
	▪ GigaSMART Group: gsport1 **4.** Click OK.
Now, create a tunnel encapsulation GigaSMART operation (**tunnelencap**) that will send traffic to IP address 21.2.9.75 on destination UDP port 10000 from source port 5000. The operation has the alias **tunnelenc**.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **2.** Click New. **3.** Configure the GigaSMART Operation: ▪ Alias: tunnelenc ▪ GigaSMART Groups: gsport1 ▪ GigaSMART Operations (GSOP): Tunnel Encapsulation **4.** Configure Tunnel Encapsulation: ▪ GMIP ▪ Destination: IPv4 ▪ Port Source: 5000 ▪ Port Destination: 10000 ▪ Destination IP: 21.2.9.75 ▪ DSCP: 0 ▪ Precision: 1 ▪ TTL: 64 **5.** Click Save.
Once you have the tunnel encapsulation operation, you can include it as part of a map rule. This map rule matches IPv4 packets and sends them to 21.2.9.75:10000 (the socket specified by the GigaSMART operation named **tunnelencap** that you created in the previous step).	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. ▪ Alias: tunnelencap ▪ Type: Regular ▪ Subtype: By Rule ▪ Source: 1/1/x3 ▪ Destination; 1/1/x1 ▪ GigaSMART Operations (GSOP): tunnelencap (gsport1) **4.** Click Add Rule. **5.** Select Pass. **6.** Select IP Version for Rule 1. **7.** Select v4 Version. **8.** Save.

Configure Receiving End of Tunnel: GigaVUE-HC3 with GigaSMART in San Francisco

Now we need to configure the receiving end of the tunnel with an IP interface associated with network port. The GigaVUE-HC3 in this location will have an IP interface associated with network port configured on network port 5/1/x2 with an IP address of 21.2.9.75 and a GigaSMART decapsulation operation that listens on UDP port 10000.

The following table summarizes the steps necessary to configure the receiving end of the tunnel using the UI:

Task	UI Steps
Start by designating port 5/1/x2 as a network port.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** In the Quick View Editor find port 5/1/x2. **4.** Set Type to Network. **5.** Select Enable **6.** OK. **7.** Close the Quick Port Editor.
Use the **IP Interfaces** page to set up the network parameters for 5/1/x2. This command sets the IP address, subnet mask, default gateway, and MTU for the IP interface associated with network port on port 5/1/x2. Note that this port uses the same IP address to which the GSOP in Reno is configured to send data (21.2.9.75).	**1.** Select Ports > IP Interfaces. **2.** Click New. **3.** Configure the IP Interface: ▪ Alias: 1_1_x2 ▪ Port: 1/1/x2 ▪ IP Address: 21.2.9.75 ▪ IP Mask: 255.255.255.0 ▪ Gateway: 21.2.9.1 ▪ MTU: 9400 ▪ GigaSMART Group: gsport5 **4.** Save.
Now, create a tunnel decapsulation GigaSMART operation (**tunnel-decap**) that will decapsulate traffic received on UDP port 10000. Recall that we configured the sending end of the tunnel to send to that UDP port. The operation has the alias **hd-decap1**.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP). **2.** Click New. **3.** Configure the GigaSMART Operation: ▪ Alias: hd-decap1 ▪ GigaSMART Groups: gsport5 ▪ GigaSMART Operations (GSOP): Tunnel Decapsulation **4.** Configure the Tunnel Decapsulation. ▪ GMIP

Task	UI Steps
	▪ GMIP Port: 10000 **5.** Save.
Once you have your tunnel decapsulation operation, you can include it as part of a map rule. This map decapsulates all traffic arriving at 5/1/x2 from IP address 21.2.9.25 (the start of the tunnel) and sends it to port 1/1/x5. This is a tool port on the chassis with box ID 1 in this cluster.	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. ▪ Alias: decapper ▪ Type: Regular ▪ Subtype: By Rule ▪ Source: 5/1/x2 ▪ Destination: 1/1/x5 ▪ GigaSMART Operations (GSOP): hd-decap1 (gsport5) **4.** Click Add Rule. **5.** Select Pass. **6.** Select IPv4 Source for Rule 1. **7.** Set the IPv4 Address to 11.1.9.75 **8.** Set the Net Mask to 255.255.255.0 **9.** Save.

GigaSMART IP Encapsulation (GigaSMART Tunnel)

Required License for IP Encapsulation: Advanced Tunneling (GigaVUE-HC3), Tunneling (GigaVUE-HC1)

GigaSMART-enabled nodes with the Advanced Tunneling license installed can encapsulate traffic and send it through a GigaSMART tunnel to a destination GigaSMART-enabled node.

1. Configure an IP interface with an IP address, subnet mask, default gateway, and MTU setting and assign it to a GigaSMART group.
2. Create a GigaSMART operation with a **Tunnel Encapsulation** component. The encapsulation settings include the IP address and listening UDP port of the IP interface associated with network port on the destination GigaVUE H Series.
3. Bind the GigaSMART operation to one or more network ports as part of a map. The network ports must be mapped to the IP interface associated with a tool port.

Refer to the sections beginning with GigaSMART IP Encapsulation/Decapsulation (GigaSMART Tunnel) for examples of the end-to-end configuration of a GigaSMART tunnel.

GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation

Required License for L2GRE Decapsulation: Base (GigaVUE-HC3), Tunneling (GigaVUE-HC1)Required License for L2GRE Encapsulation: Advanced Tunneling (GigaVUE-HC3), Tunneling (GigaVUE-HC1)

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC3 Gen 2, GigaVUE-HC1-Plus, GigaVUE-HC1 Gen 3

GigaVUE-HC1 Gen 3 module supports IPv6 for L2GRE tunneling.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

Use GigaSMART Layer 2 (L2) Generic Routing Encapsulation (GRE) tunnel encapsulation to send traffic from one GigaSMART node over the Internet to a second GigaSMART node using L2GRE encapsulation. Use GigaSMART L2GRE tunnel decapsulation at the second GigaSMART node to decapsulate the traffic before sending it to local tool ports.

GigaSMART Layer 2 GRE tunnel encapsulation/decapsulation provides the following:

- L2GRE tunnel initiation and encapsulation on the tool port at the sending end of the tunnel (for example, at a remote site)
- L2GRE tunnel termination and decapsulation on the network port at the receiving end of the tunnel (for example, at a main office site)

The GigaSMART at the remote site encapsulates the filtered packets, adds an encapsulation header, and routes it to the main office site. The encapsulation protocol is GRE and the delivery protocol is IP or IPv6, so the encapsulation header consists of Ethernet + IP + GRE or Ethernet + IPv6 + GRE headers.

The parameters of the encapsulated header are user-configurable, such as the IPv4 address of the IP interface on the destination GigaSMART node and the GRE key that identifies the source of the tunnel.

At the remote end, packets are decapsulated, the L2GRE header is stripped off, and packets are sent to the specified tool port.

IP fragmentation and reassembly are supported. Refer to IP Fragmentation and Reassembly on L2GRE and GMIP Tunnels.

Figure 94L2GRE Tunnel Encapsulation/Decapsulation shows the remote site encapsulating the filtered traffic and routing it to the main office from the remote end.

The encapsulated packet is sent out of the tool port, which is connected to the public network (the Internet). This packet is routed in the public network to reach the main office site. It ingresses at the routed network port of the GigaVUE node at the main office.

The ingress encapsulated packet is then sent to the GigaSMART at the main office, where the packet is decapsulated and sent to the tool port. The received packet's destination IP is checked against the source IP/IPv6 configured for the network port. If they match, decapsulation is applied. The Ethernet + IP + GRE or Ethernet + IPv6 + GRE header is stripped and the remaining packet is sent to the tool port.

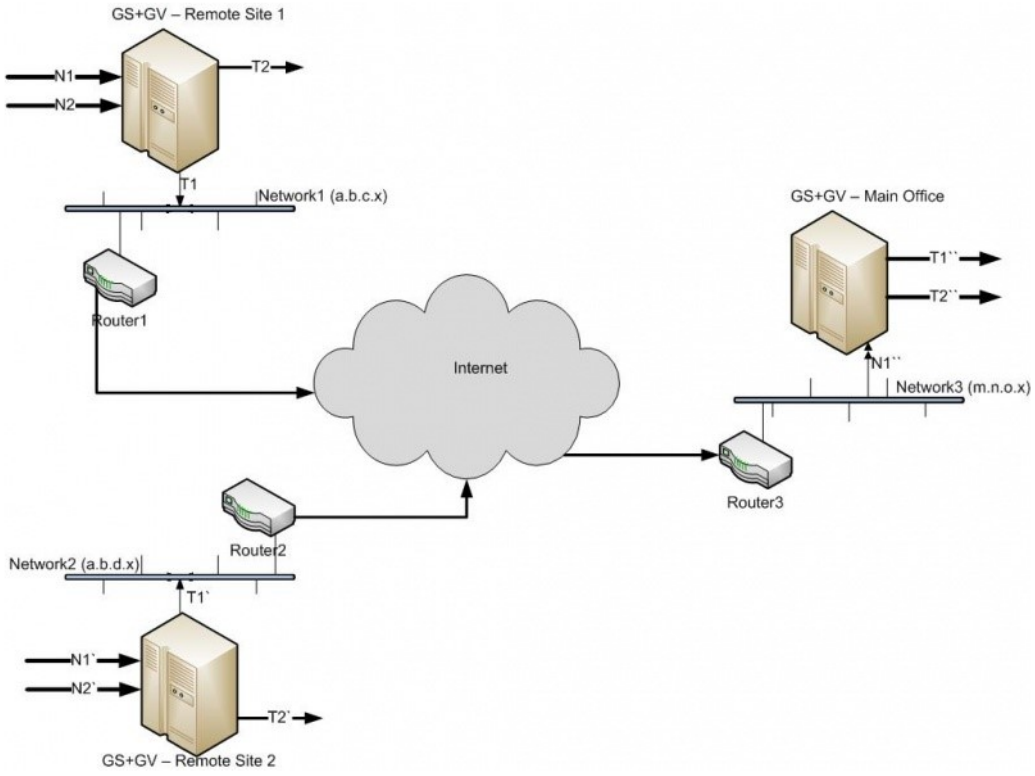


Figure 94 *L2GRE Tunnel Encapsulation/Decapsulation*

For L2GRE tunnel encapsulation/decapsulation configuration examples, refer to GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation and GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation.

For statistics for encapsulated packets, refer to Display L2GRE Tunnel Encapsulation Statistics. For statistics for decapsulated packets, refer to Display L2GRE Tunnel Decapsulation Statistics.

Layer 2 GRE Header Length

The L2GRE header length is as follows:

Header	Length in Bytes
With Key	42 bytes consisting of 14 Ethernet + 20 IP + 4 GRE + 4 GRE Key.
Without Key	38 bytes consisting of 14 Ethernet + 20 IP + 4 GRE.

Load Balancing to Multiple Destinations

Starting in software version 5.1, L2GRE tunnel encapsulation supports Load Balancing. Traffic from an IP Interface can be sent to multiple destinations Defined by IP address. The traffic is distributed using stateful Load Balancing or stateless hashing.

For information on stateful and stateless Load Balancing, refer to GigaSMART Load Balancing.

For examples of Load Balancing on L2GRE encapsulation, refer to GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation and GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation.

L2GRE IPv4 Encapsulation/Decapsulation

GigaSMART L2GRE IPv4 lets you route filtered traffic to the remote

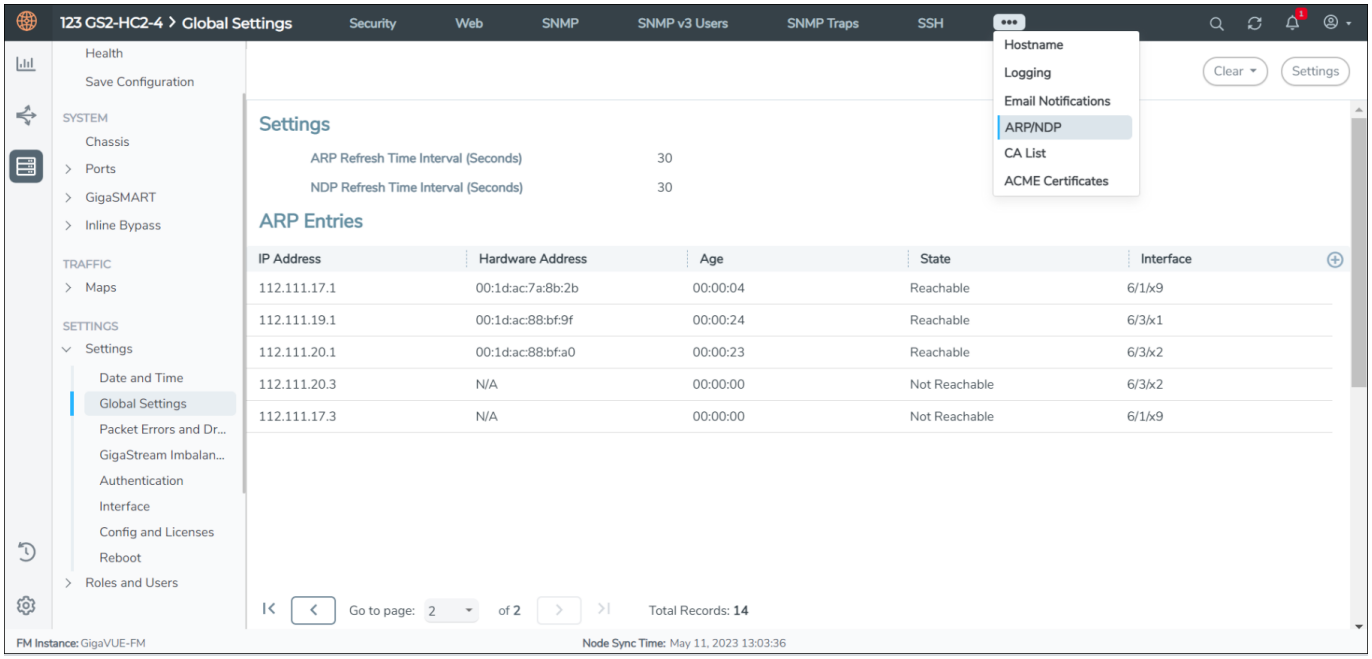
end using IPv4-based L2GRE tunneling. At the receiving end, filtered traffic is sent to GigaSMART, which adds an L2GRE header and a IPv4 header to make it routable.

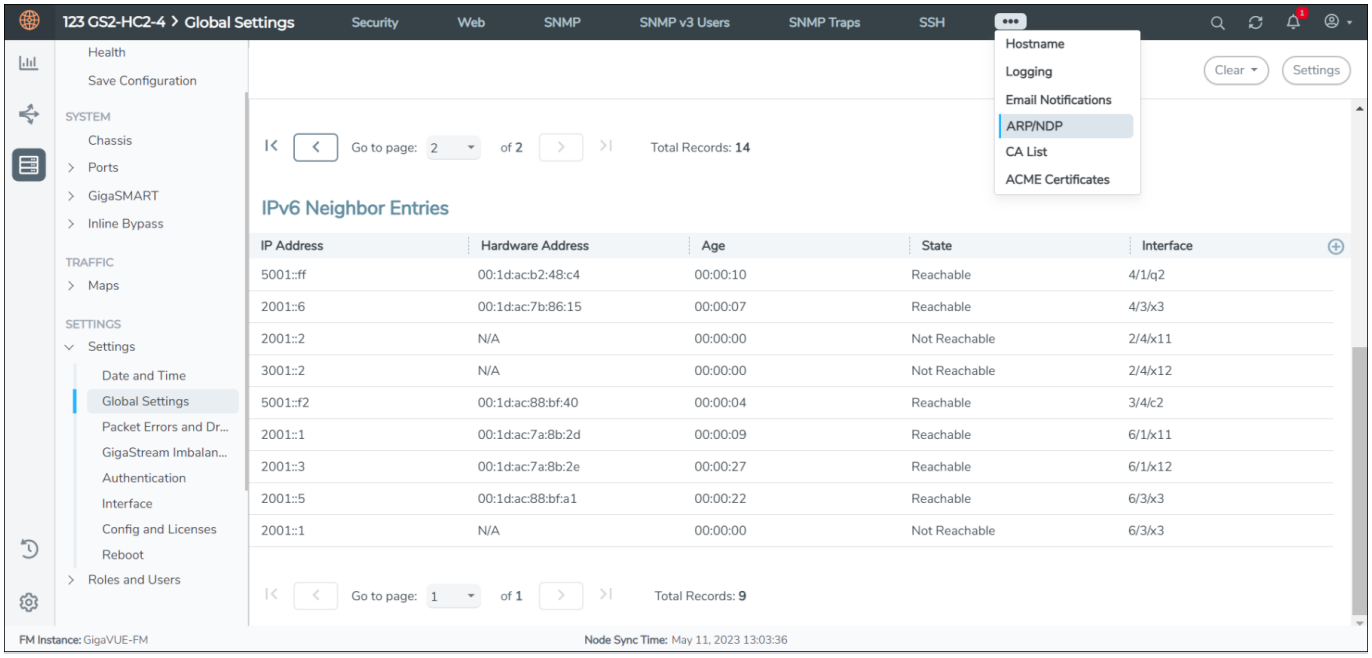
The remote end decapsulates the packet and sends it to the tool port.

GigaVUE nodes act as L2GRE encapsulation and decapsulation devices.

The IPv6 protocol is used to deliver all packets received in the encap tunnel to the termination node using the configured source and destination IPv4 address. The tunnel termination (decap) node strips the IPv4 + GRE header and sends the payload to the tool port.

The arp protocol is used by the tool port on the encapsulation node for arp request and arp response messages to resolve the gateway MAC address as well as destination tool MAC address in local network and respond to arp request messages received from the gateway in the tunnel decapsulation/termination node. arp echo request/reply messages are also sent and received.





For a configuration example, refer to GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation

L2GRE IPv6 Encapsulation/Decapsulation

GigaSMART L2GRE IPv6 lets you route filtered traffic to the remote

end using IPv6-based L2GRE tunneling. At the receiving end, filtered traffic is sent to GigaSMART, which adds an L2GRE header and a IPv6 header to make it routable.

The remote end decapsulates the packet and sends it to the tool port.

GigaVUE nodes act as L2GRE encapsulation and decapsulation devices.

The IPv6 protocol is used to deliver all packets received in the encap tunnel to the termination node using the configured source and destination IPv6 address. The tunnel termination (decapsulation) node strips the IPv6 + GRE header and sends the payload to the tool port.

The ICMPv6 protocol is used by the tool port on the encapsulation node for Neighbor Solicitation (NS) and Neighbor Advertisement (NA) messages to resolve the gateway MAC address as well as destination tool MAC address in local network and respond to NS messages received from the gateway in the tunnel decapsulation/termination node. ICMPv6 echo request/reply messages are also sent and received.

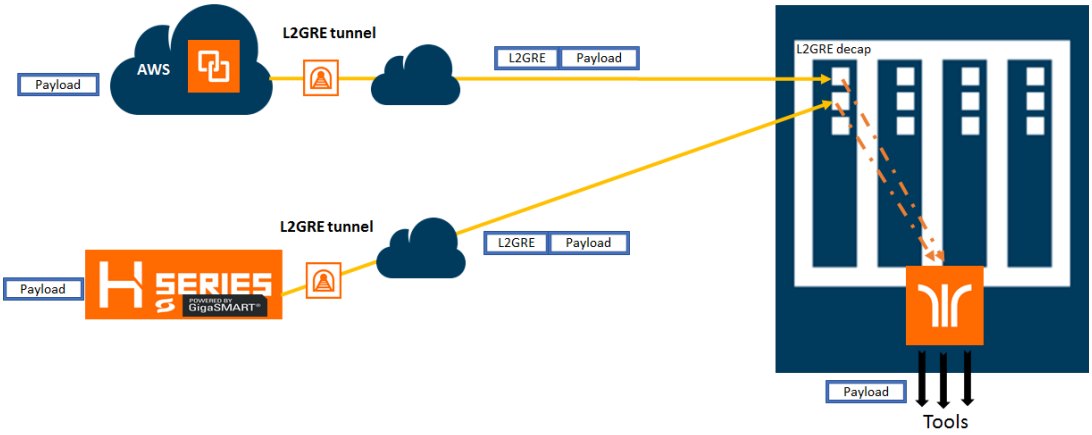
For a configuration example, refer to GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation

L2GRE Tunnel Termination

L2GRE tunnel termination is supported on physical devices, and the decapsulation happens through the GigaSMART engine. Tunneled traffic coming in the chassis is sent to the GigaSMART engine, which is sent to the tools using a hybrid port. The maps created are then applied to this decapsulated traffic.

Starting with version 5.4, tunnel termination is supported for VXLAN and L2GRE tunnel in the front panel ports of the switch. This feature provides line rate tunneling on all faceplate ports and also allows Flow Mapping® to be applied for the incoming tunneled traffic on the same ports.

The following diagram illustrates how the traffic from two sources—a GigaVUE V Series appliance running on an AWS platform and a GigaVUE HC Series device at a remote site traverses through the L2GRE tunnel and reaches the GigaVUE HC Series node in the main office site. In each case, traffic is tapped at the remote source and is then tunneled through L2GRE encapsulation across the cloud before it reaches the GigaVUE HC Series device at the main office site, which is connected to the actual tools. The L2GRE tunnel termination is executed on an ingress circuit port (IP interface) on the destination GigaVUE HC Series device. After tunnel termination, the packet is presented to the Flow Mapping® module to filter based on map rule parameters.



Configure L2GRE Tunnel Encapsulation and Decapsulation

Refer to the following configuration examples:

- GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation
- GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation
- GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation
- GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation
- GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation

Display L2GRE IPV4/IPv6 Tunnel Statistics

To view IP Interfaces statistics, **select Ports > IP Interfaces > Statistics** to open the IP Interfaces Statistics page.

The IP destinations pane displays the gateway or local tool status as **UP** if neighbor discovery is completed with gateway or local tool or **down** if neighbor discovery failed. Neighbor discovery is made only on the encapsulation node. On the decapsulation node, the gateway or local tool status will be **Not Applicable**.

Display L2GRE Tunnel Encapsulation Statistics

To display Layer 2 GRE tunnel encapsulation statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics**. The statistics for tunnel encapsulation will be in the row labeled Tunnel Encap in the GS Operations column.

Refer to Tunnel Encapsulation Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

Display L2GRE Tunnel Decapsulation Statistics

To display Layer 2 GRE tunnel decapsulation statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics** and click on the GS Operation in table to open the Quick View for GS Operation Statistics.

Refer to Tunnel Decapsulation Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

Configure GigaSMART Operation for Layer 2 GRE

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

To configure the L2GRE encapsulation/decapsulation types and options, use the GigaSMART Operations (GSOP) page:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.
2. Click **New**.
3. On the GigaSMART Operations page, do the following:
   a. Type an alias in the **Alias** field.
   b. In the **GigaSMART Groups** field, select the gsgroup for this operation.
   c. In the **GigaSMART Operations (GSOP)** field, select either **Tunnel Decapsulation** or **Tunnel Encapsulation** from the drop-down list, depending on whether you want decapsulation or encapsulation.
   d. Select **L2GRE**, and then enter options in the fields that display.
4. Click **Save**.

Example 1 – GigaSMART L2GRE Tunnel Encapsulation

In this example, an IP interface is configured on the tool port. A GigaSMART operation for tunnel encapsulation is configured to encapsulate the filtered packets. A map is configured that uses the L2GRE tunnel encapsulation GigaSMART operation, which sends packets from the remote site over the Internet to the main office using the IP interface associated with a tool port. Staring with software version 5.4 GigaSMART L2GRE Tunnel Encapsulation provides support for IPv6 with load-balancing.

Task	Description	UI Steps
1.	Configure a tool type of port and a network type of port.	**a.** Select Ports > All Ports. **b.** Click Quick Port Editor. **c.** Use Quick search to find the ports to configure. **d.** Set the type (Network or Tool) for each port and select Enable.
2.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**a.** From the device view, select GigaSMART > GigaSMART Groups. **b.** Click New. **c.** Type an alias in the Alias field and enter an engine port in the Port List field. **d.** Save.
3.	Configure the IP interface with an IP address, subnet mask, default gateway, and MTU setting. Assign it to the GigaSMART group.	**a.** Select Ports > Ports > IP Interfaces. **b.** Click New. **c.** On the IP Interfaces page, in the Alias and Description fields, enter a name and description for the IP interface. **d.** Click the Ports field and select the network or tool port from the drop-down list. **e.** Select Type: IPv4 or IPv6 **f.** Enter the IP address, subnet mask, gateway, and MTU settings in the respective fields. **g.** Click on the GigaSMART Group field to select the GigaSMART group. **h.** Save.
4.	Configure the GigaSMART operation for tunnel encapsulation and assign it to the GigaSMART group. The tunnel encapsulation settings include the IP address (IPv4) of the IP interface on the destination	**a.** From the device view, select GigaSMART > GigaSMART Operations > GigaSMART Operation. **b.** Click New. **c.** Type an alias in the Alias field. **d.** From the GigaSMART Groups drop-down list,

Task	Description	UI Steps
	GigaSMART node and the GRE key that identifies the source of the tunnel.	select the GigaSMART Group that you created in the second task. **e.** From the GigaSMART Operations (GSOP) drop-down list select Tunnel Encapsulation. **f.** Select L2GRE for the encapsulation type. **g.** Enter the IP address of the IP interface in the Destination IP field. **h.** IPv4, iPv6 or Port Group **i.** Enter the key parameter in the Key field. **j.** Save.
5.	Create a map using the tunnel encapsulation GigaSMART operation, with packets coming from the network port and being sent to the Internet through the tool port.	**a.** Select Maps > Maps. **b.** Click New. **c.** Type an alias in the Map Alias field that will help you identify this map. **d.** Select Regular for Type and By Rule for Subtype. **e.** Specify the network and tool ports that you configured in task one in the Source and Destination fields, respectively. **f.** From the GigaSMART Operations (GSOP) drop-down list, select the GigaSMART operation configured in task 4. **g.** Click Add a Rule under Map Rules and create the following rule: Select IP Version from the drop-down list and select v4 for Version, and then select Pass. **h.** Click Save.

Example 2 – GigaSMART L2GRE Tunnel Encap Stateful LB

Example 2 configures stateful load balancing of tunnel traffic to tunnel endpoints based on a metric. Each tunnel endpoint is assigned a weight.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

To configure this example:

1. Go to **Ports > Ports > All Ports**. Make sure you have one tool type of port and one network type of port enabled. Also make sure you have a GigaSMART port (eport).
2. Go to **GigaSMART > GigaSMART Groups**.
3. Click **New**, and then configure an Alias for the GigaSMART group and associate it with a GigaSMART engine port.
4. Click **OK**.
5. Go to **Ports > Ports > IP Interfaces**.
6. Click **New**, and then in the Alias and Description fields, enter the alias and description of the IP interface.
7. Select a port and configure it with an IP version Type, IP Address, IP Mask, Gateway, and MTU. Assign the IP interface to the GigaSMART group.
8. Click **OK**.
9. Go to **Ports > Tunnel Endpoints**.
10. Click **New**, then configure one or more tunnel endpoint IDs and their IP Addresses. The Alias is optional.
11. Click **OK**.
12. Go to **Ports > Port Groups**.
13. Click **New**, then type an alias for the port group, select GigaSMART Load Balancing, select the previously configured tunnel endpoints. Optionally, you can specify weights for each tunnel endpoint in the port group.
14. Click **OK**.
15. Go to **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.
16. Click **New**, then select the same GigaSMART Group and Tunnel Encapsulation for the GSOP. Under Load Balancing, select Stateful, the type Tunnel, and the metric, such as Round Robin. Under Tunnel Encapsulation, select L2GRE, Destination Port Group, select the Port Group, Session Field, and Session Position. Refer to Figure 95New GigaSMART Operation for Stateful Load Balancing.

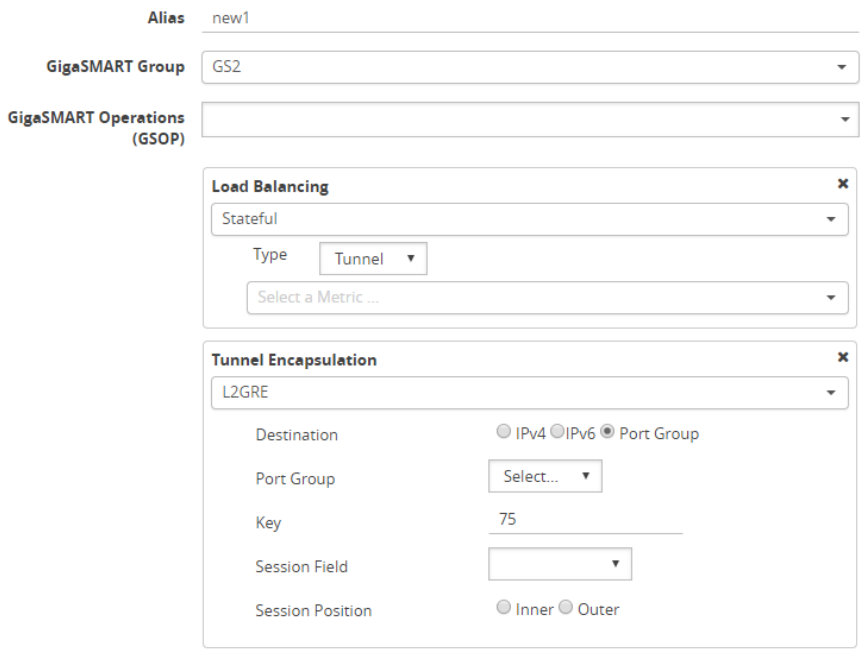


Figure 95 *New GigaSMART Operation for Stateful Load Balancing*

17. Click **OK**.
18. Go to **Maps > Maps**.
19. Click **New**, then type an alias for the map, select type Regular and subtype ByRule. Under Map Source and Destination, select a network port as the Source and a tool port as the Destination, then select the GigaSMART operation. Under Map Rules, configure a map rule. Refer to Figure 96New Map Configuration.

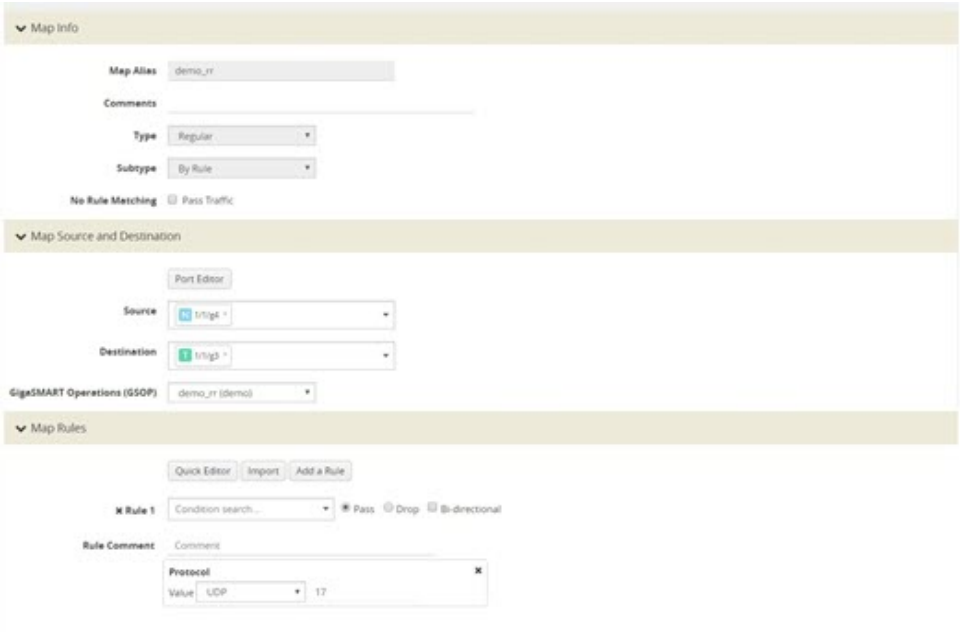


Figure 96 *New Map Configuration*

20. Click **OK**.

Example 3 – GigaSMART L2GRE Tunnel Encap Stateless LB

Example 3 configures stateless load balancing of tunnel traffic to tunnel endpoints based on a hash value.

Example 3 has the same configuration steps as Example 2 except for the GigaSMART operation (gsop) in Step 16. Under Load Balancing, select Stateless and the metric, such as Five Tuple. Under Tunnel Encapsulation, select L2GRE, Destination Port Group, and select the Port Group. Refer to Figure 97New GigaSMART Operation for Stateless.

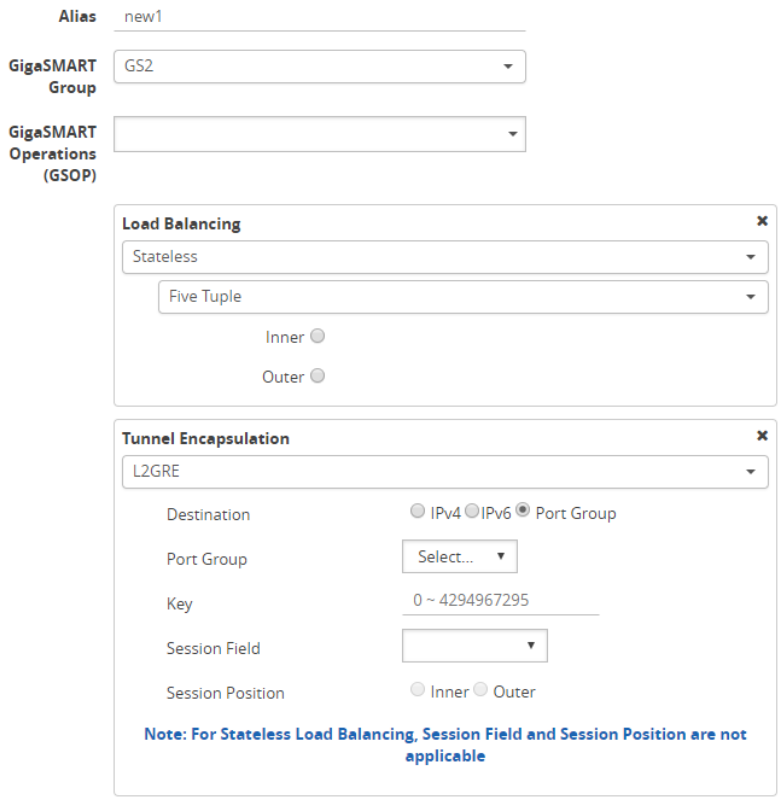


Figure 97 *New GigaSMART Operation for Stateless*

Example 4 – GigaSMART L2GRE Tunnel Decapsulation

In this example, an IP interface is configured on the network port. A GigaSMART operation for tunnel decapsulation is configured to decapsulate the filtered packets. A map is configured that uses the L2GRE tunnel decapsulation GigaSMART operation, which receives packets from the remote site over the Internet to the main office using the IP interface associated with a tool port and then forwards packets over the tool port. Staring with software version 5.4 GigaSMART L2GRE Tunnel Decapsulation provides support for IPv6 with load-balancing.

Task	Description	UI Steps
1.	Configure a network type of port and a tool type of port.	**a.** Select Ports > Ports > All Ports. **b.** Click Quick Port Editor. **c.** Use Quick search to find the ports to configure.

Task	Description	UI Steps
		d. Set the type for each port and select Enable.
2.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**a.** From the device view, select GigaSMART > GigaSMART Groups. **b.** Click New. **c.** Type an alias in the Alias field and enter an engine port in the Port List field. **d.** Click Save.
3.	Configure the IP interface with an IP address, subnet mask, default gateway, and MTU setting. Assign it to the GigaSMART group. The IP address must match the destination IP address specified at the sending end of the tunnel.	**a.** Select Ports > IP Interfaces. **b.** Click New. **c.** On the IP Interfaces page, in the Alias and Description fields, enter the name and description for the IP interface. **d.** Click the Ports field and select the port from the drop-down list. **e.** Enter the IP address, subnet mask, gateway, and MTU settings in the respective fields. **f.** Click on the GigaSMART Group field to select the GigaSMART group. **g.** Click Save.
4.	Configure the GigaSMART operation for tunnel decapsulation and assign it to the GigaSMART group. The tunnel decapsulation settings include the GRE key that identifies the source of the tunnel.	**a.** From the device view, select GigaSMART > GigaSMART Operations > GigaSMART Operation. **b.** Click New. **c.** Type an alias in the Alias field. **d.** From the GigaSMART Groups drop-down list, select the GigaSMART Group that you created in the second task. **e.** From the GigaSMART Operations (GSOP) drop-down list select Tunnel Decapsulation. **f.** Select L2GRE for the decapsulation type. **g.** Enter the GRE key in the Key field. **h.** Click Save.
5.	Create a map using the tunnel decapsulation GigaSMART operation, with packets coming from the Internet through the network port and being sent to the local tool port.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type an alias in the Map Alias field that will help you identify this map. **d.** Select Regular and By Rule for the map type and subtype. **e.** Specify the network and tool ports that you configured in task one in the Source and Destination fields, respectively.

Task	Description	UI Steps
		f. From the GSOP drop-down list, select the GigaSMART operation configured in task 4. **g.** Click Add a rule under Map Rules and create the following rule: Select IP Version from the drop-down list and select v4 for Version, and then select Pass. **h.** Click Save.

Example 5 – GigaSMART L2GRE IPv6 Tunnel Encamp/Decap with Load-Balancing

In this example, the encapsulation and decapsulation nodes are configured with IP interfaces using IPv6 addresses and load-balancing. IPv6 tunnel load-balancing feature supports the distribution of traffic across multiple IPv6 tunnel destination through the same tool port. Two types of load-balancing is supported, stateful and stateless.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

Step 1: Configure a tool type of port and a network type of port.

Create a map using the tunnel encapsulation GigaSMART operation, with packets coming from the network port and being sent to the Internet through the tool port.

1. Select **Ports > Ports > All Ports**.
2. Click **Quick Port Editor**.
3. Use **Quick search** to find the ports to configure.
4. Set the **type for each port** and select **Enable**.
   a. type: tool - port 1/3/x7
   b. type: network - 1/3/x8

Step 2: Configure a GigaSMART group and associate it with a GigaSMART engine port.

1. From the device view, select **GigaSMART > GigaSMART Groups.**
2. Click **New**.
3. Type an **alias** in the Alias field and enter an **engine port** in the Port List field.
4. Click **OK**.

Step 3: Configure the IP Interface

1. Select **Ports > IP Interfaces**.
2. Click **New**.
3. In the **Alias** and **Description** fields, enter the name and description for the IP interface.
4. Click the **Ports field** and select the port from the drop-down list.
5. Select the Port address: **IPv6**
6. Enter the **IP address, subnet mask, gateway, and MTU** settings in the respective fields.
7. From the **GS Group** drop-down list, select the required GigaSMART group.
8. Click **OK.**

Step 4: Configure the GigaSMART operation for tunnel encapsulation and load balancing and assign it to the GigaSMART group

1. From the device view, select **GigaSMART > GigaSMART Operations > GigaSMART Operation.**
2. Click **New**.

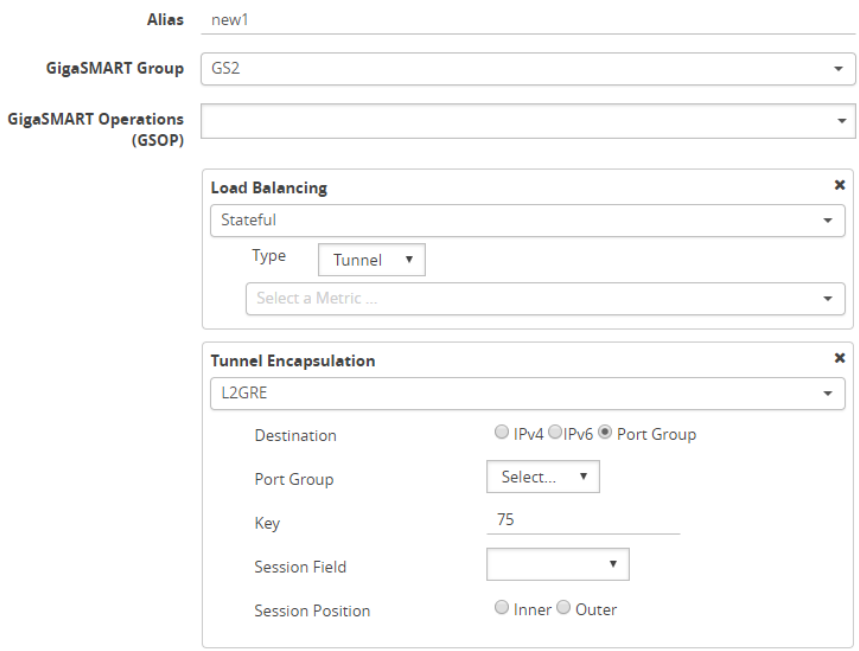


Figure 98 *GigaSMART Operation (GSOP)*

3. In the **Alias** field, enter a name for the GigaSMART operation.
4. From the **GigaSMART Group** drop-down list, select the **GigaSMARTGroup** that you created in the step 2.
5. From the **GigaSMART Operations (GSOP)** drop-down list, select **Tunnel Encapsulation**.
6. Select **L2GRE** for the encapsulation type.
7. Enter the **GRE key 123214** in the Key field.
8. Click **OK**.

Step 5: Create a map using the tunnel encapsulation GigaSMART operation

1. Select **Maps > Maps > Maps.**
2. Click **New**.

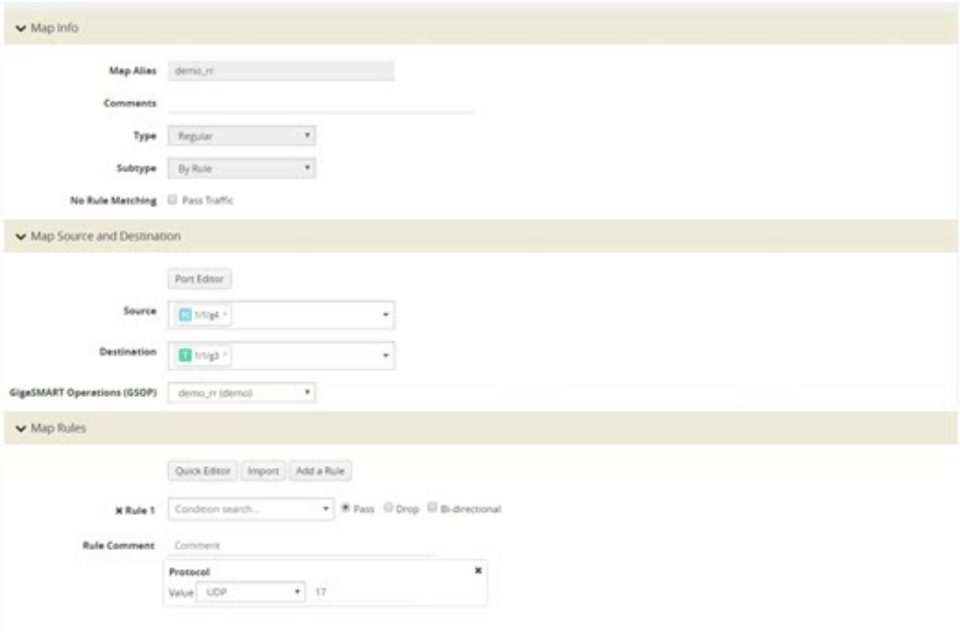


Figure 99 *Map Configuration*

3. Type an **alias** in the Map Alias field that will help you identify this map.
4. Select **Regular** and **By Rule** for the map type and subtype.
5. Specify the **network and tool ports** that you configured in step one in the Source and Destination fields, respectively.
6. From the GSOP drop-down list, select the **GigaSMART operation** configured in step 4.
7. Click **Add** a rule under Map Rules.
8. Select **IP Version** from the drop-down list and select **v4** for Version.
9. Select **Pass**.
10. Click **Add a rule** under Map Rules and create the following rule:
11. Select **IP Version** from the drop-down list and select **v6** for Version.
12. Select **Pass**.
    a. Source: **From - 1/4x24**
    b. Destination: **To: 1/4x7**
13. Click **OK**.

On the decapsulation node, configure the receiving end of the tunnel

Step 6: Configure a tool type of port and a network type of port.

1. Select **Ports > Ports > All Ports**.
2. Click **Quick Port Editor**.
3. Use **Quick search** to find the ports to configure.
4. Set the **type for each port** and select **Enable**.
   a. type: tool - port 1/4/x7
   b. type: network - 1/4/x24

Step 7: Configure a GigaSMART group and associate it with a GigaSMART engine port.

1. From the device view, select **GigaSMART > GigaSMART Groups.**
2. Click **New**.
3. Type an **alias** in the Alias field and enter an **engine port** in the Port List field.
   a. engine port: **1/3/e1**
4. Click **OK**.

Step 8: Configure the IP Interface with an IPv6 address, prefix length, default gateway, and MTU setting. Assign it to the GigaSMART group.

1. Select **Ports > IP Interfaces**.
2. Click **New**.
3. In the **Alias** and **Description** fields, enter the name and description for the IP interface.
4. Click the **Ports field** and select the port from the drop-down list.
5. Enter the **IP address**, **subnet mask, gateway, and MTU settings** in the respective fields.
6. From the **GS Group** drop-down list, select the required GigaSMART group.
7. Click **OK**.

Step 9: Configure the GigaSMART operation for tunnel decapsulation and assign it to the GigaSMART group.

1. From the device view, select **GigaSMART > GigaSMART Operations > GigaSMART Operation.**

2. Click **New**.
3. In the **Alias** field, enter a name for the GigaSMART operation.
4. From the **GigaSMART Groups** drop-down list, select the GigaSMART Group.
5. From the **GigaSMART Operations (GSOP)** drop-down list, select **Tunnel Decapsulation**.
6. Select **L2GRE** for the decapsulation type.
7. Enter the **GRE** key in the Key field.
8. Click **OK**.

Step 10: Create a map using the tunnel decapsulation GigaSMART operation.

1. Select **Maps > Maps > Maps**.
2. Click **New**.
3. Type an **alias** in the Map Alias field that will help you identify this map.
4. Select **Regular and By Rule** for the map type and subtype.
5. Specify the **network and tool ports** that you configured in task one in the Source and Destination fields, respectively.
6. From the **GSOP** drop-down list, select the GigaSMART operation.
7. Click **Add a rule** under Map Rules and create the following rule:
8. Select **IP Version** from the drop-down list and select **v4** for Version.
9. Select **Pass**.
10. Click **Add a rule** under Map Rules and create the following rule:
11. Select **IP Version** from the drop-down list and select **v6** for Version.
12. Select **Pass**.
    a. Source: **From - 1/4x24**
    b. Destination: **To: 1/4x7**
13. Click **OK**.

Orchestrated Workflow Configuration of GigaSMART L2GRE Tunnels

The GigaSMART L2GRE tunnels can now be configured through the Orchestrated Configuration page. To configure GigaSMART L2GRE Tunnels follow the below steps:

- In GigaVUE-FM go to **Traffic>Physical> Orchestrated Flows> Tunnels**.
- Click on **New**.
- Select **GigaSMART L2GRE**. The New Tunnel configuration page appears. Enter the required information as described below:

Field	Description
Tunnel Name	The name of the tunnel. **Note**: Alias must not have spaces or any of these characters: *"?;,/%@.
Tunnel Description	The description of the tunnel endpoint.
Traffic Direction	Select the traffic direction of the tunnel. Choose DECAP for decapsulation or choose ENCAP for encapsulation. • If you choose Decap, select the IP Interface from the list. You can create one using the Create an IP interface option if you do not have an IP interface. Refer to IP Interfaces to learn more about creating an IP Interface. • If you choose Encap enter the Source Port. You can select from the drop-down list. To edit the Port type or to enable the port for Admin privileges use the Port Editor window.
Nodes	Select the Nodes on which you intend to create the tunnel.
GigaSMART Group	Select the GigaSMART groups created. If you do not have any groups, you can configure them using the **Create GS Groups** option. Refer Create GigaSMART Operations – A Summary to know more about configure groups.
Key	Select the key to be used for Decapsulation tunnel. You can choose a value from 0-4297964295. For Encapsulation Tunnel the key range is from 1-4297964295.
Rules	Configure the map rules that needs to be adhered to while decapsulating or encapsulating the traffic. Click on **Rules Editor** to configure the rules. **NOTE:** Configuring rules are mandatory for Encapsulation and Decapsulation GigaSMART tunnels.
Encapsulation IP interface	Enter the interface IP address of the node (Destination IP) for an encapsulation tunnel.
Time to Live	For encapsulation tunnel, enter the value of the time interval till which the session needs to be available. The value ranges from 1 to 255.

Field	Description
DSCP	Enter the Differentiated Services Code Point (DSCP) value for an encapsulation.
Precedence	Add a value between 0-7 for an encapsulation GigaSMART L2GRE tunnel. This field is applicable only for Gen 2 cards.
Remote Tunnel IP	Enter the Remote Tunnel IP values for and encapsulation tunnel. **NOTE:** If the Tunnel is load balanced ,Remote Tunnel Port group must be configured.
Destination Port	Select the destination port for the decapsulation tunnel. You can edit the port details from the Port editor screen. **NOTE:** If the Tunnel is load balanced ,Destination Port group must be configured.
Additional GigaSMART operations	Enable this option to add GigaSMART operations. When enabled you will be provided with GigaSMART Operations section which would allow you to choose among the following: • Add Header • Add Trailer • De-duplication • Load Balancing • Masking • Remove Trailer • Slicing

The configured tunnels provide you a **Details View** and **Troubleshoot View**. Click on the tunnel profile to view the below:

- **Details View**: Displays the configured parameters of the configured tunnel.
- **Troubleshoot View**: Displays the tunnel's statistics. Use the **Clear Stats** button to refresh the statistics.

GigaSMART VXLAN Tunnel Encapsulation

Required License for VXLAN Encapsulation: Base (GigaVUE-HC1 Gen 3), Tunneling (GigaVUE-HC3 Gen 3)

Supported Devices : GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

Starting in software version 5.11, support for VXLAN tunnel origination is added to GigaSMART. After receiving packets on the network ingress ports, they are sent to the GSOP chain. When a packet egresses out of the GS card, it is encapsulated with a VXLAN header, along with all outer headers (UDP, IP, and MAC), and then sent out to the appropriate multicast ID.

Use GigaSMART VxLAN tunnel encapsulation to send traffic from one GigaSMART node over the Internet to a second GigaSMART node using VxLAN encapsulation.

GigaSMART VxLAN tunnel encapsulation provides the VxLAN tunnel initiation and encapsulation on the tool port at the sending end of the tunnel (for example, at a remote site)

The GigaSMART at the remote site encapsulates the filtered packets, adds an encapsulation header, and routes it to the main office site. The encapsulation protocol is VxLAN and the delivery protocol is IP or IPv6, so the encapsulation header consists of Ethernet + IP + VxLAN or Ethernet + IPv6 + VxLAN headers.

The encapsulated packet is sent out of the tool port, which is connected to the public network (the Internet). This packet is routed in the public network to reach the main office site. It ingresses at the routed network port of the GigaVUE node at the main office.

For statistics for encapsulated packets, refer to Display VXLAN Tunnel Encapsulation Statistics.

VXLAN Header Length

The VXLAN header length is as follows:

Header	Length in Bytes
With Key	50 bytes consisting of 14 Ethernet + 20 IP + 8 UDP + 8 VxLAN Key.

VXLAN IPv4 Encapsulation/Decapsulation

GigaSMART VXLAN IPv4 lets you route filtered traffic to the remote

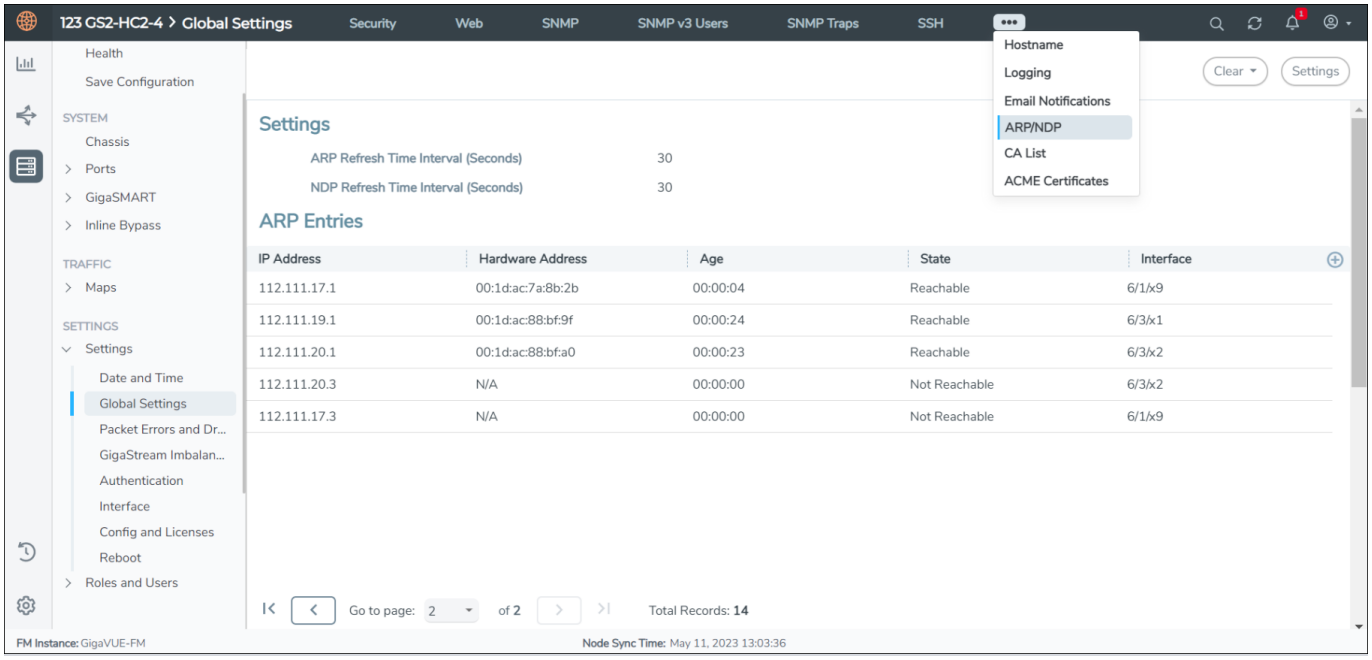
end using IPv4-based VXLAN tunneling. At the receiving end, filtered traffic is sent to GigaSMART, which adds an VXLAN header and a IPv4 header to make it routable.

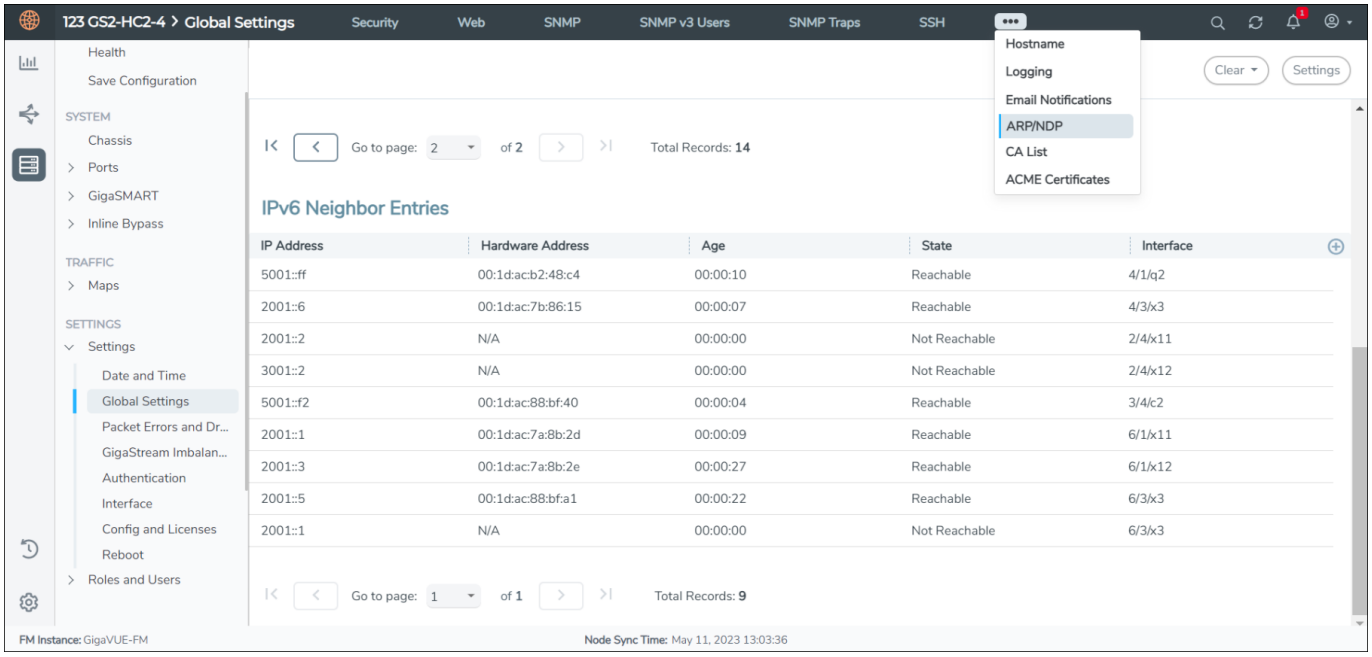
The remote end decapsulates the packet and sends it to the tool port.

GigaVUE nodes act as VXLAN encapsulation and decapsulation devices.

The IPv4 protocol is used to deliver all packets received in the encap tunnel to the termination node using the configured source and destination IPv4 address. The tunnel termination (decap) node strips the IPv4 + VXLAN header and sends the payload to the tool port.

The arp protocol is used by the tool port on the encapsulation node for arp request and arp response messages to resolve the gateway MAC address as well as destination tool MAC address in local network and respond to arp request messages received from the gateway in the tunnel decapsulation/termination node. arp echo request/reply messages are also sent and received.





VXLAN IPv6 Encapsulation/Decapsulation

GigaSMART VxLAN IPv6 lets you route filtered traffic to the remote

end using IPv6-based VxLAN tunneling. At the receiving end, filtered traffic is sent to GigaSMART, which adds an VXLAN header and a IPv6 header to make it routable.

The remote end decapsulates the packet and sends it to the tool port.

GigaVUE nodes act as VXLAN encapsulation and decapsulation devices.

The IPv6 protocol is used to deliver all packets received in the encap tunnel to the termination node using the configured source and destination IPv6 address. The tunnel termination (decap) node strips the IPv6 + VXLAN header and sends the payload to the tool port.

The ICMPv6 protocol is used by the tool port on the encapsulation node for Neighbor Solicitation (NS) and Neighbor Advertisement (NA) messages to resolve the gateway MAC address as well as destination tool MAC address in local network and respond to NS messages received from the gateway in the tunnel decapsulation/termination node. ICMPv6 echo request/reply messages are also sent and received.

Display VXLAN IPV4/IPv6 Tunnel Statistics

To view IP Interfaces statistics, **select Ports > IP Interfaces > Statistics** to open the IP Interfaces Statistics page.

The IP destinations pane displays the gateway or local tool status as **UP** if neighbor discovery is completed with gateway or local tool or **down** if neighbor discovery failed. Neighbor discovery is made only on the encapsulation node. On the decapsulation node, the gateway or local tool status will be **Not Applicable**.

Display VXLAN Tunnel Encapsulation Statistics

To display Layer 2 VXLAN tunnel encapsulation statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics**. The statistics for tunnel encapsulation will be in the row labeled Tunnel Encap in the GS Operations column.

Refer to Tunnel Encapsulation Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

Orchestrated Workflow Configuration of GigaSMART VXLAN Tunnels

The GigaSMART VXLAN tunnels can now be configured through the Orchestrated Configuration page. To configure GigaSMART VXLAN Tunnels follow the below steps:

- In GigaVUE-FM go to **Traffic>Physical> Orchestrated Flows> Tunnels**.
- Click on **New**.
- Select **GigaSMART VXLAN**. The New Tunnel configuration page appears. Enter the required information as described below:

Field	Description
Tunnel Name	The name of the tunnel. **NOTE:** Alias must not have spaces or any of these characters: *"?;,/%@.
Tunnel Description	The description of the tunnel endpoint.
Traffic Direction	Select the traffic direction of the tunnel. Choose DECAP for decapsulation or choose ENCAP for encapsulation. • If you choose Decap, select the IP Interface from the list. You can create one using the Create an IP interface option if you do not have an IP interface. Refer to IP Interfaces to learn more about creating an IP Interface. • If you choose Encap enter the Source Port. You can select from the drop-down list. To edit the Port type or to enable the port for Admin privileges

Field	Description
	use the Port Editor window. **NOTE:** For Encapsulation Tunnel, only Gen 3 nodes are supported.
Nodes	Select the Nodes on which you intend to create the tunnel.
GigaSMART Group	Select the GigaSMART groups created. If you do not have any groups, you can configure them using the **Create GS Groups** option. Refer Create GigaSMART Operations – A Summary to know more about configure groups.
Time to Live	or encapsulation tunnel, enter the value of the time interval till which the session needs to be available. The value ranges from 1 to 255
DSCP	Enter the Differentiated Services Code Point (DSCP) value for an encapsulation.
Rules	Configure the map rules that needs to be adhered to while decapsulating or encapsulating the traffic. Click on **Rules Editor** to configure the rules. **NOTE:** Configuring rules are mandatory for Encapsulation and Decapsulation GigaSMART tunnels.
Encapsulation IP interface	Enter the interface IP address of the node (Destination IP) for an encapsulation tunnel.
Remote Tunnel IP, L4 Port & Source Port	Enter the IP address, L4 Port and Source Port values for an encapsulation tunnel. **NOTE:** For Encapsulation Tunnels, the destination and L4 port can be configured only on three ports 4789, 8472 and 48879. **NOTE:** For Decapsulation Tunnel, the Destination and Source port range is from 1 - 65535.
Key	Select the key to be used for Decapsulation tunnel. You can choose a value from 0-4297964295. For Encapsulation Tunnel the key range is from 1-4297964295.
Destination Port	Select the destination port for the decapsulation tunnel. You can edit the port details from the Port editor screen.
Additional GigaSMART operations	Enable this option to add GigaSMART operations. When enabled you will be provided with GigaSMART Operations section which would allow you to choose among the following: • Add Header • Add Trailer • De-duplication • Load Balancing • Masking • Remove Trailer • Slicing

The configured tunnels provide you a **Details View** and **Troubleshoot View**. Click on the tunnel profile to view the below:

- **Details View**: Displays the configured parameters of the configured tunnel.
- **Troubleshoot View**: Displays the tunnel's statistics. Use the **Clear Stats** button to reset the statistics.

IP Fragmentation and Reassembly on L2GRE and GMIP Tunnels

Required Licenses:
IP Fragmentation on L2GRE and GMIP Tunnels-Advanced Tunneling license on GigaVUE-HC3. Referred to as “Tunneling license” on GigaVUE-HC1.
Reassembly on L2GRE and GMIP Tunnels -Base (GigaVUE-HC3, GigaVUE-HC1-Plus), Tunneling (GigaVUE-HC1).

Starting in software version 4.6, L2GRE and GMIP tunnels support IP fragmentation and reassembly of packets. IP fragmentation occurs with encapsulation. Fragmented packets are sent on the tool port at the sending end of the tunnel (for example, at a remote site). IP reassembly occurs with decapsulation. Fragmented packets reaching the network port at the receiving end of the tunnel (for example, at a main office site), are decapsulated and reassembled before being sent to the destination.

IP Fragmentation on Encapsulation

The tool port at the remote site is configured with a tunnel MTU. If a packet exceeds the tunnel MTU, the packet will be fragmented, and the fragmented packets will be sent out the IP interface.

NOTE: The first fragmented packet contains the tunnel header (Eth+IP+GRE). The rest of the fragments have the Ethernet and IP headers.

The packet size plus the tunnel header size is calculated and checked against the tunnel MTU. For example, if the tunnel MTU is 1518 and the packet is 1526, the packet exceeds the tunnel MTU. If the tunnel MTU is 1518 and the packet is 1518, the packet will also exceed the tunnel MTU due to the addition of the tunnel header.

IPv6 tunnel supports Path MTU. When a GigaSMART (GSOP) is associated to a map, Path MTU discovery message is sent to the tunnel destination and the Path MTU learnt is used for IPv6 fragmentation.

IP Reassembly on Decapsulation

The network port at the main office site receives the fragmented packets sent from the remote site. The tunnel header is removed from all fragmented packets, and they are buffered in memory. After all the fragmented packets are available, they are reassembled. The reassembled packet is then sent to the tool.

Notes and Considerations

Take into account the following notes and considerations:

Feature	Description
IPv4 and IPv6 Support	IPv4 and IPv6 packets are supported. **Note:** To avoid the overhead and improve the performance, existing GMIP IPv4 tunneling does not calculate the UDP Checksum on the Encapsulated Packets. The same will be adopted for IPv6 and IPv6 RFCs, where the checksum value will be set to 0. UDP checksum of the out header is not mandatory in the IPv6 tunneling.
Always Enabled	IP fragmentation and reassembly are always enabled. No configuration is required.
Tunnel MTU	The tunnel MTU is configured using the MTU field on the **IP Interfaces** configuration page. (Select **Ports > IP Interfaces**, and then click **New** to open the page.) The MTU is fixed at 9400 for all network/tool ports on the following platforms: o GigaVUE-HC1 o GigaVUE-HC3 o GigaVUE-TA100
Encapsulation Statistics	The encapsulation statistics count the number of fragmented packets. Refer to Display GMIP Tunnel Encapsulation Statistics and Display L2GRE Tunnel Encapsulation Statistics. For definitions, refer to Tunnel Encapsulation Statistics Definitions.
Decapsulation Statistics	The decapsulation statistics count the number of reassembled packets. Refer to Display GMIP Tunnel Decapsulation Statistics and Display L2GRE Tunnel Decapsulation Statistics. For definitions, refer to Tunnel Decapsulation Statistics Definitions.
GigaSMART Engine Ports	GigaSMART operations with a tunnel component can be assigned to GigaSMART groups consisting of multiple GigaSMART engine ports. Refer to Groups of GigaSMART Engine Ports for details.

GigaSMART ERSPAN Tunnel Decapsulation

Required License for ERSPAN Decapsulation: Advanced Tunneling (GigaVUE-HC1 Gen3, GigaVUE-HC3, GigaVUE-HC1-Plus, and GigaVUE-HCT Gen 3), Tunneling (GigaVUE-HC1 and GigaVUE-HCT Gen 3).

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus, and GigaVUE-HCT Gen 3.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

Some Cisco equipment provides the ability to mirror monitored traffic to a remote destination through an ERSPAN tunnel. Using ERSPAN tunnel decapsulation, GigaSMART can act as the receiving end of an ERSPAN tunnel, decapsulating mirrored traffic sent over the Internet from a Cisco switch or router.

Both ERSPAN Type II and Type III header decapsulation are supported. For ERSPAN Type III details, refer to ERSPAN Type III.

You can configure a GigaSMART-enabled node to act as the receiving end of an ERSPAN tunnel by configuring a GigaSMART **Tunnel Decapsulation** operation with type set to **ERSPAN** and a **Flow ID** matching the sending end of the tunnel.

The high-level steps are as follows:

1. Configure an IP interface associated with network port and assign an IP address, subnet mask, and default gateway to the IP interface. The IP address must match the destination IP address specified at the sending end of the tunnel.
2. Create a GigaSMART operation with an ERSPAN tunnel decapsulation component. The decapsulation settings include the same flow ID specified at the sending end of the tunnel. The flow ID is a value from 0 to 1023. Use this options when decapsulating traffic received over a Cisco-standard ERSPAN tunnel. A flow ID of 0 decapsulates all ERSPAN tunnel traffic regardless of flow ID.
3. For ERSPAN Type III, a trailer timestamp may be specified.
4. Bind the GigaSMART operation to the IP interface associated with network port as part of a map that distributes arriving traffic to local tool ports for analysis with local tools.

For example configurations, refer to ERSPAN Tunnel Header Removal and GigaSMART ERSPAN Tunnel Decapsulation.

For an example of APF and ERSPAN tunneling, refer to GigaSMART Adaptive Packet Filtering (APF).

ERSPAN Type III

ERSPAN Type III is similar to ERSPAN Type II but has a hardware timestamp in the packet. The hardware timestamp needs to be translated into a usable timestamp.

The UTC timestamp can be calculated, based on the reference hardware timestamp and the reference UTC timestamp carried in marker packets that are periodically sent over UDP. The calculated UTC timestamp can then be appended to the packets as a trailer.

Marker packets have a fixed length and are identified by a signature of 0xA5A5A5A5. If the marker packet session ID matches the ERSPAN session ID, the UTC timestamp can be extracted from the marker packet. An ERSPAN session is defined by a map that uses an ERSPAN GigaSMART operation (gsop).

There are three timestamp formats: **None**, **GigaSMART**, and **X12-TS** (for PRT-H00-X12TS). The timestamp options are set from the GigaSMART Group page, which is accessed by selecting **GigaSMART > GigaSMART Groups > GigaSMART Groups**, and then clicking **New** or editing an existing GigaSMART Group. Figure 100ERSPAN Type III Timestamp Formats on GigaSMART Groups Page shows the timestamp format options. If the timestamp format is **Disabled**, ERSPAN Type III packets are parsed and the ERSPAN header is removed by GigaSMART. The inner packets are forwarded to a tool port. If the timestamp format is **GigaSMART** or **X12-TS**, a trailer containing the recovered timestamp is added to the inner packets before they are forwarded to a tool port.

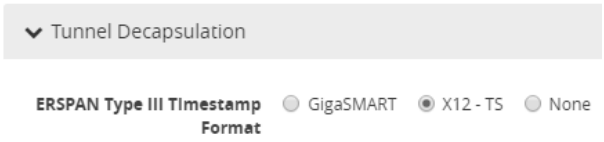


Figure 100*ERSPAN Type III Timestamp Formats on GigaSMART Groups Page*

The GigaSMART timestamp is added to the GigaSMART trailer. For the format of the GigaSMART trailer, refer to GigaSMART Trailer Reference. The x12-ts timestamp is added to the PRT-H00-X12TS trailer. For the format of the PRT-H00X12TS trailer, refer to the *GigaVUE-OS CLI Reference Guide*.

Only 10 ERSPAN sessions are supported per GigaSMART Group (gsgroup) when the timestamp format is configured to **GigaSMART** or **X12-TS**.

In summary for ERSPAN Type III encapsulation, **GigaSMART** does the following:

- strips encapsulating Ethernet + outer IP + GRE + ERSPAN Type III headers from incoming packets
- uses the timestamp field in ERSPAN packets and calculates the UTC timestamp, based on the timestamp in marker packets
- appends the UTC timestamp to the GigaSMART trailer or the PRT-H00-X12TS trailer if either **GigaSMART** format or PRT-HD00-X12TS (**X12-TS**) format is configured
- forwards packets to tool ports

ERSPAN Granularity

ERSPAN granularity is a setting that can be configured on the Cisco switch for the level of detail of the hardware timestamp in marker packets.

A marker packet will be considered overdue if it does not arrive by the following times:

- 00: Granularity—overdue after 119 hours
- 01: Granularity—overdue after 430 seconds (7 minutes)
- 10: 1588 PTP—overdue after 4.3 seconds

ERSPAN statistics include a count of overdue packets. Refer to Display ERSPAN Statistics for how to display the output and to ERSPAN Statistics Definitions for descriptions of these statistics.

PRT-H00-X12TS Unique ID

For the PRT-H00-X12TS format, you can obtain a unique ID identifying the port on which packets arrive. Use the following CLI command to display the mapping of ports to unique IDs:

```
(config) # show apps netflow port-id
============================
Port       Netflow port-id
----------------------------
1/1/x1     1
1/1/x2     2
1/1/x3     3
1/1/x4     4
1/1/x5     5
1/1/x6     6
1/1/x7     7
1/1/x8     8
1/1/x9     9
```

```
1/1/x10     10
1/1/x11     11
1/1/x12     12
---------------------------
```

Configure GigaSMART Operations for ERSPAN

Use the GigaSMART Operation (GSOP) page to configure the ERSPAN decapsulation types and options. For example, you can specify an ERSPAN flow ID, from 0 to 1023. Use this option when decapsulating traffic received over a Cisco-standard ERSPAN tunnel. Both ERSPAN Type II and Type III header decapsulation are supported.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

ERSPAN Tunnel Header Removal

To configure a tunnel to capture the ERSPAN packets, remove the ERSPAN header, and then forward the packets to a tool port, set the ERSPAN Decapsulation Flow ID to zero when creating the GigaSMART operation as shown in Figure 101Decapsulation Flow ID Set to Zero..

NOTE: A flow ID of zero is a wildcard value that matches all flow IDs.

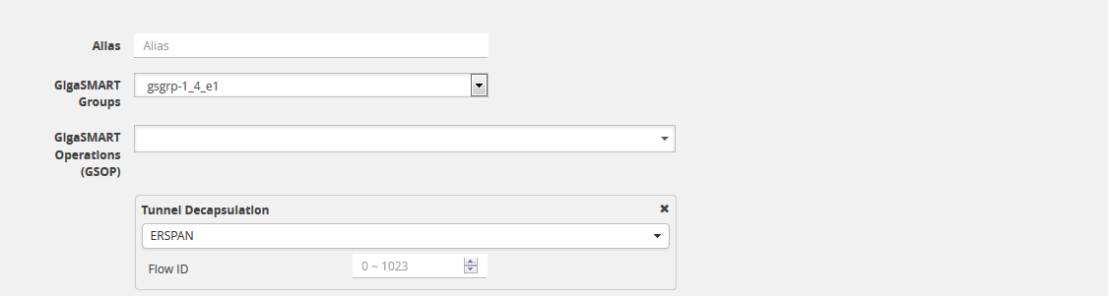


Figure 101 *Decapsulation Flow ID Set to Zero.*

In the following example, a tunnel is configured to capture ERSPAN packets, then the ERSPAN header is removed and the packets are forwarded to a tool port.

Task	Description	UI Steps
1.	Configure a tool type of port.	**a.** Select Ports > All Ports. **b.** Click Quick Port Editor. **c.** Use Quick search to find the ports to configure. For example,1/1/g1. **d.** Set the type to Tool and select Enable. **e.** Click OK.
2.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type an alias in the Alias field (for example, gsgrp1) and enter an engine port in the Port List field (for example 1/3/e1). **d.** Click Save.
3.	Configure the IP interface.	**a.** Select Ports > IP Interfaces. **b.** Click New. **c.** On the IP Interfaces page, in the Alias and Description fields, enter the name and description for the IP interface. **d.** Click the Ports field and select the port from the drop-down list. **e.** Enter the IP address, subnet mask, gateway, and MTU settings in the respective fields. For example, port 1/1/g2, IP address 10.10.10.10, mask 255.255.225.0, gateway 0.10.10.1, and MTU 1500. **f.** Click on the GigaSMART Group field to select the GigaSMART group. **g.** Click Save.
4.	Configure the GigaSMART operation and assign it to the GigaSMART group. **Note**: A flow ID of zero is a wildcard value that matches all flow IDs.	**a.** From the device view, select GigaSMART > GigaSMART Operations > GigaSMART Operation. **b.** Click New. **c.** Type an alias in the Alias field. **d.** From the GigaSMART Groups drop-down list, select the GigaSMART Group that you created in the second task. **e.** From the GigaSMART Operations (GSOP) drop-down list, select Tunnel Decapsulation. **f.** Select ERSPAN for the decapsulation type. **g.** Enter a value of 0 in the Flow ID field. The configuration should look like the example shown in Figure 101Decapsulation Flow ID Set to Zero..

Task	Description	UI Steps
		h. Click Save.
5.	Create a map.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type an alias in the Map Alias field that will help you identify this map. **d.** Select Regular and By Rule for the map type and subtype. **e.** Specify the network and tool ports in the Source and Destination fields, respectively. **f.** From the GSOP drop-down list, select the GigaSMART operation configured in task 4. **g.** Click Add a rule under Map Rules and create the following rule: Select IPv4 Protocol from the drop-down list and select GRE for Value, and then select Pass. **h.** Click Save.

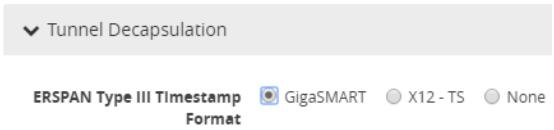
ERSPAN Type III Tunnel Header Removal

In this example, a tunnel is configured to capture ERSPAN packets. ERSPAN Type III packets are parsed, the ERSPAN header is removed, and the timestamp is calculated. A timestamp trailer is added before the packets are forwarded to a tool port.

NOTE: A flow ID of zero is a wildcard value that matches all flow IDs.

Task	Description	UI Steps
1.	Configure a port of type tool.	**a.** Select Ports > Ports > All Ports. **b.** Click Quick Port Editor. **c.** In the Quick View Editor, find the port to configure. **d.** Set Type to Tool. **e.** Select Enable **f.** Click OK. **g.** Close the Quick Port Editor.
2.	Configure a GigaSMART group and associate it with a GigaSMART	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups.

Task	Description	UI Steps
	engine port.	**b.** Click New. **c.** Enter a name in the Alias field **d.** Select the engine port in the Port List field. **e.** Click Save.
3.	Configure the IP interface.	**a.** Select Ports > IP Interfaces. **b.** Click New. **c.** On the IP Interfaces page, in the Alias and Description fields, enter the name and description of the IP interface. **d.** Click the Ports field and select the tool port from the drop-down list. **e.** Enter the IP address, subnet mask, gateway, and MTU settings in the respective fields. For example, port 1/1/g2, IP address 10.10.10.10, mask 255.255.225.0, gateway 0.10.10.1, and MTU 1500. **f.** Click on the GigaSMART Group field to select the GigaSMART group. **g.** Click Save.
4.	Configure the GigaSMART operation and assign it to the GigaSMART group. **Note**: A flow ID of zero is a wildcard value that matches all flow IDs.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Type an alias in the Alias field. **d.** From the GigaSMART Groups drop-down list, select the GigaSMART Group that you created in the second task. **e.** From the GigaSMART Operations (GSOP) drop-down list, select Tunnel Decapsulation. **f.** Select ERSPAN for the decapsulation type. **g.** Enter a value of 0 in the Flow ID field. The configuration should look like the example shown in Figure 101Decapsulation Flow ID Set to Zero.. **h.** Click Save.
5.	Configure a timestamp trailer format.	**a.** From the device view, select GigaSMART > GigaSMART Group. **b.** Select the GigaSMART Group created in Task 2. **c.** Under GigaSMART Parameters, go to Tunnel Decapsulation. **d.** For ERSPAN Type III Timestamp Format, select GigaSMART

Task	Description	UI Steps
		 **e.** Click Save.
6.	Create a map. The map contains a rule to allow marker packets (UDP) to be processed.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type an alias in the Map Alias field that will help you identify this map. **d.** Select Regular and By Rule for the map type and subtype. **e.** Specify a network ports in the Source fields. **f.** Select the tool port configured in Task 1 in the Destination field **g.** From the GSOP drop-down list, select the GigaSMART operation configured in task 4. **h.** Click Add a Rule and create the first rule. Select Pass, then select IPv4 Protocol, and then select GRE fro Value. **i.** Click Add a Rule and create the second rule. Select Pass, then select IPv4 Protocol, and then select UDP for Value. **j.** Click Save.

Display ERSPAN Statistics

To display ERSPAN statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics**. The ERSPAN statistics will be in the row labeled Tunnel Decap in the GS Operations column.

Refer to ERSPAN Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

Orchestrated Workflow Configuration of GigaSMART ERSPAN Tunnels

The GigaSMART ERSPAN tunnels can now be configured through the Orchestrated Configuration page. To configure GigaSMART ERSPAN Tunnels follow the below steps:

- In GigaVUE-FM go to **Traffic>Physical> Orchestrated Flows> Tunnels**.
- Click on **New**.
- Select **GigaSMART ERSPAN**. The New Tunnel configuration page appears. Enter the required information as described below:

Field	Description
Tunnel Name	The name of the tunnel. **NOTE:** Alias must not have spaces or any of these characters: *"?;:,/%@.
Tunnel Description	The description of the tunnel endpoint.
Traffic Direction	Select the traffic direction of the tunnel. Choose DECAP for decapsulation . **NOTE:** GigaSMART ERSPAN tunneling does not support Encapsulation.
Nodes	Select the Nodes on which you intend to create the tunnel.
IP Interface	Select the IP Interface from the list. You can create one using the Create an IP interface option if you do not have an IP interface. Refer to IP Interfaces to learn more about creating an IP Interface.
GigaSMART Group	Select the GigaSMART groups created. If you do not have any groups, you can configure them using the **Create GS Groups** option. Refer Create GigaSMART Operations – A Summary to know more about configure groups.
Key	Select the key to be used for decapsulation. You can choose a value from 0-1023.
Rules	Configure the map rules that needs to be adhered to while decapsulating or encapsulating the traffic. Click on **Rules Editor** to configure the rules.
Additional GigaSMART operations	Enable this option to add GigaSMART operations. When enabled you will be provided with GigaSMART Operations section which would allow you to choose among the following: • Add Header • Add Trailer • De-duplication • Load Balancing • Masking • Remove Trailer

Field	Description
	• Slicing
Destination Node	Select the destination node for the decapsulation tunnel.
Destination Port	Select the destination port for the decapsulation tunnel. You can edit the port details from the Port editor screen.

The configured tunnels provide you a **Details View** and **Troubleshoot View**. Click on the tunnel profile to view the below:

- **Details View**: Displays the configured parameters of the configured tunnel.
- **Troubleshoot View**: Displays the tunnel's statistics. Use the **Clear Stats** button to reset the statistics.

GigaSMART VXLAN Tunnel Decapsulation

Required License for VXLAN Decapsulation: Base (GigaVUE-HC3), Tunneling (GigaVUE-HC1)

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

Starting in software version 5.3, support for VXLAN tunnel termination is added to GigaSMART. VXLAN encapsulated packets originating from any device, such as the Gigamon cloud or a customer-specific device, will be received on a network port, then will be terminated at GigaSMART. The VXLAN payload (the inner packet) will be sent to tools. The reassembly of fragmented IP packets is also supported.

This section only includes VXLAN tunnel termination. It does not include VXLAN origination. To terminate a custom tunnel header that is not known to GigaSMART, use custom tunnel termination. Refer to GigaSMART Custom Tunnel Decapsulation

You can configure a GigaSMART-enabled node to act as the receiving end of a VXLAN tunnel by configuring a GigaSMART **tunnel-decap** operation with **type** set to **vxlan**. The high-level steps are as follows:

1. Configure an IP interface associated with network port and assign an IP address, subnet mask, and default gateway to the IP interface. The gateway forwards the encapsulated packet to the network port.

2. Create a GigaSMART operation with a **vxlan** decapsulation component.
3. Bind the GigaSMART operation to the IP interface associated with network port as part of a map.

At GigaSMART, VXLAN encap packets are received on the network port. After validation of the source port, destination port, and VXLAN Network Identifier (VNI) of the packet, the VXLAN tunnel header will be removed and the inner payload will be sent to a subsequent GSOP or to the tools. The VNI in the VXLAN header is validated against the user VNI provided. If it does not match, the packet will be dropped and counted as an error.

A VXLAN packet is identified using the **portdst** parameter. The destination port can be 4789, or any user-configured port number from 1 to 65535.

For an example configuration, refer to VXLAN Tunnel Termination Example

NOTE: GigaSMART operations with a tunnel component can be assigned to GigaSMART groups consisting of multiple GigaSMART engine ports.

VXLAN Tunnel Termination Example

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

To configure a VXLAN tunnel termination:

Step 1: Configure a Tool Port

1. Select **Ports > Ports > All Ports**.
2. Click **Quick Port Editor**.
3. Configure an available port as follows:
   - Select Tool in the **Type** field.
   - Enter an alias in the **Alias** field. For example, 1/4/x2.
   - Check the **Enable** check box.
4. Click **OK** to save the port.
5. Close the Quick Port Editor.

Step 2: Configure a GigaSMART group and associate it with a GigaSMART engine port.

1. From the device view, select **GigaSMART >GigaSMART Groups > GigaSMART Groups**.
2. Click **New**.
3. Type an aliasin the **Alias** field. For example, gsg.
4. Select an engine port in the **Port List** field. For example, 1/3/e2. (All engine ports have an 'e'.)
5. Click **OK** to save the GigaSMART group.

Step 3: Configure the IP Interface

1. Select **Ports > IP Interfaces**.
2. Click **New**.
3. In the **Alias** and **Description** fields, enter the name and description for the IP interface.
4. From the **Port** field, select any available network port. In this example, port 1/2/x1.
5. Complete the fields to configure the IP Interface:
   - Enter an **IP Address**. For example, 10.115.9.5.
   - Enter a **Mask**. For example, 255.255.255.0.
   - Enter a **Gateway**. For example, 10.115.9.1.
   - Enter the maximum transmission unit (MTU) for this port in the **MTU** field. For example, 1500.
   - Select the **GigaSMARTGroup** you created in step 2 of this process (gsg).
6. Click **OK** to save the IP interface configuration.

Step 4: Configure GigaSMART operation and assign to the GigaSMART Group

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.
2. Click **New**.
3. Type an alias in the **Alias** field to identify this operation. For example, vxlan2.
4. For **GigaSMART Groups**, select the group created in step 2 of this process (gsg).
5. For **GigaSMART Operations (GSOP)**, select Tunnel Decapsulation.

   A Tunnel Decapsulation form appears. Complete the fields as follows:

- Select “VXLAN” as the **Type**.
- Enter a **Source Port**.
- Enter a **Destination Port**.
- Enter a **VNI**.

6. Click **OK** to save the GSOP.

Step 5: Create a Map

1. Select **Maps > Maps > Maps**.
2. Click **New**.
3. Configure the map.
   - Type an alias the **Alias** field.
   - For **Type**, leave the default (Regular).
   - For **Subtype**, leave the default (By Rule).
   - For **Source**, select the IP interface you configured in step 3 (1/2/x1).
   - For **Destination**, select the tool port you configured in step 1 (1/4/x2).
   - Select the GigaSMART Operations (GSOP) associated with VXLAN decapsulation.
4. Under Map Rules, click **Add a Rule**.
   - Select **IP Version** from the drop list and set IP Version to **v4** when prompted.
   - Select **Pass** radio button for rule type.
   - Click **OK**.

Display VXLAN Tunnel GSOP

To display the VXLAN Tunnel GigaSMART operation:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP)**.
2. Select the VXLAN tunnel GSOP that you created in step 4 (vxlan2). The GSOP quick view appears.

Display VXLAN Tunnel Statistics

To display VXLAN tunnel statistics:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP)> Statistics**.

   The Statistics page displays all GSOP statistics in a table format.

2. In the table view, click the VXLAN tunnel GSOP alias that you created in step 4 (vxlan2) to display the Statistics quick view.

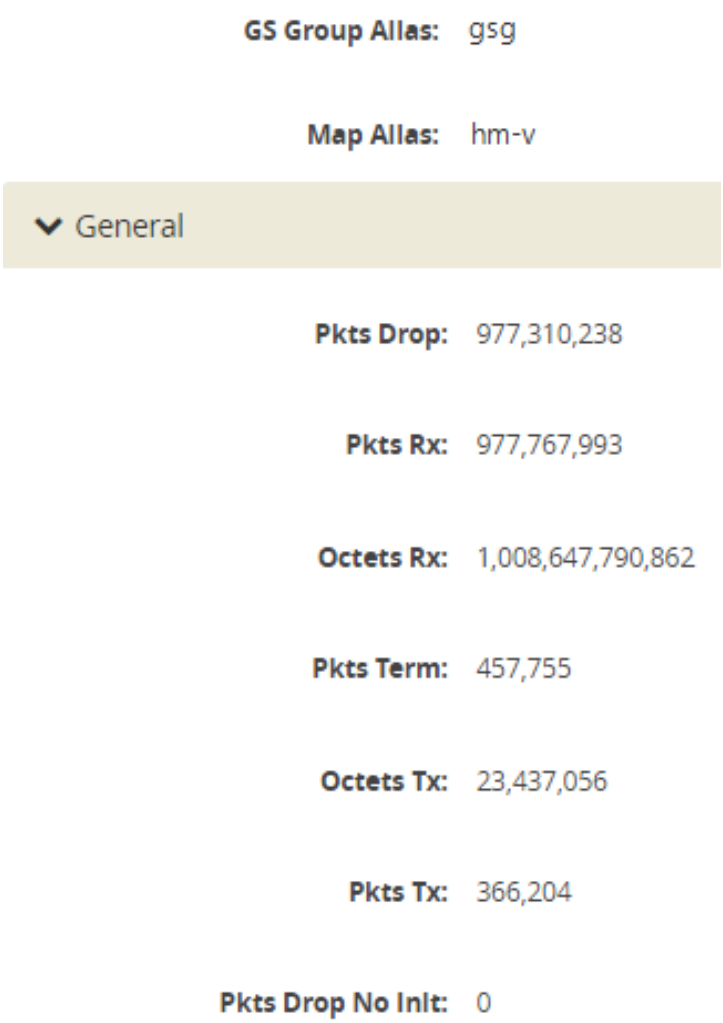


Figure 102*GSOP VXLAN Statistics Quick View*

GigaSMART Custom Tunnel Decapsulation

Required Licenses for Custom Decapsulation: Base (GigaVUE-HC3), Tunneling (GigaVUE-HC1),and Header Stripping

Starting in software version 5.3, support for custom tunnel termination is added to GigaSMART. Use custom tunnel termination to terminate a custom tunnel header that is received at the IP interface that is associated with a network port, but is not known to GigaSMART. The destination IP and MAC addresses must match the IP and MAC addresses of the network tunnel.

The packets that are successfully received at GigaSMART on a custom tunnel can be stripped, after some validations are performed, or can be sent to tools. The existing generic Header Stripping operation can be leveraged to remove the tunnel header if required. The reassembly of fragmented IP packets is also supported.

You can configure a GigaSMART-enabled node to act as the receiving end of a tunnel by configuring a GigaSMART **tunnel-decap** operation with **type** set to **custom** and Layer 4 (L4) source and destination ports. The high-level steps are as follows:

1. Configure an IP interface associated with network port and assign an IP address, subnet mask, and default gateway to the IP interface. The gateway forwards the encapsulated packet to the IP interface that is associated with a network port.
2. Create a GigaSMART operation (GSOP) with a **custom** decapsulation component.
3. (Optional) Create a chain of GigaSMART operations containing the custom tunnel decap GSOP and a generic Header Stripping GSOP.
4. Bind the GigaSMART operation to the P interface that is associated with a network port as part of a map.

The encapsulated packet will go from the P interface that is associated with a network port to GigaSMART, where basic validation against configured values will be performed. The packet will then be sent to the chained GSOPs, where the encapsulated header will be stripped off (if configured to strip) and sent to the tools.

NOTE: The generic Header Stripping operation is performed on the inner payload of the tunneled packet. Use a hybrid port for the Header Stripping GSOP.

For an example configuration, refer to GigaSMART Custom Tunnel Decapsulation

NOTE: GigaSMART operations with a tunnel component can be assigned to GigaSMART groups consisting of multiple GigaSMART engine ports.

Custom Tunnel Decapsulation Configuration Example

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

To configure a custom tunnel termination, do the following:

Configure a GigaSMART group/associate group with GigaSMART Engine Port

1. From the device view, select **GigaSMART > GigaSMART Groups**.
2. Click **New**.
3. Type an **alias** in the Alias field. For example, gsg.
4. Select **Ports > Ports.**
5. Select an engine port. For example,**1/1/g13**.
6. Click **Edit.**
7. Select **Network** for Type and enable **Admin**.
8. Click **OK**.

Configure the IP Interface

1. Select **Ports > IP Interfaces**.
2. Click **New**.
3. In the **Alias** and **Description** fields, enter the name and description for the IP interface.
4. From the **Port** drop-down list, select port 1/1/g3
5. Enter **10.115.9.5** in the IP Address field.
6. Enter **255.255.255.255** in the Mask field
7. Enter **10.115.9.1** in the Gateway field.
8. Enter **1500** in the MTU field.

9. Select the **GigaSMARTGroup** you created in the first step of this process. For example: **gsg**.
10. Click **Save**.

Configure GigaSMART operation and assign to GigaSMART Group

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.
2. Click **New**.
3. Type an **alias** in the Alias field to identify this operation. For example, custom2
4. For **GigaSMART Groups**, select **gsg**.
5. For **GigaSMART Operations (GSOP),** select **TunnelDecapsulation**.
   a. Select **Custom**
   b. Type **Source Port**
   c. Type **Destination Port**
6. Click **Save**

Create a Map

1. Select **Maps > Maps > Maps.**
2. Click **New**.
3. Configure the map.
   - Type map11 in the Alias field.
   - Select T **or tool** for **Type**.
   - Select **By Rule** for Subtype.
   - Select the from **network port 1/1/g1** for the Source.
   - Select the virtual port vp1 for the Destination.
4. Add a Rule.
   a. Click **Add a Rule**.
   b. Select **Pass**.
   c. Select **IPv4 Version** and set Version to **v4**.
5. Click **Save**.

Display Custom Tunnel GSOP

To display the custom tunnel GigaSMART operation, do the following:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP)**.
2. Select the custom tunnel from the list. The GSOP displays.

Display Custom Tunnel Statistics

To display custom tunnel statistics, do the following:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP)> Statistics**.
2. Select **Statistics**. The Statistics page displays basic tunnel termination details in a table format.
3. Click the **GSOP alias** to display all the statistics available for this GS Operation.

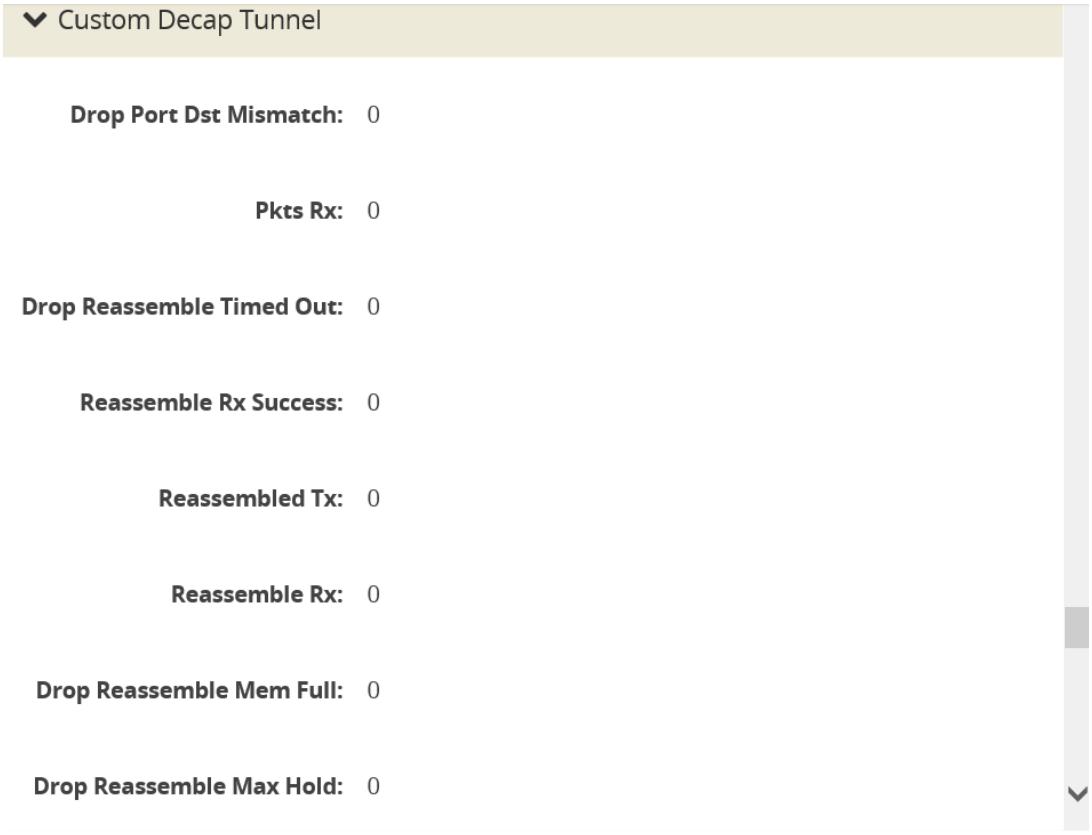


Figure 103

4. **Click** the **X** to close the **GS Operation Statistics** pane.

GigaSMART TCP tunnel

Required License- Advanced Tunneling license on GigaVUE-HC3. Referred to as “Tunneling license” on GigaVUE-HC1.

TCP tunnel feature routes the mirrored traffic from GigaVUE-VM to remote GigaVUE HC Series node reliably and without any reorder issues. TCP tunnel encapsulation is supported in the GigaVUE-VM node and the TCP tunnel decapsulation is supported in the GigaVUE HC Series node. Tunnel decapsulation can terminate more than one TCP connection initiated by the GigaVUE-VM node.

Configuration

The following are the steps to configure TCP tunnel between GigaVUE-VM and GigaVUE HC Series:

Steps	Refer to..
Configure GigaVUE-VM For Encapsulation	Configure TCP Tunnel in GigaVUE-VM User's Guide
Configure GigaVUE H-Series for Decapsulation	Configure GigaVUE HC Series for Decapsulation
Configure vMap for VMware	Configure vMap for VMware in GigaVUE-VM User's Guide

Supported Devices

TCP tunnel Decapsulation is supported only in GigaVUE-HC1, and GigaVUE-HC3 platforms.

Secure Tunnels

Secure Tunnels allow the traffic to be transported securely between a GigaVUE V Series to a GigaVUE HC Series or a GigaVUE HC Series to a GigaVUE HC Series device. The encapsulated traffic is sent to the GigaVUE HC Series device via secure tunnels in PCAPng format over a TLS connection.

Refer to the PCAPng Application section for more.

Secure Tunnels in GigaVUE HC Series has three primary use cases:

Secure Tunnels from GigaVUE V Series to GigaVUE HC Series

In this scenario, the encrypted traffic originating from a GigaVUE V Series node is transmitted to a remote GigaVUE HC Series device. This device then strips off the encryption, processes the traffic by utilizing GigaSMART applications (optional) and then forwards it to the tool. The GigaVUE HC Series can have a single map, multiple map rules, or multiple tools, as shown below.

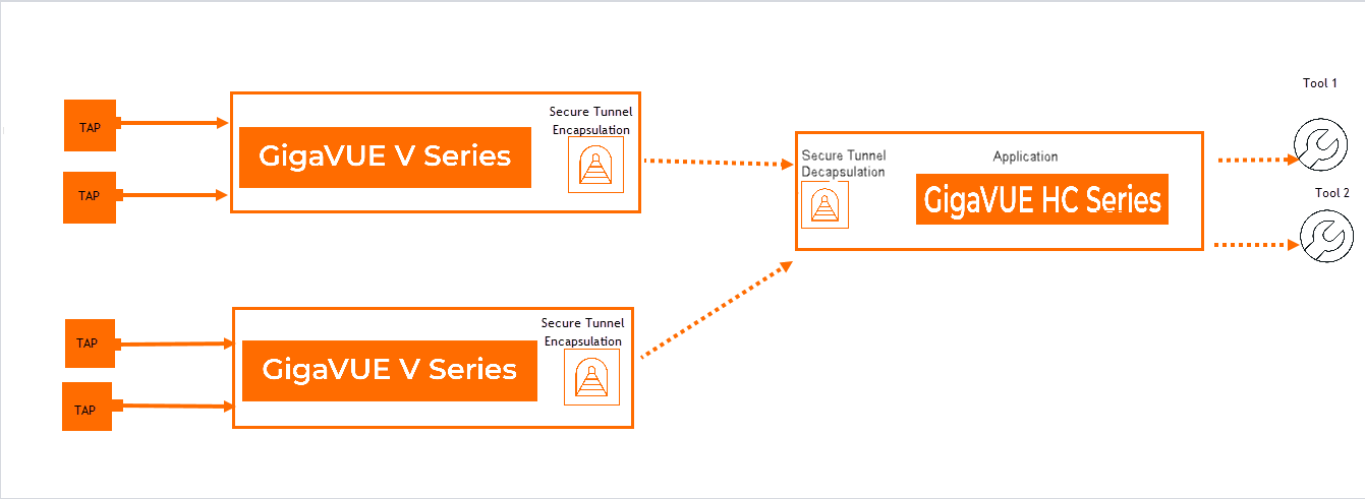


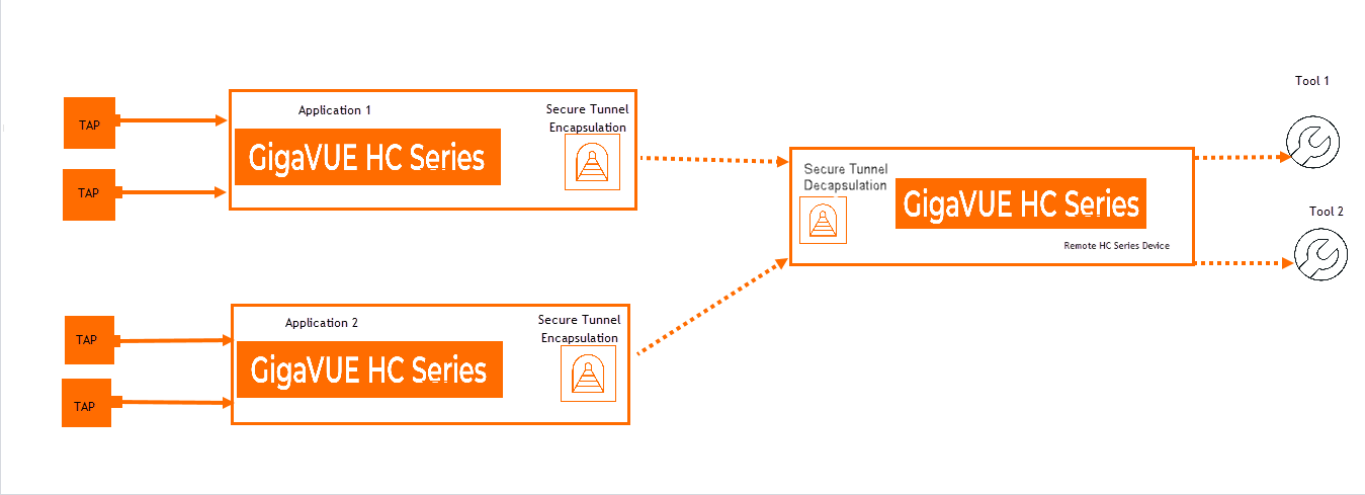
Figure 104*GigaVUE HC Series with a single IP Interface and multiple Map rules*

NOTE: Utilizing GigaSMART applications on the traffic is optional.

Secure Tunnels between the GigaVUE HC Series and a remote GigaVUE HC Series

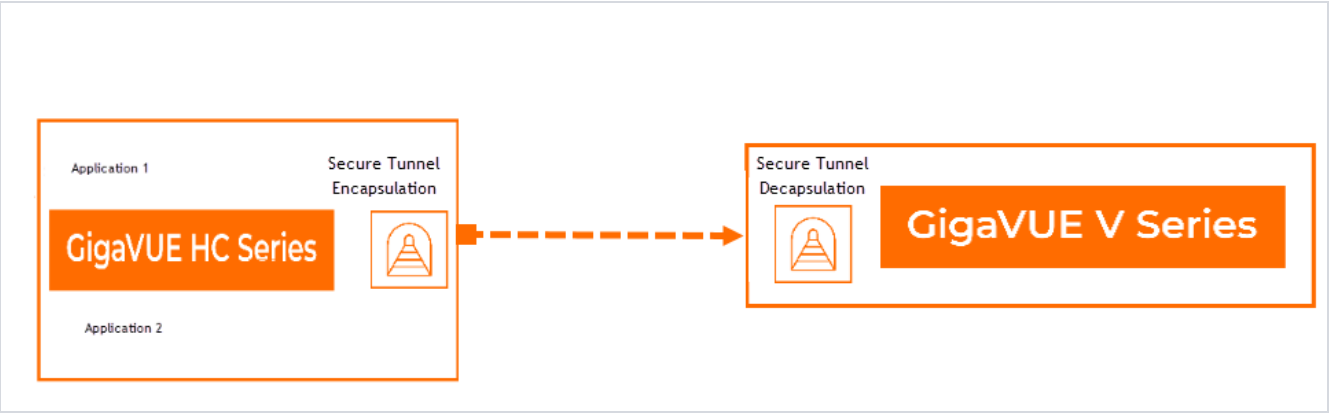
In this scenario, encrypted traffic is sent from the GigaVUE HC Series devices to a remote GigaVUE HC Series via secure tunnels, where it is decrypted, and forwarded to the tool.

In locally available GigaVUE HC Series devices you can utilize physical connectivity to securely send the traffic.



Secure Tunnels between GigaVUE HC Series to a GigaVUE V Series.

In this scenario , the encapsulated traffic from the GigaVUE HC Series is transported to GigaVUE V Series node. The GigaVUE V Series node decapsulates and processes the packet per the configuration.



Supported Platforms

Secure Tunnels is supported in the following platforms:

- GigaVUE-HC3 Gen3 (SMT-HC30C08)
- GigaVUE-HC1 Gen3 (SMT-HC1-S)
- GigaVUE-HCT (SMT-HC1-S)
- GigaVUE-HC1-Plus (onboard Gen3 GigaSMART engine)

Configure Secure Tunnels

To configure Secure Tunnels, follow the below steps:

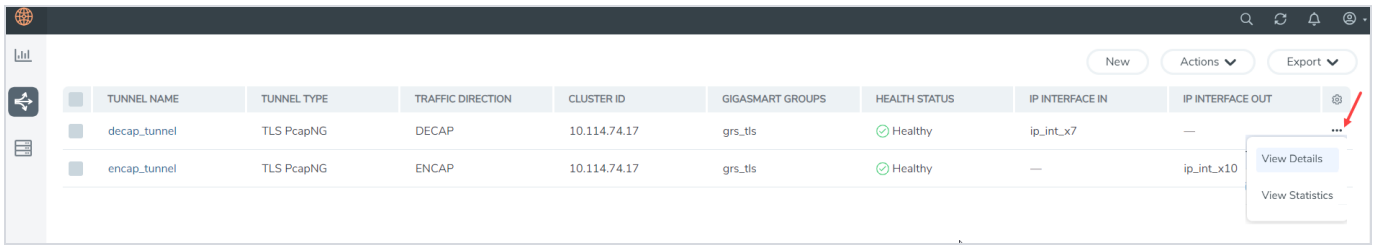
- In GigaVUE-FM, go to **Traffic>Physical> Orchestrated Flows> Tunnels**.
- Click on **New**.
- Select **GigaSMART TLS PCAPNG**. The New Tunnel configuration page appears. Enter the required information as described below:

Field	Description
Tunnel Name	The name of the tunnel. **NOTE:** Alias must not have spaces or any of these characters: *"?;:,/%@.
Tunnel Description	The description of the tunnel endpoint.
Tunnel Type	The tunnel type TLS-PCAPNG is selected by default.
Traffic Direction	Select the traffic direction of the tunnel. Choose DECAP for decapsulation or choose ENCAP for encapsulation. • If you choose Decap, select the IP Interface from the list. You can create one using the Create an IP interface option if you do not have an IP interface. Refer to IP Interfaces to learn more about creating an IP Interface. • If you choose Encap enter the Source Port. You can select from the drop-down list. To edit the Port type or to enable the port for Admin privileges use the Port Editor window. **NOTE:** In the scenario where secure tunnels needs to be established between GigaVUE V Series to a GigaVUE HC Series , you can utilize the **Configure Physical Tunnel** option provided at the GigaVUE V Series Secure Tunnel page. This allows you to conveniently configure secure tunnels at your physical device . Refer to Create Ingress and Egress Tunnels section in the respective cloud guides.
Nodes	Select the Nodes on which you intend to create the tunnel.
GigaSMART Group	Select the GigaSMART groups created. If you do not have any groups, you can configure them using the **Create GS Groups** option. Refer Create GigaSMART Operations – A Summary to know more about configure groups. **NOTE:** Secure tunnels can be configured only for Gen 3 nodes.
TCP Profile	Configure the TCP profile parameters such as: • TCP Keep Alive Timer - Enter the time duration between keep alive messages. The value ranges from 30-7200. The default value is 60. • Selective Acknowledgment- Choose Enable to turn on the TCP selective acknowledgments. • SYN Retries- Enter the value for number of times the SYN has to be tried. The value ranges from 1 to 6. The default value is 3.
SSL Profile	Enter the following SSL Profile parameters: • MTLS- MTLS is disabled by default. • Cipher - Only SHA 256 is supported. • TLS Version- Only TLS Version1.3 is supported. • Key Alias- Select the key to be used for decapsulation. You can choose a key

Field	Description
	from Node or GigaVUE-FM. o Add Keychain Password o Add Key Alias • Trust Store - Select the Trust Store key or click Append to add a new key for encapsulation.
Listener Port	For a decapsulation , specify the Listener port range. The range value is between 1- 30000. **NOTE:** If you update the Listener Ports in an existing secure tunnel destination, the tunnel will be re-established by deleting and recreating the maps.
Destination Node	Select the destination node for the decapsulation tunnel.
Destination Port	Select the destination port for the decapsulation tunnel. You can edit the port details from the Port editor screen.
DSCP	Enter the Differentiated Services Code Point (DSCP) value for an encapsulation.
Time to Live	For encapsulation tunnel, enter the value of the time interval till which the session needs to be available. The value ranges from 1 to 255. The default value is 64. **NOTE:** This is ignored for IPv6 based tunnel.
Encapsulation IP interface	Enter the interface IP address of the node (Destination IP) for an encapsulation tunnel.
Remote Tunnel IP, L4 Port & Source Port	Enter the IP address , L4 Port and Source Port values for an encapsulation tunnel. The Source Port range is from 30001 - 65535. **NOTE:** For decapsulation tunnel, It is recommended to match your L4 destination port rule to your listener port range. **NOTE:** The 50 ports from the configured exporter Source Port will be reserved for multi-engine GigaSMART group.
Rules	Configure the map rules that needs to be adhered to while decapsulating or encapsulating the traffic. Click on **Rules Editor** to configure the rules. For detailed information about map rules, refer to Map Rules. You can configure Inner Header qualifiers and MPLS Header qualifiers for GIgaVUE-TA400 device. Refer to Inner Header and MPLS Header Filtering
Additional GigaSMART operations	Enable this option to add GigaSMART operations. When enabled you will be provided with GigaSMART Operations section which would allow you to choose among the following: • Add Header • De-duplication • Load Balancing • Masking • Slicing

- Click on **Save.**

The configured Secure tunnels will be displayed. To get a quick view of the secure tunnel configuration go to ••• **>View Details** or click on the tunnel name. To know if your Secure tunnel is healthy refer the **Health Status** column.



NOTE: The Tunnels page accessible here at **Traffic>Physical> Orchestrated Flows> Tunnels** is solely for configuring Secure tunnels in GigaVUE HC Series and is an Early Access page. To learn more about configuring other Tunneling operations, refer to the Tunneling Operations section.

Edit Secure Tunnel

To edit a Secure Tunnel follow the below steps:

- Select the secure tunnel from the tunnels page.
- Click on **Actions > Edit** .
- Edit your tunnel parameters.
- Click **Save**.

Delete Secure Tunnel

Select the secure tunnel from the tunnels page.

- Click on **Actions> Delete**.
- Delete your tunnel.

Export Secure Tunnels

To export the configured Secure Tunnels, click **Export**. The secure tunnels is downloaded in either CSV or XLSX format.

View Tunnel Statistics

The configured tunnels provide you with a **Details View** and a **Troubleshoot View**. Click on the tunnel profile to view the below:

- **Details View**: Displays the configured parameters of the configured tunnel.
- **Troubleshoot View**: Displays the tunnel's statistics. Use the **Clear Stats** button to reset the statistics.

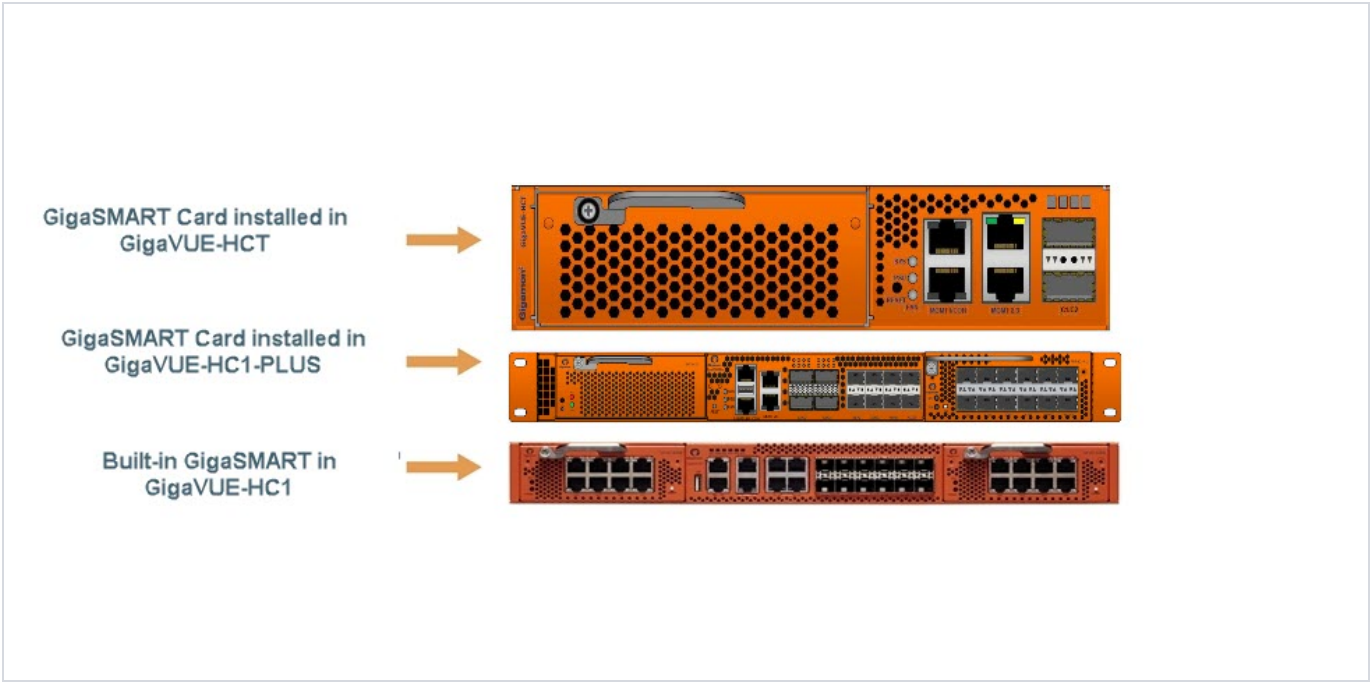
Limitations

The following limitations apply when configuring Secure Tunnel in GigaVUE HC Series.

- Secure-tunnel is not supported with the following applications:
  - GTP Correlation
  - Passive SSL
  - Inline SSL
  - Internet Content Adaptation Protocol (ICAP)
- GigaSMART operations with Secure Tunnel configuration is supported only with a regular map.
- The Tunnel Decapsulation listener port cannot be same as the Tunnel Encapsulation Source Port.

Work with GigaSMART Operations

This chapter describes how to use GigaSMART operations – advanced processing features available for use on GigaVUE-HC1 nodes, GigaVUE-HC1-Plus nodes with GigaSMART modules, and GigaVUE-HCT nodes with SMT-HC1-S module (Work with GigaSMART Operations).



Refer to the following sections for details:

- Access GigaSMART from GigaVUE-FM**
- Create GigaSMART Operations – A Summary
- Engine Watchdog Timer in GigaSMART
- GigaSMART Rules and Tips
- Virtual Ports
- GigaSMART Operations in Clusters
- How to Combine GigaSMART Operations
- Supported GigaSMART Operations
- GigaSMART Operations ***
- Order of GigaSMART Operations
- View GigaSMART Statistics

**Refer to the Access GigaSMART from GigaVUE-FM for an overview of how to access GigaSMART from GigaVUE-FM.

***Refer to GigaSMART Operations for comprehensive HowTo's on using all GigaSMART operations.

GigaSMART Licensing

Refer to the GigaVUE Licensing Guide for details.

NOTE: Contact your Gigamon Sales Representative to learn more about the available floating license options for your fabric configuration.

Access GigaSMART from GigaVUE-FM

You can access GigaSMART operations from within GigaVUE-FM, by accessing a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation pane of the device view on supported devices.

To access GigaSMART from the GigaVUE-FM interface:

1. Select **Physical** from the top navigation menu.
2. Select **Physical Nodes** from the side panel. This displays the list of Devices/Cluster Nodes managed by this instance of GigaVUE-FM.
3. Click the Cluster ID of any node to open the node. Once you are in the node, this part of the interface should behave just like HVUE for devices that support GigaSMART. Refer to About GigaSMART® Applications.
4. Click **GigaSMART** from the side navigation pane.

Refer to:

- About GigaSMART® Applications for devices that support GigaSMART.
- Create GigaSMART Operations – A Summary to get started with GigaSMART.
- GigaSMART Operations to learn how to use GigaSMART operations.

Quick Glance- How to Configure a GigaSMART Application

This section provides you a quick glance on document sections that can get you started on configuring the GigaSMART Applications.

GigaSMART Application	Refer to
Application Intelligence	
Application Intelligence	Create an Application Intelligence Session in Physical Environment
Application Filtering Intelligence	Create Application Filtering Intelligence for Physical Environment

GigaSMART Application	Refer to
Application Metadata Intelligence	Create Application Metadata Intelligence for Physical Environment
Application Session Filtering	Define ASF Session
Subscriber Intelligence	
GTP Correlation	• GTP Correlation Configuration Examples • GigaSMART GTP Whitelisting and GTP Flow Sampling Examples • Display Flow Ops Reports • GTP Engine Grouping Configuration Examples • GTP Stateful Session Recovery
SIP/RTP Correlation	• Configure SIP/RTP Correlation Engine
FlowVUE	• Configure FlowVUE o FlowVUE Configuration Examples
5G CUPS	• 5G Load Balancing • Configure CPN-UPN Communication Solution using Ansible • About Flow Sampling Rules and Maps • Monitoring CUPS Solution
GigaSMART TLS/SSL Decryption for Inline and Out-of-Band Tools	• Gigamon Inline Solutions
GigaSMART TLS/SSL Decryption for Out-of-Band Tools	• GigaSMART Passive TLS/SSL Decryption
Traffic Intelligence	
GigaSMART Adaptive Packet Filtering (APF)	• Implement APF Through the UI o Adaptive Packet Filtering Examples
Advanced Load Balancing	• Stateful Load Balancing • Stateless Load Balancing • Enhanced Load Balancing
De-duplication	• De-duplication Configuration Steps o Configure GigaSMART Parameters for Packet De-duplication o Example – GigaSMART De-duplication
Flow Masking	• GigaSMART Encapsulated Traffic Performance Enhancement
Header Stripping	• GigaSMART Header Stripping • Generic Header Stripping
Masking	• GigaSMART Masking

GigaSMART Application	Refer to
NetFlow Generation	• Configure Netflow Generation
Slicing	• Create Advanced Flow Slicing Profile • GigaSMART Packet Slicing
Tunneling	• Custom Tunnel Decapsulation Configuration Example • ERSPAN Tunnel Header Removal o ERSPAN Type III Tunnel Header Removal • GigaSMART IP Encapsulation/Decapsulation (GigaSMART Tunnel) • Configure L2GRE Tunnel Encapsulation and Decapsulation • GigaSMART VXLAN Tunnel Encapsulation • GigaSMART VXLAN Tunnel Decapsulation • Configuration

Create GigaSMART Operations – A Summary

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

GigaSMART operations require the following steps:

1. **Create GigaSMART groups on the GigaSMART Groups page (select GigaSMART > GigaSMART Groups > GigaSMART Groups)**

   Use GigaSMART groups to manage and budget GigaSMART processing power. Use the ***New*** button on the GigaSMART Group page to create groups of the available GigaSMART ports in a given chassis.

   NOTE: The GigaSMART engine ports in a GigaSMART group can be on different line cards or modules in the same GigaVUE-HC3 chassis. However, all GigaSMART engine ports must be on the same chassis.

   The number of GigaSMART engine ports are as follows:
   - The GigaSMART-HC0 module includes one GigaSMART engine port.
   - The GigaVUE-HC1 node includes one GigaSMART engine port.

   The number of GigaSMART engine ports available in a chassis depends on the number of GigaSMART line cards in the chassis

   The processing power of the GigaSMART engine ports is as follows:

- Each GigaSMART port on the GigaSMART-HC0 module can process packets at **40Gb**.
- GigaSMART port on the GigaVUE-HC1 node can process packets at **20Gb**.

GigaSMART engine ports are numbered with an **e** prefix using **<bid/sid/e1..e2>** nomenclature – **1/1/e1**, for example.

> **NOTE:** The ports in a GigaSMART group can be on different line cards or modules in the same GigaVUE-HC3 chassis. However, they must all be on the same chassis.

> **NOTE:** The slot ID for a GigaVUE-HC1 chassis is fixed at **1**.

Each GigaSMART operation you create in the next step must be assigned to one of the GigaSMART groups you create in this step. You can select GigaSMARToperation only when the card supports the operation.

2. **Create GigaSMART operations using the GigaSMART *Operations page (select* GigaSMART > GigaSMART *Operations (GSOP))***

   Give your GigaSMART operation a name, include a valid combination of GigaSMART operations, and assign it to one of the GigaSMART groups created in the previous step.

   Refer to How to Combine GigaSMART Operations for details on supported combinations of GigaSMART operations.

   You can also configure how (or, in some cases, whether) a GigaSMART operation attaches a trailer that indicates where a packet arrived in the Gigamon Deep Observability Pipeline and how it was modified. This trailer can be interpreted using a recent version of the Wireshark® Protocol Analyzer. Refer to GigaSMART Trailer Reference for details on the GigaSMART Trailer.

3. **Apply GigaSMART operations to network ports in maps**

   The New Map and Edit Map pages contain a **GigaSMART Operations (GSOP)** field that lets you select a GigaSMART operation to be used in a map. Refer to Manage Maps for details. Keep in mind, however, that GigaSMART operations **must** be selected before destination tool ports or collector destinations.

Groups of GigaSMART Engine Ports

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

Use the GigaSMART Group page to create groups of GigaSMART engine ports in a given GigaVUE chassis. (To go to the GigaSMART Group page, select **GigaSMART > GigaSMART Groups > GigaSMART Groups > New**.)

NOTE: You cannot create a new group of GigaSMART engine ports when no GigaSMART engine ports are available on the device.

The GigaSMART engine ports in a GigaSMART group can be on different line cards or modules in the same GigaVUE-HC3 chassis. However, all GigaSMART engine ports must be on the same chassis.

Use groups of GigaSMART engine ports to increase the processing power of GigaSMART. The processing power of the GigaSMART engine ports is as follows:

- Each of the two GigaSMART engine ports in an SMT-HC3-C08 module on GigaVUE-HC3 can process packets at up to **200Gb**.
- The GigaSMART engine port in the GigaVUE-HC1 node can process packets at up to **20Gb**.
- The GigaSMART engine port in the GigaVUE-HC1 Gen3 can process packets at up to **80Gb**.
- The GigaSMART engine port in the GigaVUE-HC1-Plus node can process packets at up to **100Gb.**
- The GigaSMART engine port in the GigaVUE-HCT Gen3 can process packets at up to **80Gb**.

The number of GigaSMART engine ports are as follows:

- Each GigaSMART-HC0 module includes one GigaSMART engine port.
- Each SMT-HC3-C08 module on GigaVUE-HC3 includes two GigaSMART engine ports.
- The GigaVUE-HC1 node includes one GigaSMART engine port.
- The GigaVUE-HC1-Plus node includes one fixed GigaSMART module.

The number of GigaSMART engine ports available in a chassis depends on the number of GigaSMART line cards or modules in the chassis as follows:

- Up to four modules in the GigaVUE-HC3 (eight GigaSMART engine ports).
- One module in GigaVUE-HC1 (one engine port).
- Up to two modules in GigaVUE-HC1 Gen 3 (two GigaSMART engine ports).
- One module in GigaVUE-HCT Gen3 (one engine port).

The following table provides a summary:

GigaVUE Node	Maximum GigaSMART Line Cards or Modules per Node	Number of GigaSMART engine ports per Line Card or Module	Maximum Number of GigaSMART engine ports in a GigaSMART group (gsgroup)
GigaVUE-HC3	4	2	8
GigaVUE-HC1	1 built-in	1	1
GigaVUE-HC1 Gen3	2	1	2
GigaVUE-HC1-Plus	1 built-in rear module, 1	2	2
GigaVUE-HCT Gen3	1 built-in front	1	1

NOTE: Grouping Generation 3 and Generation 2 cards is not supported.

NOTE: On supported GigaVUE HC Series platforms with Generation 3 GigaSMART cards, you can use the GigaSMART Groups page to reload one or more GigaSMART engine ports in a group. This provides a granular way to reset only the engines in a selected group without reloading the entire chassis or GigaSMART card. Reload is not available on Generation 2 cards.

Engine grouping for GTP is a special case, which is described in GTP Engine Grouping.

GigaSMART engine ports are numbered with an **e** prefix using **<bid/sid/e1..e2>** nomenclature, such as 1/1/e1. For example:



Reload GigaSMART engine ports in a group

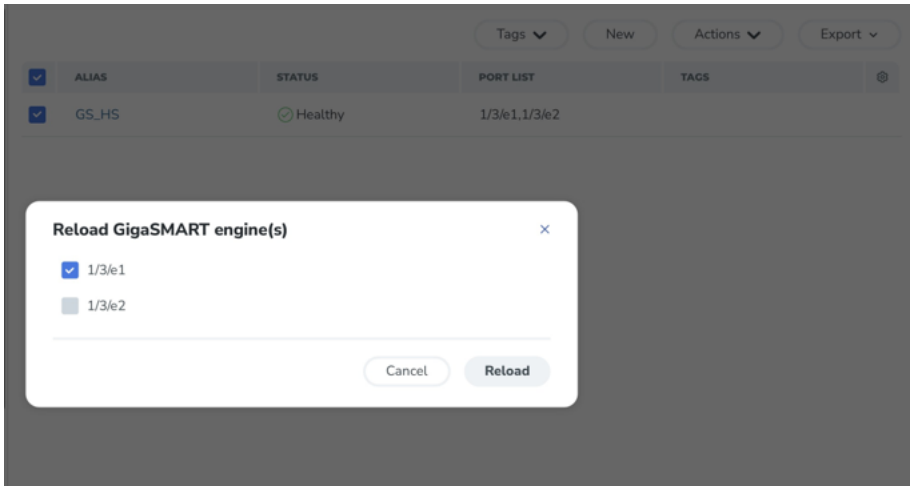
On supported GigaVUE HC Series devices with Generation 3 GigaSMART cards, you can reload one or more GigaSMART engine ports in an existing GigaSMART group. Reloading a group restarts the selected engines in that group without rebooting the chassis or the GigaSMART card.

Use this operation to recover or troubleshoot a group while minimizing impact to other traffic.

To reload GigaSMART engine ports in a group:

1. In GigaVUE-FM, navigate to the device and select **GigaSMART > GigaSMART Group.**

2. On the GigaSMART Groups page, select the GigaSMART group (gsgroup) that you want to reload.
3. From the **Actions** menu, choose **Reload**.
4. In the reload dialog:
   - Review the list of GigaSMART engine ports that belong to the selected group.
   - Review the list and select the engine that needs to be reloaded.
5. If the group is associated with inline traffic configuration , a warning is displayed indicating that the reload will impact traffic on that group. Proceed only if you have planned for a brief disruption.
6. Click **Reload** to start the reload.



Only the selected GigaSMART engine ports in the group are reloaded; other groups and unrelated traffic on the device continue to run. The reload is a software-level restart of the engines and does not reboot the chassis or the GigaSMART card.

NOTE: After you apply the reload, use existing GigaSMART statistics and monitoring views in GigaVUE-FM to confirm that traffic resumes for the affected group. If the Reload action is not available for a group, verify that the device uses Gen 3 GigaSMART cards on a supported GigaVUE HC Series platform; reload is not supported on Gen 2 cards .

How to Use GigaSMART Operations – Example

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

The following procedure summarizes the major steps in creating and using a GigaSMART operation.

1. From the device view, start by selecting **GigaSMART > GigaSMART Groups > GigaSMART Groups** and clicking **New** to create a **GigaSMART group** – a collection of one or more internal GigaSMART ports available in a given chassis. GigaSMART groups are used to process GigaSMART operations – each GSOP you create is assigned to a GigaSMART group.

In this example, a GigaSMART group called **GS1** is created that uses virtual port **e1** on the GigaSMART-HD0 line card in slot 2 of box 16 (**16/2/e1**):

On the GigaSMART Group configuration page:

a. Enter GS1 in the **Alias** field.

b. Click in the **Port List** field to select the port.

   On the GigaSMART Group configuration page, you can also set parameters for specific types of GigaSMART operations.

c. Click **Save**.

2. Next, you can create a **GigaSMART operation** – a combination of packet modification actions that can be used in a map – and assign it to a GigaSMART group for processing.

In this example, a GigaSMART operation called **tcpmask** is created that will overwrite 16 bytes of packet data starting 64 bytes after the end of the TCP header using a hexadecimal **ee** pattern. The GigaSMART operation is assigned to the **GS1** GigaSMART group created in the first step.

To create a GigaSMART operation called **tcpmask**:

a. From the device view, select **GigaSMART > GigaSMART Operations**, and then click **New**.

b. In the **Alias** field, enter **tcpmask**.

c. Click in the **GigaSMART Groups** field and select **GS1** from the list of GigaSMART groups.

d. Click in the **GigaSMART Operations (GSOP)** field and select **Masking** from the list. The configuration dialog for Masking displays.

e. Configure Masking as follows:

   - **Protocol**: TCP
   - **Offset**: 64
   - **Pattern**: ee
   - **Length**: 16

3. Once you have set up a GigaSMART operation, you can include it as part of a map with the **GigaSMART Operations (GSOP)** field in the Map configuration page. In this example, the **tcpmask** GigaSMART operation is combined with an IP Version pass rule so that all IPv4 traffic processed is masked according to the GSOP created in the previous step.

If you are not sure which GigaSMART operation you want to use, click in the GSOP field to display a list of the operations you have already configured.

To configure the map:

a. Select **Maps > Maps > Maps**, and then click **New** to open the New Map page.

b. On the New Map page, configure the map as follows:

- **Alias**: gsmap
- **Type**: Regular
- **Subtype**: By Rule
- **Source**: select the network ports (for example: 16/3/x7, 16/3/x8, 16/3/x9, 16/3/x10, 16/3/x11, 16/3/x12)
- **Destination**: select the tool port (for example: 16/3/x1)
- **GigaSMART Operations (GSOP)**: tcpmasking (GS1)

c. Click **Add Rule** and specify the following for the rule:

- Select **Pass**
- Click in the **Rule** field and select **IP Version**
- Select **v4** for **Version**.

d. Click **Save**.

Here, a map named **gsmap** is created that forwards IPv4 traffic from network ports 16/3/x7..x12 to tool port 16/3/x1. The traffic will be masked using the **tcpmask** GigaSMART operation created in Step 2.

Engine Watchdog Timer in GigaSMART

In rare scenarios, a packet processing core in the CPU of a GigaSMART engine can enter a deadlocked state. The engine watchdog timer detects the issue and reloads the GigaSMART engine after a specified number of seconds. The engine watchdog timer is enabled by default with a value of 60 seconds. The maximum number of seconds is 600 seconds.

NOTE: If a core is in a deadlocked state, all packets are dropped.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

To configure the engine watchdog timer, do the following:

1. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.
2. Click **New**.

NOTE: If you are modifying an existing GigaSMART Group, select the GS Group and click **Edit**. Otherwise, .

3. In the **Alias** field, enter an alias for this GS Group.
4. In **Port List** field, select the engine port for this GS Group.
5. Under **Engine Timer**, do the following:
   a. Select **Enable** to enable the time or clear the checkbox to disable the timer.
   b. In the **Engine Watchdog Time field**, set the number of seconds to wait before reloading the engine. The minimum is 60. The maximum is 600.

In the following example, the timer is enabled and set to 100 seconds.

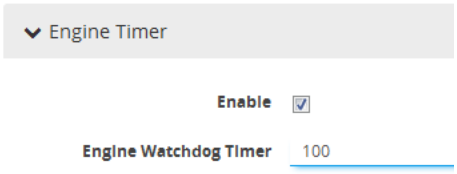


Figure 105*Engine Watchdog Timer*

Tunnel Health Checks

Starting in software version 5.3, there are tunnel health checks. The reachability of IP destinations is checked and, if the destinations are not reachable, packets will not be sent or will stop being sent.

The tunnel health check on the GigaSMART card defines destinations as follows:

- IP destinations used for sending packets from a single IP interface with tool port to a single IP destination
- tunnel endpoints used for load balancing from a single IP interface with tool port to multiple IP destinations

An SNMP notification can be sent when the status of a IP destination or tunnel endpoint changes, either from Up to Down or from Down to Up.

> **NOTE:** Tunnel Health-check is not supported for GMIP tunnel.

To enable the SNMP notification, refer to Enable or Disable Events for SNMP Notifications..

To configure ICMP health check parameters for the GigaSMART group, refer to Figure 106Configure ICMP Health Check Parameters.

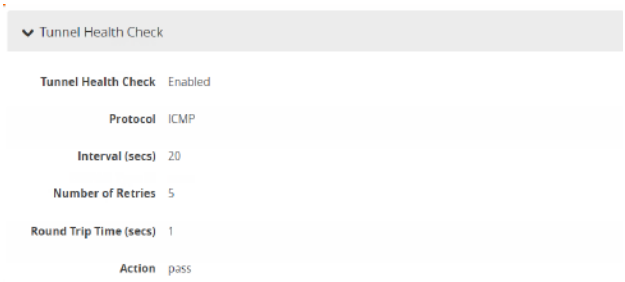


Figure 106*Configure ICMP Health Check Parameters*

To configure UDP health check parameters for the GigaSMART group, refer to Figure 107Configure UDP Health Check Parameters.

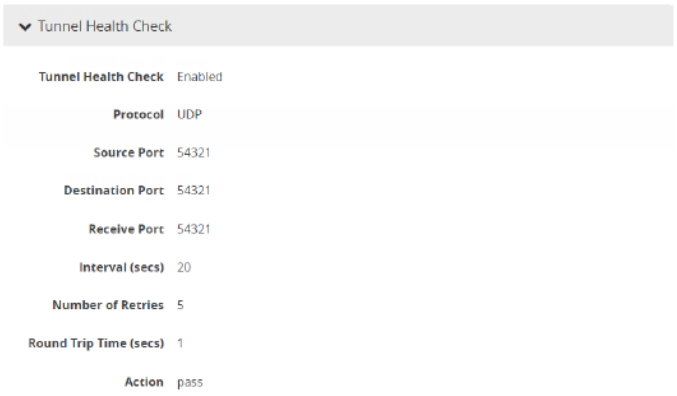
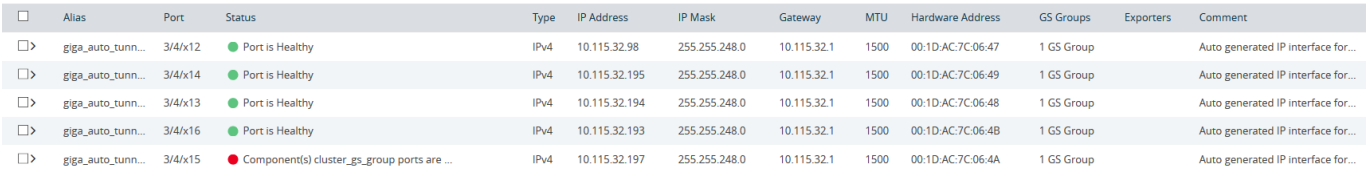


Figure 107*Configure UDP Health Check Parameters*

To view IP interface status, refer to Figure 108View IP Destination Status.



Alias	Port	Status	Type	IP Address	IP Mask	Gateway	MTU	Hardware Address	GS Groups	Exporters	Comment
giga_auto_tunn...	3/4/x12	Port is Healthy	IPv4	10.115.32.98	255.255.248.0	10.115.32.1	1500	00:1D:AC:7C:06:47	1 GS Group		Auto generated IP interface for...
giga_auto_tunn...	3/4/x14	Port is Healthy	IPv4	10.115.32.195	255.255.248.0	10.115.32.1	1500	00:1D:AC:7C:06:49	1 GS Group		Auto generated IP interface for...
giga_auto_tunn...	3/4/x13	Port is Healthy	IPv4	10.115.32.194	255.255.248.0	10.115.32.1	1500	00:1D:AC:7C:06:48	1 GS Group		Auto generated IP interface for...
giga_auto_tunn...	3/4/x16	Port is Healthy	IPv4	10.115.32.193	255.255.248.0	10.115.32.1	1500	00:1D:AC:7C:06:4B	1 GS Group		Auto generated IP interface for...
giga_auto_tunn...	3/4/x15	Component(s) cluster_gs_group ports are ...	IPv4	10.115.32.197	255.255.248.0	10.115.32.1	1500	00:1D:AC:7C:06:4A	1 GS Group		Auto generated IP interface for...

Figure 108*View IP Destination Status*

To view tunnel endpoint status, refer to Figure 109View Tunnel Endpoint Status.

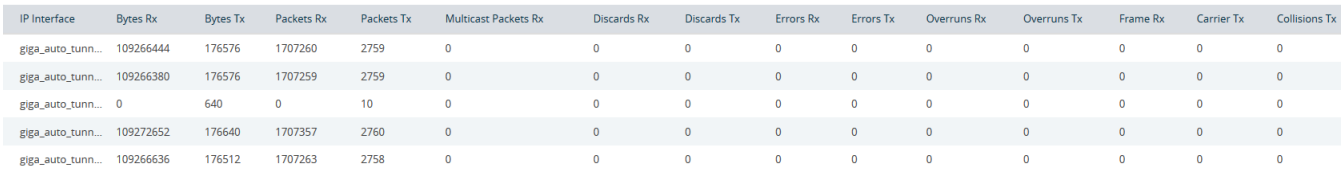


Tunnel Endpoint ID	Alias	IP Address	Status
te1		192.168.1.145	Up
te2		192.168.1.188	Down
te3		192.168.1.123	Up

Total Items : 3

Figure 109*View Tunnel Endpoint Status*

To view IP interface statistics, refer to Figure 110View IP DestinationStatistics.



IP Interface	Bytes Rx	Bytes Tx	Packets Rx	Packets Tx	Multicast Packets Rx	Discards Rx	Discards Tx	Errors Rx	Errors Tx	Overruns Rx	Overruns Tx	Frame Rx	Carrier Tx	Collisions Tx
giga_auto_tunn...	109266444	176576	1707260	2759	0	0	0	0	0	0	0	0	0	0
giga_auto_tunn...	109266380	176576	1707259	2759	0	0	0	0	0	0	0	0	0	0
giga_auto_tunn...	0	640	0	10	0	0	0	0	0	0	0	0	0	0
giga_auto_tunn...	109272652	176640	1707357	2760	0	0	0	0	0	0	0	0	0	0
giga_auto_tunn...	109266636	176512	1707263	2758	0	0	0	0	0	0	0	0	0	0

Figure 110 *View IP DestinationStatistics*

To view GigaSMART operation statistics, refer to Figure 111View GigaSMART Operation Statistics.



Session Current Total: 0
Reassemble Rx: 0
Drop Reassemble Mem Full: 0
Session Lookup: 0
Session Alloc: 0
Drop Key Mismatch: 0
Drop Wrong Addr: 0
Fragment Rx: 0
Drop No Ndp: 0
Pkts Rx: 0
Fragment Tx: 0
Drop Unknown Proto: 0
Session Lookup Success: 0
Drop Reassemble Max Hold: 0
Drop Destination Down: 0

Figure 111 *View GigaSMART Operation Statistics*

Configure Hashing

The **Hash** option in the GigaSMART Group page allows you to select the required hashing option for the GigaSMART Groups. The following options are available:

- **Advanced**: Advanced hashing. Refer to the "*Fabric Advanced Hashing*" section in the *GigaVUE Administration Guide*.
- **IPSrcIPDst**: Fixed two tuple hashing based on outer IP, Src and Dst.

GigaSMART Rules and Tips

When using GigaSMART operations, keep in mind the following rules and tips:

Note	Description
Use on Any GigaSMART-Enabled Node	Maps including GigaSMART operations can be bound to **any network port** on a GigaSMART-enabled node: o Standalone GigaVUE-HC3 with SMT-HC3-C08 module. o Standalone GigaVUE-HC1 nodes. o Any GigaVUE H Series node operating in a cluster with one of these node types. Clustered nodes can use maps with GigaSMART operations so long as there is at least one node available in the cluster with GigaSMART capabilities. The GigaSMART group used to power the GigaSMART operation **does not** need to reside on the same physical chassis as the network or tool ports for the map. Refer to GigaSMART Operations in Clusters for some illustrations of how this works.
Use in Maps with Standard Rule Criteria	Combine GSOPs with map rules carefully to ensure selective application. For example, headers in an IPv4 packet end at a different offset than those in an IPv6 traffic. You can create maps for the following: o Identify IPv4 traffic, slice it at a 64-byte offset and forward the results to tool port 5. o Identify IPv6 traffic, slice it at an 82-byte offset and forward the results to tool port 6. GSOPs are applied to **all** packets matching **any** rule in the map in which the GSOP is included.
Editing Maps and GSOPs	Maps containing GigaSMART operations (GSOPs) should not be edited. Also GSOPs should not be edited once they are associated with maps. **Note**: In Gen3 platform, if there are multiple maps with GSOPs sourcing from the same network ports, when you edit, delete, or reconfigure a map, traffic drop from one map to another maps GSOP will be observed.
Rules for Tunneling Operations	The rules for tunneling operations are as follows: o A map including a GigaSMART operation with a tunnel decapsulation component (**tunnel-decap**) cannot also include a collector rule. The system prevents situations that would violate this rule. o GigaSMART operations with a tunnel decapsulation component can only be assigned to GigaSMART groups consisting of a single GigaSMART port. o IP interfaces cannot be shared with map-passalls, tool-mirrors, port-pairs, or other regular maps. o For devices involved in tunneling using GMIP. the recommendation is that they run the same software version on both sides of the tunnel (encapsulating/decapsulating).

Note	Description
Combine Multiple Components in a Single Operation	You can combine multiple GigaSMART components into a single operation. For example, you could set up a single GigaSMART operation that masks a packet, strips its VLAN header, and applies a trailer. **Note**: With the exception of slicing and masking, most GigaSMART components can be combined in a single operation. Slicing and masking can be combined with other components **but not with each other**. Refer to How to Combine GigaSMART Operations for details on the combinations of GigaSMART operations. **Note**: The **[trailer <remove>]** argument cannot combined with others – it must be used by itself. Refer to Remove GigaSMART Trailers for details. **Note**: NetFlow can only be combined with de-duplication and with APF (using second level maps).
GigaVUE-HC3 Nodes with Multiple GigaSMART Modules	The GigaVUE-HC3 chassis supports up to four GigaSMART SMT-HC3-C08 modules. Each module has two GigaSMART engine ports. Each GigaSMART engine port can process packets at up to 100Gb.

Virtual Ports

Virtual ports (vports) are used in flow maps to aggregate and redirect traffic to the GigaSMART ports. It is an aggregation point for traffic to be directed to the GigaSMART second level maps. Second level maps are used for configuring filtering rules enabled through GTP correlation and Adaptive Packet Filtering (APF).

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

Virtual Port Rules

The following rules apply to single virtual ports:

1. A given vport can only belong to one GigaSMART group.
2. Different first level maps with the same network ports can use the same vport. However, you must keep in mind the limitation described in Rule 1.
3. Different first level maps with different network ports can use the same vport. However, you must keep in mind the case described in Rule 2.

4. Different vports can be configured on the same GigaSMART group, but must be used in different maps.
5. A GigaSMART operation can only belong to one GigaSMART engine group.
6. In a first level map, you can specify a vport but not a GigaSMART operation.
7. The vport and the GigaSMART operation used in a second level map must be defined on the same GigaSMART group.
8. In a second level map, a maximum of 5 maps are allowed to be attached to a vport. The maximum number of gsrules in each map is 5. The maximum number of flowrules in each map is 32.

   NOTE: For GTP: refter to GigaSMART GTP Whitelisting and GTP Flow Sampling. For APF and ASF: refer to GigaSMART Adaptive Packet Filtering (APF) and ASF and Buffer ASF Examples.

9. Multiple GTP flow sampling maps can receive traffic from the same virtual port when **GTP Overlap** is enabled.

Create Virtual Port

To create a virtual port, use the following procedure:

1. From the device view, select **GigaSMART > Virtual Ports**.
2. Click **New**.
3. On the Virtual Ports page, do the following:
   a. In the **Alias** field, enter a name for the virtual port. For example, gsTraffic.
   b. From the GigaSMART Group drop-down list, select the GigaSMART group to associate with the virtual port. For example, gsgrp1.
   c. Select the GTP Overlap check box to enable multiple GTP flow sampling maps to receive traffic from the same virtual port.
   d. In the Inline Failover Action drop-down list, select one of the following options:
   - **Tool Bypass** — When the inline tool fails, all traffic coming to the respective inline tool is directed via the bypass path.
   - **Network Bypass** — When the inline tool fails, the traffic is directed to multiple inline tools associated with an inline network or inline network group using rule-based inline maps.
   - **Tool Drop** — When the inline tool fails, all traffic coming to the respective inline tool is dropped.
   - **Network Drop**—When the inline tool fails, all traffic coming to the respective inline tool is dropped.

- **Network Port Forced Down**—When the inline tool fails, the inline network ports of the respective inline network are forced as "down".

e. Click **OK**.

4. Create a GigaSMART operation with an Adaptive Packet Filtering (APF) component.
   a. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.
   b. Click **New**.
   c. In the **Alias** field, enter g1 for the GigaSMART Operation.
   d. From the **GigaSMART Groups** list, select a GigaSMART group.
   e. From the GigaSMART Operations (GSOP) list, select **Adaptive Packet Filtering**.
5. Create a first-level map name map1 to direct traffic to the virtual port gsTraffic.
   a. Go to the Maps page and click **New**. The New Map page opens.
   b. Enter map1 in the **Alias** field.
   c. For **Type**, select **First Level** and for **Subtype** select **By Rule** map
   d. For **Source**, select a network port. For example, a network port with the alias N1.
   e. For **Destination**, select the virtual port (gsTraffic).
   f. Click Add a Rule, and create rule with the following conditions:
      - Pass
      - VLAN 20
      - IPv4 Protocol UDP
      - Port Destination 2152
   g. Click **Save**.
6. Create second-level maps named map2 and map3 to direct traffic from the virtual port to the GigaSMART operation.

Configure the map map2 with the **Source** as virtual port gsTraffic, the **Destination** as port T1, and select g1 for **GigaSMART Operation (GSOP)**. For map2, add the following two rules.

Rule 1:
- Pass
- IPv4 Destination 65.128.7.21 Cidr 32
- IPv4 Protocol TCP
- Port Destination 80

Rule 2:
- Pass

- IPv4 Destination 98.43.132.70 Cidr 32
- IPv4 Protocol TCP
- Port Destination 80

Configure the map map3 with the **Source** as virtual port gsTraffic, the destination **Destination** as port T2, and g1 for **GigaSMART Operation (GSOP)** For map3, add the following two rules.

Rule 1:

- Pass
- IPv4 Destination 65.128.7.21 Cidr 32
- IPv4 Protocol TCP
- Port Destination 443

Rule 2:

- Pass
- IPv4 Destination 98.43.132.70 Cidr 32
- IPv4 Protocol TCP
- Port Destination 443

If there are other GigaSMART applications defined in the GigaSMART operation, filtering will be done on the packets before sending the GigaSMART applications.

7. Create a shared collector map name mapC1 to direct traffic not matching the second-level maps to the tool port with the alias T3. The collector for the first-level map named map1 uses the standard collector available in prior releases.
   a. Go the Maps page and click **New**. The New Page page opens.
   b. In the **Alias** field, enter mapC1.
   c. For **Type**, select **Regular**.
   d. For **Subtype**, select **Collector**.
   e. For **Source**, select the virtual port gsTraffic.
   f. For **Destination**, select the tool port with the alias T3, and the click **Save**.

Task	Navigation Path
Clone the configurations of an existing virtual port	In the device view select **GigaSMART** > **Virtual Ports** > **Actions** > **Clone**

Edit an existing virtual port	In the device view select **GigaSMART** > **Virtual Ports** > **Actions** > **Edit**
Delete a virtual port	In the device view select **GigaSMART** > **Virtual Ports** > **Actions** > **Delete**

Multiple Virtual Ports for First Level Map

A first level map can have multiple virtual ports (vports). When multiple vports are configured on a first level map, data is sent from network ports to the multiple vports destined to specific GigaSMART groups.

The tool ports of a first level map can be a combination of vports and tool ports. Each vport is bound to a GigaSMART group (gsgroup). Multiple vports are bound to multiple gsgroups. However, in a single first level map, all the vports must be bound to different gsgroups.

NOTE: For a second level map, only a single vport can be configured.

When a second level map is configured using a vport, the data that is sent to the gsgroup is forwarded to the tool ports according to the gsrules or flow rules configured on the map.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

To configure multiple vports for a first level map:

1. Create gsgroups on the GigaSMART engine, using the following steps to create a gsgroup with the alias gsgrp1 and one with the alias gsgroup2:

a. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.

b. Click **New**.

c. Enter an alias for the GigaSMART Group in the **Alias** field.

d. Click in the **Port List** field to select an engine port.

e. Enable parameters on the GigaSMART Group as needed. For example, to enable GTP correlation, enter the Timeout value under **GTP Flow**.

f. Click **Save**.

g. Repeat steps b through f to add another GigaSMART Group.

2. Enable GTP correlation on the GigaSMART groups.

a. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.

b. Click **New**.

c. Enter a name for the GigaSMART Operation in the **Alias** field (for example gs1).

d. Select the GigaSMART groups from the **GigaSMART Groups** list.

e. Select **Flow Filtering** from the GigaSMART Operations (GSOP) list.

f. Click **Save**.

g. Repeat steps a through f to create the second GigaSMART Operation.

3. Create vports (for example, vp1 and vp2) and assign them to the gsgroups. For the steps to create a virtual ports, refer to Virtual Ports.

1. Create a **First Level By Rule** map and direct traffic to both vports and a tool port by navigating to the Maps page and clicking **New**.

a. Enter a name for the map in the **Alias** field. For example, map1.

b. Select **First Level** for **Type** and **By Rule** for **Subtype**.

c. Specify networks ports in the **Source** field.

d. Select the virtual ports and a tool port in the **Destination** field.

e. Click **Add Rule**.

f. Select **Pass** and **Port Destination** for ports 251 to 252.

g. Click **Save**.

4. Create a second level map named to direct traffic from the first vport (vp1) to the GTP correlation GigaSMART operation by navigate to the Maps page and clicking **New**.

a. Enter a name for the map in the **Alias** field. For example, map2.

b. Select **Second Level** for **Type** and **Flow Filtering** for **Subtype**.

c. Specify the first virtual port (vp1) in the **Source** field.

d. Select a tool port in the **Destination** field.

e. Click in the **GigaSMART Operation (GSOP)** field and select GSOP from the list. For example, gs1.

f. Click **Add Rule**.

g. Select **Pass** and **GTP IMSI**.

h. Enter 302720* in the IMSI field and select **Version V1**.

i. Click **Save**.

5. Create a shared collector map to direct traffic not matching the second level map to a tool port by navigating to the Maps page and clicking **New**.

   a. Enter a name for the map in the **Alias** field. For example, sc1.

   b. Select **Second Level** for **Type** and **Collector** for **Subtype**.

   c. Select the first virtual port (vp1) in the **Source** field.

   d. Select a tool port in the **Destination** field.

6. Create a second level map named **map3** to direct traffic from the second vport (vp2) to the GTP correlation GigaSMART operation by navigating to the Maps page and clicking **New**.

   a. Enter a name for the map in the **Alias** field. For example, map3.

   b. Select **Second Level** for **Type** and **Flow Filtering** for **Subtype**.

   c. Specify the second virtual port (vp2) in the **Source** field.

   d. Select a tool port in the **Destination** field.

   e. Click in the **GigaSMART Operation (GSOP)** field and select GSOP from the list. For example, gs2.

   f. Click **Add Rule**.

   g. Select **Pass** and **GTP IMSI**.

   h. Enter * in the IMSI field.

   i. Click **Save**.

Multiple Virtual Port Rules

The following rules apply to multiple virtual ports for first level maps:

1. Multiple vports with different GigaSMART groups (gsgroups) can be used on a first level map.
2. Different vports can be configured on the same GigaSMART group, but must be used in different maps.
3. A gsgroup can have multiple GigaSMART operations (gsops).
4. Only one vport is allowed on egress (second level) maps.
5. A vport on a first level map cannot be edited. To make a map change, delete the vport from the **Destination** field first and then add a new vport by clicking in the **Destination** field and selecting another vport.
6. In a standalone system, the number of vports for a first level map is limited by the number of GigaSMART engines (eports) in the chassis.
7. In a cluster environment, the number of vports for a first level map is limited by the number of eports in the cluster.

Multiple Virtual Port with Other GigaSMART Applications

The example in Multiple Virtual Ports for First Level Map uses the GTP correlation GigaSMART operation. You can also use multiple vports with other GigaSMART operations, such as Adaptive Packet Filtering (APF). You can also chain multiple GigaSMART applications. This allows you to perform different functions and filtering with the same packets.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

To configure multiple vports for a first level map:

1. Create GigaSMART Groups on the GigaSMART engines, using the following steps to create three gsgroups with the aliases gg2, gg3, and gg5 and associate with gsops gs2, gs3, and gs5, respectively.
   a. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.
   b. Click **New**.
   c. Enter an alias for the GigaSMART Group in the **Alias** field. (In this example, gg2, gg3, or gg5.)
   d. Click in the **Port List** field to select an engine port.

e. Enable parameters on the GigaSMART Group as needed. For example, to enable GTP correlation, enter the Timeout value under **GTP Flow**.

f. Click **Save**.

g. Repeat steps b through f to add another GigaSMART Group.

2. Enable APF on the GigaSMART groups, using the following steps:

   a. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.

   b. Click **New**.

   c. Enter an alias for the GigaSMART Operation in the **Alias** field. (In this example, gs2, gs3, or gs5.)

   d. Select the GigaSMART group from the **GigaSMART Groups** list. (For gs2, select gg2. For gs3, select gg3. For gs5, select gg3.)

   e. Select **APF** from the GigaSMART Operations (GSOP) list and select **Enable**.

   f. Select **Slicing** from the GigaSMART Operations (GSOP) list, and then set **Protocol** to none. For gs2 set the **offset** to80. For gs3, set the **offset** to 90. For gs5, set the **offset** to 100.

   g. Click **Save**.

   h. Repeat steps a through f to create another GigaSMART Operation.

3. Create vports and assign them to the GigaSMART Groups, using the following steps to create three virtual ports named vp2, vp3, and vp4, assigning them to gsgroup gg2, gg3, and gg5 respectively.

   a. From the device view, select **GigaSMART > Virtual Ports**.

   b. Click **New**.

   c. Enter an alias for the virtual in the **Alias** field. (In this example, vp2, vp3, or vp5.)

   d. Select the gsgroup from the **GigaSMART Group** field (gg2 for vp2, gg3 for vp2, and gg5 for vp5).

   e. Click **Save**.

   f. Repeat step b through step e to create another virtual port.

4. Create a first level map, and direct traffic to two vports and a tool port, using the following steps:

   a. Select **Maps > Maps > Maps**.

   b. Click **New**.

   c. Enter m1 in the **Alias** field.

d. Select **First Level** for **Type** and **By Rule** for **Subtype**.

e. Select a network port in the **Source** field.

f. Select virtual ports vp2 and vp3 plus a tool port in the **Destination** field.

g. Use the **Add a Rule** button to add two rules to the map.

For the first rule, select **pass** and select **VLAN** for the condition. Set the VLAN value to 100.

For the second rule, select **pass** and select **VLAN** for the condition. Set the VLAN value to 200.

h. Click **Save**.

5. Create a second level map named to direct traffic from the vports to the GigaSMART operation, using the following steps:

a. Select **Maps > Maps > Maps**.

b. Click **New**.

c. Enter m2 in the **Alias** field.

d. Select **Second Level** for **Type** and **By Rule** for **Subtype**.

e. Select a virtual port vp2 the **Source** field.

f. Select a tool port in the **Destination** field.

g. Select gs2 for **GigaSMART Operation (GSOP)**.

h. Use the **Add a Rule** button to add a rules to the map.

For the rule, select **pass** and select **VLAN** for the condition. Set the VLAN value to 100 with **Subset** set to none and **Position** set to 0.

i. Click **Save**.

6. Create another second level map named to direct traffic from the vports to the GigaSMART operation, using the following steps:

a. Select **Maps > Maps > Maps**.

b. Click **New**.

c. Enter m4 in the **Alias** field.

d. Select **Second Level** for **Type** and **By Rule** for **Subtype**.

e. Select a virtual port vp3 in the **Source** field.

f. Select a tool port in the **Destination** field.

g. Select gs3 for **GigaSMART Operation (GSOP)**.

h. Use the **Add a Rule** button to add a rules to the map.

For the rule, select **pass** and select **VLAN** for the condition. Set the VLAN value to 200 with **Subset** set to none and **Position** set to 0.

i. Click **Save**.

7. Create another second level map to direct traffic from the vports to the GigaSMART operation.
   a. Select **Maps > Maps > Maps**.
   b. Click **New**.
   c. Enter m6 in the **Alias** field.
   d. Select **Second Level** for **Type** and **By Rule** for **Subtype**.
   e. Select a virtual port vp3 in the **Source** field.
   f. Select a tool port in the **Destination** field.
   g. Select gs3 for **GigaSMART Operation (GSOP)**.
   h. Use the **Add a Rule** button to add a rules to the map.

      For the rule, select **pass** and select **VLAN** for the condition. Set the VLAN value to 200 with **Subset** set to none and **Position** set to 0.

   i. Click **Save**.
8. Create a shared collector map named to direct traffic not matching the maps to a tool port.
   a. Select **Maps > Maps > Maps**.
   b. Click **New**.
   c. Enter sc1 in the **Alias** field.
   d. Select **Second Level** for **Type** and **Collector** for **Subtype**.
   e. Select a virtual port vp1 in the **Source** field.
   f. Select a tool port in the **Destination** field.
   g. Click **Save**.

Virtual Port Statistics

The Virtual Ports page displays statistics about the configured virtual ports. To view the statistics select **GigaSMART > Virtual Ports > Statistics**.

The following table describes virtual port statistics:

Statistic	Description
Rx Packets	The number of packets received into the virtual port.
Tx Packets	The number of packets transmitted out of the virtual port.
Rx Octets	The number of bytes received into the virtual port.
Tx Octets	The number of bytes transmitted out of the virtual port.
Packets Drops	The number of packets dropped at the virtual port.
Packet Drops No Init	For internal debugging.

Differences in GigaSMART Nomenclature Between the CLI and GigaVUE-FM

The CLI and the Web-basedGigaVUE-FM interface occasionally use different names to refer to the same GigaSMART-related functionality. The following table summarizes the differences:

Documentation Term	CLI Term	GigaVUE-FM Term	Description
GigaSMART Port	e1, e2	GigaSMART Engine	Internal ports on GigaSMART line card, module used to power GigaSMART features.
GigaSMART Group	gsgroup	GigaSMART Engine Group	Group of internal ports on GigaSMART line card, module used to power GigaSMART features. You configure these in the CLI with the **gsgroup** command.
GigaSMART Component	gsop	GigaSMART Operation	One of the available GigaSMART packet processing features (de-duplication, slicing, and so on)
GigaSMART Operation	gsop	GigaSMART Operation Group	A combination of GigaSMART packet-processing features into a single entity used in a map. You configure these in the CLI with the **gsop** command.

GigaSMART Operations in Clusters

Clustered environments with at least one GigaSMART-enabled chassis can take advantage of GigaSMART operations in maps on network ports elsewhere in the cluster. As shown in Figure 112GigaSMART Group on Different Chassis than Network/Tool Ports and Figure 113GigaSMART Operations in Clusters, the GigaSMART group providing the packet processing power for a GigaSMART operation does not have to be on the same chassis as the network or tool ports for a map.

- A map's network ports, tool ports, and the GigaSMART group ports can all be on different nodes in a clustered environment.
- GigaSMART operations do not require any specific role for the chassis with the GigaSMART group – it can be a leader, standby, or normal node.
- The main requirements to keep in mind is that the GigaSMART ports in a GigaSMART group must all be on the same chassis. So, for example, if you had two GigaSMART-HD0 line cards in a single chassis, you could create a single GigaSMART group out of the four GigaSMART engine ports available across the two line cards. However, if the two line cards were in two different chassis, they could not be combined into a single GigaSMART group.

NOTE: Inline SSL is not supported in clusters nor on any nodes that are part of a cluster. Do not attempt to enable inline SSL on individual nodes that are part of a cluster or have inline networks and inline tools distributed among various nodes in a cluster.

Figure 112GigaSMART Group on Different Chassis than Network/Tool Ports and Figure 113GigaSMART Operations in Clusters illustrate examples of network, tool, and GigaSMART group ports on different nodes in a clustered environment. For example, Figure 112GigaSMART Group on Different Chassis than Network/Tool Ports shows a map accepting ingress packets on GV1, sending them to the GigaSMART group on GV3 for GigaSMART processing (for example, de-duplication), and sending the results to a tool port on GV4.

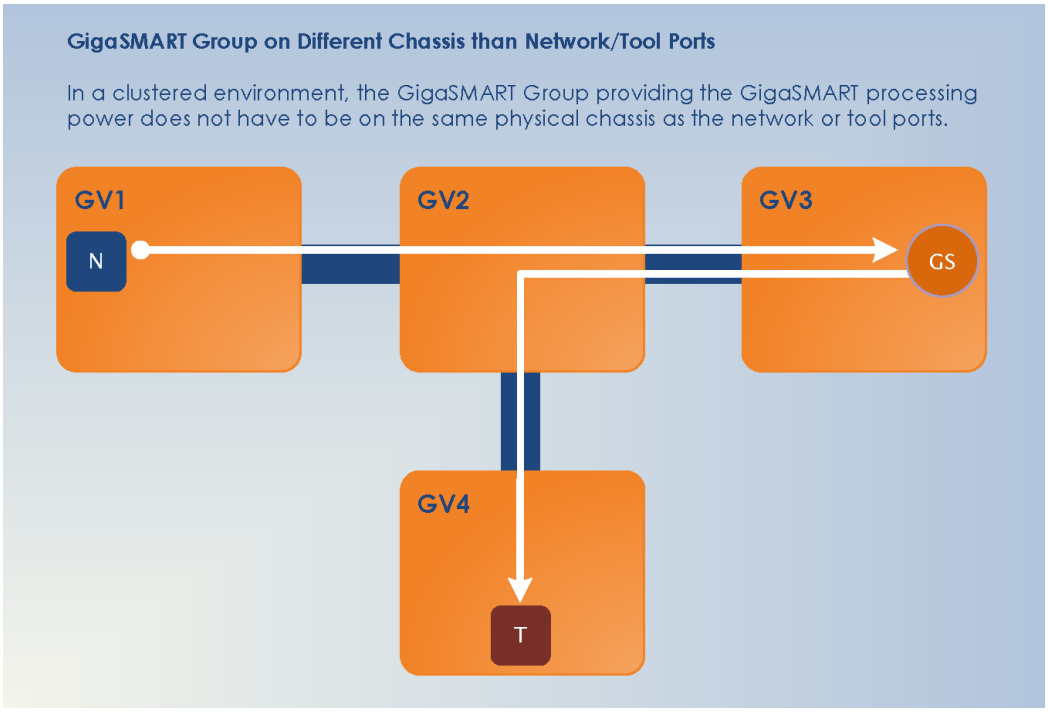


Figure 112 *GigaSMART Group on Different Chassis than Network/Tool Ports*

You could also have the GigaSMART group on the same chassis as either the network or tool ports for the map. An example of this is shown in Figure 113GigaSMART Operations in Clusters. The GigaSMART group on GV3 is on the same chassis as the egress port.

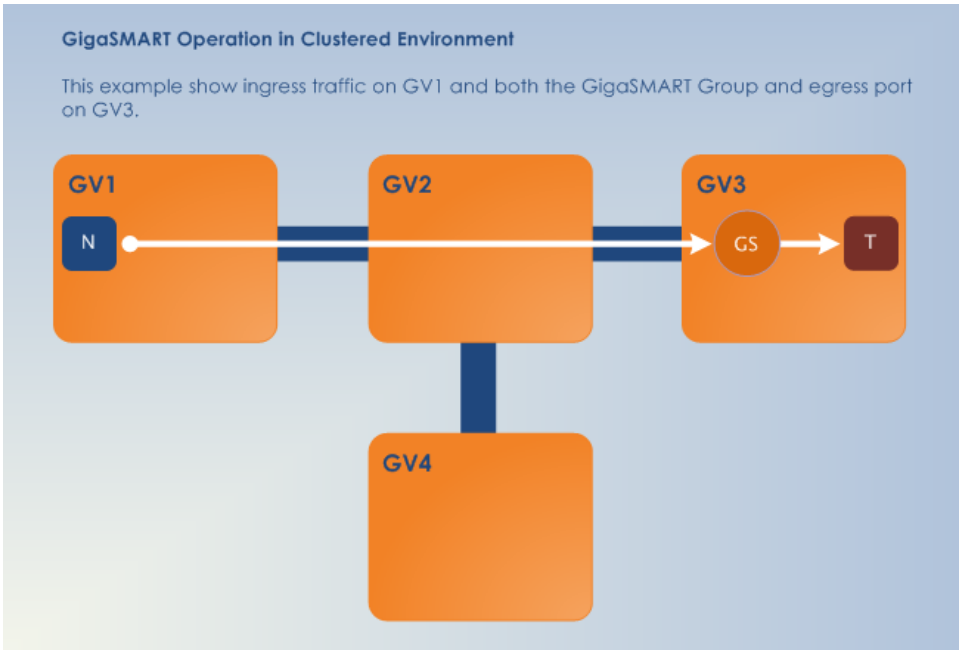


Figure 113 *GigaSMART Operations in Clusters*

How to Combine GigaSMART Operations

You can combine multiple GigaSMART components into a single operation. For example, you can set up a single GigaSMART operation that masks a packet, strips its VLAN header, and inserts a field in the GigaSMART Trailer.

The following table summarizes the valid combinations of GigaSMART operations.

GigaSMART Operations Compatibility Chart

Supported GSOP Combinations		Masking	Trailer Addition (Source Port Labeling)	Trailer Removal (Source Port Labeling)	De- Dup	Load Balance	Adaptive Packet Filtering (APF)	App Intelligence		Flow-Ops				Headers			NetFlow/IPFIX[1]		Slicing		SSL /TLS			Tunnels	
								AFI (APF+ASF)	AMI[2]	FlowVUE	GTP Flow Filter	GTP Whitelist	GTP Flow Sampling	Strip Headers	Add Headers	Remove H/T	NetFlow 1st level	NetFlow 2nd level	Slicing	Advanced Flow Slicing	SSL Decrypt	iSSL	ICAP	Tunnel Encap	Tunnel Decap
Masking			S	S	S	S	S	S	NS	S	NS	NS	NS	S	S	S	NS	NS	S	S	NS	NS	NS	S	S
Trailer Addition (Source Port Labeling)		S			S	S	S	NS	NS	S	NS	NS	NS	S	S	S	NS	NS	S	S	NS	NS	NS	S	S
Trailer Removal (Source Port Labeling)		S			S	S	S	NS	NS	S	NS	NS	NS	S	S	S	NS	NS	NS	NS	NS	NS	NS	S	S
De-Dup		S	S	S		S	S	S	S	S	NS	NS	NS	S	S	S	S	S	S	S	S	NS	NS	S	S
Load Balance		S	S	S	S		S	S	NS	S	S	S	S	S	S	S	NS	NS	S	S	NS	NS	NS	S	S
Adaptive Packet Filtering	APF	S	S	S	S	S		NA	NS	S	NS	NS	NS	S	S	NS	NS	NS	S	S	NS	NS	NS	S	NS
App Intelligence	AFI (APF+ASF)	S	NS	NS	S	S	NA		S	NS	NS	NS	NS	NS	NS	NS	NS	S	S	S	NS	NS	NS	NS	NS
	AMI[2]	NS	NS	NS	S	NS	NS	S		NS	NS	NS	NS	NS	NS	NS	NA	NA	NS	NS	NS	NS	NS	NS	NS
Flow- Ops	FlowVUE	S	S	S	S	S	S	NS	NS		NS	NS	NS	S	S	S	NS	NS	S	NS	NS	NS	NS	NS	NS
	GTP Flow Filter	NS	NS	NS	NS	S	NS	NS	NS	NS		S	NS	NS	NS	NS	NS	NS	S	NS	NS	NS	NS	NS	NS
	GTP Whitelist	NS	NS	NS	NS	S	NS	NS	NS	NS	S		NS	NS	NS	NS	NS	NS	S	NS	NS	NS	NS	NS	NS
	GTP Flow Sampling	NS	NS	NS	NS	S	NS	NS	NS	NS	NS	NS		NS	NS	NS	NS	NS	S	NS	NS	NS	NS	NS	NS
Headers	Strip Headers	S	S	S	S	S	S	NS	NS	S	NS	NS	NS		S	S	NS	NS	S	S	NS	NS	NS	S	S
	Add Headers	S	S	S	S	S	S	NS	NS	S	NS	NS	NS	S		S	NS	NS	S	S	NS	NS	NS	NS	S
	Remove H/T	S	S	S	S	S	NS	NS	NS	S	NS	NS	NS	S	S		NS	NS	NS	NS	NS	NS	NS	S	NS
NetFlow/IPFIX[1]	NetFlow 1st level	NS	NS	NS	S	NS	NS	NS	NA	NS	NS	NS	NS	NS	NS	NS		NS	NS	NS	NS	NS	NS	NS	NS
	NetFlow 2nd level	NS	NS	NS	S	NS	NS	S	NA	NS	NS	NS	NS	NS	NS	NS	NS		NS	NS	NS	NS	NS	NS	NS
Slicing	Slicing	S	S	NS	S	S	S	S	NS	S	S	S	S	S	S	NS	NS	NS		S	NS	NS	NS	S	S
	Advanced Flow Slicing	S	S	NS	S	S	S	S	NS	S	S	S	S	S	S	NS	NS	NS	S		NS	NS	NS	S	S
SSL/ TLS	SSL Decrypt	NS	NS	NS	S	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS		NS	NS	NS	NS
	iSSL	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS		NS	NS	NS
	ICAP	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS		NS	NS
Tunnels	Tunnel Encap	S	S	S	S	S	S	NS	NS	NS	NS	NS	NS	S	NS	S	NS	NS	S	S	NS	NS	NS		NS
	Tunnel Decap	S	S*	S	S	S	NS	NS	NS	NS	NS	NS	NS	S	S	NS	NS	NS	S	S	NS	NS	NS	NS	

S = Supported; S* = Supported for Gen 3 cards only; NS = Not Supported; NA = Not Applicable

1 = NetFlow/IPFIX is specific to Gen2 GigaSMART card

2 = AMI includes support for application metadata and NetFlow/IPFIX

NOTE: Masking operations cannot be combined with slicing unless the offset of the slicing is after the offset of the masking.

How to Read GigaSMART Operations Table

The GigaSMART operations table is read both across and down. The following is an example of how to read the GigaSMART operations table:

1. Begin in the left-most column and select a GigaSMART operation, for example: Add Header.
2. Move to the right along the Add Header row. Add Header can be combined with Slicing or Masking, Source Id, Header/Trailer Remove, De-duplication, Tunnel decapsulation, and Strip Header. It cannot be combined with Tunnel Encap.
3. Move to the right another square to the gray square at end of the Add Header row. This is the Add Header column.
4. Move down the Add Header column, below the gray square at the end of the Add Header row. Add Header can be combined with Flow-Ops (FlowVUE), APF, ASF, and Load Balance. It cannot be combined with Flow-Ops (Flow-Filter GTP), Flow-Ops (GTP Whitelist), Flow-Ops (GTP Flow Sampling), NetFlow (1st Level Maps), NetFlow (2nd Level Maps), SSL Decryption (for Out-of-Band Tools), Inline-SSL (SSL Decryption for Inline Tools), Flow-Ops (SIP Flow Sampling), or Flow-Ops (SIP Flow Whitelist).

Work with GigaSMART Operation Combinations in GigaVUE-FM

When combining GigaSMART operations in GigaVUE-FM, the drop down option in the GSOP screen will gray out if a set of combinations is invalid.

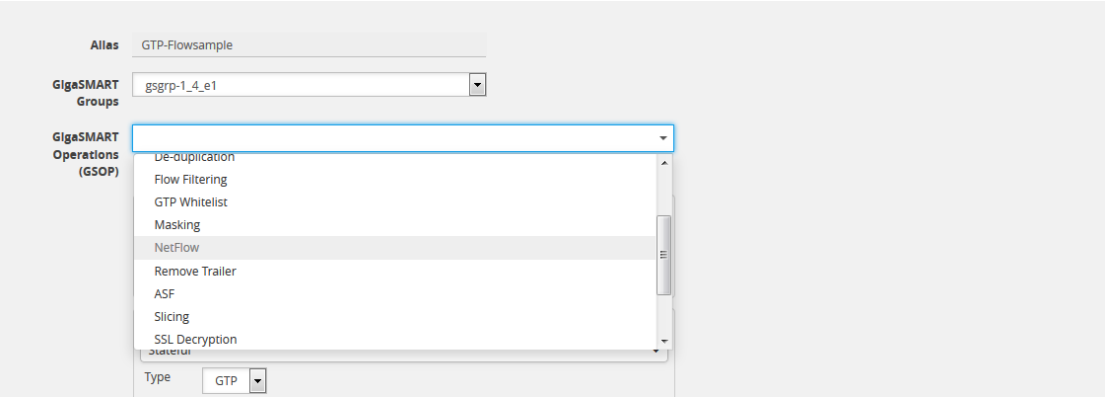


Figure 114 *GigaSMART Operation Combinations Not Available*

However, in certain cases you may not see the combinations grayed out but when trying to save the combination a pop-up is displayed stating that the combination is invalid.

To get the valid combinations, refer to the table at the beginning of this section that summarizes the combinations.

Supported GigaSMART Operations

The following topics lists the supported GigaSMART operations by GigaVUE physical and virtual nodes.

- Supported GigaSMART Operations on GigaVUE HC Series
- Supported GigaSMART operations on GigaVUE V Series Node

Supported GigaSMART Operations on GigaVUE HC Series

NOTE: An initialization/deinitialization delay is encountered while configuring or removing configuration of a Passive SSL GSOP in new the Generation 3 GigaSMART card.

GigaSMART Operation	GigaVUE-HC1		GigaVUE-HC3	GigaVUE-HC1 Plus		GigaVUE-HCT
	Gen 2 On-board	Gen 3 SMT-HC1-S	Gen 3 (C08)	Rear Gen3 SMT-HC1A-R	Front Gen 3 SMT-HC1-S	Gen 3 SMT-HC1-S
GigaSMART Masking	✓	✓	✓	✓	✓	✓
GigaSMART Packet Slicing	✓	✓	✓	✓	✓	✓
GigaSMART Advanced Flow Slicing	✓	✓	✓	✓	✓	✓
Source ID (add Trailer)	✓	✓	✓	✓	✓	✓
Header/Trailer Remove	✓	✓	✓	✓	✓	✓

GigaSMART Operation	GigaVUE-HC1		GigaVUE-HC3	GigaVUE-HC1 Plus		GigaVUE-HCT
	Gen 2 On-board	Gen 3 SMT-HC1-S	Gen 3 (C08)	Rear Gen3 SMT-HC1A-R	Front Gen 3 SMT-HC1-S	Gen 3 SMT-HC1-S
GigaSMART 5G CUPS	✗	✗	✓	✗	✗	✗
GigaSMART De-Duplication	✓	✓	✓	✓	✓	✓
L2GRE Tunnel Encapsulation	✓	✓	✓	✓	✓	✓
GigaSMART VXLAN Tunnel Encapsulation	✓	✓	✓	✓	✓	✓
L2GRE Tunnel Decapsulation	✓	✓	✓	✓	✓	✓
GigaSMART VXLAN Tunnel Decapsulation	✓	✓	✓	✓	✓	✓
GigaSMART ERSPAN Tunnel Decapsulation	✓	✓	✓	✓	✓	✓
GigaSMART IP Encapsulation (GigaSMART Tunnel)	✓	✗	✗	✗	✗	✗
GigaSMART IP Decapsulation	✓	✗	✗	✗	✗	✗
GigaSMART Header Stripping	✓	✓	✓	✓	✓	✓
GigaSMART Header Addition	✓	✓	✓	✓	✓	✓
GigaSMART IP FlowVUE	✓	✓	✓	✓	✓	✗
GTP Flow Filtering	✗	✗	✗	✗	✗	✗
GigaSMART Rotational Sampling	✗	✗	✓	✓	✓	✗
GTP Whitelisting (GigaSMART Rotational Sampling)	✗	✗	✓ Non-CUPS GTP Correlation is supported. LTE	✓	✓	✗

GigaSMART Operation	GigaVUE-HC1		GigaVUE-HC3	GigaVUE-HC1 Plus		GigaVUE-HCT
	Gen 2 On-board	Gen 3 SMT-HC1-S	Gen 3 (C08)	Rear Gen3 SMT-HC1A-R	Front Gen 3 SMT-HC1-S	Gen 3 SMT-HC1-S
			CUPS is released as BETA* in 6.1.			
GTP Flow Sampling (GigaSMART Rotational Sampling)	✗	✗	✓ Non-CUPS GTP Correlation is supported. LTE CUPS is released as BETA* in 6.1.	✓	✓	✗
GigaSMART Adaptive Packet Filtering (APF) **NOTE:** Generation 3 nodes do not support RegEx masking.	✓	✓	✓	✓	✓	✓
Application Session Filtering (ASF) and Buffer ASF	✓	✓	✓	✓	✓	✓
Application Filtering Intelligence (AFI)	✓	✓	✓	✓	✓	✓
Application Metadata Intelligence (AMI)	✓	✓	✓	✓	✓	✓
GigaSMART NetFlow Generation (Application Intelligence)	✓	✓	✓	✓	✓	✓
Application Visualization	✓	✓	✓	✓	✓	✓
GigaSMART NetFlow Generation (Traffic Intelligence)	✓	✗	✗	✗	✗	✗
GigaSMART Load Balancing (Stateless)	✓	✓	✓	✓	✓	✓
GigaSMART Load Balancing (Stateful)	✓	✓	✓	✓	✓	✓

GigaSMART Operation	GigaVUE-HC1		GigaVUE-HC3	GigaVUE-HC1 Plus		GigaVUE-HCT
	Gen 2 On-board	Gen 3 SMT-HC1-S	Gen 3 (C08)	Rear Gen3 SMT-HC1A-R	Front Gen 3 SMT-HC1-S	Gen 3 SMT-HC1-S
SSL Decryption for Out-of-Band Tools (Passive SSL)	✓	✓	✓	✓	✓	✓
SSL Decryption for Inline Tools	✓	✓	✓	✓	✓	×
SIP Flow Sampling	×	×	×	✓	✓	×
SIP Flow Whitelist	×	×	×	✓	✓	×
4G/5G Traffic Monitoring using UPN	×	×	✓	×	×	×
GigaSMART TCP tunnel	✓	×	×	×	×	×
Secure Tunnels Secure Tunnels	×	✓	✓	✓	✓	✓
Secure Tunnels Secure Tunnels	×	✓	✓	✓	✓	✓

*BETA features are available for functional Proof of Concepts only. Scalability and Performance limits will available when the feature is released for formal deployments.

NOTE: The combination of Application Visualization, Application Filtering Intelligence (AFI), De-duplication, and Application Metadata Intelligence (AMI) is not supported for GigaVUE-HC1 Gen2 platform.

Supported GigaSMART operations on GigaVUE V Series Node

GigaSMART Operation	GigaVUE Cloud Suite for AWS	GigaVUE Cloud Suite for Azure	GigaVUE Cloud Suite for OpenStack	GigaVUE Cloud Suite for VMware (VMware vCenter)	GigaVUE Cloud Suite for VMware (NSX-T)	GigaVUE Cloud Suite for Nutanix
Masking	✓	✓	✓	✓	✓	✓
Slicing	✓	✓	✓	✓	✓	✓
De-duplication	✓	✓	✓	✓	✓	✓
Application Metadata Exporter	✓	✓	✓	✓	✓ (Only when deploying GigaVUE V Series Node using Third party Orchestration)	✓
L2GRE Tunnel Encapsulation Refer to *Create Ingress and Egress Tunnels* section on the retrospective GigaVUE Cloud Suite Deployment Guide.	✓	✗	✓	✓	✓	✓
VXLAN Tunnel Encapsulation Refer to *Create Ingress and Egress Tunnels* section on the retrospective GigaVUE Cloud Suite Deployment Guide.	✓	✓	✓	✓	✓	✓
L2GRE Tunnel Decapsulation	✓	✗	✓	✓	✓	✓

GigaSMART Operation	GigaVUE Cloud Suite for AWS	GigaVUE Cloud Suite for Azure	GigaVUE Cloud Suite for OpenStack	GigaVUE Cloud Suite for VMware (VMware vCenter)	GigaVUE Cloud Suite for VMware (NSX-T)	GigaVUE Cloud Suite for Nutanix
Refer to *Create Ingress and Egress Tunnels* section on the retrospective GigaVUE Cloud Suite Deployment Guide.						
VXLAN Tunnel Decapsulation Refer to *Create Ingress and Egress Tunnels* section on the retrospective GigaVUE Cloud Suite Deployment Guide.	✓	✓	✓	✓	✓	✓
ERSPAN Tunnel Decapsulation Refer to *Create Ingress and Egress Tunnels* section on the retrospective GigaVUE Cloud Suite Deployment Guide.	✓	✗	✓	✓	✓	✓
UDPGRE Tunnel Decapsulation Refer to *Create Ingress and Egress Tunnels* section on the retrospective GigaVUE Cloud Suite Deployment	✓	✗	✓	✓	✓	✗

GigaSMART Operation	GigaVUE Cloud Suite for AWS	GigaVUE Cloud Suite for Azure	GigaVUE Cloud Suite for OpenStack	GigaVUE Cloud Suite for VMware (VMware vCenter)	GigaVUE Cloud Suite for VMware (NSX-T)	GigaVUE Cloud Suite for Nutanix
Guide.						
GENEVE Decapsulation	✓	×	×	×	✓	×
Header Stripping	✓	✓	✓	✓	✓	✓
Adaptive Packet Filtering (APF) without RegEx	✓	✓	✓	✓	✓	✓
Application Session Filtering (ASF)	✓	✓	✓	✓	✓	✓
Application Filtering Intelligence	✓	✓	✓	✓	✓	✓
Application Metadata Intelligence	✓	✓	✓	✓	✓	✓
GigaSMART NetFlow Generation	✓	✓	✓	✓	✓	✓
Application Visualization	✓	✓	✓	✓	✓	✓
Load Balancing	✓	✓	✓	✓	✓	✓
SSL Decrypt	✓	✓	✓	✓	×	✓
5G-Service Based Interface Application	×	×	✓	✓	✓	×
5G-Cloud Application	×	×	✓	✓	×	×
Secure Tunnels Encapsulation	×	✓	×	✓	✓	✓
Secure Tunnels Decapsulation	×	✓	×	✓	✓	✓
GRE-In-UDP Tunnel Decapsulation	×	×	✓	✓	×	×

Order of GigaSMART Operations

When combining multiple GigaSMART components into a single operation, the components are applied in the following order:

1. [apf/asf/buffer asf] OR [flow-ops <flow-filtering>] or [flow-ops <gtp-whitelist>] or [flow-ops <gtp-flowsampling>] AND [lb] (Stateful load balancing with GTP or ASF)
2. [apf/asf/buffer asf (AFI)]
3. [tunnel-decap]
4. [dedup]
5. [ssl-decrypt]
6. [flow-ops <flow-sampling>]
7. [strip-header]
8. [slicing]
9. [enhanced slicing]
10. [masking]
11. [trailer]
12. [tunnel-encap]
13. [add-header]
14. [flow-ops <netflow>]
15. [lb] (Stateless load balancing)
16. metadata (AMI)

View GigaSMART Statistics

To view the GigaSMART statistics from the UI, select **GigaSMART Operations (GSOP) > Statistics**. The Statistics page displays as shown in Figure 115GigaSMART Operations Statistics Page.

GSOP Alias	GS Group Alias	Map Alias	Rx Packets	Tx Packets	Rx Octets	Tx Octets	Packet Drops	Packet Drops No Init	Packets Terminated	Packets Parse Errors	GS Operations
GigavueVM_Tunnel	gsgrp-1_4_e1	OpenStack_vTraffic_toWireshark	3.23 M	60	814.48 M	7.26 K	3.23 M	0	0	0	Tunnel Decap
gtp_flow_sampling_elias_test	gsgrp-1_4_e1	GTP-Sampling-2	0	0	0	0	0	0	0	0	Flow Sampling

Figure 115 *GigaSMART Operations Statistics Page*

The Statistics page shows the aliases for GS Operation, GS Group, the alias of the map that is using GS Operation, and the GS Operations being used. In Figure 115GigaSMART Operations Statistics Page, Flow Sampling and Load Balance operations are assigned to the alias GTP-Flowsample in the gsgrp-1_4_e1 GS Group. The GS operations is used by the map with the alias map-gtpFS. For a description of the other columns, refer to GigaSMART Operations Statistics Definitions.

Other statistics available for viewing from the UI are the following:

- GigaSMART Groups Statistics (select **GigaSMART > Statistics**)
- NetFlow / IPF Generation Statistics for Exporters and Monitors (select **NetFlow / IPFIX > Exporter Statistics or NetFlow / IPFIX Generation > Monitor Statistics**)
- IP Interface Statistics (select **Ports > IP Interfaces > Statistics**)

NOTE: When tunnel load balancing is configured in Generation 3 GigaSMART card (SMT-HC1-S) with multiple maps using the same port-group, the GSOP statistics is the cumulative statistics of all the maps as the individual GSOP level statistics are not maintained in Generation 3 GigaSMART card (SMT-HC1-S).

Definitions of GigaSMART Statistics

The following sections provide definitions for the statistics displayed. Refer to the following:

- *NetFlow Monitor Statistics Definitions*
- *NetFlow Exporter Statistics Definitions*
- IP Interfaces Statistics Definitions
- GigaSMART Group Statistics Definitions
- GigaSMART Group Flow Ops Report Statistics Definitions
- GigaSMART Operations Statistics Definitions

NetFlow Monitor Statistics Definitions

To view NetFlow Monitor statistics, select **GigaSMART > NetFlow / IPFIX Generation > Monitor Statistics** to open the Monitor Statistics page.

The following table describes NetFlow Monitor statistics displayed on the Monitor Statistics page:

Statistic	Description
No Entries	The current number of flows in the monitor cache.
High Watermark	The maximum number of flows that have ever been in the monitor cache.
Flows Added	The sum of all flows added to the monitor cache.
Flows Aged	The sum of the flows that have aged due to the following: o Active Timeout—The configured active timeout for the monitor cache was exceeded. o Inactive Timeout—The configured inactive timeout for the monitor cache was exceeded. o Event Aged—The number of entries aged from the cache because a TCP FIN/RST flag was received. o Watermark Aged—The number of entries aged from the cache because the CPU utilization of the cache exceeded 75%. o Emergency Aged—The number of entries aged from the cache because a user requested a forced flush through the CLI.
No of Active Timeout	The number of times the configured active timeout for the monitor cache was exceeded.
No of Inactive Timeout	The number of times the configured inactive timeout for the monitor cache was exceeded.
No of Event Aged	The number of entries aged from the cache because a TCP FIN/RST flag was received.
No of Watermarked Aged	The number of entries aged from the cache because the CPU utilization of the cache exceeded 75%.
No of Emergency Aged	The number of entries aged from the cache because a user requested a forced flush through the CLI.

NetFlow Exporter Statistics Definitions

To view NetFlow Exporter statistics, select **GigaSMART > NetFlow / IPFIX Generation > Monitor Statistics** to open the Exporter Statistics page.

The following table describes NetFlow Exporter statistics:

Statistic	Description
Templates Added	The number of data templates added to the exporter.
Records Added	The number of data records added to the exporter.
Filtered Records Removed	The number of records filtered by exporter filters
Empty Records Not Added	The number of records not added to NetFlow because they were empty or blank. In the NetFlow record, if all the collect fields contain only enterprise extensions such as URL, HTTP, or DNS, and if during run-time, the records are blank or empty, they will be counted as Empty Records Not Added.
Packets Sent	The number of packets sent from the exporter to the collector.
Packet Dropped	The number of packets dropped, which could be due to the inability to send packets, such as there is no network connection or the port is down.
Packet Transmit Errors	The number of packets dropped, which could be due to the inability to send packets, such as there is no network connection or the port is down.

IP Interfaces Statistics Definitions

To view IP Interfaces statistics, select **Ports > IP Interfaces > Statistics** to open the IP Interfaces Statistics page. The statistics of the control traffic such as ARP, ICMP, and ICMPv6 for the physical node will be displayed. Use the **Clear** button to remove the selected IP interface statistics. Use the **Clear All** button to remove all the IP interface statistics.Figure 116IP Interfaces Statistics shows an example.

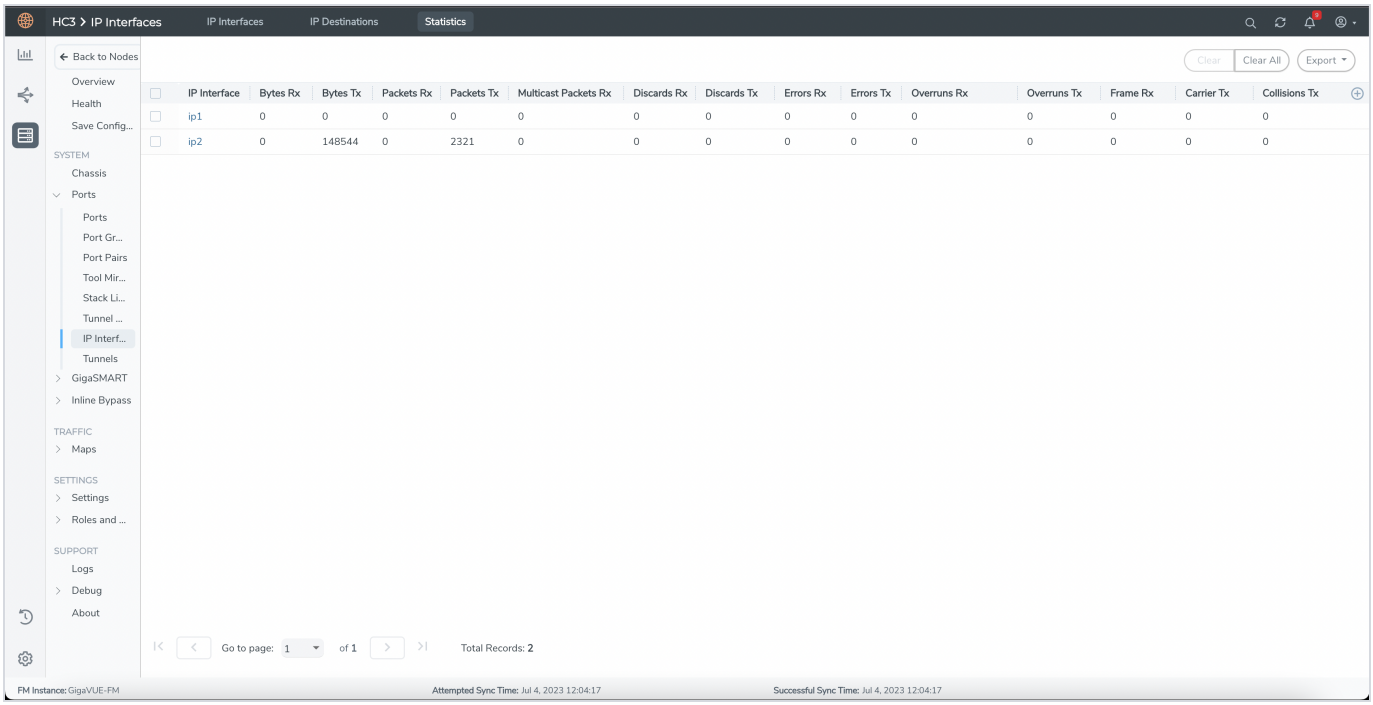


Figure 116 *IP Interfaces Statistics*

The following table describes IP interfaces statistics:

Statistic	Description
IP Interface	The IP interface for which the statistics is displayed.
Bytes Rx	The number of bytes received into the IP interface.
Bytes Tx	The number of bytes transmitted out of the IP interface.
Packets Rx	The number of packets received into the IP interface.
Packets Tx	The number of packets transmitted out of the IP interface.
Multicast Packets Rx	The number of multicast packets received into the IP interface.
Discards Rx	The number of packets that were received and discarded by the IP interface.
Discards Tx	The number of packets that were transmitted out of the IP interface and were discarded.
Errors Rx	The number of packets with errors that were received into the IP interface.
Errors Tx	The number of packets with errors that were transmitted out of the IP interface.

Statistic	Description
Overruns Rx	The number of first-in-first-out (FIFO) buffer errors that were received into the IP interface.
Overruns Tx	The number of first-in-first-out (FIFO) buffer errors that were transmitted out of the IP interface.
Frame Rx	The number of frames received into the IP interface.
Carrier Tx	The number of packets in which carrier losses were detected when transmitted out of the IP interface.
Collisions Tx	The number of packets that were colliding when transmitted out of the IP interface.

TLS/SSL Application Statistics Definitions

The following table describes TLS/SSL application statistics:

Statistic	Description
active_sessions	The number of current active sessions.
num_session_ids_cached	The number of current session IDs that have been cached.
num_tls_tickets_cached	The number of current TLS tickets that have been cached.
total_sessions	The cumulative total number of sessions.
handshaked_sessions	The number of sessions that passed TLS/SSL handshake between a client and server to establish a session.
failed_sessions	The cumulative total number of failed sessions.
num_session_resumed_ session_id	The total number of resumed sessions with a session ID.
num_session_resumed_ticket	The total number of resumed sessions with TLS tickets.
num_session_timedout	The number of TLS/SSL sessions that have timed out.
avg_in_pkts_per_session	The average number of input packets per session.
avg_out_pkts_per_session	The average number of output packets per session.

ASF Statistics Definitions

The following table describes ASF statistics:

Statistic	Description
Packet In	The number of packets received by the ASF application.
Session Created	The number of sessions created.
Session Freed	The number of sessions deleted.
Session Timeout	The number of sessions that were deleted due to inactivity (expiry of the session timer).
Packet Match Session	The number of incoming packets matching ASF sessions. This count does not include the packets that triggered the creation of the session.
Packets Buffered	The number of packets stored by buffer ASF prior to the APF match.
APF Match Pass	The number of packets matching a pass GigaSMART rule.
APF Match Drop	The number of packets matching a drop GigaSMART rule.
Packet Pass	The number of packets sent to tool ports due to a pass GigaSMART rule. (This includes packets belonging to pass sessions for ASF.)
Packet Drop	The number of packets dropped due to a drop GigaSMART rule. (This includes packets belonging to drop sessions for ASF.)
Packet No Match	The number of packets that did not match any GigaSMART rule. (This includes packets belonging to non-pass/drop sessions for ASF.)
Packets masking failures in this Map	Increases when a packet matches a regex rule with masking enabled, but the configured masking range (start/end offsets or length) does not result in a valid area to overwrite. The counter is updated at most once per packet, even if multiple regex gsrules in the same map match that packet.
Buffered Packet Pass	The number of buffered packets whose sessions matched a pass GigaSMART rule.
Buffered Packet Drop	The number of buffered packets whose sessions matched a drop GigaSMART rule.
Buffered Packet No Match	The number of buffered packets whose sessions did not

Statistic	Description
	match any GigaSMART rule.
Session Exceed Buffer Count	The number of sessions without an APF match after buffering that exceeded the configured buffer count.
Packet Exceed Buffer Count	The number of packets belonging to sessions that exceeded the configured buffer count.
Out of Buffer	The number of times GigaSMART ran out of packet buffer when trying to store a packet in a session.
Out of Session	The number of times GigaSMART ran out of session entries when trying to create a session.

GigaSMART Group Statistics Definitions

To view GigaSMART Group statistics, select **GigaSMART > GigaSMART Groups > Statistics** to open the Statistics page. Figure 117GigaSMART Group Statistics shows an example.

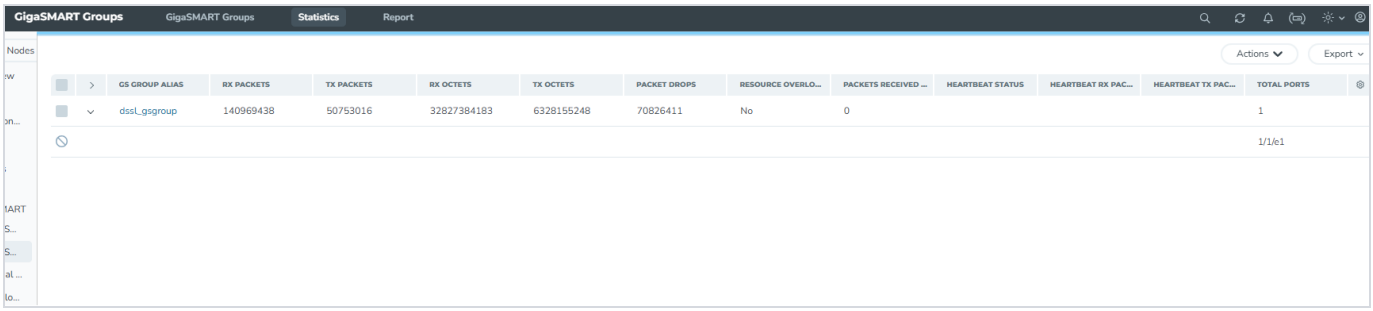


Figure 117 *GigaSMART Group Statistics*

The following table describes GigaSMART group statistics:

Statistic	Description
Rx Packets	The cumulative number of packets coming into a GigaSMART group.
Tx Packets	The cumulative number of packets going out of a GigaSMART group.
Rx Octets	The cumulative number of bytes coming into a GigaSMART group.
Tx Octets	The cumulative number of bytes going out of a GigaSMART group.
Packet Drops	The cumulative number of packets dropped at a GigaSMART group.
Packet Received	The cumulative number of received packets with errors at a

Statistic	Description
Errors	GigaSMART group.
Heartbeat Status of eport	The status of heartbeat. The valid values are up, down, or NA. This heartbeat status is for a GigaSMART engine port on a GigaSMART group for inline TLS/SSL decryption On a supported platform, such as GigaVUE-HC3, the status will be either up or down, followed by number of packets sent and received. On an unsupported platform, the status will be NA and the number of packets will not be displayed. The heartbeat status for the GigaSMART cards is fully operational only when it is configured with the inline TLS/SSL application. In all other cases, the default value of the heartbeat status is Up and the status is generally ignored.
Heartbeat Rx Packets	The heartbeat receive packet count.
Heartbeat Tx Packets	The heartbeat transmit packet count.

NOTE: The K, M, G and T in the statistics page denote KiloBytes, MegaBytes, GigaBytes and TeraBytes, respectively. You can click on the GSOP alias to view the detailed statistics.

GigaSMART Group Flow Ops Report Statistics Definitions

The following sections provide definitions for the statistics that are reported for the GigaSMART group. Refer to the following:

- Flow Ops Report Statistics Definitions for FlowVUE
- Flow Ops Report Statistics Definitions for GTP
- *Flow Ops Report Statistics Definitions for GTP Overlap*
- Flow Ops Report Definitions for SIP/RTP Correlation
- Flow Ops Report Statistics for Passive TLS/SSL Decryption

Flow Ops Report Statistics Definitions for FlowVUE

The following table describes Flow Ops report statistics for FlowVUE:

Statistic	Description
From GigaVUE-OS-CLI	
Device IP	The IP address of the flow.
In Sample	The sample selected for pass (1) or drop (0).
Num Packets	The number of packets seen for the flow.
Num Octets	The sum of the packet lengths of all packets seen for the flow.
From GigaVUE-FM	
num_devices	The number of devices.
num_devices_in_sample	The number of devices in the sample.

Flow Ops Report Statistics Definitions for GTP

The following table describes Flow Ops report statistics for GTP, including GTP flow filtering, GTP whitelisting, and GTP flow sampling:

Statistic	Description
Tunnel[Ver]	The tunnel type (CTRL for control plane, USER for user data plane) and version (1 or 2).
Interface EBI:LBI[QCI]	The interface and Evolved Packet System (EPS) Bearer Identifier (EBI), Linked Bearer Identity (LBI), and QoS Class Identifier.
IP:Tunnel-ID ==> IP:Tunnel-ID	The IP addresses and tunnel identifiers of both sides of the tunnel. **Note**: For LTE nodes only, the IP addresses can be both IPv4 and IPv6.
IMSI	The International Mobile Subscriber Identity (IMSI) value.
APN	The Access Point Name (APN) value.
WL	The IMSI had a match in the forward list (WL) or did not. The values are N for no and Y for yes.
FS	The IMSI was flow sampled or not. The values are A for accepted, R for rejected, and N for no match.
ID	The rule ID of the flow sample rule.

Statistic	Description
LB port	The load balancing port number.
Pkts	The number of packets.
Timestamp	The internal clock time of when the session was created.
GTP Resource Summary	
Num Sessions In Use	The number of correlated session in use.
Num Tunnels In Use	The number of used tunnels in the tunnel resource pool.
Tunnels Available	The number of available tunnels in the tunnel resource pool.
UPN CTunnels Available	The number of available tunnels in the UPN C tunnels resource pool.
Tunnels Pending Free	The number of tunnels marked free, ready to be returned to the tunnel resource pool.
Tunnels Marked Free	The number of times used tunnels are marked free.
Tunnels Returned	The number of times used tunnels are returned to the tunnel resource pool.
Current Time	The current time (in CPU cycles). Used for debugging.
Flow Filtering Report Summary	
Control Tunnels	The total number control tunnels.
Control & User Tunnels	The total number of control and user tunnels.
GTP Session Stats	The interface type: Gn/Gp, S1U/S11, S5/S8, S3/S4, or Other.
Sessions	The number of sessions by interface. Note the following: ▪ The counter for S3/S4 will always be 0 (not supported). ▪ The S10 interface is counted under Other.
Tunnels	The number of tunnels by interface.
Pending Session	The number of sessions waiting for control message response.
Control Only Session	The number of sessions without user bearers.
Flow Sampling Report Summary	
Total Devices	The total number devices in the flow sample

Statistic	Description
Number of Device in Sample	The number of devices used in each sample session
Flow SIP Report Summary	
SIP Sessions	This graph displays the total sessions, total sessions in use and the number of parse errors.
RTP Sessions	This graph displays the total number of sessions and the total data pool in use.
GTP Interface Stats For overlap maps, refer to the notes below the table for Flow Ops Report Statistics Definitions for GTP Overlap	
Rx Pkts	The received (Rx) packets for GTP correlation statistics for GTP traffic by interface type: S11, S1u, S5S8 (control and user), Gn (control and user), total (control and user), collector (control and user). **Note**: If traffic does not match any map rules, it will be sent to the collector.
Rx Bytes	The received (Rx) bytes for GTP correlation statistics for GTP traffic by interface type.
Sample/WL/Filter (Tx) Pkts	The transmitted (Tx) packets sampled in for flow sampling, forward listing, and flow filtering by interface type. For example, if sampling is 60%, then 60% is sampled in.
Sample/WL/Filter (Tx) Bytes	The transmitted (Tx) bytes sampled in for flow sampling, forward listing, and flow filtering by interface type.
Sample Out (Dropped) Pkts	The packets sampled out (dropped) by interface type. For example, if sampling is 60%, then 40% is sampled out.
Sample Out (Dropped) Bytes	The bytes sampled out (dropped) by interface type.
Xaui Drop	The total traffic (Rx bytes and packets) dropped due to oversubscription on the interface into the GigaSMART.
Statistics for Control Message (GTP-c)	
GTPV1	The GTPv1 (version 1) message type.
Cre PDP Req	Create Packet Data Protocol (PDP) context request.
Cre PDP Rsp	Create PDP context response.
Upd PDP Req	Update PDP context request.

Statistic	Description
Upd PDP Rsp	Update PDP context response.
Del PDP Req	Delete PDP context request.
Del PDP Rsp	Delete PDP context response.
GTPV2	The GTPv2 (version 2) message type.
Cre Ssn Req	Create session request.
Cre Ssn Rsp	Create session response.
Mod Bea Req	Modify bearer request.
Mod Bea Rsp	Modify bearer response.
Del Ssn Req	Delete session request.
Del Ssn Rsp	Delete session response.
Cre Bea Req	Create bearer request.
Cre Bea Rsp	Create bearer response.
Upd Bea Req	Update bearer request.
Upd Bea Rsp	Update bearer response.
Del Bea Req	Delete bearer request.
Del Bea Rsp	Delete bearer response.
Mod Bea Cmd	Modify bearer command.
Mod Bea Fai	Modify bearer failure indication.
Bea Rsr Cmd	Bearer resource command.
Bea Rsr Fai	Bearer resource failure indication.
Message Counters	
Tool Pass	The number of control messages passed to the tools.
Col NoSess	The number of control messages sent to the collector without matching sessions.
Col NoTnlx	The number of “out of tunnels” requests sent to the collector for the control message.
Col ParseEr	The number of control messages sent to the collector with unsupported options.
Col NoRule	The number of control messages sent to the collector

Statistic	Description
	without matching rules.
Col Other	The number of control messages sent to the collector with other conditions, for example, the message matched a drop rule.
Statistics for User Data Message (GTP-u)	
Tool Pass	The number of user data messages passed to the tools.
Collector	The number of user data messages sent to the collector without matching sessions.
Drop	The number of user data messages dropped.

Flow Ops Report Statistics Definitions for GTP Overlap

The following table describes Flow Ops report statistics for GTP overlap flow sampling. Since most fields are the same as *Flow Ops Report Statistics Definitions for GTP* on page 719, they are not repeated. Refer to the note below the table for GTP Interface Stats for overlap maps.

Statistic	Description
Tunnel[Ver]	The tunnel type (CTRL for control plane, USER for user data plane) and version (1 or 2).
Interface EBI:LBI[QCI]	The interface and Evolved Packet System (EPS) Bearer Identifier (EBI), Linked Bearer Identity (LBI), and QoS Class Identifier.
IP:Tunnel-ID ==> IP:Tunnel-ID	The IP addresses and tunnel identifiers of both sides of the tunnel.
IMSI	The International Mobile Subscriber Identity (IMSI) value.
APN	The Access Point Name (APN) value.
Overlap Result	The overlap result. A map group for GTP overlap flow sampling maps contains six maps. Each character in the result represents one of the maps. The result is in alphabetical order for each map within the map group. (There is no map priority.) Refer to the following legend: ▪ A—Flow Sample Accept ▪ R—Flow Sample Reject

Statistic	Description
	▪ W—Whitelist Accept ▪ N—No Match
Pkts	The number of packets.
Timestamp	The internal clock time of when the session was created.

The following notes are for the combination of overlap and non-overlap maps:

- If at least one flow sample map accepts the packets among all the pairs of overlap maps, the Sample (Tx) counters in the GTP Interface Stats will be incremented. If more than one pair of maps accepts the packets the Sample (Tx) counters in the GTP Interface Stats will still be incremented only once.
- If at least one forward list map matches the packets among all the pairs of overlap maps, the Sample (Tx) counters in the GTP Interface Stats will be incremented. If more than one pair of maps matches the packets the Sample (Tx) counters in the GTP Interface Stats will still be incremented only once.
- If there are no forward list maps and all flow sample maps are no-rule-match, the Sample (Tx) and Sample Out counters in the GTP Interface Stats will not be incremented.

Flow Ops Report Definitions for SIP/RTP Correlation

The following table describes Flow Ops report statistics for SIP/RTP correlation.

Statistic	Description
PROTO	The protocol, such as SIP, RTP, or RTCP. **Note**: The RTCP port number is assumed to be the corresponding RTP port number plus 1. This results in even numbered RTP ports and odd numbered RTCP ports.
TRANSPORT	The transport layer, such as UDP or TCP.
METHOD	The SIP method (or SIP message).
CALLER: IP	The corresponding IP address of the caller.
CALLEE: IP	The corresponding IP address of the callee.

Statistic	Description
PDU	The number of packets seen for this session.
CALL-ID	The call identifier.
WL	The caller/callee ID had a match in the whitelist (WL) or did not. The values are N for no and Y for yes.
FS	The caller ID was flow sampled or not. The values are A for accepted, R for rejected, and N for no match.
ID	The rule ID
LB port	The load balancing port number, which is the port over which the session has been load balanced. All SIP and corresponding RTP packets will be sent to this port.
Timestamp	The internal clock time of when the session was created.
Message counters	
SIP messages	All the SIP messages that have been correlated, such as, ACK, BYE 200, CANCEL, INFO, and so on. These counters are cumulative
Tool Pass	The number of SIP messages passed to the tools
NoSess	The number of SIP messages sent to the collector without matching sessions
NoRule	The number of SIP messages sent to the collector without matching rules.
NoMatch	Reserved
Other	The number of SIP messages sent to the collector with other conditions.
SIP Resource Summary	
Num Sessions In Use	The number of SIP sessions.
Sessions Available	The number of remaining SIP sessions available. (This varies by GigaVUE node.)
RTP Resource Summary	
Num Data Pools In Use	The number of RTP sessions.
Data Pools Available	The number of remaining RTP sessions available. (This varies by GigaVUE node.)

Flow Ops Report Statistics for Passive TLS/SSL Decryption

The following table describes Flow Ops report statistics for Passive S SL decryption:

Statistic	Description
Session Summary	
GsGroup	The GigaSMART group associated with the passive TLS/SSL decryption.
Total Sessions	The total number of passive TLS/SSL decryption sessions.
SSLv3 Sessions	The cumulative total number of SSL 3.0 sessions.
TLS1.0 Sessions	The cumulative total number of TLS 1.0 sessions.
TLS 1.1 Sessions	The cumulative total number of TLS 1.1 sessions.
TLS 1.2 Sessions	The cumulative total number of TLS 1.2 sessions.
Session IDs	The number of current session IDs.
Tickets	The number of current TLS tickets.
Report Summary	The link to the Report Summary page that displays the graphical representation of the passive TLS/SSL decryption session summary.
Session Details	
Server IP	The IP address of the server.
Server Port	The port number of the server.
Client IP	The IP address of the client.
Client Port	The port number of the client.
Version	The version of the SSL/TLS protocol.
First Error	The first error encountered on a session. After the first error, subsequent packets are dropped. If a session encounters errors, the packets for that session are ignored. The session will be cleared after the session times out.
First Error Reason	The reason for the first error encountered on a session.
In Packets	The number of packets going into the session.
Out Packets	The number of packets transmitted out of the session.

Statistic	Description
SNI	The TLS server name indication (SNI).
Decryption Status	The passive TLS/SSL decryption status.
Start Time	The session start time.
Duration	The duration of the session.
Finger Print	The finger print is applicable only for certificate-based TLS/SSL decryption session.
Cipher Suite	The cipher suite details used to decrypt the session.

GigaSMART Operations Statistics Definitions

To view GigaSMART Operations statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics** to open the Statistics page.

The following table describes the statistics that are common to all GigaSMART operations:

Statistic	Description
Pkts Drop	The cumulative number of packets dropped at a GigaSMART operation for a map.
Pkts Rx	The cumulative number of packets coming into a GigaSMART operation for a map.
Octets Rx	The cumulative number of bytes coming into a GigaSMART operation for a map.
Pkts Term	The cumulative number of packets of a terminated session of a GigaSMART operation for a map.
Octets Tx	The cumulative number of bytes going out of a GigaSMART operation for a map.
Pkts Tx	The cumulative number of packets going out of a GigaSMART operation for a map.
Pkts Drop No Init	For internal debugging.
Pkts Parse Err	The cumulative number of packets with invalid or unsupported header types of a GigaSMART operation for a map.

The following sections provide definitions for the statistics that are specific to a particular GigaSMART operation. Refer to the following:

- Passive TLS/SSL Decryption Statistics Definitions
- Inline TLS/SSL Decryption Statistics Definitions
- *De-duplication Statistics Definitions*
- ERSPAN Statistics Definitions
- Tunnel Decapsulation Statistics Definitions
- Tunnel Encapsulation Statistics Definitions
- *APF Statistics Definitions*
- ASF Statistics Definitions
- Masking Statistics Definitions
- *Slicing Statistics Definitions*
- Header Stripping Statistics Definitions
- Generic Header Stripping Statistics Definitions
- Trailer Statistics Definitions
- FlowVUE Statistics Definitions
- NetFlow Statistics Definitions

Passive TLS/SSL Decryption Statistics Definitions

The following table describes GigaSMART operations statistics for Passive TLS/SSL decryption:

Statistic	Description
Sessions Total	The cumulative total number of sessions.
Sessions Active	The number of currently active sessions.

Inline TLS/SSL Decryption Statistics Definitions

The statistics for inline TLS/SSL decryption are described in Work With Inline TLS/SSL Decryption.

De-duplication Statistics Definitions

The following table describes GigaSMART operations statistics for de-duplication (including the statistics that are displayed in a cluster environment):

Statistic	Description
Ip 6 Missed Op Busy	For IPv6 packets, the number of non-duplicate packets passed through without storing because the GigaSMART engine was too busy processing other packets. **Note**: Applies only to non-duplicate packets.
Non Ip Dupl	The number of non-IPv4 and non-IPv6 duplicate packets detected.
Non Ip	The number of non-IPv4 and non-IPv6 packets received for de-duplication.
Ip 6 Missed Op Space	For IPv6 packets, the number of non-duplicate packets passed through without storing because the GigaSMART engine did not have enough storage space. **Note**: Applies only to non-duplicate packets.
Non Ip Missed Op Busy	For non-IPv4 and non-IPv6 packets, the number of non-duplicate packets passed through without storing because the GigaSMART engine was too busy processing other packets. **Note**: Applies only to non-duplicate packets.
Non Ip Missed Op Space	For non-IPv4 and non-IPv6 packets, the number of non-duplicate packets passed through without storing because the GigaSMART engine did not have enough storage space. **Note**: Applies only to non-duplicate packets.
Ip 4 Missed Op Busy	For IPv4 packets, the number of non-duplicate packets passed through without storing because the GigaSMART engine was too busy processing other packets. **Note**: Applies only to non-duplicate packets.
Ip 4 Missed Op Space	For IPv4 packets, the number of non-duplicate packets passed through without storing because the GigaSMART engine did not have enough storage space. **Note**: Applies only to non-duplicate packets.
Ip 4 Ipid Tcp Rst	The number of TCP RESETs for IPv4 plus TCP packets. (IPv4 plus TCP RESET packets are not de-duplicated.)
Ip 6 Dupl	The number of IPv6 duplicate packets detected.
Ip 4 Dupl	The number of IPv4 duplicate packets detected.

ERSPAN Statistics Definitions

The following table describes GigaSMART operations statistics for ERSPAN tunnel:

Statistic	Description
Drop Type 3 Marker Bad Sig	The number of ERSPAN Type III packets that have a bad marker packet signature. The expected marker packet signature is 0xa5a5a5a5.
Pkts Too Big	The number of ERSPAN packets that are larger than 9600 bytes after the timestamp trailer is added. These packets will be dropped.
Pkts Rx Type 3	The total number of ERSPAN Type III packets received.
Pkts Tx Type 3 Marker	The total number of ERSPAN Type III marker packets received. (These packets are not forwarded.)
Pkts Rx Type 2	The total number of ERSPAN Type II packets received.
Type 3 Marker Overdue	The number of ERSPAN Type III marker packets that are overdue. Based on granularity, if a marker packet does not arrive by the time specified, it is considered overdue. Refer to ERSPAN Granularity.
pkts Rx	The total number of ERSPAN packets received into the IP interface.
Drop Unknown Proto	The number of packets dropped because they were not recognized as ERSPAN packets or ERSPAN marker packets.
Drop Id No Hit	The number of ERSPAN packets with a wrong flow ID/ERSPAN ID.

Tunnel Decapsulation Statistics Definitions

The following table describes GigaSMART operations statistics for tunnel decapsulation:

Statistic	Description
GMIP Tunnel	
Drop Wrong Addr	The number of GMIP tunneled packets dropped whose destination UDP port does not match the configured value.
Drop Other	The number of GMIP tunneled packets dropped because of fragmented packets.
Pkts Rx	The number of packets received into the IP interface.
Reassemble Rx Success	The number of incoming packets successfully

Statistic	Description
	reassembled.
Drop Unknown Proto	The number of packets through the network IP interface dropped if they are neither IPv4 or UDP packets.
Reassemble Rx	The number of incoming packets to be reassembled.
Drop Reassemble Mem Full	The number of packets not successfully reassembled and dropped due to lack of memory for reassembly.
Sliced Mtu	Not valid for tunnel decapsulation.
Drop Reassemble Max Hold	The number of packets not successfully reassembled and dropped due to the arrival of too many fragments not yet reassembled. The maximum fragment sessions are already held for reassembly.
Drop No Arp	The number of packets dropped because ARP was not resolved on the IP interface (in particular, on a tool IP interface).
Drop Reassemble Overlap Frag	The number of packets not successfully reassembled and dropped due to overlapping reassembly fragments.
L2GRE Decap Tunnel	
Drop Key Mismatch	The number of packets dropped due to key mismatch.
Pkts Rx	The number of packets received into the IP interface.
Reassemble Rx Success	The number of incoming packets successfully reassembled.
Drop Unknown Proto	The number of packets dropped due to an unknown protocol.
Reassemble Rx	The number of incoming packets to be reassembled.
Drop Reassemble Mem Full	The number of packets not successfully reassembled and dropped due to lack of memory for reassembly.
Drop Reassemble Max Hold	The number of packets not successfully reassembled and dropped due to the arrival of too many fragments not yet reassembled. The maximum fragment sessions are already held for reassembly.
Drop Other	The number of packets dropped due to other reasons.
Drop Reassemble Overlap Frag	The number of packets not successfully reassembled

Statistic	Description
	and dropped due to overlapping reassembly fragments.
Pkts Too Big	The number of packets that are larger than 9000 bytes after the timestamp trailer is added. These packets will be dropped.

Tunnel Encapsulation Statistics Definitions

The following table describes GigaSMART operations statistics for tunnel encapsulation:

Statistic	Description
GMIP Encap Tunnel	
Fragment Rx	The number of incoming packets to be fragmented.
Drop Other	Not valid for tunnel encapsulation.
Pkts Rx	The number of packets received into the IP interface.
Fragment Tx	The number of outgoing packets sent to the tunnel after fragmentation.
Sliced Mtu	The number of packets that are sliced to the MTU size of the tool IP interface.
Drop No Arp	The number of packets dropped because ARP was not resolved on the IP interface (in particular, on a tool IP interface).
L2GRE Encap Tunnel	
Fragment Rx	The number of incoming packets to be fragmented.
Pkts Rx	The number of packets received into the IP interface.
Fragment Tx	The number of outgoing packets sent to the tunnel after fragmentation.
Session Current Total	The number of currently active sessions.
Session Alloc Fail	The number of session allocations that failed.
Session Lookup	The number of lookups in a session.
Session Lookup Success	The number of lookups in a session that were a success.

Statistic	Description
Session Alloc	The number of sessions allocated.
Session Timedout	The number of sessions that timed out after a configured timer value.
Sliced Mtu	The number of packets that are sliced to the MTU size of the tool IP interface.
Drop No Arp	The number of packets dropped because ARP was not resolved on the IP interface (in particular, on a tool IP interface).
CUSTOM Decap Tunnel	
rx_packets	The number of packets received into the IP interface.
pkts_drop_unknown_protocol	The number of packets dropped due to an unknown protocol.
pkts_drop_portsrc_mismatch	The number of packets dropped due to source port mismatch.
pkts_drop_portdst_mismatch	The number of packets dropped due to destination port mismatch.
pkts_in_reassemble	The number of incoming packets to be reassembled.
pkts_in_reassemble_success	The number of incoming packets successfully reassembled. For example, if four (4) packets are reassembled into one (1), 4 is displayed in this field.
pkts_out_reassembled	The actual number of packets sent out the tool port. For example, if four (4) packets are reassembled into one (1), 1 is displayed in this field.
pkts_drop_reassemble_overlap_frag	The number of packets not successfully reassembled and dropped due to overlapping reassembly fragments.
pkts_drop_reassemble_max_hold	The number of packets not successfully reassembled and dropped due to the arrival of too many fragments not yet reassembled. The maximum fragment sessions are already held for reassembly.
pkts_drop_reassemble_mem_full	The number of packets not successfully reassembled and dropped due to lack of memory for reassembly.
pkts_drop_reassemble_timed_out	The number of packets dropped due to timeout in the reassembly queue.
VXLAN Decap Tunnel	
rx_packets	The number of packets received into the IP interface.
pkts_drop_unknown_protocol	The number of packets dropped due to an unknown protocol.
pkts_drop_portsrc_mismatch	The number of packets dropped due to source port mismatch.
pkts_drop_portdst_mismatch	The number of packets dropped due to destination port mismatch.
pkts_in_reassemble	The number of incoming packets to be reassembled.

Statistic	Description
pkts_in_reassemble_success	The number of incoming packets successfully reassembled. For example, if four (4) packets are reassembled into one (1), 4 is displayed in this field.
pkts_out_reassembled	The actual number of packets sent out the tool port. For example, if four (4) packets are reassembled into one (1), 1 is displayed in this field.
pkts_drop_reassemble_overlap_frag	The number of packets not successfully reassembled and dropped due to overlapping reassembly fragments.
pkts_drop_reassemble_max_hold	The number of packets not successfully reassembled and dropped due to the arrival of too many fragments not yet reassembled. The maximum fragment sessions are already held for reassembly.
pkts_drop_reassemble_mem_full	The number of packets not successfully reassembled and dropped due to lack of memory for reassembly.
pkts_drop_reassemble_timed_out	The number of packets dropped due to timeout in the reassembly queue.
pkts_too_big	The number of packets that are larger than 9000 bytes after the timestamp trailer is added. These packets will be dropped.

APF Statistics Definitions

The following table describes GigaSMART operations statistics for APF:

Statistic	Description
Apf Drop	The number of packets matching the GigaSMART drop rules.
Apf Pass	The number of packets matching the GigaSMART pass rules.
Rule Not Match	The number of packets not matching any GigaSMART rules in the map.
Masking Err	The number of masking errors in the map. This number is usually zero (0), which indicates no issues with the masking operation. If the number is non-zero, it indicates there is some issue with the masking operation.

ASF Statistics Definitions

The following table describes GigaSMART operations statistics for ASF:

Statistic	Description
Session Created	The number of sessions created.
Session Deleted	The number of sessions deleted.
Session Timeout	The number of sessions that were deleted due to inactivity (expiry of the session timer).
Session Matched (pkt)	The number of incoming packets matching ASF sessions. This count does not include the packets that triggered the creation of the session.
Exceed Count Drop	The number of packets dropped, even if they matched a flow session, because a packet-count was configured and exceeded.

Masking Statistics Definitions

The following table describes GigaSMART operations statistics for masking:

Statistic	Description
No Header	The number of packets with no configured masking protocol.
Too Short	The number of packets with a length less than the masking offset.

Slicing Statistics Definitions

The following table describes GigaSMART operations statistics for slicing:

Statistic	Description
No Header	The number of packets with no configured header for slicing.
Too Short	The number of packets with a length less than the slicing length.
Min Len	The number of packets that are sliced to less than 64 bytes.

Header Stripping Statistics Definitions

The following table describes GigaSMART operations statistics for Header Stripping:

Statistic	Description
Id No Hit	The number of packets that do not match the configured VXLAN ID to be stripped.
Fm 6000 Pkts Ts	The number of packets received with the FM6000 timestamp.
Unknown Next	The number of packets not stripped of their configured header type as the packets will have an unknown header after the header is stripped.
No Header	The number of packets with no configured header type to be stripped.
Fm 6000 Data Pkt Too Big	The number of FM6000 packets that are larger than 9600 bytes after the timestamp trailer is added. These packets will be dropped.
Fm 6000 Keyframe Overdue	The number of FM6000 key frames that are overdue. By default, the key frame rate is 1 packet per second. If a key frame is not received in 1 second, this statistic is incremented.
Fm 6000 Keyframe	The number of key frames received from the FM6000 device. Key frames are marker packets that carry information to convert the FM6000 timestamp to UTC time.

Generic Header Stripping Statistics Definitions

The following table describes GigaSMART operations statistics for generic Header Stripping:

Statistic	Description
first_anchor_header_not_present	The number of packets with no configured first anchor header.
offset_beyond_anchor_header_size	The number of packets that are not stripped as the configured offset size exceeds the size of the first anchor header.
second_anchor_header_not_present	The number of packets with no configured second anchor header.

Statistic	Description
strip_length_beyond_pkt_len	The number of packets that are not stripped as the configured strip length exceeds the packet length.
strip_success	The number of packets that are successfully stripped.
incompatible_anchor_headers	The number of packets that are dropped as the first anchor header is incompatible with the second anchor header after the stripping operation is complete.

Trailer Statistics Definitions

The following table describes GigaSMART operations statistics for trailers:

Statistic	Description
pkts_too_big	The number of packets for which the size of the packet has become greater than the maximum supported size (9600) when adding the trailer. This count is incremented with one trailer added to the packet.

FlowVUE Statistics Definitions

The following table describes GigaSMART operations statistics for FlowVUE:

Statistic	Description
Dev Ip Src Match	The number of packets with a source IP matching the range defined in the GigaSMART parameters (gsparams).
Exceed License Warn	The number flows that exceed the installed license.
Exceed License Err	The number of flows that exceed the installed license error limit. (The limit is the number of sessions in the license plus 5%.)
Dev Ip No Match	The number of packets with a source and destination IP that does not match any of the ranges defined in the GigaSMART parameters (gsparams).
Dev Ip Non Ip	The number of packets with no IP headers (and hence, not sampled).

Statistic	Description
Dev Ip Dst Match	The number of packets with a destination IP matching the range defined in the GigaSMART parameters (gsparams).
Out of Resource	The number of times there was a failure to allocate resources for recording a new flow. **Note**: The maximum supported flows for each engineport is 2 million.
Dev Ip Drop Not In Sample	The number of packets dropped by the sampling application, not because of errors, but because the flow was sampled to be dropped.

NetFlow Statistics Definitions

The following table describes GigaSMART operations statistics for NetFlow:

Statistic	Description
Out of Resource	The number of times there was a failure to allocate resources for recording a new flow. **Note**: The maximum supported flows for each engine port is 2 million.
Non Ip	The number of packets received that are not IPv4 or IPv6.
Non Configured	Packets received when NetFlow is not enabled in the GigaSMART group.
TLS/SSL Active Sessions	The number of currently active TLS/SSL sessions monitored by NetFlow.
Total TLS/SSL Sessions	The cumulative total number of TLS/SSL sessions monitored by NetFlow.

Display GigaSMART Application Resource Usage

GigaSMART applications, such as De-duplication and inline and Passive SSL decryption, use memory resources on the GigaSMART line card or module. As new GigaSMART applications are configured, the total resources on the GigaSMART line card or module can become fully used.

Starting in software version 4.4, you can display the GigaSMART application resource usage, which provides information about the applications that use resources. With this information, you can choose to free up resources on one application to use them on another.

Table 12: GigaSMART Application Resource Information

Name	Format
GSgroup	The alias of the GigaSMART group associated with the GigaSMART application.
Application	The list of licensed GigaSMART applications that use resources, including Adaptive Packet Filtering (APF), de-duplication, GPRS Tunneling Protocol (GTP), NetFlow Generation, Passive SSL decryption, Adaptive Session Filtering (ASF) with buffering, and inline SSL decryption. Other GigaSMART applications (such as flow sampling, header addition, Header Stripping, ERSPAN tunnel decapsulation, slicing, masking, and others) do not have databases that store data, therefore they do not use resources and are not displayed with the **show gsgroup gsapp-resource** command.
% of Total	The percentage of the total amount of memory used by each GigaSMART application.
Configured Resource	The amount of resources configured for each GigaSMART application. The valid values are as follows: ▪ app-max—Indicates the maximum amount of memory configured for the application. It is a pre-allocated amount. ▪ integer, such as 2—Indicates the number of sessions configured, in the units specified. M indicates millions.
Installed Resource	The amount of resources installed for each GigaSMART application.
Licensed Quantity	The amount of resources licensed for each GigaSMART application.
Units	The units, such as sessions, or millions (M) of sessions.

Overview of GigaSMART Application Resources

GigaSMART application resources are managed per GigaSMART group (gsgroup). A gsgroup can be configured with one or more GigaSMART engine ports on one or more GigaSMART line cards or modules.

For most GigaSMART applications, resources are allocated automatically, based on configuration. For some GigaSMART applications, resources are allocated when they are

configured in the gsgroup. For other applications, resources are allocated when a GigaSMART operation (gsop) using the application is created. For buffer ASF however, you can explicitly configure resources, in millions of sessions.

The allocation of resources for a new application will be successful if the application is licensed and if there is sufficient space for the new application.

If there is insufficient space, then the resources need to be managed to free up memory for the new application. Managing resources includes deleting applications that are no longer used.

However, deleting a GigaSMART application does not result in the immediate deletion of application resources. Once a resource has been allocated, it remains allocated. To delete resources for APF, out-of-band or inline SSL decryption, de-duplication, and GTP, remove the configuration related to the application, then reload the GigaSMART line card or module.

To remove the configuration related to an application, delete the gsop first, then delete the gsgroup. If the gsop or gsgroup is bound to a map, you will also have to delete the map.

Resources for Buffer ASF

The resources for buffer ASF depend on the number of sessions and the type of node. For example, GigaVUE-HC3, 5 million sessions uses 26%.

For GigaVUE-HC3, the resources for buffer ASF for the number of sessions is as follows:

- 2 million—11% of total resources
- 3 million—16% of total resources
- 4 million—21% of total resources
- 5 million—26% of total resources

Reload GigaSMART Line Card or Module

Occasionally, the GigaSMART line card or module will need to be reloaded for changes to take effect and to allocate resources accordingly. Reloading also provides applications with contiguous memory.

The following message displays at the bottom of the output of the **show gsgroup gsapp-resource** command when the GigaSMART line card or module needs to be reloaded:

```
*Resource allocation changes have been made that require GigaSMART card 2/1/1 to be reloaded in
order for them to take effect.
```

When this message is displayed, you cannot change the configuration relating to that application until after the reload. For example, you cannot use the gsop, associated with the gsgroup, in a map.

Use the following command to reload a GigaSMART line card or module:

(config) # card slot <slot ID> down

Use the following command to bring the GigaSMART line card or module back up:

(config) # no card slot <slot ID> down

GigaSMART CPU Utilization Statistics

You can display CPU utilization statistics for GigaSMART. The statistics indicate the performance of GigaSMART, improve visibility, and help identify high load conditions.

Show commands display instantaneous CPU utilization as well as historical, providing trends for CPU utilization.

You can also configure a rising threshold, as a percentage, to indicate when high CPU utilization occurs. When the aggregate CPU utilization percentage exceeds the rising the threshold, an SNMP notification can be triggered.

This feature is supported on all products that support GigaSMART: GigaVUE-HC1 and GigaVUE-HC3.

The GigaSMART engine port (**e** port) numbers are e1 on nodes with one GigaSMART engine and e1 and e2 on nodes with two GigaSMART engines, for example: 10/1/e1 or 8/1/e1 and 8/1/e2.

Refer to the following sections for viewing statistics, configuring the threshold, and configuring a notification that can be sent when the threshold is exceeded:

- Display GigaSMART CPU Utilization
- Configure Threshold
- Configure Threshold Crossing Notification

Display GigaSMART CPU Utilization

Use the **show gsgroup stats all** command to display the statistics on all GigaSMART groups on the node.

The statistics are displayed for 1 second, 1 minute, 5 minute, 10 minute, and 15 minute intervals. The 1 second interval displays the statistics for the previous second. The 1 minute, 5 minute, 10 minute, and 15 minute intervals display statistics containing history.

Statistics are displayed in an aggregate form. For example, if there are two GigaSMART**e** ports: e1 and e2, there will be two aggregates. One aggregate will be for e1, the other will be for e2. The term aggregate refers to aggregation across all packet processing cores (up to 31) in the CPU. It does not refer to an aggregate across CPUs.

The statistics are as follows:

- Useful Time—Amount of time during which the CPU is processing packets, in milliseconds (ms) or seconds (s).
- Idle Time—Amount of time during which the CPU is not processing packets, for example, when it is busy looping, in milliseconds (ms) or seconds (s).
- In Packets (pkts/s)—Number of packets per second coming into the CPU. For the 1 second interval, In Packets is the actual number of incoming packets for that second. For the 1 minute, 5 minute, 10 minute, and 15 minute intervals, In Packets is an average number of incoming packets per second.
- Packets Drop (pkts/s)—Number of packets per second dropped by the CPU. For the 1 second interval, Packets Drop is the actual number of dropped packets for that second. For the 1 minute, 5 minute, 10 minute, and 15 minute intervals, Packets Drop is an average number of dropped packets per second.
- Packets Recv Error—Number of received packets per second with errors. For the 1 second interval, Packets Recv Error is the actual number of errored packets for that second. For the 1 minute, 5 minute, 10 minute, and 15 minute intervals, Packets Recv Error is an average number of errored packets per second.
- CPU Usage %—Percentage of time during which the CPU is processing packets. CPU Usage % plus CPU Idle % equals 100.
- CPU Idle %—Percentage of time during which the CPU is not processing packets. CPU Idle % plus CPU Usage % equals 100.

NOTE: When the node is restarted, the 1 minute, 5 minute, 10 minute, and 15 minute statistics will not be exactly for 1 minute, 5 minute, 10 minute, and 15 minute intervals, until the full interval has elapsed and the history is available.

Configure Threshold

An upper threshold (rising) can be configured. When the aggregate value of the CPU utilization on the GigaSMART engine exceeds the threshold, an SNMP notification can be triggered. Refer to Configure Threshold Crossing Notification.

Configure Threshold Crossing Notification

When the aggregate value of the CPU utilization exceeds the upper (rising) threshold, a message is logged, and optionally, an SNMP notification is sent to all configured destinations.

When enabled, the SNMP notification is sent when the rising threshold is exceeded in any 5-second interval over which the CPU utilization is averaged.

NOTE: Once the rising threshold is exceeded for 5 seconds, the SNMP notification is generated. However, if the CPU utilization falls below the upper threshold but does not remain below that threshold continuously for 5 seconds, a new notification is not generated when the upper threshold is exceeded again. A new notification is generated only when the CPU utilization falls below the threshold, stays below the threshold continuously for 5 seconds, exceeds the threshold again, and stays above it for 5 seconds.

How to Use GigaSMART Trailers

Required License: Base

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigavUE-HC3 Gen 3, GigaVUE-HC1-Plus, and GigaVUE-HCT

GigaSMART operations can add the GigaSMART Trailer to packets, providing metadata on where the packet was processed.

GigaSMART Trailers can be combined with other GigaSMART operations. Refer to How to Combine GigaSMART Operations for the valid combinations.

If a trailer is included, it can optionally include the original packet's CRC as one field and a Source ID (Source port label) as another. The Source ID indicates where the packet entered the GigaVUE H Series node and how it was processed. The modified packet's actual CRC is always recalculated to reflect its new length. Refer to the Source ID Field for information.

Trailer operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to Groups of GigaSMART Engine Ports for details.

Refer to Supported GigaSMART Operations Supported GigaSMART Operationsfor more details on the devices that support GigaSMART operations.

Source ID Field

When you enable the **Source ID** field for the GigaSMART Trailer, the trailer includes an additional field that identifies the platform type, box ID, slot number, and port number of the network port where the packet entered the GigaVUE H Series node. Refer to Format of the GIGAMON_SRCID TLV for the exact details.

Keep in mind the following when configuring GigaSMART operations with a **Source ID** argument:

Feature	Description
Source ID Field in GigaSMART Trailer	The **Source ID** field in the Gigamon Trailer includes the following values: o **Platform Type –** The type of GigaVUE node where the packet was first seen. o **Group ID** – The cluster ID/group ID configured for the node on which the packet was received. The GigaVUE H Series uses cluster IDs. o **Box ID** – The box ID is configured for the GigaVUE node on which the packet was received. Box IDs are used for the unique identification of nodes in a cluster. **NOTE:** The box ID field in the Gigamon Trailer supports box ID values from 1-63, inclusive. o **Slot ID** – The slot ID for the port on which the packet was received. o **Port ID** – The physical port number on which the packet was received.

Example – GigaSMART Source Port Labeling with a GigaSMART Trailer

This example creates a GigaSMART operation named **src_headermask** with **Masking** and **Trailer** components. This operation will mask packets using a static masking offset of 148 bytes that continues for the next 81 bytes, writing over the existing data with an FF pattern. Then it attaches a GigaSMART Trailer indicating the original size of the packet before

masking, the original packet's CRC, and the box ID, slot ID, and port ID of the physical input port on the GigaVUE H Series node. Figure 118GigaSMART Operation with Masking and Trailer Components shows the GigaSMART operation with masking and trailer components.

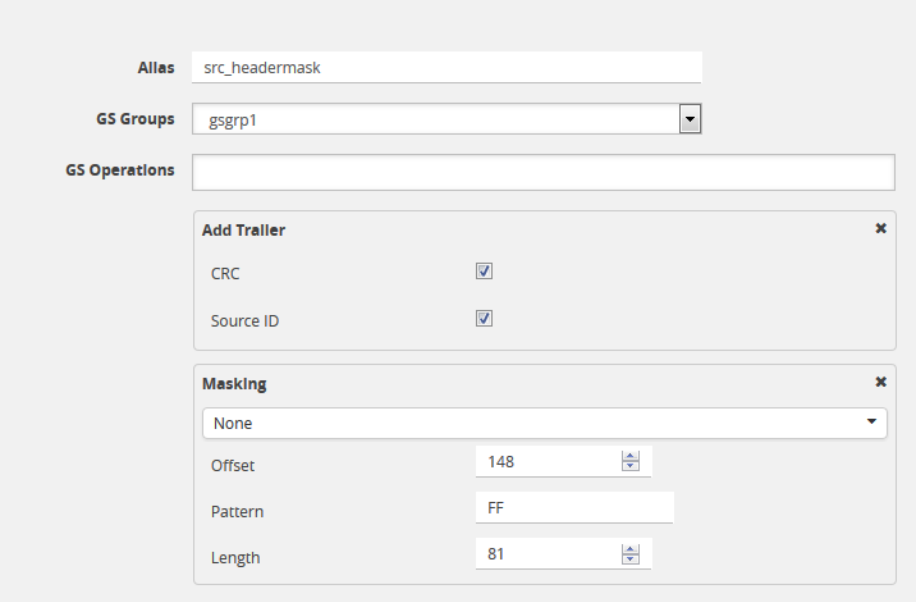


Figure 118 *GigaSMART Operation with Masking and Trailer Components*

Remove GigaSMART Trailers

You can also construct GigaSMART operations that remove the GigaSMART Trailer from packets. These operations are useful in cases where you have cascade connections – a tool port receiving packets with a GigaSMART trailer is physically cabled to a GigaVUE H Series network port, sending the packets received on the tool port back into a GigaVUE H Series node. You may want to remove the GigaSMART trailer before the packets are forwarded to other tools – that is when the special **Remove Trailer** argument comes in handy. Figure 119Remove Trailer Enabled shows Remove Trailer enabled for a GigaSMART operation.

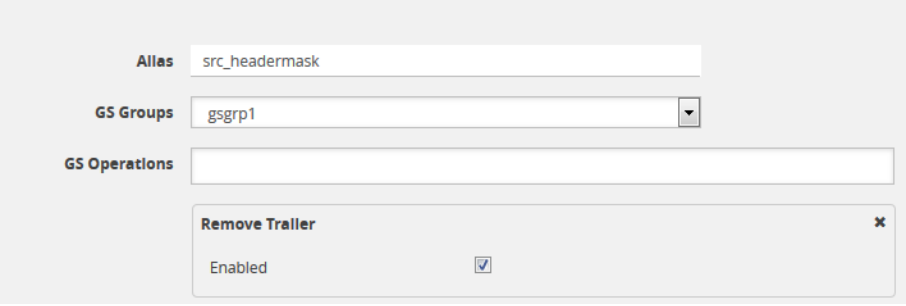


Figure 119 *Remove Trailer Enabled*

Example: Removing GigaSMART Trailers in a Cascade

Figure 120Cascade Physically Cabled from Tool to Network Port and Figure 121Removing the GigaSMART Trailer from a Cascade Connection illustrate a situation where a GigaSMART operation that removes the GigaSMART trailer would be useful. Consider the physical deployment shown in Figure 120Cascade Physically Cabled from Tool to Network Port:

- The green illustrates a one-way cascade between a tool (output) port (port 1/1/x12) and a network (input) port (port 1/1/x9).
- If traffic arriving on port 1/1/x12 includes a GigaSMART Trailer, you may want to use a **Trailer Remove** GigaSMART operation to remove it before logically forwarding traffic from port 1/1/x9 to another port on the GigaVUE.

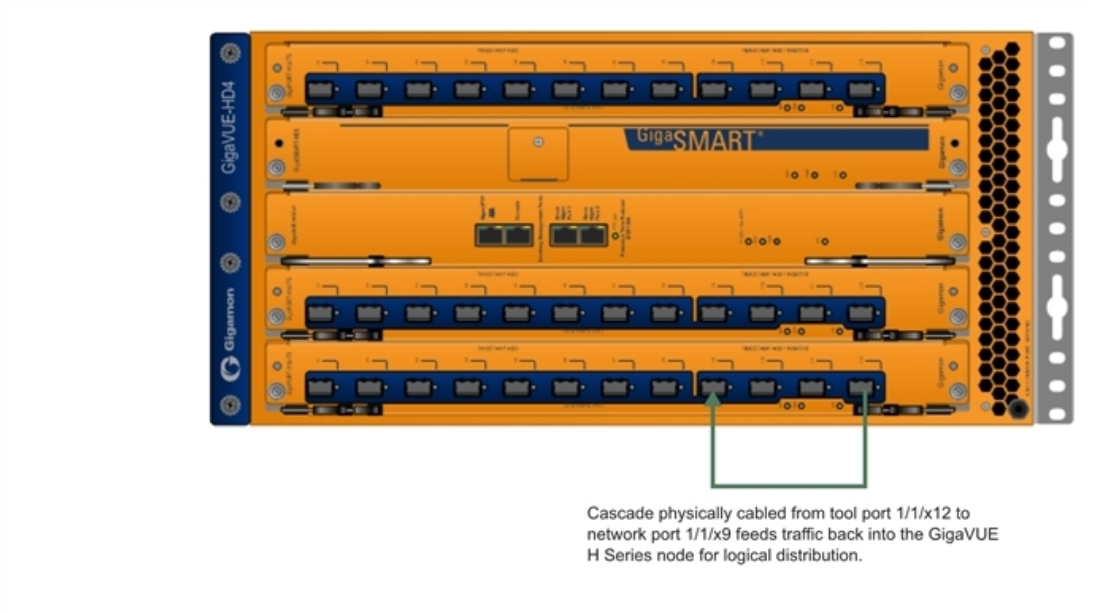


Figure 120*Cascade Physically Cabled from Tool to Network Port*

For example, consider the packet distribution shown in Figure 121Removing the GigaSMART Trailer from a Cascade Connection:

- The map named **add_trailer** is bound to network port 1/1/x1..x2. It adds the GigaSMART trailer and sends it to tool ports 1/1/x5..x8.
- Tool ports 1/1/x5..x7 are all connected to tools that expect the extra data in the GigaSMART Trailer.
- Tool port 1/1/x8 is physically cabled to network port 1/1/x4 in a cascade. To remove the GigaSMART Trailer from packets arriving on this port before they are forwarded to tool port 1/1/x9, we have bound a map named **no_trailer** to network port 1/1/x4 that is configured to remove the GigaSMART Trailers from all arriving packets.
- Tool port 1/1/x9 receives packets without the GigaSMART Trailer attached.

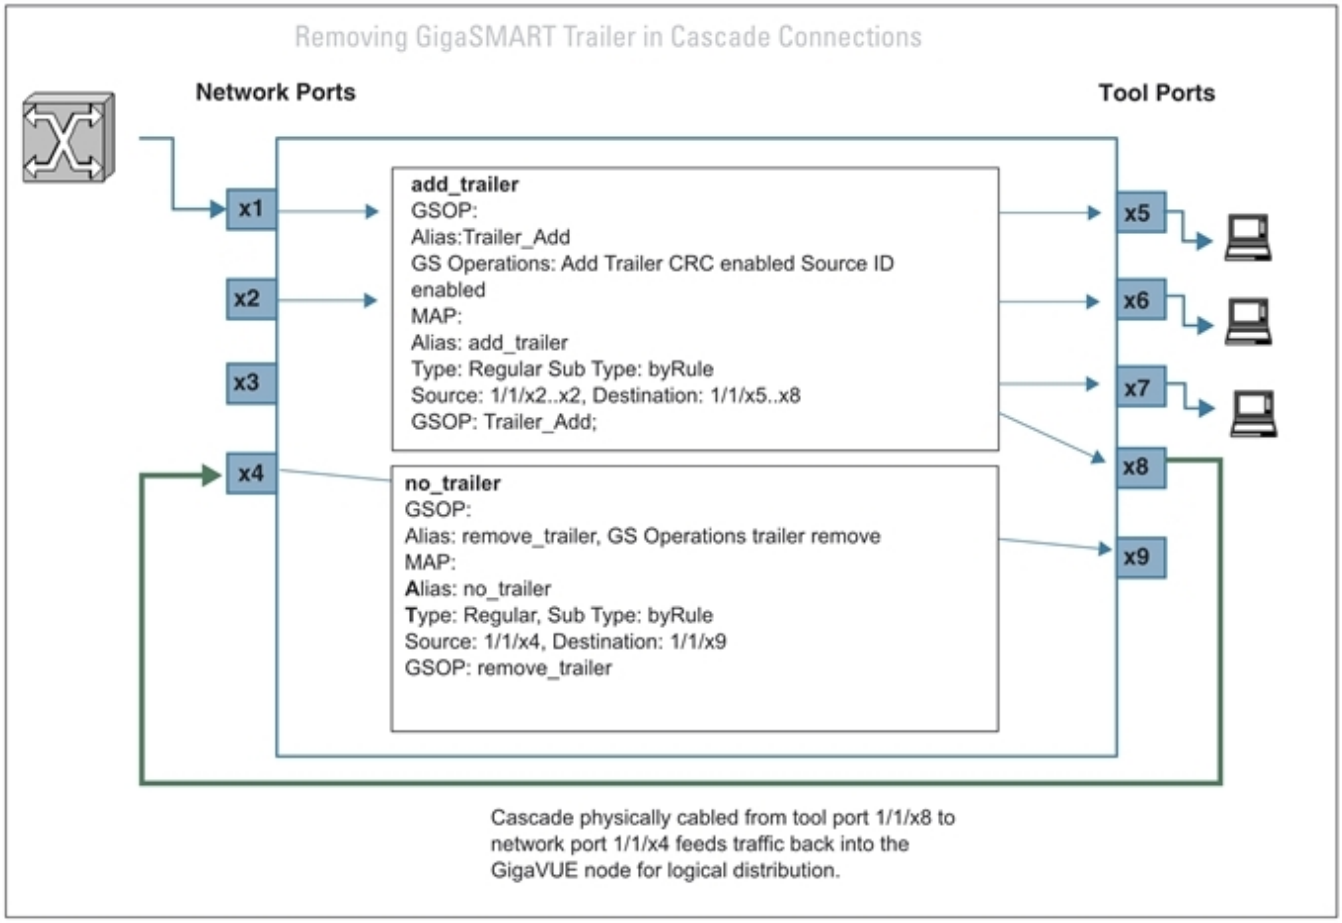


Figure 121 *Removing the GigaSMART Trailer from a Cascade Connection*

Multiple GigaSMART Trailers and Cascade Connections

Cascade connections also make it possible for multiple GigaSMART trailers to be attached to the same packet. For example, consider the cascade shown in Figure 121Removing the GigaSMART Trailer from a Cascade Connection and suppose that instead of the **no_trailer** map removing the GigaSMART Trailer on packets arriving over the cascade physically cabled from tool port 1/1/x8, there is a second GigaSMART operation adding another trailer. In cases like this, the GigaSMART adds the most recent trailer at the end of the packet.

The same principle works for the **Remove Trailer** operations:

- The most recent trailer is removed from the end of the packet. Any other trailers are left intact by a single Remove operation.

Interpret GigaSMART Trailer

The trailer inserted by the GigaSMART line card can be interpreted using a recent version of the Wireshark® Protocol Analyzer. Refer to the GigaSMART Trailer Reference for details on the GigaSMART Trailer and its TLVs.

GigaSMART Trailer Reference

This section provides reference information on the format, position, and contents of the Gigamon Ethertype and GigaSMART Trailer fields in a packet processed by the GigaSMART-HD0 line card.

Refer to How to Use GigaSMART Trailers for details on how the GigaSMART Trailer is used in packets.

GigaSMART Trailer Format

This section describes the format of the GigaSMART Trailer. Figure 122GigaSMART Trailer Format summarizes the position and contents of the GigaSMART Trailer and the Gigamon Ethertype field (0x22E5).

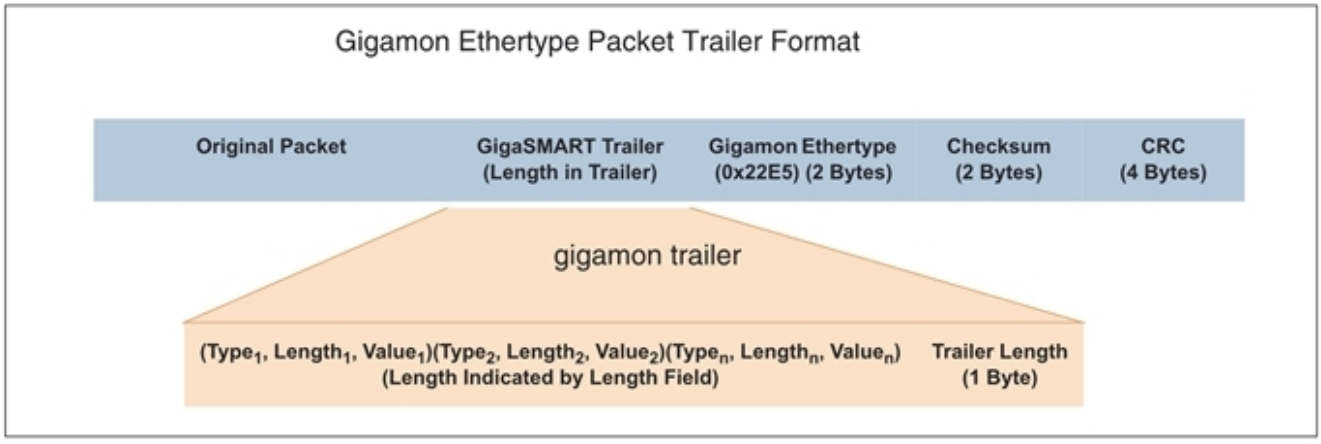


Figure 122 *GigaSMART Trailer Format*

GigaSMART Trailer Format

GigaSMART operations can insert metadata as a trailer at the end of the packet. As shown in the following figure, the GigaSMART Trailer consists of one or more TLVs, followed by the Trailer Length Field. The GigaSMART Trailer is followed by the Gigamon Ethertype, a

Checksum, and a recalculated CRC for the packet.

(Type1, Length1, Value1)(Type2, Length2, Value2)(Type(n), Length(n), Value(n))(Length Indicated by Length Field)	Trailer Length(1 Byte)	Gigamon Ethertype (0x22E5) (2 Bytes)	Checksum(2 Bytes)

Table 13: Gigamon Trailer Format lists and describes each of the fields in the GigaSMART Trailer:

Table 13: Gigamon Trailer Format

Field	Description
TLVs	TLVs are used to carry the metadata inserted by GigaSMART operations. Within a TLV, the Type and Length fields are each one byte long. The length field indicates the size of the Value field. Refer to GigaSMART Trailer TLVs for a summary of the available TLVs and their formats.
Length	Specifies the size of the GigaSMART trailer. The Trailer Length does not includes the length bytes
Gigamon Ethertype	Two-byte field identifying packet modified by GigaSMART line card. Value is 0x22E5. Protocol analysis equipment can use the Gigamon Ethertype in the trailer to find the start of the trailer: • Move one byte to the left to find the Trailer Length Field. • Read the number of bytes for the Trailer Length and move the number of bytes specified to the left to find the start of the Trailer.
Checksum	Two-byte field used to validate that the extra data is, in fact, a trailer and not random data with the Gigamon Ethertype.

GigaSMART Trailer TLVs

This section lists and describes the format of the Gigamon TLVs used in this release. TLVs are used to carry the metadata inserted by GigaSMART operations. Within a TLV, the Type and Length fields are each one byte long. The size of the Value field is indicated by the Length field.

Table 14: Gigamon TLVs

Tag Type	TLV ID	Value Field Length (Bytes)	Description
GIGAMON_PKT_LEN	1	2	**Original Packet Length** This TLV is included in any packet with a GigaSMART Trailer – adding the Trailer changes the original packet length.
GIGAMON_SRCID_G	2	3	**Original Packet Source Identifier** This TLV is included in any packet processed by a GigaSMART operation configured to include the Source ID (Source port label) field as part of its trailer. Refer to Format of the GIGAMON_SRCID TLV for a description of how the physical input source is encoded.
GIGAMON_CRC	7	4	Original Packet CRC This TLV is included in any packet processed by a GSOP configured to include the original packet's CRC as part of its trailer (If the **CRC** option is selected in the **Trailer Add** argument).
GIGAMON_SRCID	8	4	Original Packet Source Identifier (GigaVUE H Series) This TLV is included in any packet processed by a GigaSMART line card on a GigaVUE H Series node configured to include the Source ID (Source port label) field as part of its trailer (If the **SRCID** option is selected in the **Trailer Add** argument). Refer to Format of the GIGAMON_SRCID TLV for a description of how the physical input source is encoded.

Format of the GIGAMON_SRCID TLV

The GIGAMON_SRCID TLV is included in any packet processed by a **Source ID**(Source port label) GigaSMART operation. The GIGAMON_SRCID TLV consists of 3 bytes indicating the platform type, group ID, box ID, and port ID for the physical port where the packet entered the GigaVUE H Series node:

GIGAMON_SRCID TLV (32 Bits/4 Bytes)				
Platform(6 Bits)	**Group ID (4 Bits)**	**Box ID(6 Bits)**	**Slot ID(6 Bits)**	**Port ID(10 Bits)**

Name	Description	Bits
Platform	The type of GigaVUE node where the packet was first seen. Can be one of the following: o Unknown o GigaVUE-HCT o GigaVUE-HC1 o GigaVUE-HC1-Plus o GigaVUE-HC3 o GigaVUE-TA25 o GigaVUE-TA25E o GigaVUE-TA100 o GigaVUE-TA200 o GigaVUE-TA200E o GigaVUE-TA400 o GigaVUE-TA400E For bitmapping details, refer to the Platform Mapping	6
Group ID	The “Group Id” ranges between 0 and 15. For standalone devices, it is set to 0 and for the devices in a cluster, it reports the corresponding cluster ID.	4
Box ID	Box IDs are used to uniquely identify nodes in a cluster. Standalone systems typically have a default box ID of **1** here. **NOTE:** The box ID field in the Gigamon Trailer supports box ID values from 0-63, inclusive. Box IDs within the range 1 to 63 will be displayed in the SRC TLV. For any other value, the Box ID will be shown as zero.	6
Slot ID	The number of the slot including the physical port for the packet.	6
Port ID	The physical input port number for the packet. Refer to Port ID Values by Line Card Type for a summary of the values used by GigaVUE H Series line card.	10

Platform Mapping

The following table summarizes the bitmapping details for the platforms:

Platform	Ports SCRID Value for Platform Type
GMON_GS_CHASSIS_TYPE_HCT	26
GMON_GS_CHASSIS_TYPE_HC1	15
GMON_GS_CHASSIS_TYPE_HC1P	16
GMON_GS_CHASSIS_TYPE_HC3	17
GMON_GS_CHASSIS_TYPE_TA25	21
GMON_GS_CHASSIS_TYPE_TA25E	25
GMON_GS_CHASSIS_TYPE_TA100	13
GMON_GS_CHASSIS_TYPE_TA200	19
GMON_GS_CHASSIS_TYPE_TA200E	24
GMON_GS_CHASSIS_TYPE_TA400	23

Port ID Values by Line Card Type

The following table summarizes the values inserted in the GIGAMON_SRCID TLV for the port ID by GigaVUE device and the line card type:

GigaVUE Device / Line Card	Ports	Port IDs Inserted in Gigamon Trailer
GigaVUE-HC1 Gen 2		
HC1-X12G4	x1..x12	5..16
	g1..g4	1..4
BPS-HC1-D25A24	x1..x4	53..56
	x5..x8	49..52
PRT-HC1-X12	x1..x12	49..60
GigaVUE-HC3 Gen 2		
PRT-HC3-X24	x1..x24	1..24
PRT-HC3-C08Q16	c1..c8	1, 5, 9, 13, 17, 21, 25, 29

GigaVUE Device / Line Card	Ports	Port IDs Inserted in Gigamon Trailer
PRT-HC3-C08Q16–4x10Gb Mode	c1x1–x4.. c8x1–x4	1..32
PRT-HC3-C08Q16–4x25Gb Mode on Control Card Version 2	c1x1–x4.. c8x1–x4	1..32
BPS-HC3-C25F2G	x1..x16	1..16
	c1..c4	17, 21, 25, 29
BPS-HC3-C35F2G	x1..x16	1..16
	c1..c4	17, 21, 25, 29
BPS-HC3-Q35F2G	x1..x16	1..16
	q1..q4	17, 21, 25, 29
PRT-HC3-C16	c1..c16	1, 5, 6, 10, 11, 15, 16, 20, 21, 25, 26, 30, 31, 35, 36, 40
PRT-HC3-C16–4x10Gb Mode	c1x1..x4	1..4
	c3x1..x4	6..9
	c5x1..x4	11..14
	c7x1..x4	16..19
	c9x1..x4	21..24
	c11x1..x4	26..29
	c13x1..x4	31..34
	c15x1..x4	36..39
PRT-HC3-C16–4x25Gb Mode	c1x1..x4	1..4
	c3x1..x4	6..9
	c5x1..x4	11..14
	c7x1..x4	16..19
	c9x1..x4	21..24
	c11x1..x4	26..29
	c13x1..x4	31..34
	c15x1..x4	36..39

GigaVUE Device / Line Card	Ports	Port IDs Inserted in Gigamon Trailer
GigaVUE-HC1 Gen 3		
HC1-X12G4	g1..g4	1..4
	x1..x12	5..16
PRT-HC1-X12	x1..x12	1..12
PRT-HC1-Q04X08	c1..c4	25, 29, 37, 33
	x1..x8	41..48
GigaVUE-HC3 Gen 3		
PRT-HC3-X24	x1..x24	1..24
PRT-HC3-C08Q08	c1..c8	1, 5, 9, 13, 17, 21, 25, 29
SMT-HC3-C08	c1..c8	1, 5, 9, 13, 17, 21, 25, 26
PRT-HC3-C16	c1..c16	1, 5, 6, 10, 11, 15, 16, 20, 21, 25, 26, 30, 31, 35, 36, 40
GigaVUE-HC1-Plus		
HC1P-C04X08	c1..c4	1, 5, 9, 13
	x1..x8	17..24
PRT-HC1-X12	x1..x12	1..12
PRT-HC1-Q04X08	c1..c4	25, 29, 33, 37
	x1..x8	41..48
GigaVUE-HCT		
HCT-C02	c1..c2	1, 5
PRT-HC1-Q04X08	q1..q4	1, 5
	x1..x8	
PRT-HC1-G12	g1..g12	1..12
GigaVUE TA Series		
GigaVUE-TA25	x1..x48	1..48

GigaVUE Device / Line Card	Ports	Port IDs Inserted in Gigamon Trailer
	c1..c8 c1x1..c1x4 c5x1..c5x4	49,53,57,61,65,69,73,77 49..52 65..68
GigaVUE-TA25E	x1..x48	1..48
	c1..c4	49, 53, 57, 61, 65, 69, 73, 77
GigaVUE-TA100	c1..c32	1, 5, 9, 13, 17, 21, 25, 29, 33, 37, 41, 45, 49, 53, 57, 61, 65, 69, 73, 77, 81, 85, 89, 93, 97, 101, 105, 109, 113, 117, 121, 125
	c1x1..c32x4	1..128
GigaVUE-TA200	c1..c32	1, 5, 9, 13, 17, 21, 25, 29, 33, 37, 41, 45, 49, 53, 57, 61, 65, 69, 73, 77, 81, 85, 89, 93, 97, 101, 105, 109, 113, 117, 121, 125
	c33..c64	129..160
GigaVUE-TA200E	c1..c32	1, 5, 9, 13, 17, 21, 25, 29, 33, 37, 41, 45, 49, 53, 57, 61, 65, 69, 73, 77, 81, 85, 89, 93, 97, 101, 105, 109, 113, 117, 121, 125
	c33..c64	1..128
GigaVUE-TA400	d1..d32	129..160
GigaVUE-TA400E	d1..d32	1, 5, 9, 13, 17, 21, 25, 29, 33, 37, 41, 45, 49, 53, 57, 61, 65, 69, 73, 77, 81, 85, 89, 93, 97, 101, 105, 109, 113, 117, 121, 125
	x1..x12	129,130

GigaSMART Trailer Example

The following figures show a sample GigaSMART trailer with the packet length and source ID (Source port label) TLV included. The total length of the trailer in this example is 10 bytes, plus another 5 bytes for the Trailer Length, Ethertype, and Checksum fields.

GIGAMON_PKT_ LEN TLV	GIGAMON_SRCID TLV	Trailer Length, Ethertype, Checksum

GIGAMON_PKT_LEN=1 (1 Byte)	TAG_LEN=2 (1 Byte)	PKT_LEN (2 Bytes)	GIGAMON_SRCID Type (1 Byte)	GIGAMON_SRCID Length (1 Byte)	Platform (6 Bits)	Group ID (4 Bits)	Box ID (6 Bits)	Slot ID (6 Bits)	Port ID (10 Bits)	Length =10 (0a) (1 Byte)	Ethertype=0x22E5 (2 Bytes)	Checksum (2 Bytes)

GIGAMON_SRCID Value (4 bytes)

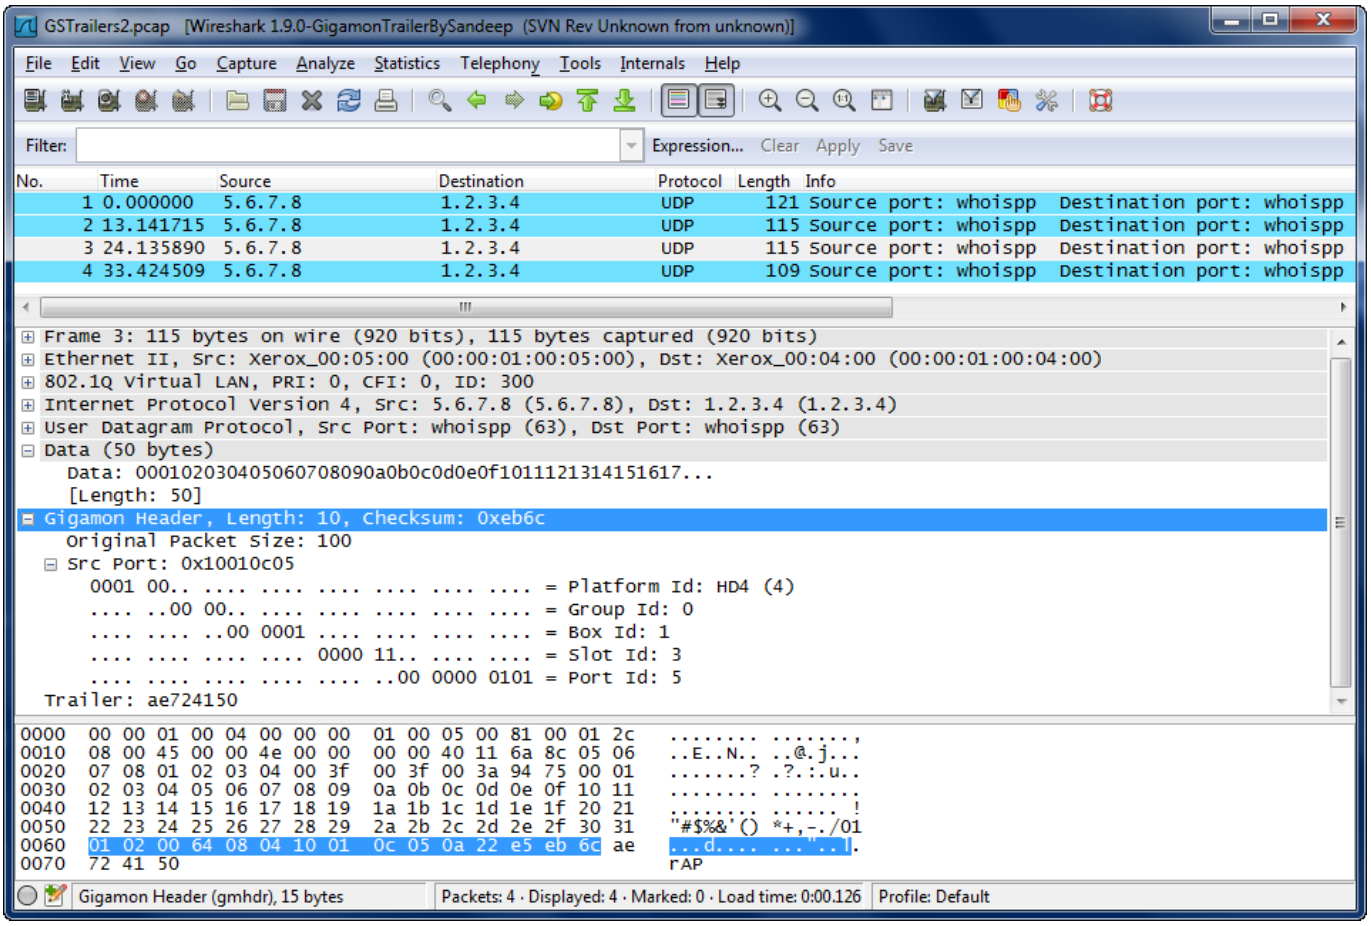


Figure 123 *Wireshark Decode of GigaSMART Trailer*

Limitations

- Group ID will be set to zero even for the devices in a cluster.
- The Trailer Addition/Removal GSOP does not support E-Tag clusters. GigaVUE-TA400, GigaVUE-TA400E and GigaVUE-HCT platforms support only E-Tag clustering. However, you can configure the GSOP on GigaVUE-HCT platform when operating in standalone mode.

- In Gen2 cards, application chaining, a combination of Trailer Addition with Tunnel Decapsulation, is currently unsupported within the same GSOP. To achieve the intended functionality, you can configure two distinct GSOPs, one for decapsulation and another for trailer addition linked by maps.

GigaSMART Logs

As of 5.4, GS Log files enable you to generate and download application-specific logging information to use for troubleshooting problematic applications. Gigamon Support can use these files for root cause analysis.

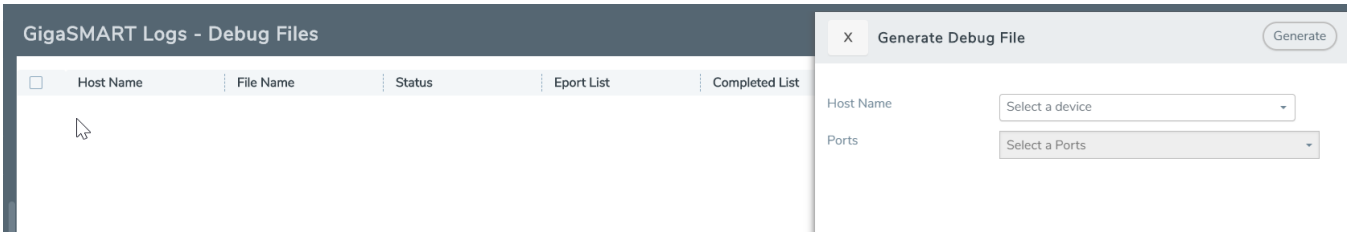
GS Log levels are applied at the process level by specifying **info** or *debug* when configuring the logging level for the application's processes. Logging levels modified at the application level will take priority over the system level setting.

To access Logs, click  on the top navigation bar. On the left navigation pane, select **System > Logs**.

Create GS Log file

To create a log file that Gigamon can use for analysis, do the following:

1. Access a device: Physical > Physical Node > click the Cluster ID for a node.
2. Under Support in the navigation pane, select **Debug > GigaSMART Logs**.
3. Click **Generate**.
4. Select the Host Name and Ports.



   The system generates a new GS Log file.
5. Select the GS Log file to download, and then click **Download**.

   You can only download one file at a time.

The system downloads the file to your local environment with a name like, gsdump_<hostname>_<date>_<time>.tgz.gpg. The file is in a compressed and encrypted format that you can provide to Gigamon.

Delete Log File

To delete the GS Log files for clearing up the disk space:

1. Select **Debug** > **GS Logs**.
2. Select the GS Logs that you want to delete and click **Delete**.

Upgrade GigaSMART Cards

This section describes the steps to upgrade the software on GigaSMART cards using GigaVUE-FM for GigaVUE ® HC Series. You can upgrade your GigaSMART card using an image located on an external image server or on your local server or machines.

Refer to the following sections:

- Upgrade GigaSMART card using External Image Server
- Upgrade GigaSMART card using Internal Image Server

Upgrade GigaSMART card using External Image Server

This section provides the steps for upgrading the GigaSMART card from an image stored on an external server. The image can be transferred from the server using the SCP file protocol.

To upgrade a GigaSMART card with an image stored on an external image server, you must sequentially perform the tasks in the following table:

S.No	Task	Refer to
1	Upload the image to the external image server	Upload image to external image server
2	Add the external server to GigaVUE-FM	Add the external server to GigaVUE-FM
3	Upgrade GigaSMART using external image	Upgrade GigaSMART card using External Image Server

Upload image to external image server

1. Upload the image to the external image server to make it available to GigaVUE-FM. To obtain software images and download the software, register on the Gigamon Customer Portal
2. Add the image server to GigaVUE-FM. This will store the server's credentials, image file name, and IP address on GigaVUE-FM.

Add the external server to GigaVUE-FM

To add the image from the external server, perform the following steps:

1. Go to **Settings**.
2. Go to **System** > **Images** > **External Servers**.
3. Click **Add**. The **Add External Server** page appears.
4. In the Add External Server page enter the following details:
   - An alias to help identify the image server.
   - The host IP address of the server.
   - The protocol to use for the download: SCP.
   - The user name and password.
5. Click **OK**.

   The External Server page displays the newly added external server.

Upgrade GigaSMART card using External Image

1. Go to **Inventory** > **Physical** > **GigaSMART Upgrade**.
2. Select one or more nodes or clusters to upgrade.
3. Click **Upgrade** to view the **GigaSMART Upgrade** page.

   The **GigaSMART Upgrade** page provides the node details such as **Host Name**, **IP Address**, **Role**, and **Version along with GigaSMART** card details.
4. Enter the name in the **Task Name** field.
5. Select the radio button **External Image Server** from the **Image Server**.
6. Select the server in which the image is stored from the **Image Server** drop-down list box.

NOTE: If the image server is not available in the Image Server drop-down list box, Click **Add External Server** and enter the details provided in **Step 4** in the section Add the external server to GigaVUE-FM

7. Enter the path used for storing the image as required.

For Example:

- **GigaVUE-HC3 (GEN2 GigaSMART File Path)**
- **GigaVUE-HC1 (GEN3 GigaSMART File Path)**

8. Enter the time for performing the upgrade. The following are the options available:

- Immediate—The upgrade is performed immediately.
- Scheduled—The upgrade is performed at a scheduled time. Select the date and time.

9. Select the required cards from the devices to upgrade.

10. . Click **Upgrade**.

GigaVUE-FM notifies the various stages of the upgrade and also the upgrade process's completion.

Upgrade GigaSMART card using Internal Image Server

This section provides the steps for upgrading GigaSMART card when the image is available on the internal server.

To upgrade a GigaSMART using internal image files, you must sequentially perform the tasks in the following table:

S.No	Task	Refer to
1	Download the Image	Download Image
2	Upload the image file to GigaVUE-FM	Upload the image file to GigaVUE-FM
3	Upgrade GigaSMART using Internal image	Upgrade GigaSMART card using Internal Image Server

Download Image

Download the images from the Gigamon website and place them where they can be uploaded to GigaVUE-FM. Register on the Gigamon Customer Portal to obtain software images and download the software.

Upload the image file to GigaVUE-FM

To upload the images file to GigaVUE-FM, follow these steps:

1. Go to **Settings**.
2. Go to **System > Images > Internal Image Files** and do the following:

a. On the Internal Image File page, click **Upload**.

b. Click **Choose** to locate the image file on the **Upload Internal Image Files** page.

c. Click **OK** to upload the file. The page displays the progress of the upload.

After the upload, the GigaSMART card image for the upgrade is available on the Internal Images Files page.

Upgrade GigaSMART card using Internal image

1. Go to **Inventory** > **Physical** > **GigaSMART Upgrade**.
2. Select one or more nodes or clusters to upgrade Nodes that need to be upgraded.
3. Click **Upgrade** to view the **GigaSMART Upgrade** page.

   The **GigaSMART Upgrade** page provides the node details such as **Host Name**, **IP Address**, **Role**, and **Version** along with GigaSMART card details.
4. Select the radio button **Internal Image Server** from the **Image Server Type**.
5. Select the version you that you are upgrading from the **Version** drop-down list box.
6. Enter the time for performing the upgrade. The following are the options available:

- Immediate—The upgrade is performed immediately.
- Scheduled—The upgrade is performed at a scheduled time. Select the date and time.

7. Select the required cards from the devices to upgrade.
8. Click **Upgrade**.

NOTE: If the image server is not available in the Image Server drop-down list box, Click **Add Internal Server** and perform the actions provided in **Step 2** in the section Upload the image file to GigaVUE-FM

NOTE: You can only upgrade to another instance of the current version. Downgrading to a lower version is not supported through the UI.

GigaVUE-FM notifies the various stages of upgrade and upgrade process's completion.

NOTE: When you change the IP address of the GigaVUE-FM instance using the jump-start configuration, the internal database and the in-memory caches of the GigaVUE-FM instance are not updated. The Database continues to have the IP address of the old GigaVUE-FM, and the image upgrade using the internal server option does not work. To fix this, you must restart the GigaVUE-FM instance after upgrade.

Go to **Settings** and select **Events**, to monitor the progress and status of the upgrade. Also, email notifications are sent if email notifications have been configured.

Additional Sources of Information

This appendix provides additional sources of information. Refer to the following sections for details:

- Documentation
- Documentation Feedback
- Contact Technical Support
- Contact Sales
- The VÜE Community

Documentation

©This table lists all the guides provided for GigaVUE software and hardware. The first row provides an All-Documents Zip file that contains all the guides in the set for the release.

NOTE: In the online documentation, view What's New to access quick links to topics for each of the new features in this Release; view Documentation Downloads to download all PDFs.

Table 1: Documentation Set for Gigamon Products

GigaVUE 6.14 Hardware and Software Guides
DID YOU KNOW? If you keep all PDFs for a release in common folder, you can easily search across the doc set by opening one of the files in Acrobat and choosing **Edit > Advanced Search** from the menu. This opens an interface that allows you to select a directory and search across all PDFs in a folder.
Hardware how to unpack, assemble, rackmount, connect, and initially configure ports the respective GigaVUE devices; reference information and specifications for the respective GigaVUE devices
GigaVUE-HC1 Hardware Installation Guide
GigaVUE-HC3 Hardware Installation Guide
GigaVUE-HC1-Plus Hardware Installation Guide
GigaVUE-HCT Hardware Installation Guide
GigaVUE-TA25 Hardware Installation Guide
GigaVUE-TA25E Hardware Installation Guide
GigaVUE-TA100 Hardware Installation Guide

GigaVUE 6.14 Hardware and Software Guides
GigaVUE-TA200 Hardware Installation Guide
GigaVUE-TA200E Hardware Installation Guide
GigaVUE-TA400 Hardware Installation Guide
GigaVUE-TA400E Hardware Installation Guide
GigaVUE-OS Installation Guide for DELL S4112F-ON
G-TAP A Series 2 Installation Guide
GigaVUE M Series Hardware Installation Guide
GigaVUE-FM Hardware Appliances Guide
Software Installation and Upgrade Guides
GigaVUE-FM Installation, Migration, and Upgrade Guide
GigaVUE-OS Upgrade Guide
GigaVUE V Series Migration Guide
Fabric Management and Administration Guides
GigaVUE Administration Guide covers both GigaVUE-OS and GigaVUE-FM
GigaVUE Fabric Management Guide how to install, deploy, and operate GigaVUE-FM; how to configure GigaSMART operations; covers both GigaVUE-FM and GigaVUE-OS features
GigaVUE Application Intelligence Solutions Guide
GigaVUE Inline Solutions Guide(**NEW**) (previously included in the **GigaVUE Fabric Management Guide**)
Cloud Guides how to configure the GigaVUE Cloud Suite components and set up traffic monitoring sessions for the cloud platforms
GigaVUE V Series Applications Guide
GigaVUE Cloud Suite Deployment Guide - AWS
GigaVUE Cloud Suite Deployment Guide - Azure
GigaVUE Cloud Suite Deployment Guide - OpenStack
GigaVUE Cloud Suite Deployment Guide - Nutanix
GigaVUE Cloud Suite Deployment Guide - VMware (ESXi)
GigaVUE Cloud Suite Deployment Guide - VMware (NSX-T)

GigaVUE 6.14 Hardware and Software Guides
GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration
Universal Cloud TAP - Container Deployment Guide
Gigamon Containerized Broker Deployment Guide
GigaVUE Cloud Suite Deployment Guide - AWS Secret Regions
GigaVUE Cloud Suite Deployment Guide - Azure Secret Regions
Reference Guides
GigaVUE-OS CLI Reference Guide library of GigaVUE-OS CLI (Command Line Interface) commands used to configure and operate GigaVUE HC Series and GigaVUE TA Series devices
GigaVUE-OS Security Hardening Guide
GigaVUE Firewall and Security Guide
GigaVUE Licensing Guide
GigaVUE-OS Cabling Quick Reference Guide guidelines for the different types of cables used to connect Gigamon devices
GigaVUE-OS Compatibility and Interoperability Matrix compatibility information and interoperability requirements for Gigamon devices
GigaVUE-FM REST API Reference in GigaVUE-FM User's Guide samples uses of the GigaVUE-FM Application Program Interfaces (APIs)
Factory Reset Guidelines for GigaVUE-FM and GigaVUE-OS Devices Sanitization guidelines for GigaVUE Fabric Management Guide and GigavUE-OS devices.
Release Notes
GigaVUE-OS, GigaVUE-FM, GigaVUE-VM, G-TAP A Series, and GigaVUE Cloud Suite Release Notes new features, resolved issues, and known issues in this release ; important notes regarding installing and upgrading to this release **Note**: Release Notes are not included in the online documentation. **Note**: Registered Customers can log in to My Gigamon to download the Software and Release Notes from the Software and Docs page on to My Gigamon. Refer to How to Download Software and Release Notes from My Gigamon.
In-Product Help
GigaVUE-FM Online Help how to install, deploy, and operate GigaVUE-FM.

How to Download Software and Release Notes from My Gigamon

Registered Customers can download software and corresponding Release Notes documents from the **Software & Release Notes** page on to My Gigamon. Use the My Gigamon Software & Docs page to download:

- Gigamon Software installation and upgrade images,
- Release Notes for Gigamon Software, or
- Older versions of PDFs (pre-v5.7).

To download release-specific software, release notes, or older PDFs:

1. Log in to My Gigamon.
2. Click on the **Software & Release Notes** link.
3. Use the **Product** and **Release** filters to find documentation for the current release. For example, select Product: "GigaVUE-FM" and Release: "5.6," enter "pdf" in the search box, and then click **GO** to view all PDF documentation for GigaVUE-FM 5.6.xx.

NOTE: My Gigamon is available to registered customers only. Newer documentation PDFs, with the exception of release notes, are all available through the publicly available online documentation.

Documentation Feedback

We are continuously improving our documentation to make it more accessible while maintaining accuracy and ease of use. Your feedback helps us to improve. To provide feedback and report issues in our documentation, send an email to: documentationfeedback@gigamon.com

Please provide the following information in the email to help us identify and resolve the issue. Copy and paste this form into your email, complete it as able, and send. We will respond as soon as possible.

Documentation Feedback Form		
About You	**Your Name**	
	Your Role	
	Your Company	

For Online Topics	**Online doc link**	*(URL for where the issue is)*
	Topic Heading	*(if it's a long topic, please provide the heading of the section where the issue is)*
For PDF Topics	**Document Title**	*(shown on the cover page or in page header)*
	Product Version	*(shown on the cover page)*
	Document Version	*(shown on the cover page)*
	Chapter Heading	*(shown in footer)*
	PDF page #	*(shown in footer)*
How can we improve?	**Describe the issue**	*Describe the error or issue in the documentation.* *(If it helps, attach an image to show the issue.)*
	How can we improve the content? **Be as specific as possible.**	
	Any other comments?	

Contact Technical Support

For information about Technical Support: Go to **Settings**  **> Support > Contact Support** in GigaVUE-FM.

You can also refer to https://www.gigamon.com/support-and-services/contact-support for Technical Support hours and contact information.

Email Technical Support at support@gigamon.com.

Contact Sales

Use the following information to contact Gigamon channel partner or Gigamon sales representatives.

Telephone: +1.408.831.4025

Sales: inside.sales@gigamon.com

Partners: www.gigamon.com/partners.html

Premium Support

Email Gigamon at inside.sales@gigamon.com for information on purchasing 24x7 Premium Support. Premium Support entitles you to round-the-clock phone support with a dedicated Support Engineer every day of the week.

The VÜE Community

The VÜE Community is a technical site where Gigamon users, partners, security and network professionals and Gigamon employees come together to share knowledge and expertise, ask questions, build their network and learn about best practices for Gigamon products.

Visit the VÜE Community site to:

- Find knowledge base articles and documentation
- Ask and answer questions and learn best practices from other members.
- Join special-interest groups to have focused collaboration around a technology, use-case, vertical market or beta release
- Take online learning lessons and tutorials to broaden your knowledge of Gigamon products.
- Open support tickets (Customers only)
- Download the latest product updates and documentation (Customers only)

The VÜE Community is a great way to get answers fast, learn from experts and collaborate directly with other members around your areas of interest.

Register today at community.gigamon.com

Questions? Contact our Community team at community@gigamon.com.

Glossary

D

decrypt list

need to decrypt (formerly blacklist)

decryptlist

need to decrypt - CLI Command (formerly blacklist)

drop list

selective forwarding - drop (formerly blacklist)

F

forward list

selective forwarding - forward (formerly whitelist)

L

leader

leader in clustering node relationship (formerly master)

M

member node

follower in clustering node relationship (formerly slave or non-master)

N

no-decrypt list

no need to decrypt (formerly whitelist)

nodecryptlist

no need to decrypt- CLI Command (formerly whitelist)

P

primary source

root timing; transmits sync info to clocks in its network segment (formerly grandmaster)

R

receiver

follower in a bidirectional clock relationship (formerly slave)

S

source

leader in a bidirectional clock relationship (formerly master)