Gigamon®

GigaVUE Fabric Management Guide

GigaVUE-FM

Product Version: 6.9

Document Version: 1.2

(See Change Notes for document updates.)

Copyright 2025 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
6.9	1.2	04/04/2025	The original release of this document with 6.9 GA. This update includes bug fixes and minor cosmetic changes for improved usability and document consistency.
6.9	1.1	03/03/2025	This update includes bug fixes and minor cosmetic changes for improved usability and document consistency. • Updated the images in the following sections: o About Cluster o Example 1: Unprotected Inline Bypass with an Inline Tool Group o Sample Stack-Link Configurations o Orchestrated Configurations o Links o View Clusters and Nodes in the Topology o GigaSMART IP Encapsulation/Decapsulation (GigaSMART Tunnel) o Configure GigaSMART IP Encapsulation/Decapsulation o GigaVUE Modules for Inline SSL Decryption o Policy Profile Options • A note is added for more clarity on UDE in the GigaVUE Fabric Management Guide. • Updated theLicensing information of Application Intelligence section.
6.9	1.0	12/06/2024	The original release of this document with 6.9.00 GA.

Contents

GigaVUE Fabric Management Guide **1**
Change Notes 3
Contents 4
Fabric Management **47**
GigaVUE-FM **47**
Supported Nodes **48**
Device and Configuration Management **49**
How to enable Web Server for Node Management **50**
Supported Browsers for GigaVUE-FM 50
GigaVUE-FM Features and Benefits **51**
Get Started with GigaVUE-FM **53**
Log In to GigaVUE-FM **54**
Log Out of GigaVUE-FM **54**
GigaVUE-FM **55**
Homepage 55
Left Navigation Pane 55
Top Navigation Bar 56
Page-Level Header 56
GigaVUE-FM GUI Navigation 57
Footer 57
Theme Settings 57

Configure a Custom Banner **58**
Quick Views **58**
Recently Viewed **58**
Find Pages **59**
Return to the Dashboard **60**
Table View Customization **60**
Notifications Panel **61**
Long-term Notifications **62**
Notification Type Icons **63**
Notification Banners **63**
Notification Banners for Disk Space Utilization 64
How to Add the GigaVUE-FM Instance Name **66**
How to Search in GigaVUE-FM **67**
Performing a Search **68**
Search Examples **69**
Searching Maps 70
Example 1: Searching for a Map by Alias 70
Example 2: Searching for a Map with an IP Address 71
Example 3: Searching for a Map with a MAC Address 73
Example 4: Searching for Maps with Down Ports 73
Searching for Roles and Users 74
Example 1: Searching for Monitor Role 74
Example 2: Searching for a User 76
Searching Ports 77
Example 1: Searching for Down Ports 77
Example 2: Searching for Port Details of Devices Managed by GigaVUE-FM 78

Filtering Search Results **79**
How to Apply Filters **79**
Dashboard **81**
Physical and Virtual Dashboard **81**
Overview of the Physical and Virtual Dashboard **82**
Physical Dashboard Profiles **82**
Physical Dashboard Quick Views **84**
Physical Dashboard Widgets **84**
Highest Traffic **85**
Lowest Traffic **88**
Traffic Comparison By Tags **89**
Most Utilized Traffic **93**
Least Utilized Traffic **96**
Inventory **97**
Nodes by Model 97
Nodes by Software Version 98
Status Summary **100**
Nodes' Status Summary 100
Port Link Status Summary 101
Unhealthy Maps 102
Unhealthy Flows 103
Port Drops and Errors 103
Unhealthy Fabric Maps 106
Audit Logs 106
Events 109

FM Health Dashboard **112**
Overview of GigaVUE-FM Health Dashboard **113**
CPU Utilization **114**
Memory Utilization **114**
Storage Utilization **115**
Alarm Thresholds and Notifications **116**
GigaVUE-FM Reference Materials **117**
Disk Size on GigaVUE-FM **118**
Increase Disk Size on a New or Existing GigaVUE-FM Installation on KVM **118**
How to Clean up Disk Space on a GigaVUE-FM Instance **121**
Data Transfer Rate Units **122**
Open Ports in GigaVUE-FM **122**
Open Ports in HC Series Devices **126**
Health Status **128**
How GigaVUE-FM Computes Health Status **128**
Node Health Status **130**
Port Health Status **132**
Map Health Status **134**
GigaSMART Map Health Status **135**
GigaSMART Group Health Status 135
vPort Health Status 137
GigaSMART Operations Health Status 137
IP Interfaces Health Status 137
Flow Health Status **138**
Priority Map Set Health **139**
Flow Health Computation **140**
GigaVUE-FM APIs **141**
Supported Cloud Environments **141**
Analytics **143**
Rules, Notes, and Limitations for Analytics 143
GigaVUE-FM Statistics and Data Roll-up 144
Control Filters in Analytics Dashboards 145

Get Started with Analytics UI **146**
Work with the Analytics User Interface **149**
Share 149
Clone 150
Reporting 150
Set as Default 150
Copy Dashboard Path 150
Auto Refresh Tags in Visualizations 151
Filter Data Using Tags in Control Filters 151
Search Data 154
Filter Data in Visualizations **154**
Visualizations **157**
View System Visualization 158
Create Custom Visualizations **158**
Create a Visualization 158
Clone a Visualization 161
Visualizations - Example Work Flows **161**
TSVB Chart Displaying Traffic Trend 161
Pie Chart Displaying Alarms Summary 163
Metric Displaying Card Count 165
Bar Chart for Alarms by Severity 166
Dashboards **168**
Clone Dashboard **169**
Create New Dashboard **170**
Edit Dashboard **173**
Reports **174**
Discover **176**
Filter 178
Save Search 178
Find Data 179
Default Dashboards **179**
Dashboard **193**
Dashboards for Volume-based Licenses Usage **194**
Virtual Inventory Statistics and Cloud Applications Dashboard **197**
Rules and Notes 203

FHA Dashboards for 5G-Cloud Applications **203**
GigaSMART Mobility Session and Flow Filtering Dashboards **213**
How to Access the Dashboards 213
Rules and Notes 216
GigaSMART Inline TLS/SSL Dashboards **216**
Basic Dashboards 217
Advanced Dashboards 224
System Requirements 224
Configure Advanced dashboard 225
Rules and Notes 225
OpenSearch Indices and Fields Used in Analytics **230**
Deprecated Fields in Analytics **233**
Analytics Appendix **234**
Ports 234
Maps 235
Map Rule 236
Nodes and Clusters **238**
GigaVUE Nodes and Clusters **238**
GigaVUE® HC Series and TA Series Overview **238**
Notes on TA Series Nodes 241
About Cluster **242**
Overview of Seed Node **242**
Manage GigaVUE® Nodes and Clusters **243**
Configure Physical Nodes **244**
Node Control Options 246
Add New Physical Node or Cluster to GigaVUE-FM **247**
Add Nodes Manually 248
Add Nodes From an Excel Spreadsheet 249
Device Level Tagging **251**
NAT Behind Nodes **251**
Prerequisites 251
Add NAT Behind Nodes 252
Limitations 252
Cluster Discovery Behavior **252**
ARP/NDP Timer Settings **253**
Change the ARP/NDP Timer Settings 253

Enable Gratuitous ARP on Management Interface **254**
SNMPv3 Support **255**
Enable SNMPv3 on Nodes 255
Create and Manage Clusters **255**
About Cluster **256**
Cluster Node Limit **257**
Cluster Scaling 258
Cluster Topologies 258
Separate Paths for Cluster Control and Stack Traffic **258**
About Cluster Roles **260**
Setting a Node's Priority in the Leader Election Process 260
About the "Unknown" Cluster Role 261
Sample Cluster Control Connections **262**
Sample Cluster Control Configurations 262
Zeroconf for Cluster Management Ports 262
Keep Cluster Management Ports Connected! 262
Sample Stack-Link Configurations **263**
Creating Clusters: A Roadmap **264**
IPv6 Based Clustering **266**
Rules, Notes, and Limitations 266
Create IPv6 based Cluster 266
Switch IPv4 based Cluster to IPv6 based Cluster 267
Rules and Recommendations for Nodes and Clusters **269**
Using Command Line Interface for Managing Clusters 271
GigaVUE TA Series and GigaVUE-HC3 Clustering Recommendations **274**
Cluster Rules **274**
Best Practices for OOB Clusters with IGMP Snooping **275**
About IGMP Snooping in a Cluster 275
Allow IGMP Traffic 276
Enable an IGMP Querier 276
E-Tag Clustering **277**
Stacking Mode 278
Switching Stacking Mode 279
Notes 279
New 6.0 280

Cluster Safe and Limited Modes **280**
Safe Mode **281**
Limited Mode **281**
Support for Cluster Types **282**
Create Clusters **282**
Regular Cluster Formation Workflow **282**
Deployment Checklist 282
Pre-deployment checklist 282
Create Regular Cluster Formation 283
Select Cluster Type 283
Select Devices 283
Customize Devices 284
Customize Stack Links 285
Edit Cluster **286**
Prerequisites: 287
Inband Cluster Management **287**
Inband Cluster Management Topologies **288**
Inband Cluster Management Stack Ports **288**
Inband Cluster Management Stack Ports Example 289
Inband Cluster Management Configuration Flow Chart **290**
Inband Cluster Management Configuration **292**
Enable Cluster Management for GigaVUE TA Series Nodes **292**
Add Nodes to a Cluster **292**
Remove Nodes from a Cluster **293**
Edit Cluster Parameters **294**
Check Cluster Status **295**
Cluster Management Events 295
Audit Logs 296

Export Nodes and Clusters **296**
Problems with SCP? **297**
Events **297**
Alarms **298**
Audit Logs **298**
Search for Specific Nodes Using Keywords **299**
Search for Ports on a GigaVUE Node **299**
Overview Page **301**
Systems Information **301**
Ports Information **302**
Traffic **303**
Workflows **303**
Overview of Workflows **303**
How to Use Workflows **304**
Chassis Table View **306**
Live Graphing **307**
Safe and Limited Modes **307**
Safe Mode **308**
Limited Mode **309**
Enable SNMP Trap for Safe Mode and Limited Mode **309**
Collect Information for Technical Support **310**
Manage Not-reachable Nodes in Cluster **310**
Rules, Notes, and Limitations 312
Remove Offline Chassis 312
Alarms and Health Status of Not-reachable Nodes 313
Upgrade Cluster with Not-reachable Nodes 314
Backup and Restore of Cluster with Not-reachable Nodes 315
Multi-Path Leaf and Spine **316**
Introduction to Multi-Path Leaf and Spine **316**
Path Protection **318**
Leaf Node Failure 318
Restoration 318
Affected Time 318
Stack Link Failure on Leaf (TAP Connected) 319
Restoration 319

Affected Time 319
Spine Node Failure 319
Restoration 320
Affected Time 320
Stack Link Failure on Leaf (Tool Connected) 320
Configuration Overview 321
Notes and Considerations 322
Leaf-Spine Cluster Deployment 323
Deployment Checklist 323
Pre-deployment checklist 323
Formation Scenario 323
Leaf-Spine Cluster Formation Workflow 324
Create a Leaf-Spine Cluster 324
Select the Cluster Type 325
Select Devices 325
Customize Devices 326
Customize Stack Links for your Leaf-Spine configuration 327
Edit a Cluster 328
Prerequisites 328
Delete a Node from a Cluster 330
How to Change the Leader Preference of a Device 330
Spine to Spine and Leaf 331
Introduction to Spine to Spine and Leaf 331
Configuration Overview 333
Notes and Considerations 335
Configuration of Spine to Spine and Leaf Architecture 335
Limitations 335
Leaf-Spine Cluster Deployment 336
Deployment Checklist 336
Documentation 336
Pre-deployment checklist 336

Formation Scenario **337**
Fabric Statistics **337**
About Fabric Statistics **338**
Display Fabric Statistics for All Ports **338**
Display Fabric Statistics for a Single Port **341**
Port Quick View **341**
Export Fabric Statistics **342**
Filter Fabric Statistics **342**
Clear Fabric Statistics Counters 344
Enable Discovery Protocols **344**
Enable Gigamon Discovery on Chassis **345**
Enable LLDP and CDP on Chassis **345**
Enable LLDP, CDP, and Gigamon Discovery on Ports **346**
Limitations of Gigamon Discovery **347**
Topology Visualization **347**
Edit Topologies 353

Configure Network Devices **361**
Configure Connections **363**
Configure Tools **364**
Flows **367**
About Flows **368**
View Flows **371**
View the Flow Summary and Statistics **372**
View Maps and Ports **375**
View Total Ports **376**
View Total Unhealthy Ports **376**
View Total Maps **377**
View Unhealthy Maps **377**
Filter Flows **377**
How to Change the Flow Layout **378**
How to Update Flows **378**
View Events **379**
Set Notifications **379**
Limitations of Flows **380**
Device Logs and Event Notifications **380**
Stream Device Logs to GigaVUE-FM **381**
Cluster Behavior **381**
Standardized Logs **381**
Device Log Categories 382
Device Log Message Types 382
Device Logging Levels 383
Device Logging Processes 384
View Device Logs **385**
Arrange Columns in the Logs View **387**
Device Log Host Servers **388**
Add an External Logging Host Server to a Node **388**
Host Server Options 389

Edit Host Server Settings **390**
Storage Management for Device Logs **391**
Access Storage Management **391**
Manage Device Log Output **392**
Traffic Filtering **394**
Ports and GigaStreams **395**
About Ports **395**
About Network and Tool Ports **395**
Network (Ingress) Ports Defined 396
Tool (Egress) Ports Defined 396
Ports on GigaVUE® TA Series Traffic Aggregator Nodes 397
Hybrid Ports 399
Stack Ports 399
Inline Network Ports 400
Inline Tool Ports 400
Circuit Ports 400
GigaSMART Engine Ports 400
Port Lists **401**
Port Aliases **401**
Work with Hybrid Ports **401**
Port Filters **404**
Port Filter—Rules and Notes 404
Port-Filter Maximums 404
How to Apply Port Filters 406
View Port Filter Statistics 407
View Filter Resources for a Slot 408
Status of Line Cards/Nodes and Ports **409**
How to Check Port Status with Ports Page 409
How to Check Port Status with Chassis Page 410
Managing Ports **410**
Ports **411**
All Ports 411
Port Quick View 414
Port List Filter 415
Quick Port Editor 417
Configure Ports 418
Ports Discovery 420
Statistics 421

Port Groups **422**
All Port Groups 422
Create Port Groups 423
Edit Port Group 423
Clone Port Group 423
Port Group Statistics 425
Port Pairs **425**
Create Port Pair 425
Tool Mirrors **427**
Create Tool Mirror 427
Edit Tool Mirror Description 428
Clone Tool Mirror 428
Stack Links **429**
IP Interfaces **430**
Configure IP Interface 430
Prerequisites 430
IP Interface Statistics 432
Circuit Tunnels **432**
Port Discovery **432**
Port Discovery with LLDP and CDP **433**
Enable Port Discovery **434**
Limits of Discovery Information **435**
Port Discovery Support **435**
Port Discovery for a Cluster 435
Port Discovery Supported for SNMP 435
Ingress and Egress VLAN **436**
About Ingress Port VLAN Tagging **437**
Ingress Port VLAN Tagging 437
Adding VLAN Tags 438
Deleting VLAN Tags 438
Using VLAN Tags in Maps **438**
Ingress Port VLAN Tag Limitations **439**
Second Level Maps 439
Double-Tagged Packets 439
IP Interfaces 440
Local Tool Port Ingress VLAN Tag 440
Configure Egress Port VLAN Stripping **440**
Enable Egress Port VLAN Stripping 441

Disable Egress Port VLAN Stripping 441
Display Egress Port VLAN Stripping 441
Egress Port VLAN Stripping Limitations 441
How to Use Both Ingress Tagging and Egress Stripping 442
How to Use GigaStream 443
GigaStreams 443
Accessing the GigaStream page 443
About GigaStream 445
Regular GigaStream 445
Controlled GigaStream 445
Regular GigaStream 446
Regular Tool GigaStream 446
Regular Stack GigaStream 447
Configure Regular GigaStream 448
Edit Regular GigaStream 449
Edit Regular Stack GigaStream 450
Traffic Distribution Across Regular GigaStream 450
Regular GigaStream Failover Protection—Resiliency 451
Regular Circuit GigaStream 451
Controlled GigaStream 452
Notes and Considerations for Controlled GigaStream 454
Controlled GigaStream Configuration 455
Edit Controlled GigaStream 457
Traffic Distribution Across Controlled GigaStream 458
Failover and Controlled GigaStream 460
Advanced Hashing 461
How to Change Advanced Hash Criteria 461
Advanced Hash Settings 461
Advanced Hash Examples 463
Hashing Behavior 465
Notes and Considerations for Advanced Hashing 467
Advanced Hashing with MPLS 467
Advanced Hashing with GTP TEID 469
Packet Distribution and the Advanced Hash Algorithm 471
Weighted GigaStream 471
GigaStream Rules and Maximums 472
Maximum Ports per GigaStream 473

Port Statistics and Counters **474**
Display Port Statistics **474**
Display Port Statistics for a Single Port 474
Display Statistics for All Ports 476
How to Clear Traffic Counters **478**
Header Stripping **478**
About VXLAN Header Stripping **479**
VXLAN Tunnel Decapsulation Versus VXLAN Header Stripping 480
VXLAN Header Stripping – Rules, Notes, and Limitations 480
Configure VXLAN Header Stripping **482**
Configure Ageing Interval for VXLAN Header Stripping **483**
View Header Stripping Statistics **483**
About MPLS Header Stripping **484**
Advanced MPLS Header Stripping for GigaVUE-TA400 485
MPLS Header Stripping – Rules, Notes, and Limitations 486
Configure MPLS Header Stripping **488**
Create a New Map 488
Configure MPLS Labels on a Chassis **491**
Enable Header Stripping Protocol on Ports **491**
Tunnels **493**
About Circuit-ID Tunnels **494**
Circuit-ID Tunnels—Rules and Notes **495**
Circuit-ID Tunnel Encapsulation **495**
Circuit-ID Tunnel Decapsulation **496**
About Layer 2 Generic Routing Encapsulation (L2GRE) Tunnels **496**
L2GRE Tunnel Configuration—Rules and Notes 497
Limitation 499
Configure L2GRE Tunnel to Encapsulate Traffic **499**
Configure L2GRE Tunnel to Decapsulate Traffic **500**
Orchestrated Workflow Configuration of L2GRE Tunnels **501**
About Virtual Extensible LAN (VXLAN) Tunnels **503**
VXLAN Tunnel Configuration—Rules and Notes 503

Configure VXLAN Tunnel to Encapsulate Traffic **505**
Configure VXLAN Tunnel to Decapsulate Traffic **506**
Orchestrated Workflow Configuration of VXLAN Tunnels **506**
Create Tunnel **508**
Create VXLAN / L2GRE Group **509**
Configure L2GRE / VXLAN Identifier **510**
Configure L2GRE / VXLAN Identifier for a Chassis 510
Configure L2GRE / VXLAN Identifier for a Network Port 510
View VXLAN / L2GRE ID Statistics **511**
Tunnel Monitoring **513**
Rules, Notes, and Limitations 514
View Tunnel Monitoring 515
Packet Capture (PCAP) **517**
Rules, Notes, and Limitations 518
Configure PCAP Profile 519
View PCAP 522
Delete PCAP 522
View PCAP Files 522
Flow Mapping® **523**
About Flow Mapping® **523**
Flow Mapping® Overview **524**
Get Started with Flow Mapping® **525**
Check Status of Nodes and Ports 525
Designated Port Types **526**
About Shared Collectors 526
No Map Statistics for Shared-Collector Only 527
Shared Collector Configuration 527
About Map-passall Maps 528
Map-Passall and Regular Byrule Map 529
Map-Passall and Shared-Collector Only 529
Map-passalls Configuration 529
Map-passalls 529
Define Map Source Port Lists 529
Share Network Ports Between Maps 530
Share Tool Ports Between Maps 530
Map Priority 530
Adjust Map Priority in GigaVUE-FM 531

Flow Map Syntax and Construction **531**
Map Types 531
Define Map Source Port Lists for First Level Maps 532
Null Port in Maps 535
Map Subtypes 535
Map Type and Subtype Modification 536
Backwards Compatibility 537
Minimum Requirements for Map Creation 537
Map Rules 537
Other Types of Map Rules for GigaSMART Operations 537
IPv4/IPv6 and Map Rules 537
Set Map Rules for TCP Control Bits 539
How to Use Bit Count Netmasks 541
How to Combine Rules and Rule Logic 543
How to Mix Pass and Drop Rules 543
Work with User-Defined Pattern Match Rules 544
How to Handle Q-in-Q Packets in Maps **548**
Comparison of Q-in-Q Tagging 550
Priority Tagged Packets 550
Flow Mapping® on Inner VLAN Tags 551
Inner VLAN Limitation 552
Work with Map-Passalls and Port Mirroring in GigaVUE-FM **552**
Syntax for Maps-passalls and Port Mirroring 553
Rules for Map-Passalls and Port Mirroring 553
View and Delete Map-passalls 553
Port Access and Map Sharing **554**
Port-based Access Levels 554
How to share Maps 555
Map Examples **556**
Example: How to Create a Simple Map 556
Example: How to Handle Overlaps when Sending VLANs and Subnets to Different Tools 558
Manage Maps **560**
Map Views **560**
List View 560
Map Topology View 561
Manage Maps **561**
Map Templates **569**
Create Map Templates 569

Edit Map Templates ... 570
Map Template Quick View ... 570
Manage Map Rule Resources ... 570
Template Groups ... 570
Add Tags to Map Rules ... 572
Notes ... 574
Map Rule Statistics Dashboard ... 574
Filter Templates ... 575
Filter Template Qualifiers and Defaults ... 576
Filter Template Configuration ... 576
Filter Template Limits ... 577
How to Understand Map Filter Resources ... 577
Filter Template Rules and Recommendations ... 577
Filter Template Best Practices ... 578
Filter Templates in a Cluster ... 578
Filter Templates Formulas ... 579
Review Map Statistics with Map Rule Counters ... 580
Viewing Map Statistics with the Statistics Page ... 581
Viewing Map Statistics with Quick View ... 581
Flow Mapping® FAQ ... 581
How Many Map Rules are Supported? ... 581
How Many Rules Can Each Map Have? ... 586
How Many Maps Can Run at Once? ... 587
What Criteria can be Filtered in Q-in-Q Packets? ... 587
How Many Maps Can Share a Network Port? ... 587
How Many Network Ports and Tool Ports Can Be in a Map? ... 587
Are Port-Filters Supported? ... 587
Does Flow Mapping® Support Passalls? ... 588
Does Flow Mapping® Support port-pairs? ... 588
Does Flow Mapping® Support UDA Pattern Matches? ... 589
Can a GigaStream Act as a Shared Collector? ... 589
What Are the GigaStream Maximums? ... 589
What order are the map rules displayed in "show running config"? ... 590
Active Visibility ... 590
MAC Address Rewrite ... 593
Configuring MAC Address Re-write ... 594
License ... 595
Limitations ... 595
IP Address Rewrite ... 596
Configuring IP Address Re-write ... 597

License 598
Limitations 598
VLAN Manipulation 598
Configuring VLAN manipulation 599
License 600
Limitations 600
Monitor Port Utilization 601
Port Utilization Availability by Port Type 601
Set Port Utilization Thresholds 602
Utilization Alarm/SNMP Trap Generation 603
Configure Alarm Buffer Thresholds 604
Set Alarm Buffer Thresholds 605
Microburst 606
Best Practices to Improve Burst Tolerance 607
Logical Grouping of Ports in GigaVUE-HC3, GigaVUE-TA100, and GigaVUE-TA200 607
Flexible Inline Arrangements 608
Flexible Inline Arrangements 609
Flexible Inline Arrangement vs Inline Bypass Solution 610
Supported Platforms 611
Flexible Inline Arrangement License 612
Software Version 612
GRIP Supported by Flexible Inline Arrangements 612
Flexible Inline Solution Supported in Clustered Nodes 613
Limitations 614
Functionalities Not Supported by Flexible Inline Arrangement 614
Benefits of Flexible Inline Arrangements 614
About Flexible Inline Maps 615
Types of Flexible Inline Maps 616
Configure Flexible Inline Maps 617
Flexible Inline Arrangements—Rules and Notes 617
Visualize the Flexible Inline Arrangements Canvas 620
Configure Flexible Inline Flows 620
Configure Inline Network Ports and Inline Network 621
Network Port Link Status Propagation Parameter 622

Heartbeat Support Between GigaVUE Nodes 622
Configure IP Interface 624
Configure Inline Network Link Aggregation Group (LAG) 625
About Inline Network LAG 626
Inline Network LAG—Rules and Notes 626
Configure Inline Network LAG 629
Configure Inline Network Bundle 630
Configure Inline Tool Ports and Inline Tools 631
Configure Inline Tool Group 633
Configure Inline Single Tag 637
Configure Resilient Inline Arrangement 638
Resilient Inline Arrangement 639
Resilient Inline Arrangement—Classic 640
Resilient Inline Arrangement With Single VLAN Tag 640
Inter-broker Pathway (IB-P) 641
Resilient Inline Arrangement—Rules and Notes 642
Deploy Resilient Inline Arrangement 642
Configure Hardware Security Model (HSM) 647
About HSM 648
HSM Group - Limitations 649
Supported Platforms 649
Configure HSM Group 649
Configure Entrust nShield HSM: 650
Configure Thales-Luna HSM: 652
Modifying a HSM Decryption Deployment 653
Configure Flexible Inline TLS/SSL Decryption Solution 653
Flexible Inline TLS/SSL Decryption Solution 654
Benefits of Flexible Inline TLS/SSL Decryption Solution 655
Flexible Inline TLS/SSL Decryption Solution—Rules and Notes 656
Configure Flexible Inline TLS/SSL Decryption Solution 658
Inline SSL App—Field References 659
Support for unattended restart of TLS/SSL decryption in managed nodes 667
Single VLAN Tagging (SVT) in iSSL 667
Rules and Notes 668
Supported Platforms 668
Limitations 669
Configure Internet Content Adaptation Protocol (ICAP) 669
About ICAP 670

Supported Platforms 671
Configure ICAP Client 671
ICAP Client—Field References 673
ICAP - Rules, Notes, and Limitations 674
Configure Gigamon Resiliency for Inline Protection 676
How to Handle Recovery 679
Both Nodes Go Down and Only Secondary Comes Up 679
How to Cable GigaVUE Nodes 680
Configure GRIP Solution Software 680
Limitation for Suspended Role 681
Configure Synchronization 681
Display Redundancy Control State 682
How to Use Suspended Role for Maintenance 682
GRIP for Mixed Topologies 683
Upgrade Procedure 683
Pros and Cons 684
Upgrade Secondary Node followed by Primary Node 684
Upgrade both nodes together 684
Rules and Notes 684
Limitations 685
Troubleshoot 685
Signaling Ports Down 685
Traffic outage in Inline Tool 686
Network Traffic Outage 686
FAQs 686
Example: Gigamon Resiliency for Inline Protection 687
Configure Primary Role GigaVUE-HC1 687
Configure Secondary Role GigaVUE-HC1 689
Troubleshoot Flexible Inline Flows 690
Status 691
Rules and Notes 691
Statistics 692
HSM Statistics 693
ICAP Show Stats 694
Troubleshoot 696

Example: Troubleshoot Traffic Issues Between Side A and Side B **697**
View the Forwarding States of Inline Networks **698**
Import and Export Flexible Inline Solution **705**
Backup and Restore Flexible Inline Flows **706**
Timestamps **706**
About Timestamps **707**
Using Timestamps **708**
Why PTP? **708**
Synchronization of the PTP and Local System Clock 710
Using PTP to Timestamp Packets **711**
Configuring PTP Globally **711**
Enabling PTP on a Network Port **712**
Enabling Timestamp on a Port **713**
Viewing PTP Details **714**
Viewing the PTP Configuration on a Port **715**
Viewing the PTP Clock Details **717**
Viewing the PTP Foreign Source Details **721**
Viewing PTP Statistics **722**
PTP and Timestamp—Rules and Notes **723**
Fabric Maps **724**
Supported Topologies **724**
Multihop Topology 725
Supported Multihop Topology 726
Unsupported Multihop Topology 726
Leaf and Spine Topology 727
Dual Multipath Leaf and Spine Topology 728
Supported Failovers 728
Scenario 1 729
Scenario 2 730
UnSupported Failovers 730
Scenario 1 731
Scenario 2 732

Fabric Maps Prerequisites **732**
Notes on Circuit Ports and Circuit GigaStream in Fabric Maps **733**
Create Links between Clusters **734**
Create Manual Links 734
Create Gigamon Discovery Protocol (GDP) Based Links 735
Create Fabric Maps **736**
Edit and Delete Fabric Maps **740**
Edit Fabric Maps **740**
Edit Fabric Map Component 740
Delete Fabric Maps **741**
Prioritize Fabric Maps **742**
Share Circuit ID Resource **742**
Fabric Maps Statistics **743**
Display Fabric Map Statistics **743**
Display Fabric Map Details **743**
Filter Fabric Maps List View **744**
Fabric Port Group **745**
How Fabric Port Groups Work? 745
Rules and Notes **746**
Create Fabric Port Group **747**
Configure Fabric Map Using Fabric Port Group 748
Troubleshooting **749**
How to Troubleshoot Partial Success Errors **752**
Example 1: 752
Config Audit **754**
Limitations of Fabric Maps **755**
Backup and Restore Fabric Maps and Orchestrated Configurations **755**
Orchestrated Configurations **755**
About Intent Based Orchestrated Configurations 756
Benefits of Orchestrated Configurations 758
Orchestrated Configuration: Examples 759
Priority-Free Policies 759
Overlapping Sources in Policies 761
Supported Topologies 762

Rules and Notes for Orchestrated Configurations 762
Create Orchestrated Policies 763
Prerequisites 763
How to Create a Policy 764
Tagging and RBAC Support 766
Drop Rules 767
Import and Export Orchestrated Policies 769
Rules and Notes 771
Import and Export Orchestrated Configuration 771
Egress Filters for Additional Filtering Capabilities 772
GigaSMART® 774
GigaSMART Operations 774
About GigaSMART® Applications 775
Quick Glance- How to Configure a GigaSMART Application 775
Application Intelligence 777
Application Intelligence 777
Licensing 778
About Application Intelligence 779
Application Intelligence Capabilities 780
Supported Platforms 780
Application Intelligence Feature Compatibility 781
User Defined Application 781
Create Rules under User Defined Application 781
Supported Protocols and Attributes 782
Mindata 787
Supported RegExp Syntax 787
Limitations 788
How Application Intelligence Works 788
GigaVUE-FM Application Protobook 789
Application Intelligence—Rules and Notes 790
Application Intelligence Session 791
Create an Application Intelligence Session in Physical Environment 792
Upgrading the Protocol Signature 794
Fast Mode 795
View the Details of an Application Intelligence Session 796
Application Filtering Intelligence 798
Elephant Flows in Application Filtering Intelligence 799

Create Application Filtering Intelligence for Physical Environment **801**
Adding Rules 804
AFI with Slicing and Masking **805**
View the Application Intelligence Dashboard **807**
About De-duplication **810**
Application Metadata Intelligence **811**
Create Application Metadata Intelligence for Physical Environment **812**
Create NetFlow Session for Physical Environment **816**
NetFlow Dashboard 820
Tool Templates **821**
Import and Export Tool Templates 825
RTT Attribute in AMI **826**
View the Health Status of a Solution **826**
GigaSMART Application Session Filtering (ASF) and Buffer ASF **829**
Session-Aware APF (SAPF) Renaming and Licensing Change to ASF 831
Application Session Filtering with Buffering 831
ASF and Buffer ASF Session Definitions 832
Define ASF Session 833
Quick Session Delete for Buffer ASF 834
Specify Resources for Buffer ASF 834
ASF and Buffer ASF Session Notes 835
Buffer ASF Packet Processing Special Cases 836
Map Statistics 836
ASF and Buffer ASF Examples 836
Example 1: ASF, Forward TCP Traffic 837
Example 2: ASF, Forward VNC Traffic 839
Example 3: ASF, Forward Traffic Matching a Pattern 841
Example 4: ASF, Forward GTP Traffic 843
Example 1: Buffer ASF, Drop Netflix Traffic 845
Example 2: Buffer ASF, Drop YouTube Traffic 847
Example 3: Buffer ASF, Drop Windows Update Traffic 849
Example 4: Buffer ASF, Forward VNC Traffic 851
Example 5: Buffer ASF, Forward HTTPS Traffic on Non-Standard Port 853
Display ASF Statistics 855
Enhanced Application Session Filtering 856

GigaSMART Enhanced Application Session Filtering 856
GigaVUE Enriched Metadata for Mobile Networks 857
Subscriber Intelligence 860
Installation and Configuration of Subscriber Intelligence Solution using Ansible 860
System Requirements 861
Installation and Configuration of Gigamon Ansible Module 861
Deployment Report 862
Check Mode 862
Reapplying Golden Payload 863
Rules and Notes 863
Configuration of Non-CUPS using Ansible 864
Prerequisite 868
Site 869
cpNode 869
upNode 869
5GPolicy 869
LTEPolicy 869
Deployment of Non-CUPS Solution 869
GigaSMART GTP and CUPS Correlation 870
Filtering on Subscriber IDs and Version 871
Session Correlation 874
Supported Interfaces 875
Conditional S10 Support 877
GTP Session Timeout 877
Priorities for Flow Rules and Maps 878
GTP Correlation Configuration Examples 881
GigaSMART GTP Rotational Sampling 900
Supported Platforms 902
Configure Rotational Sampling 902
Viewing the Configuration of Rotational Sampling for non-CUPS LTE 903
Rotational Sampling support in Mobility Solution 903
GigaSMART 4G RAN Correlation 906
Rules and Notes 906
GigaSMART SIP/RTP Correlation 907
CallerID Tracking 908
Support for SIP, RTP, and RTCP 909
SIP/RTP Correlation Engine 909
Configure SIP/RTP Correlation Engine 909

SIP Whitelist ... 910
RTP Flow Sampling ... 911
Support for Sessions ... 911
Support for IPv4 and IPv6 ... 911
Support for Content Masking ... 912
Behaviors of Some SIP Methods ... 912
SIP Whitelisting in a Cluster ... 913
Support for NAT ... 913
Not Supported by SIP/RTP Correlation ... 913
Display SIP/RTP Reports ... 914
Display SIP Map Statistics ... 915
SIP/RTP Support for Tool Throttling ... 915
Admission Control ... 916
SIP/RTP Examples ... 916
SIP/RTP Load Balancing Example ... 916
SIP/RTP Minimum Configuration Example ... 917
GigaSMART IP FlowVUE ... 925
Configure FlowVUE ... 926
FlowVUE Configuration Examples ... 927
Display FlowVUE Statistics ... 927
FlowVUE_Examples ... 927
GigaSMART GTP Whitelisting and GTP Flow Sampling ... 934
GTP Whitelisting ... 934
GTP Flow Sampling ... 941
GTP Subscriber Aware Random Sampling ... 946
GigaSMART GTP Whitelisting and GTP Flow Sampling Examples ... 947
4G/5G Traffic Monitoring using UPN ... 973
Configure 4G/5G Traffic Monitoring using UPN ... 975
PFCP Messages ... 975
PFCP Load Balancing ... 976
Map Configuration ... 976
Stateful Session Recovery ... 977
5G Stateful Session Recovery ... 978
Configure 5G Stateful Session Recovery ... 979
View Backup and Restore Information ... 979
Configure GTP Overlap Mapping ... 980
Configuration Considerations ... 981
About Standalone UPN Flow Sampling Map Mode and Port Groups ... 981
Standalone UPN Flow Sampling Map Priority ... 982
Virtual Port Configuration in Standalone UPN Mode ... 982

Overlap Support ... 982
Standalone UPN ... 982
PFCP (Packet Forwarding Control Protocol) ... 983
PFCP Messages ... 983
PFCP Load Balancing ... 984
Map Configuration ... 984
Stateful Session Recovery ... 985
5G Stateful Session Recovery ... 986
Configure 5G Stateful Session Recovery ... 986
View Backup and Restore Information ... 987
Configure GTP Overlap Mapping ... 988
Configuration Considerations ... 989
About Standalone UPN Flow Sampling Map Mode and Port Groups ... 989
Standalone UPN Flow Sampling Map Priority ... 990
Virtual Port Configuration in Standalone UPN Mode ... 990
Overlap Support ... 990
Interface Filtering and APN/DNN Filtering ... 991
Interface Filtering ... 991
APN (Access Point Name)/DNN (Data Network Name) Filtering ... 992
Custom Interface Selection ... 993
GigaSMART 5G CUPS ... 995
Overview ... 996
Supported Platforms ... 998
5G Correlation ... 999
5G Load Balancing ... 999
Configure 5G Load Balancing ... 1000
5G RAN Correlation ... 1000
Configure 5G RAN Correlation ... 1001
5G Network Slice Correlation ... 1001
Configure 5G Network Slice Correlation ... 1003
5QI Correlation ... 1003
Rules and Notes ... 1003
5G Flow Sampling and Filtering ... 1003
About Flow Sampling Rules and Maps ... 1004
GTP Overlap Flow Sampling for 4G and 5G ... 1007
Configure GTP Overlap Mapping ... 1007
Configuration Considerations ... 1007
Overlap Map Statistics ... 1015

5G Stateful Session Recovery **1015**
Configure 5G Stateful Session Recovery 1016
View Backup and Restore information 1016
Remove Backup files 1018
5G Whitelisting **1019**
Configure Whitelist Maps 1019
Change Priority of Whitelist Maps 1020
Delete Whitelist Maps 1020
User Plane Node Traffic Monitoring **1022**
Rules and Notes 1023
Configure Stand-Alone User Plane Node Traffic Monitoring 1023
Control Plane Metadata **1024**
Configuration of 3G/4G Control Plane Metadata using Ansible **1026**
Creating Inventory Directory 1026
Creating fmInfo.yml 1026
Creating ansible_inputs.json 1027
Creating 3G/4G control plane metadata inventory file 1027
Creating host_vars directory 1029
Creating host_vars files 1030
Prerequisite 1031
Site 1031
cpNode 1031
upNode 1032
5GPolicy 1032
LTEPolicy 1032
Deployment of 3G/4G Control Plane Metadata Solution 1032
Set up of additional variable for Single GigaVUE-FM instance 1032
Execute the Playbook 1032
Configuration of 5G Control Plane Metadata using Ansible **1033**
Creating Inventory Directory 1033
Creating fmInfo.yml 1034
Creating ansible_inputs.json 1034
Creating 5G control plane metadata inventory file 1034
Creating host_vars directory 1037
Creating host_vars files 1038
Prerequisite 1038
Site 1039
cpNode 1039

upNode 1039
5GPolicy 1039
LTEPolicy 1039
Deployment of 3G/4G Control Plane Metadata Solution 1039
Set up of additional variable for Single GigaVUE-FM instance 1039
Execute the Playbook 1039
CPN-UPN Communication for Support of RAN and Network Slice Attributes 1040
CPN-UPN Communication Configuration 1042
Configure CPN-UPN Communication using Ansible 1042
Configure CPN-UPN Communication using CLI 1043
Configure CPN-UPN Communication Solution using Ansible 1044
Create Inventory Directory 1044
Create fmInfo.yml 1044
Create ansible_inputs.json File 1044
Create an Inventory File for CPN-UPN Communication 1045
Create host_vars directory 1046
Create host_vars files 1046
Deploy CPN-UPN Communication Solution in Ansible 1052
View CPN-UPN Communication Solution in GigaVUE-FM 1053
Remove CPN-UPN Communication 1054
Map Rules for CPN-UPN Communication 1054
Flow Sample Map Rules for CPN-UPN Communication 1054
Rules and Notes: 1054
Forward List Map Rules for CPN-UPN Communication 1056
Rules and Notes: 1056
View CUPS Communication Dashboard 1059
Upgrade Standalone UPN to CPN-UPN Communication Solution 1064
Rollback from CPN-UPN Communication Solution to Standalone UPN 1064
Quick Rollback from CPN-UPN Communication Solution to Standalone UPN 1065
Monitoring of Subscriber Intelligence Solutions 1066
Monitoring CUPS Solution 1066
Health Status of Solution 1067
Display Flow Ops Reports 1068
Flow Ops Report - Field Reference 1071

Export Flow SIP Session Reports 1074
Export Flow Filtering Reports 1075
Generate Delta Reports 1076
GTP Overlap Flow Sampling Maps 1078
Configure GTP Overlap Mapping 1080
Configuration Considerations 1080
GTP Overlap Flow Sampling Maps Example 1084
This section contains: 1084
Example 1: GTP Overlap Mode 1084
GTP Stateful Session Recovery 1089
GTP Scaling 1092
GigaSMART Cards in GigaVUE-OS Devices 1093
GTP Engine Grouping 1093
Passing GTP Control Traffic 1094
Configure GTP Engine Grouping 1095
Display Statistics 1095
GTP Engine Grouping Configuration Examples 1096
GigaSMART TLS/SSL Decryption for Inline and Out-of-Band Tools 1107
Limitations 1107
Inline TLS/SSL Decryption 1107
About Inline TLS/SSL Decryption 1108
What Inline TLS/SSL Decryption Provides 1108
Important Inline TLS/SSL Rules and Notes 1109
Supported Ciphers 1110
More About Inline TLS/SSL Decryption 1113
TLS/SSL Decryption for Inline Tools 1113
Example Inline TLS/SSL Decryption 1114
TLS/SSL Terminology and Acronyms 1116
InlineTLS/SSL show command Field Descriptions 1118
TLS/SSL Sessions 1123
TLS/SSL Handshake 1124
TLS/SSL Handshake Steps 1125
TLS/SSL Session, Inbound Deployment 1127
TLS/SSL Session, Outbound Deployment 1128
TCP Transition States during TLS/SSL Session 1128
TLS/SSL Session Resumption 1129
TLS/SSL Session Search 1130
StartTLS and HTTP CONNECT 1130

Inline TLS/SSL Decryption Behavior with StartTLS 1130
TLS/SSL Keys and Certificates 1131
Key Store 1132
Set Up Key Store Certificate Management 1132
Generate and Add a Certificate to Key Store 1133
Display Key Store Certificates 1133
Trust Store 1134
Certificate Validation 1134
Client Authentication 1136
Re-Signed Certificates 1136
Checking Certificate Revocation Status 1137
Certificate Revocation List (CRL) 1137
Online Certificate Status Protocol (OCSP) 1138
CRL and OCSP 1138
Policy Profile 1138
Policy Evaluation 1139
Network Phase 1139
SNI Phase 1140
Certificate Validation 1140
Cert Phase 1141
Policy Profile Options 1141
Inline TLS/SSL Decryption Port Map 1142
Enable or Disable Tool Bypass 1142
High Availability Active Standby 1142
Inline Network Group Multiple Entry 1143
Tool Early Engage 1144
One-Arm Mode 1144
Configuring One-Arm Mode 1144
Failover Support 1146
Important Rules and Notes 1146
Tool Early Inspect 1147
Caches 1148
Cache Persistence 1149
GigaSMART Overload Bypass 1149
CPU Overload Threshold 1150
Inline TLS/SSL Monitor Mode 1150
Configure the Inline TLS/SSL Monitor Mode 1151
Inline Tool Configurations 1152
Inline Bypass Restriction 1154
Forwarding 1154

Failover 1154
Out-of-Band and Inline Tools 1155
Service Chaining of Decrypted Traffic 1156
Inline TLS/SSL Decryption Deployments 1156
GigaVUE Modules for Inline SSL Decryption 1158
Packet Flows 1159
Modules Matrix 1161
Inline SSL Traffic Filtering 1161
No-decrypt Listing Policy 1162
Decrypt Listing Policy 1162
Rules and Notes while configuring a No-Decrypt/Decrypt List Policy 1163
IP Address Subnet with Longest Prefix Match(LPM) 1163
URL Categorization 1163
URL Look-ups and Caching 1164
Inline SSL URL categories 1165
Proxy Server Profile for URL Categorization and Certificate Revocation status 1173
Get Started with Inline TLS/SSL Decryption 1173
Before You Begin 1174
Supported Platforms 1174
GigaSMART Licensing 1174
GigaSMART Compatibility 1174
Installation 1175
Install GigaVUE Modules 1175
Install Software Version 1175
Install U-Boot Version on GigaVUE-HC3 1175
Install MitM Certificates in Client Trust Store 1176
Initial Configuration 1176
Set up the Stack Port Interface 1176
About the Stack Port Interface 1176
Configure Stack Port Interface 1177
Set up GigaSMART for Inline TLS/SSL Decryption on GigaVUE-HC1 1177
Configure Keychain Password 1178
Configure Keychain Password 1178
Configure Primary Certificate and Key 1178
Configure an Inline TLS/SSL Session Logging Server 1179
Configure Inline TLS/SSL Decryption 1181
Introduction to Inline TLS/SSL Map Workflows 1181
Flow A 1183
Flow B 1183

Flow C 1184
Flow D 1185
Flow E 1186
Flow F 1187
Flow G 1188
Configure Inline TLS/SSL Decryption Using GigaVUE-FM 1189
Workflow Overview 1190
Inline TLS/SSL Configuration Workflow Steps 1190
Inline TLS/SSL Policy Profile—Field References 1193
Inline TLS/SSL Map Workflow Steps (for Flow B) 1195
View Statistics 1201
Inline TLS/SSL Session Statistics 1202
Monitor Statistics 1202
Certificate Statistics 1204
GigaSMART Inline TLS/SSL Dashboards 1205
Basic Dashboards 1206
Advanced Dashboards 1213
System Requirements 1213
Configure Advanced dashboard 1214
Rules and Notes 1214
Resilient Inline Arrangement with GigaSMART Flex Inline Solution 1218
Symmetric Traffic in RIA 1219
Asymmetric Traffic in RIA 1220
VLAN Tagging Behavior for Decrypted Traffic. 1222
VLAN Tagging Behavior for Non-Decrypted Traffic. 1222
Setup Resilient Inline TLS/SSL 1222
Limitations 1223
GigaSMART Passive TLS/SSL Decryption 1224
About Passive TLS/SSL Decryption 1225
Supported Protocols, Algorithms, and Ciphers 1228
Limitations 1230
Licensing 1230
Create and Reset TLS/SSL Keychain Passwords 1230
Work with Keys and Services 1232
Set Up Key Store Certificate Management for Passive TLS/SSL 1232
Upload TLS/SSL Private Keys 1233
Delete TLS/SSL Key 1234
Display Key Store Details for Passive TLS/SSL 1234
Create TLS/SSL Service 1235
ECODES for Troubleshooting Passive TLS/SSL Decryption 1236

View Passive TLS/SSL Decryption Flow Ops Report 1237
Configuring Passive TLS/SSL Decryption Examples - Command Line Reference 1238
Example 1: TLS/SSL Decryption with a Regular Map 1238
Example 2: TLS/SSL Decryption with De-duplication 1240
Other Usage Examples 1240
Entrust nShield and Thales-Luna Network HSM for TLS/SSL Decryption for Out-of-Band Tools(Passive) 1241
Entrust nShield and Thales-Luna HSM for TLS/SSL Decryption for Out-of-Band Tools—Rules and Notes 1242
Configure HSM for TLS/SSL Decryption for Out-of-Band Tools 1243
Configure HSM Group 1243
Configure Set Key Handler 1245
Configure Passive TLS/SSL Network Access 1246
Use RFS to Manage Encrypted Keys 1247
Configure a GigaSMART Group 1249
Create a GigaSMART Operation (GSOP) 1249
Create TLS/SSL Keychain Password 1250
Upload TLS/SSL Private Keys 1251
Configure TLS/SSL Service 1252
Configure Maps 1253
Entrust nShield and Thales- Luna HSM for TLS/SSL Decryption for iSSL 1256
Configure HSM for TLS/SSL Decryption for iSSL 1257
Traffic Intelligence 1258
GigaSMART Adaptive Packet Filtering (APF) 1259
Feature Overview 1259
Implement APF Through the UI 1260
Content-based Filtering 1262
Encapsulation Awareness 1265
Pattern Matching on Gen 2 GigaSMART modules 1266
Cross-Packet Pattern Matching 1272
Pattern Matching on Gen 3 GigaSMART modules 1274
Feature Parity with Gen 2 GigaSMART module and Gen 3 GigaSMART module 1277
Map Statistics 1277
Display APF Statistics 1278
Adaptive Packet Filtering Examples 1278
GigaSMART Load Balancing 1319
Stateful Load Balancing 1320
Examples 1324

Stateless Load Balancing 1329
Examples 1332
Enhanced Load Balancing 1338
GigaSMART De-Duplication 1340
Feature Overview 1341
De-duplication Configuration Steps 1342
Configure GigaSMART Parameters for Packet De-duplication 1342
Example – GigaSMART De-duplication 1343
Display De-duplication Statistics 1345
GigaSMART Traffic Performance Enhancement 1345
Flow Masking 1345
GigaSMART Encapsulated Traffic Performance Enhancement 1346
GigaSMART Header Stripping 1349
Header Stripping Illustrated 1353
Example – Stripping MPLS Headers and Adding a VLAN ID 1353
Example – Stripping Layer 3 GRE IP Encapsulated Packets 1354
Example – Stripping Layer 2 GRE Ethernet Encapsulated Packets 1355
Display Header Stripping Statistics 1356
Example – FM6000 Timestamping 1356
Generic Header Stripping 1358
Custom 1360
PPPoE 1362
Cisco LISP 1363
L2 MPLS 1365
VXLAN 1366
TRILL 1368
Avaya SPB 1369
GigaSMART Header Addition 1371
GigaSMART Masking 1372
Examples – GigaSMART Masking 1374
GigaSMART NetFlow Generation 1376
NetFlow Generation Components 1379
Network Ports 1380
Map(s) 1380
GigaSMART Group 1380
GSOP 1381
NetFlow Records 1381
NetFlow Monitors 1381
Sampled NetFlow Data 1382
NetFlow Exporters 1382
IP Interface with Tool Ports 1383

Enhancements to NetFlow 1383
IPFIX Extension: HTTP Response Code 1387
Display Exporter Statistics 1406
Display Monitor Statistics 1407
Display IP Interfaces Statistics 1407
NetFlow Generation Configuration Modification and Removal 1407
Modify NetFlow Generation Monitor Configuration 1407
Modify NetFlow Generation Record Configuration 1408
Remove NetFlow Generation Configuration 1409
V5 Fixed Record Template 1410
NetFlow Generation Match/Key and Collect/Non-Key Elements 1411
Match/Key Syntax 1412
Collect/Non-Key Syntax 1415
NetFlow Cacheless export using TCP protocol 1423
Rules and Notes 1423
Configure Netflow Generation 1424
GigaSMART Packet Slicing 1432
Example – GigaSMART Packet Slicing 1433
Display Slicing Statistics 1434
GigaSMART Advanced Flow Slicing 1434
Limitations 1434
Create Advanced Flow Slicing Profile 1435
Example-GigaSMART Advanced Slicing 1436
Display Slicing Statistics 1436
GigaSMART SSL Decryption 1437
PCAPng Application 1437
Tunneling Operations 1438
GigaSMART IP Encapsulation/Decapsulation (GigaSMART Tunnel) 1439
Configure Sending End of Tunnel: GigaVUE-VM vMap in Sydney 1441
Configure Receiving End of Tunnel: GigaVUE H Series with GigaSMART in Melbourne 1441
Configure IP Interfaces for GigaSMART Tunnels 1443
About IP Interface Centralization 1443
Configure GigaSMART IP Encapsulation/Decapsulation 1444
Configure Sending End of Tunnel: GigaVUE-HC1 in Reno 1447
Configure Receiving End of Tunnel: GigaVUE-HC3 with GigaSMART in San Francisco 1450

GigaSMART IP Encapsulation (GigaSMART Tunnel) 1451
GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation 1452
Layer 2 GRE Header Length 1453
Load Balancing to Multiple Destinations 1454
L2GRE IPv4 Encapsulation/Decapsulation 1454
L2GRE IPv6 Encapsulation/Decapsulation 1455
L2GRE Tunnel Termination 1456
Configure L2GRE Tunnel Encapsulation and Decapsulation 1457
Display L2GRE IPV4/IPv6 Tunnel Statistics 1457
Display L2GRE Tunnel Encapsulation Statistics 1457
Display L2GRE Tunnel Decapsulation Statistics 1458
Configure GigaSMART Operation for Layer 2 GRE 1458
Example 1 – GigaSMART L2GRE Tunnel Encapsulation 1459
Example 2 – GigaSMART L2GRE Tunnel Encap Stateful LB 1460
Example 3 – GigaSMART L2GRE Tunnel Encap Stateless LB 1463
Example 4 – GigaSMART L2GRE Tunnel Decapsulation 1464
Example 5 – GigaSMART L2GRE IPv6 Tunnel Encamp/Decap with Load-Balancing 1466
Orchestrated Workflow Configuration of GigaSMART L2GRE Tunnels 1472
GigaSMART VXLAN Tunnel Encapsulation 1473
VXLAN Header Length 1474
VXLAN IPv4 Encapsulation/Decapsulation 1474
VXLAN IPv6 Encapsulation/Decapsulation 1476
Display VXLAN IPV4/IPv6 Tunnel Statistics 1476
Display VXLAN Tunnel Encapsulation Statistics 1477
Orchestrated Workflow Configuration of GigaSMART VXLAN Tunnels 1477
IP Fragmentation and Reassembly on L2GRE and GMIP Tunnels 1479
IP Fragmentation on Encapsulation 1479
IP Reassembly on Decapsulation 1480
Notes and Considerations 1480
GigaSMART ERSPAN Tunnel Decapsulation 1481
ERSPAN Type III 1482
Configure GigaSMART Operations for ERSPAN 1484
ERSPAN Tunnel Header Removal 1484
ERSPAN Type III Tunnel Header Removal 1486
Display ERSPAN Statistics 1488

Orchestrated Workflow Configuration of GigaSMART ERSPAN Tunnels **1489**
GigaSMART VXLAN Tunnel Decapsulation **1490**
GigaSMART Custom Tunnel Decapsulation **1495**
Custom Tunnel Decapsulation Configuration Example 1496
GigaSMART TCP tunnel **1499**
Configuration 1499
Supported Devices 1499
GRE-In-UDP Tunnel Decapsulation 1499
Secure Tunnels **1500**
Supported Platforms 1502
Configure Secure Tunnels 1502
Edit Secure Tunnel 1505
Delete Secure Tunnel 1505
Export Secure Tunnels 1505
View Tunnel Statistics 1505
Limitations 1505
Work with GigaSMART Operations **1506**
GigaSMART Licensing **1507**
Access GigaSMART from GigaVUE-FM **1507**
Quick Glance- How to Configure a GigaSMART Application **1508**
Create GigaSMART Operations – A Summary **1509**
Groups of GigaSMART Engine Ports **1511**
How to Use GigaSMART Operations – Example 1512

Engine Watchdog Timer in GigaSMART **1514**
Tunnel Health Checks **1515**
Configure Hashing **1517**
GigaSMART Rules and Tips **1518**
Virtual Ports **1519**
Virtual Port Rules **1519**
Create Virtual Port **1520**
Multiple Virtual Ports for First Level Map **1523**
Multiple Virtual Port Rules **1525**
Multiple Virtual Port with Other GigaSMART Applications **1526**
Virtual Port Statistics **1529**
Differences in GigaSMART Nomenclature Between the CLI and GigaVUE-FM **1530**
GigaSMART Operations in Clusters **1531**
How to Combine GigaSMART Operations **1533**
How to Read GigaSMART Operations Table **1534**
Work with GigaSMART Operation Combinations in GigaVUE-FM **1534**
Supported GigaSMART Operations **1535**
Supported GigaSMART Operations on GigaVUE HC Series 1535
Supported GigaSMART operations on GigaVUE V Series Node 1539
Order of GigaSMART Operations **1543**
View GigaSMART Statistics **1544**
Definitions of GigaSMART Statistics **1545**
NetFlow Monitor Statistics Definitions 1545
NetFlow Exporter Statistics Definitions 1546
IP Interfaces Statistics Definitions 1547
TLS/SSL Application Statistics Definitions 1548
ASF Statistics Definitions 1549
GigaSMART Group Statistics Definitions 1550
GigaSMART Group Flow Ops Report Statistics Definitions 1551
Flow Ops Report Statistics Definitions for FlowVUE 1552
Flow Ops Report Statistics Definitions for GTP 1552
Flow Ops Report Statistics Definitions for GTP Overlap 1556
Flow Ops Report Definitions for SIP/RTP Correlation 1557

Flow Ops Report Statistics for Passive TLS/SSL Decryption 1559
GigaSMART Operations Statistics Definitions 1560
Passive TLS/SSL Decryption Statistics Definitions 1561
Inline TLS/SSL Decryption Statistics Definitions 1562
De-duplication Statistics Definitions 1562
ERSPAN Statistics Definitions 1563
Tunnel Decapsulation Statistics Definitions 1564
Tunnel Encapsulation Statistics Definitions 1565
APF Statistics Definitions 1568
ASF Statistics Definitions 1568
Masking Statistics Definitions 1568
Slicing Statistics Definitions 1569
Header Stripping Statistics Definitions 1569
Generic Header Stripping Statistics Definitions 1570
Trailer Statistics Definitions 1570
FlowVUE Statistics Definitions 1571
NetFlow Statistics Definitions 1571
Display GigaSMART Application Resource Usage 1572
Overview of GigaSMART Application Resources 1573
Resources for Buffer ASF 1574
Reload GigaSMART Line Card or Module 1574
GigaSMART CPU Utilization Statistics 1574
Display GigaSMART CPU Utilization 1575
Configure Threshold 1576
Configure Threshold Crossing Notification 1576
How to Use GigaSMART Trailers 1577
About the Source ID Field 1577
Remove GigaSMART Trailers 1579
Multiple GigaSMART Trailers and Cascade Connections 1581
Interpret GigaSMART Trailer 1582
GigaSMART Trailer Reference 1582
GigaSMART Trailer Format 1582
GigaSMART Trailer Format 1582
GigaSMART Trailer TLVs 1583
Format of the GIGAMON_SRCID TLV 1584
GigaSMART Logs 1589
Create GS Log file 1590
Delete Log File 1590
Upgrade GigaSMART Cards 1590
Upgrade GigaSMART card using External Image Server 1591
Upload image to external image server 1591

Add the external server to GigaVUE-FM 1591
Upgrade GigaSMART card using External Image 1592
Upgrade GigaSMART card using Internal Image Server 1593
Download Image 1593
Upload the image file to GigaVUE-FM 1593
Upgrade GigaSMART card using Internal image 1593
Additional Sources of Information 1595
Documentation 1595
How to Download Software and Release Notes from My Gigamon 1598
Documentation Feedback 1598
Contact Technical Support 1599
Contact Sales 1600
Premium Support 1600
The VÜE Community 1600
Glossary 1601

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.

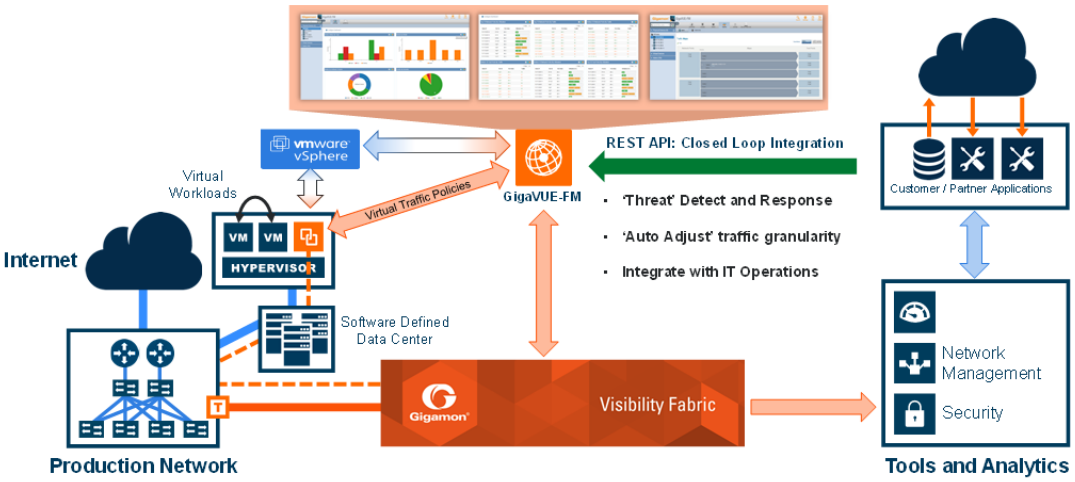

Figure 1 *GigaVUE-FM Overview*

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 Versions

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
	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
	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
	GigaVUE-HCT	v6.5.xx, v6.6.xx, v6.7xx, v6.8xx, v6.9xx
	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
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
	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
	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
	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

GigaVUE Series	Node	Compatible GigaVUE-OS versions
	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
	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™	132.0.1
	Microsoft Edge	130.0.2849.68
	Google® Chrome®	130.0.6723.117
Mac OSX	Google® Chrome®	127.0.6533.100
	Mozilla Firefox™	129.0
	Apple® Safari®	17.5
	Microsoft Edge	127.0.2651.98
Linux	Google® Chrome®	130.0.6723.116
	Mozilla Firefox™	128.0

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 2: 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.

Benefit	Descriptions
	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 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.

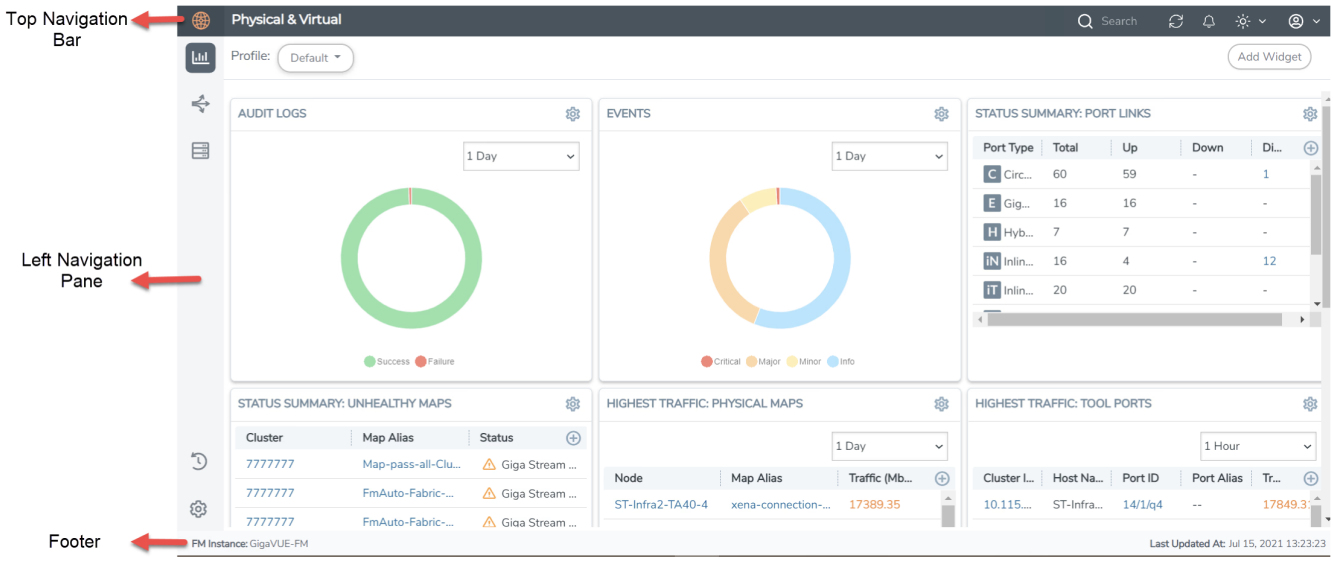

Left Navigation Pane

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.

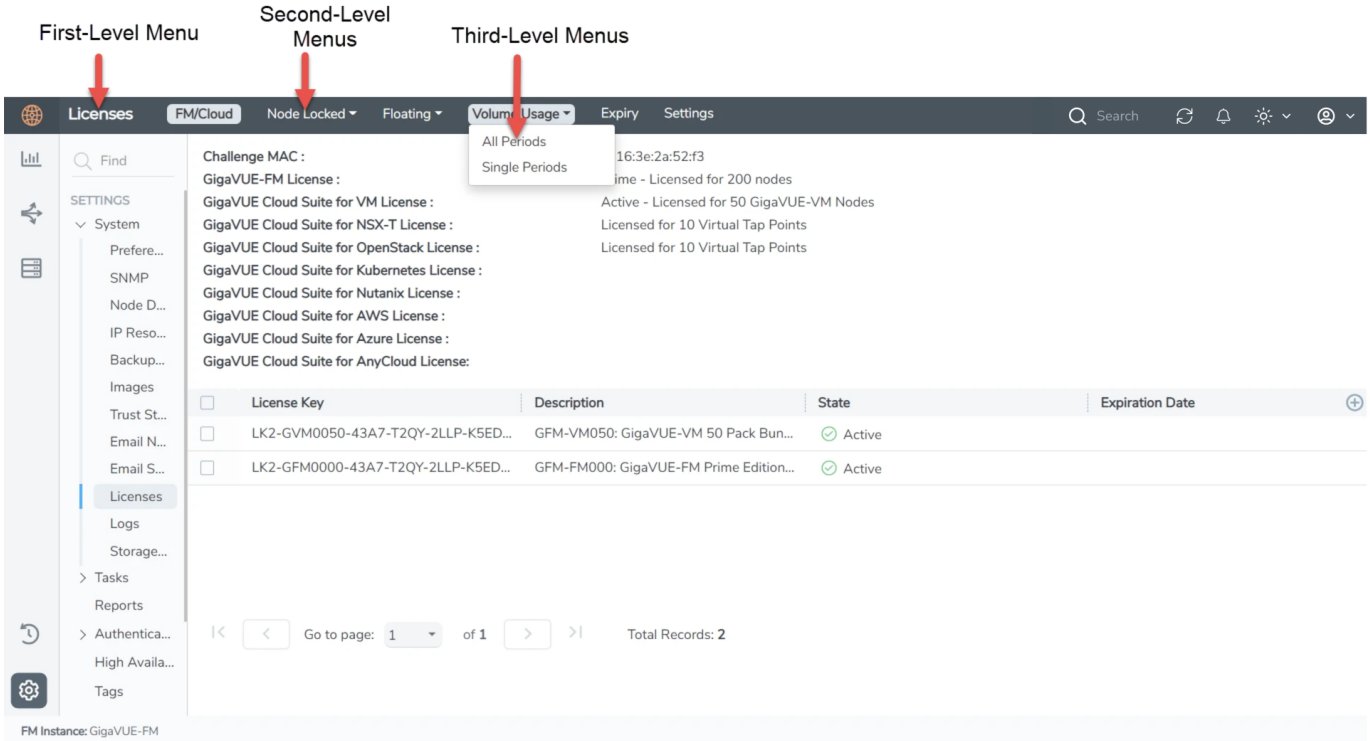

- The first-level menu is displayed as the Main Header in the top navigation bar.
- The second-level menus are displayed next to the first-level menu.

  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.

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. On the left navigation pane, click  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.

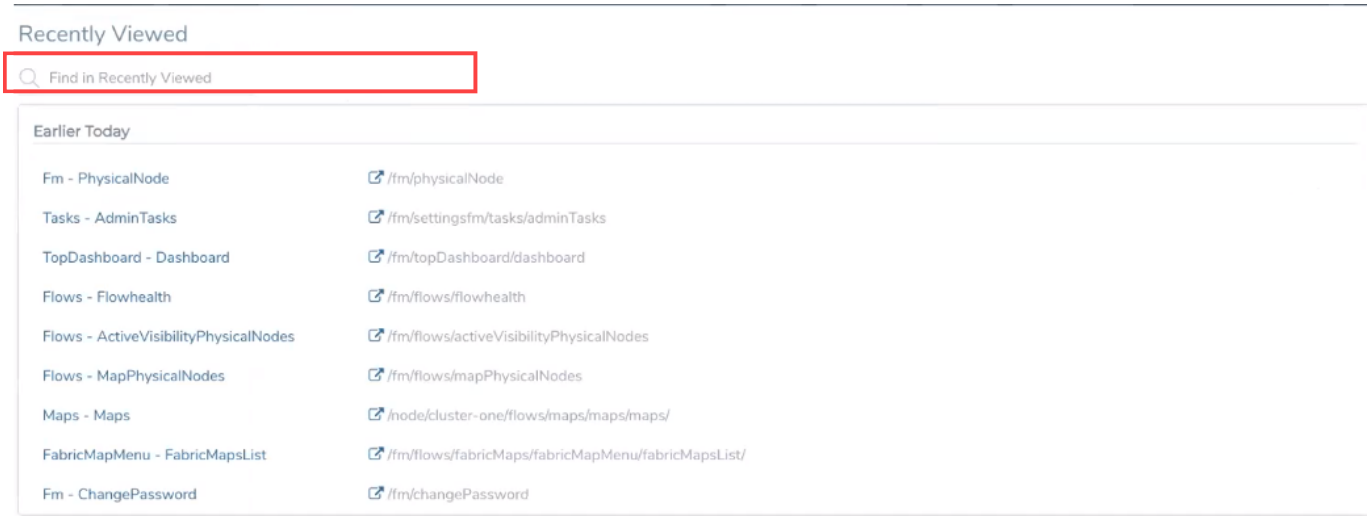

Find Pages

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.

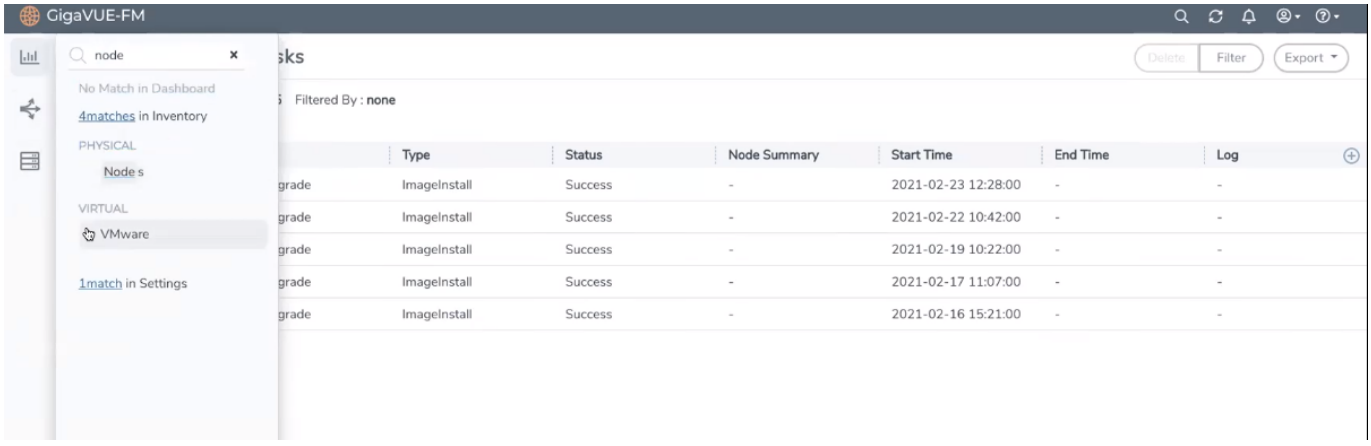

.

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.

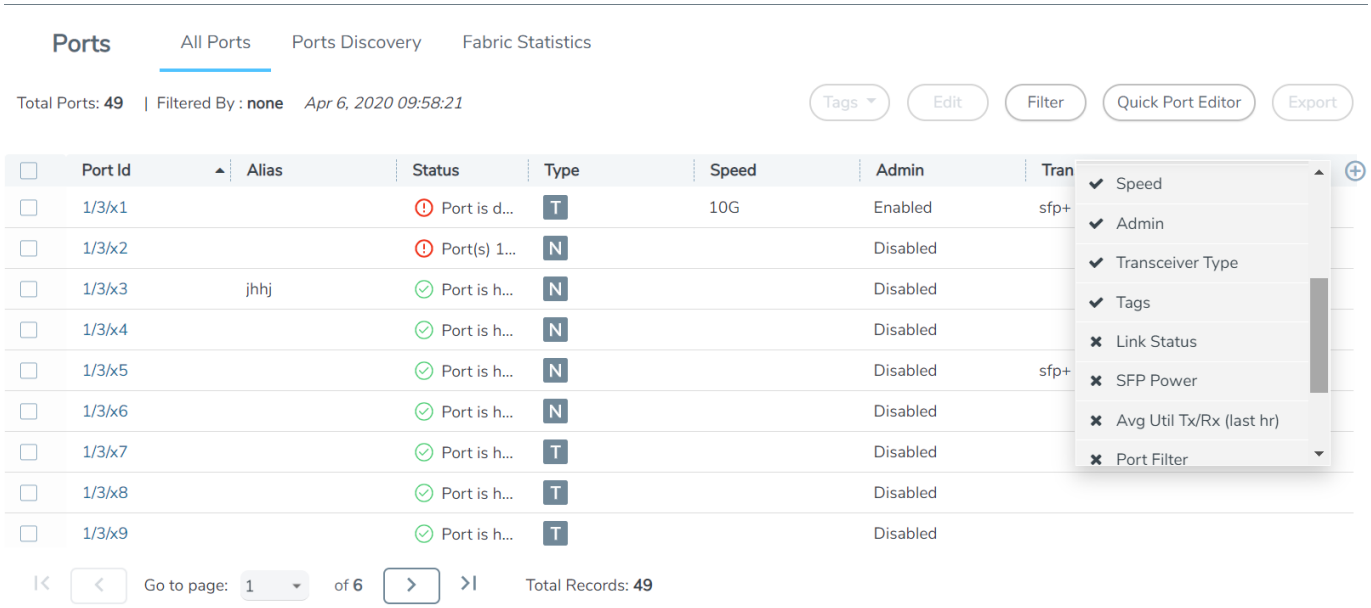

Figure 2 *Table menu to configure columns*

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.

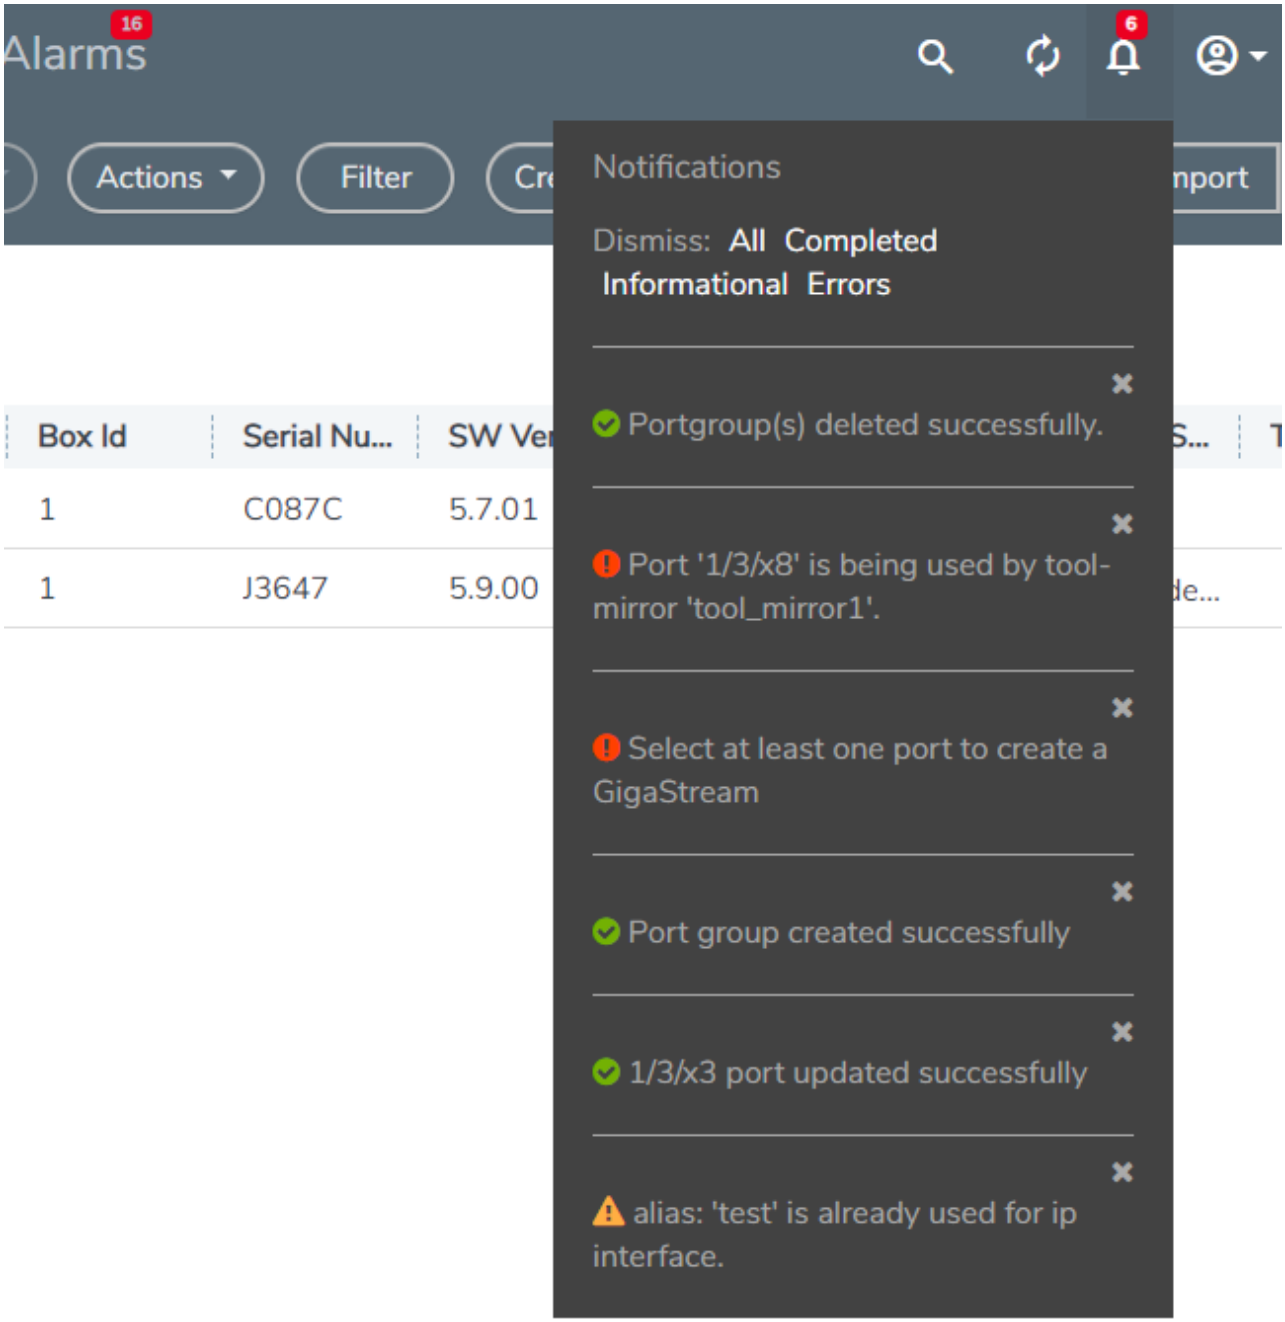

Figure 3 *Pop-up messages*

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.

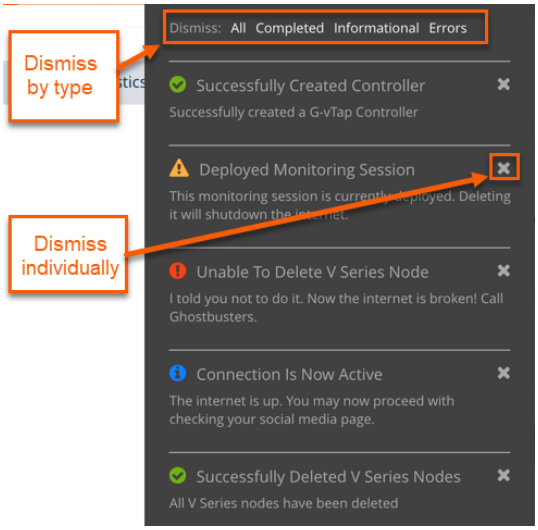

Figure 4 *Long-term messages*

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.

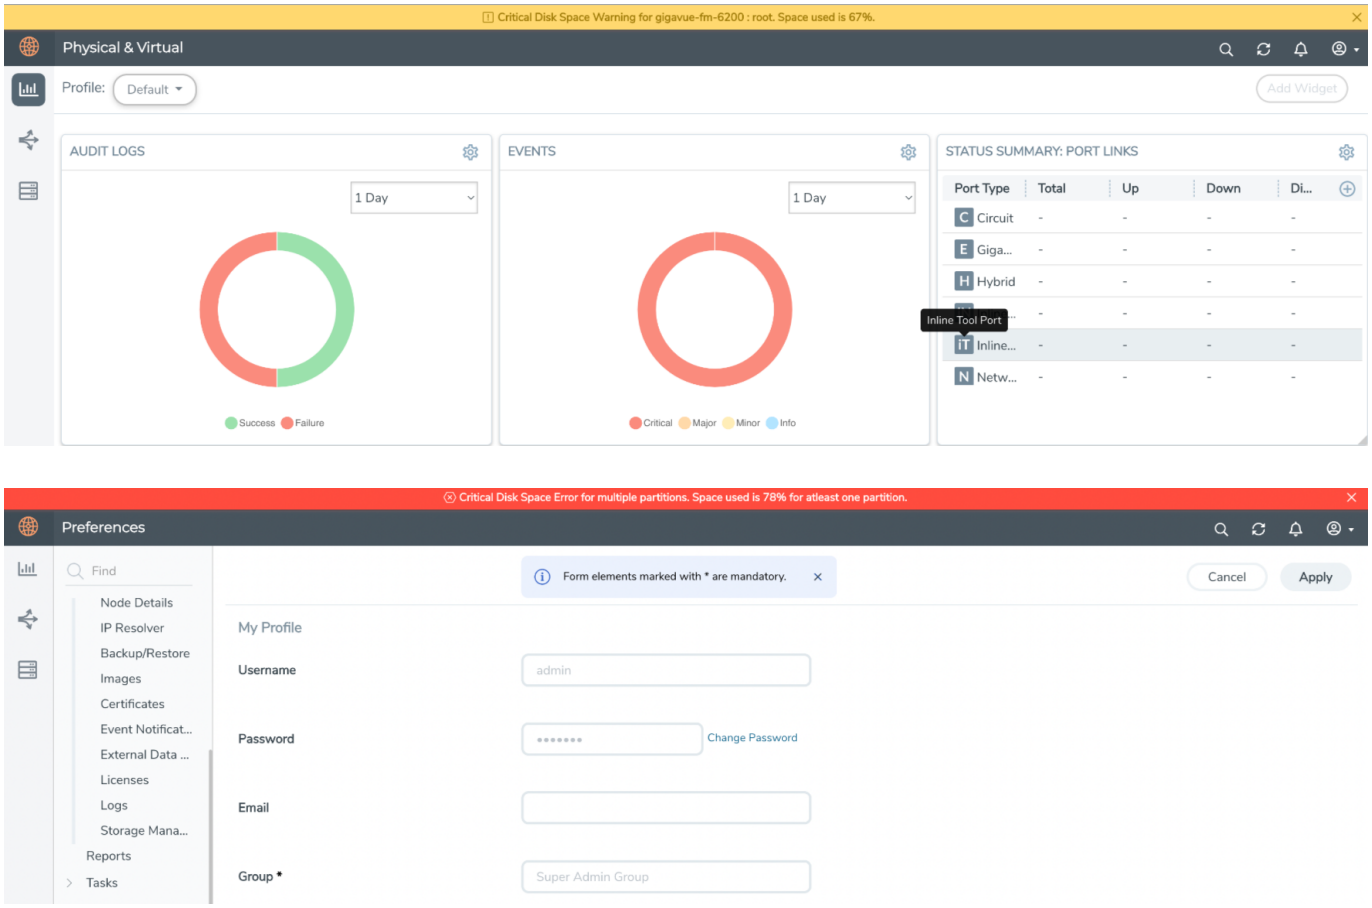

If you click on the notification banner ribbon, a pop-up with details about the disk space utilization is displayed. See the following example:

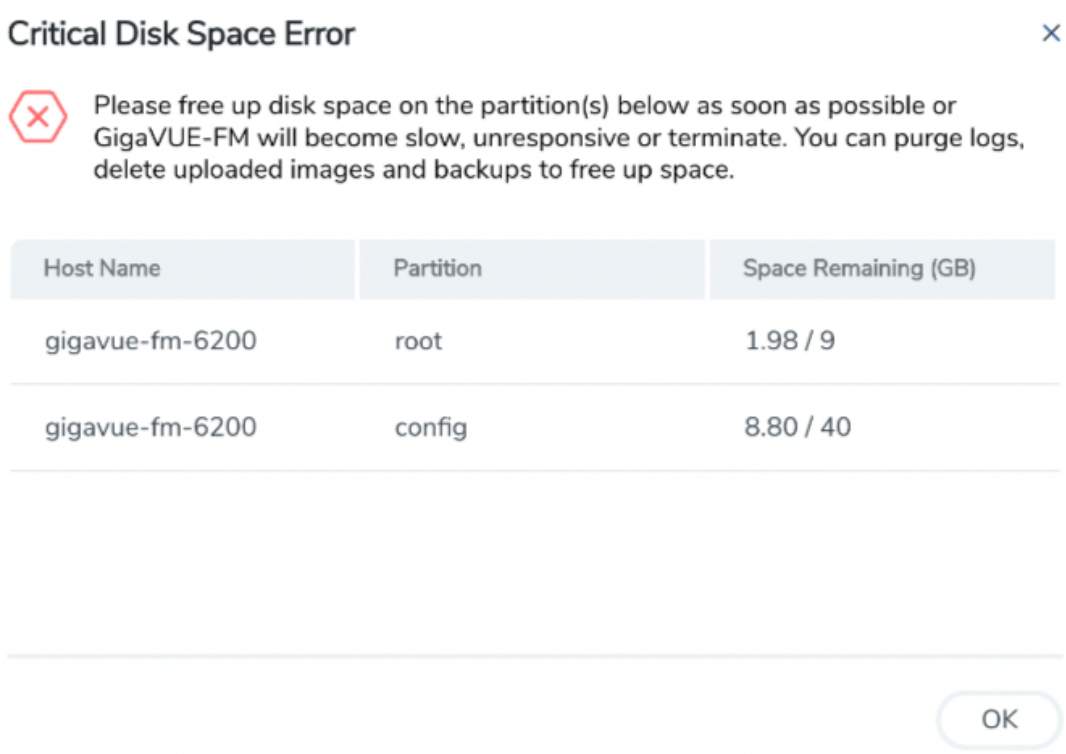

You may experience a slight delay while the notification banner pops-up and clears.

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. On the left navigation pane, click 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. 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
- 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.

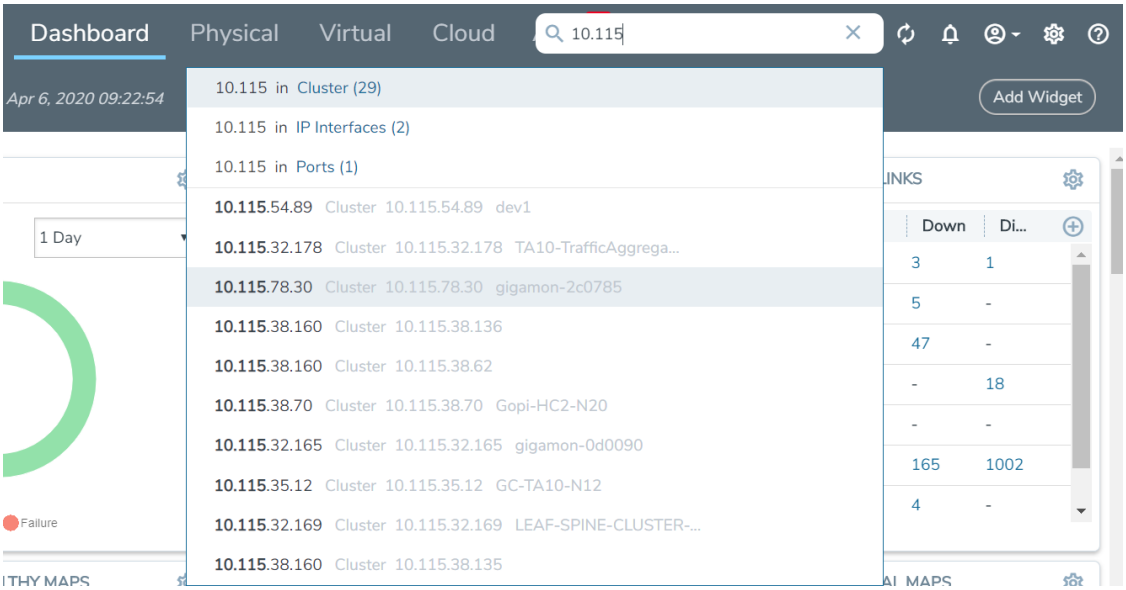

Figure 5 *GigaVUE-FM Search*

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.

Search Examples

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

- Searching Maps
- Searching for Roles and Users
- Searching Ports

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

Example 1: Searching 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.

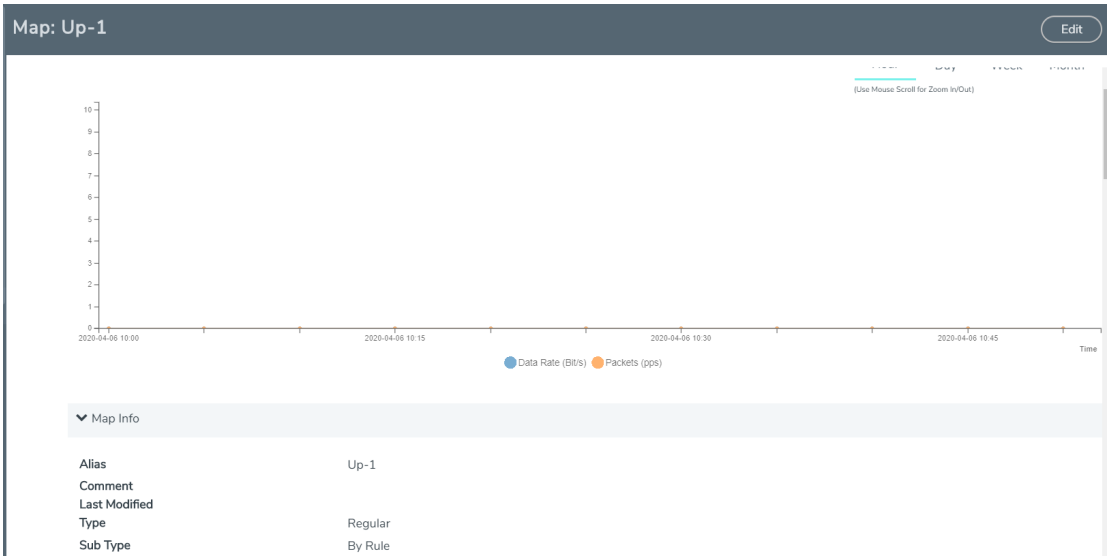

Figure 7 *Map Page*

Example 2: Searching 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.

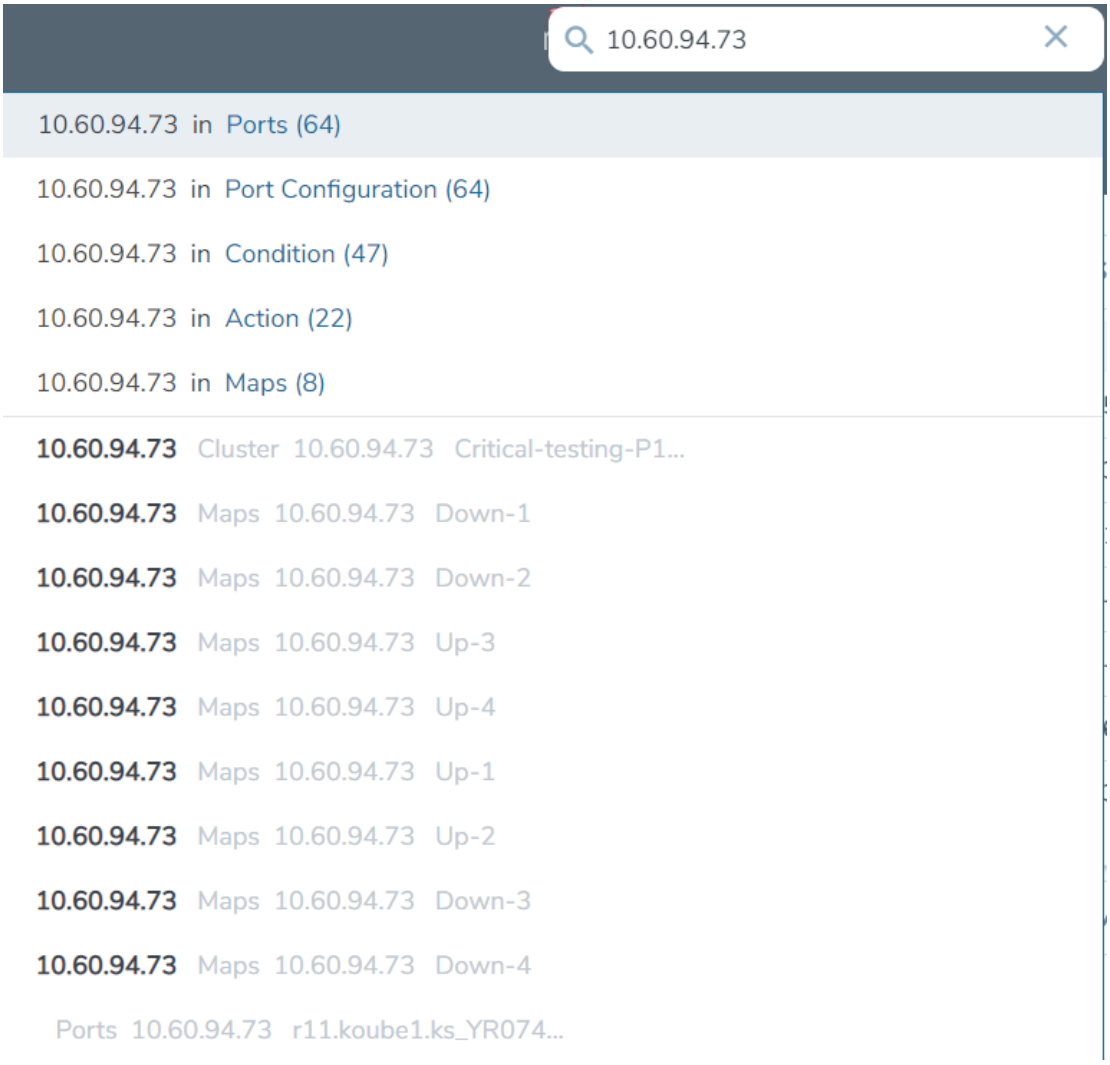

Figure 8 *Search for IP Address*

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.

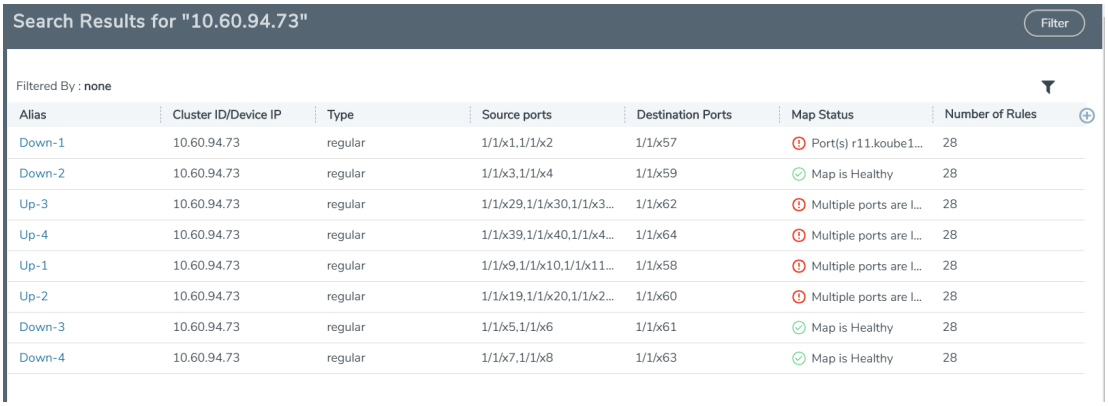

Alias	Cluster ID/Device IP	Type	Source ports	Destination Ports	Map Status	Number of Rules
Down-1	10.60.94.73	regular	1/1/x1,1/1/x2	1/1/x57	Port(s) r11.koube1...	28
Down-2	10.60.94.73	regular	1/1/x3,1/1/x4	1/1/x59	Map is Healthy	28
Up-3	10.60.94.73	regular	1/1/x29,1/1/x30,1/1/x3...	1/1/x62	Multiple ports 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: Searching 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: Searching 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.

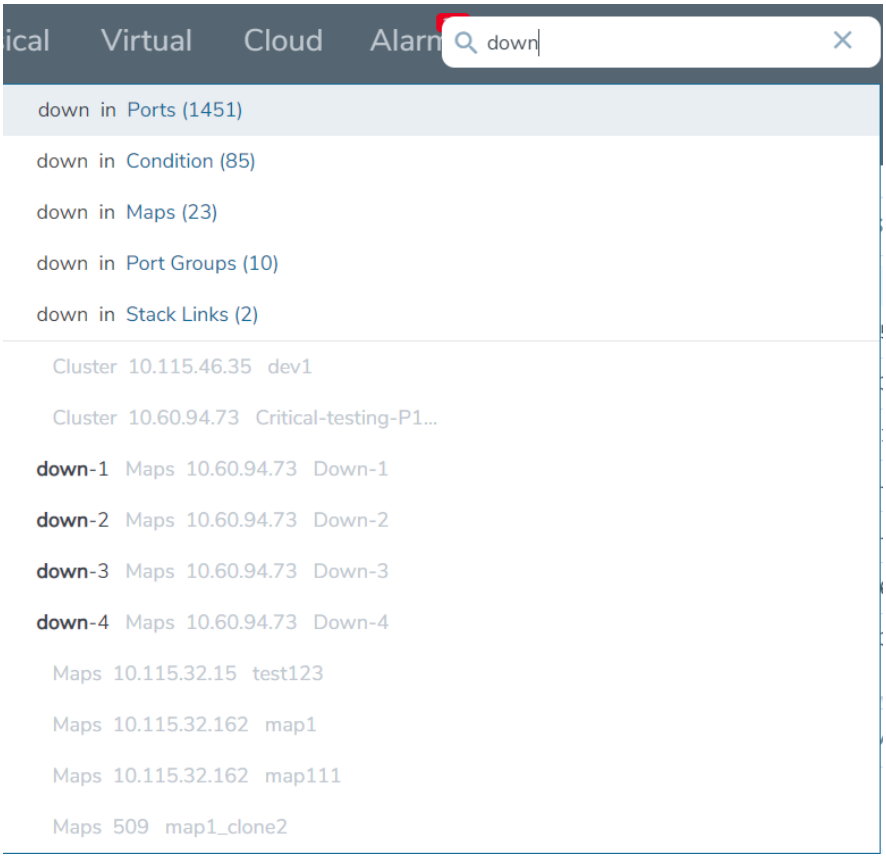

Figure 11 *Search Results for down Keyword*

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

Searching for Roles and Users

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

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

Example 1: Searching 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.

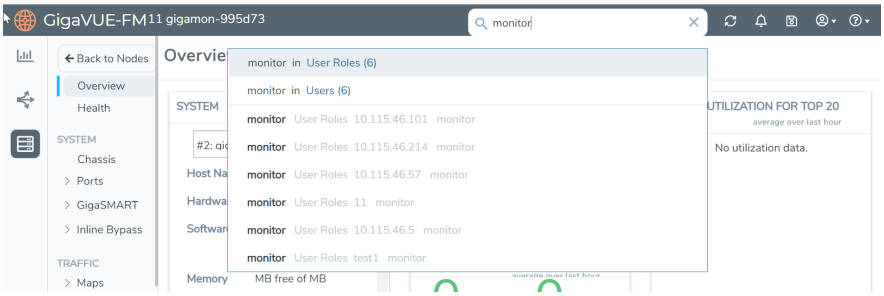

Figure 12 *Categories for Monitor.*

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

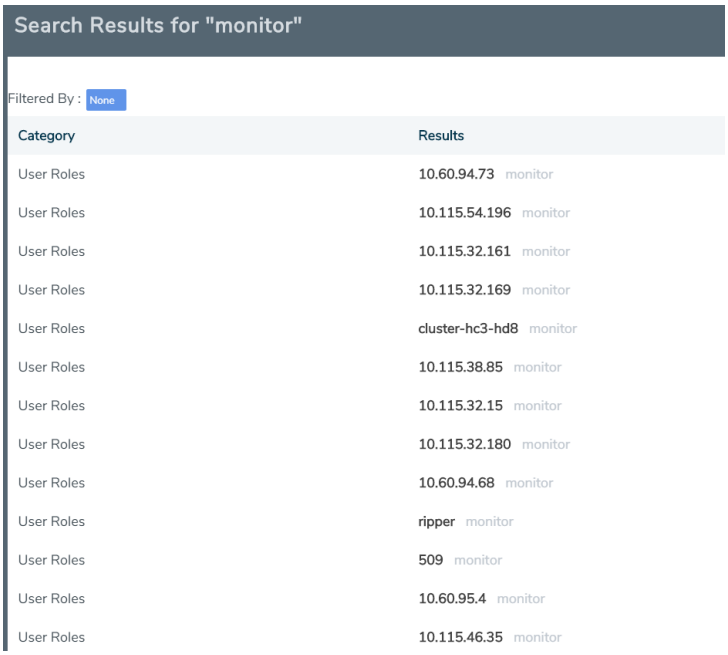

Figure 13 *Search Results for Monitor.*

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.

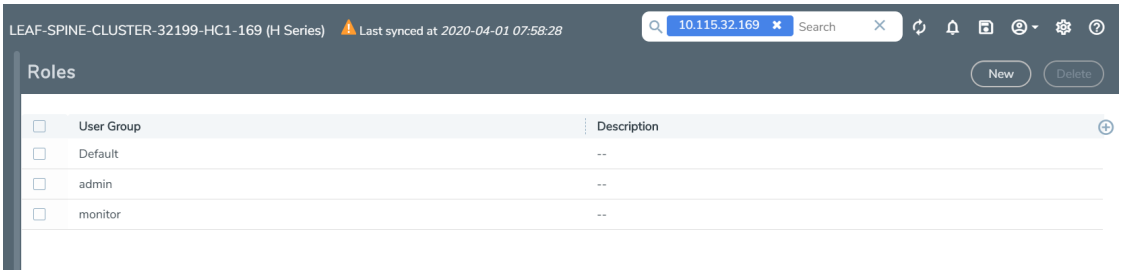

Figure 14 *Role From Search Result*

Example 2: Searching 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*.

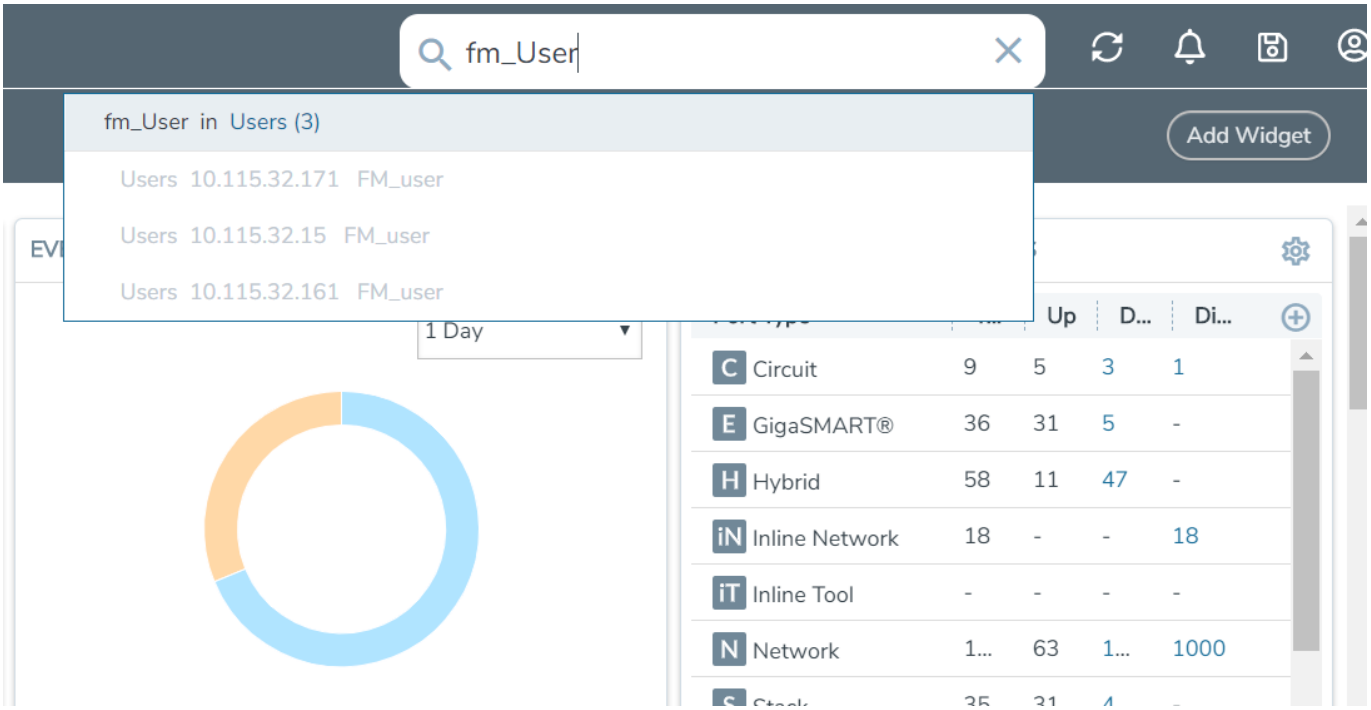

2. Click on *fm_User* in the categories and items list. The User page on the node with the *fm_User* opens a shown in the following figure. The Keywords field also indicates further searches are restricted to the current node.

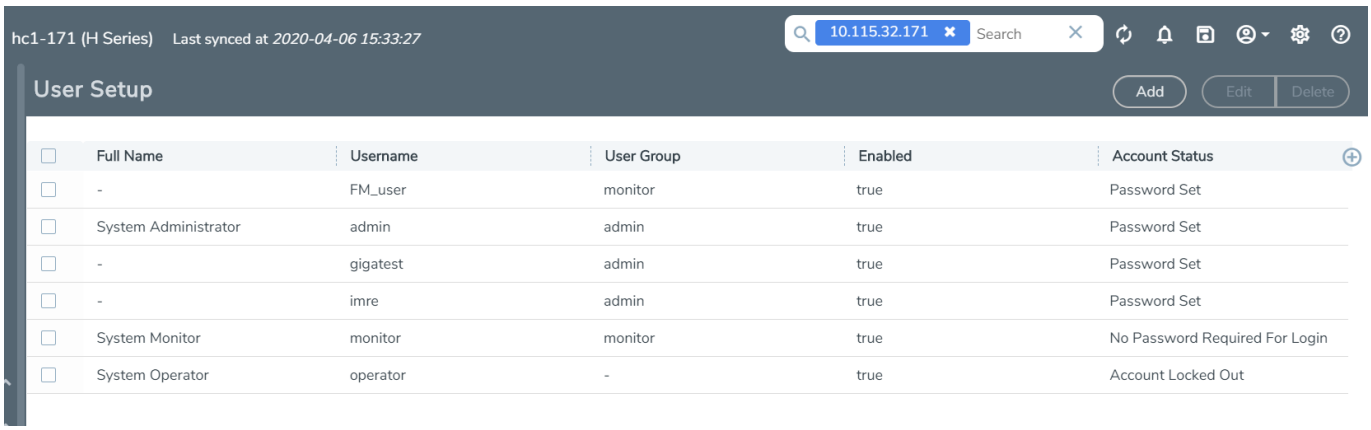

Searching Ports

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: Searching for Down Ports
- Example 2: Searching for Port Details of Devices Managed by GigaVUE-FM

Example 1: Searching 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.

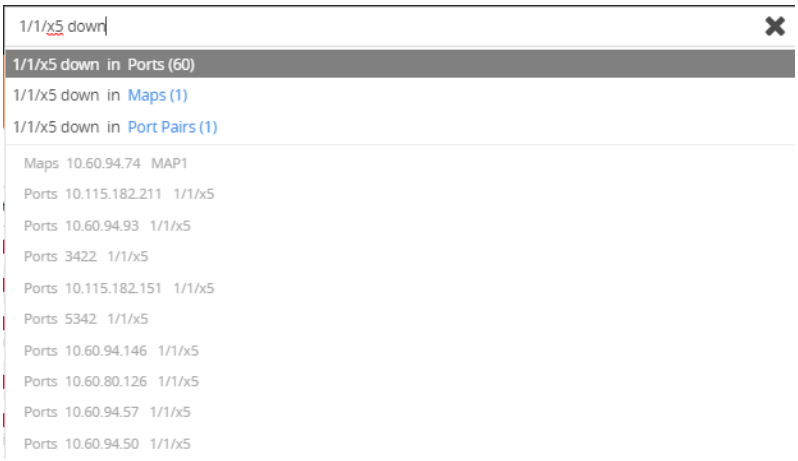

Figure 15 *Categories Returned for Port ID and Down*

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.

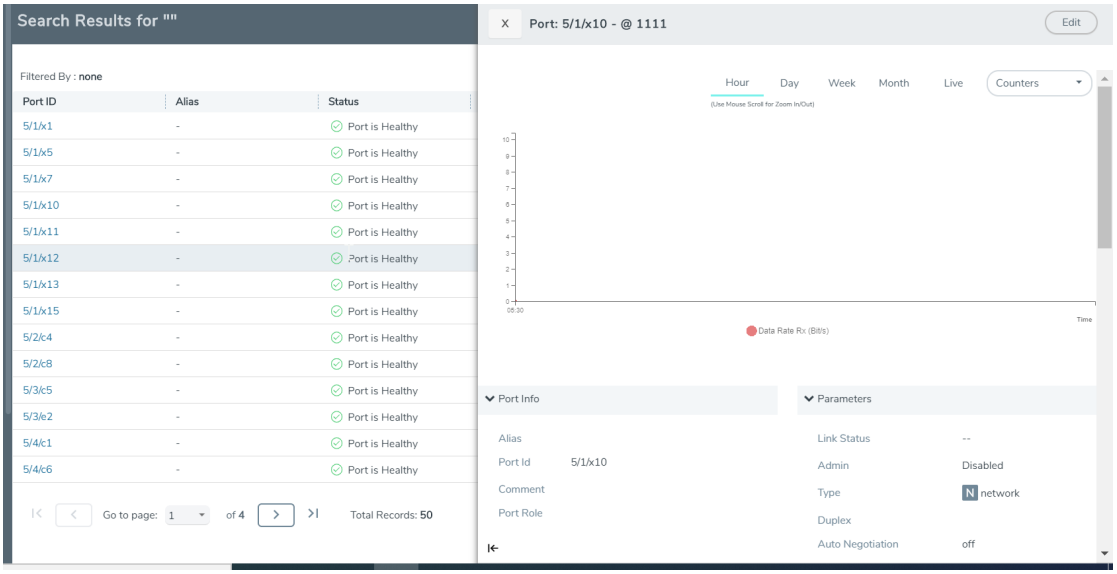

Figure 16 *Port Quick View for Search Results*

Example 2: Searching 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.

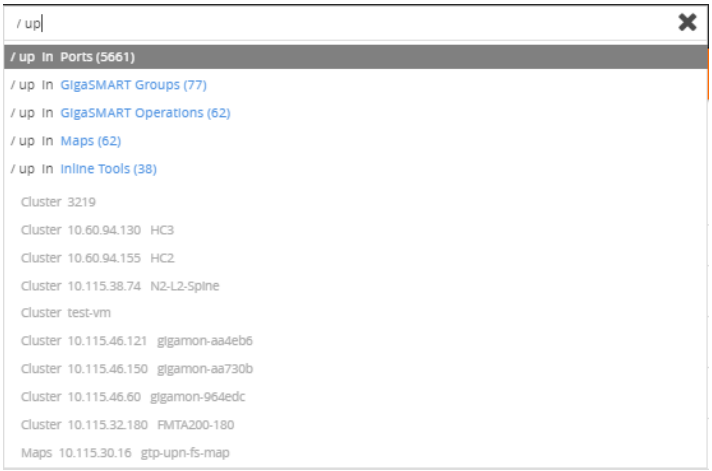

Figure 17 *Categories Returned for Port Details of All Devices and up*

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

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:

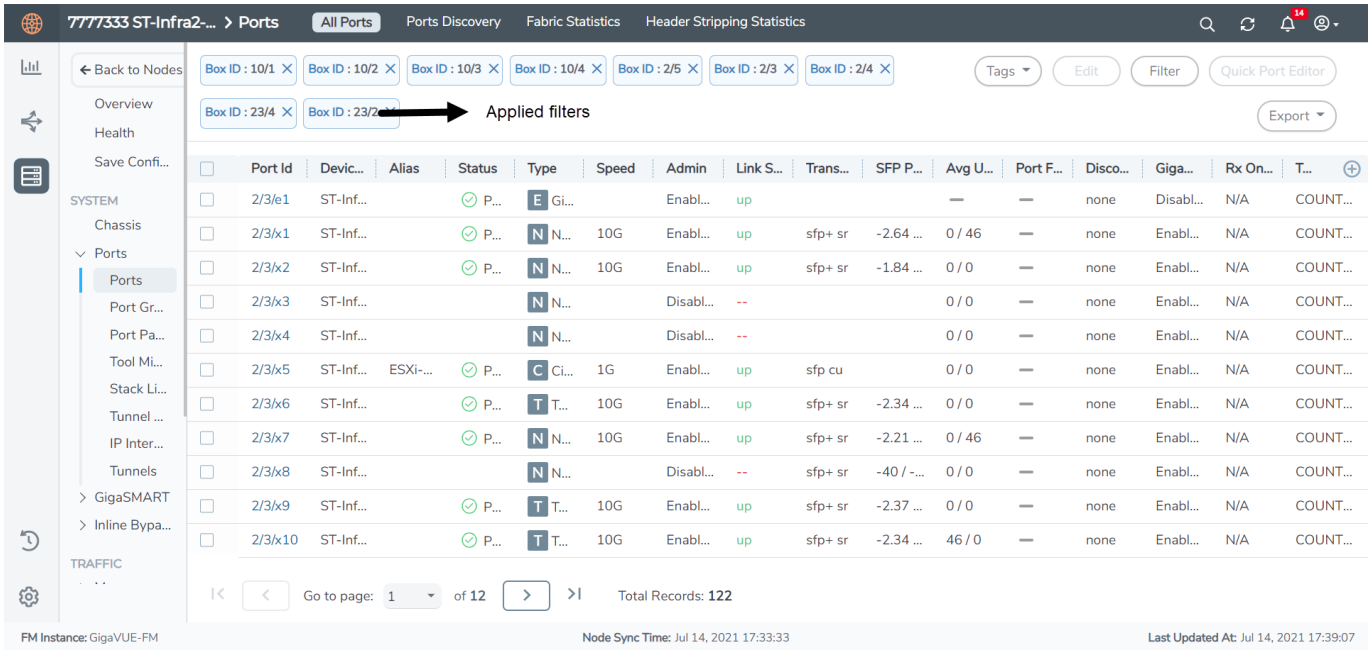
7777333 ST-Infra2-... > Ports
All Ports
Ports Discovery
Fabric Statistics
Header Stripping Statistics
Back to Nodes
Overview
Health
Save Confi...
SYSTEM
Chassis
Ports
Ports
Port Gr...
Port Pa...
Tool Mi...
Stack Li...
Tunnel ...
IP Inter...
Tunnels
GigaSMART
Inline Bypa...
TRAFFIC
Box ID : 10/1
Box ID : 10/2
Box ID : 10/3
Box ID : 10/4
Box ID : 2/5
Box ID : 2/3
Box ID : 2/4
Box ID : 23/4
Box ID : 23/2
Applied filters
Tags
Edit
Filter
Quick Port Editor
Export
Port Id
Devic...
Alias
Status
Type
Speed
Admin
Link S...
Trans...
SFP P...
Avg U...
Port F...
Disco...
Giga...
Rx On...
T...
2/3/e1 ST-Inf... P... E Gi... Enabl... up — — none Disabl... N/A COUNT...
2/3/x1 ST-Inf... P... N N... 10G Enabl... up sfp+ sr -2.64 ... 0 / 46 — none Enabl... N/A COUNT...
2/3/x2 ST-Inf... P... N N... 10G Enabl... up sfp+ sr -1.84 ... 0 / 0 — none Enabl... N/A COUNT...
2/3/x3 ST-Inf... N N... Disabl... -- 0 / 0 — none Enabl... N/A COUNT...
2/3/x4 ST-Inf... N N... Disabl... -- 0 / 0 — none Enabl... N/A COUNT...
2/3/x5 ST-Inf... ESXi-... P... C Ci... 1G Enabl... up sfp cu 0 / 0 — none Enabl... N/A COUNT...
2/3/x6 ST-Inf... P... T T... 10G Enabl... up sfp+ sr -2.34 ... 0 / 0 — none Enabl... N/A COUNT...
2/3/x7 ST-Inf... P... N N... 10G Enabl... up sfp+ sr -2.21 ... 0 / 46 — none Enabl... N/A COUNT...
2/3/x8 ST-Inf... N N... Disabl... -- sfp+ sr -40 / -... 0 / 0 — none Enabl... N/A COUNT...
2/3/x9 ST-Inf... P... T T... 10G Enabl... up sfp+ sr -2.37 ... 0 / 0 — none Enabl... N/A COUNT...
2/3/x10 ST-Inf... P... T T... 10G Enabl... up sfp+ sr -2.34 ... 46 / 0 — none Enabl... N/A COUNT...
Go to page: 1 of 12
Total Records: 122
FM Instance: GigaVUE-FM
Node Sync Time: Jul 14, 2021 17:33:33
Last Updated At: Jul 14, 2021 17:39:07

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.

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**.

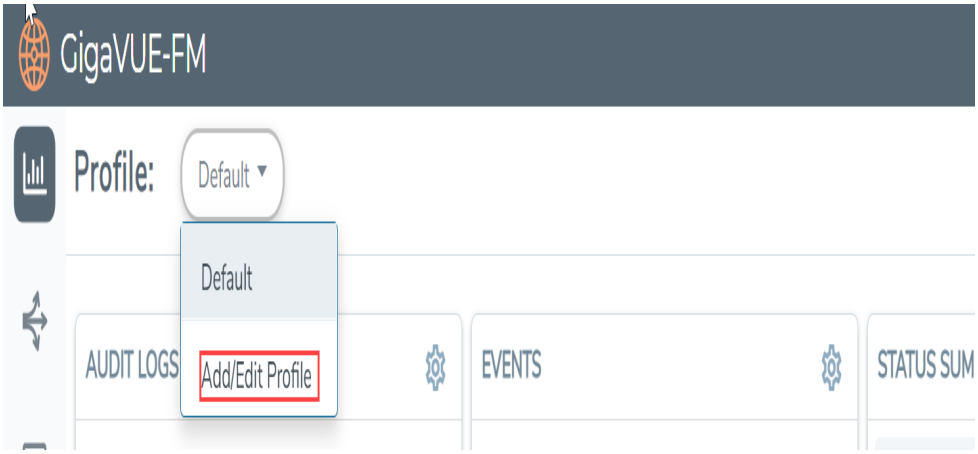

2. In the **Add/Edit Profile**... box, enter the name of the new profile and click **Enter**.

   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.

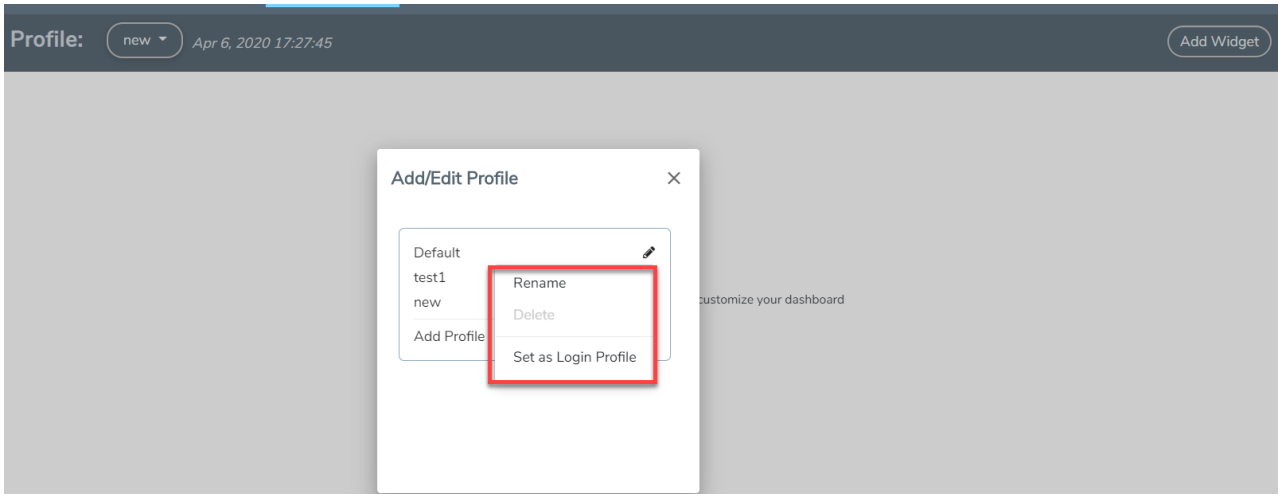

Figure 18 *Profile Settings*

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.

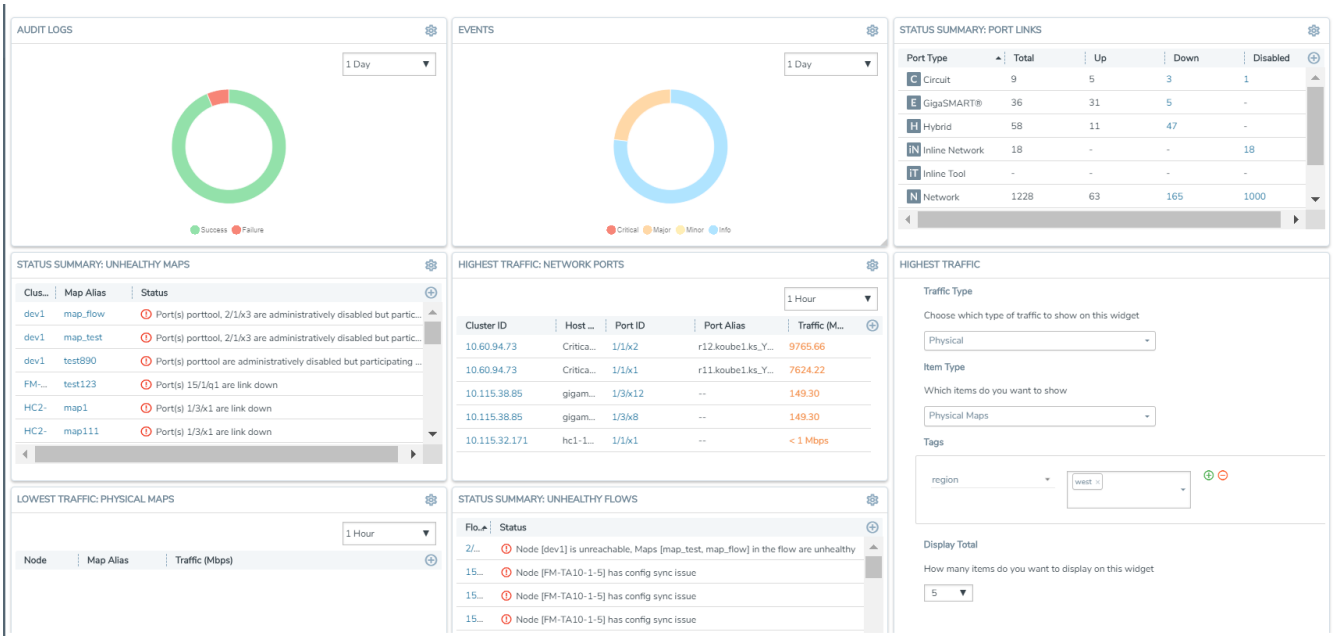

Note: If the percentage displayed in a pie-chart is negligible or less , then it would be difficult to click on the pie-chart arc and view the details .

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.

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.

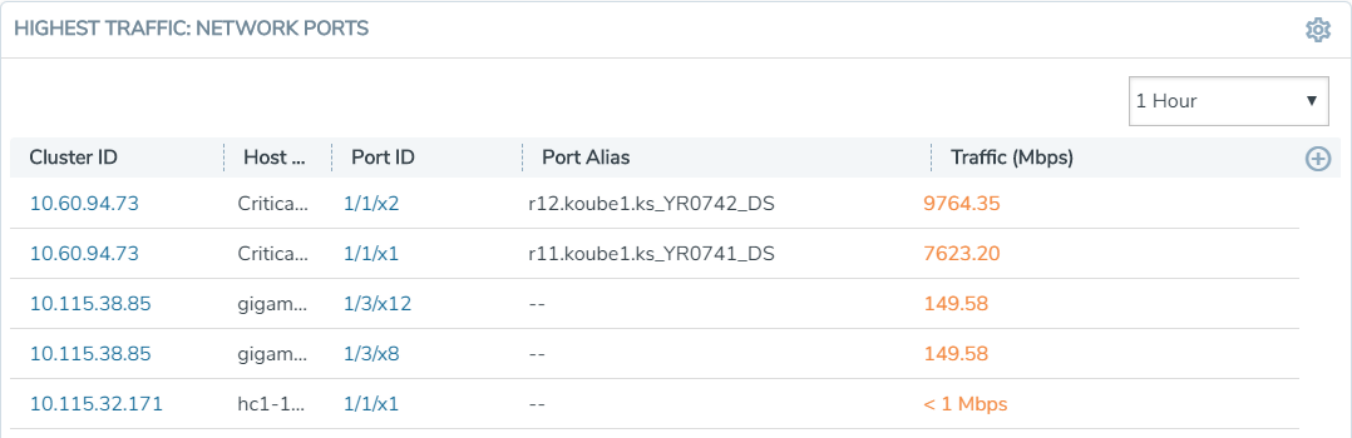

Figure 19 *Highest Traffic: Example*

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.

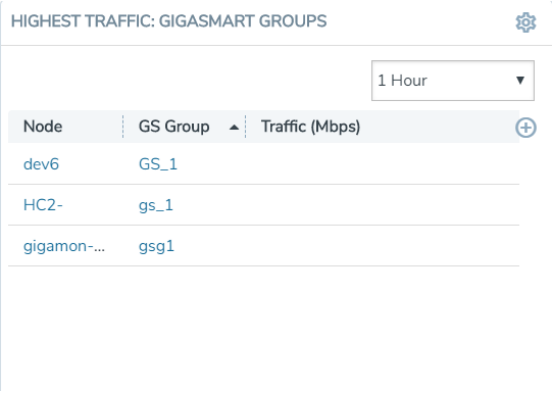

Figure 20 *Highest Traffic GigaSMART*

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.

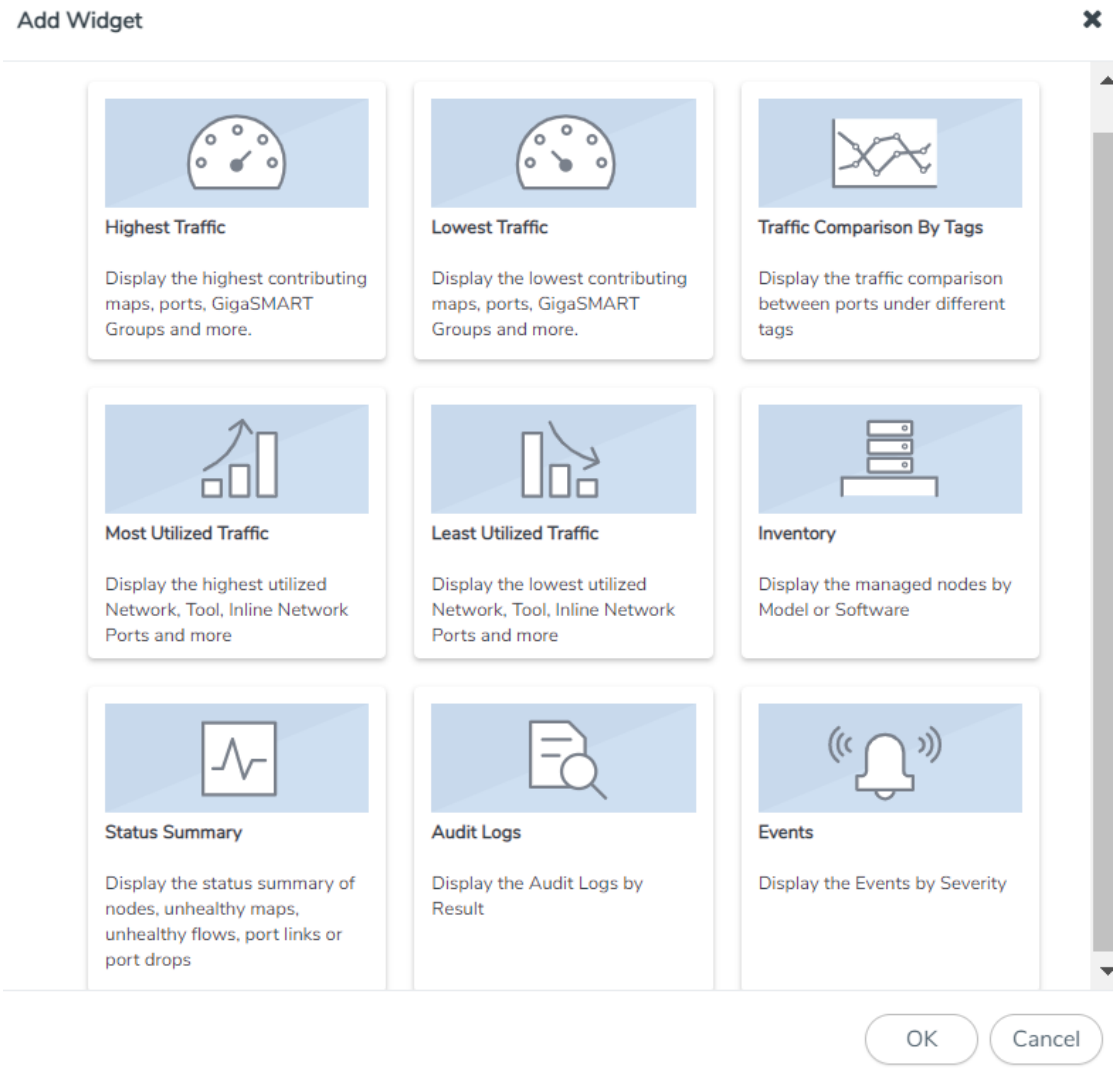

Figure 21 *Add New Widget*

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.

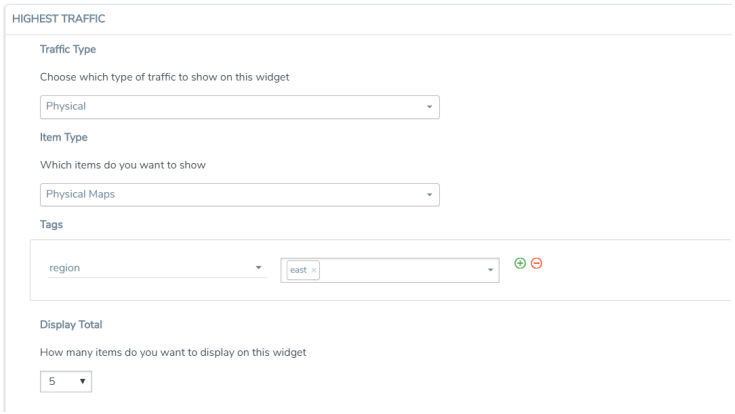

Figure 22 *Highest Traffic Configuration*

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.

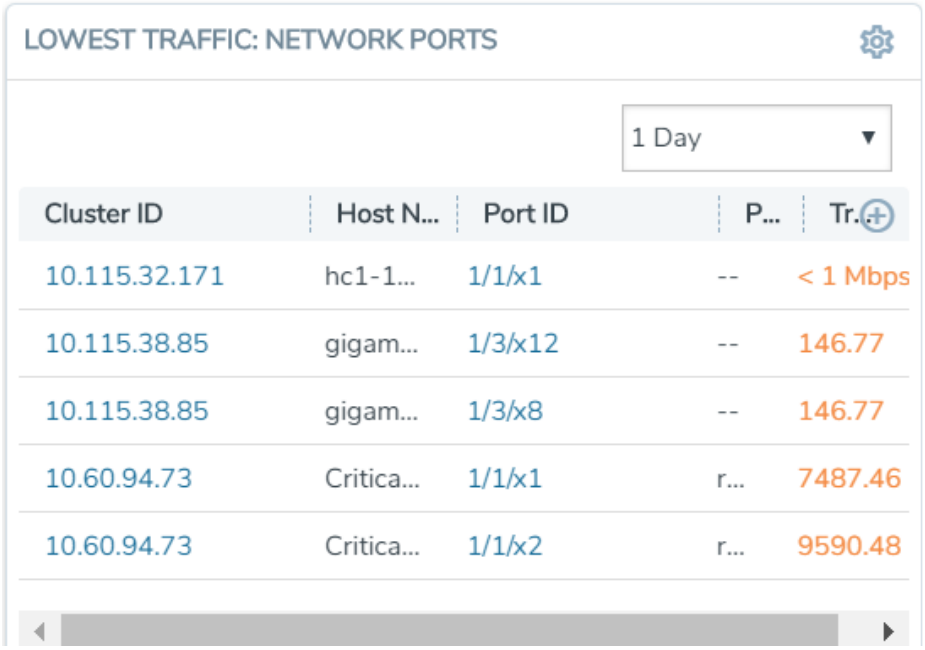

Figure 23 *Lowest Traffic*

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.

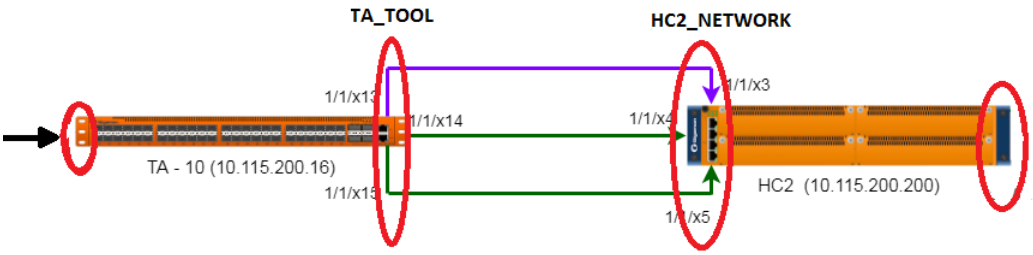

Figure 24 *Example for* Traffic Comparison By Tags Widget

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

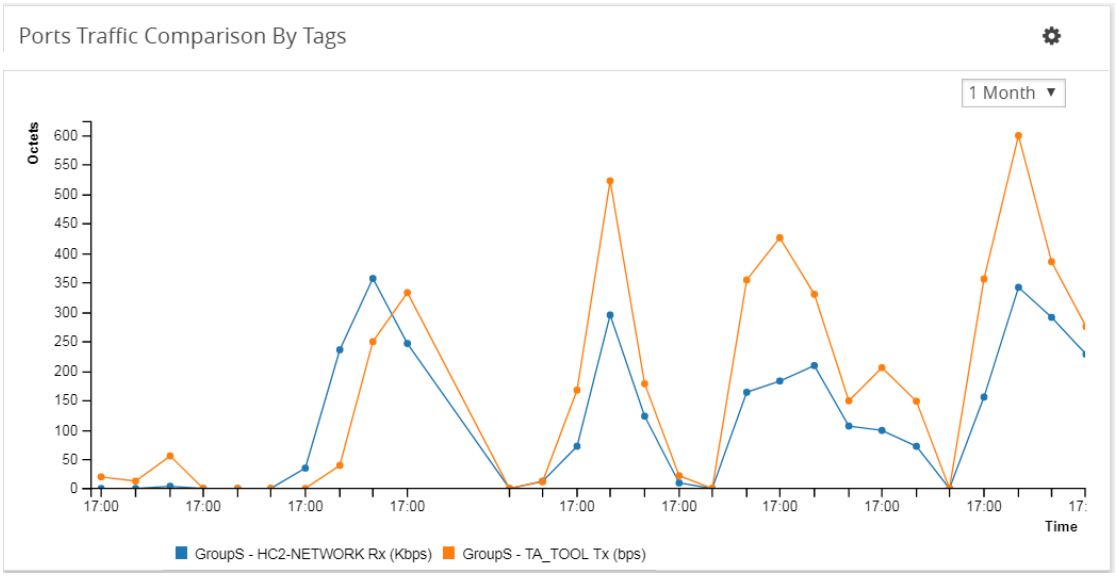

Figure 25 *Traffic Comparison By Tags*

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):

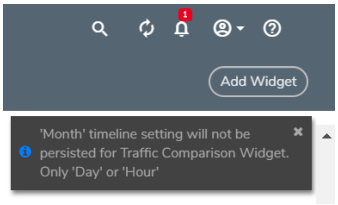

Figure 26 *Notifications*

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.

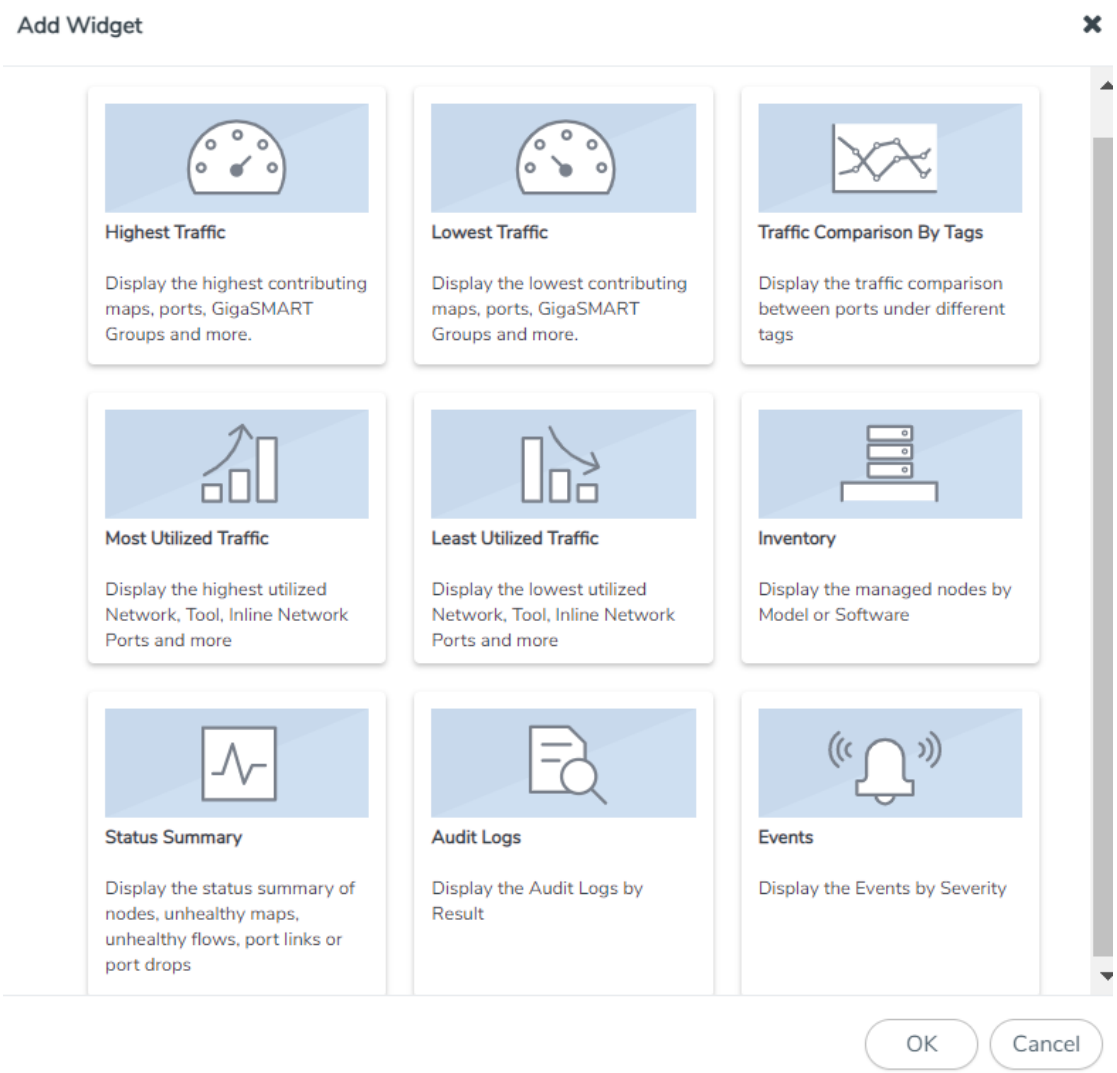

Figure 27 *Add New Widget*

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.

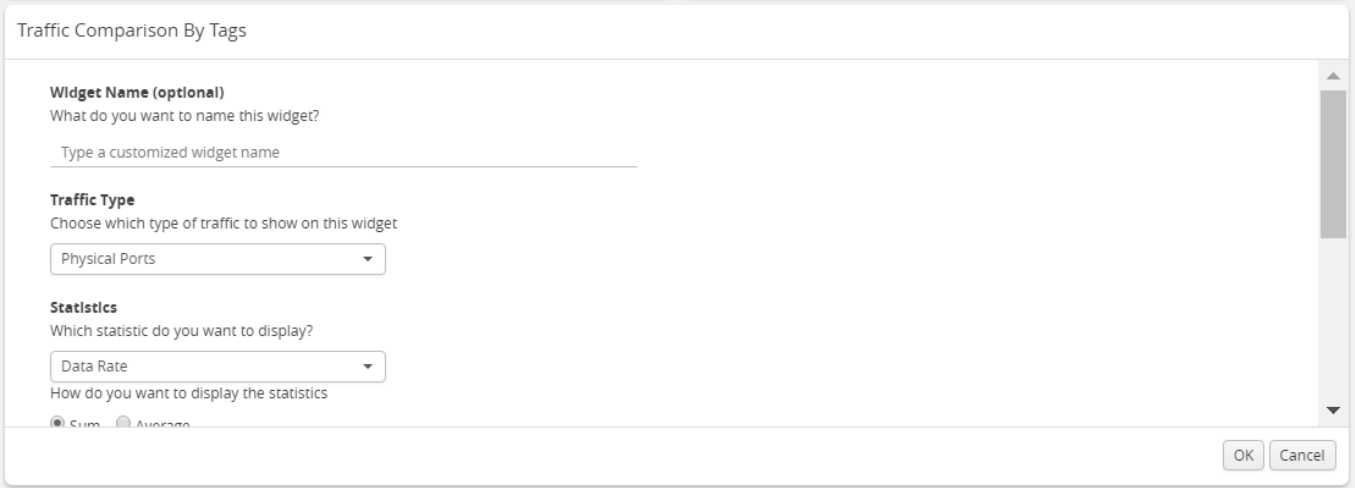

Figure 28 *Traffic Comparison By Tags Configuration*

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.

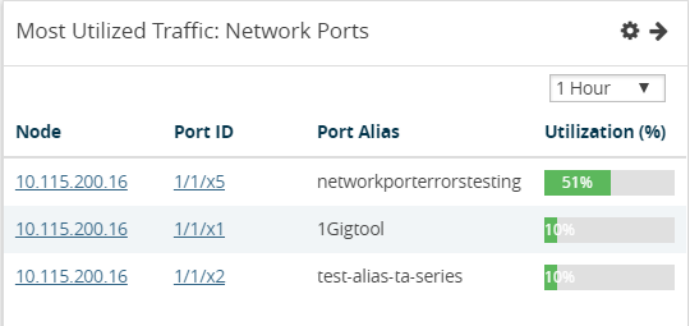

Figure 29 *Most Utilized Traffic Widget*

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.

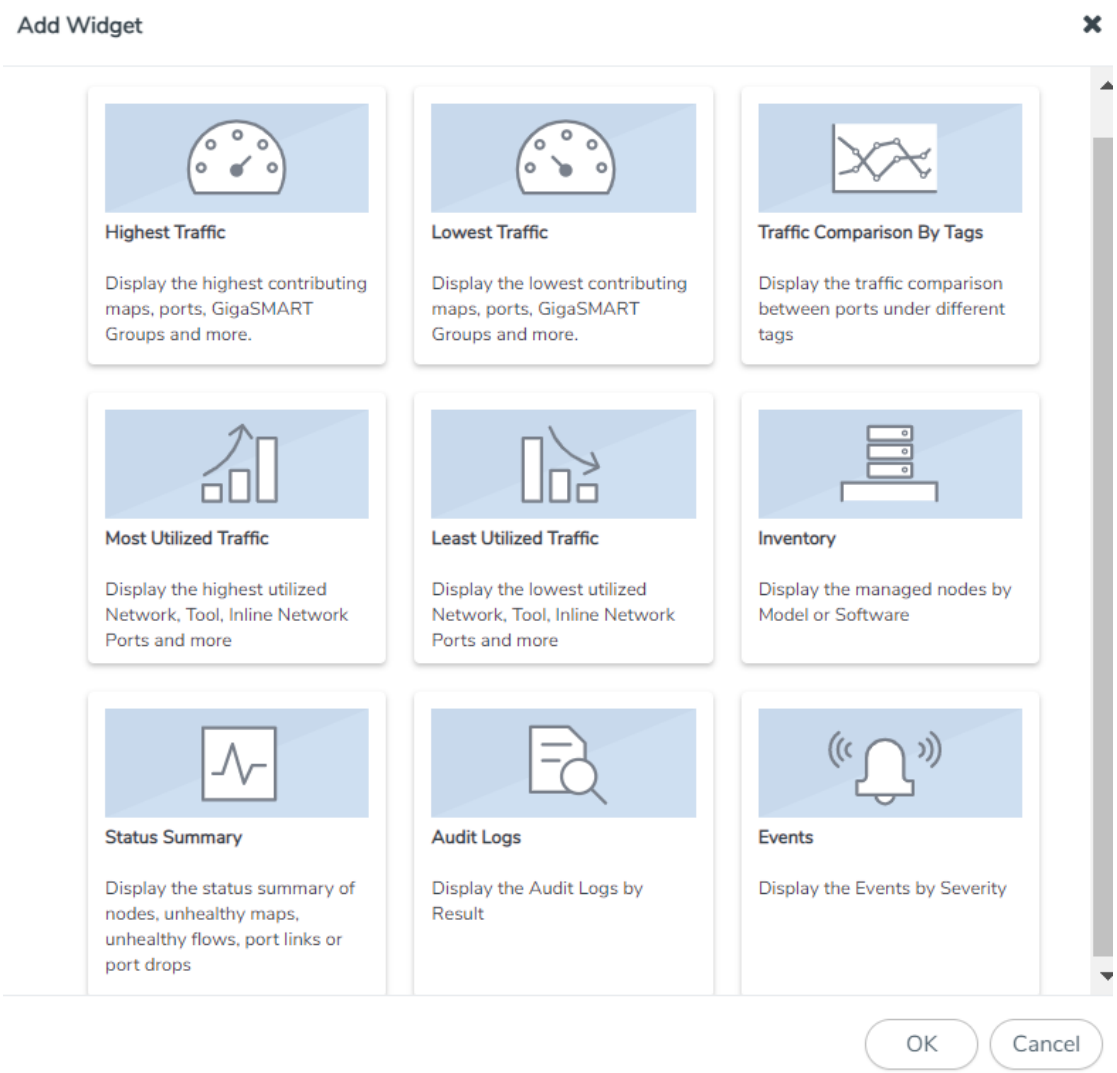

Figure 30 *Add New Widget*

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.

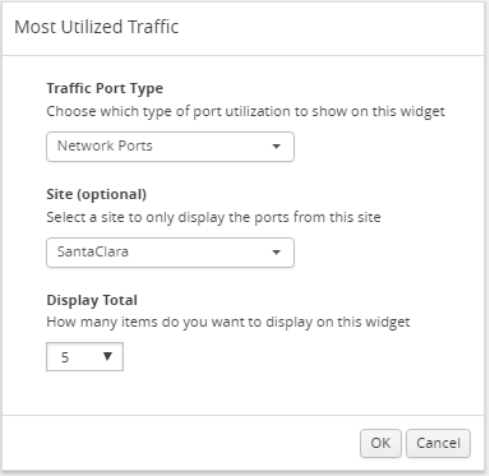

Figure 31 *Most Utilized Traffic Configuration*

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.

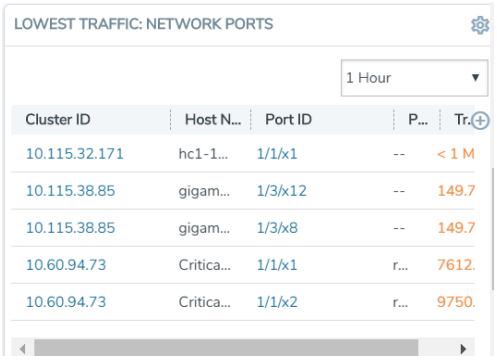

Figure 32 *Least Utilized Traffic*

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.

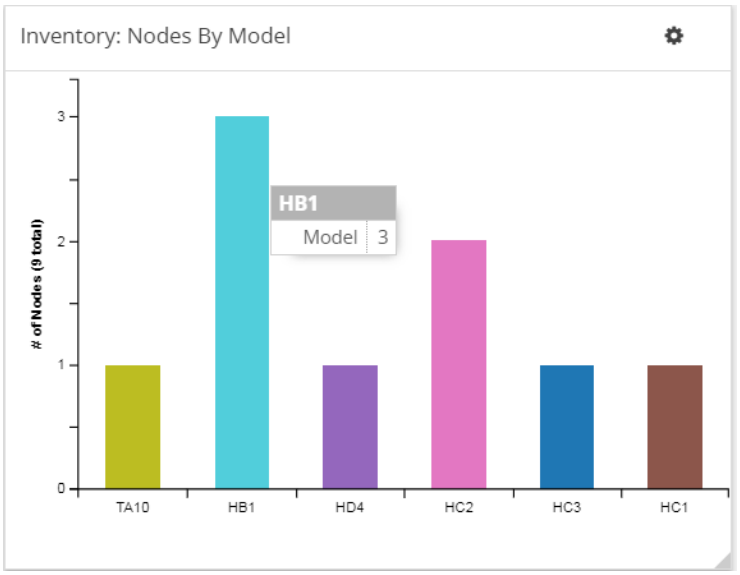

Figure 33 *Nodes by Model*

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.

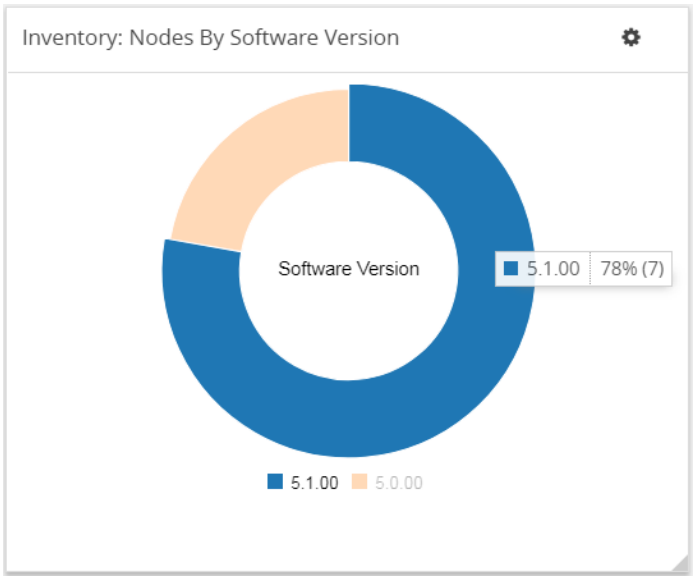

Figure 34 *Nodes by Software Version*

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.

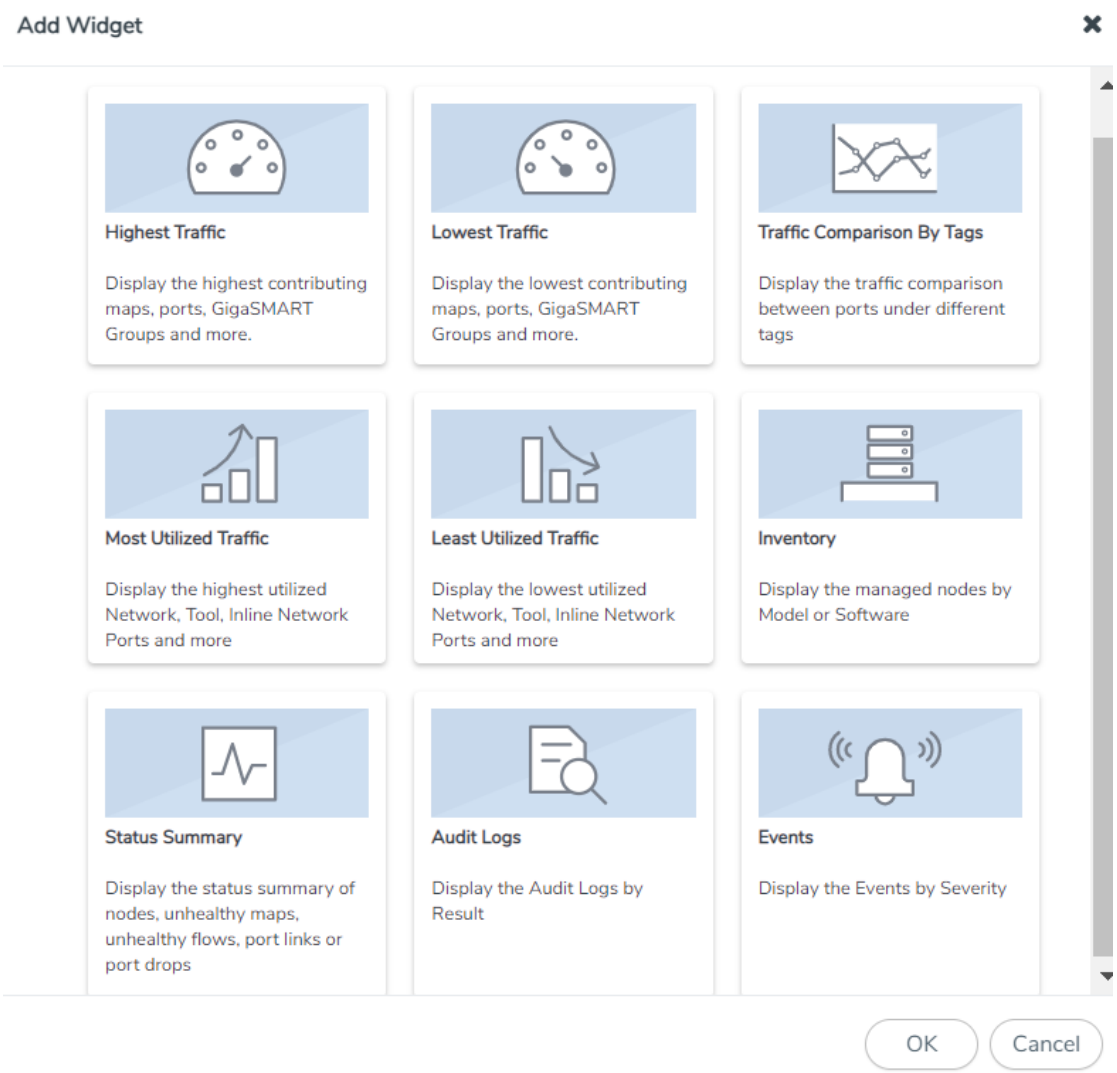

Figure 35 *Add New Widget*

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.

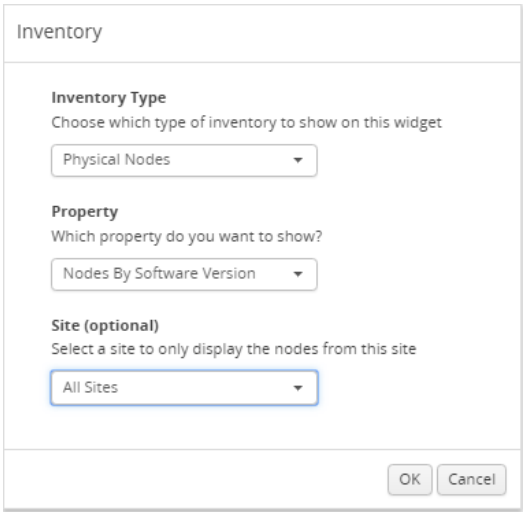

Figure 36 *Inventory Configuration*

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.

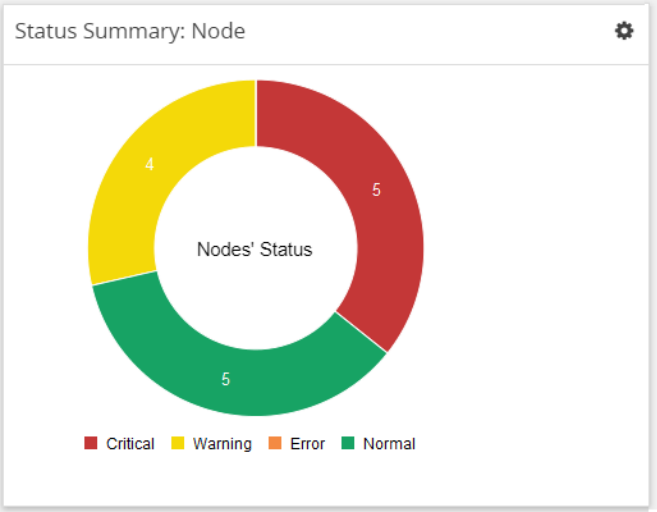

Figure 37 *Nodes' Status Widget*

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.

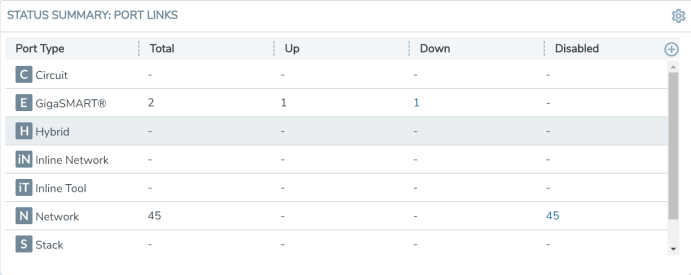

Figure 38 *Status Summary: Port Links Widget*

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.

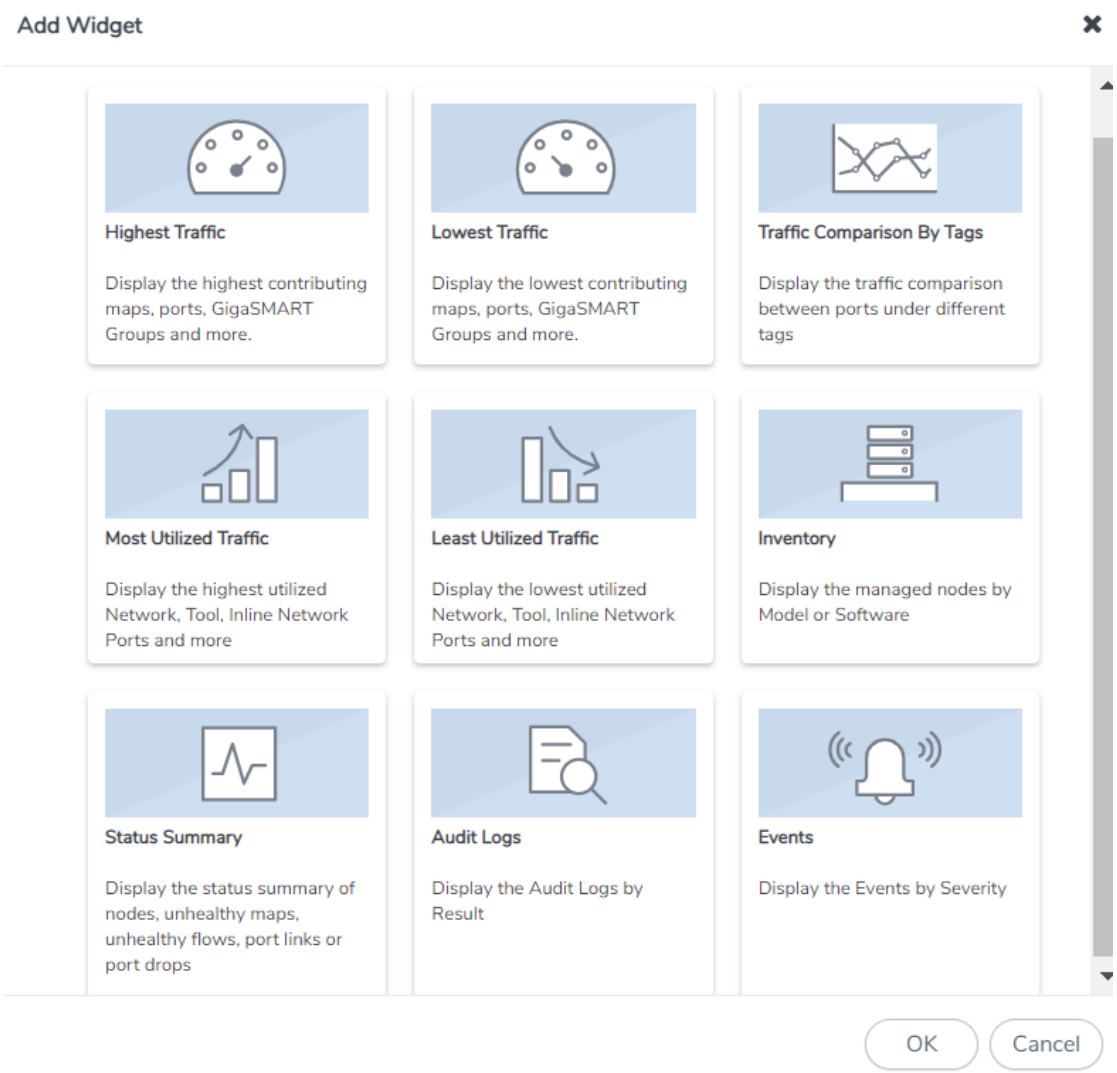

Figure 39 *Add New Widget*

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

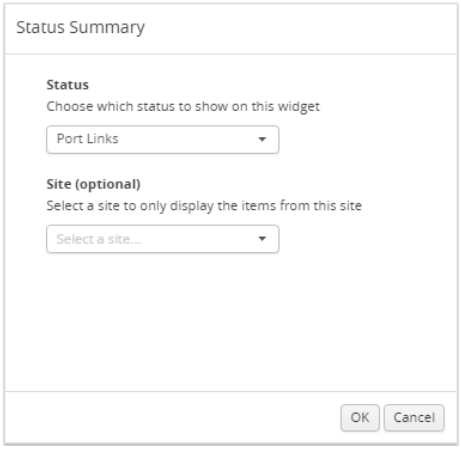

Figure 40 *Status Summary Configuration*

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.

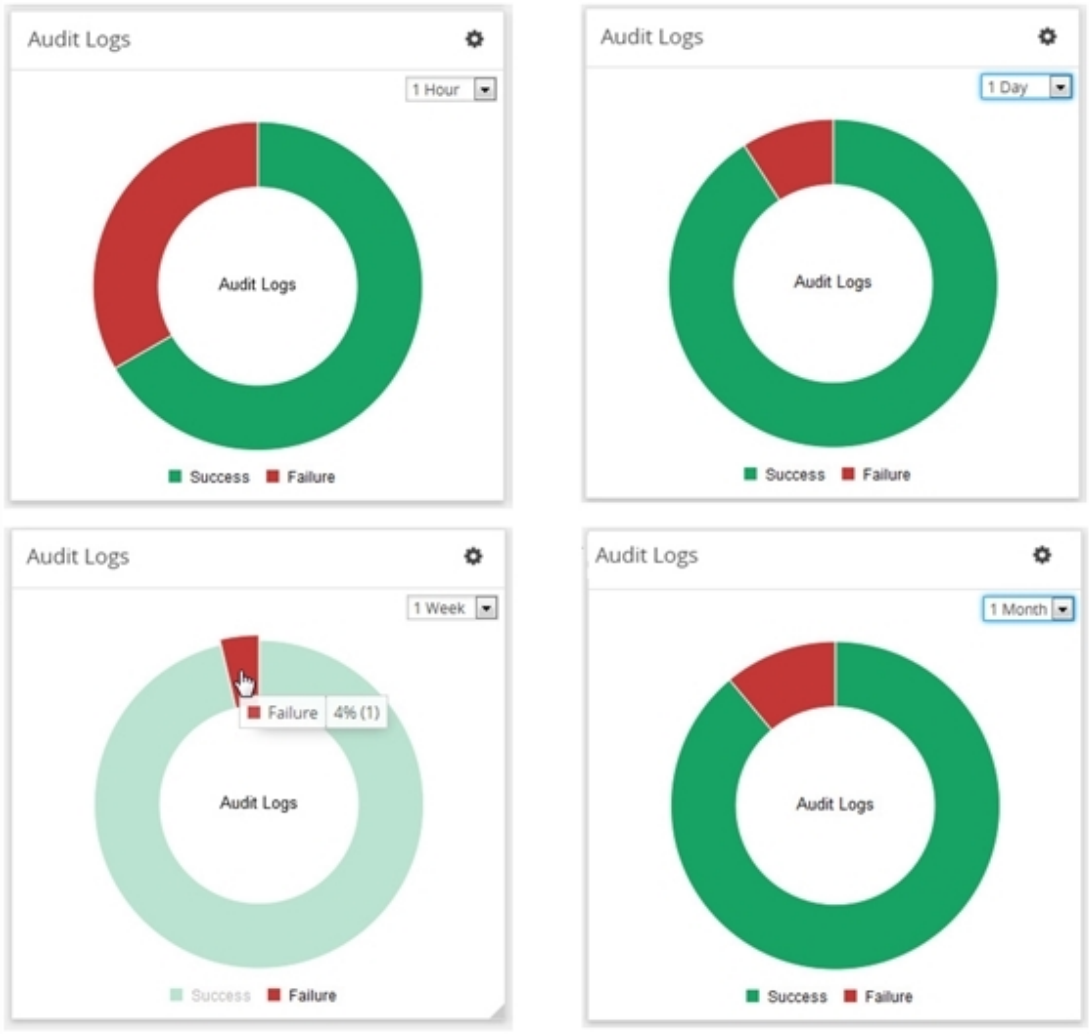

Figure 41 *Audit Logs by Result*

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.

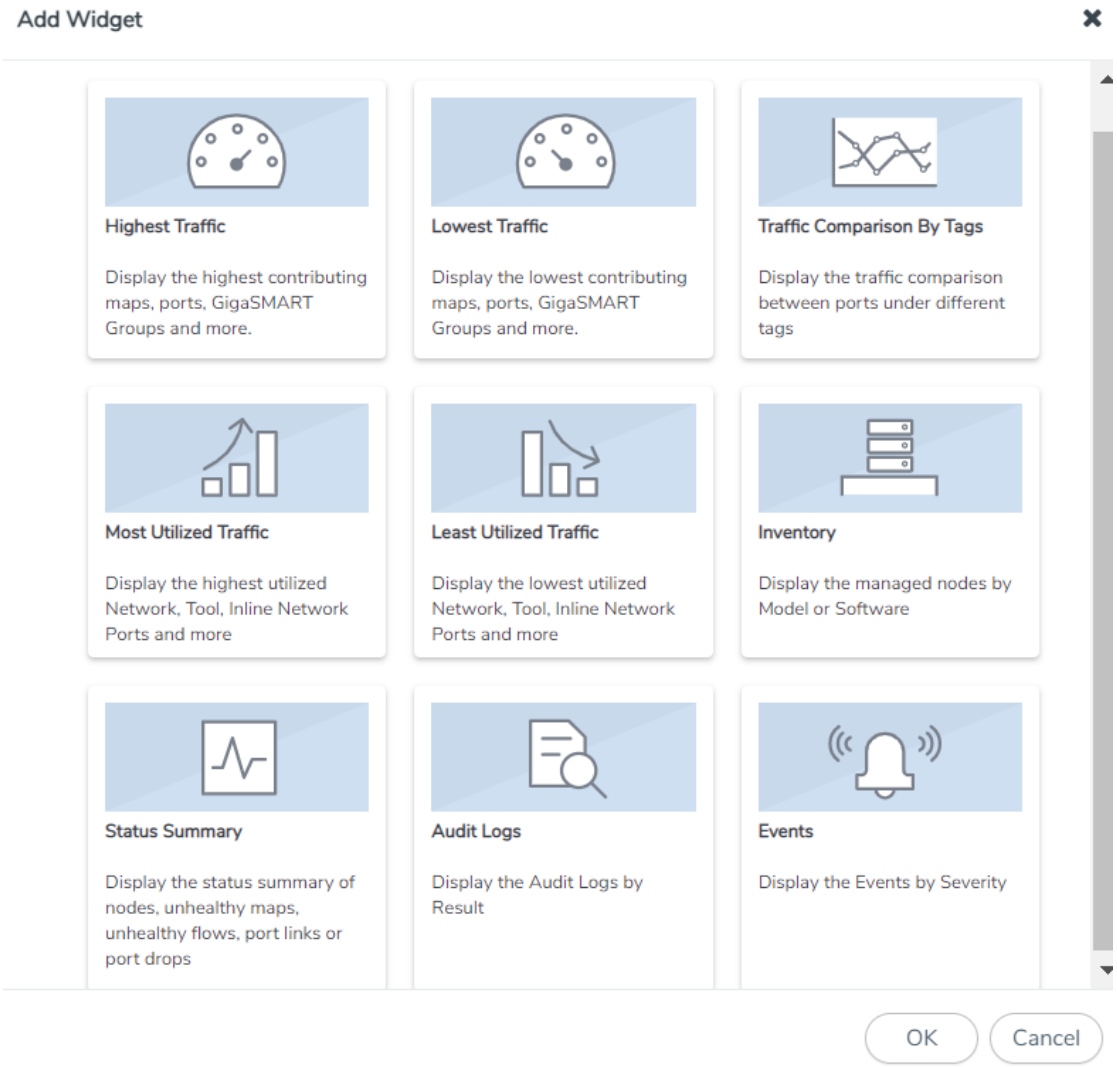

Figure 42 *Add New Widget*

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.

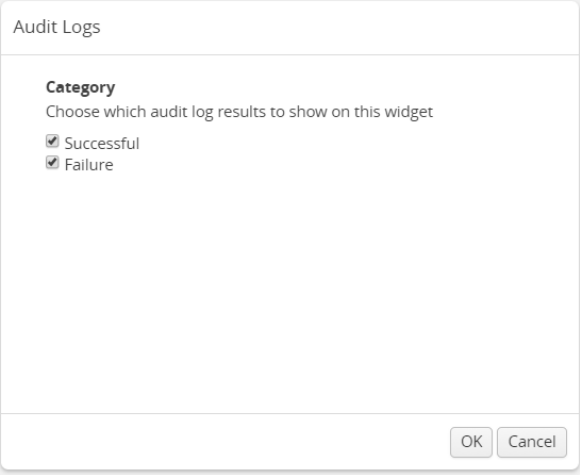

Figure 43 *Status Summary Configuration*

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.

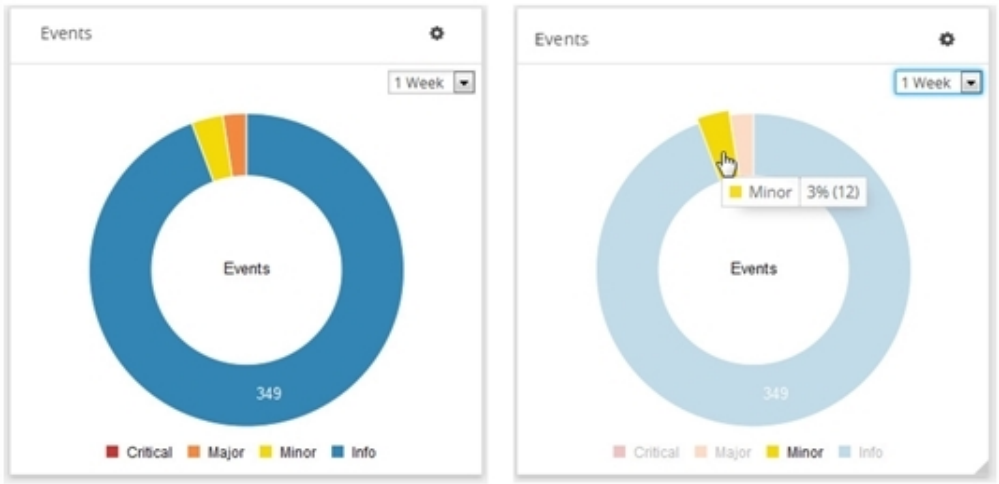

Figure 44 *Events Widget*

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.

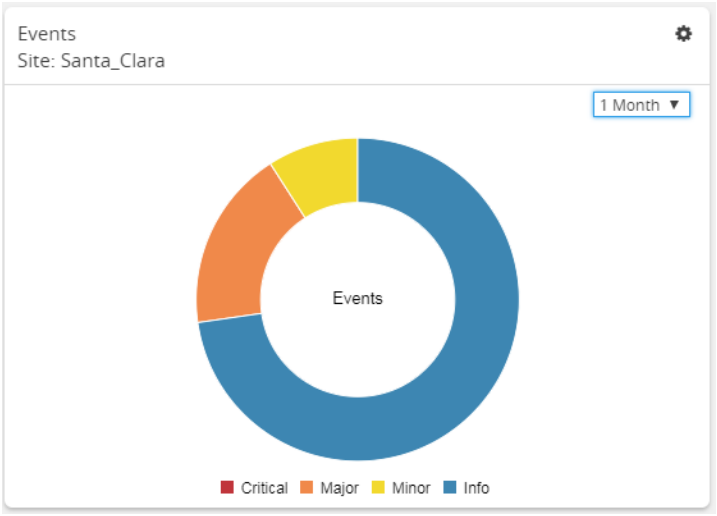

Figure 45 *Events Widget for Santa Clara*

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.

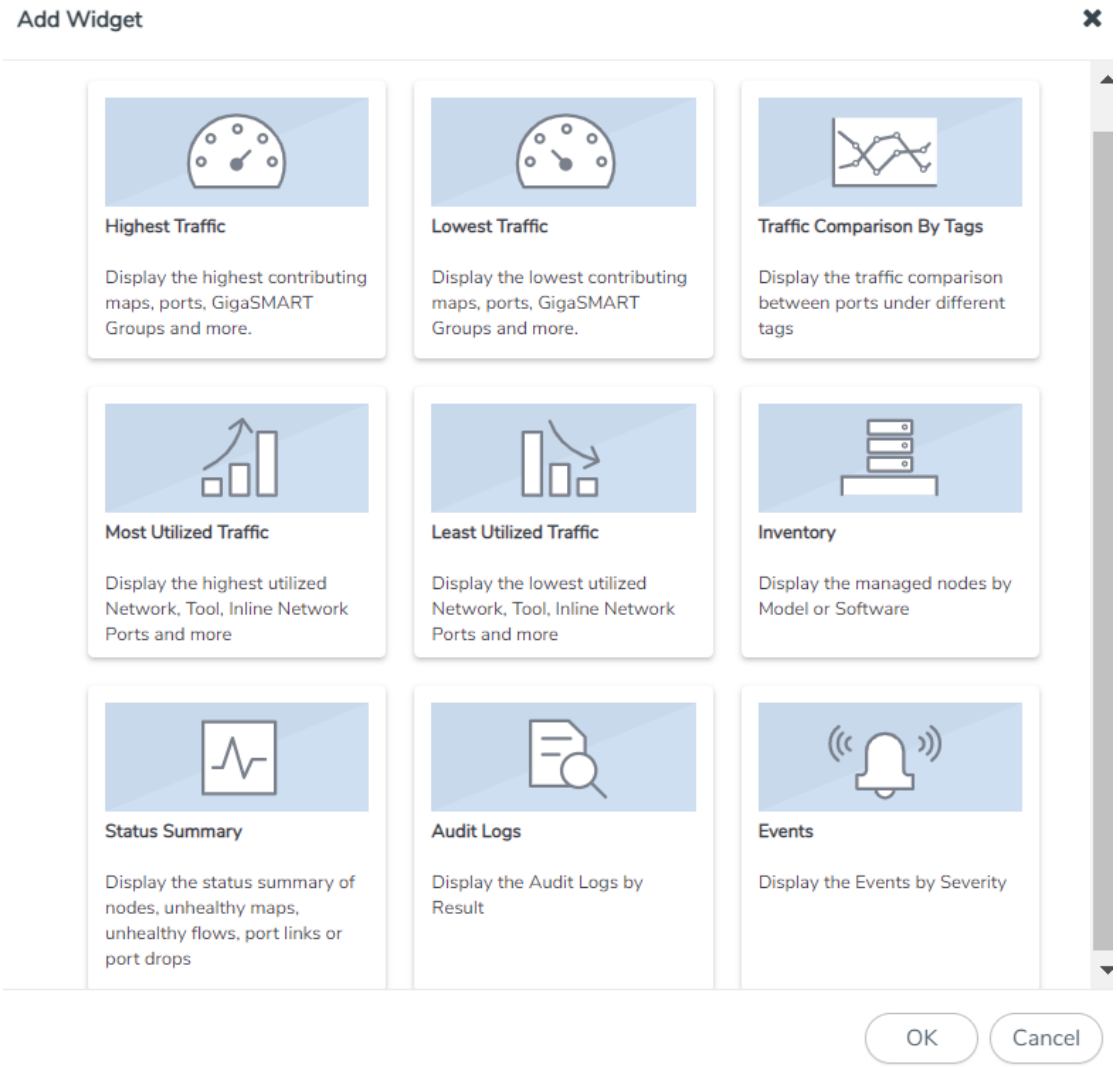

Figure 46 *Add New Widget*

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.

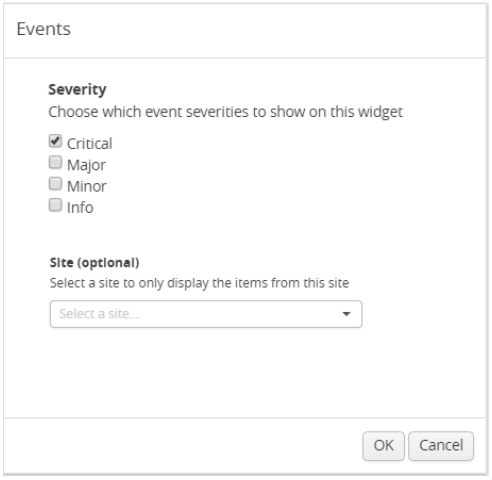

Figure 47 *Status Summary Configuration*

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

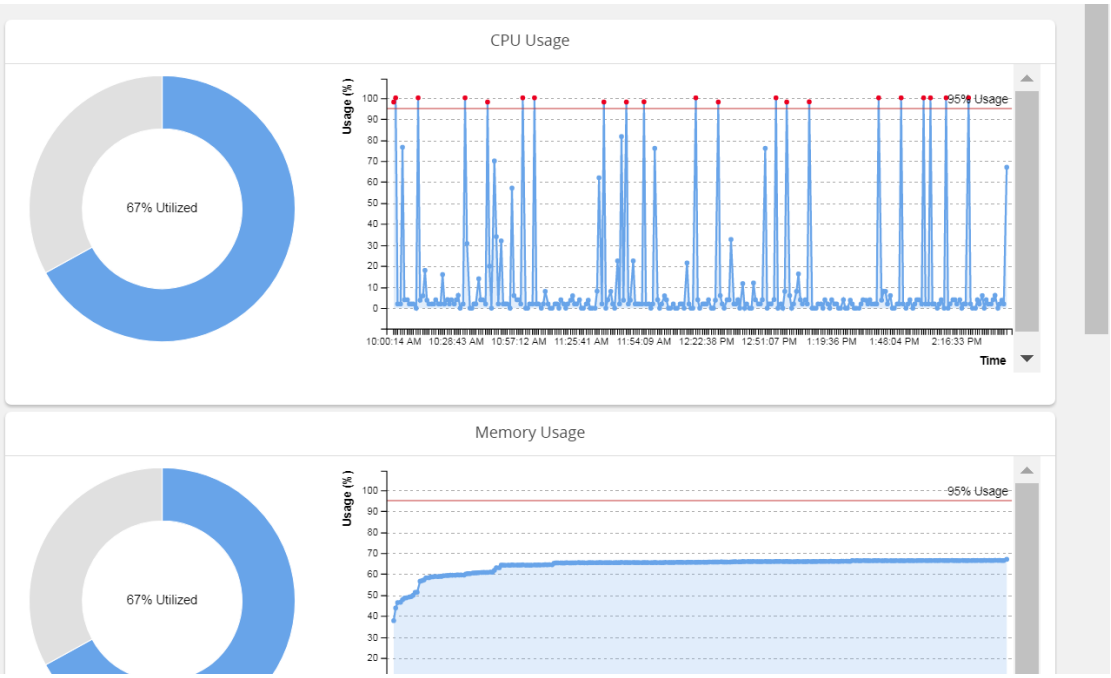

Figure 48 *Health Monitor Dashboards*

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.

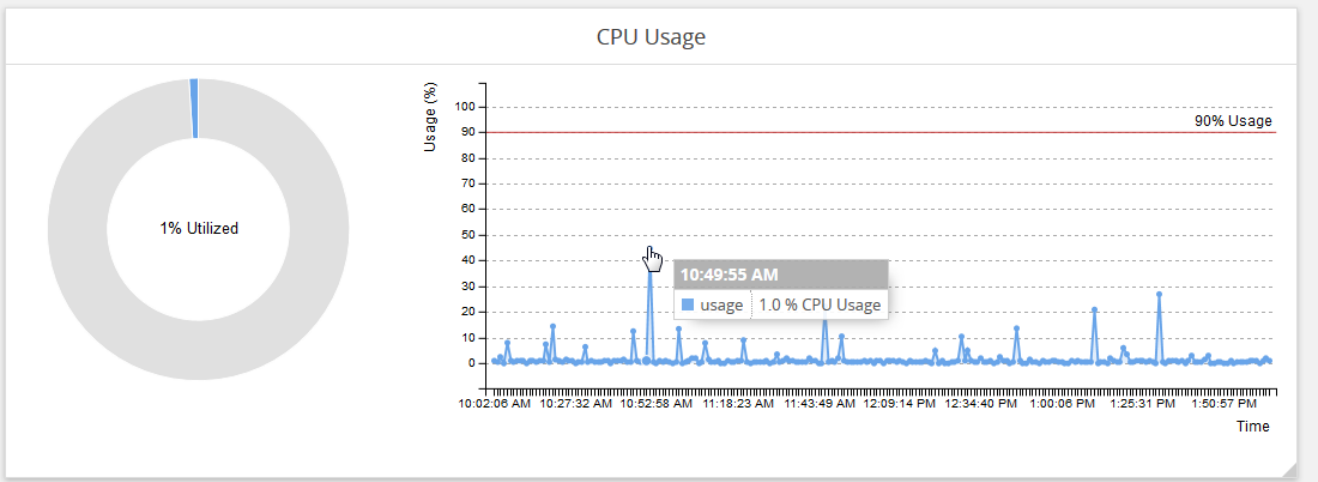

Figure 49 *CPU Utilization Monitor*

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.

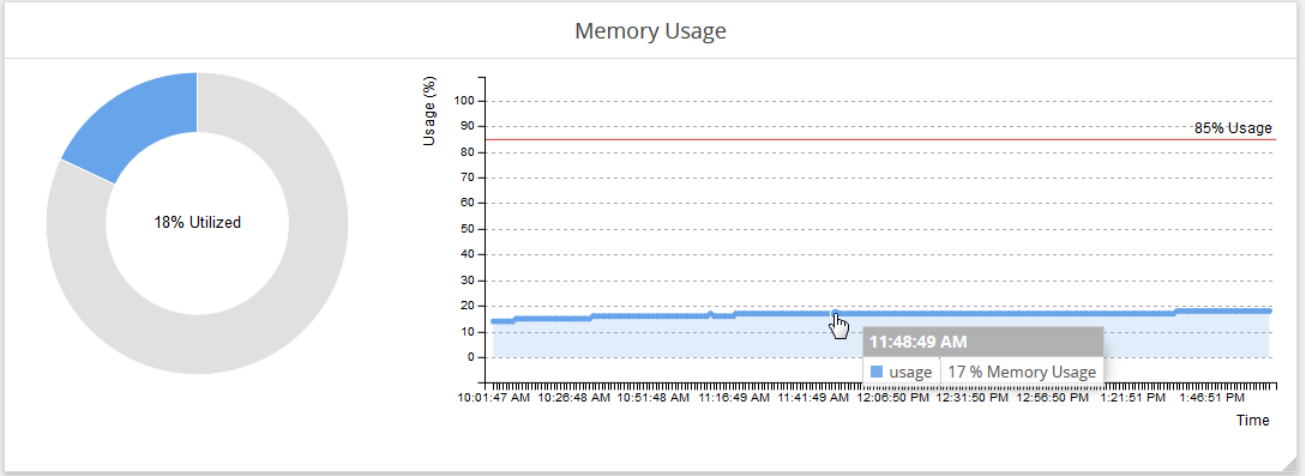

Figure 50 *Memory Utilization Monitor*

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%.

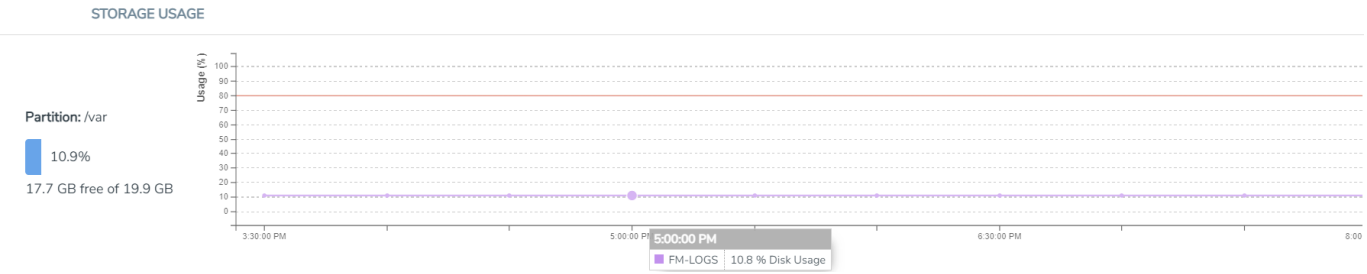

Figure 51 *Storage Utilization Monitor*

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**.

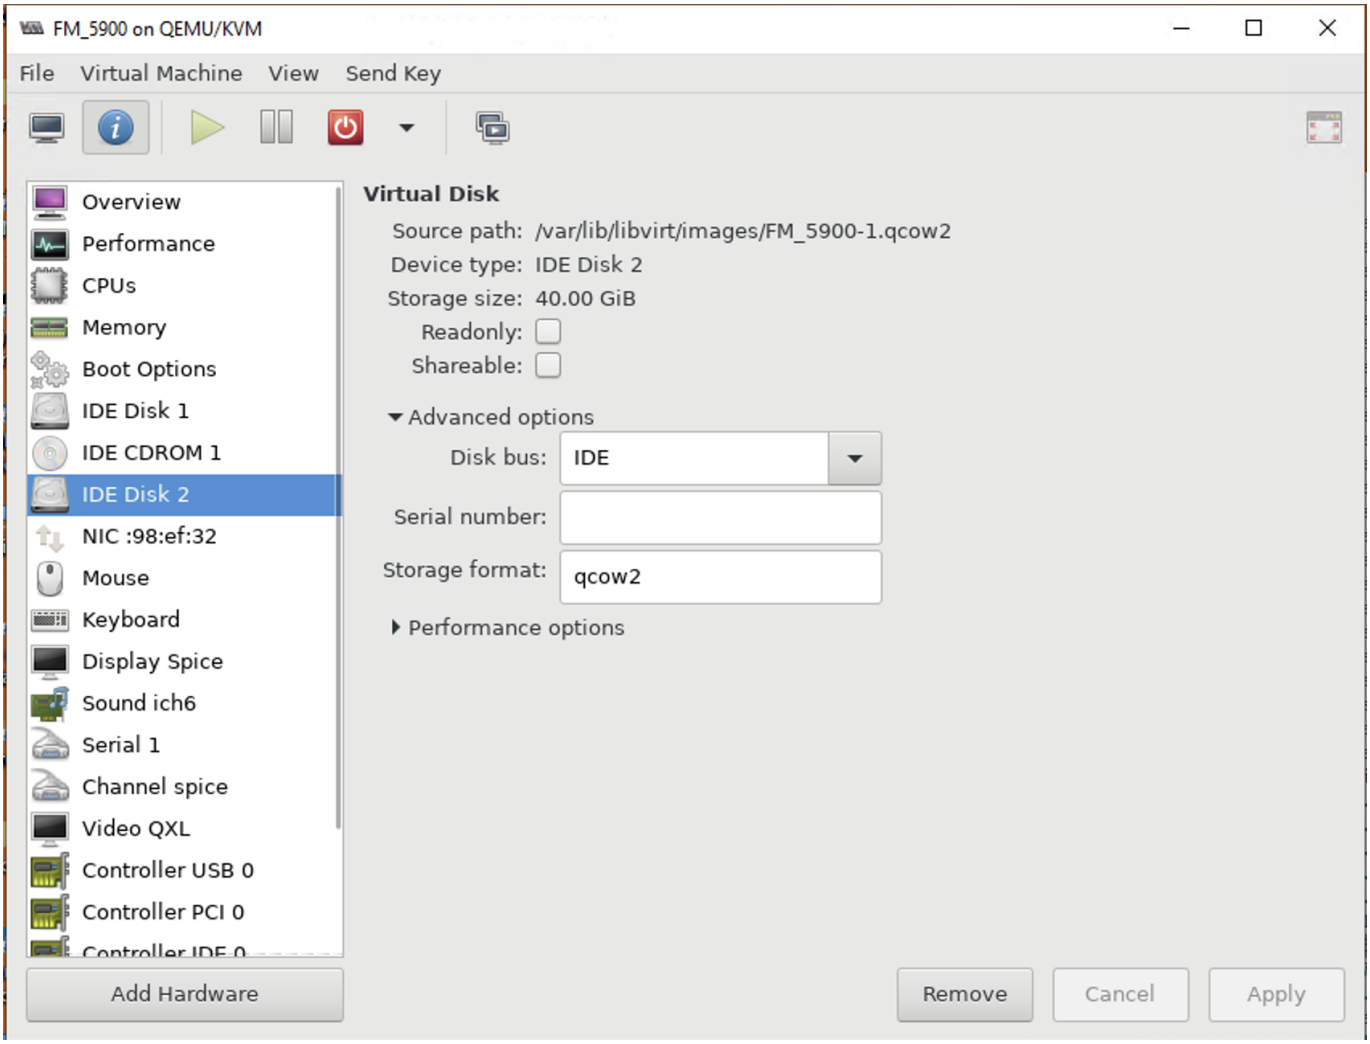

3. Open the console, and run the following command to increase the disk size:

   /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 / GigaVUE-OS/ Cloud Fabrics	Allows GigaVUE-FM to accept Management connection using REST API from user and Gigamon devices. 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 Seriesnodes to send Application Intelligence.
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.
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.
TCP	9600	StepCA	Cloud Fabric Nodes	Allows GigaVUE-FM to reach the NTP server for time synchronization.

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

Direction	Protocol	Port Number	Service	Source CIDR	Purpose
Bidirectional	TCP	443	HTTPS	GigaVUE-FM	REST API communication between HA members.
Bidirectional	TCP	8300	Consul	GigaVUE-FM	RPC communication between Consul members.
Bidirectional	TCP	8301	Consul	GigaVUE-FM	Heartbeat and Gossip between Consul members.
Bidirectional	UDP	8301	Consul	GigaVUE-FM	Heartbeat and Gossip between Consul members.
Bidirectional	TCP	8302	Consul	GigaVUE-FM	Heartbeat and Gossip between Consul members over WAN.
Bidirectional	UDP	8302	Consul	GigaVUE-FM	Heartbeat and Gossip between Consul members over WAN.
Bidirectional	TCP	27071	MongoDB	GigaVUE-FM	Used for data replication across HA members and data access through GigaVUE-FM CLI.
Bidirectional	TCP	9300	OpenSearch	GigaVUE-FM	Used for data replication across cluster members.
Bidirectional	TCP	30865	CSync2	GigaVUE-FM	Used for the Synchronization of files / directories across HA members. For example, Image files during GigaVUE-FM HA Upgrade.
Bidirectional	TCP	24224	FluentD	GigaVUE-FM	Used for receiving / forwarding the packets from / to other HA members.
Bidirectional	UDP	24224	FluentD	GigaVUE-FM	Used for receiving / forwarding the packets from / to other HA members.

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.
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	
AnyCloud	GigaVUE Cloud Suite for AnyCloud Guide
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 (Bits/s), Tx Port Statistics (Bits/s), Tool GigaStream Maximum Rx Rate (pps) 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.

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.

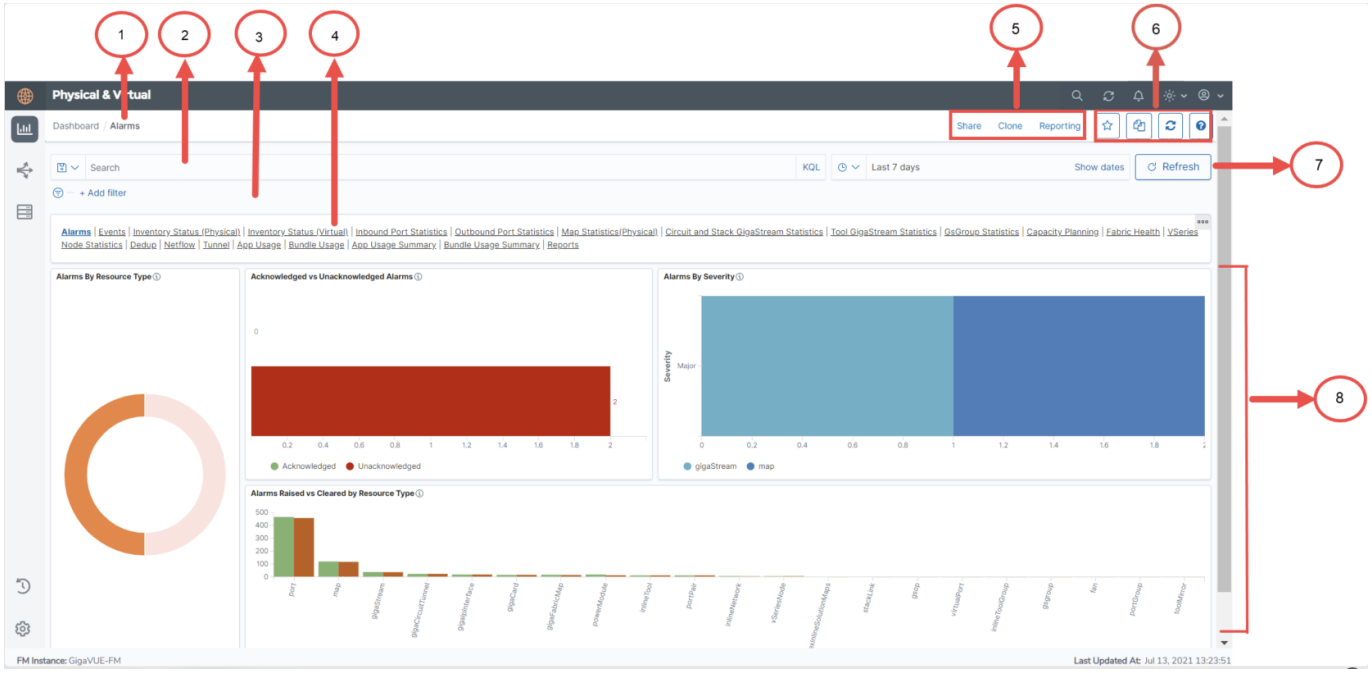

Refer to the following table for details:

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.

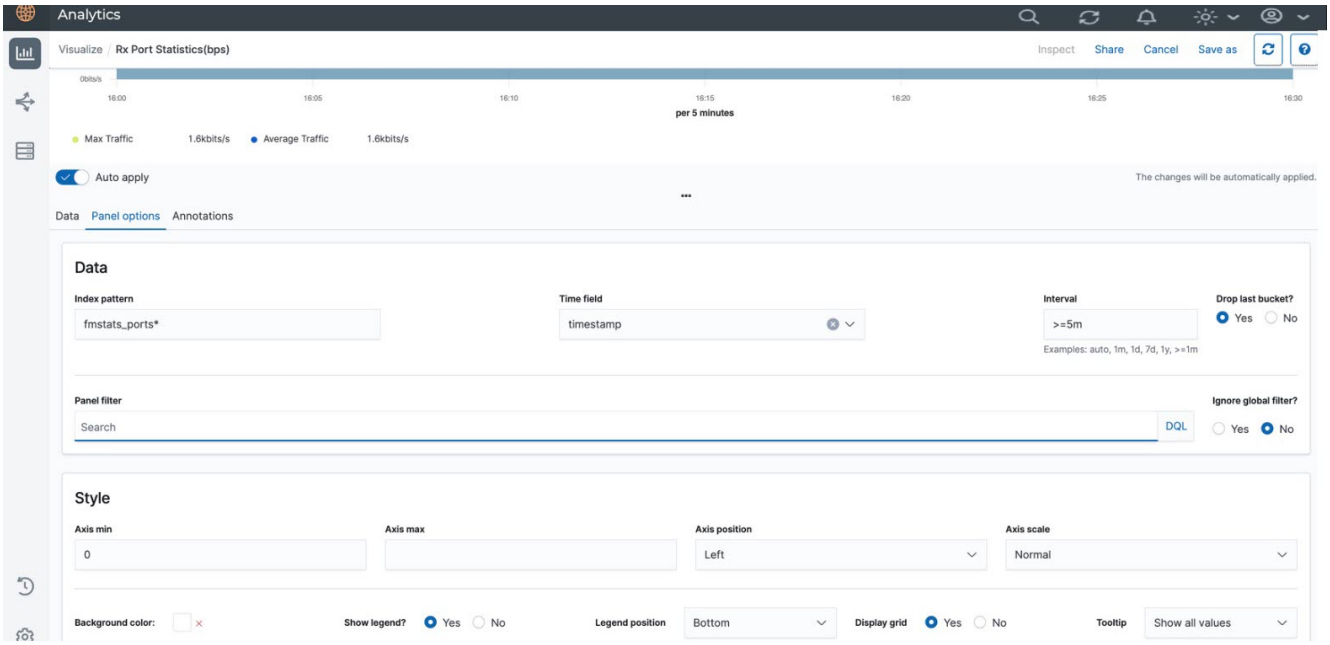

Work with the Analytics User Interface

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.

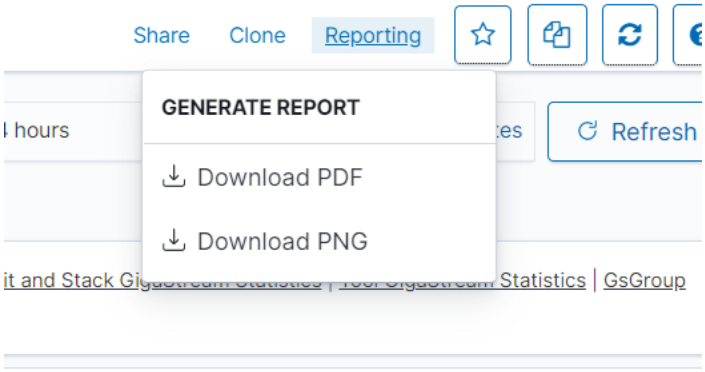

Set as Default

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.

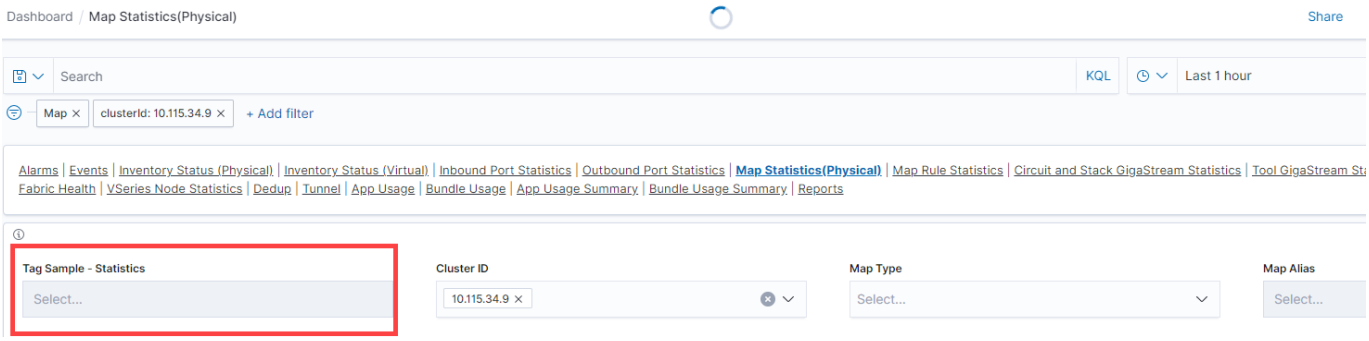

Pre-requisites

- 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**.

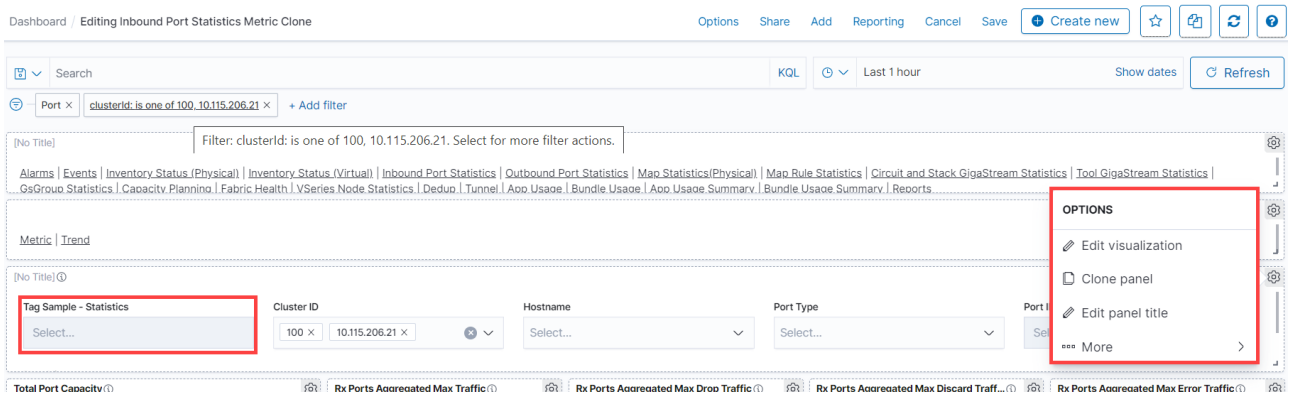

4. In the **Control** tab, configure the following. This is for filtering the data based on the tag value SITE.

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

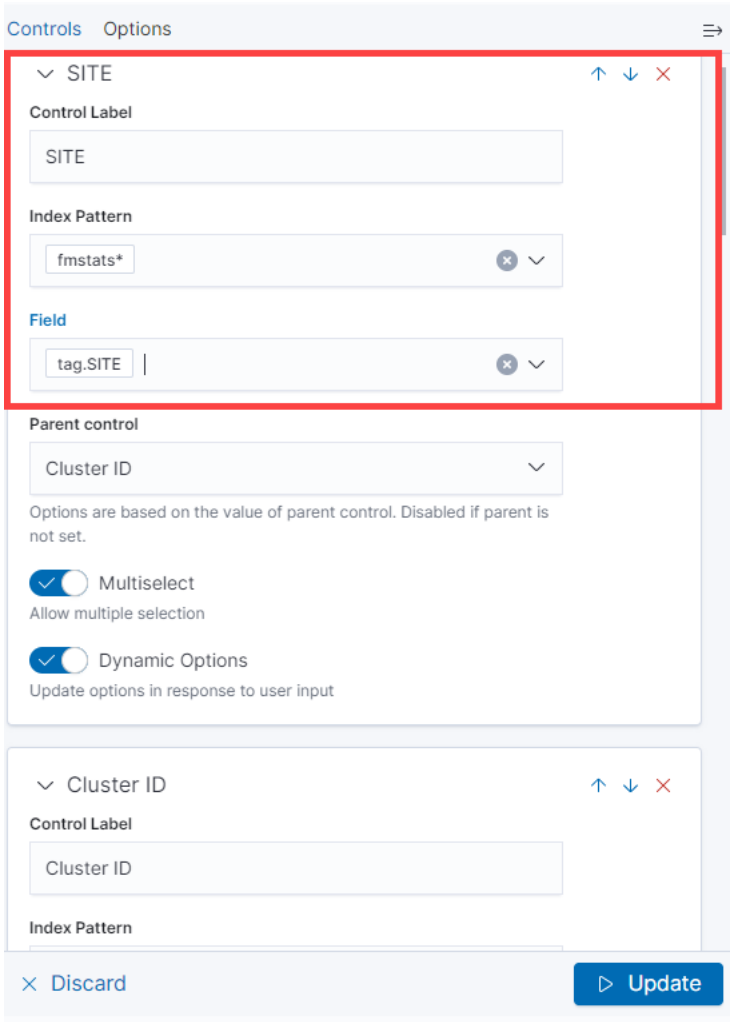

5. Configure the existing tags as required:

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.

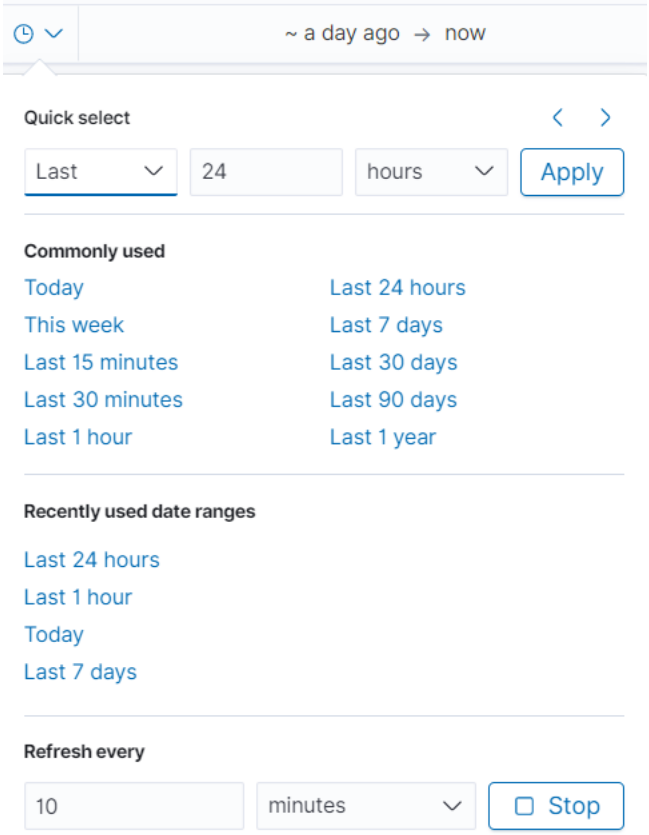

  Use the CHANGE ALL FILTERS to configure the following options:

  - **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

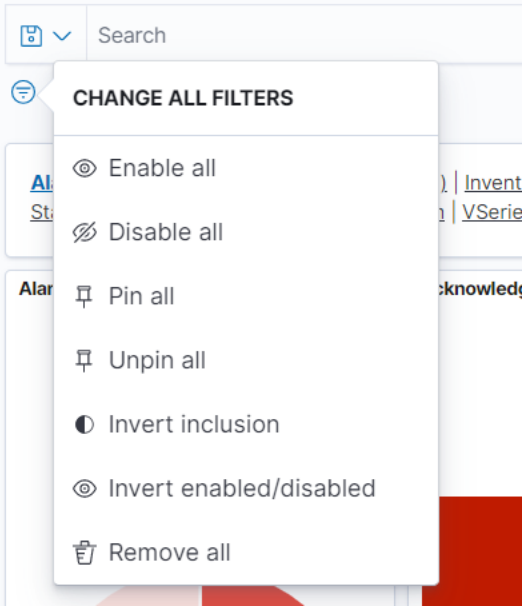

Filtering in visualizations:

- 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.

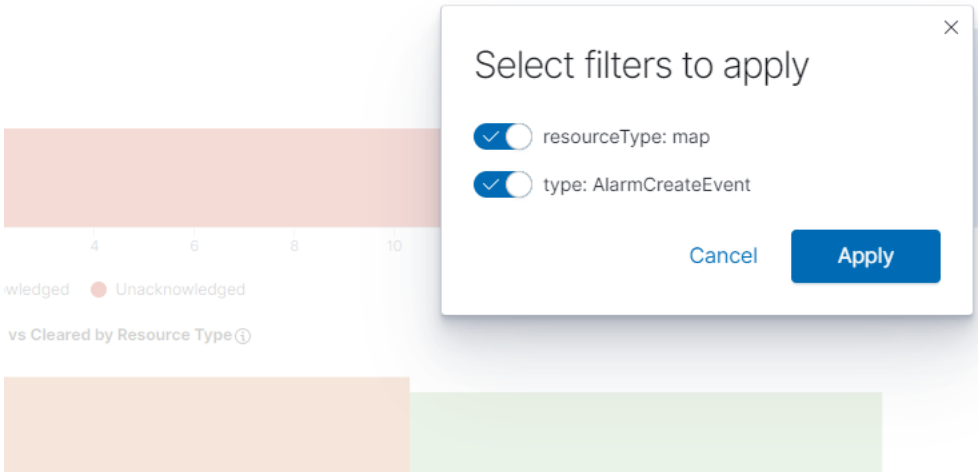

- **Control Visualizations:** Use Control Visualizations to filter the data based on tags. For example, in the Inbound Port Statistics visualization, you can filter the data based on the cluster id, port number, port id or port alias. Refer to the following sections:
  - Auto Refresh Tags in Visualizations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

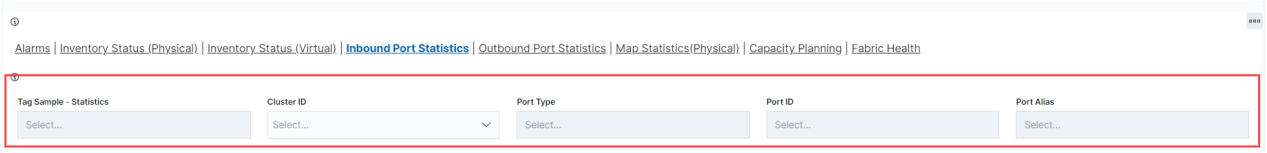

.

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.

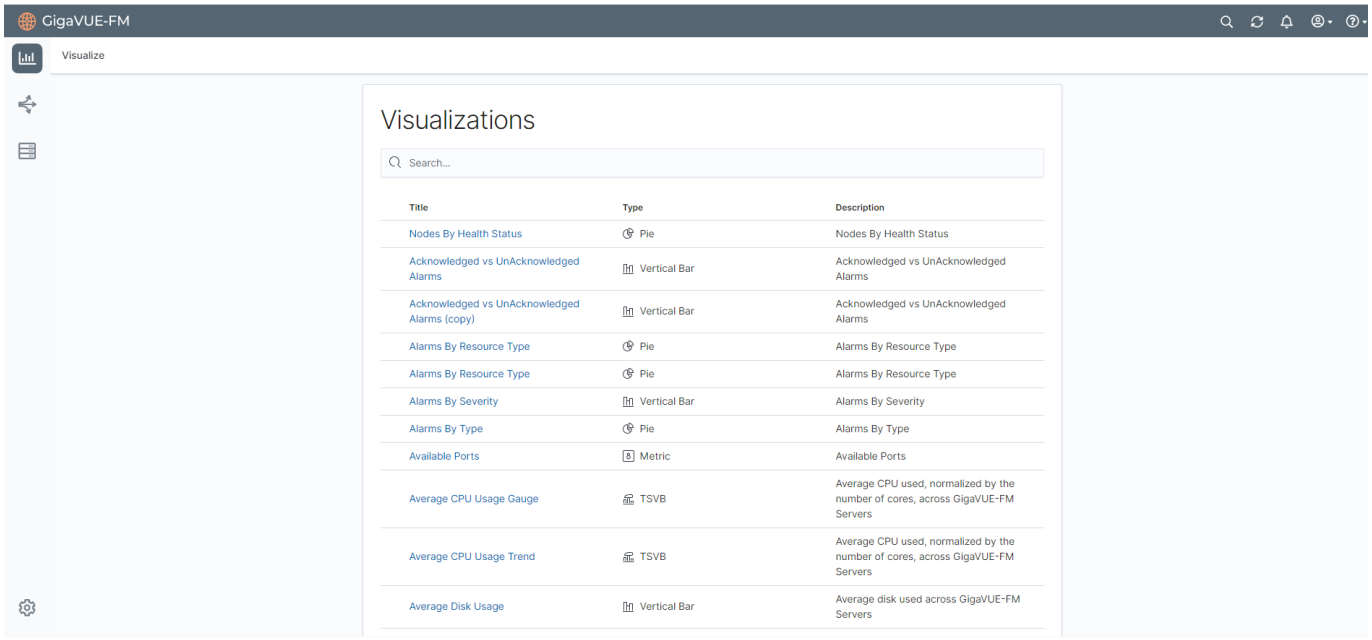

Create Custom Visualizations

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**.

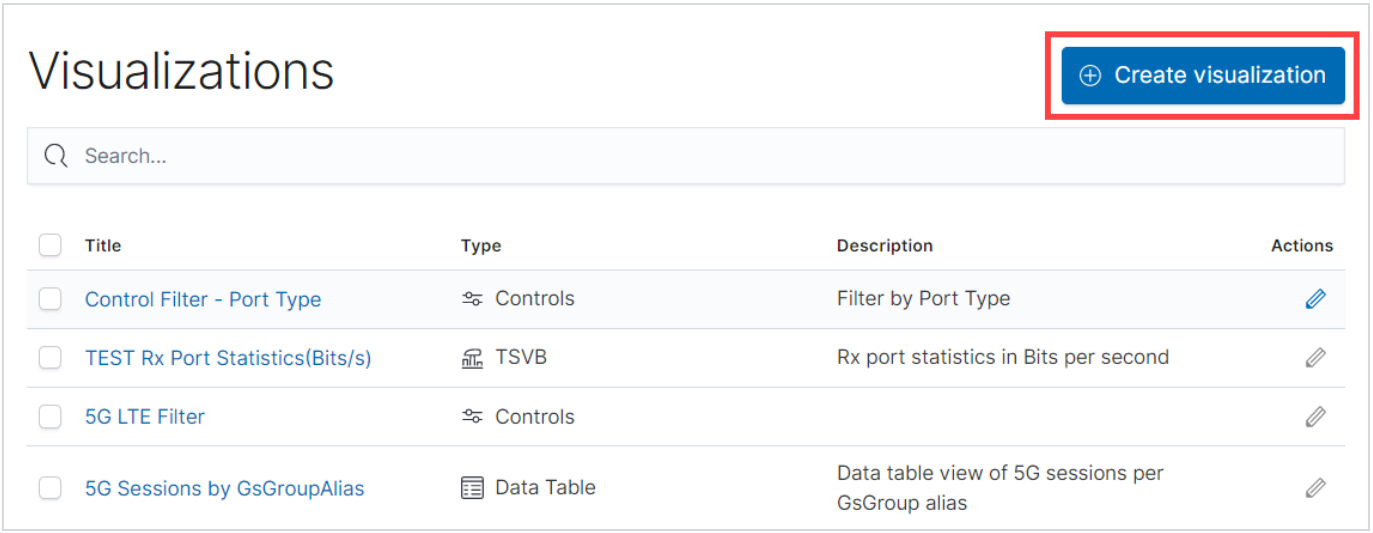

3. In the New Visualization page, select the required **Visualization Type.**

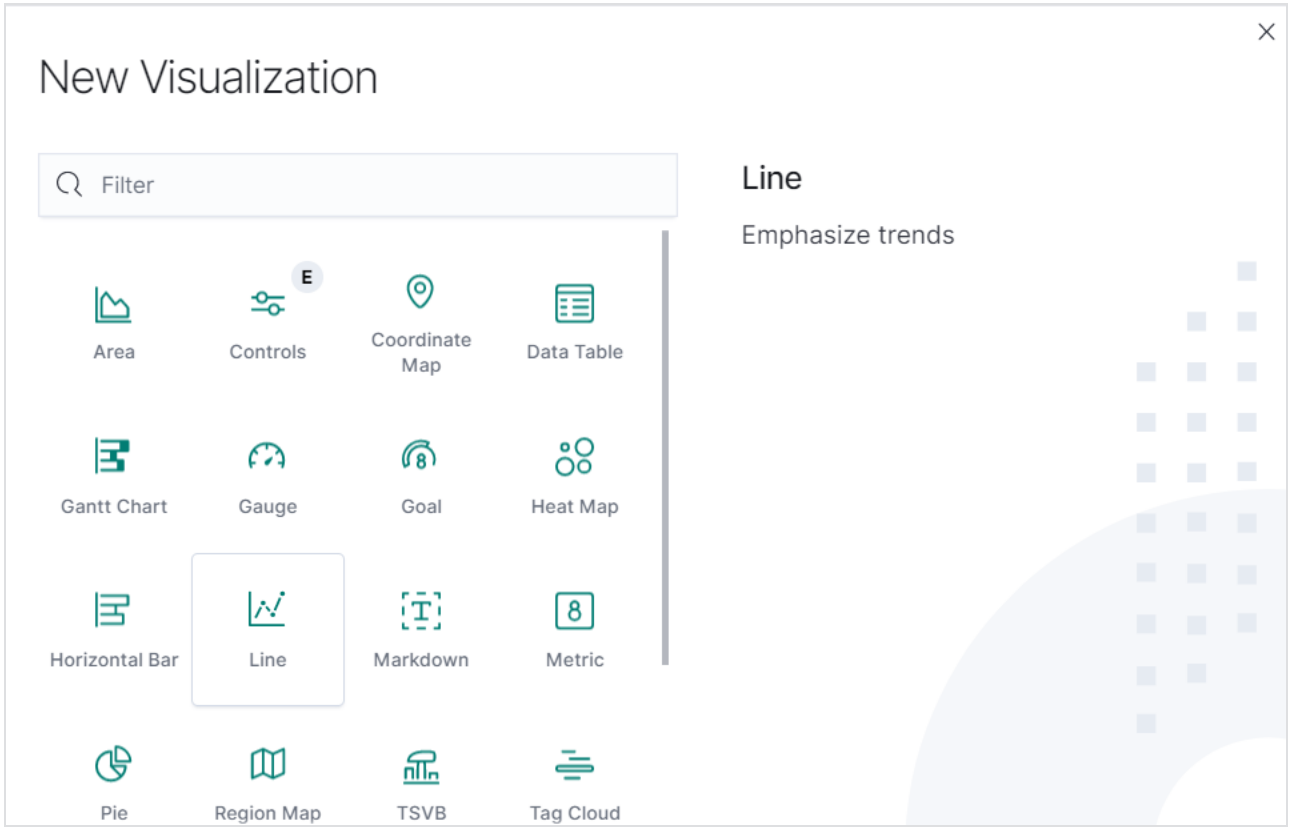

4. Select the data source for the visualization.

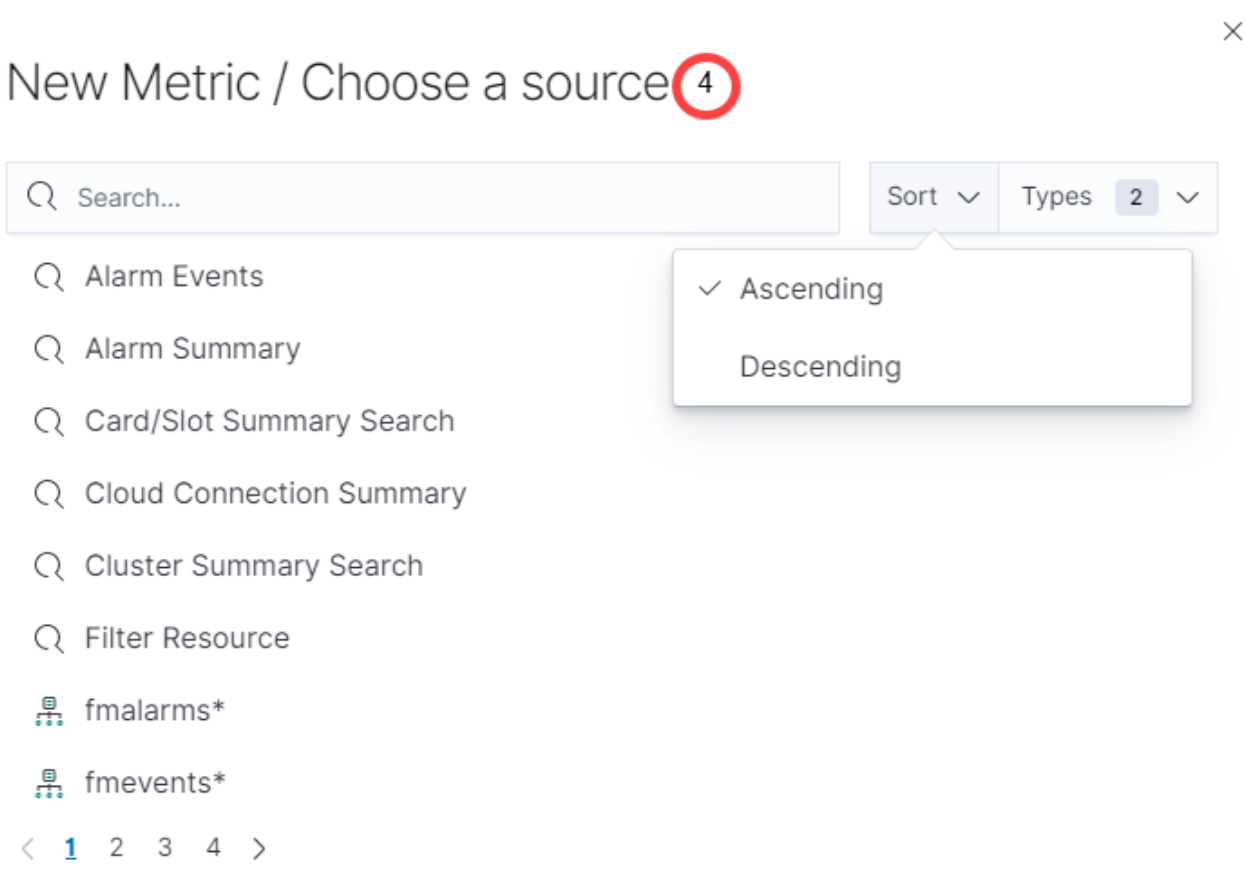

5. Enter the required details for the type of visualization selected. Refer to the table below for more details.
6. Click **Save**.
7. In the **Save Visualization** dialog enter the Title and Description for the visualization and click **Save**.

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.**

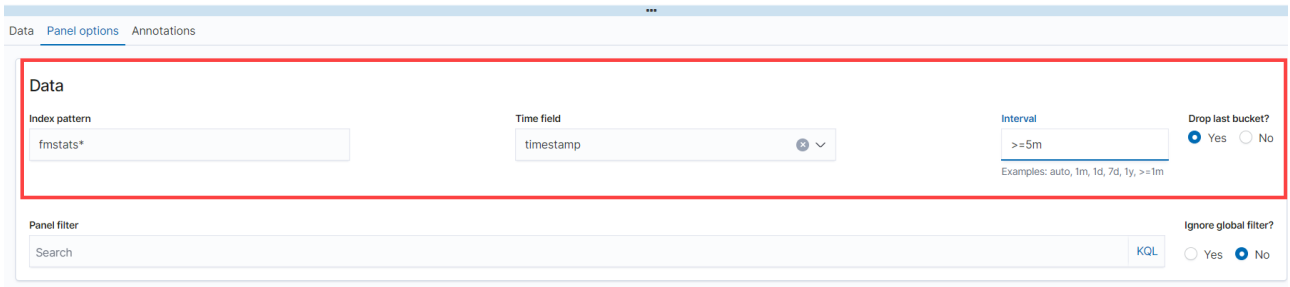

3. Select or enter the following details under Data:

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

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

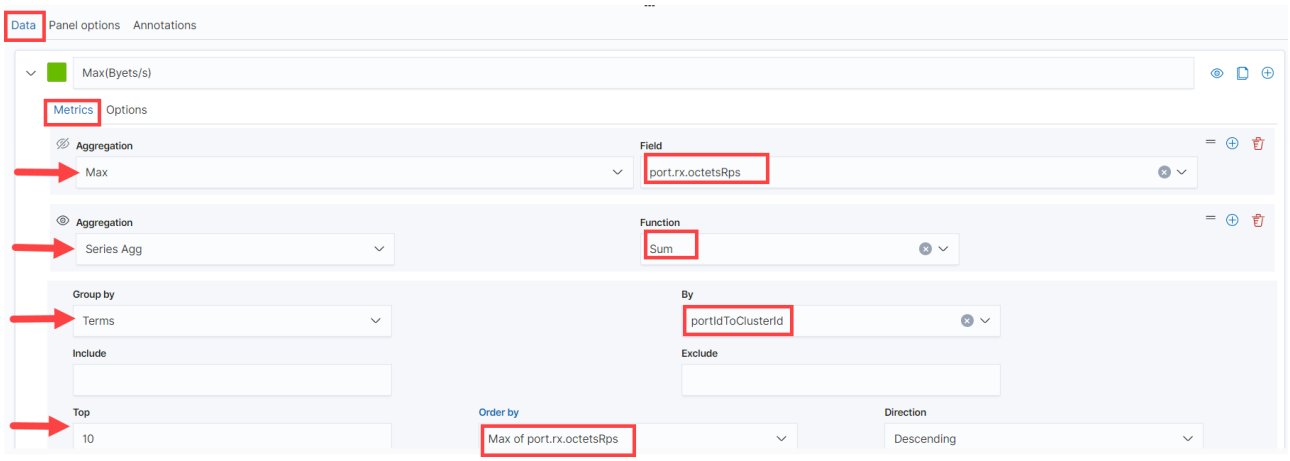

Select or enter the following details:

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.**

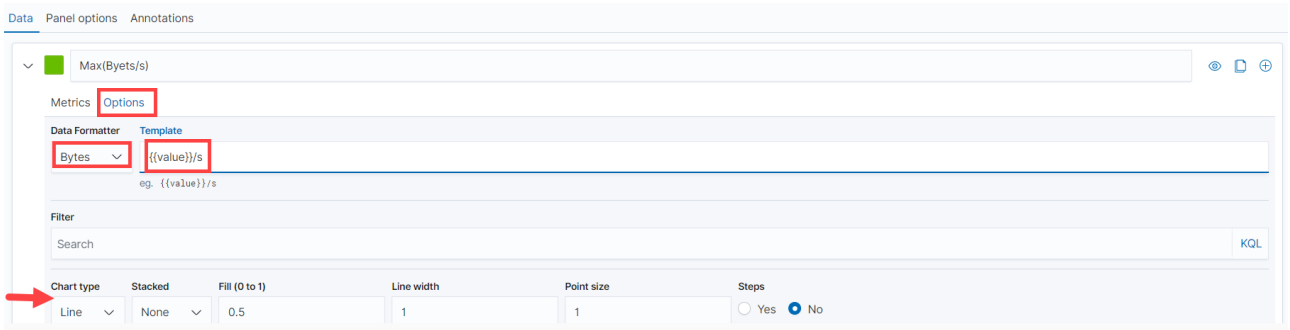

Data Formatter	*Bytes*
Template	*Values/s*

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.

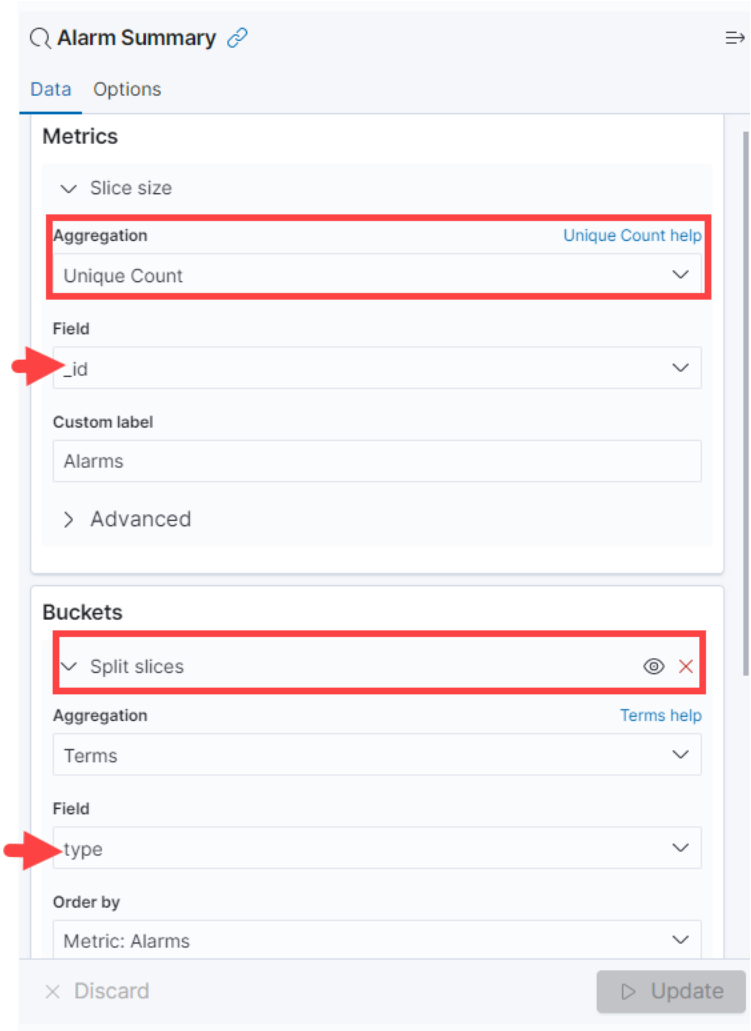

- **Aggregation function**: Select Unique Count. This returns the number of unique values in a field.
- **Field values**: _Id and Type can be changed as per your requirements
  - **Field value _Id**: The field that you want to visualize (in the example, id is being used because it is unique for each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

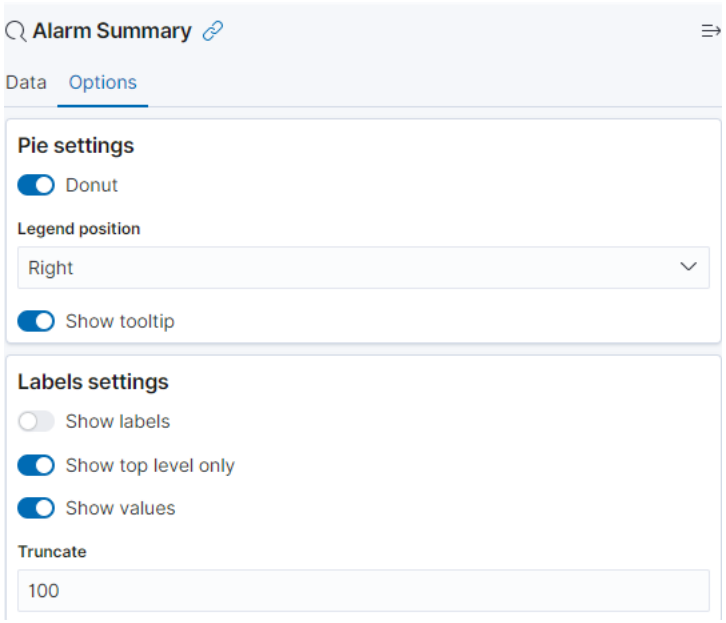

Metric Displaying Card Count

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:

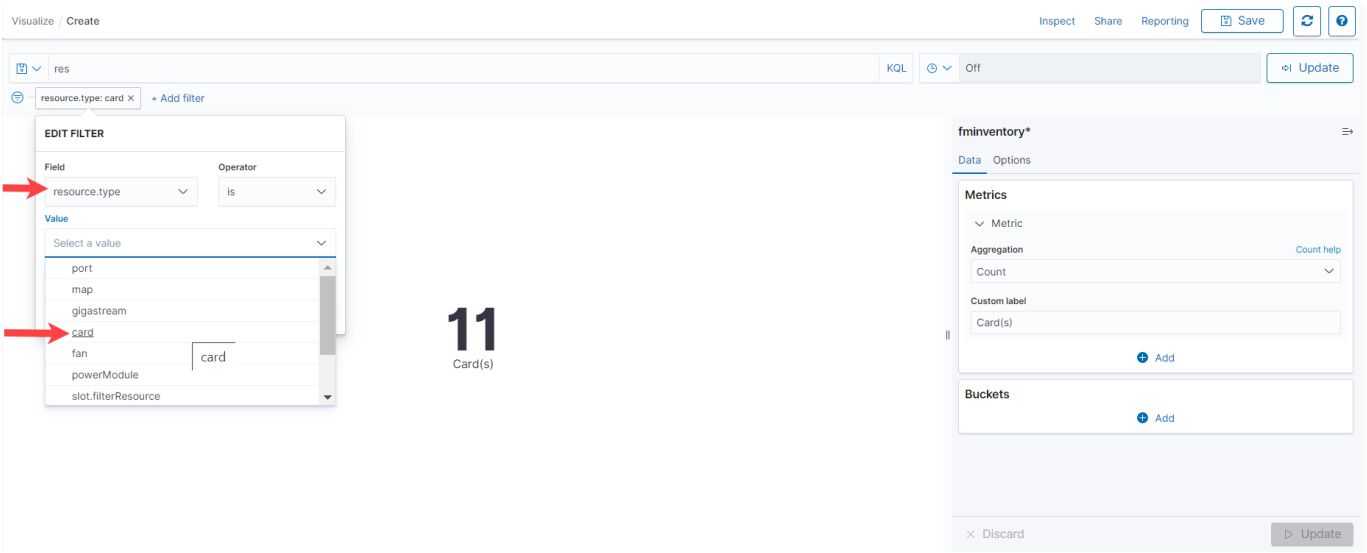

NOTE: Use the **Options** tab to configure the required visual effects such as adjusting the font size.

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*

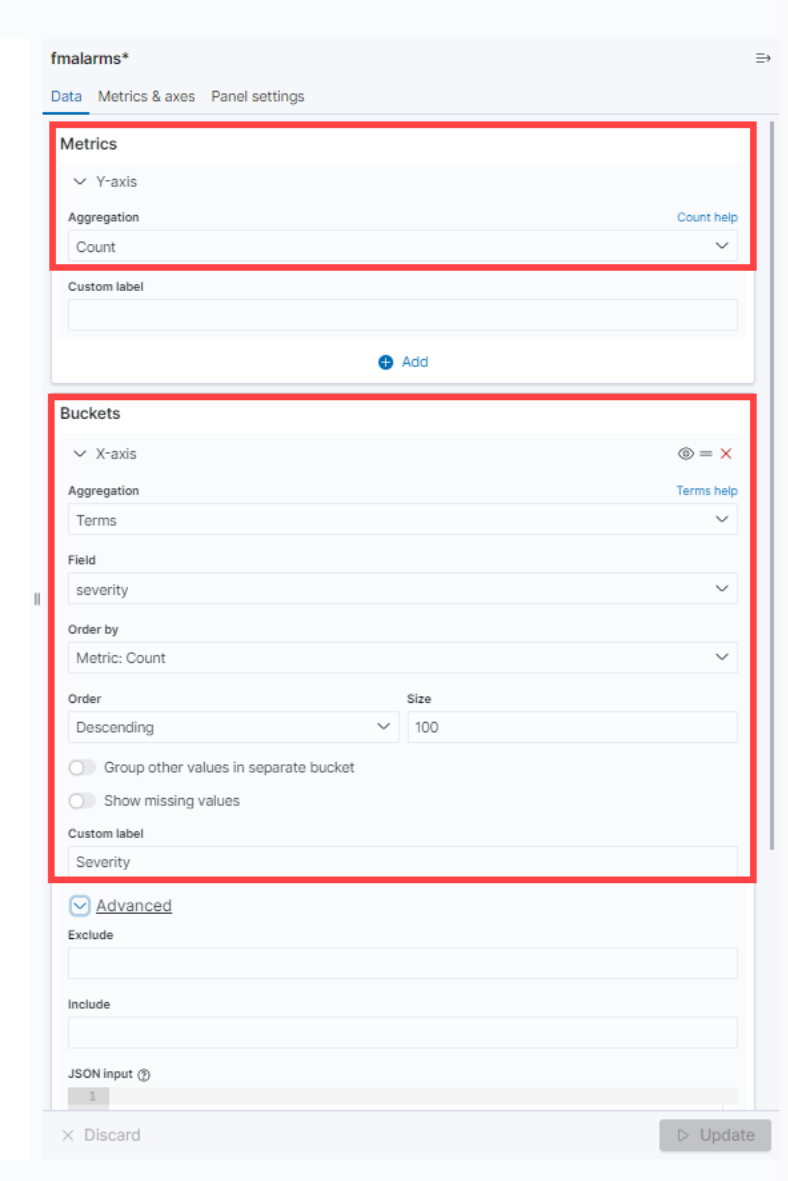

4. Under **Advanced**, select **Split series** and configure the following:
   - *Sub aggregation: Terms*
   - *Field: resourceType*
   - *Order by: Metric Count*

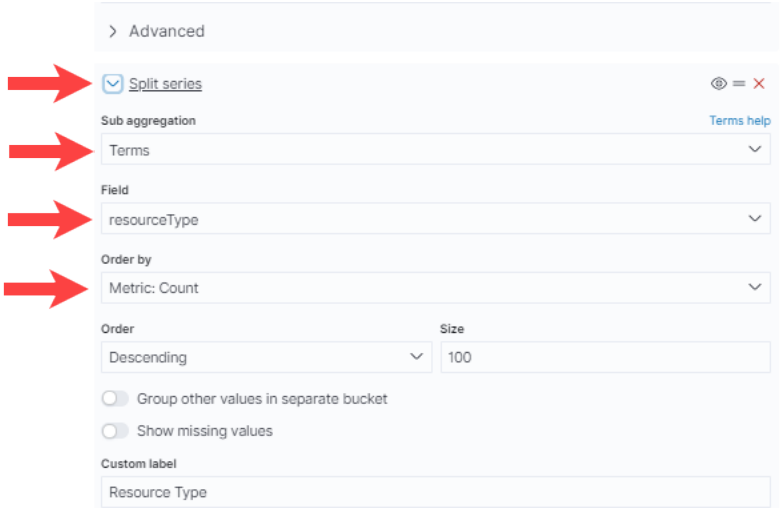

5. Under Metrics & Axes, configure the following:

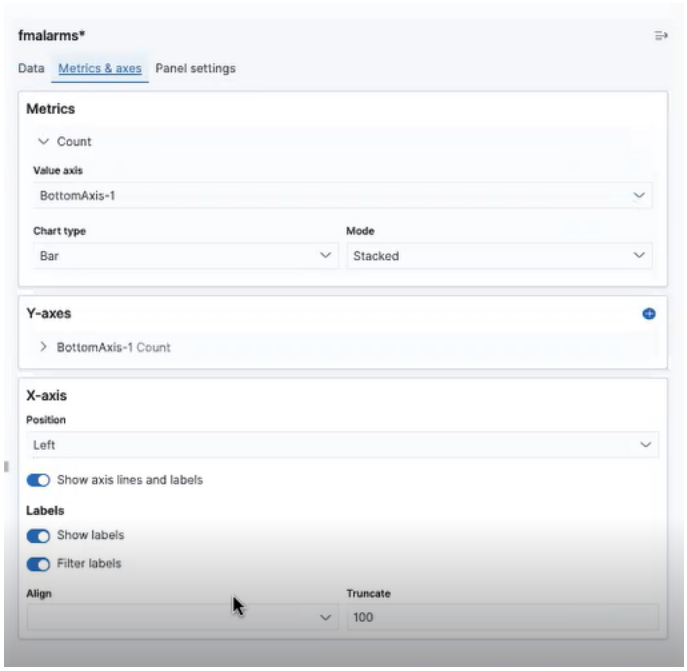

6. Click **Update**.

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.

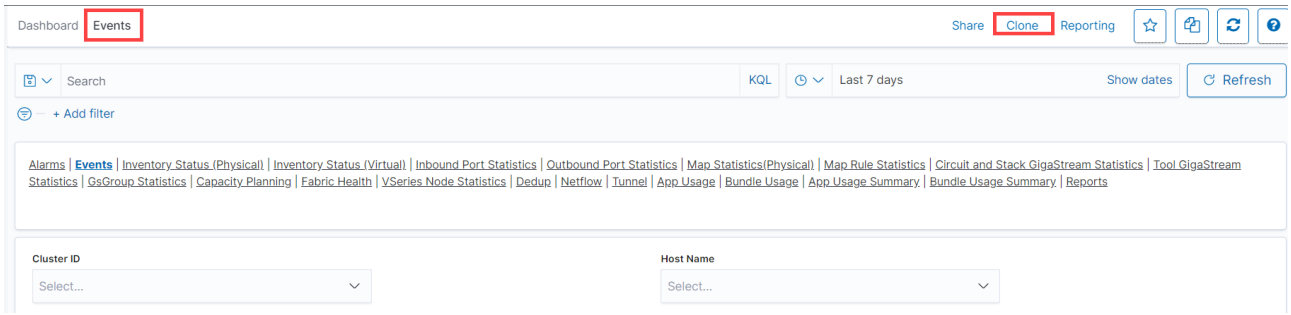

Refer to the Edit Dashboard section for details on editing the dashboard.

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**.

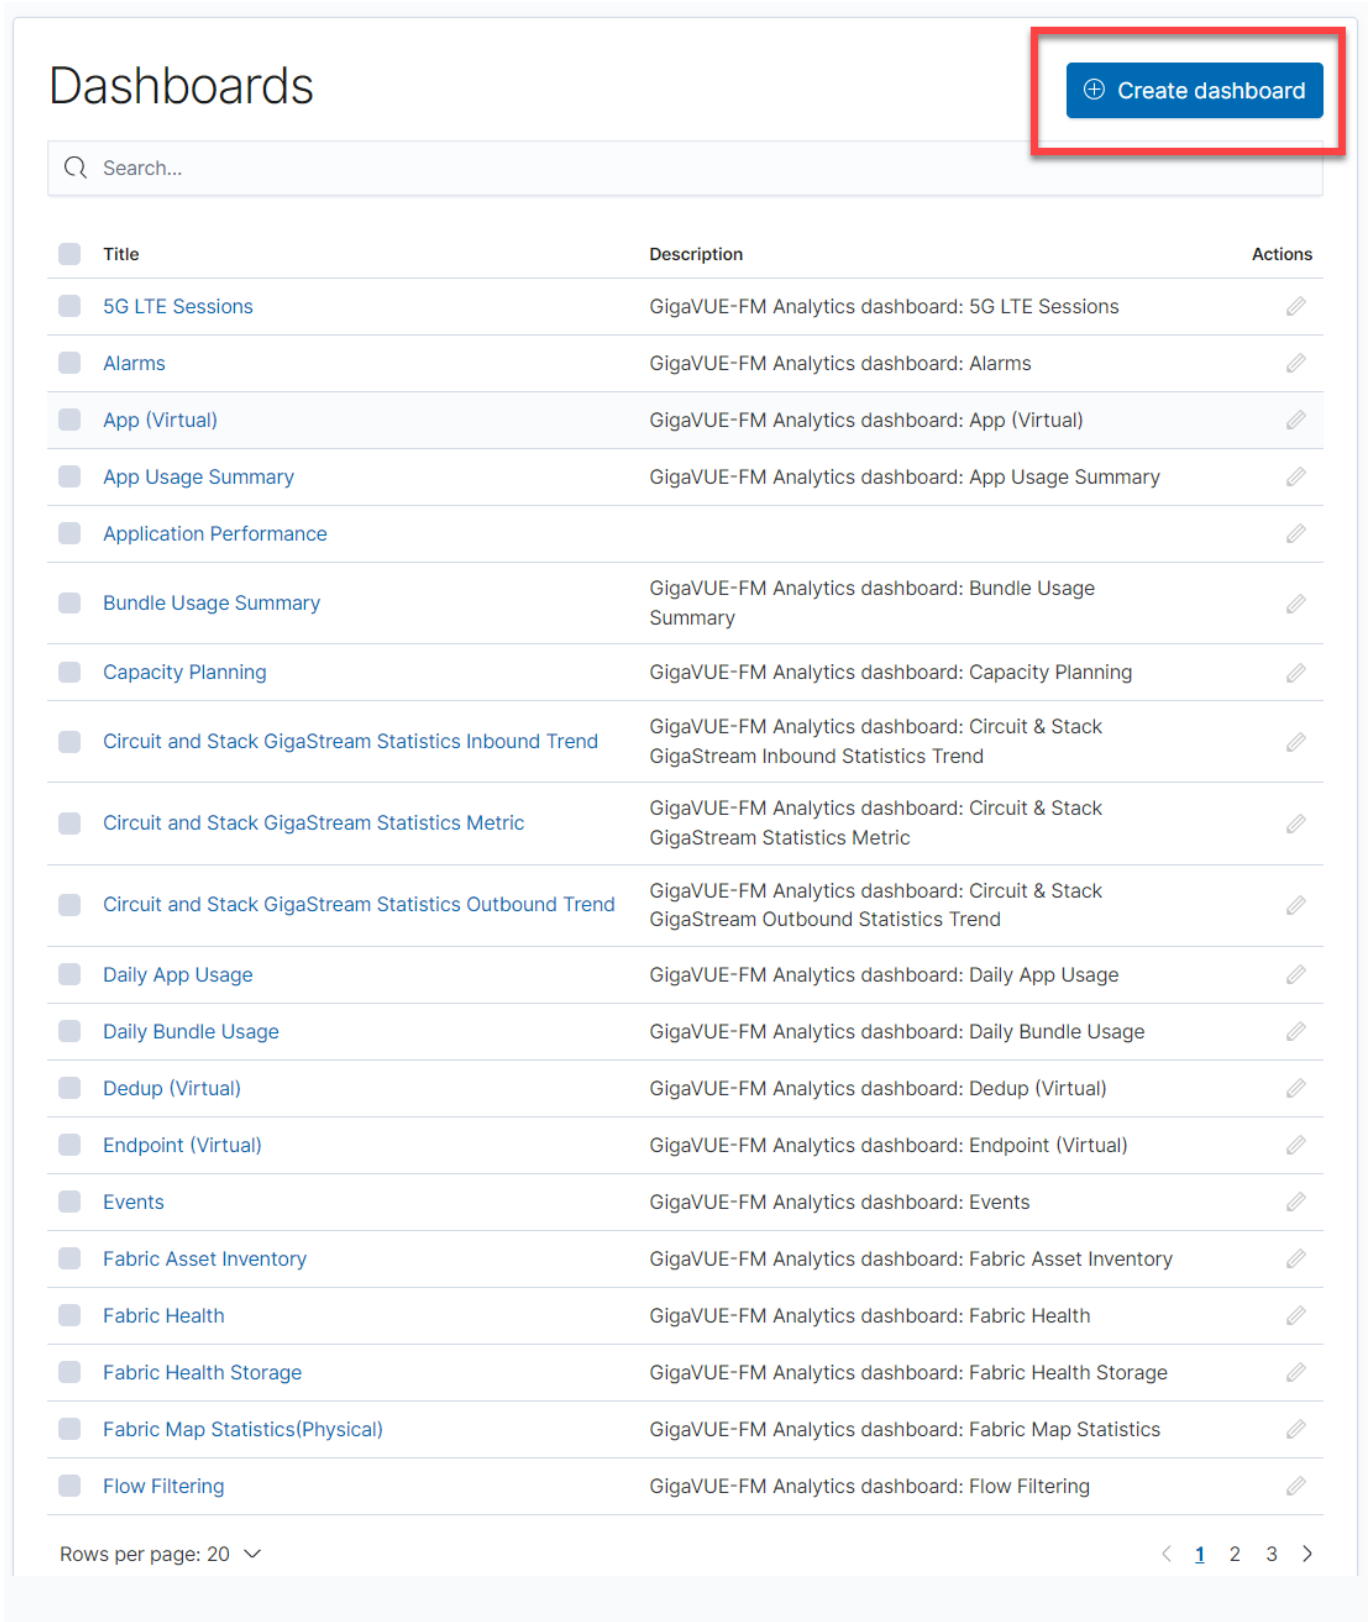
Dashboards
Create dashboard
Search...
Title
Description
Actions
5G LTE Sessions
GigaVUE-FM Analytics dashboard: 5G LTE Sessions
Alarms
GigaVUE-FM Analytics dashboard: Alarms
App (Virtual)
GigaVUE-FM Analytics dashboard: App (Virtual)
App Usage Summary
GigaVUE-FM Analytics dashboard: App Usage Summary
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

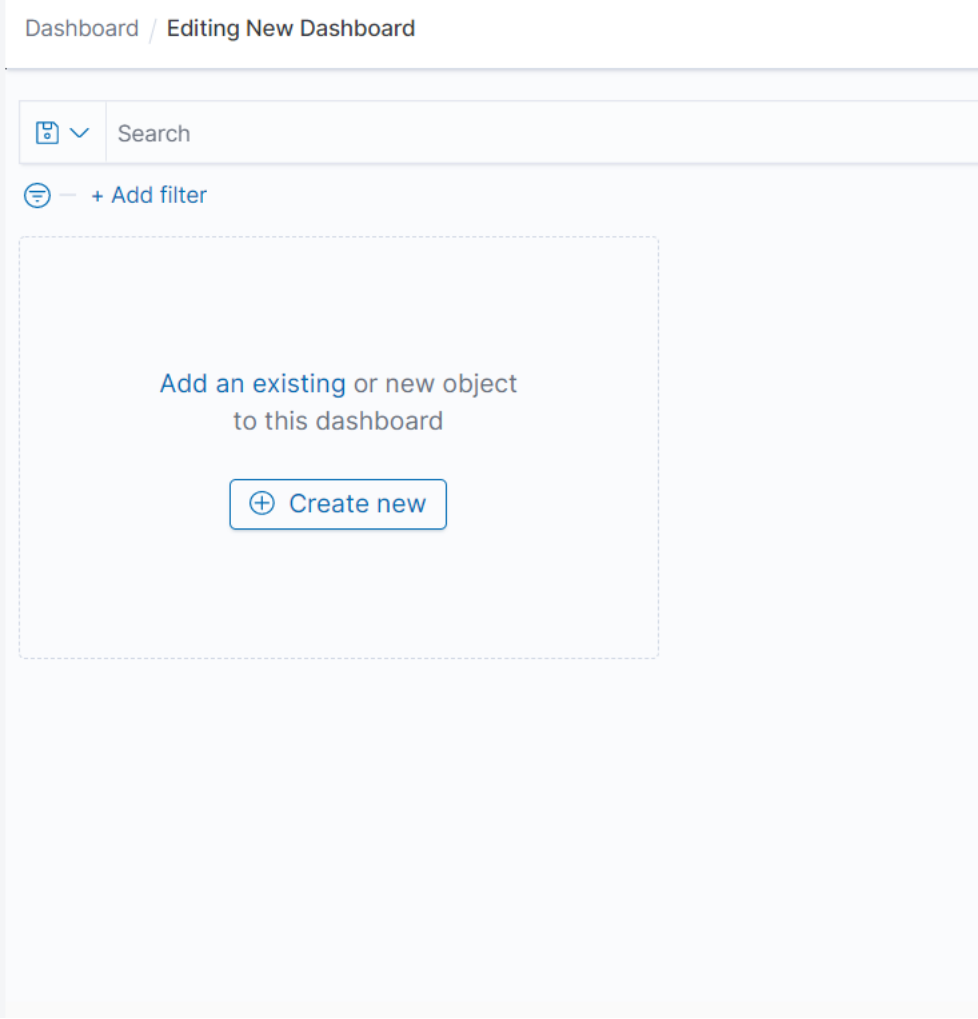

5. Click **Save**.
6. In the **Save dashboard** dialog box, enter the **Title** and **Description** for the dashboard and click **Save**.

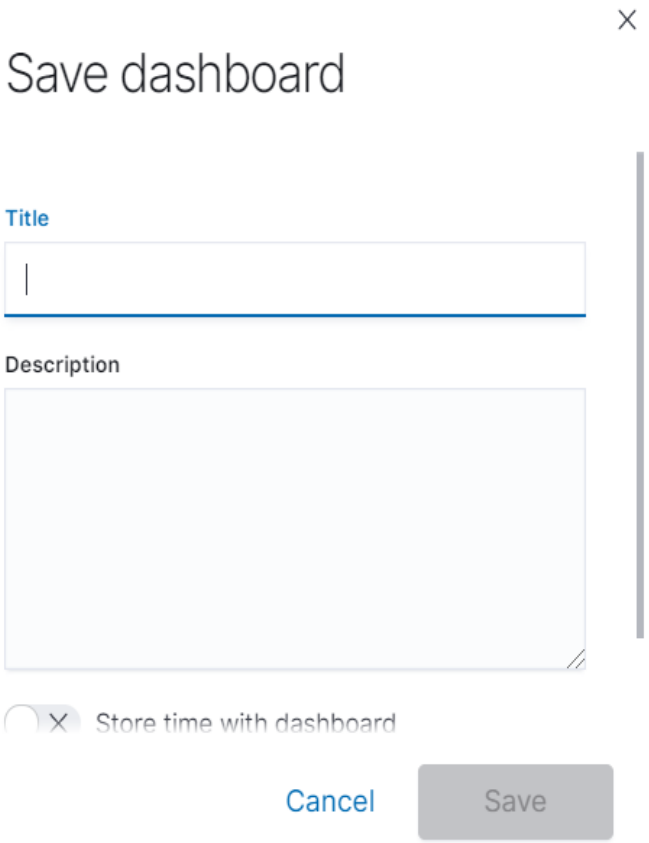

To make further changes to the dashboard, refer to the Edit Dashboard section.

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

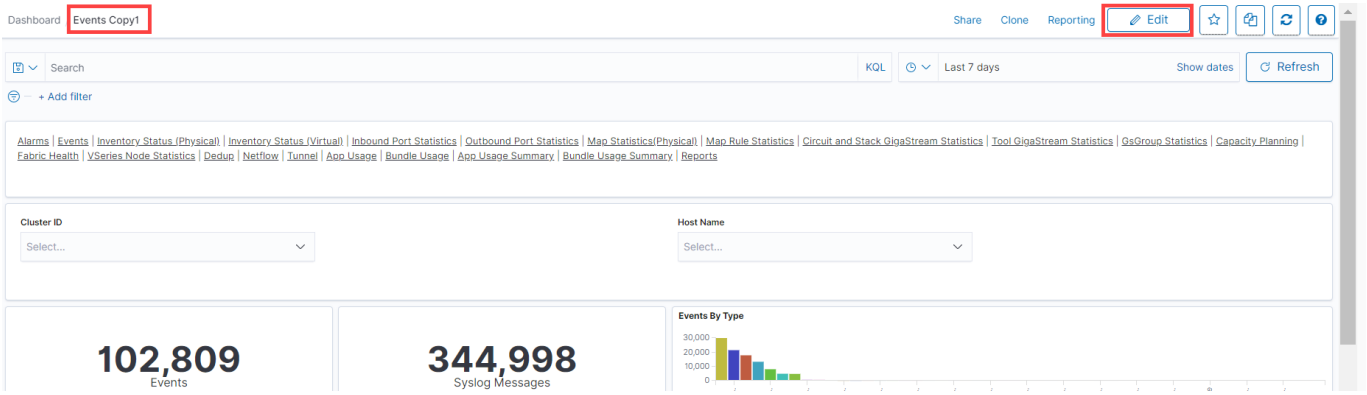

Click **Save** to save the changes to the dashboard. In the **Save Dashboard** dialog box, use the toggle option to save the changes to a new dashboard. Click **Cancel** to discard the changes.

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**.

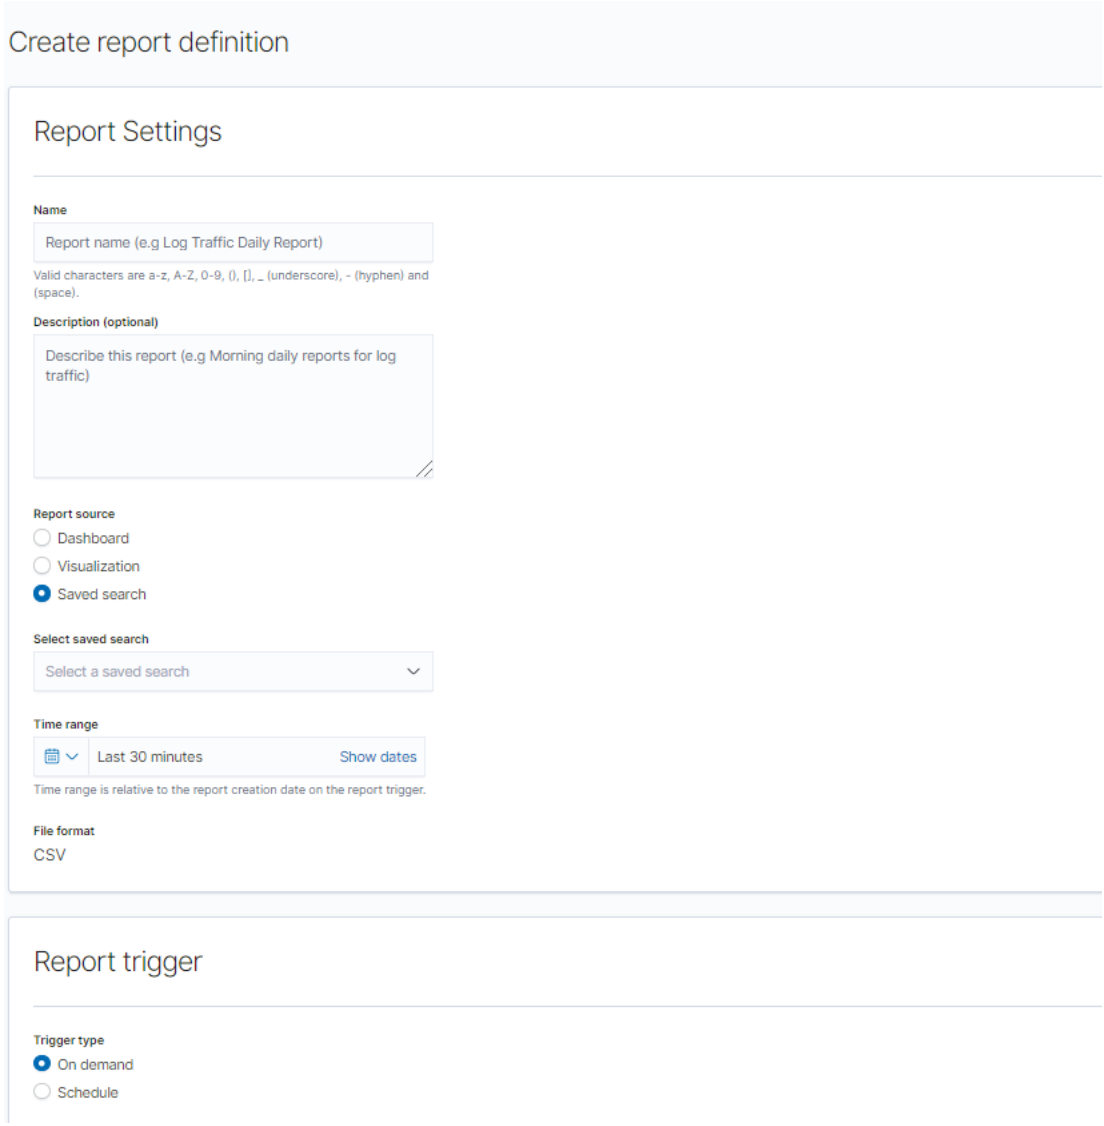

2. Select or enter the following details under **Report Settings**:

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

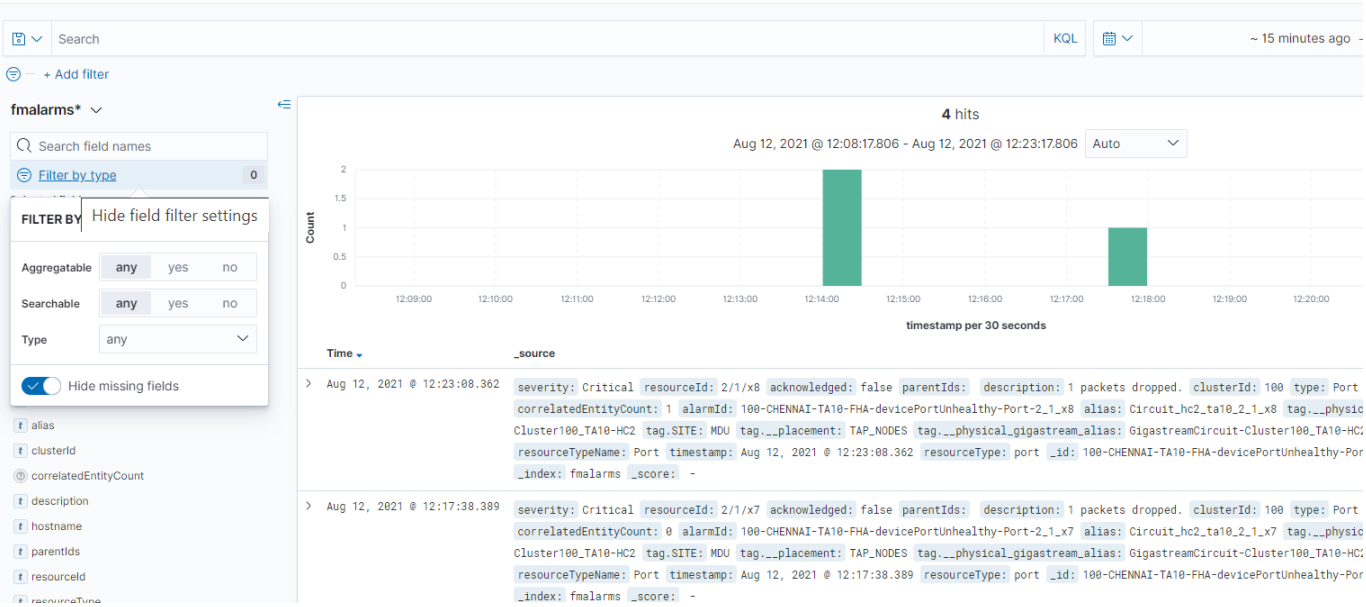

You can perform the following operations from the Discover page:

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.

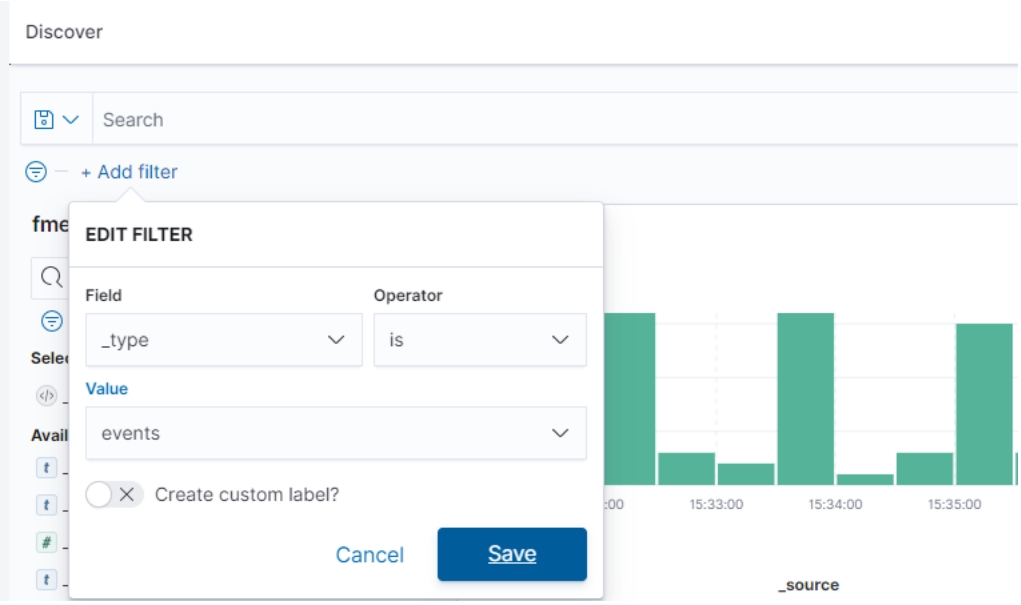

Save Search

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	*Average Disk Usage (/var)*

Dashboard	Details	Visualizations
	Use the Server Control Filter to view the memory and storage metric and trend for the particular server.	*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 (/)
		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	*Top Rx ports by Max Rate (bps)*

Dashboard	Details	Visualizations
	• 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 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 Statistics	Displays statistics of the transmitting (Tx) 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 • 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.	*Top Tx ports by Max 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)*

Dashboard	Details	Visualizations
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: • 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 (bps)* *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)*

Dashboard	Details	Visualizations
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 • Cluster ID • Host Name • GigaStream	*Circuit and Stack GigaStream Tx Traffic Rate (bps)* *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*

Dashboard	Details	Visualizations
		Maximum Rx Rate (bps) by member port Alias
		Circuit and Stack 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*

Dashboard	Details	Visualizations
		packets Discard Rate (pps)
		Tool GigaStream Tx packets Drop Rate (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
Fabric Asset Inventory	Displays details about the devices managed by GigaVUE-FM: • Devices • Cards • Fans	

Dashboard	Details	Visualizations
	• Power • Power Modules • 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	

Dashboard	Details	Visualizations
	data: • Sample Tag • 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	

Dashboard	Details	Visualizations
	data: • Sample Tag • Cluster ID • 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:	

Dashboard	Details	Visualizations
	• Session Overall • Session Engine Overview • Traffic Insights • Session Insights • Session Table • 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*

Dashboard	Details	Visualizations
		• *Average CPU Per Engine* • *Max CPU Per Engine* • *Max CPS Per 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 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 the 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 V Series nodes do not have service cores, therefore the CPU usage is reported as 0.
		V Series Node with Most CPU Usage For Past 5 minutes	Line chart that displays Maximum CPU usage of the V

Dashboard	Displays	*Visualizations*	Displays
			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 the 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 to V Series Node, Grouping by tunnel identifier comprising {monDomain, conn, VSN, tunnelName}, before aggregation.
		V Series Node Tunnel Tx Packets/Errors	TX is for output tunnels from VSN. V Series Node Tunnel Tx Packets/Errors
Dedup	Displays visualizations related to De-duplication application. You can filter the visualizations based on the following control filters: • Platform • Connection	*Dedup Packets Detected/Dedup Packets Overload*	Statistics of the total de-duplicated packets received (ipV4Dup, ipV6Dup and nonIPDup) against the de-duplication application overload.

Dashboard	Displays	*Visualizations*	Displays
	• V Series Node		
		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. The following statistics are displayed for the tunnel: • Received Bytes • Transmitted Bytes • Received Packets • Transmitted Packets	*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
	• 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-Cloud Applications

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

5G-Cloud applications dashboard provides comprehensive insights into the performance and statistics of GVHTTP2 and 5G-Cloud 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:

- OpenStack
- VMware ESXi
- 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 listed in the table below.
3. Select the required 5G-Cloud Monitoring Session name from the drop-down menu. Once selected, the linked GVHTTP2 Monitoring Sessions will be dynamically displayed in the "Gvhttp2 Monitoring Session Name" drop-down menu.
4. Select the applicable GVHTTP2 Monitoring Sessions from the drop-down menu.
5. Select the required V Series Names from the drop-down menu.

The following tables lists the various 5G Apps visualizations:

Table 4: Overall 5G Apps Dashboard

Dashboard	Description	Visualizations	Details
Overall 5G Apps	Displays the overall visualization details of 5G Apps.	**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.

Dashboard	Description	Visualizations	Details
		GVHTTP2 URI's	Displays the cumulative count of the number of HTTP2 requests with Valid and Invalid paths for 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 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.
		5G-Cloud Tool Synthesized Requests & Responses	Displays the cumulative count of Mirrored Transaction HTTP/2 requests and responses synthesized towards the tool port

Dashboard	Description	Visualizations	Details
			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.

Table 5: Detailed 5G Apps Dashboard

Dashboard	Description	Visualizations	Details
Detailed 5G Apps	Displays the detailed visualizations details of 5G Apps.	**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 transaction are not applicable.
		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 a HTTP2 streams are timed out on the GVHTTP2 server.

Dashboard	Description	Visualizations	Details
		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.
		5G-Cloud Rx Packets	Displays the cumulative count of the total received packets from

Dashboard	Description	Visualizations	Details
			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.

Dashboard	Description	Visualizations	Details
		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.

Table 6: NokiaSCP Statistics

Dashboard	Description	Visualizations	Details
NokiaSCP Statistics	Displays the visualization details of Nokia SCP statistics	**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.
		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.

Table 7: OracleSCP Statistics

Dashboard	Description	Visualizations	Details
OracleSCP Statistics	Displays the visualization details of Oracle SCP statistics	**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.

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:

Table 8: Default Mobility Dashboards.

Dashboard	Description	*Visualizations*	Details
5G LTE Sessions	Displays visualizations for 5G and LTE Session statistics	***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

Dashboard	Description	*Visualizations*	Details
			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	Displays visualizations related to flow filtering statistics	***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	Displays visualizations related to flow filtering interface statistics	***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	Displays Max aggregation of GTPC packets for all the interface types in all GigaSMART groups

Dashboard	Description	*Visualizations*	Details
		Statistics - GTPC Packets	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.
Flow Filtering Correlation Statistics	Displays visualizations related to flow filtering correlation statistics.	***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	Displays visualizations related to flow filtering session statistics.	***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.

Refer to the following screenshot

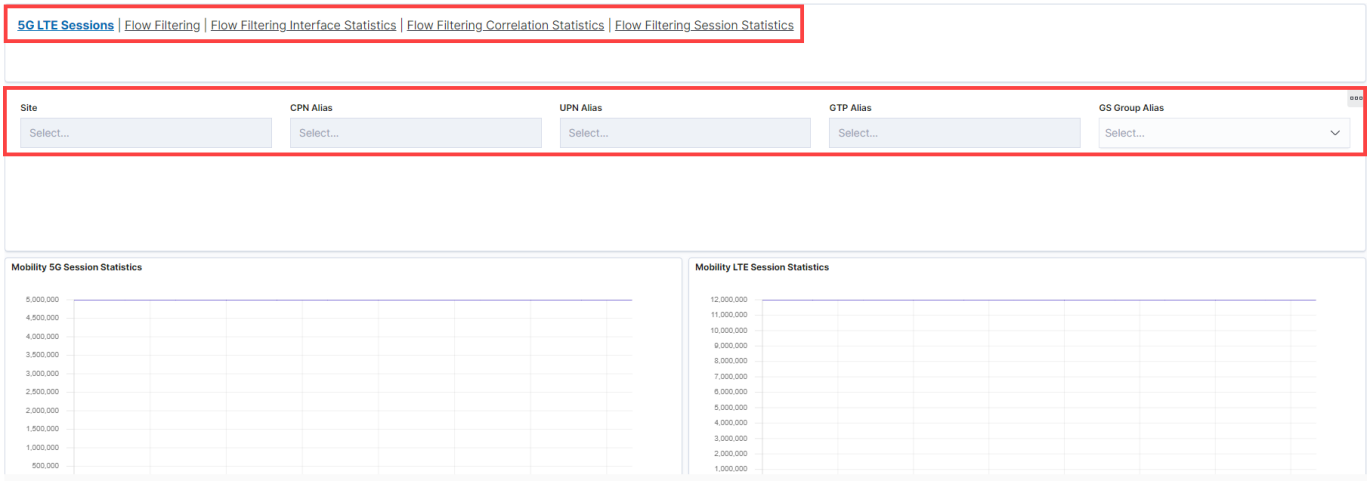

Rules and Notes

- 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 are the basic dashboards and its visualizations:

Table 9: Session Overall Dashboard

Dashboard	Description	Visualizations	Details
Session Overall	Displays visualizations on the overall details of encrypted traffic.	***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.

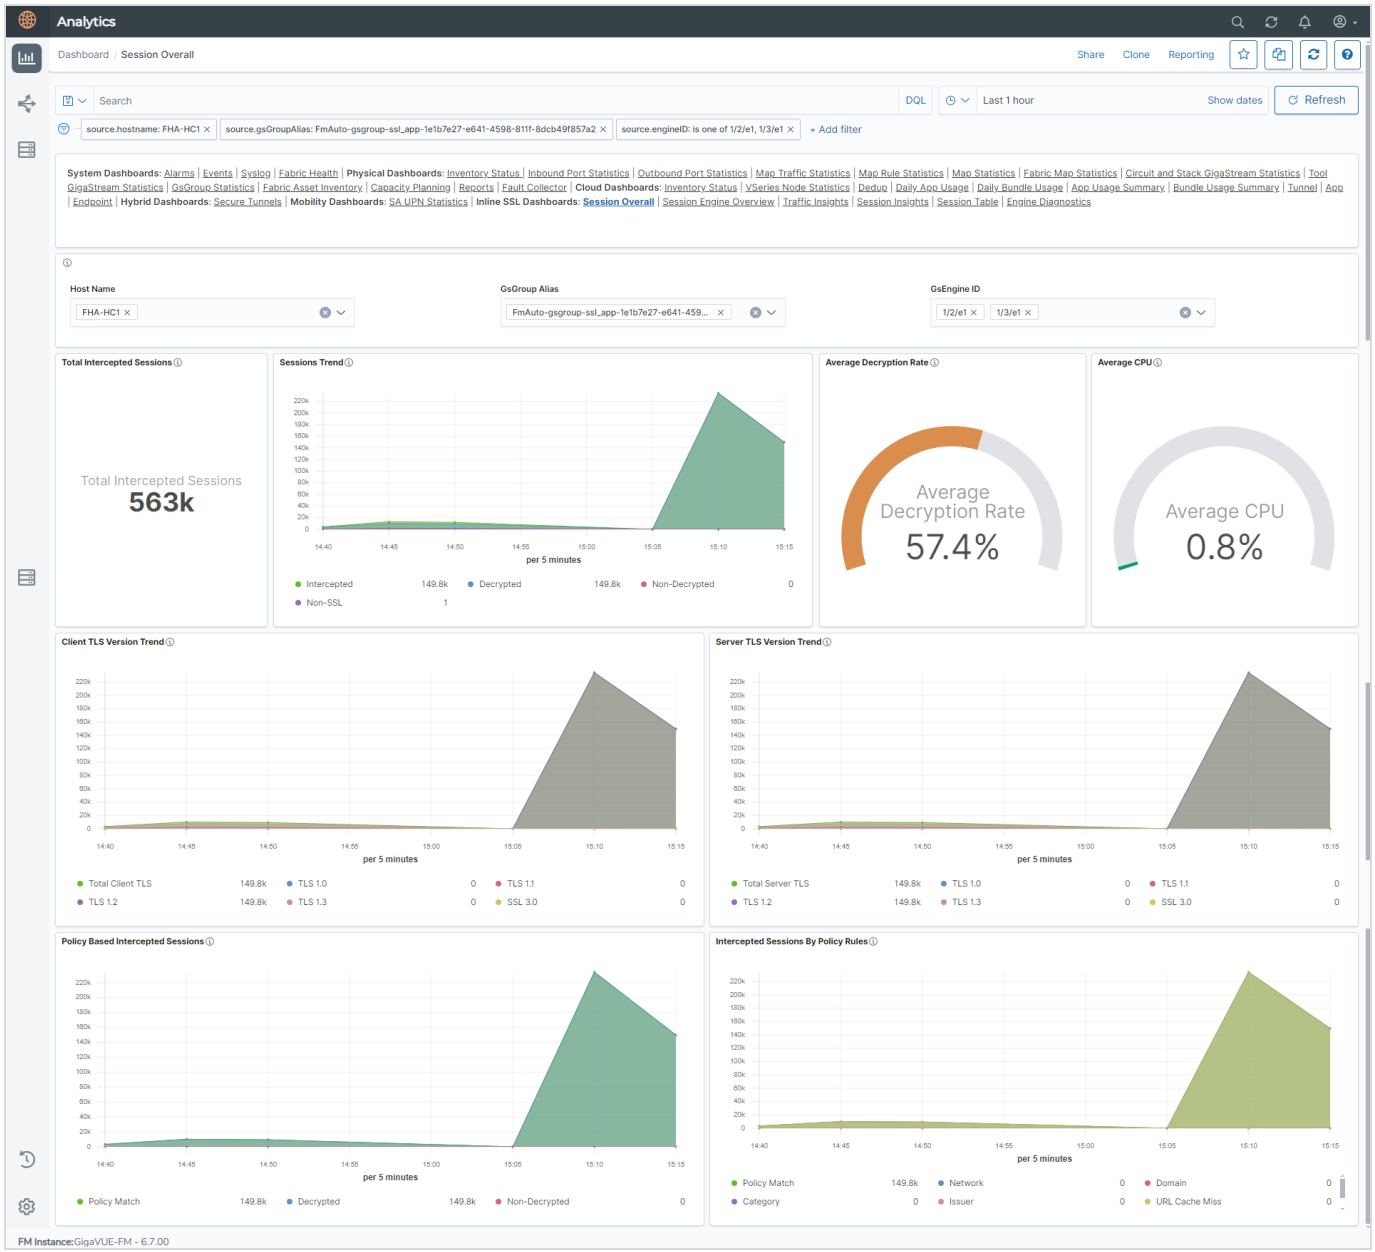

Table 10: Session Engine Overview Dashboard

Dashboard	Description	Visualizations	Details
Session Engine Overview	Displays visualizations related to Inline TLS/SSL Sessions per Engine	***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.
		Average CPS per Engine	Displays the average Connections per Second

Dashboard	Description	Visualizations	Details
			(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)

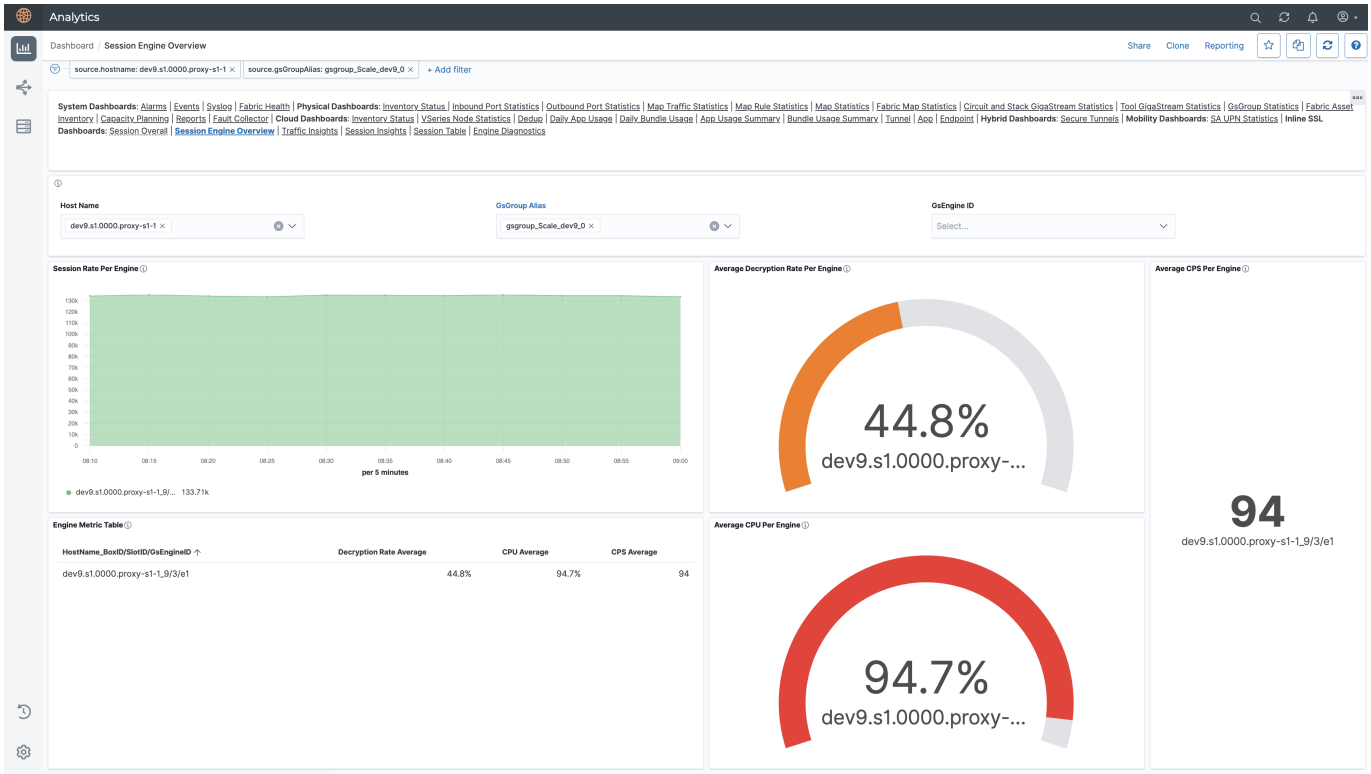

Table 11: Traffic Insights Dashboard

Dashboard	Description	Visualizations	Details
Traffic Insights	Displays visualizations related to the traffic that is handled with a Inline TLS/SSL sessions	***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 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.

Dashboard	Description	Visualizations	Details
		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.

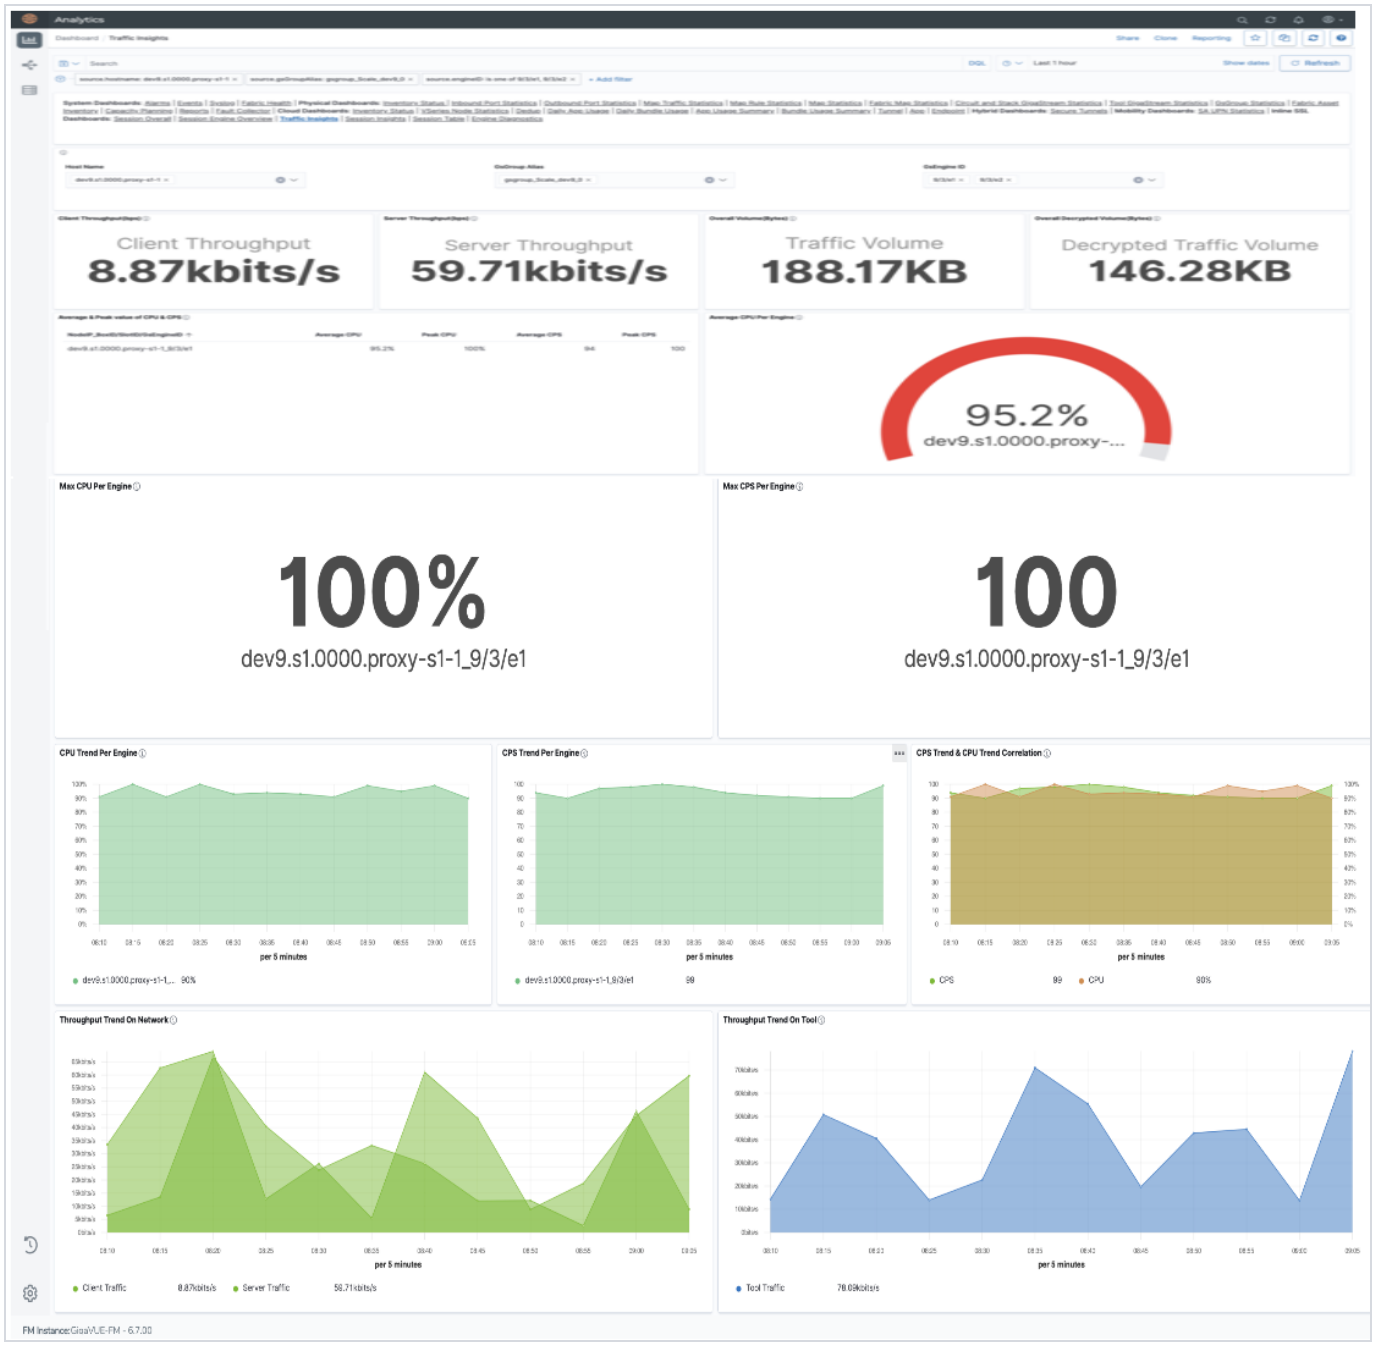
Analytics
Client Throughput
8.87kbits/s
Server Throughput
59.71kbits/s
Traffic Volume
188.17KB
Decrypted Traffic Volume
146.28KB
95.2%
dev9.s1.0000.proxy-...
Max CPU Per Engine
100%
dev9.s1.0000.proxy-s1-1_9/3/e1
Max CPS Per Engine
100
dev9.s1.0000.proxy-s1-1_9/3/e1
CPU Trend Per Engine
CPS Trend Per Engine
CPS Trend & CPU Trend Correlation
per 5 minutes
CPS
CPU
Throughput Trend On Network
Throughput Trend On Tool
Client Traffic
Server Traffic
Tool Traffic
FM Instance:GigaVUE-FM - 6.7.00

Table 12: Engine Diagnostics Dashboard

Dashboard	Description	Visualizations	Details
Engine Diagnostics	Displays the certificates and SSL alerts related to a GigaSMART engine	***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.

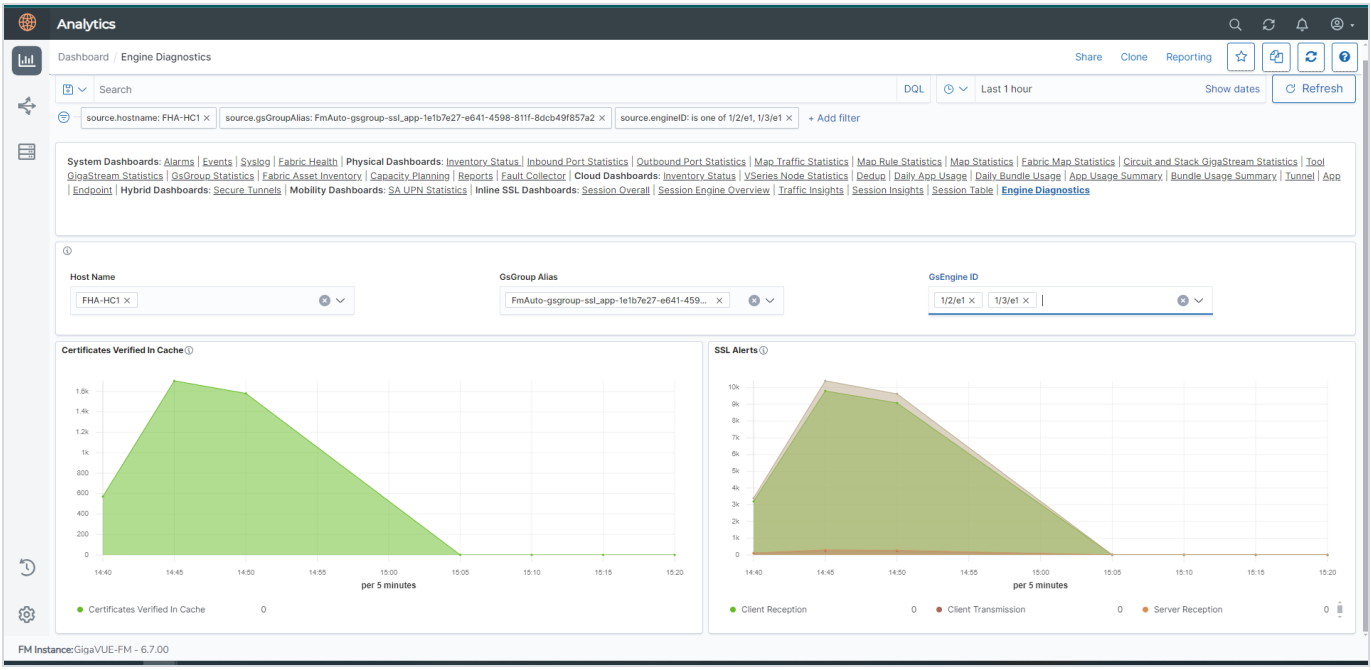

Advanced Dashboards

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 eight 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 are the advanced dashboards and its visualizations:

Table 13: Session Insight Dashboard

Dashboard	Description	Visualizations	Details
Session Insights	Displays visualizations on the details of an Inline TLS/SSL session.	***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

Dashboard	Description	Visualizations	Details
			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.

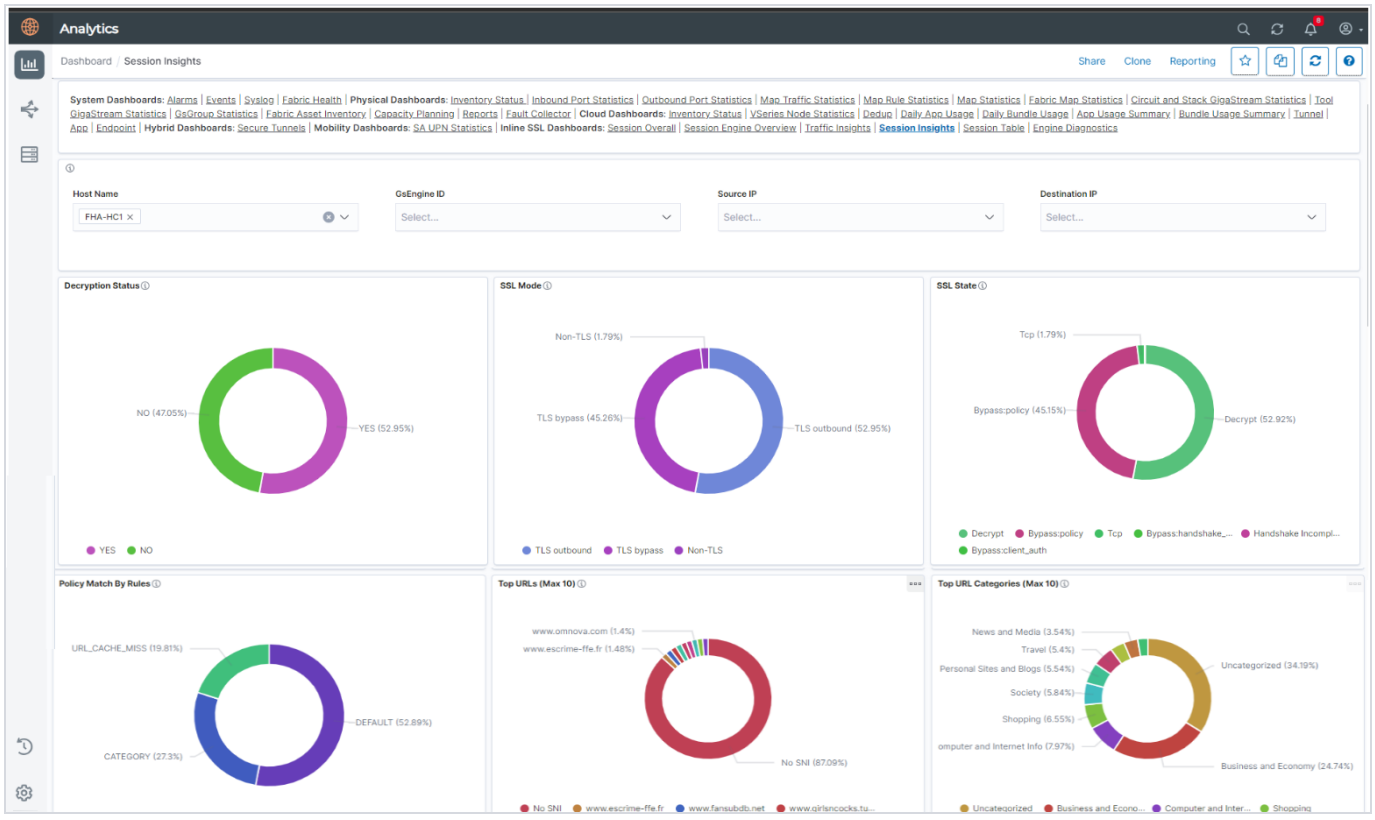

Table 14: Session Table Dashboard.

Dashboard	Description	Visualizations	Dashboard
Session Table	Displays visualizations related to Inline TLS/SSL Sessions per Engine in a tabular format.	***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

Dashboard	Description	Visualizations	Dashboard
			option by using (+) (−)l button.

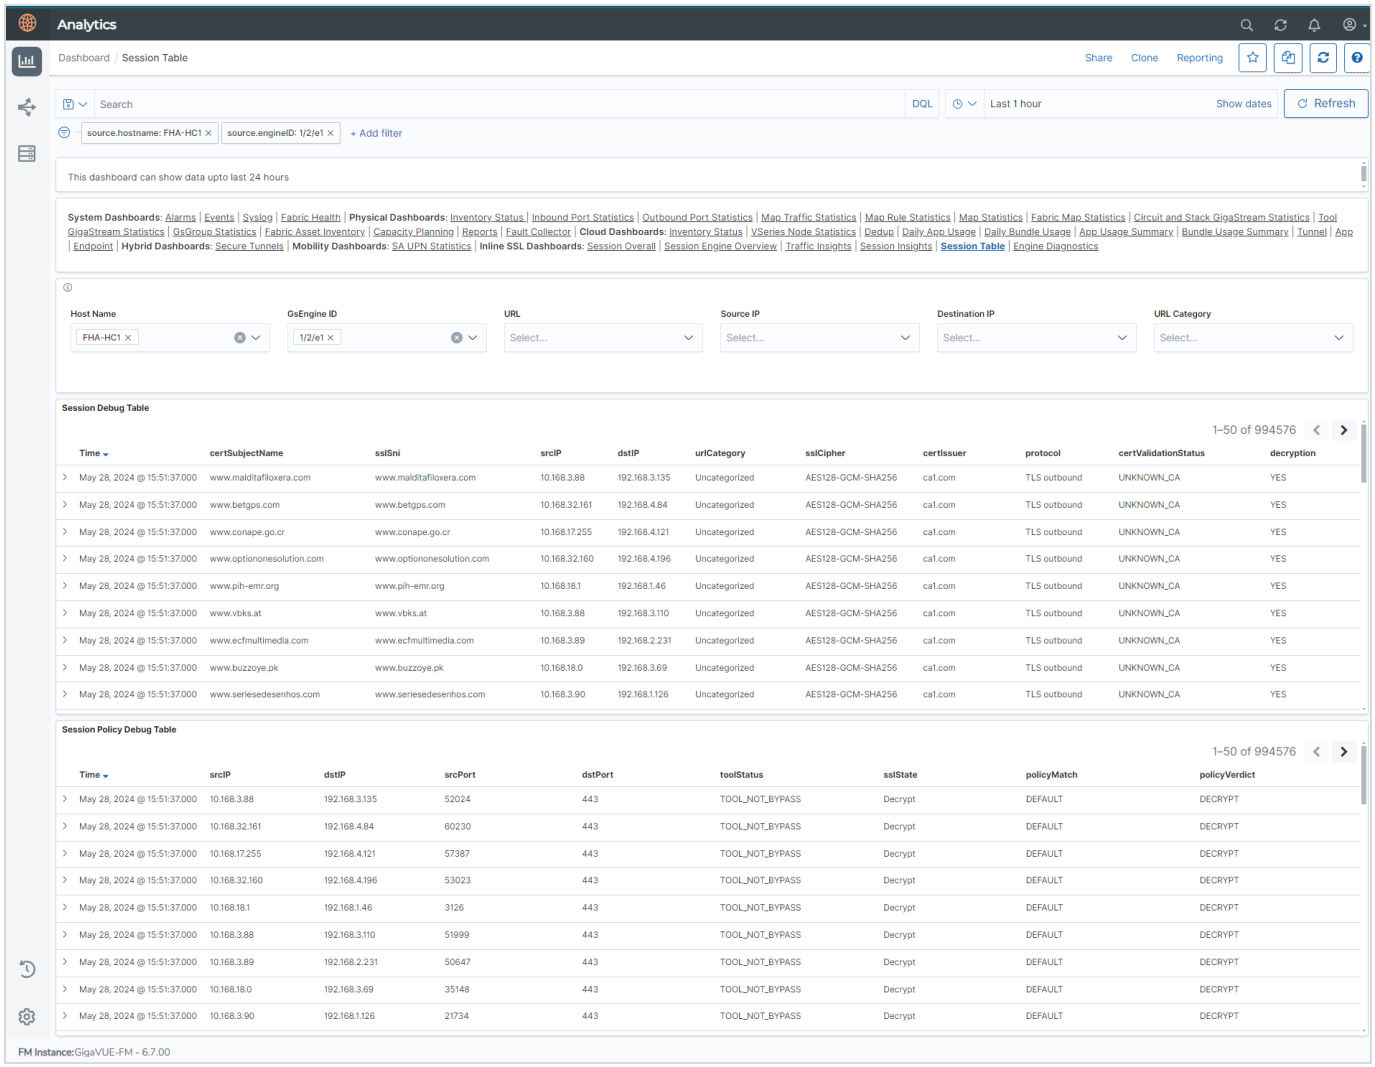

OpenSearch Indices and Fields Used in Analytics

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
- **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 15:

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**	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

[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.

New Field	Applicable Indices	Type	Description

	• **fmstats:** Applicable for the following document types:
	  • **gsgroup stats**
	  • **gsgroup port rate core rate stats**
	  • **vport stats**
	  • **Map and MapRule stats**
	  • **gsop stats**		

resource.id.clusterId	• **fmalarms**
	• **fmstats:** Applicable for the following document types:
	  • **vport stats**
	  • **gsop stats**
		text	Represents the **clusterId**. If the **clusterId** does not exist in the document, then **resource.id.clusterId** also will not exist in the document

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.

New Field	Applicable Indices	Type	Description

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.

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

New Field	Applicable Indices	Type	Description
gsVport.alias	• **fmstats:** Applicable for the following document types: • **vport stats** • **gsop stats**		Represents the alias name of the vport alias
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 16:

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

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

Fields	Description
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.
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.

Fields	Fields
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.
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

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.	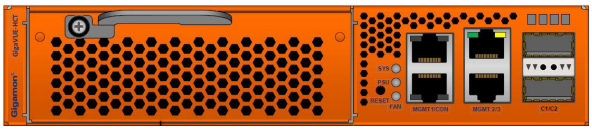
GigaVUE-TA25	• 1 RU Footprint • 1Gb/10Gb/48 x 25Gb/ ports and 8 x 100Gb/40Gb ports, dual hot-pluggable power supplies (AC/DC), four rear hot swappable fan modules, two console ports, and a 10Mb/100Mb/1Gb management port • Optional patch or breakout panel support • Cluster with GigaVUE HC Series and GigaVUE 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).	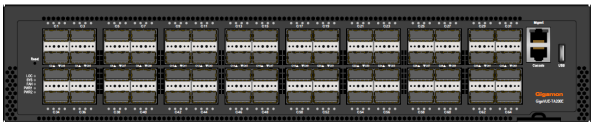
GigaVUE-TA400	• 1RU Footprint • 32 x 400Gb /100Gb/ 40Gb QSFP-DD/ QSFP28/QSFP+ ports, dual hot-pluggable power supplies (AC/DC), seven rear hot swappable fan modules, console port and a 10M/100M/1G management port. • Cluster with GigaVUE HC Series and GigaVUE TA Series Nodes.	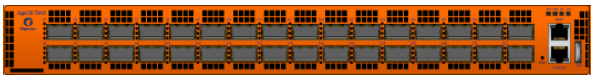

Notes on TA Series Nodes

- 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, click on  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.

Field	Description
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. **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. • **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.	Refer to the Tags section in the *GigaVUE Administration Guide* for details on how to create, edit, delete tags.
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.

Button	Description	Reference
	• 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.	
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, click on 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 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

Field	Description
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
10.114.34.113	admin	admin123A!!

To add physical nodes by importing from a spreadsheet, do the following:

1. On the left navigation pane, click on  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, click on 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.

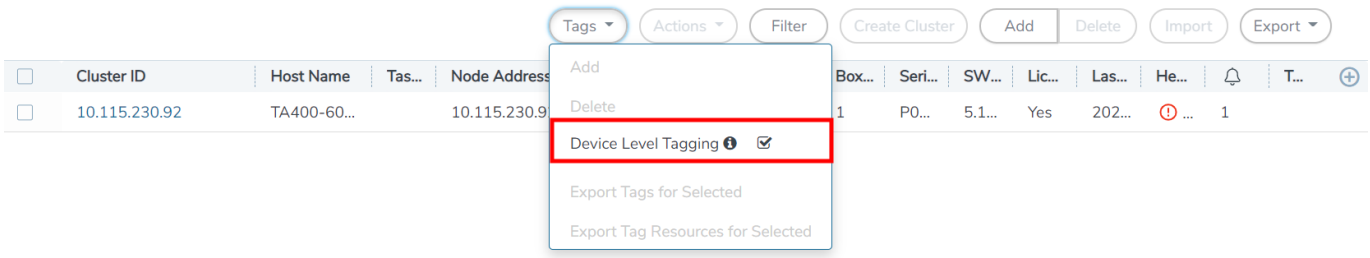

The nodes in a cluster are now associated to tags which can be used for creating the required topology.

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, click on  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, click on  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. On the left navigation pane, click  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.

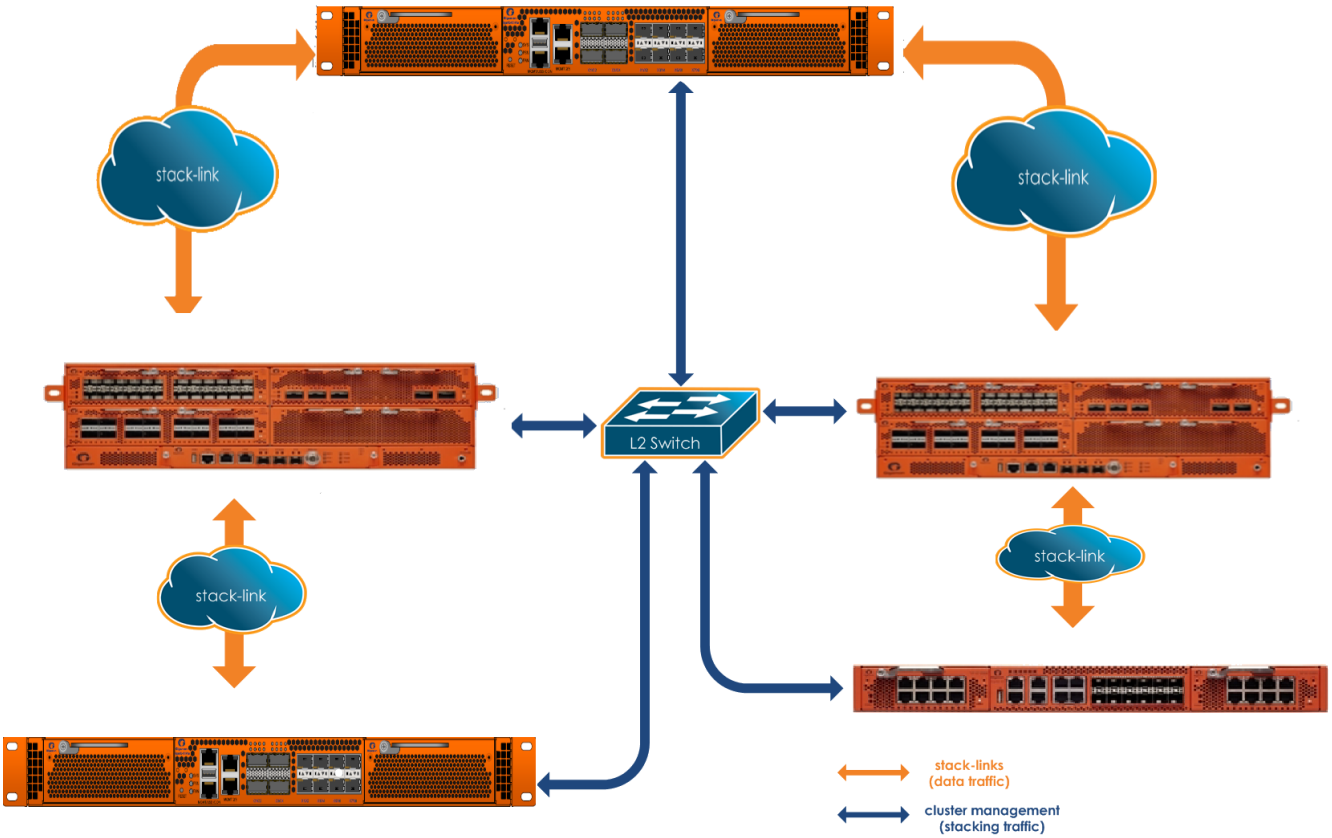

Figure 2 *Sample Cluster*

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-C05 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.

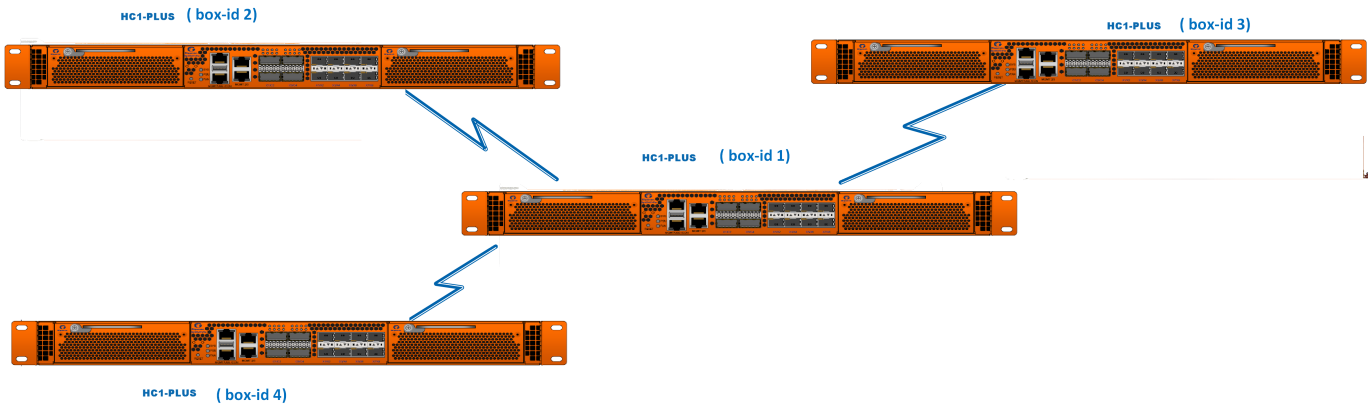

Figure 3 *Star Configuration*

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:

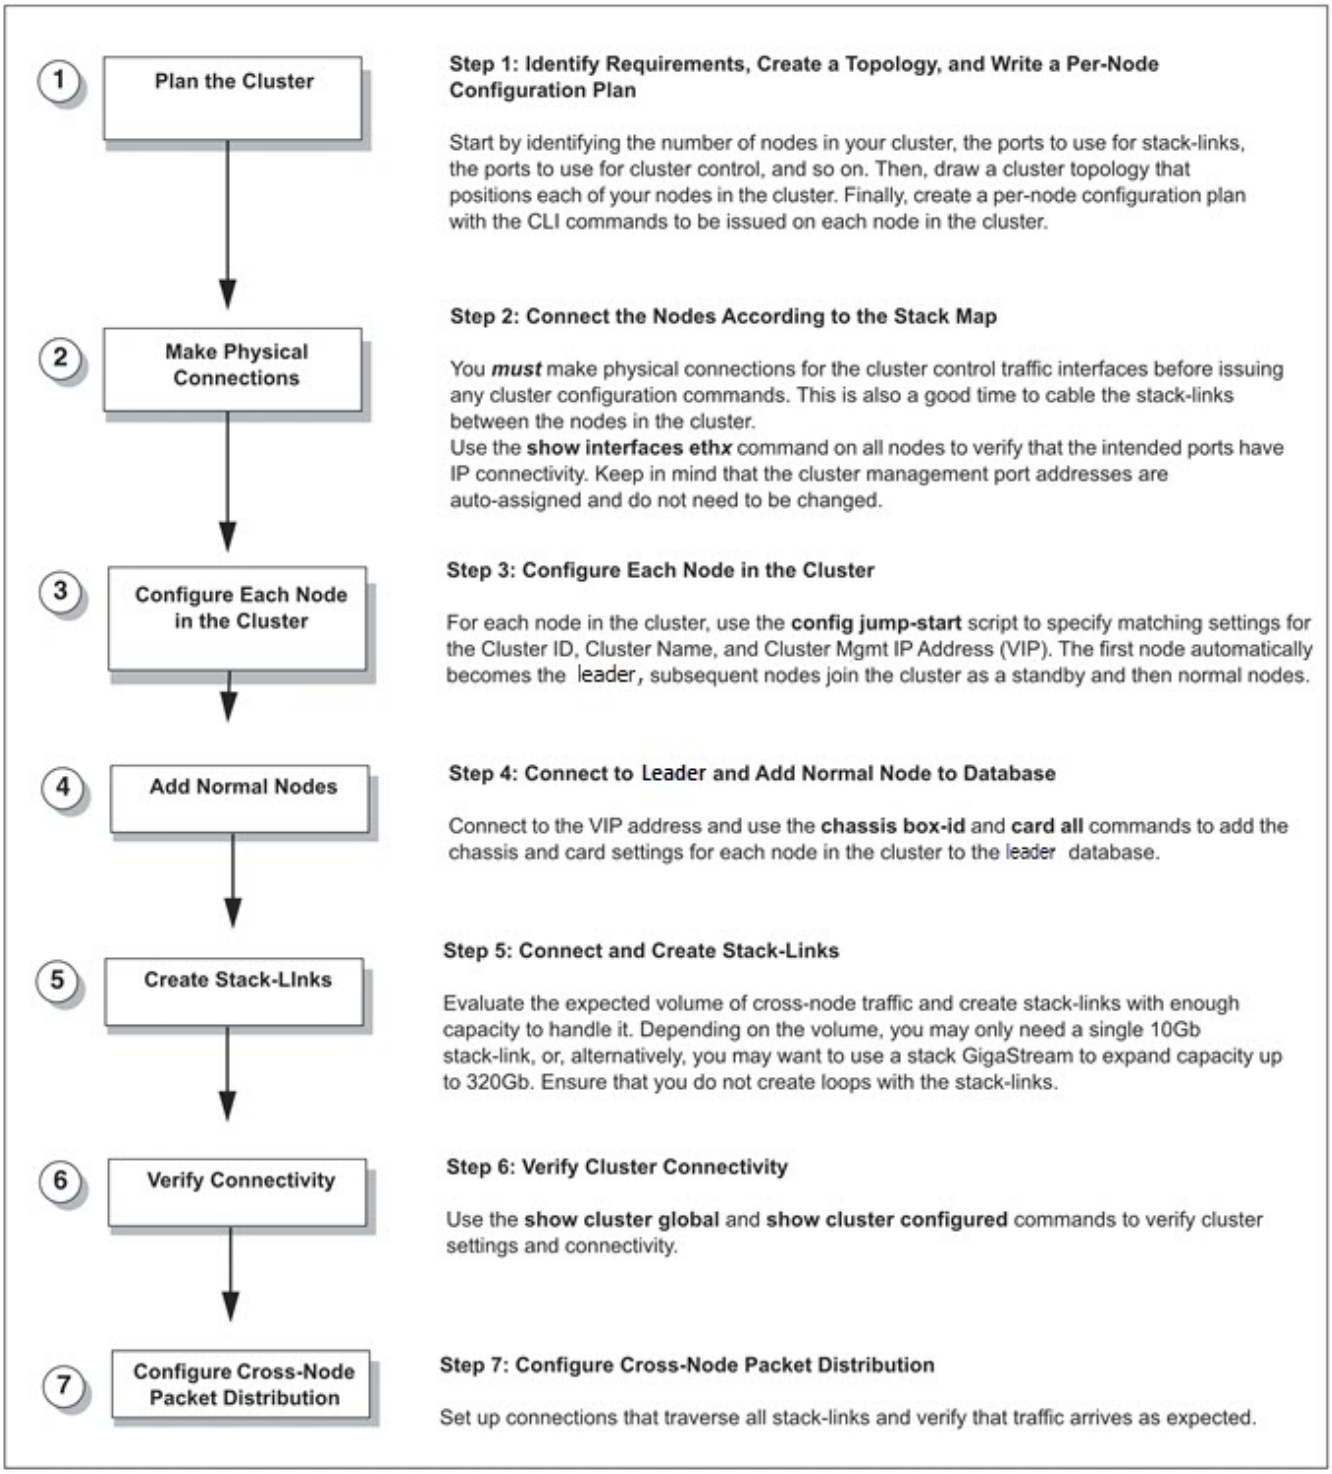

Figure 5 *Creating Clusters: Major Steps*

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.

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** option. 5. Enter the following details: • 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 **no traffic all** or **reset factory** 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	• **If the node is already managed by GigaVUE-FM:** Config sync will fail and is notified with appropriate messages. You

Process	Until Software Version 5.10.00	From Software Version 5.11.00
	updated. • **If the node is not already managed by GigaVUE-FM:** The node will be added to GigaVUE-FM.	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 device is moved out is triggered first.

Process	Until Software Version 5.10.00	From Software Version 5.11.00
		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.
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.

Rule
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. All legacy cluster topologies (leaf-spine, tree, etc.) with OOB cluster management are supported with E-tag stacking mode. Refer to the example below.

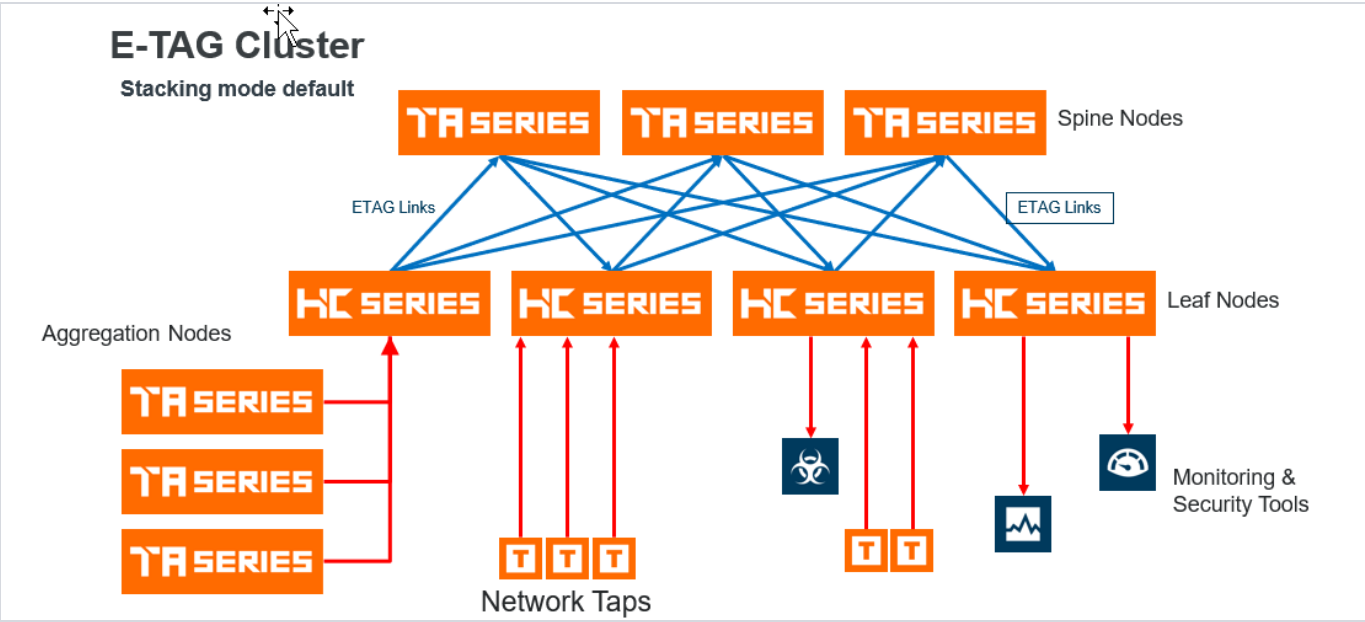

The new E-TAG stacking mode is supported on all the platforms listed below.

- GigaVUE-HC3 CCv1 & CCv2
- GigaVUE-HC1
- GigaVUE-HC1-Plus
- GigaVUE-HCT
- GigaVUE-TA25
- GigaVUE-TA25E
- GigaVUE-TA100
- GigaVUE-TA200
- GigaVUE-TA200E
- GigaVUE-TA400
- DELL S4112F-ON

The BiDi optics on stack ports are also supported in E-Tag stacking mode.

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.

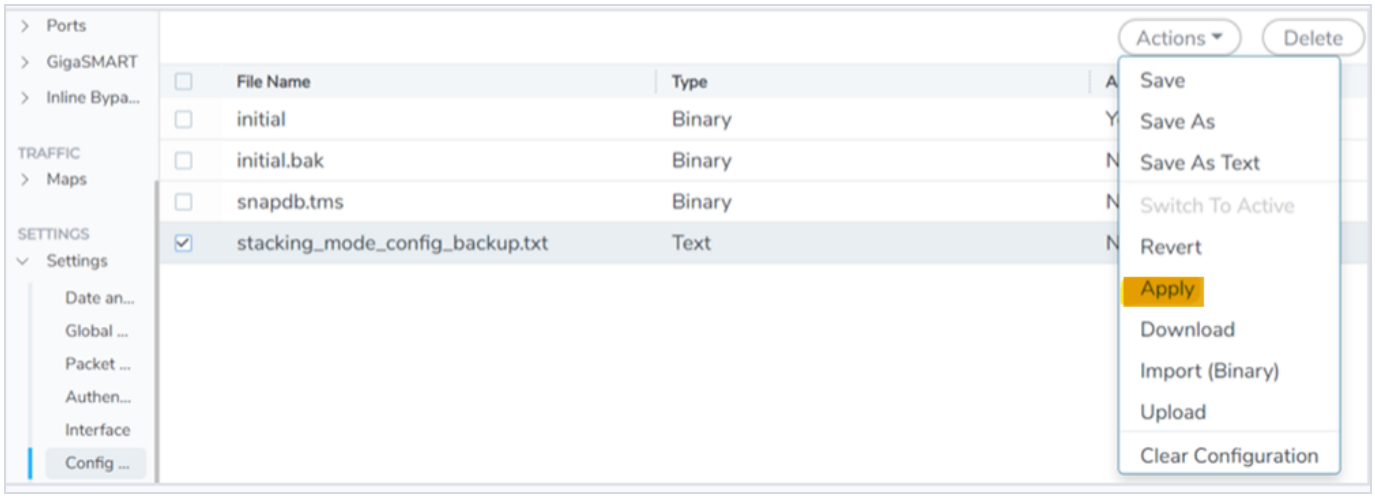

Switching the stacking mode will not be allowed if the cluster includes an unsupported node.

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 any number of devices with different combinations of 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 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.

Create Regular Cluster Formation

To create a cluster:

1. On the left navigation pane, click  , and then select **Physical > Nodes**..
2. Click **Create Cluster**.

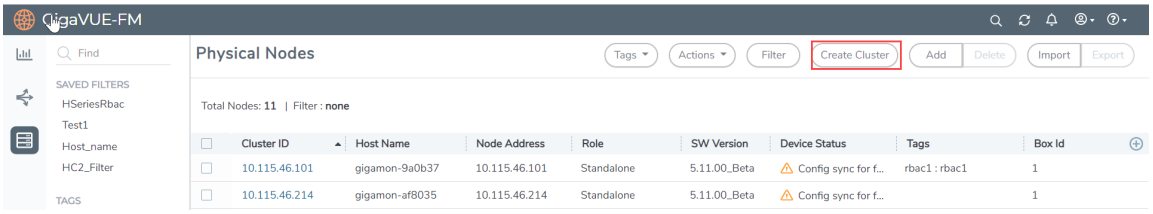

Figure 6 *Create Cluster*

Select Cluster Type

3. The Create a Cluster screen opens with two options:
   - Create a Cluster
   - Create a Leaf Spine Cluster
4. Hover over the Create a Cluster option and click **Let's Begin** to start the wizard.

Select Devices

The wizard guides you through the cluster set-up. The first step is to select the devices in your cluster.

5. The Select Devices page displays a list of standalone devices with filter options:
   - *Software Version:* Filter the nodes based on the software version for which the cluster will be formed.
   - *Model*: Filter the nodes based on a Gigamon model.
   - *HostName/Node IP*: Enter the HostName of the Gigamon Nodes to specify a device.
   - *Stacking Mode*: Select the required Stacking Mode. The following options are available:
     - ***Legacy***: Selects the legacy stacking mode.
     - ***Default***: Selects the Default stacking mode.

Notes:

- GigaVUE-TA400 and GigaVUE-HCT devices are not supported in Legacy stacking mode.

- *Protocol*: Select the required protocol based on which the devices will communicate with each other. The following options are available:
  - ***IPv4***: Cluster leader and the member nodes communicate through IPv4.
  - ***IPv6***: Cluster leader and the member nodes communicate through IPv6.

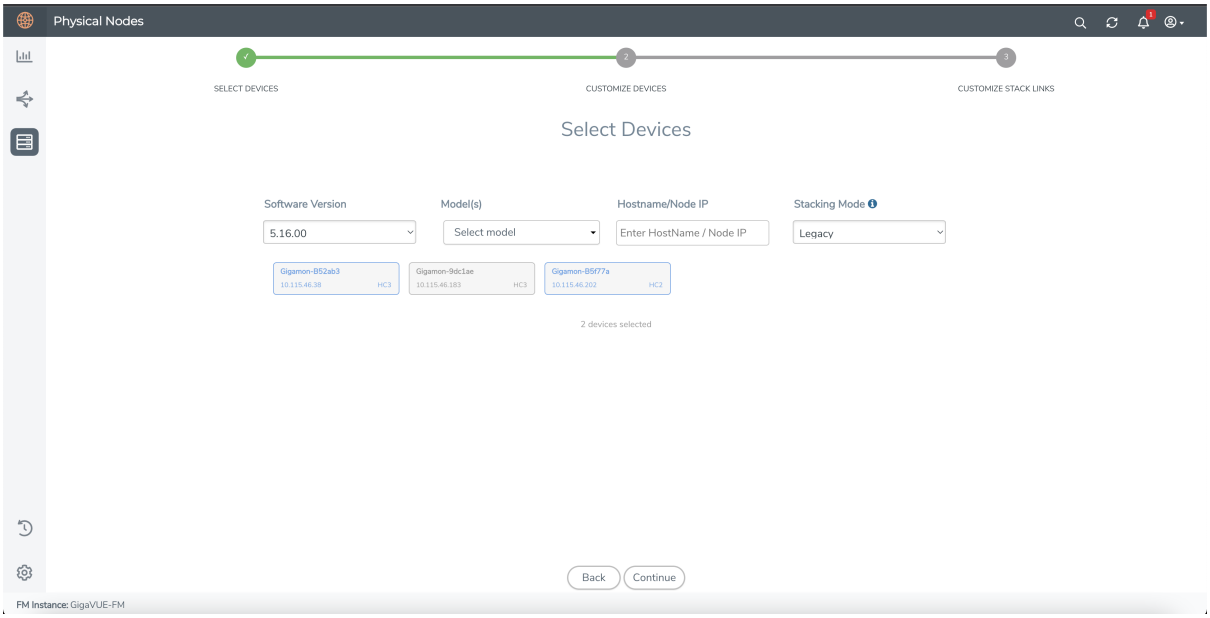

6. Select the nodes to include in this cluster and click **Continue**.

   Click a device to select it; click it again to deselect it. Selected devices are highlighted.

Customize Devices

Use the Cluster Configuration window to customize your devices.

7. Enter a valid **Cluster ID** and **VIP** and select the leader in the **Seed Node** list.

   NOTE: The leader cluster preferences in GigaVUE-FM determines which of the nodes will be the default the seed node. TA devices cannot be a leader.

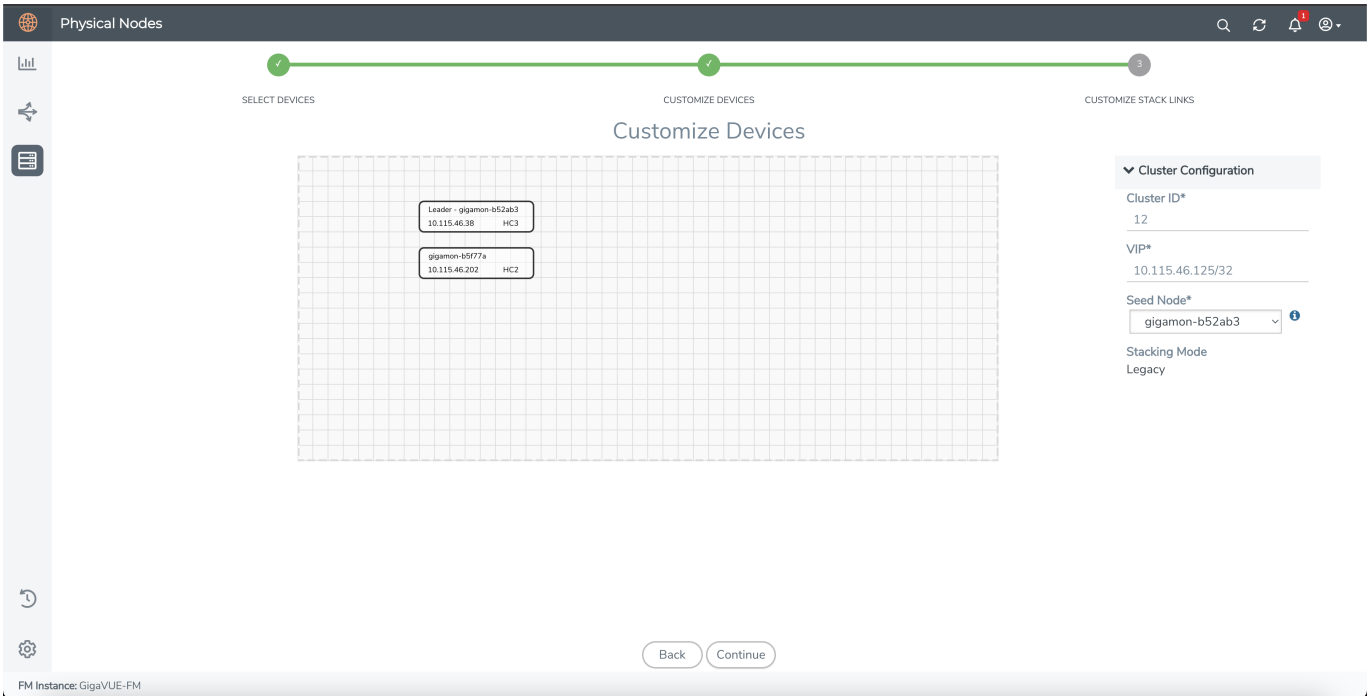

8. After completing the Cluster Configuration details, click **Continue**.

   NOTE: Use the **Back** button to return to the Select Devices page to revise the selection of devices for this cluster.

Customize Stack Links

Finally, customize the stack links to define how the nodes should be connected.

9. Connect any two devices to create a stack links 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 will illustrate the connection.

10. Configure the formed links in the **Stack Links** table as follows:
    - Select ports in each device that are compatible, for example: x-x ports ,x-q ports, q-c ports, x-c ports.
    - Select two or more ports in each device to create a stack GigaStream.
    - After selecting the ports, save the stack link by clicking the **Save** button enabled in the right of stack link table.

    The alias for each stack link and GigaStream is auto generated by GigaVUE-FM. This alias can be edited as needed.

11. After the required Stack Links and GigaStreams are created and saved, click **Continue** to start the cluster creation process.

    The Creating Cluster page 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.

NOTE: Refer to Check Cluster Status for Events.

When the cluster formation process is complete the notification window will display a, "Create Cluster Competed," message.

12. Click **Go to Cluster** to view the cluster overview.

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.
- 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.

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:

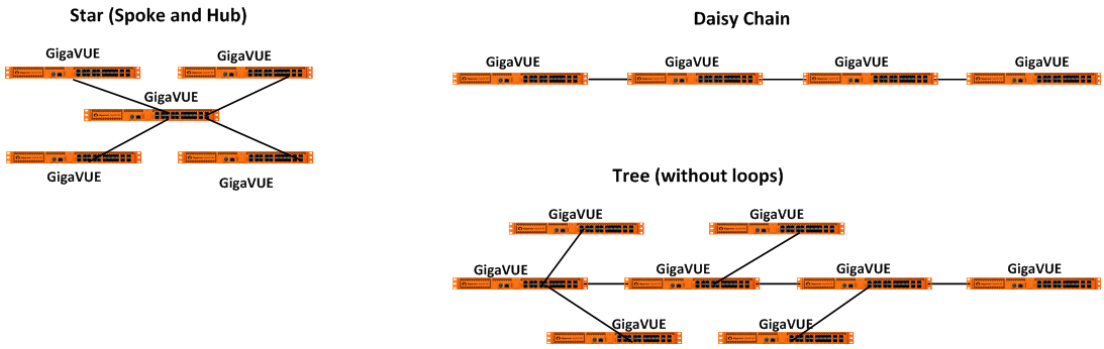

NOTE: Subsets or aggregations of these topologies may be created; however, it is important not to create a loop within these specified topologies.

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.

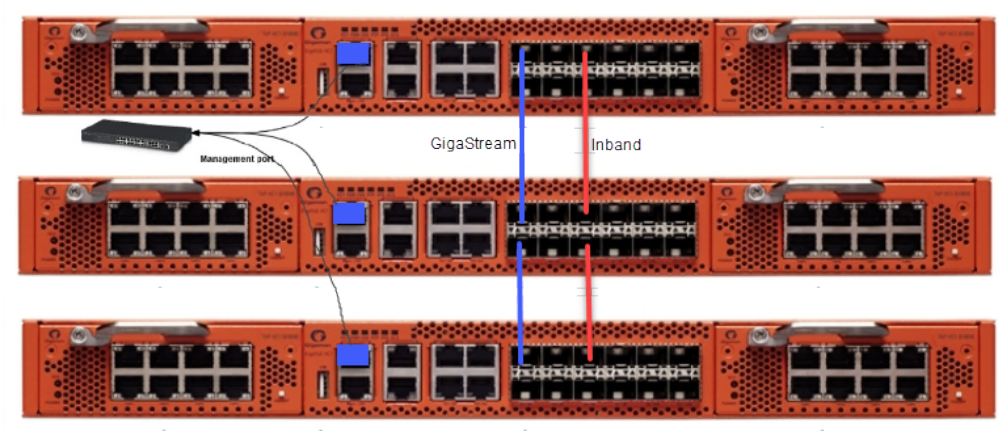

NOTE: Ensure that there is a physical connection between the stack ports of the two nodes that are being added to the Inband cluster.

Inband Cluster Management Configuration Flow Chart

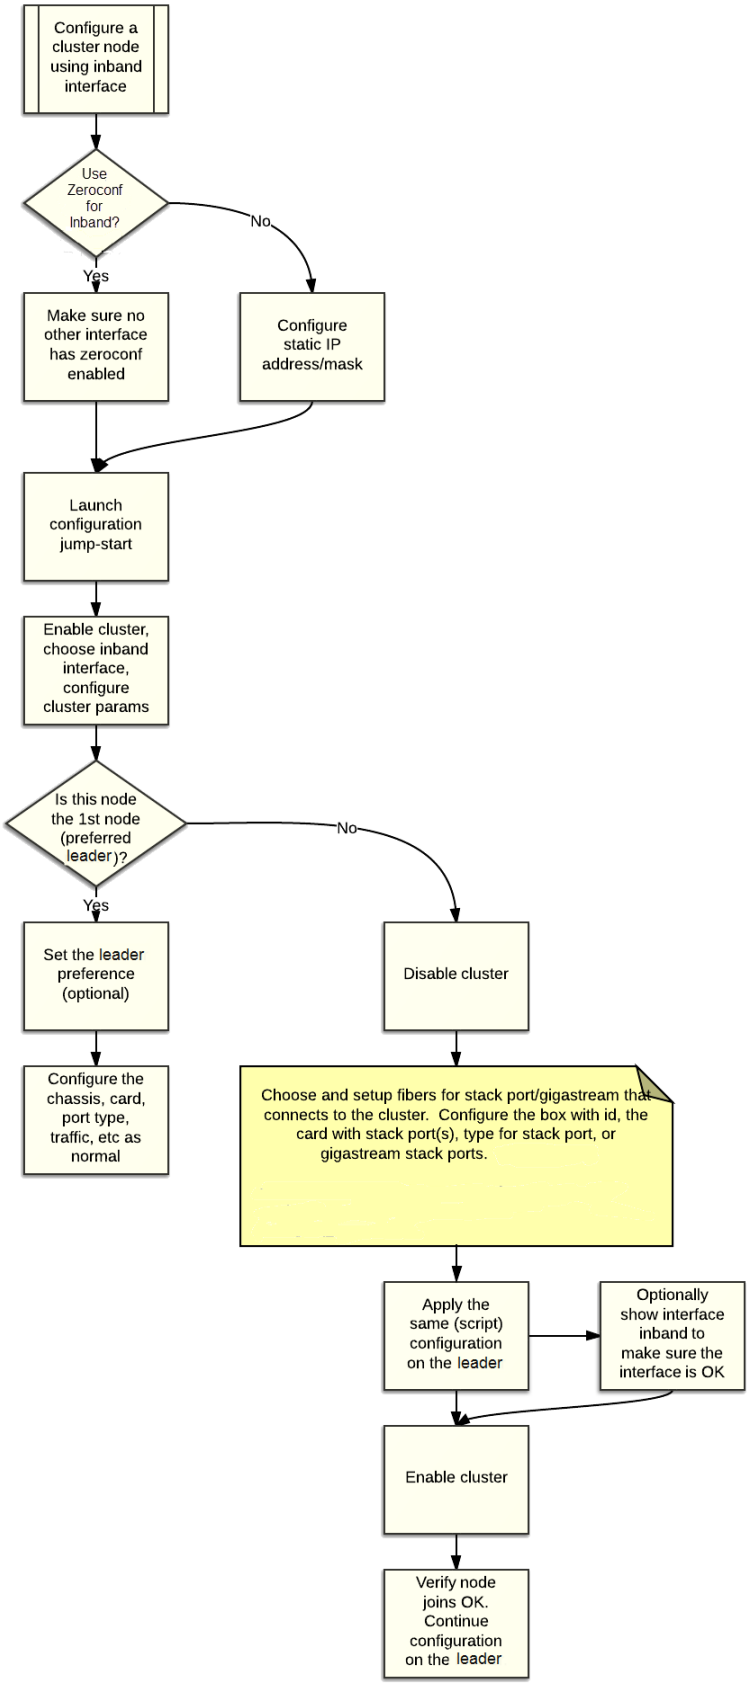
Configure a cluster node using inband interface
Use Zeroconf for Inband?
No
Yes
Make sure no other interface has zeroconf enabled
Configure static IP address/mask
Launch configuration jump-start
Enable cluster, choose inband interface, configure cluster params
Is this node the 1st node (preferred leader)?
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. On the left navigation pane, click  , and then select **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. On the left navigation pane, click  , and then select **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, click on 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.

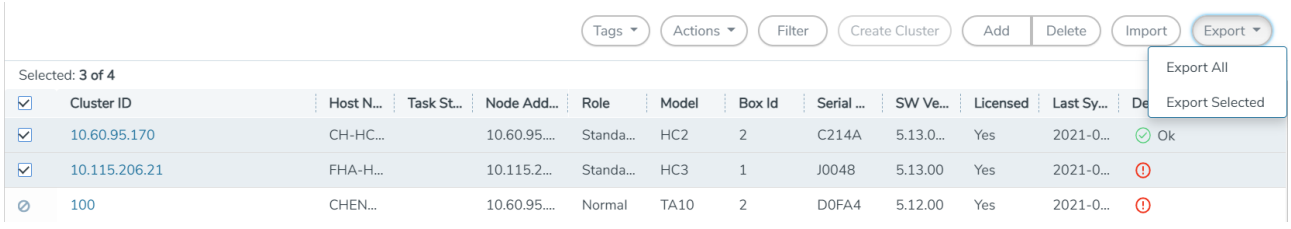

However, the member nodes will not get downloaded as they do not get selected. To select and download the member nodes:

1. Enable **Device Level Tagging** option under the Tags menu.
2. Click **Export Selected** to export the selected nodes, including the member nodes.

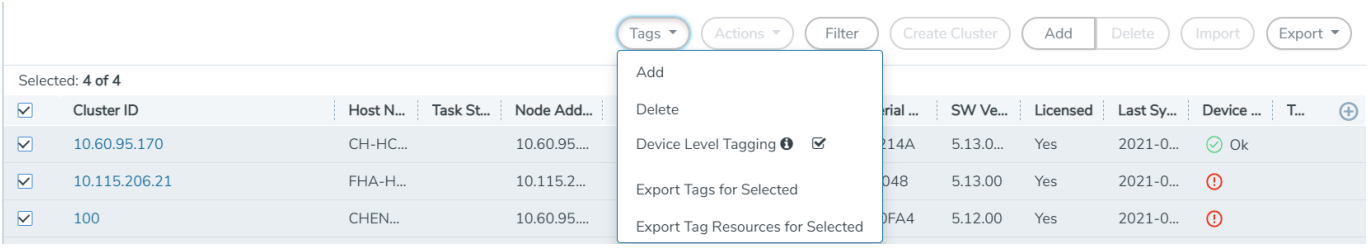

Problems with SCP?

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

On the left navigation pane, click  , and then select **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 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.

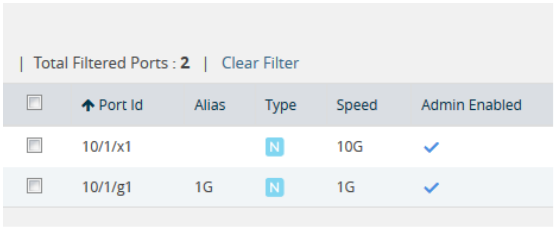

Overview Page

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).

Field	Description
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. **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. The wizard selects Symmetric for Traffic Type.

Map Type	Map Wizard	Description
	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**.

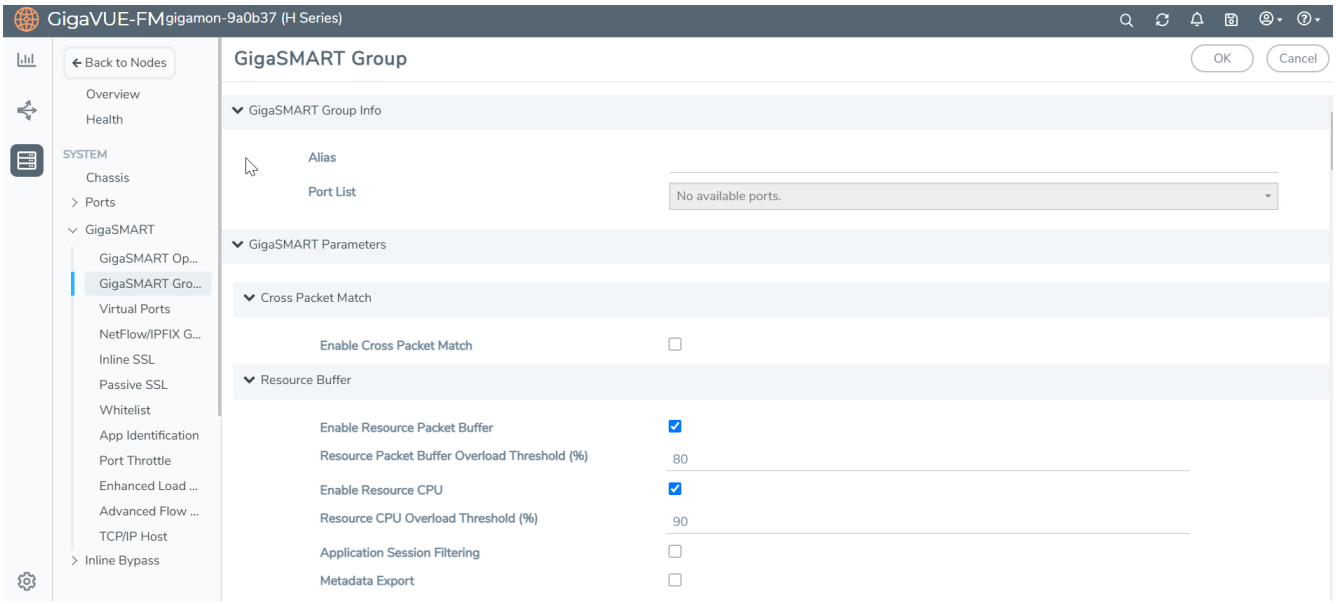

Figure 8 *GS Group Configuration in Workflow*

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.

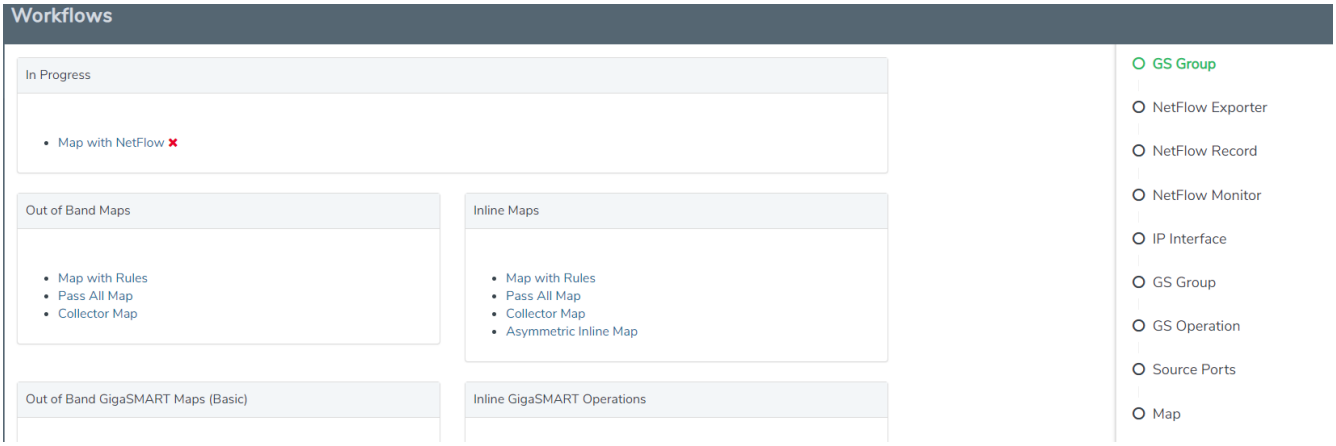

Figure 9 *Workflow in Progress*

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.

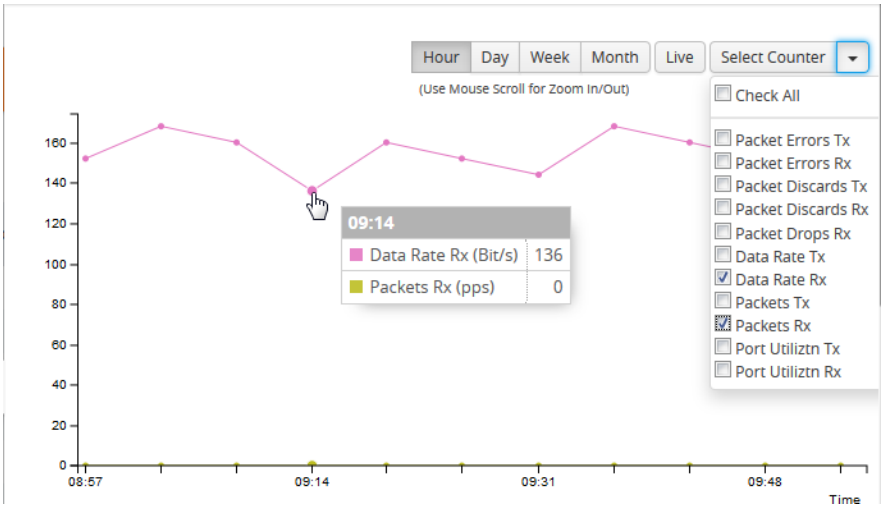

Safe and Limited Modes

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. On the left navigation pane, click  , and then select **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.

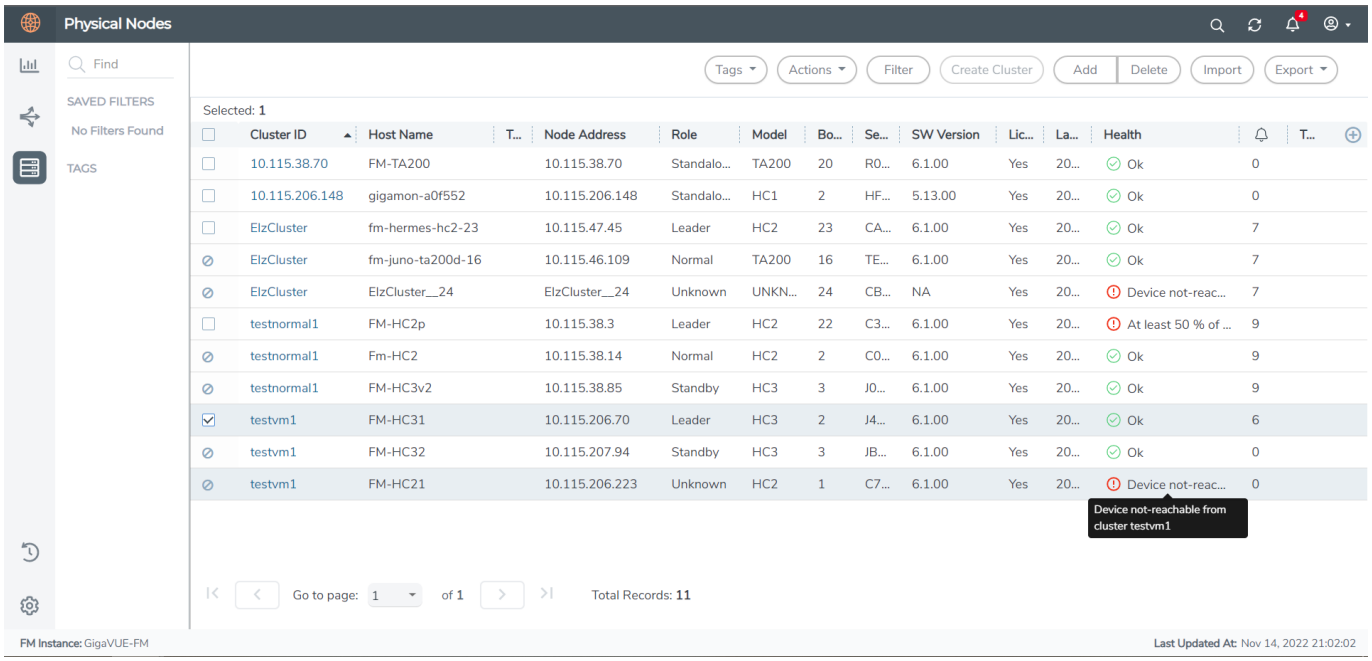

To view only clusters with not-reachable node from the Physical Nodes page:

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**.

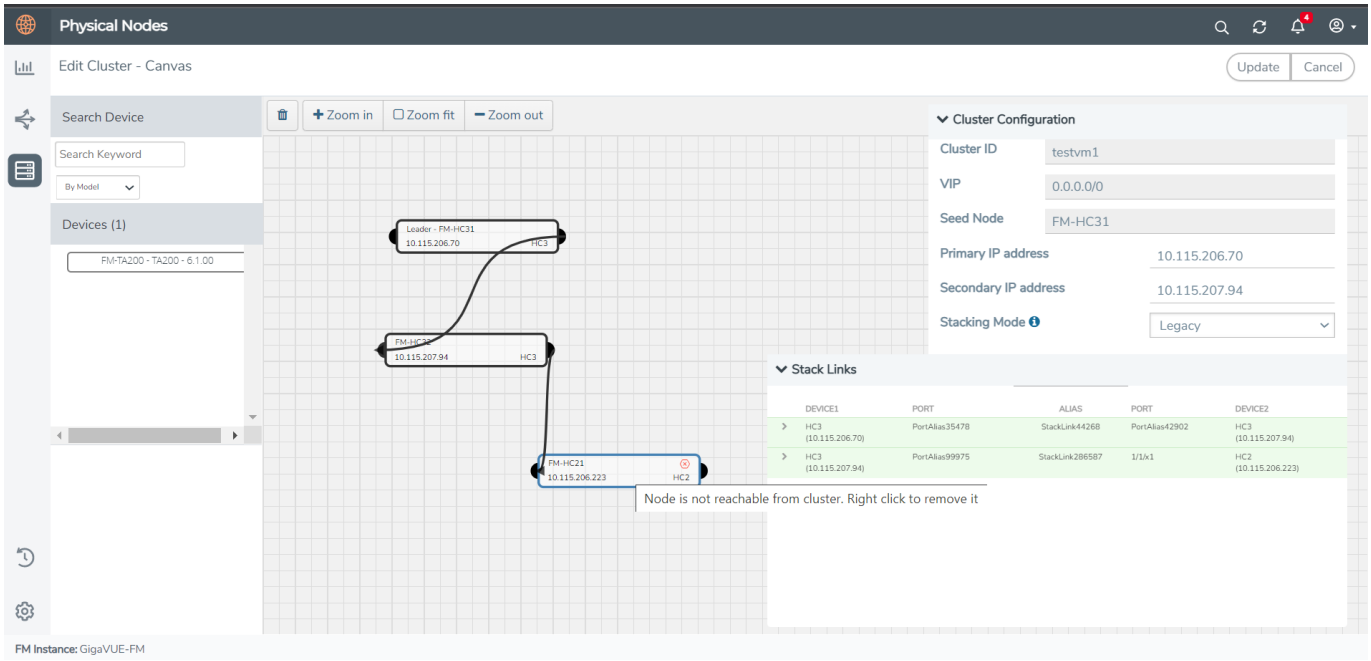

To perform this action:

- 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.

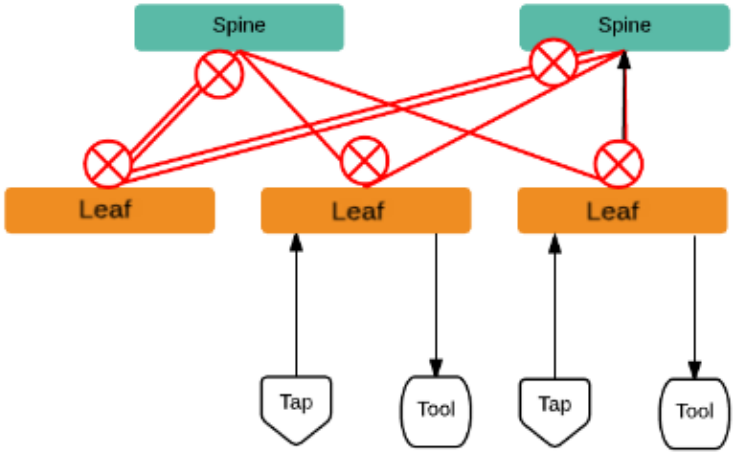

In a cluster, the number of leaf nodes is typically greater than the number of spine nodes. In the above figure, there are three leaf nodes and two spine nodes. The leaf nodes aggregate to a fewer number of spine nodes.

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.

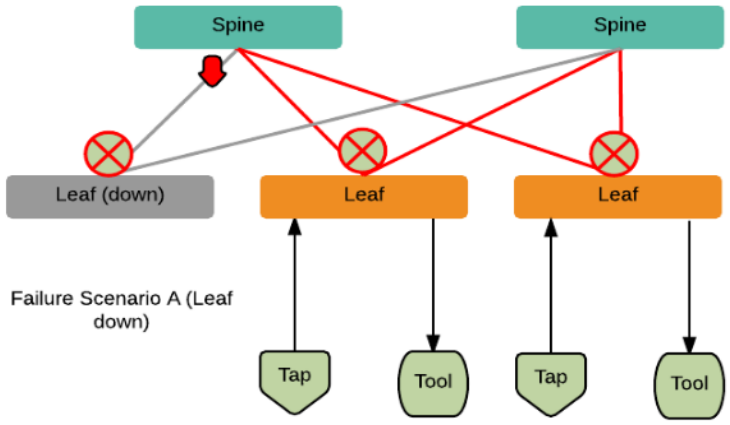

Restoration

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.

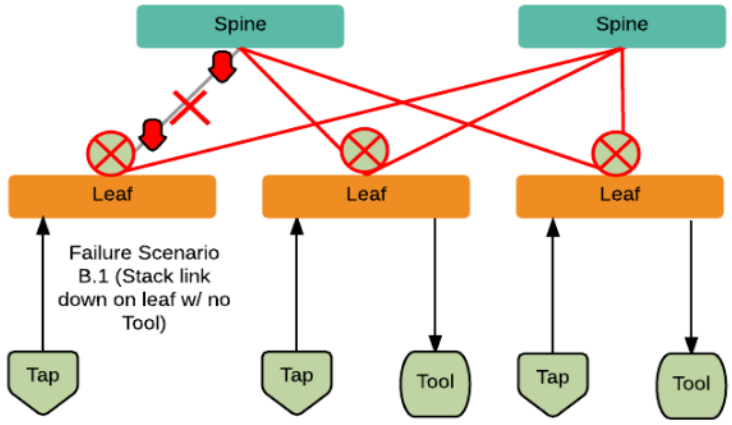

Restoration

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.

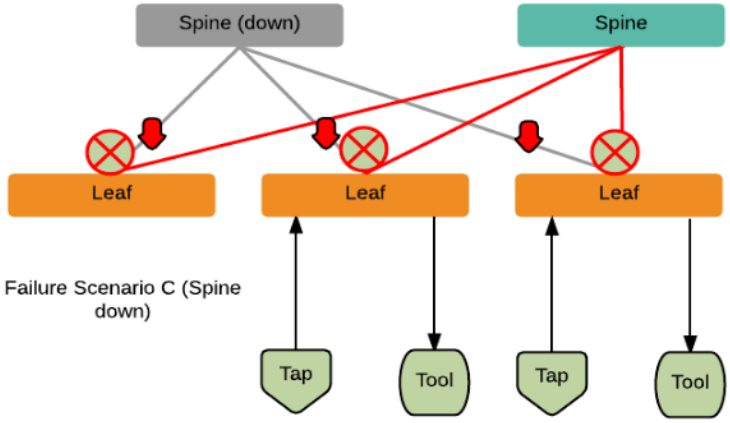

Restoration

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.

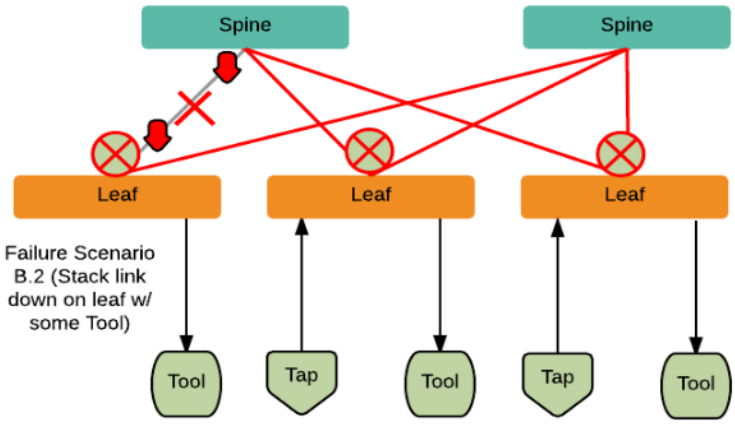

Configuration Overview

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.

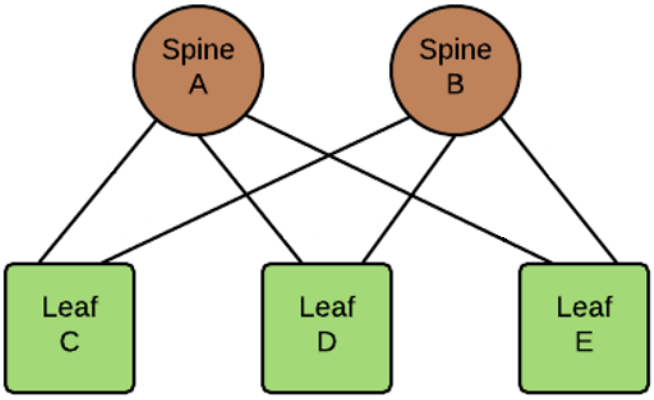

The configuration steps are as follows:

- 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.

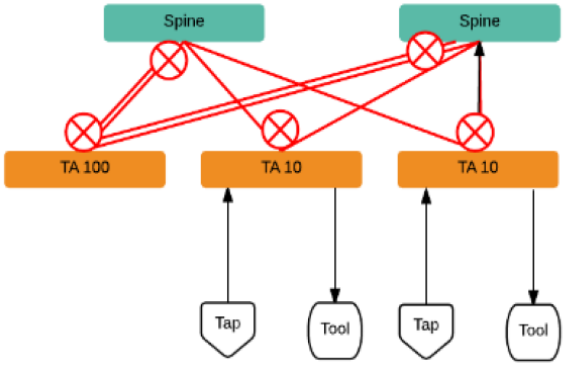

- Configure spine links. On each leaf node, there is one spine link that contains the list of GigaStream connecting the leaf nodes to the spine nodes. The spine links contain multiple stack GigaStream that are bundled together. The spine links are: {c1,c2}, {d1,d2}, and {e1,e2}. The total number of spine 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.

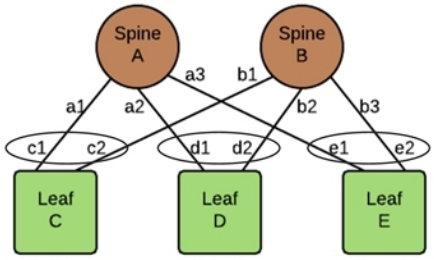

NOTE: For the spine links, make sure that all paths are reachable.

- 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.

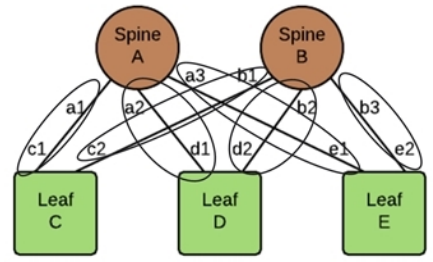

These configuration steps ensure that the spine and leaf nodes are fully meshed.

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.

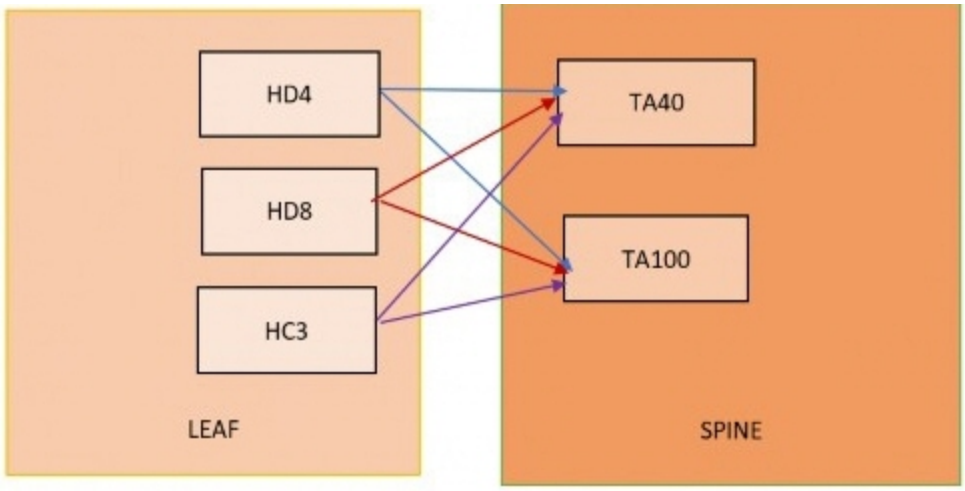

Figure 10 *Leaf spine cluster overview,*

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 5.3 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. On the left navigation pane, click  , and then select **Physical > Nodes**..
2. Click **Create Cluster**.

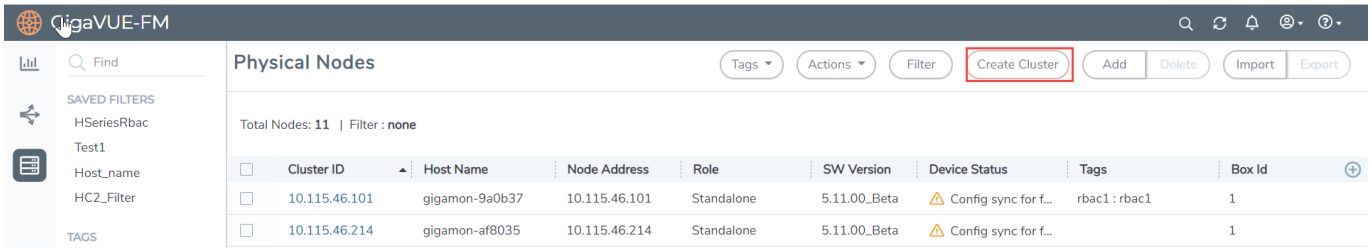

Figure 11 *Create Cluster*

Select the Cluster Type

3. The Create a Cluster screen opens with two options:
   - CREATE A CLUSTER
   - CREATE A LEAF SPINE CLUSTER
4. Hover over the CREATE A LEAF SPINE CLUSTER option and click **Let's Begin** to start the wizard.

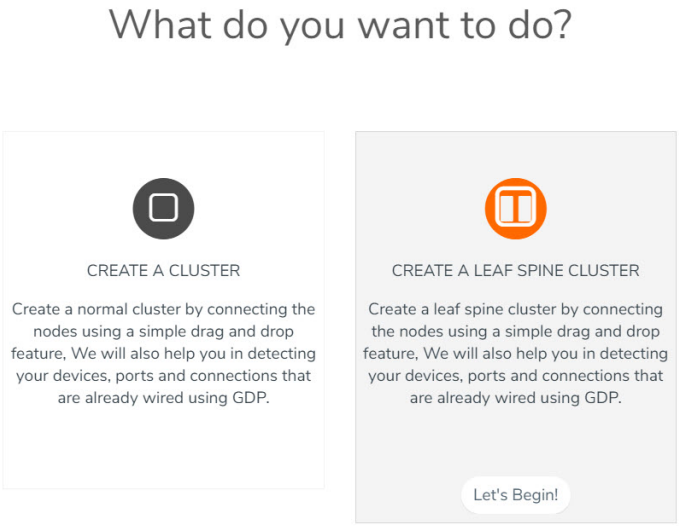

Select Devices

The wizard guides you through the cluster set-up. The first step is to select the devices in your cluster.

5. The Select Devices page displays a list of standalone devices with filter options:
   - **Software**: Filter the nodes based on the software version for which the cluster will be formed.
   - **Model**: Filter the nodes based on a Gigamon model.
   - **HostName**: Enter the HostName of the Gigamon Nodes to specify a device.

- **Stacking Mode**: Select the required Stacking Mode. The following options are available:
  - ***Legacy***: Selects the Legacy stacking mode.
  - ***Default***: Selects the Default stacking mode.

> **Notes:**
>
> - GigaVUE-TA400 device is not supported in Legacy stacking mode.

- **Protocol**: Select the required protocol based on which the devices will communicate with each other. The following options are available:
  - ***IPv4***: Cluster leader and the member nodes communicate through IPv4.
  - ***IPv6***: Cluster leader and the member nodes communicate through IPv6.

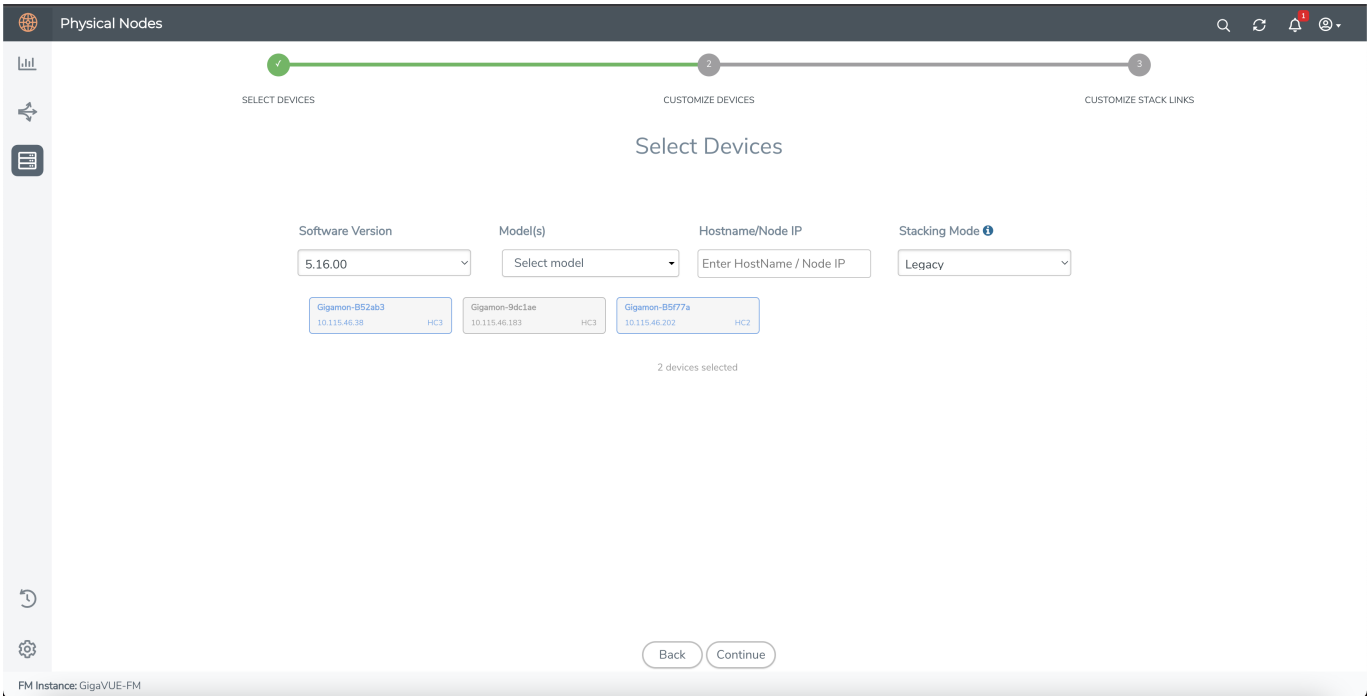

Select the nodes to include in this cluster and click **Continue**.

6. Click a device to select it. Click it again to deselect it. Selected devices are highlighted.

Customize Devices

Use the Cluster Configuration window to customize your devices.

7. Enter a valid **Cluster ID** and **VIP** and select the leader in the **Seed Node** list.

NOTE: TA devices cannot be a leader.

8. After completing the Cluster Configuration details, click **Continue**.

   NOTE: Use the **Back** button to return to the Select Devices page to revise the selection of devices for this cluster.

Customize Stack Links for your Leaf-Spine 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 will 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 will illustrate the connection.
11. Configure the formed links in the Stack Links table as follows:
    - Select ports in each device that are compatible, for example: x-x ports ,x-q ports, q-c ports, x-c ports.
    - Select two or more ports in each device to create a stack GigaStream.
    - After selecting the ports, save the stack link by clicking Save button enabled in the right of stack link table.

Stack GigaStreams and stack links can act as a spine link between the devices. The alias for stack link and GigaStream is auto generated by GigaVUE-FM. This alias can be edited as needed.

12. After the required Stack Links and GigaStreams are created and saved, click **Continue** to start the cluster creation process.

    The Creating Cluster page 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 the notification window will display a, “Create Cluster Competed,” message.

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 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 wizard.
- 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 cluster.

1. Select a cluster and choose **Actions > Edit cluster**.

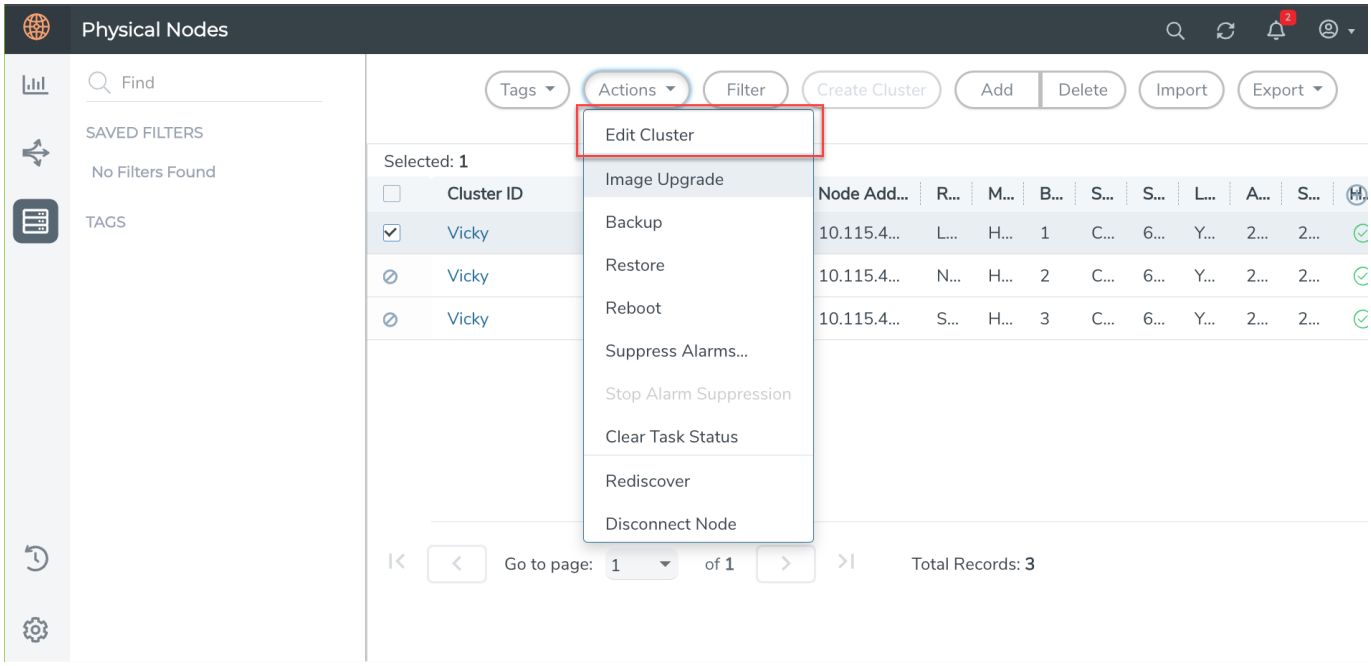

2. The Edit Leaf Spine Cluster Canvas displays the existing stack link configuration details in the cluster wizard canvas. Standalone devices are listed under the Devices pane.
3. Drag the required devices from the devices pane into the cluster wizard under the leaf or spine container.
4. Draw the links 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. After selecting the ports, save the stack link by clicking **Save** in the stack link table.
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 will appear 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 completed, “Manage Cluster Completed,” will appear in the Manage Cluster notification window.
10. Click **Go to Cluster** to go to view the cluster overview.

The created GigaStreams will appear in the device Port Groups page, and the created stack links will 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, right-click the device to be removed from the canvas and click **Remove**.

   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 “Manage Cluster Completed,” message will appear in the notification window.
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, right-click the device and click the **Edit Details** option.
3. Click a device in the Edit Leaf Spine Cluster canvas to display the device configuration quick view.
4. Update the leader preference text box and click **Update** to proceed.

   The Manage Cluster notification will appear to show the progress of the cluster update. When the process is complete, a “Manage Cluster Completed,” message will appear in the notification window.
5. 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:

1. Leaf nodes to the TAPs or tools.
2. First level spine nodes to the leaf nodes and the second level spine nodes.
3. Second 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.

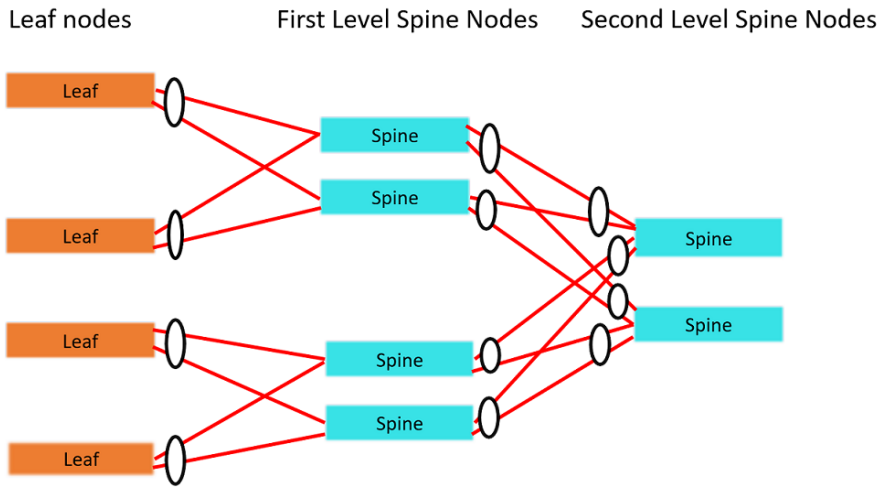

In a Spine-Leaf cluster, the number of leaf nodes is typically higher than the number of spine nodes. In the figure, there are four leaf nodes, four first level spine nodes, and two second level spine nodes. The leaf nodes aggregate to a fewer number of spine nodes.

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.

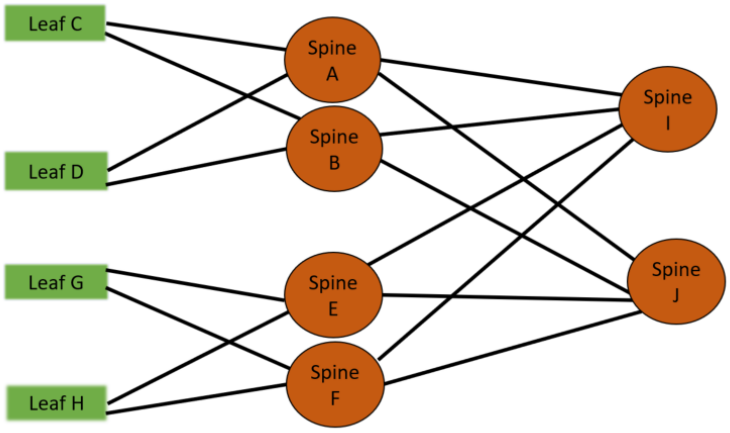

Configure Stack GigaStream

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.

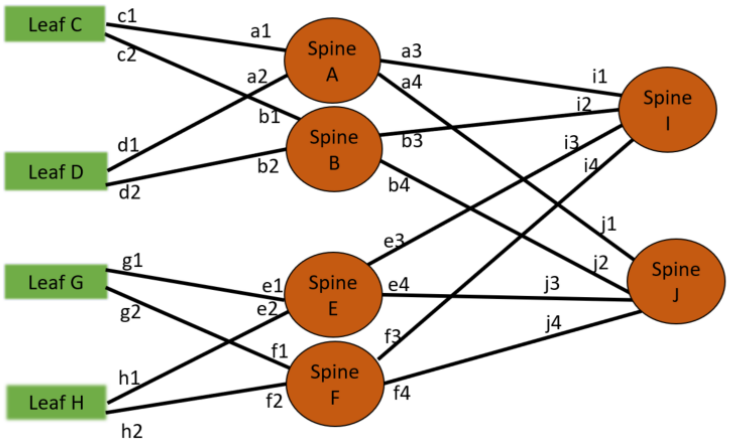

Configure Spine Links

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.

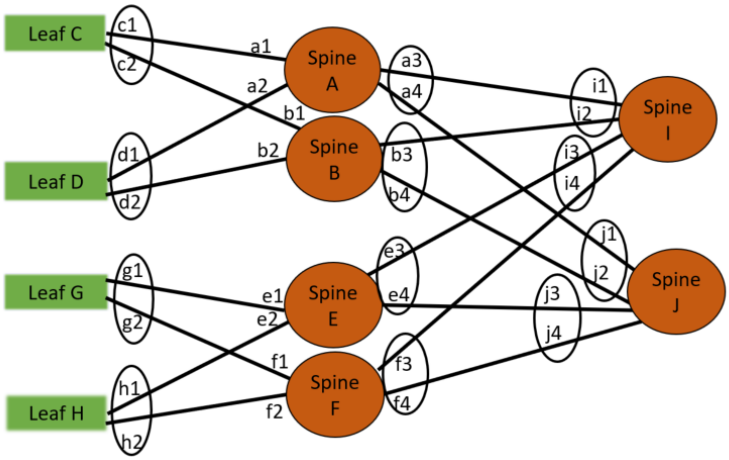

NOTE: You must not configure any spine links from spine nodes to leaf nodes. For example, in the figure, a1 and a2 should not be configured as spine links.

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.

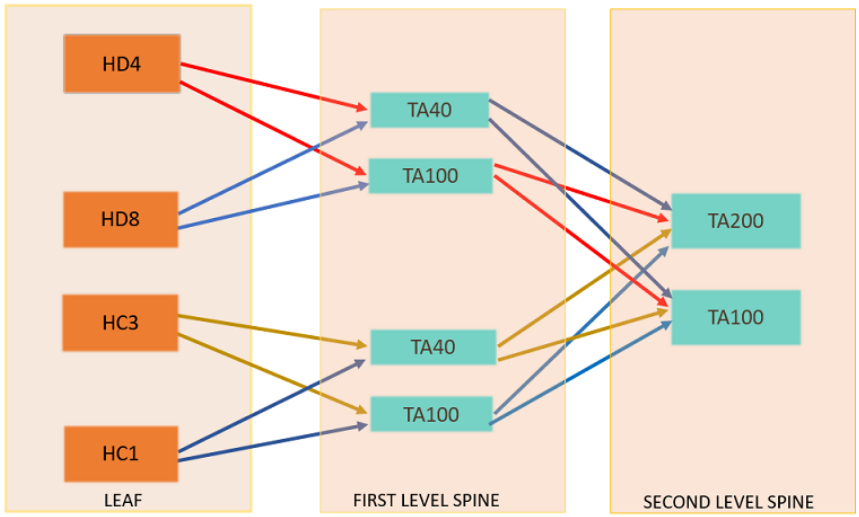

NOTE: GigaStreams support different speeds, as indicated by the different colored connector lines in the figure.

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. On the left navigation pane, click  , and then select **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	Excludes undersize frames.

Counter	Definition	Notes
	columns.	
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.

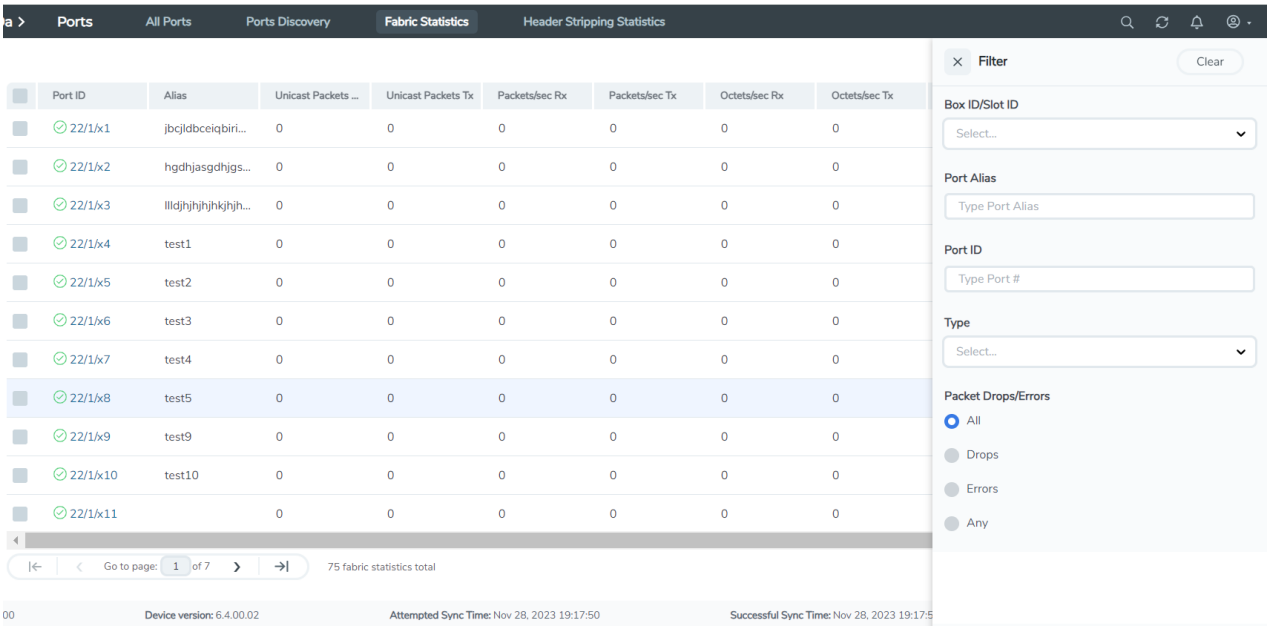

Figure 12 *Fabric Statistics Filter*

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. On the left navigation pane, click  , and then select **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. On the left navigation pane, click  , and then select **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. On the left navigation pane, click  , and then select **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 Gigamon Discovery

Following are the limitations of 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 disabled:

  (config) # ! GDP is not supported with port type inline-too. 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.

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, click on 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.

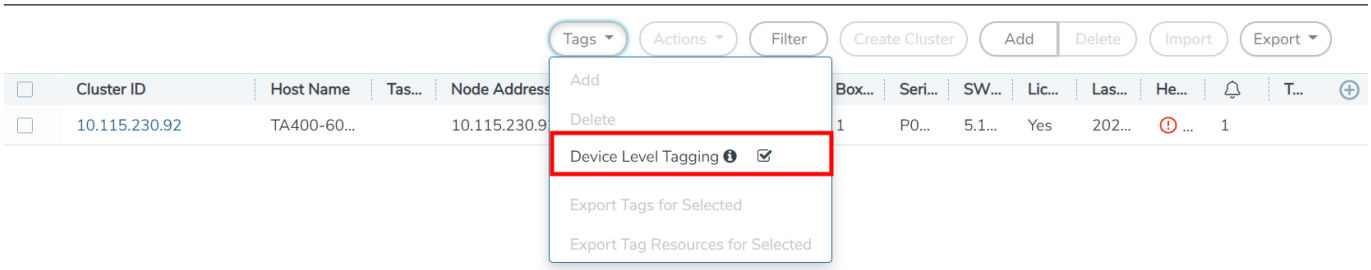

The nodes in a cluster are now associated to tags which can be used for creating the required topology.

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.

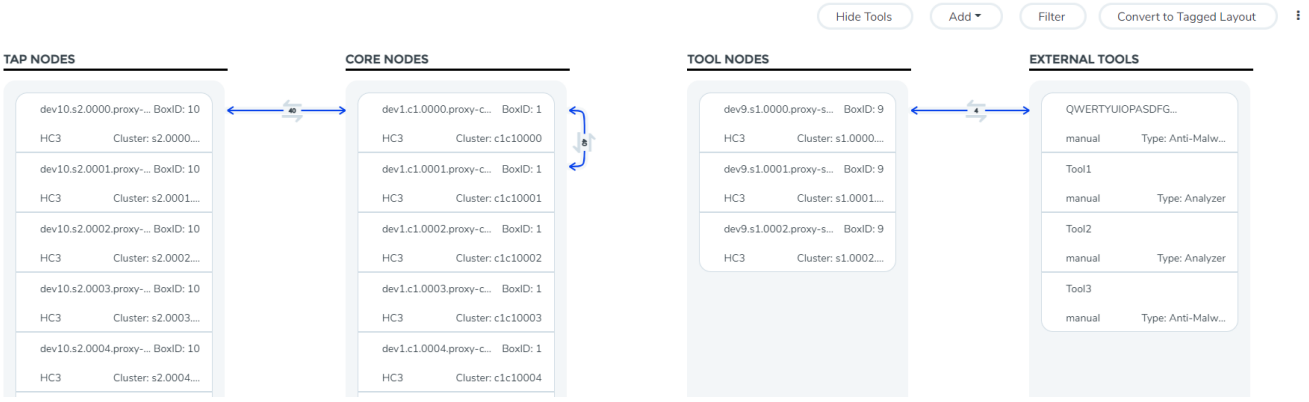

- Click on a link or a device. The **Connections Nodes and Analytics** pane is displayed at the bottom. Refer to Smart Layout Topology section for details.

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, click on  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.

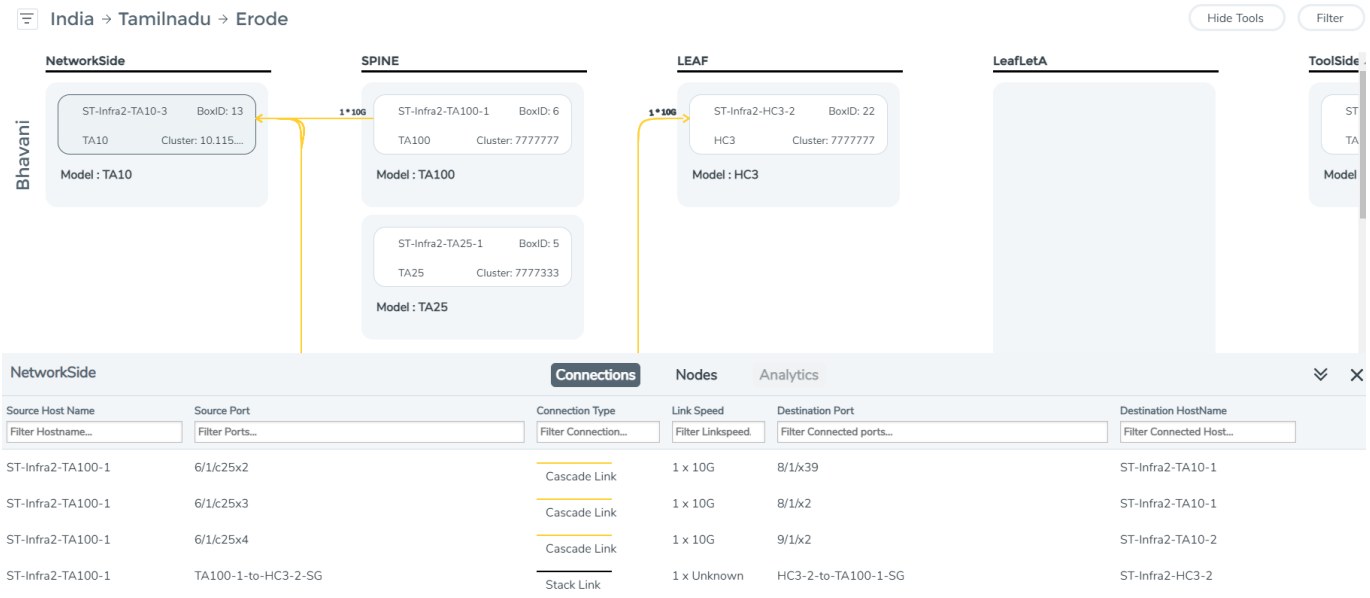

How to Filter

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

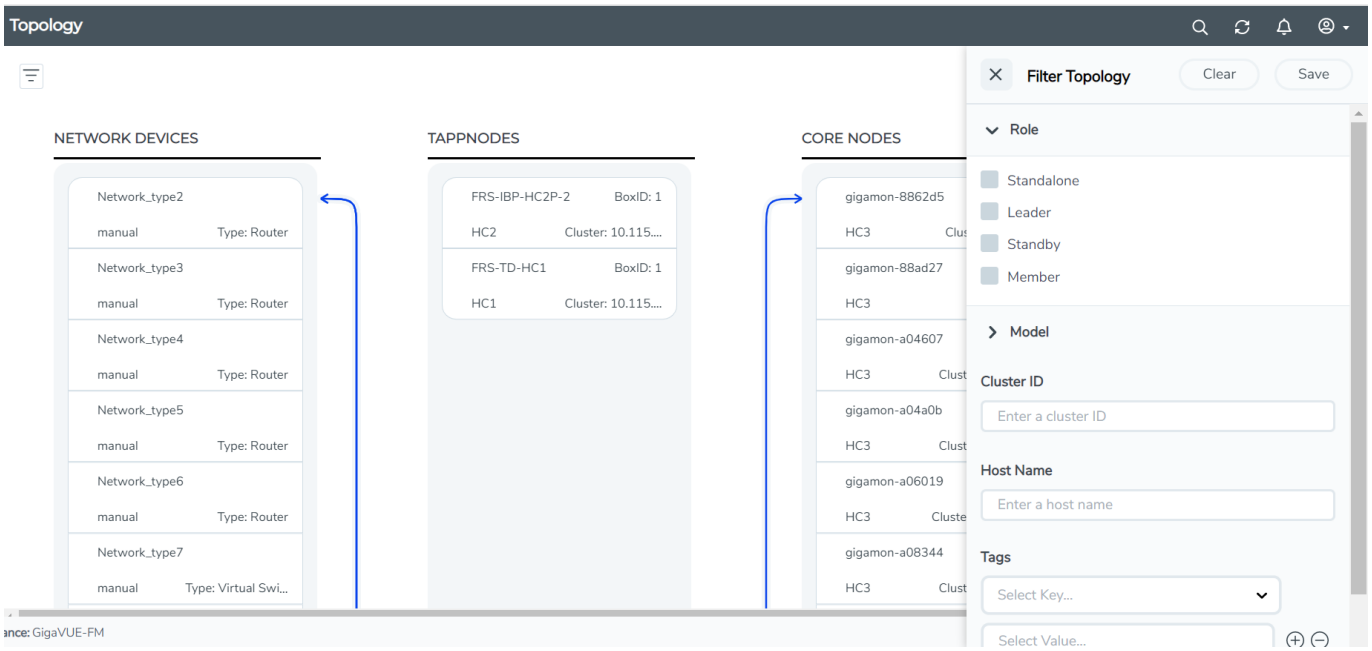

The selected filter criteria appear as info panes on the topology page and the filtered node(s) will get highlighted. Click **Save** to save the filter. Click Clear to clear the filter.

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.

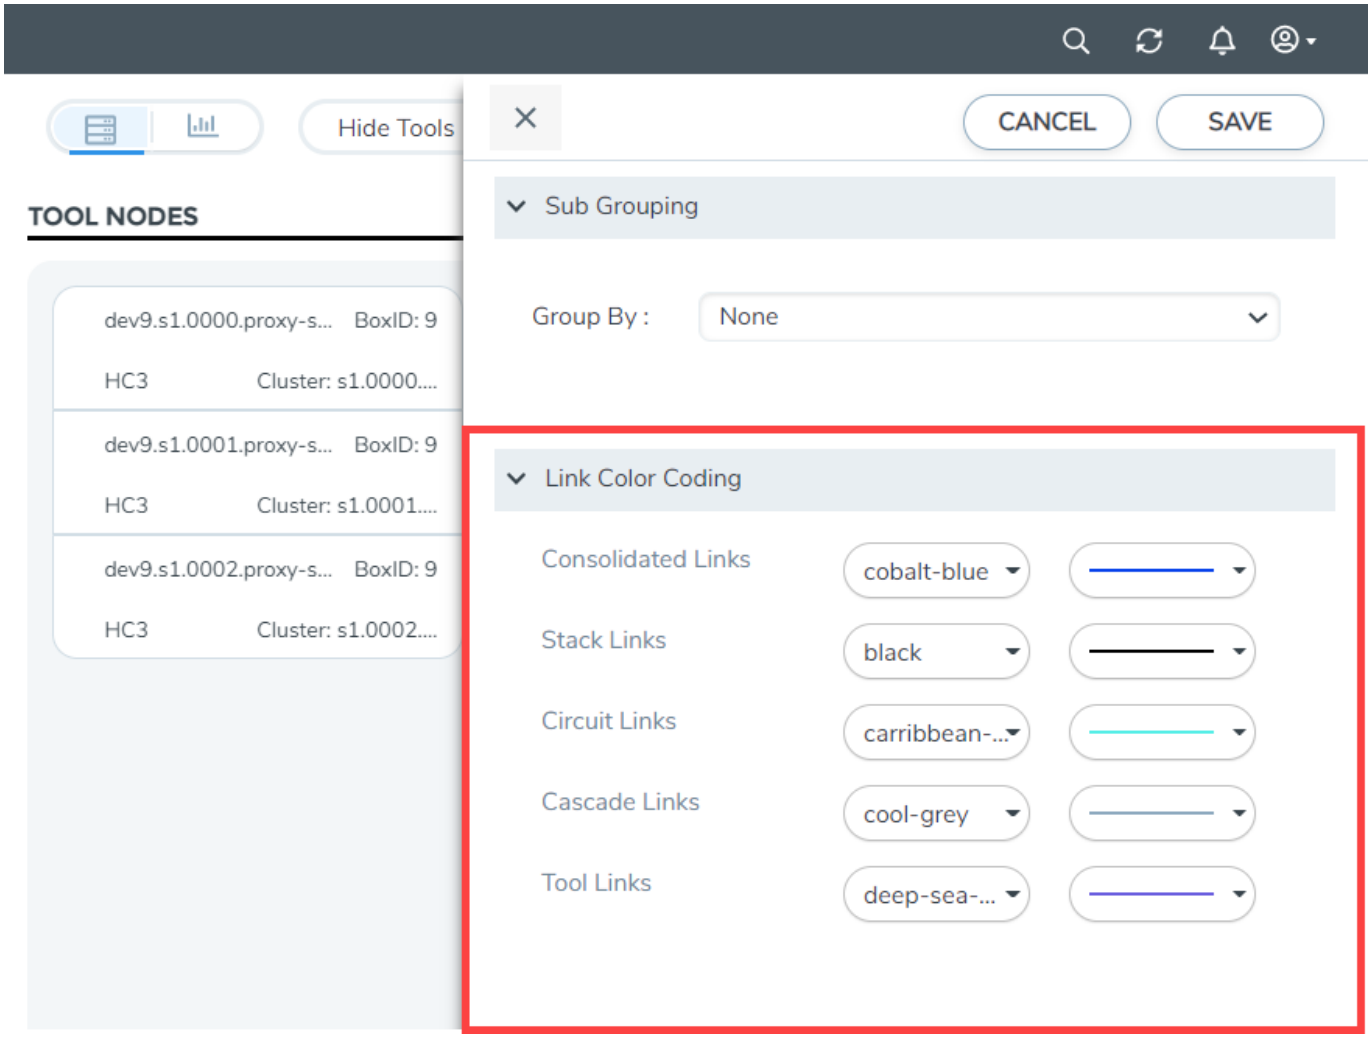

Sub-Grouping of Nodes

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.

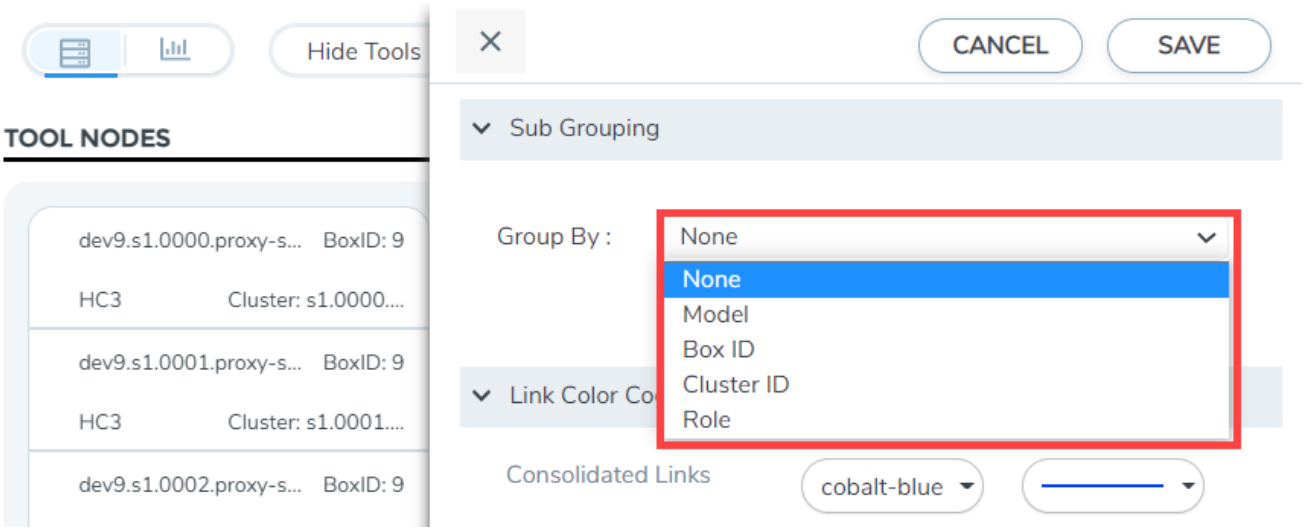

Based on the current user who has logged in, the sub-grouping is saved.

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:

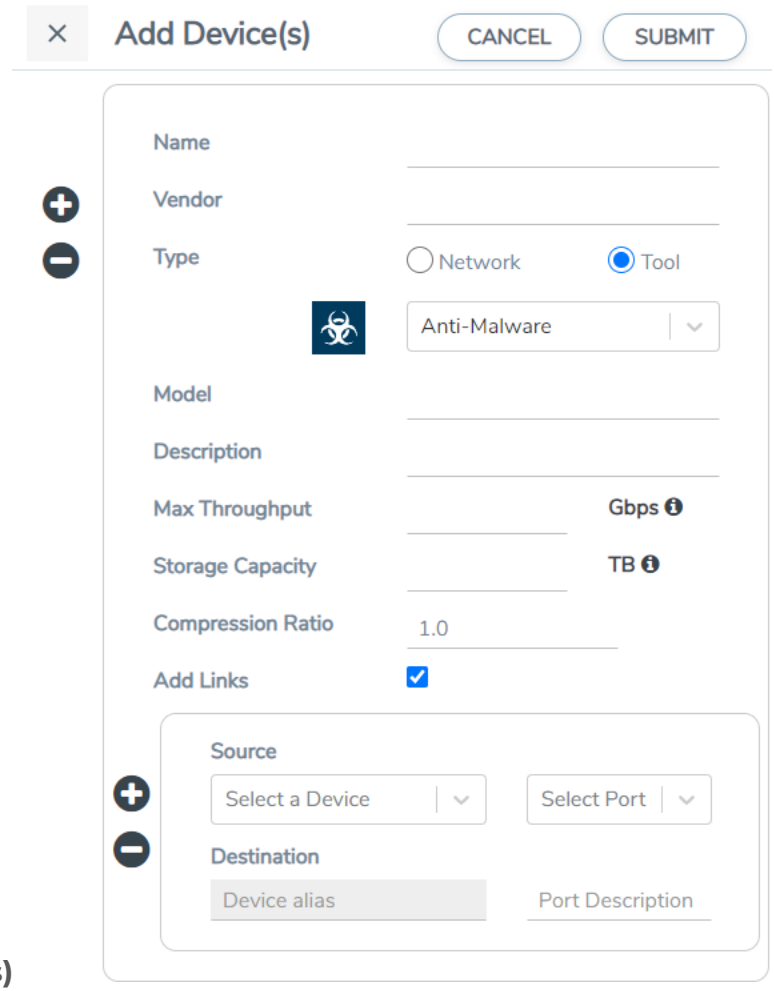

1. Select **Add > Add Device(s)**.
2. You can add additional nodes by clicking the **+** button. To remove a node, click the **-** button.
3. Enter the following information for the node:
    - **Name**—the name of the node.
    - **Vendor**—the node's manufacturer. For Gigamon nodes, you cannot edit 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:

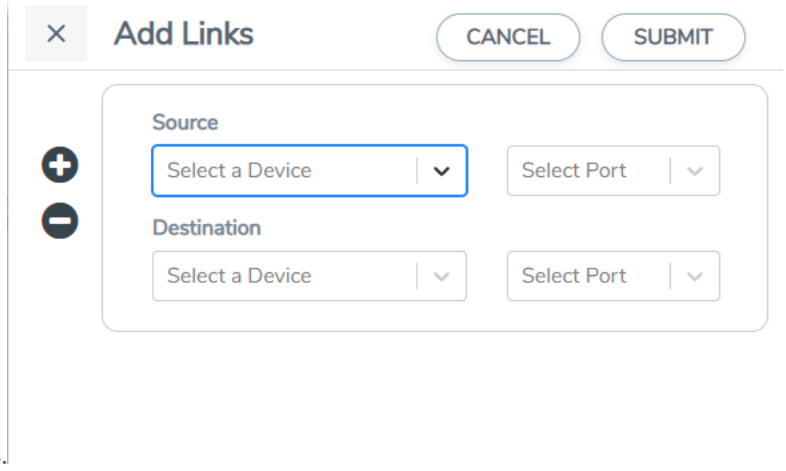

1. Click Add Links.
2. Select the **Source** device and the required source port from the drop-down list.
3. Select the **Destination** device and the required destination port from the drop-down list.

   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, click on 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, click on  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, click on  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.
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. On the left navigation pane, click  and select **System** > **Preferences**.
2. Click **Edit**.
3. On the Edit Preferences page, under **Flows**, change the status to Enabled.

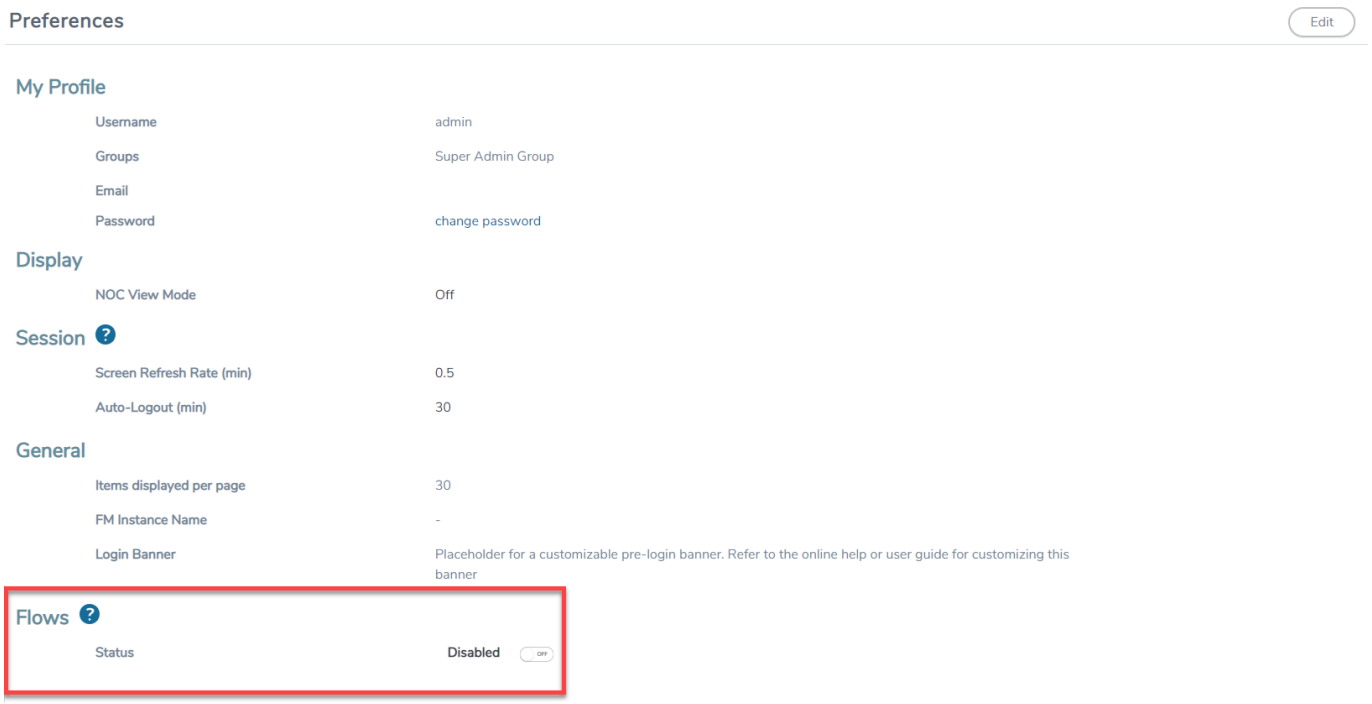

NOTE: You cannot disable flows once it is enabled. Contact Customer Support to disable Flows.

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.

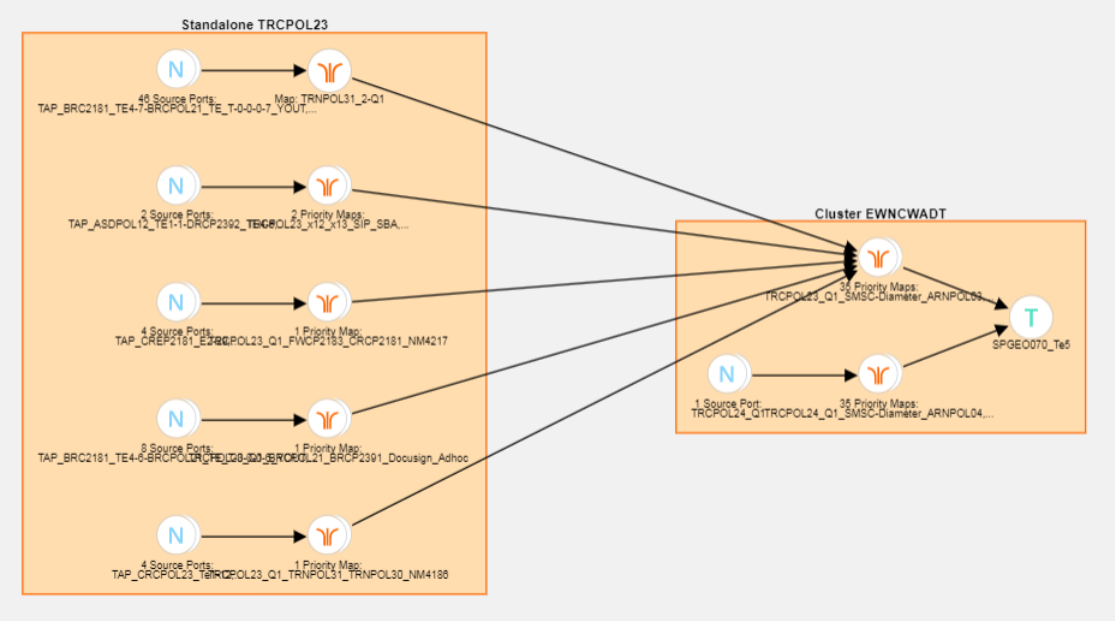

Figure 13 *Flow - All Sites*

Figure 14Standalone GigaVUE Nodes shows a number of GigaVUE nodes in the Topology page that are not connected to each other.

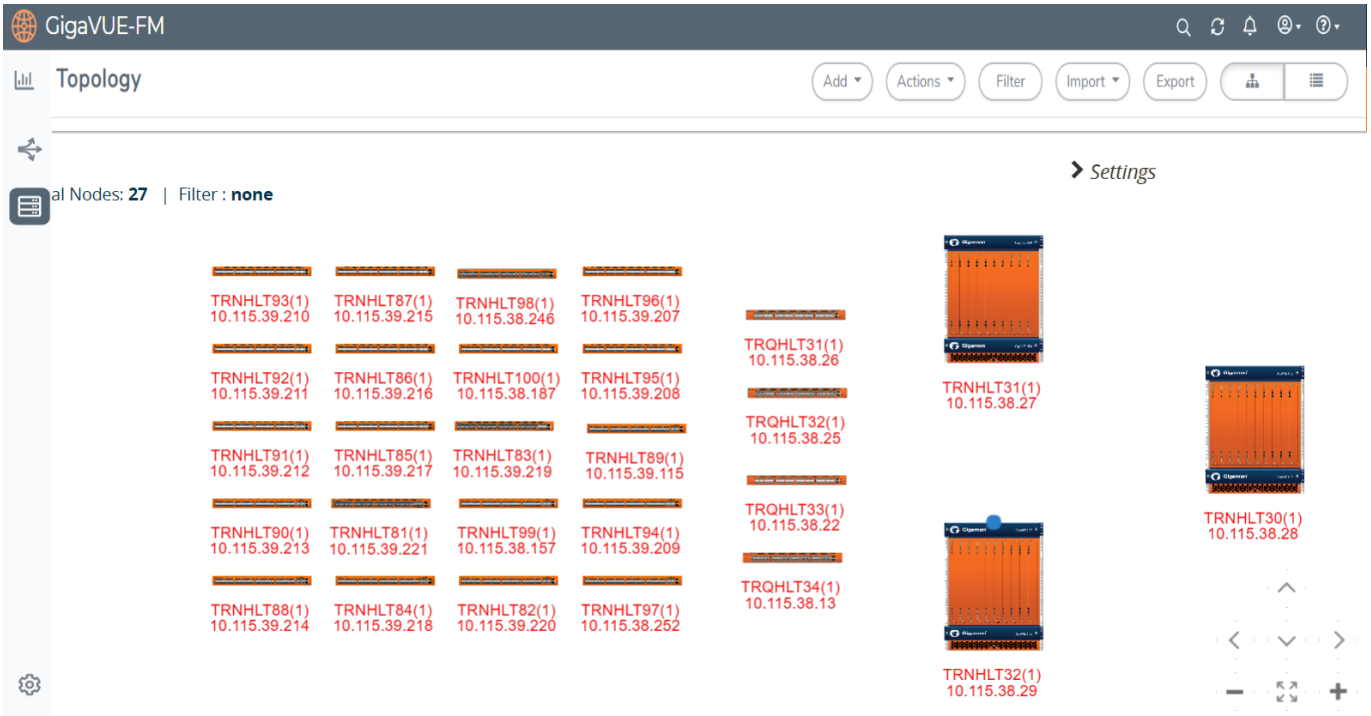

Figure 14 *Standalone GigaVUE Nodes*

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&Te2	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.

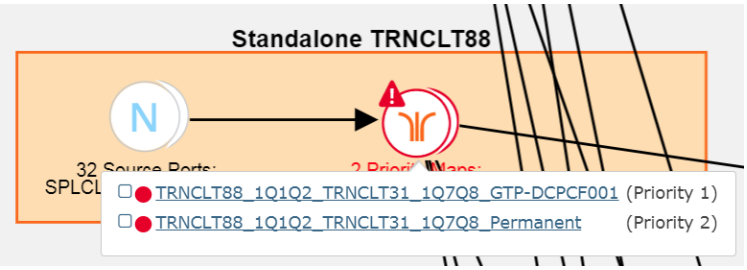

Figure 17 *Unhealthy Maps*

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.
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.

Option	Description
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.

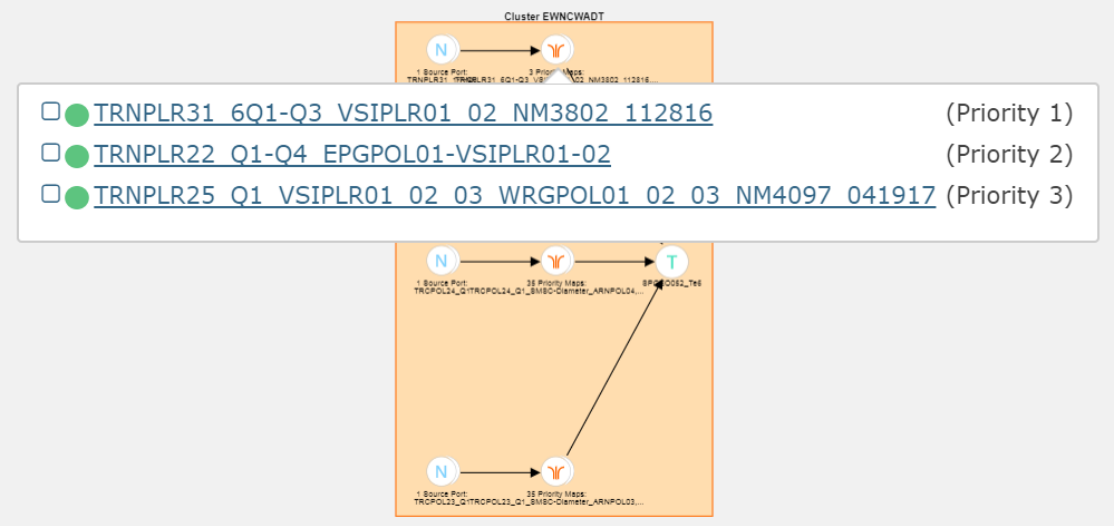

Figure 18 *Selecting Maps From Maps List*

   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.

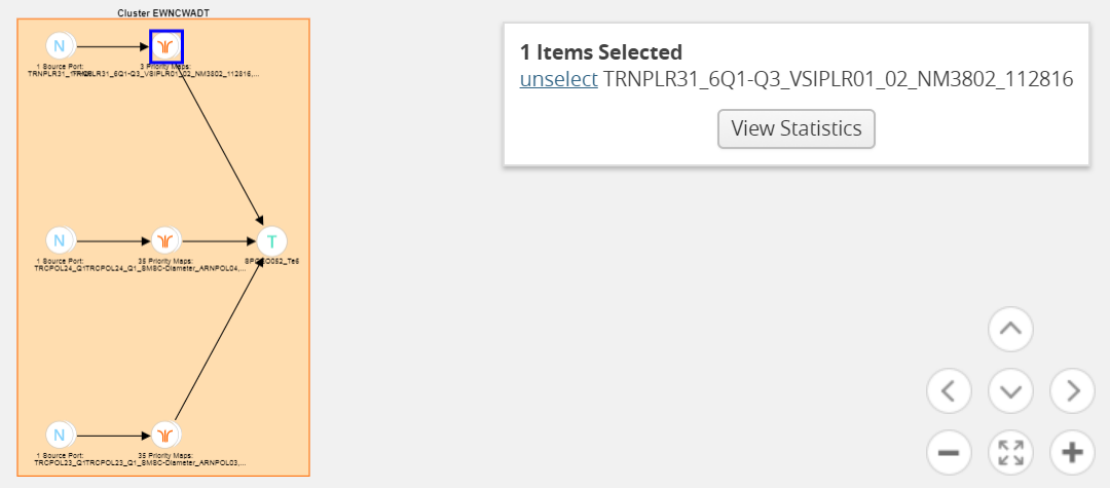

Figure 19 *Selected Maps in Items Selected Pop-up*

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.

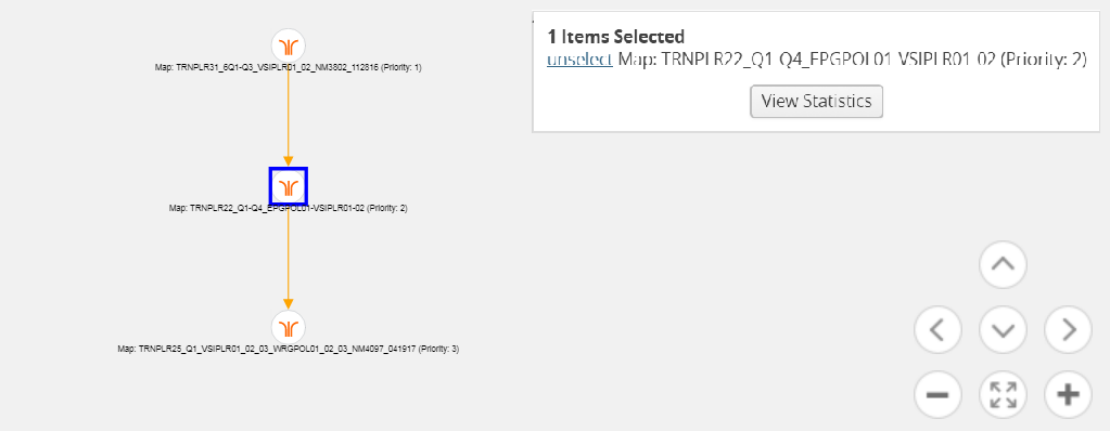

Figure 21 *Maps Selected*

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.

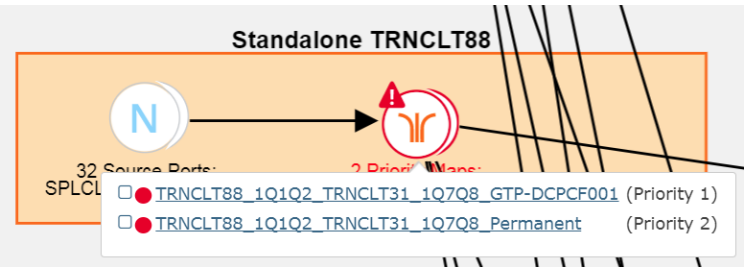

Figure 22 *Unhealthy priority Maps*

Figure 23Unhealthy Pass-all Map shows how a pass-all map is illustrated in a flow view page.

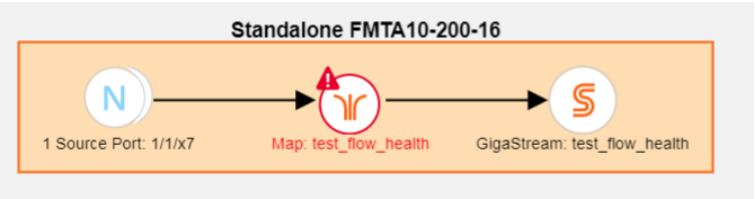

Figure 23 *Unhealthy Pass-all Map*

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.

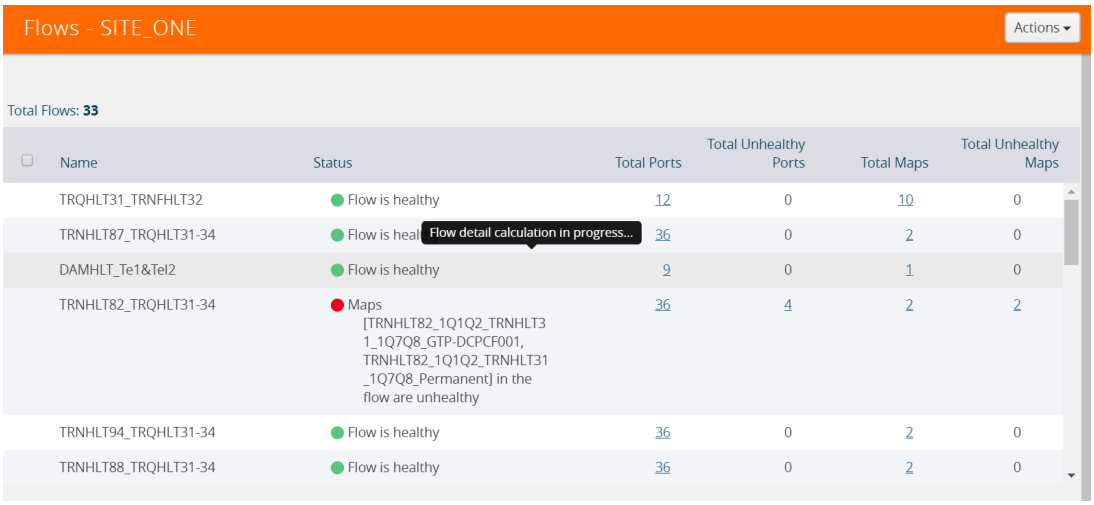

View Events

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.

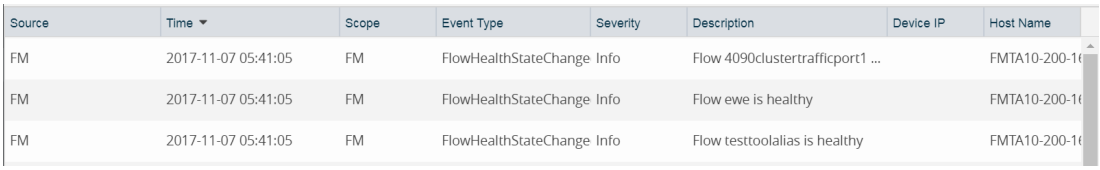

Set Notifications

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.

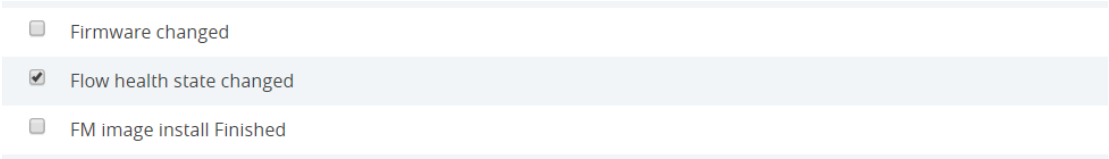

Figure 24 *Flow Health State Changed Notification*

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, click on 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**.

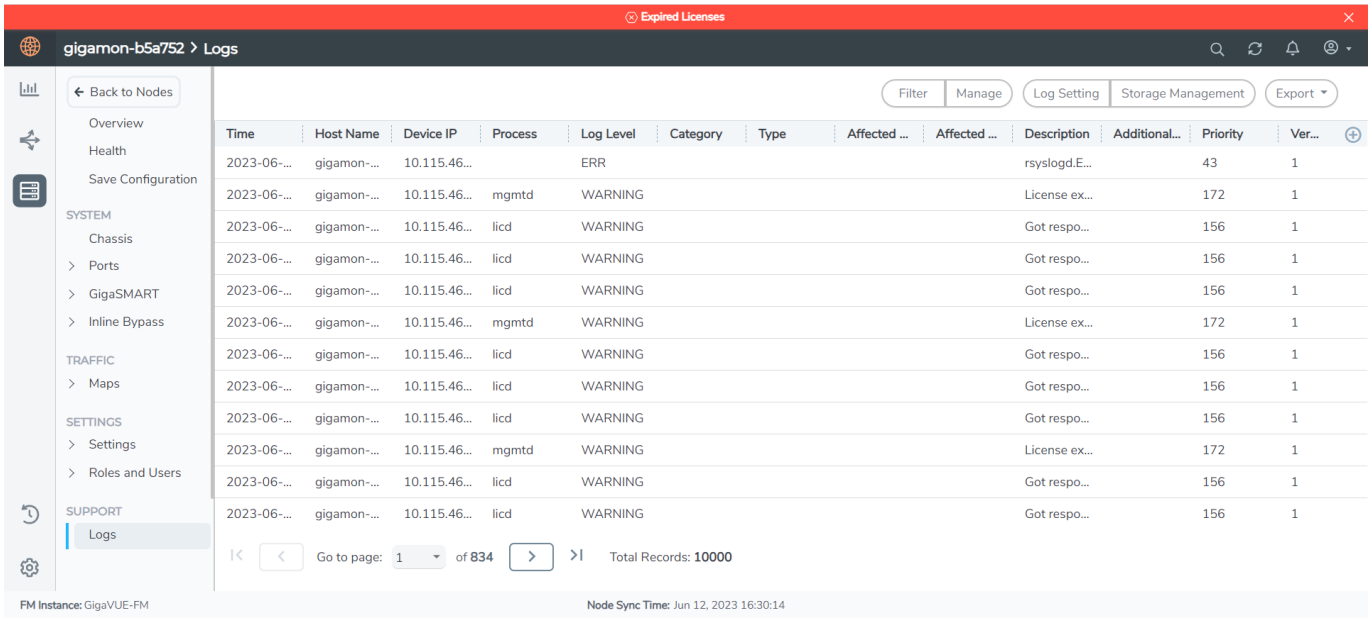

The Logs view displays the latest log information for the selected GigaVUE-OS node in table format.

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**.

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, click on 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, click on  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 more details.
From the device logs page:	1. On the left navigation pane, click  , and then select **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.

Path	Steps
	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, click on  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 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:

Setting		Description
		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.

On the left navigation pane, click on  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.

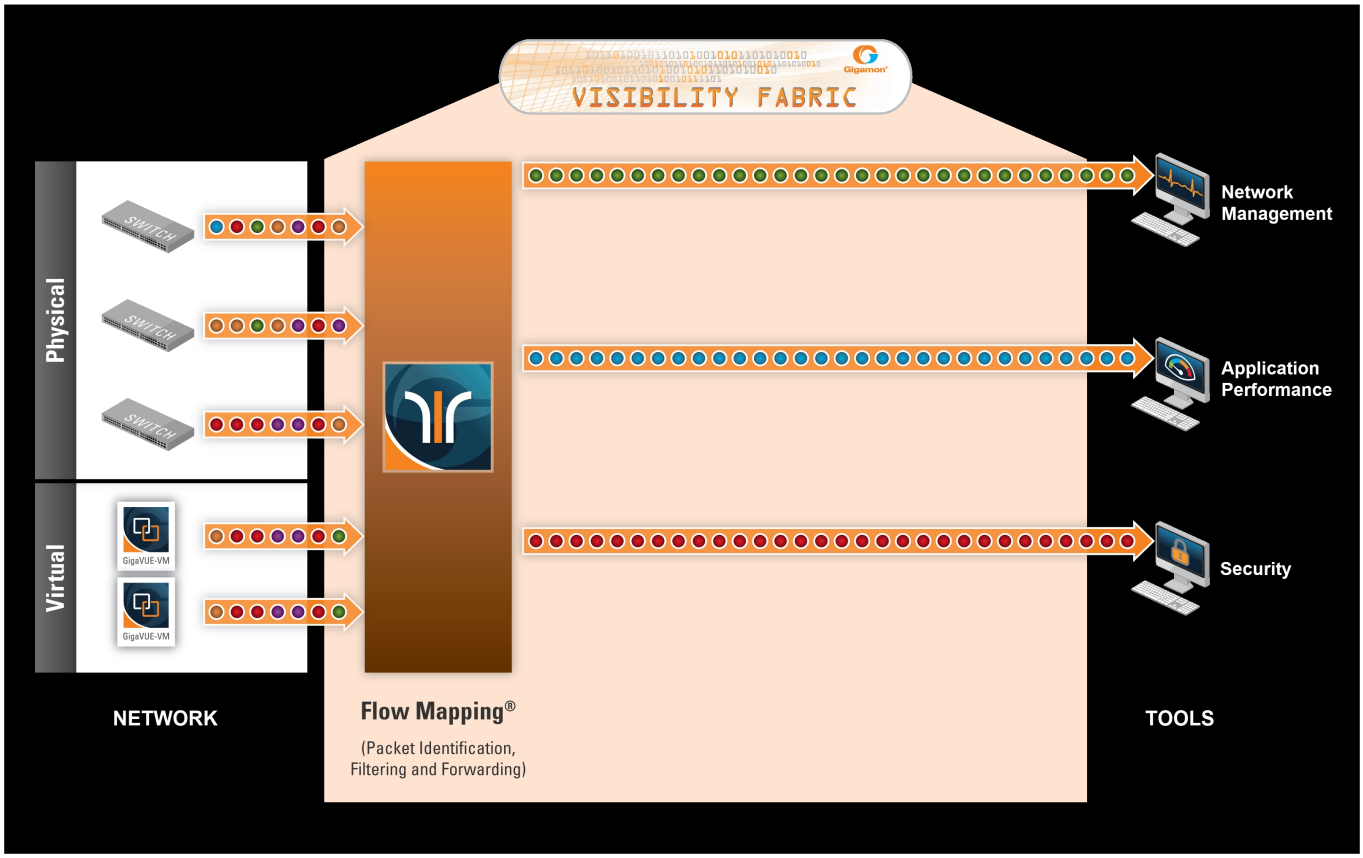

Figure 1 *GigaVUE-OS Packet Distribution*

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.

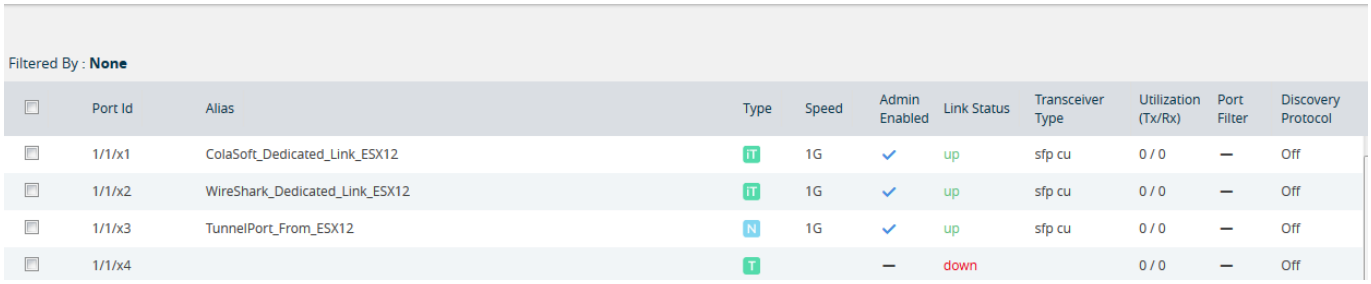

Filtered By : **None**

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.

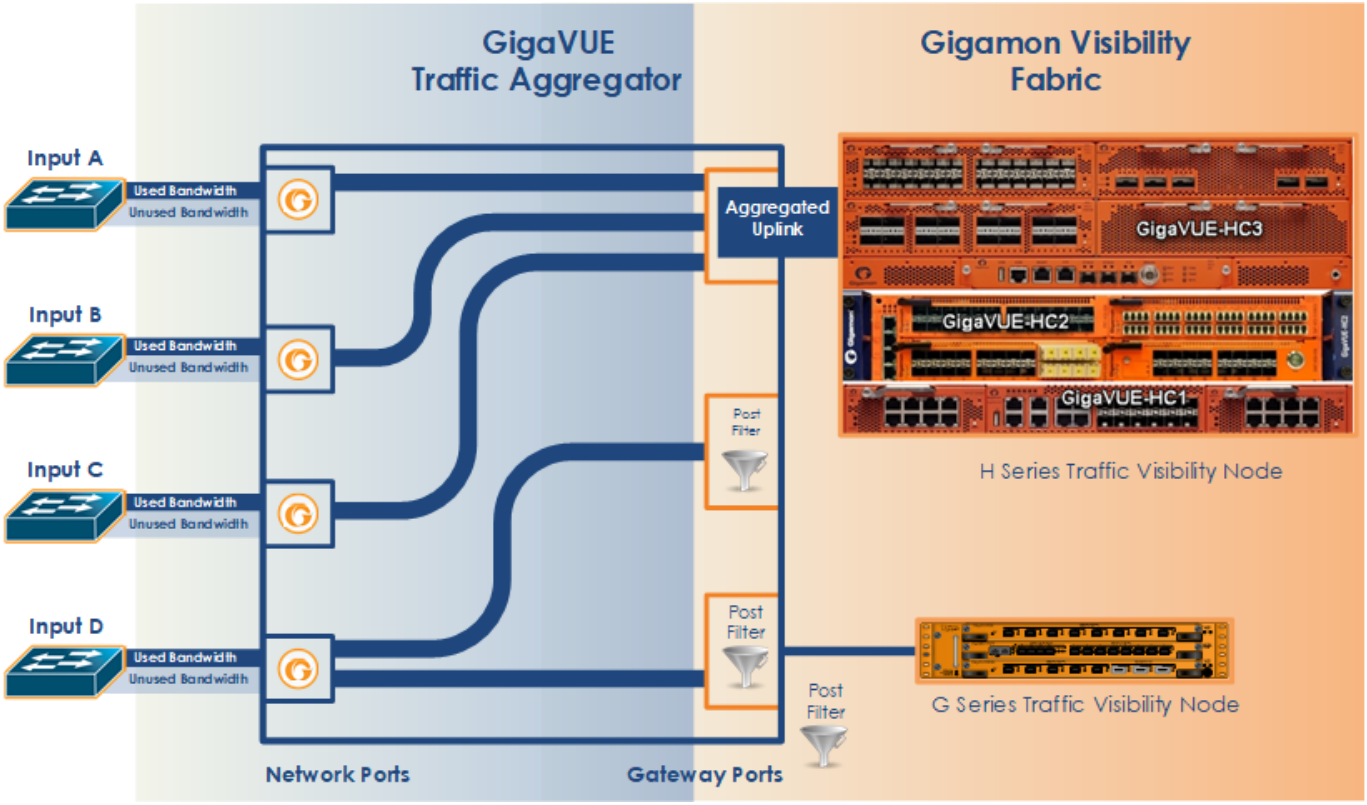

Figure 3 *GigaVUE-TA Packet Distribution*

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, and GigaVUE-TA200.

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.

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.

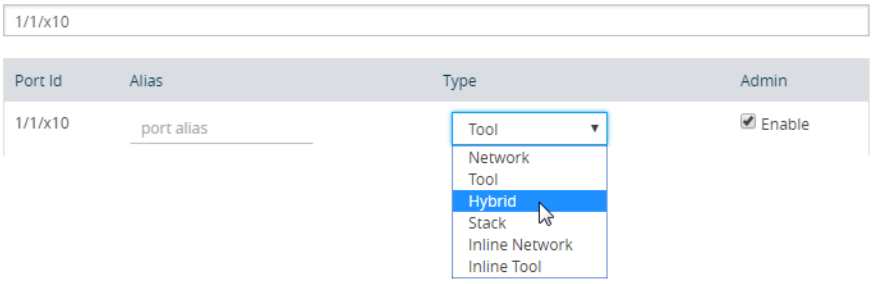

4. Click **Save**.

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 and GigaVUE-TA200E.
- 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.
- In the release 5.14 on GigaVUE-TA400, the outer VLAN tool port filter cannot be used to match ingress VLAN tag that is configured on the source port.
- 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). 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 GS 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 **GigaVUE-HC3 (CCv1 and CCv2)**	• 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-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. **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 Node	Maximum Number of Port-Filters
GigaVUE-TA200	
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 chassis **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
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

Rule	Action
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:

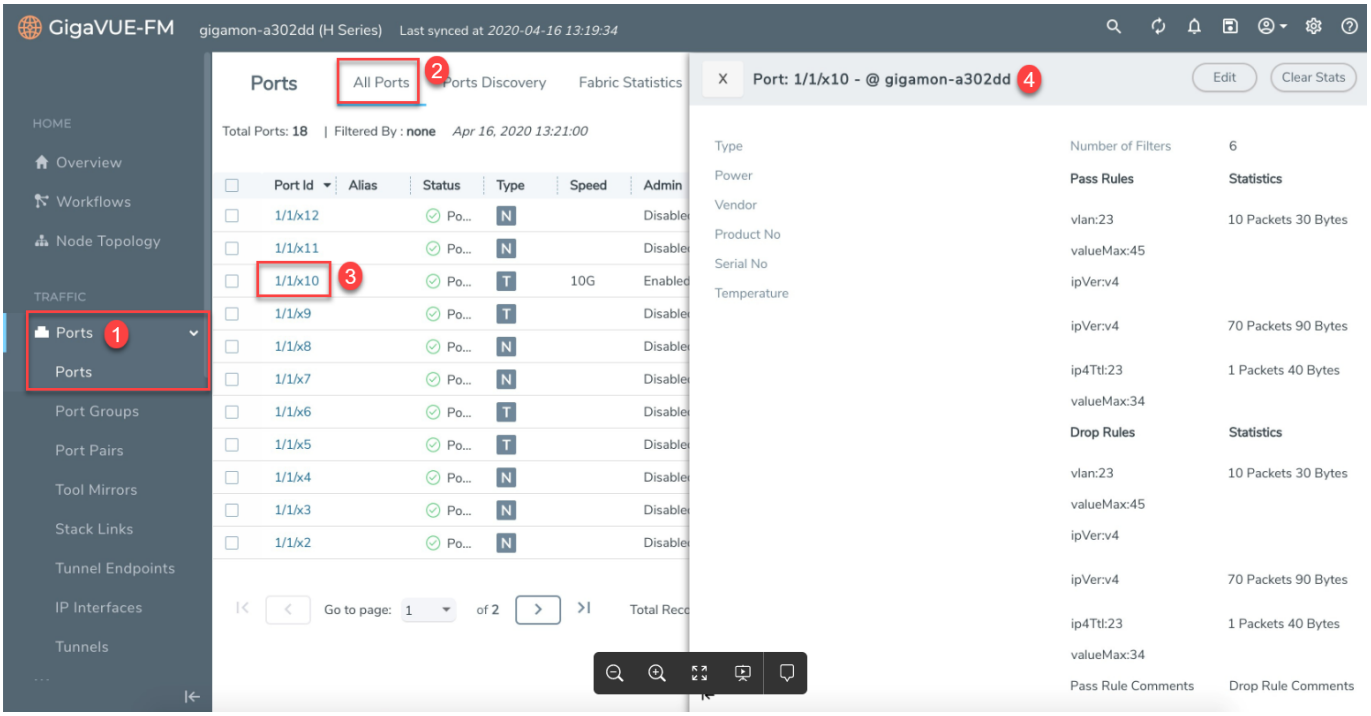

View Filter Resources for a Slot

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:

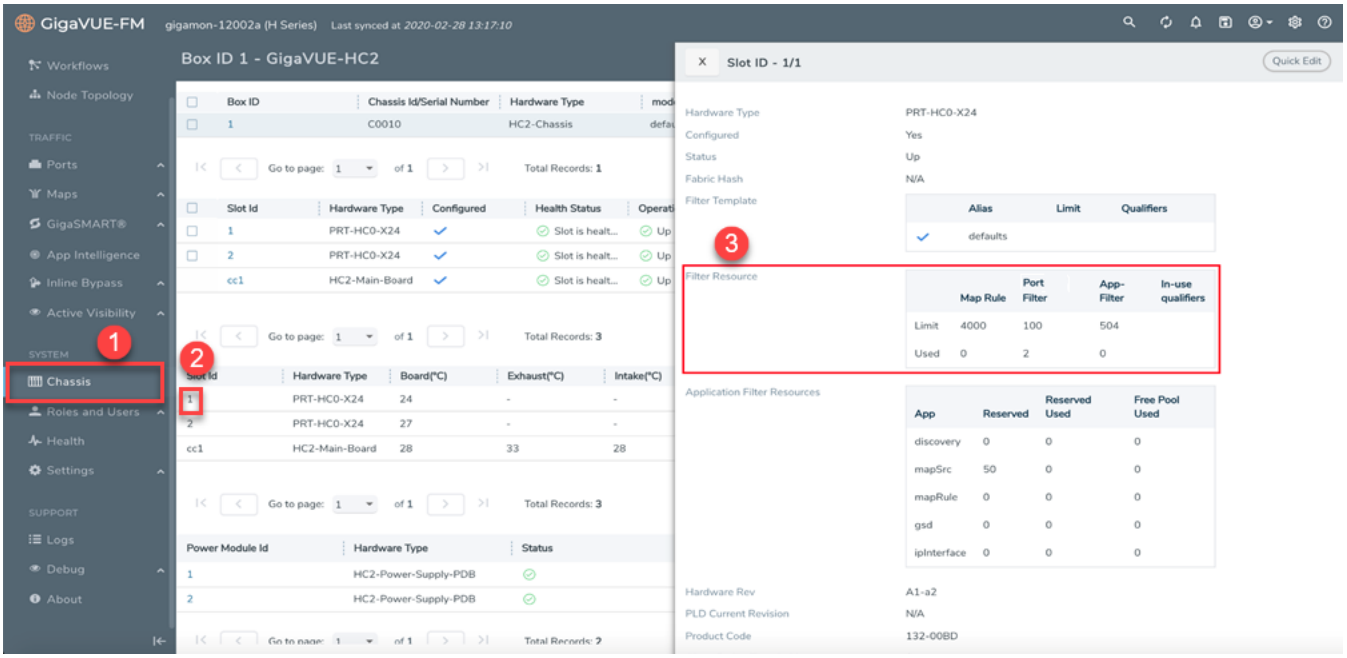

Status of Line Cards/Nodes and Ports

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*.

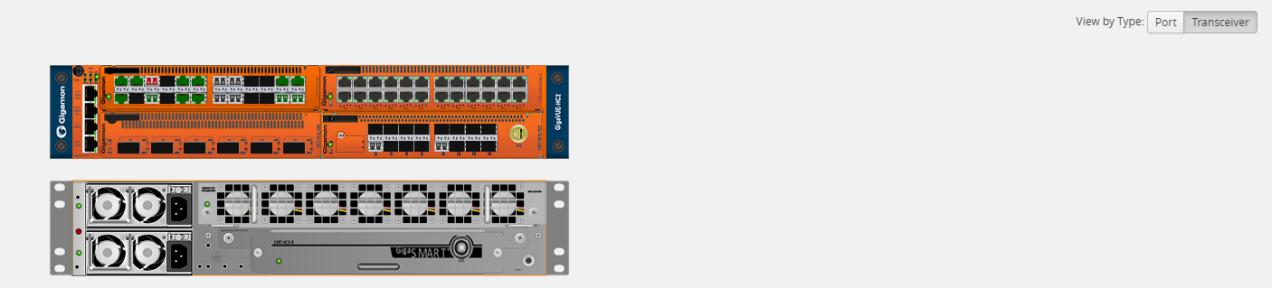

Figure 4 *Chassis Page*

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.

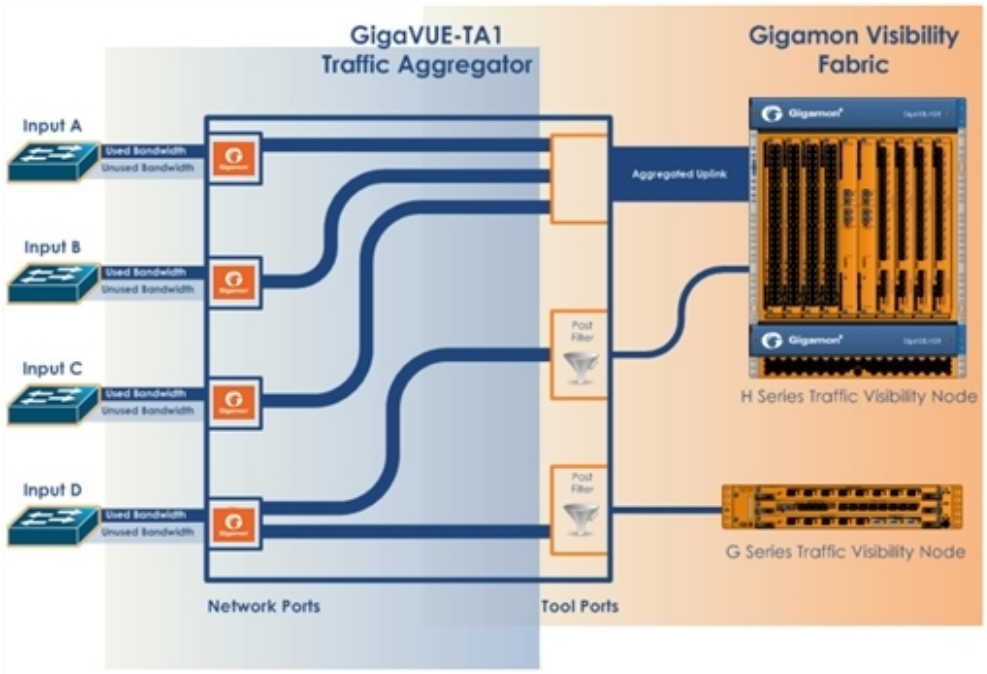

Figure 5 Ports Page

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.

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. • **CL74**— Enables FEC on the port . • **CL108** — Enables FEC on the port for GigaVUE-TA25 and GigaVUE-TA25E. • **OFF**—Disables FEC on the port. This option is available only on 25Gb and 100Gb transceivers.
Link Status	The current status of the link connected to the port, either port link up or port link down.

Column	Description
Transceiver Type	The type of transceiver installed in this port.
SFP Power	SFP power for copper transceivers **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.

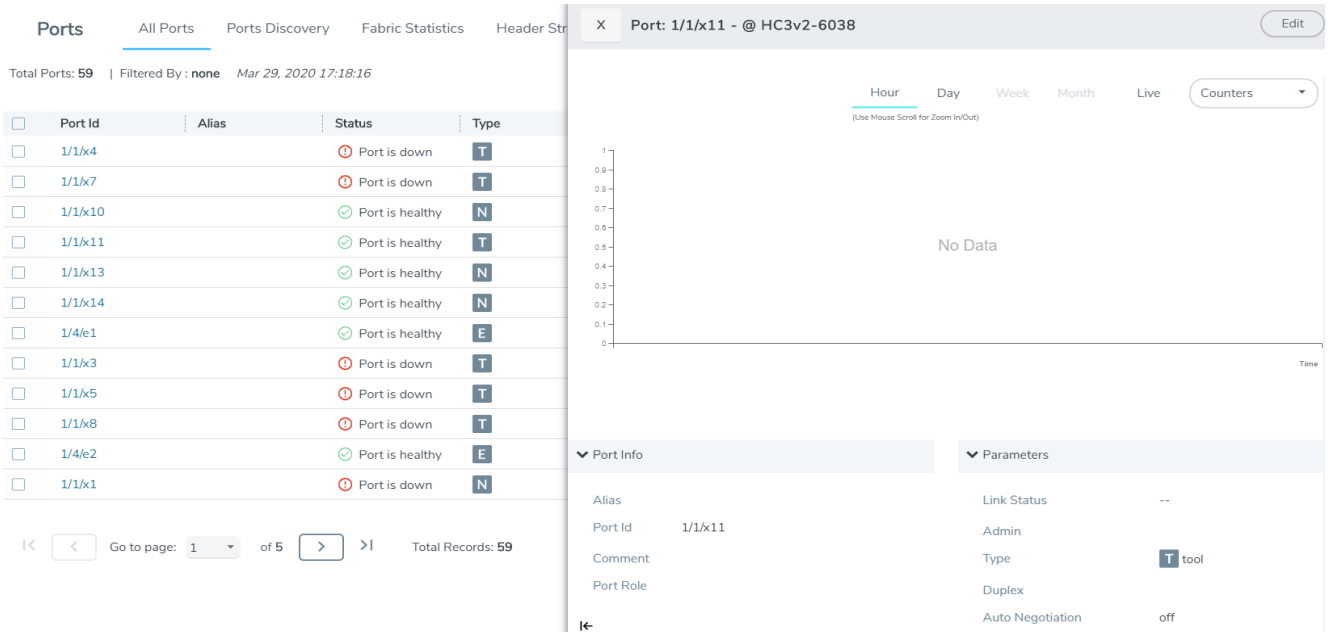

Figure 6 *Ports Quick View*

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.

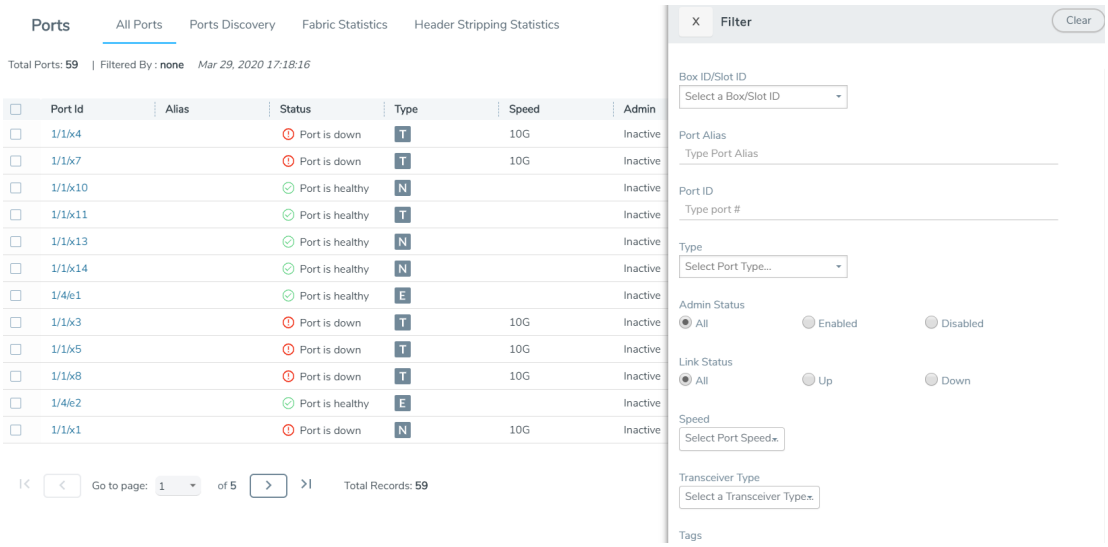

Figure 7 *Port List Filter*

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

Criteria	Description
	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. **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.

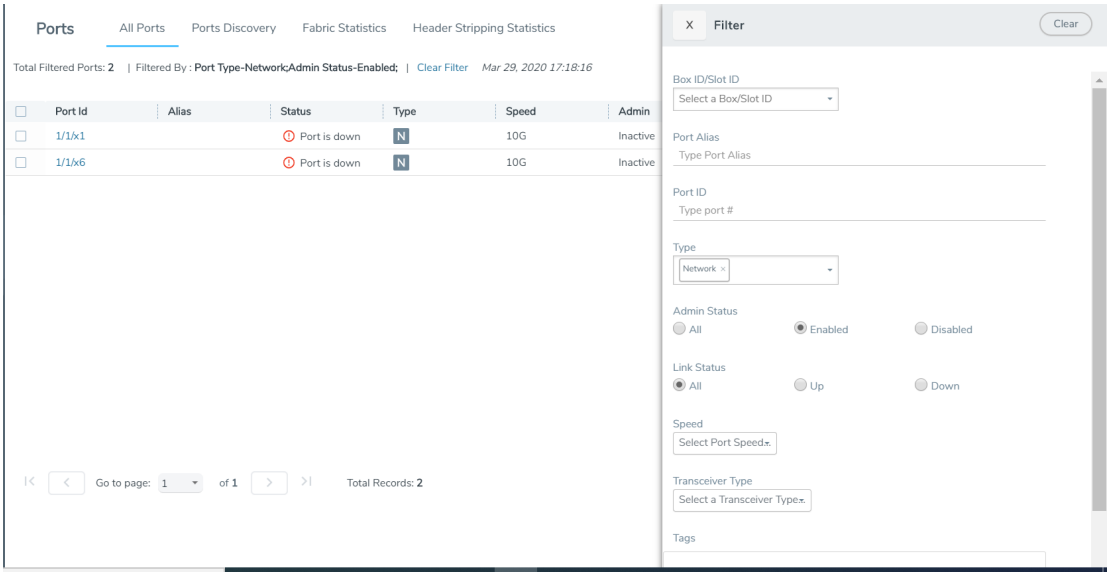

Figure 8 *Filtering by Network Port Type and Admin Status Enabled*

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.

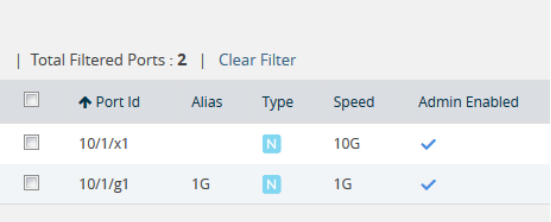

Figure 9 *Filtered Ports List*

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.

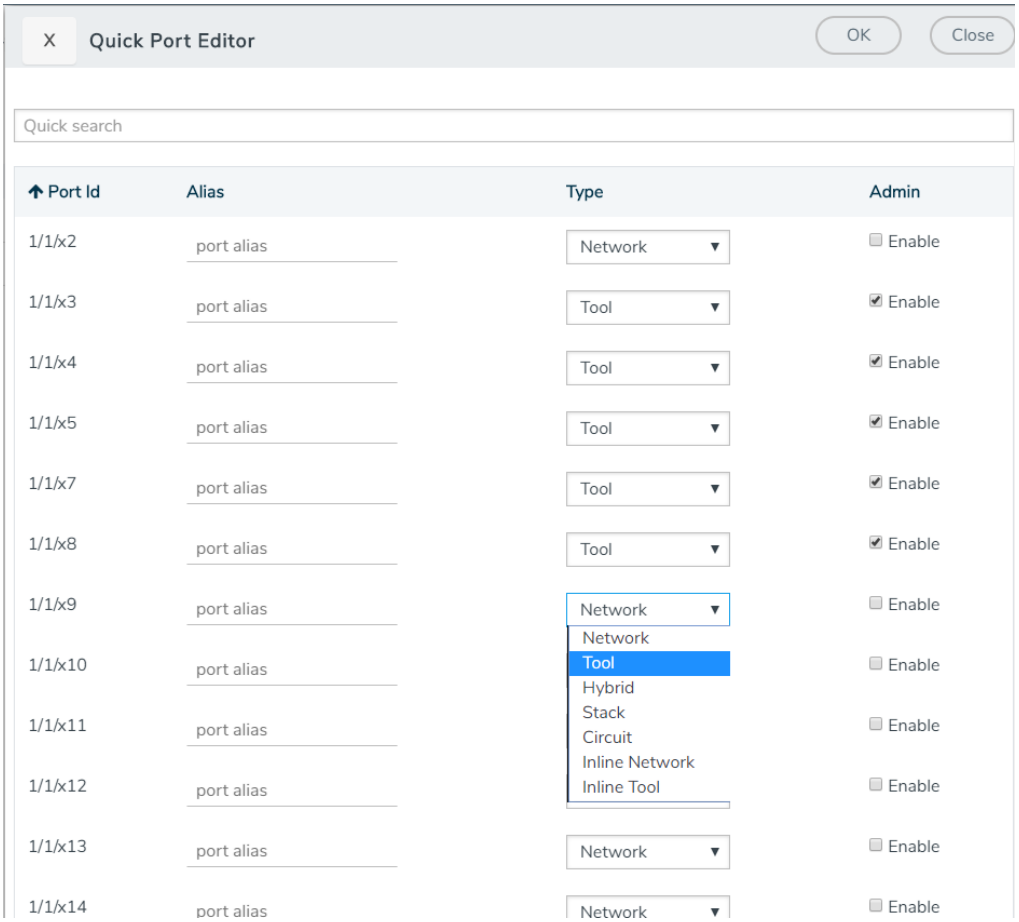

Figure 10 *Port Type Selection*

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 and GigaVUE-TA400 platforms with 100Gb BiDi (QSB-512) transceiver. If

Field	Description
	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 . **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

Field	Description
	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 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 .
   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.**
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 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**. **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.

Button	Description
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

Button	Description
	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 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. On the left navigation pane, click  , and then select **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.
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 and GigaVUE-TA200E

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. On the left navigation pane, click  , and then select **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.

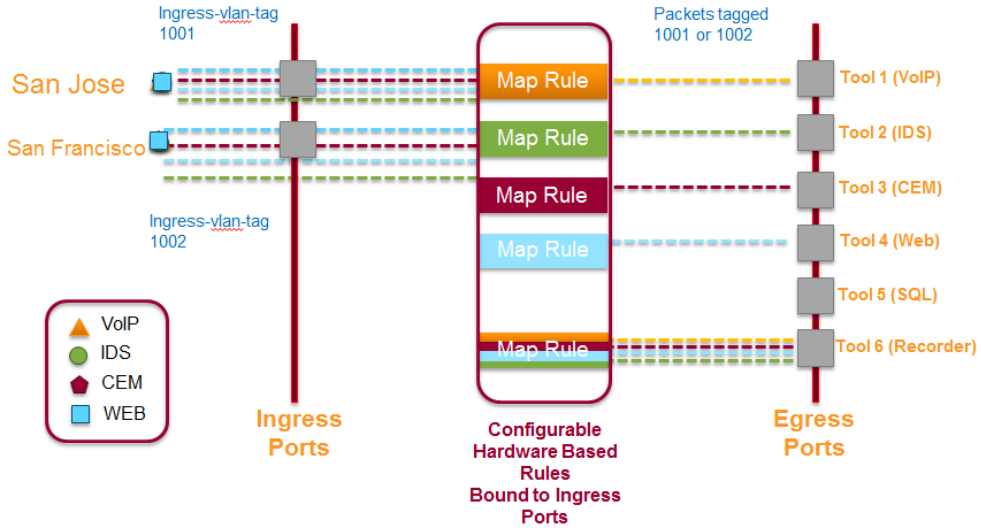

Figure 11 *Using Ingress Port VLAN Tagging*

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. From the left navigation pane, 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**.

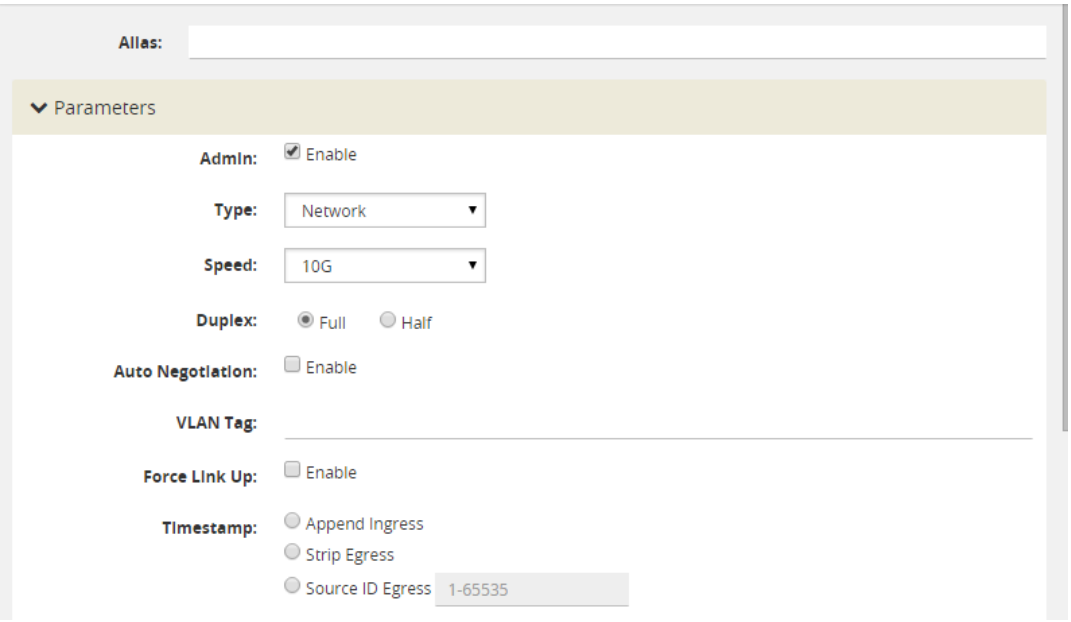

Figure 12 *Port-ID Configuration*

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).

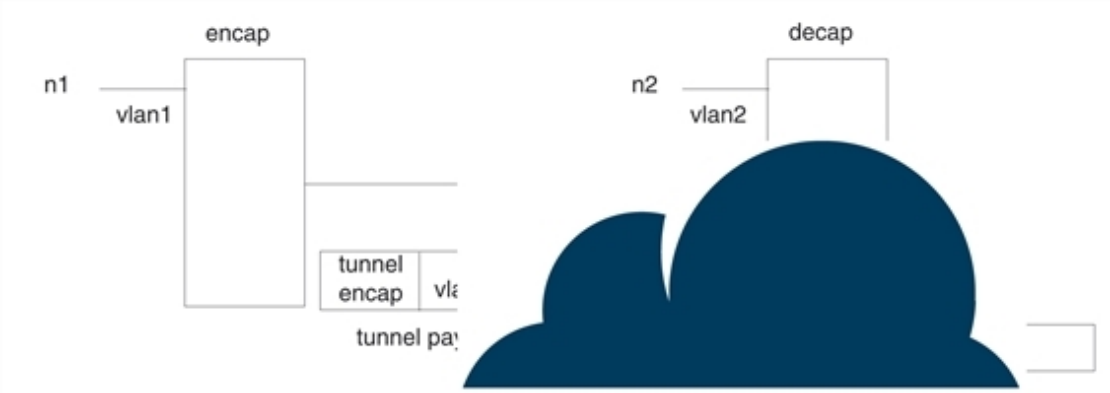

Figure 13 *IP Interfaces*

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.

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.

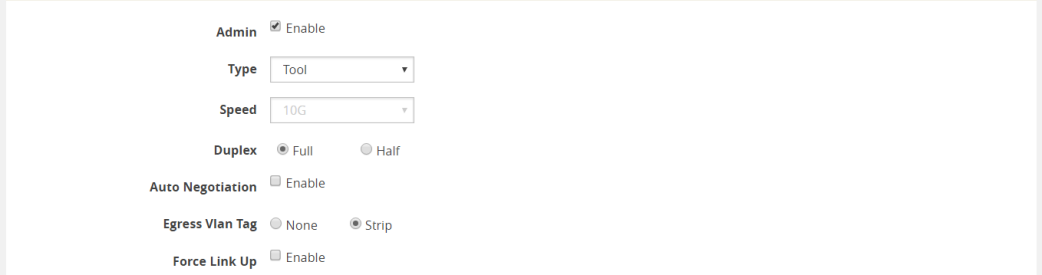

Figure 14 *Select Strip Egress Vlan Tag*

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. On the left navigation pane, click  , and then select **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 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.

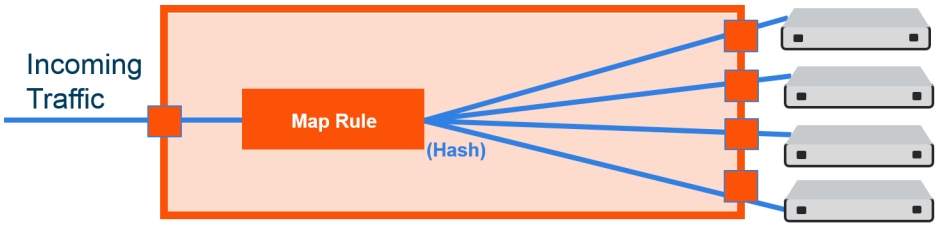

Figure 15 *Regular GigaStream Overview*

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.

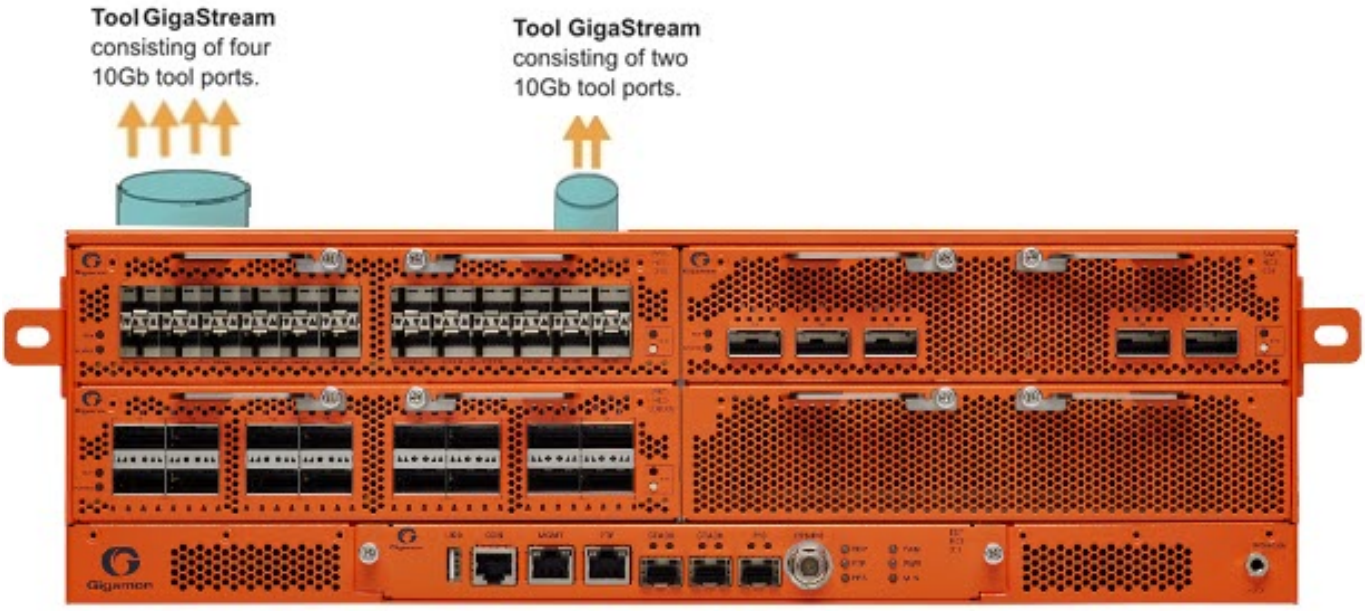

Figure 16 *Regular Tool GigaStream Illustrated*

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.

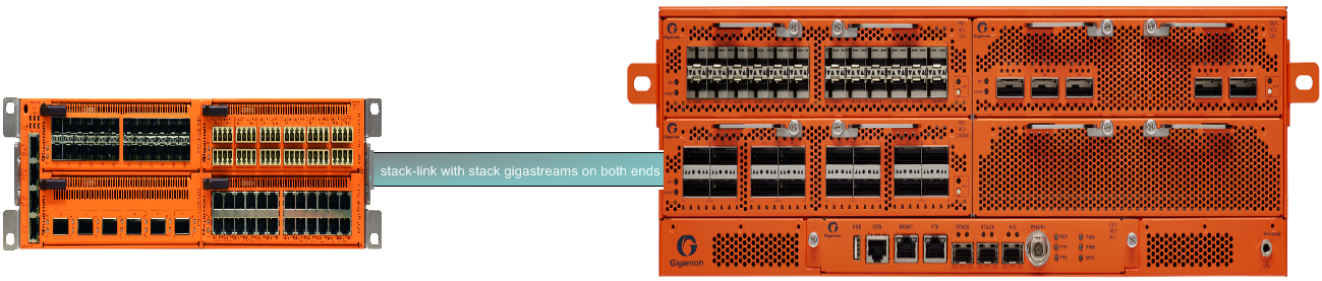

Figure 17 *Regular Stack GigaStream Illustrated*

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. On the left navigation pane, click  , and then select **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. On the left navigation pane, click  , and then select **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.

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.

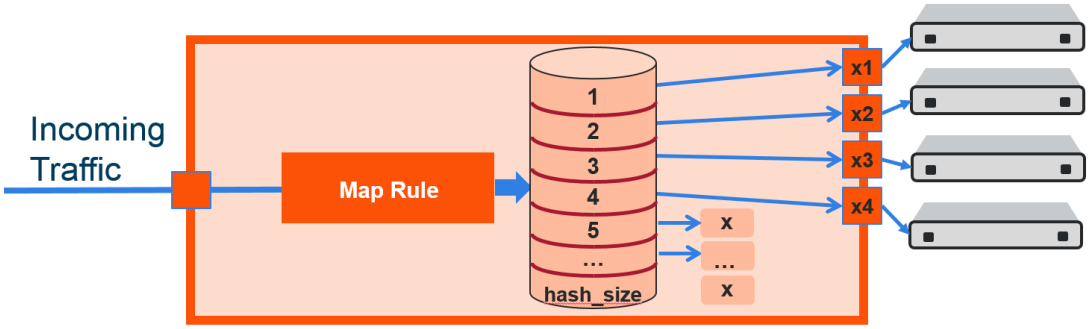

Figure 18 *Controlled GigaStream Overview*

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.

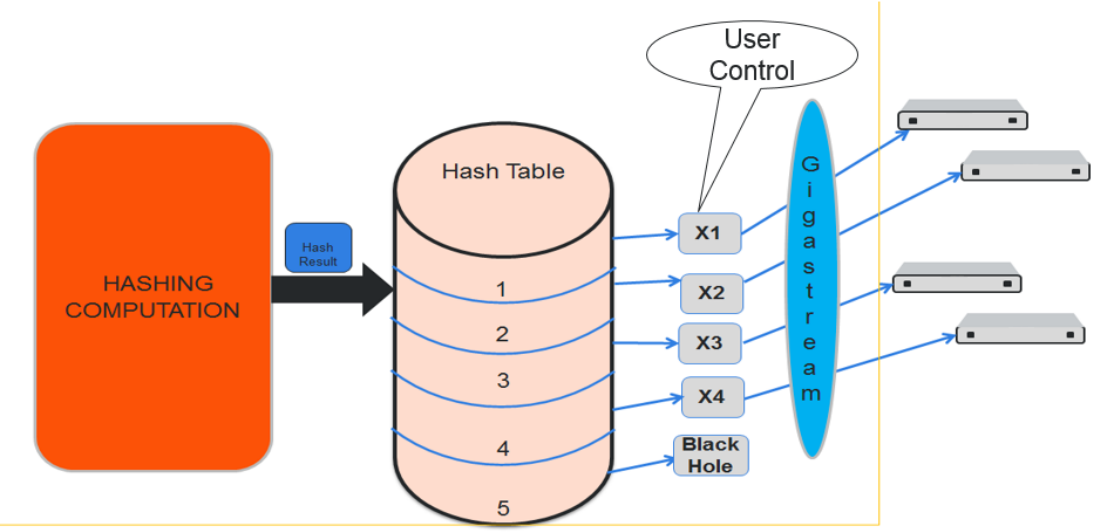

Figure 19 *Controlled GigaStream Example*

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.

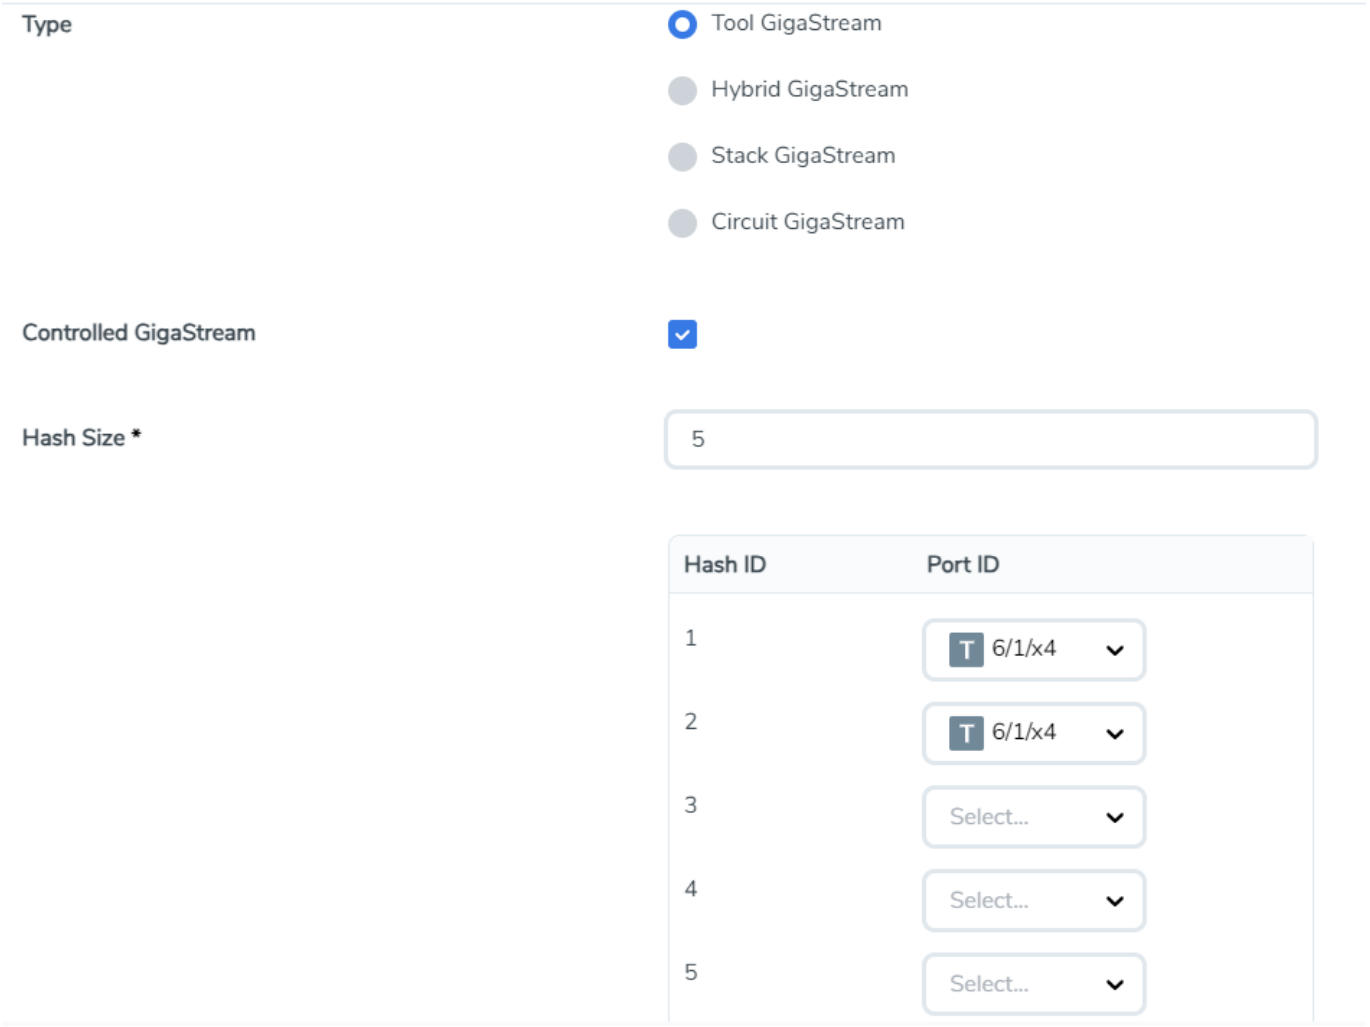

Figure 20 *Controlled GigaStream with Five Hash IDs and Port IDs*

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.

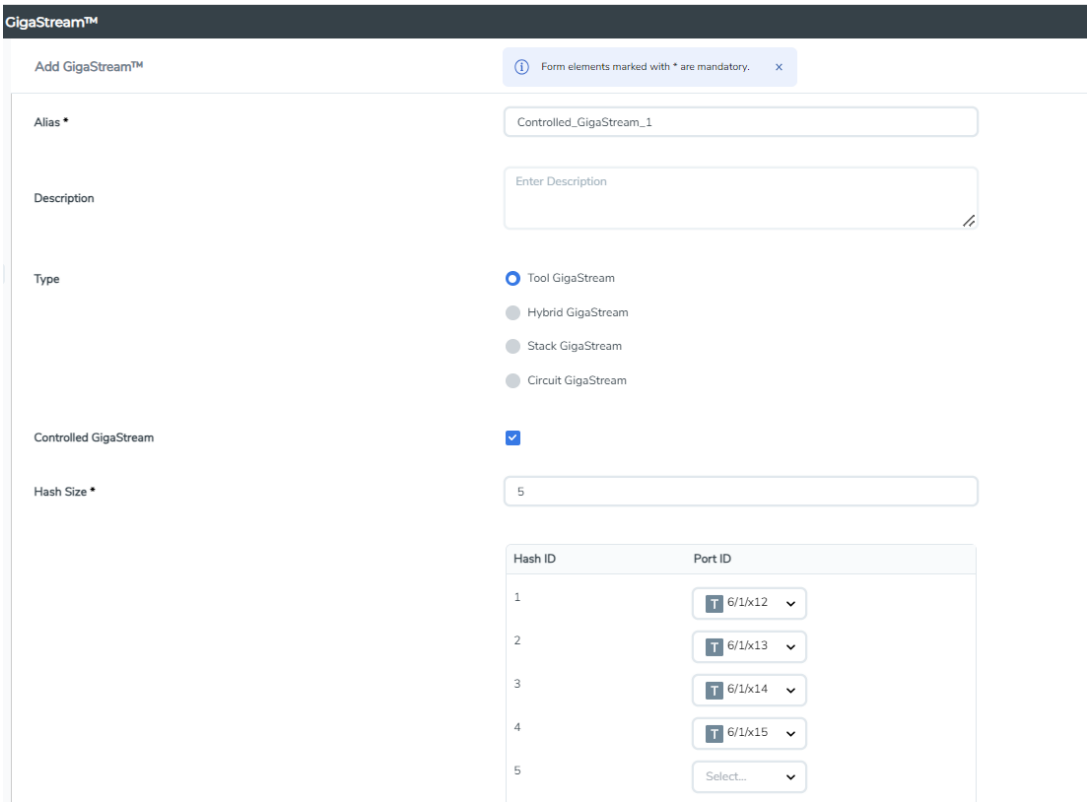

Figure 21 *Ports Assigned to Hash Bucket IDs*

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: • Ingress Ports—Adds ingress port

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.

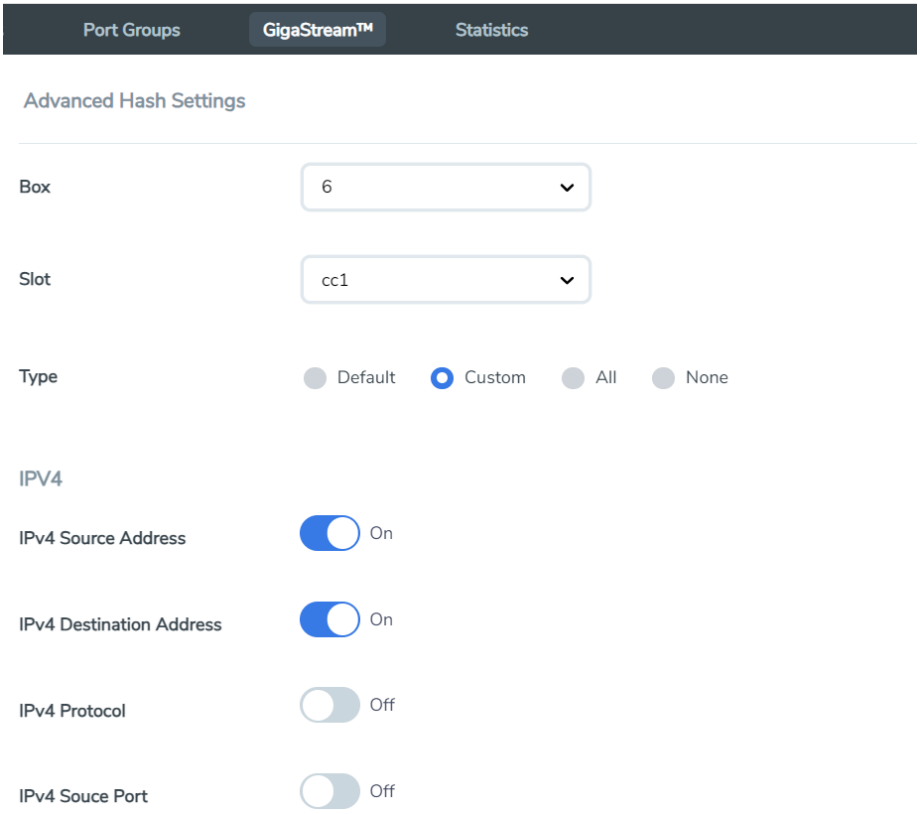

Figure 22 *Advanced Hash with IPv4 Source and Destination Addresses*

The example in Figure 23Advanced Hash Default Criteria sets the advanced hash for slot cc1 in box ID 6 to the **Default** criteria.

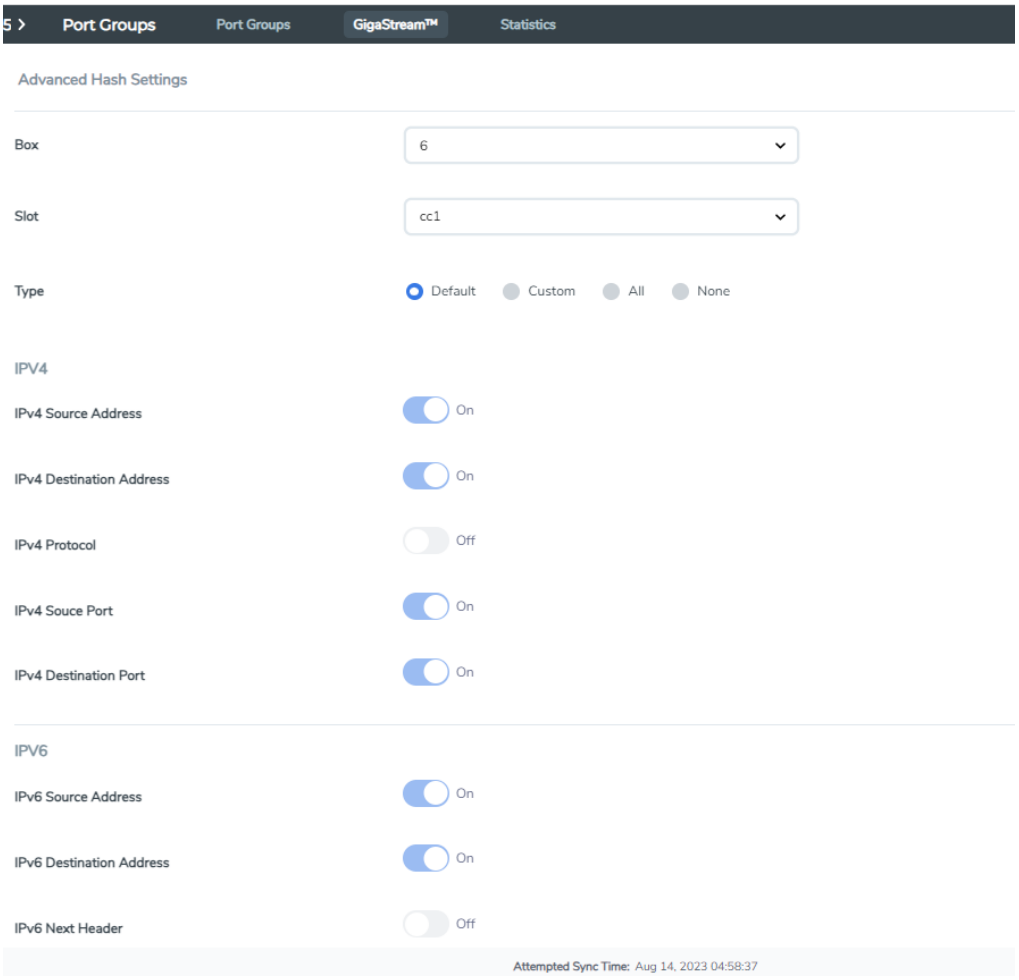

Figure 23 *Advanced Hash Default Criteria*

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.

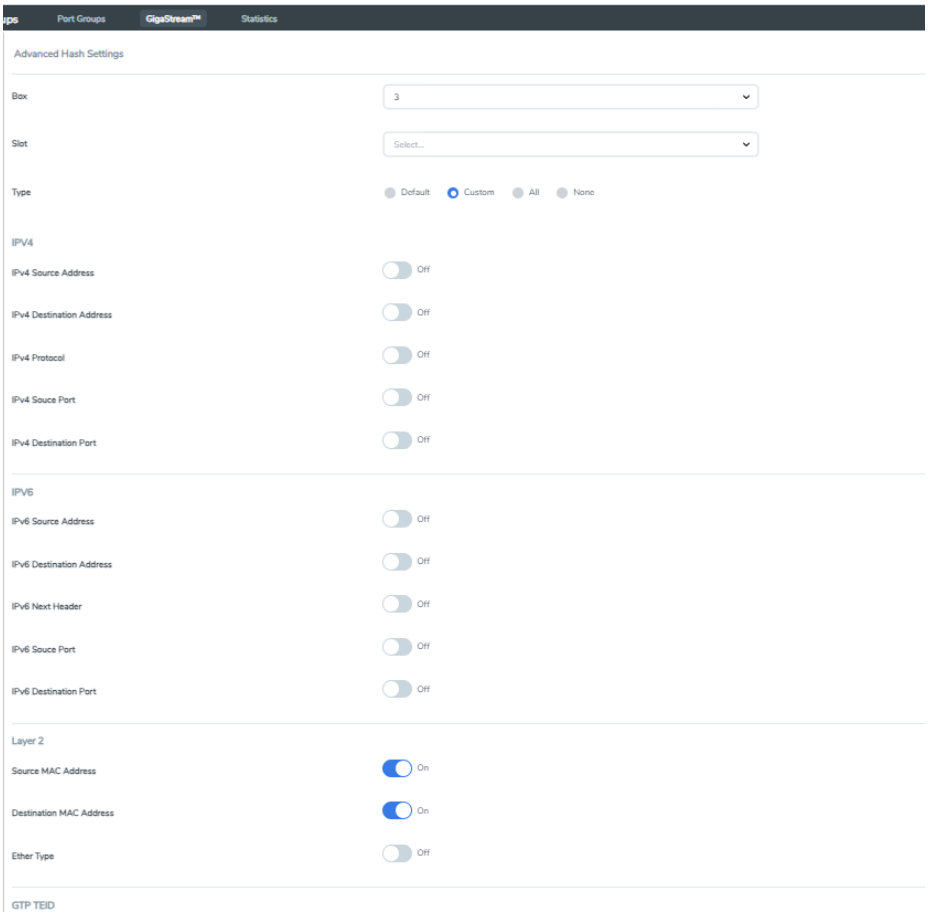

Figure 24 *Advanced Hash with Source and Destination MAC Address*

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 (except GigaVUE-TA25 and GigaVUE-TA400).
  - 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.
- ASICs used in 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-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.

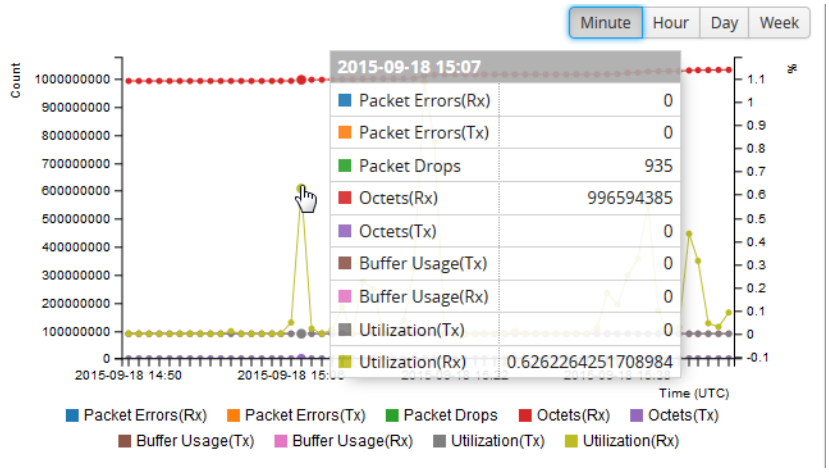

Figure 25 *Numerical Values for Data Points on the Statistics Graph*

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. Figure 26Numerical Values for the Data Points Measured by Hour shows the data points measured by hour.

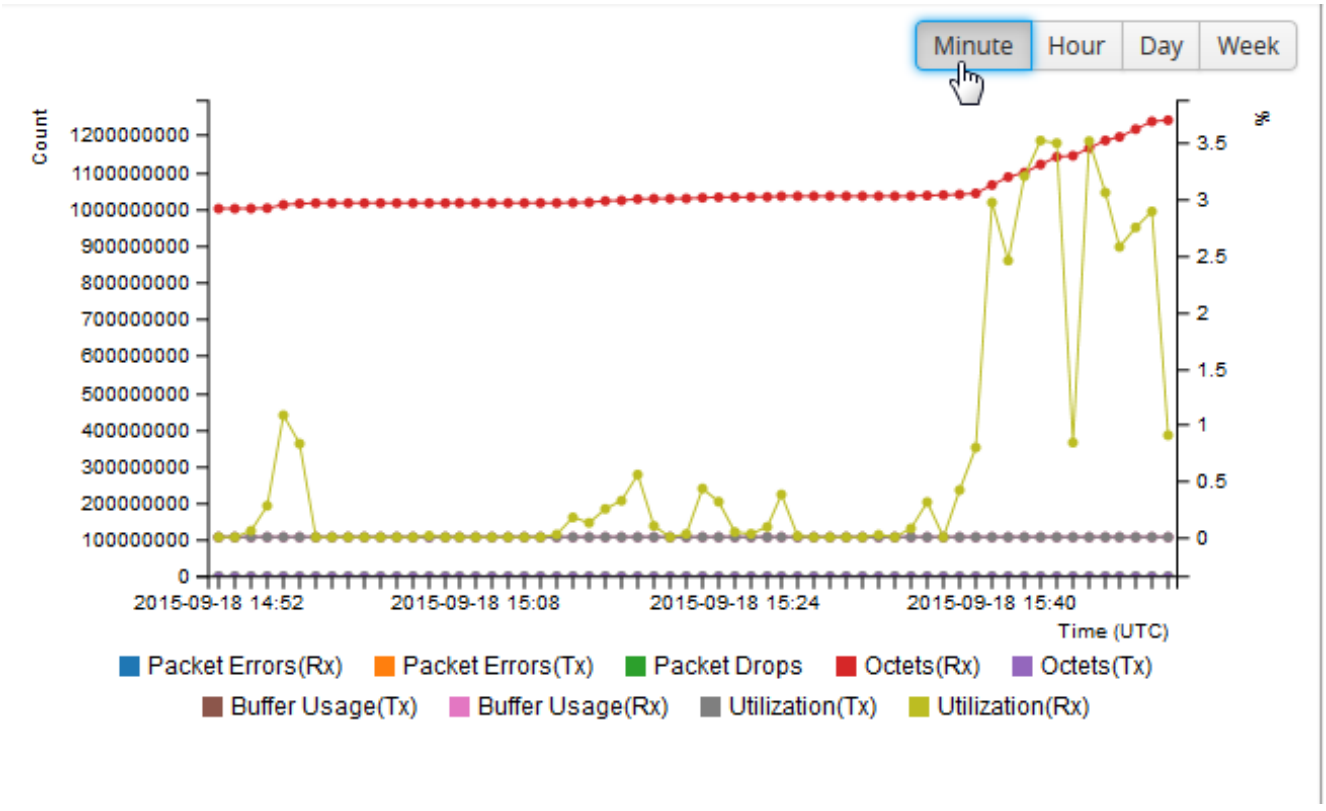

Figure 26 *Numerical Values for the Data Points Measured by Hour*

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. So 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/map rules applied on a network port. ▪ In packets on a tool port. ▪ 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 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 during a microburst. If the buffer reaches 100 percent utilization, packets may 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 27Port Statistics for All Ports. For the definitions of the statistics shown in the table, refer to Table

10: Port Statistics Definitions.

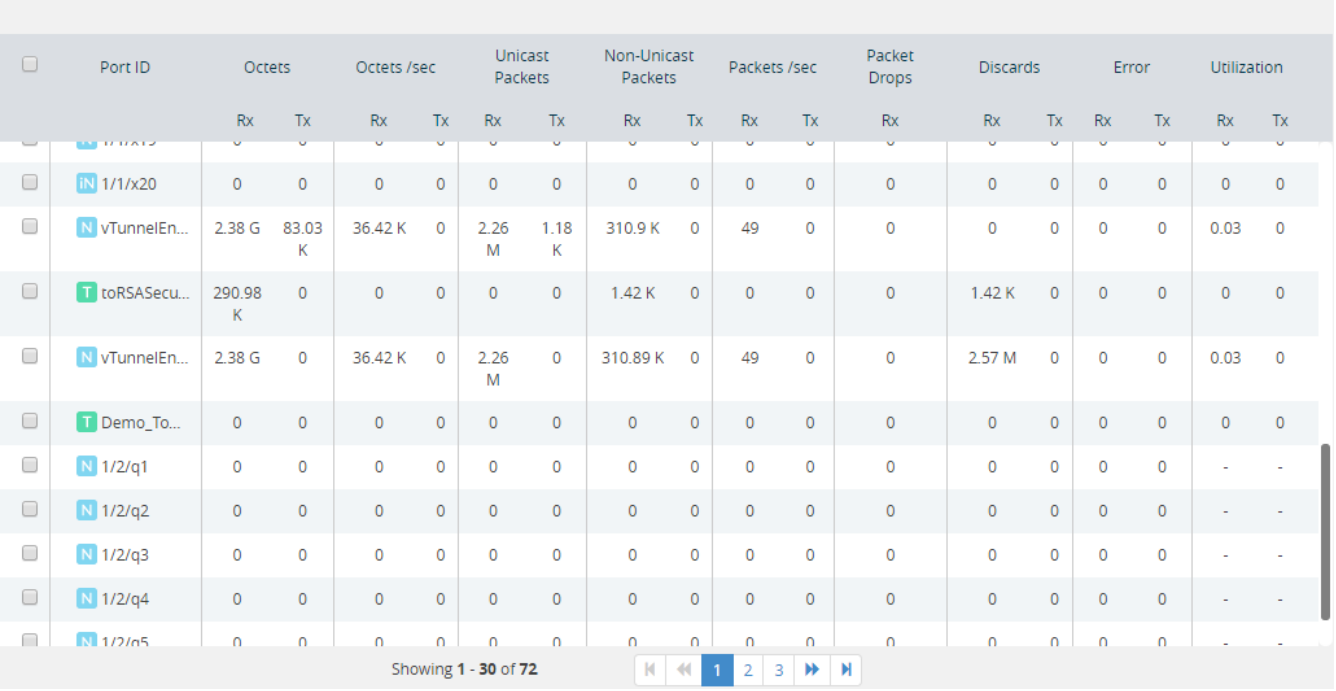

Port ID	Octets		Octets /sec		Unicast Packets		Non-Unicast Packets		Packets /sec		Packet Drops	Discards		Error		Utilization	
	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Rx	Tx	Rx	Tx	Rx	Tx
1/1/x20	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
vTunnelEn...	2.38 G	83.03 K	36.42 K	0	2.26 M	1.18 K	310.9 K	0	49	0	0	0	0	0	0	0.03	0
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	-	-

Figure 27 *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 • Tool

Criteria	Description
	• Inline Network • Inline Tool • GigaSMART • Hybrid • 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. • Enabled — display ports with a link status of up. • Disabled — 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:

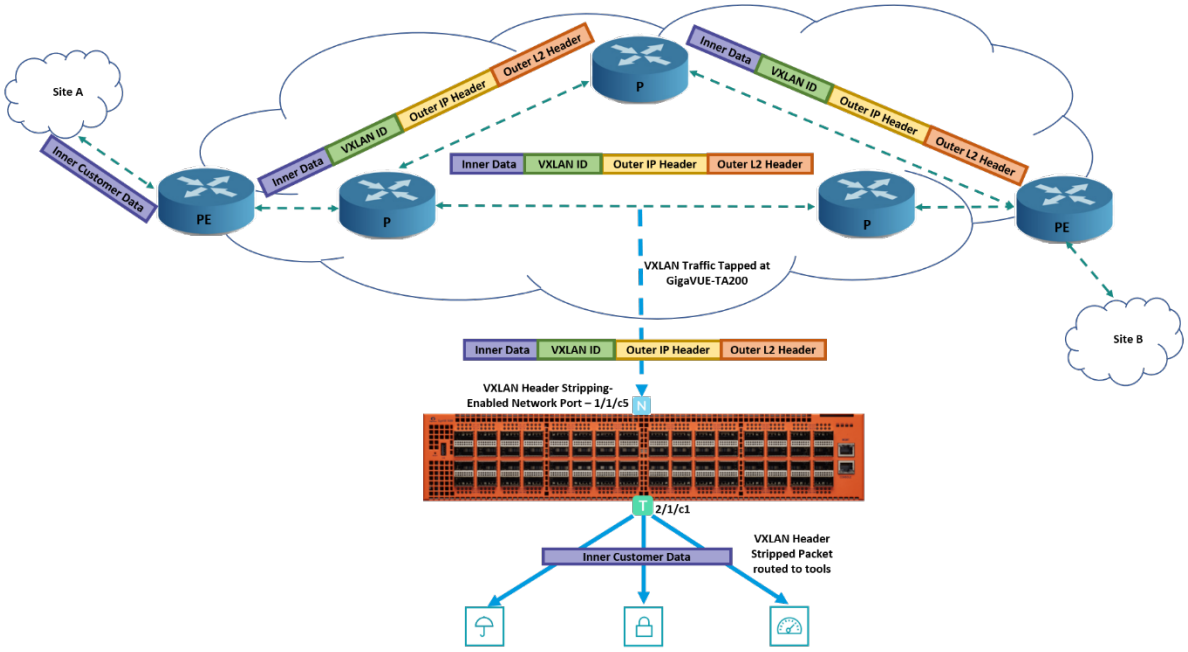

In this diagram, the VXLAN encapsulated traffic from the network is tapped on the network port, 1/1/c5 in the GigaVUE-TA200 device. The VXLAN Header Stripping functionality is enabled in this network port. Based on the traffic intelligence capability, the GigaVUE-TA200 device strips the VXLAN header 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 and GigaVUE-TA400 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, 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**.

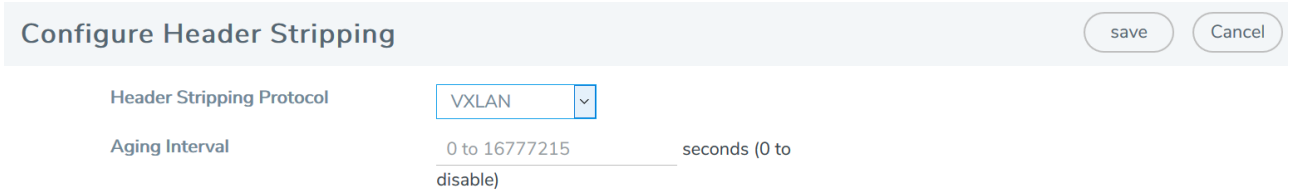

View Header Stripping Statistics

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.

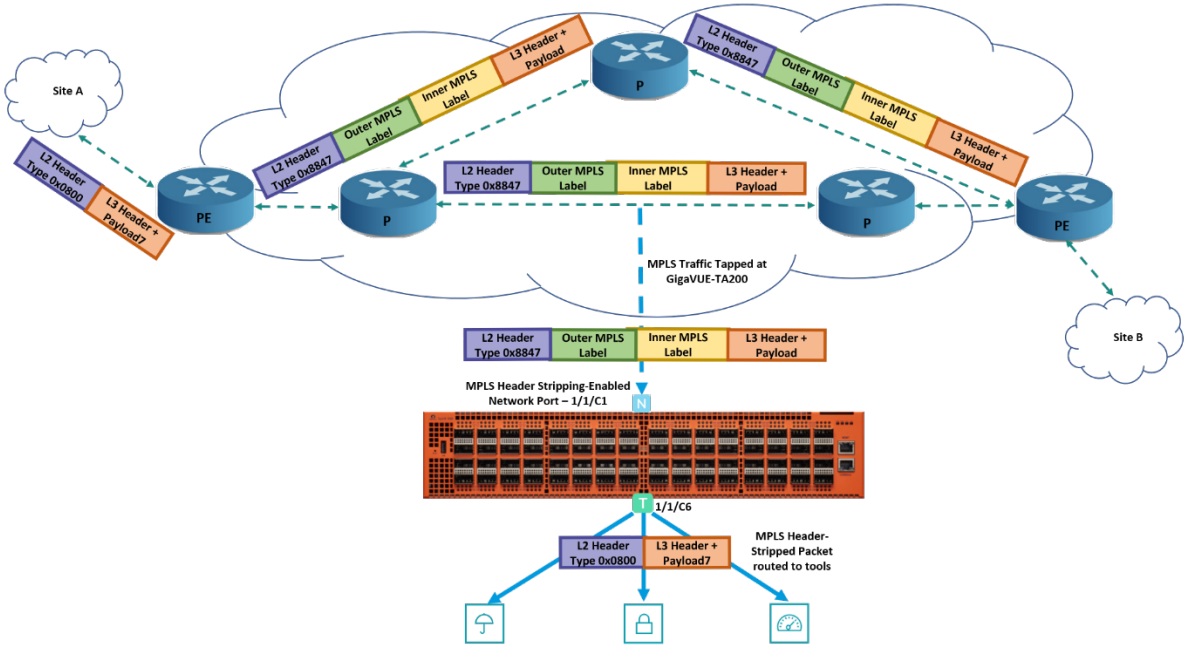

In this diagram, the MPLS encapsulated traffic from the network is tapped on the network port, 1/1/c1 in the GigaVUE-TA200 device. The MPLS Header Stripping functionality is enabled on this network port. Based on the MPLS label configured for the GigaVUE-TA200 device, the device strips the MPLS header 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 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 Layer 2 MPLS packets without Control Word (CW) or Associated Channel Header (ACH). This enhancement is particularly beneficial for customer deployments with prevalent L2 MPLS packets without CW/ACH. Header Stripping applies only to the outer MPLS header in nested MPLS scenarios. 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.

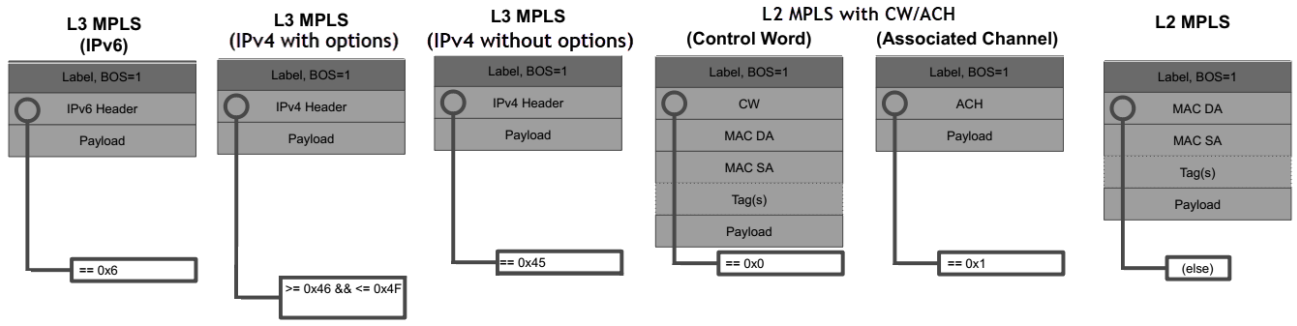

Below are the processing rules for MPLS packets based on specific header values after the Bottom of Stack (BOS) label:

L3 MPLS 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.

L3 MPLS IPv6 - If the first nibble after the BOS label is 0x6, it strips MPLS headers.

L2 MPLS with CW/ACH

- If the first nibble after the BOS label is 0x0, it strips the outer ethernet and MPLS headers along with Control Word (CW).

- If the first nibble after the BOS label is 0x1, it strips the outer ethernet and MPLS headers along with Associated Channel Header(ACH).

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 and GigaVUE-TA400 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. For GigaVUE-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, 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 *Advanced MPLS Header Stripping for GigaVUE-TA400*.

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 *Advanced MPLS Header Stripping for GigaVUE-TA400*.	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. Refer to *Advanced MPLS Header Stripping 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. This option is only applicable in networks that exclusively handle Layer 3 MPLS traffic.
   - **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. 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. 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. 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. 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. 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. 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.

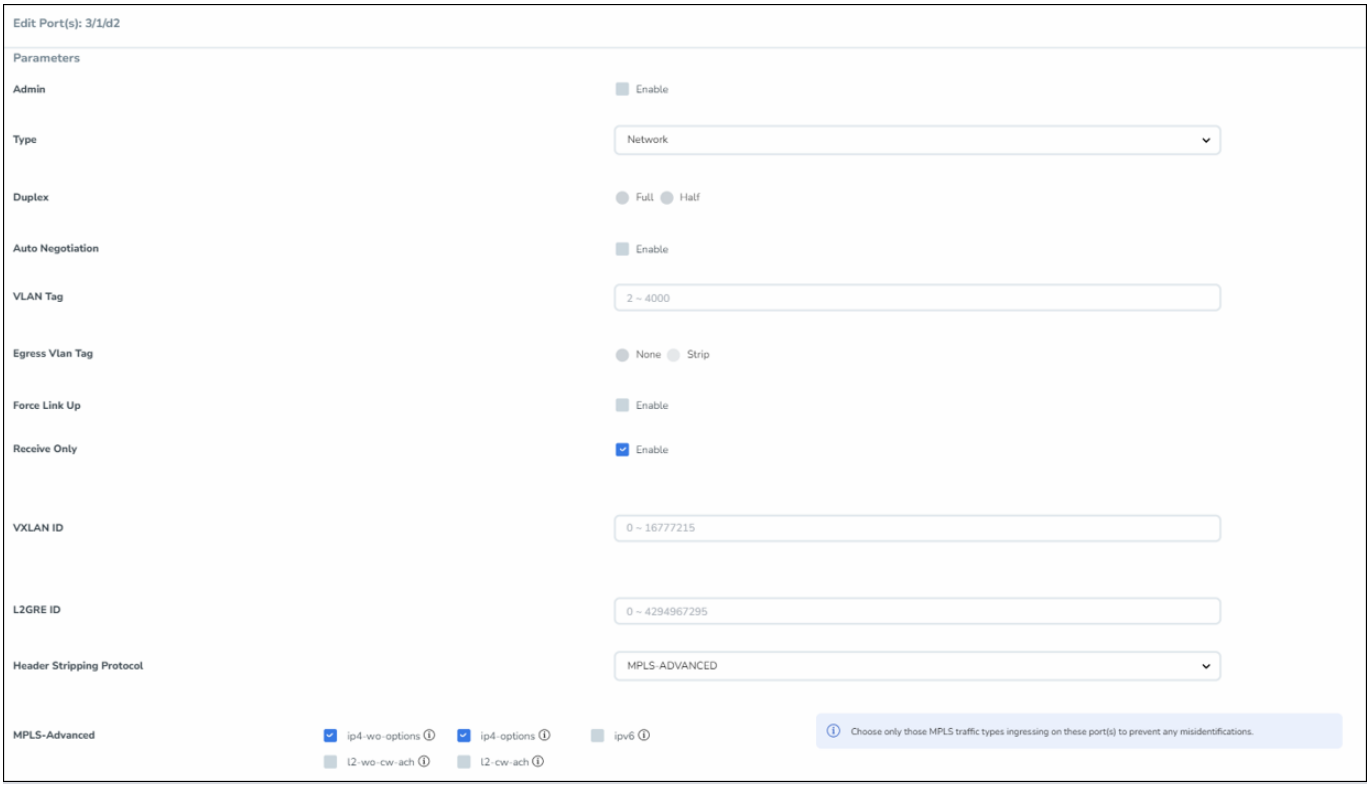

Tunnels

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.

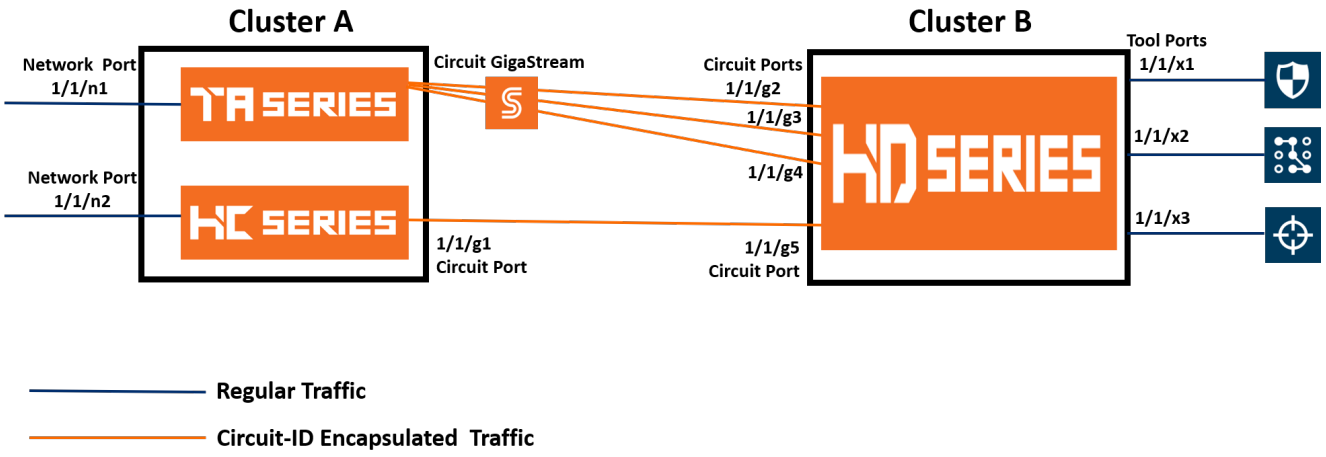

Figure 28 *Circuit Flow Mapping®*

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 512 circuit-IDs are supported within a cluster 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 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.

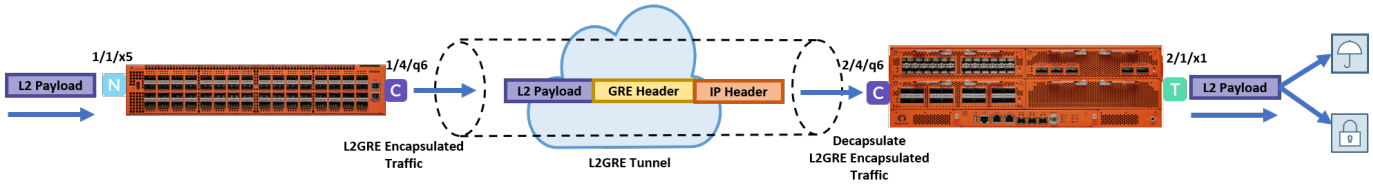

Figure 29 *L2GRE Tunnel Encapsulation and decapsulation*

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 and DELL S4112F-ON devices.
- A maximum of 1500 L2GRE IDs are supported.
- IPv6 protocol is not supported with L2GRE tunnels.
- 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, 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:

Template	Platform	
	GigaVUE-HC1	GigaVUE-HC3/GigaVUE-TA100/TA200/TA200E/TA25//TA25E/ TA400
IPv4	No	Yes
IPv6	Yes	Yes
IPv4+UDA	No	Yes
IPv4+MAC	Yes	Yes
UDA	Yes	Yes

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

S.No	Task	Refer to...
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	Create Tunnel

S.No	Task	Refer to...
	interface that you created in task 2. Ensure that you select the Type as **L2GRE** and Mode as **Decapsulation**.	
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. • 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.

Field	Description
	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.

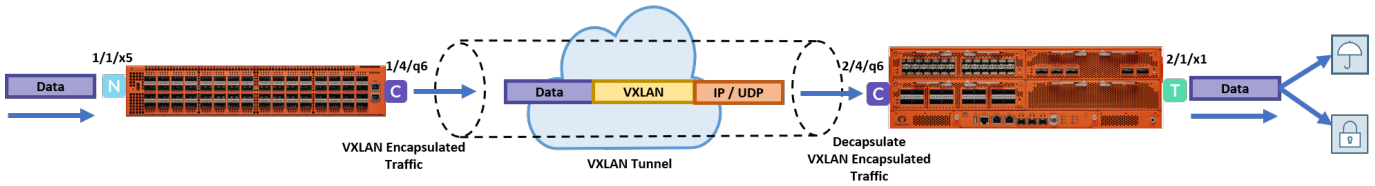

Figure 30 *VXLAN Tunnel Encapsulation and Decapsulation*

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-TA100, GigaVUE-TA200, GigaVUE-TA200E,GigaVUE-TA400 and GigaVUE-TA25, GigaVUE-TA25E 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.
- 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, 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.
- 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:

Template	Platform	
	GigaVUE-HC1	GigaVUE-HC3/ GigaVUE-TA100/TA200/TA200E/ TA25/TA25E/ TA400
IPv4	No	Yes
IPv6	Yes	Yes
IPv4+UDA	No	Yes
IPv4+MAC	Yes	Yes
UDA	Yes	Yes

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
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 VXLAN Tunnel Termination—Rules and Notes.

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
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. • 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. On the left navigation pane, click  , and then select **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. On the left navigation pane, click  , and then select **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 > 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 31VXLAN / L2GRE ID Statistics illustrates the statistics.

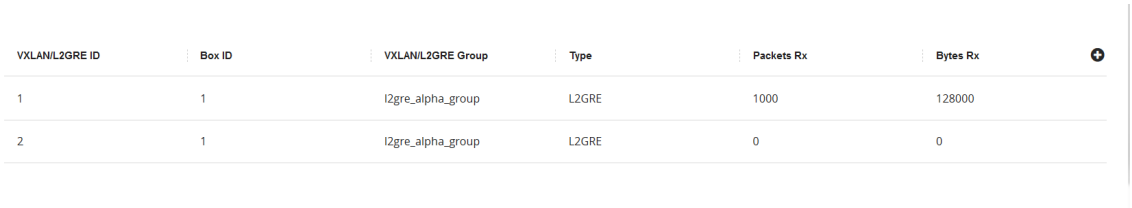

VXLAN/L2GRE ID	Box ID	VXLAN/L2GRE Group	Type	Packets Rx	Bytes Rx
1	1	l2gre_alpha_group	L2GRE	1000	128000
2	1	l2gre_alpha_group	L2GRE	0	0

Figure 31 *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. On the left navigation pane, click on 
2. Go to **Physical > Tunnel Monitoring.**

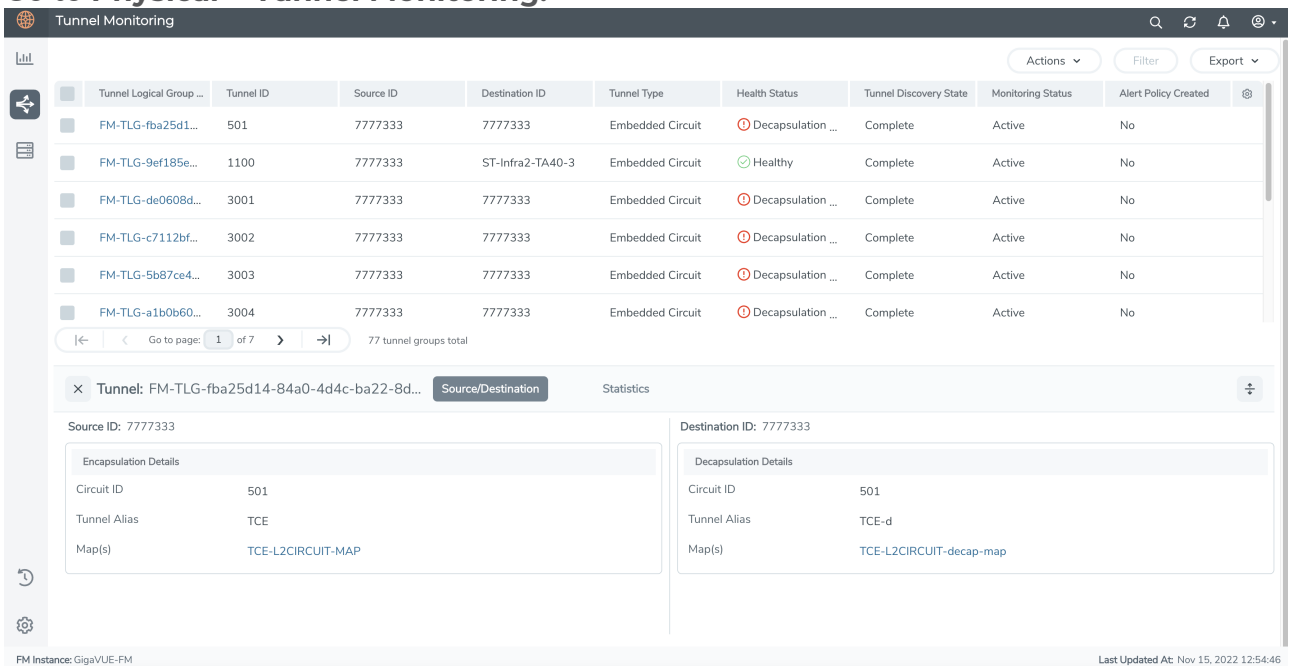

The list of tunnels are displayed.

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

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

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. The GigaVUE-TA400 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 is not supported on ports associated with the IP interface.

- The PCAP feature does not support 'vlan' and 'inner-vlan' filter rules on a tool or hybrid 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, and GigaVUE-TA400 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.

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:

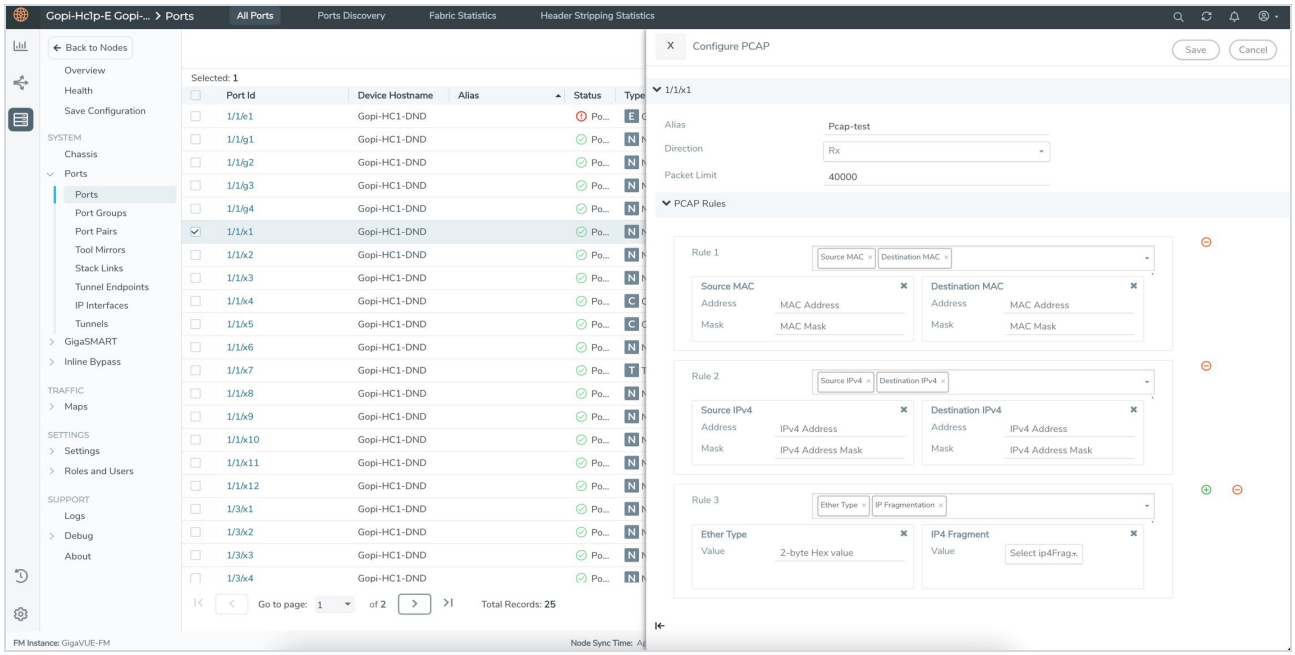

Field	Description
Alias	Name of the packet capture filter
Direction	The direction of traffic. Can be: • Rx • Tx • Both
Channel Port	The channel port identifier for the packet capture filter. The channel port is any unused port that does not have any map configuration. The channel port must be on the 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

Field	Description
	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.
	• **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.

Field	Description

	• **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. On the left navigation pane, click  , and then select **Physical > Nodes**.
2. Select a cluster ID, and then from the left navigation pane, go to **Support > Debug > PCAP**.

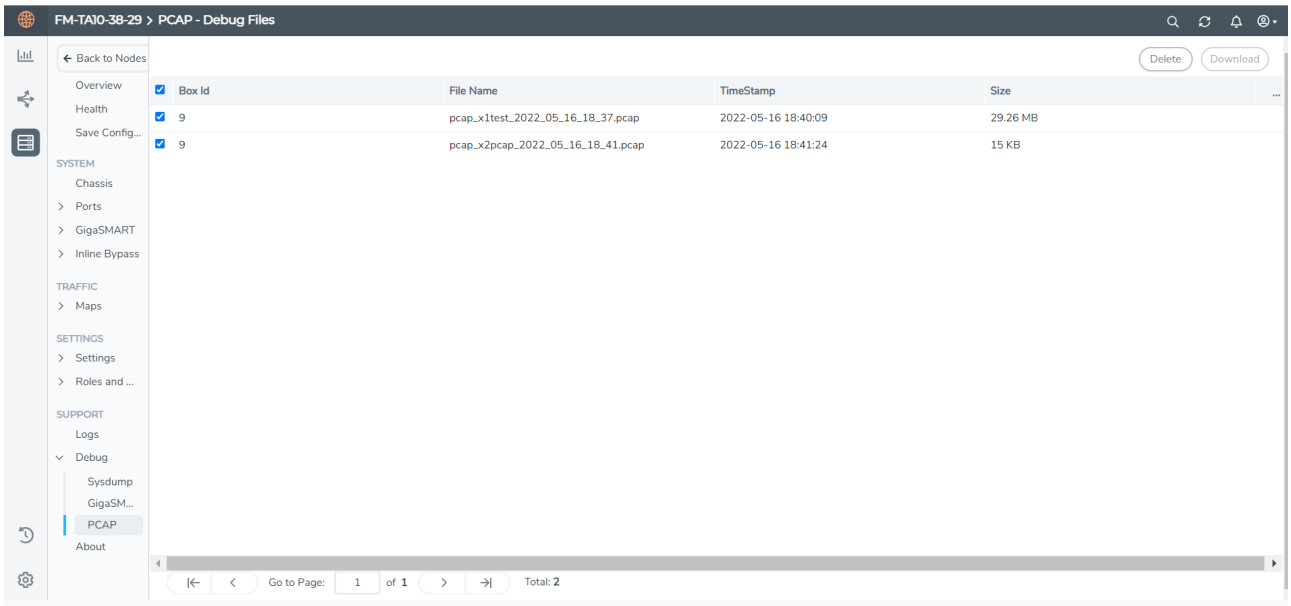

3. Select the required PCAP file(s):
   - Click **Download** to download the file. You can download only one file at a time.
   - Click **Delete** to delete the PCAP files.

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. On the left navigation pane, click on  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*. From the left navigation pane, 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 and GigaVUE-TA200.

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.

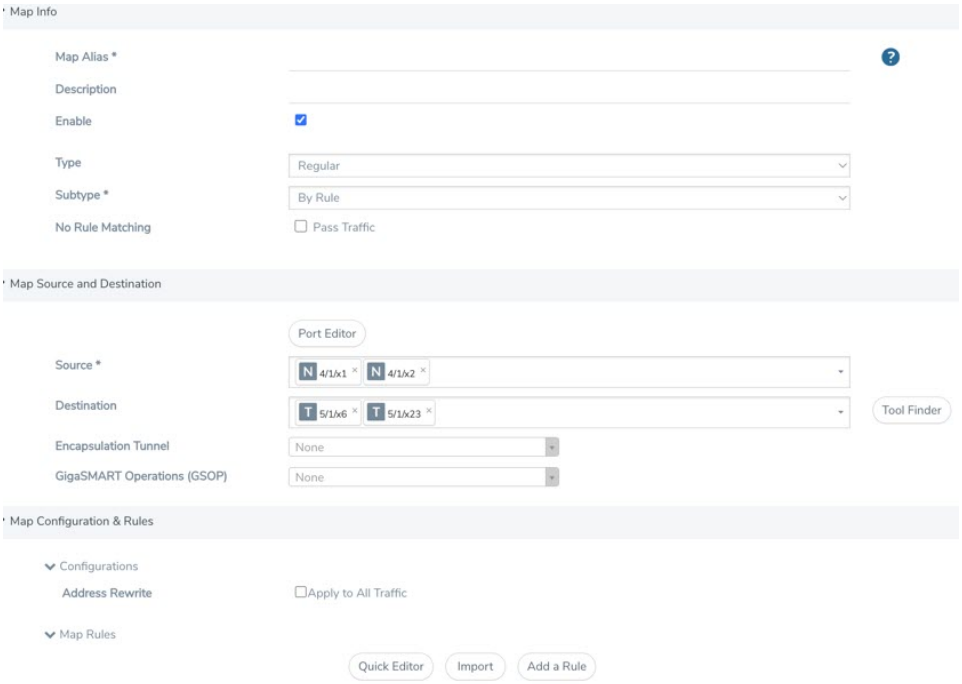

NOTE: Shared-collector maps do not include any rules, priority settings, or GigaSMART operations. These are grayed out in the UI when Collector is select for the map's subtype.

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.

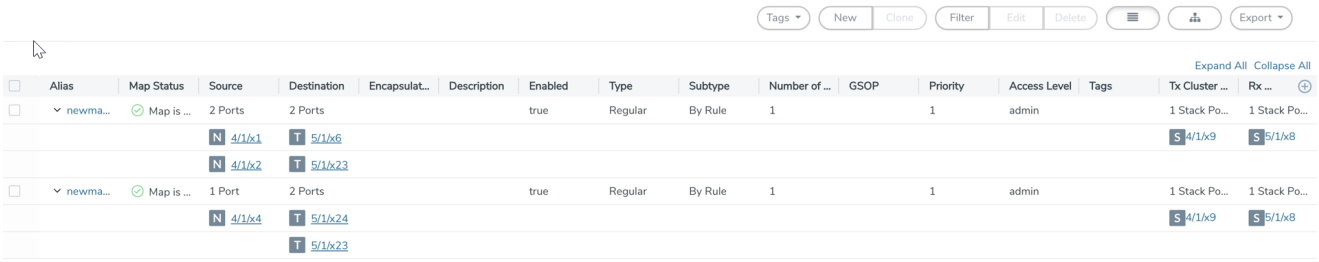

Figure 1 *Shared Collector Map*

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 can also be configured

- map3 **Source** ports 1/1/x3, 1/1/x4, 1/1/x5 cannot be configured because ports x4 and x5 are in both maps (they overlap)

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, SMT-HC3-C05 of HC3v2, 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 SMT-HC3-C05 of GigaVUE-HC3v2	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	gsrule	yes	virtual ports	tool ports, or

Type	Subtype	Map Rule	GSOP	Source	Destination
	Collector				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
- Work with User-Defined Pattern Match Rules

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.
	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

Flag	Bit Position	Pattern	TCP Control Mask
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:

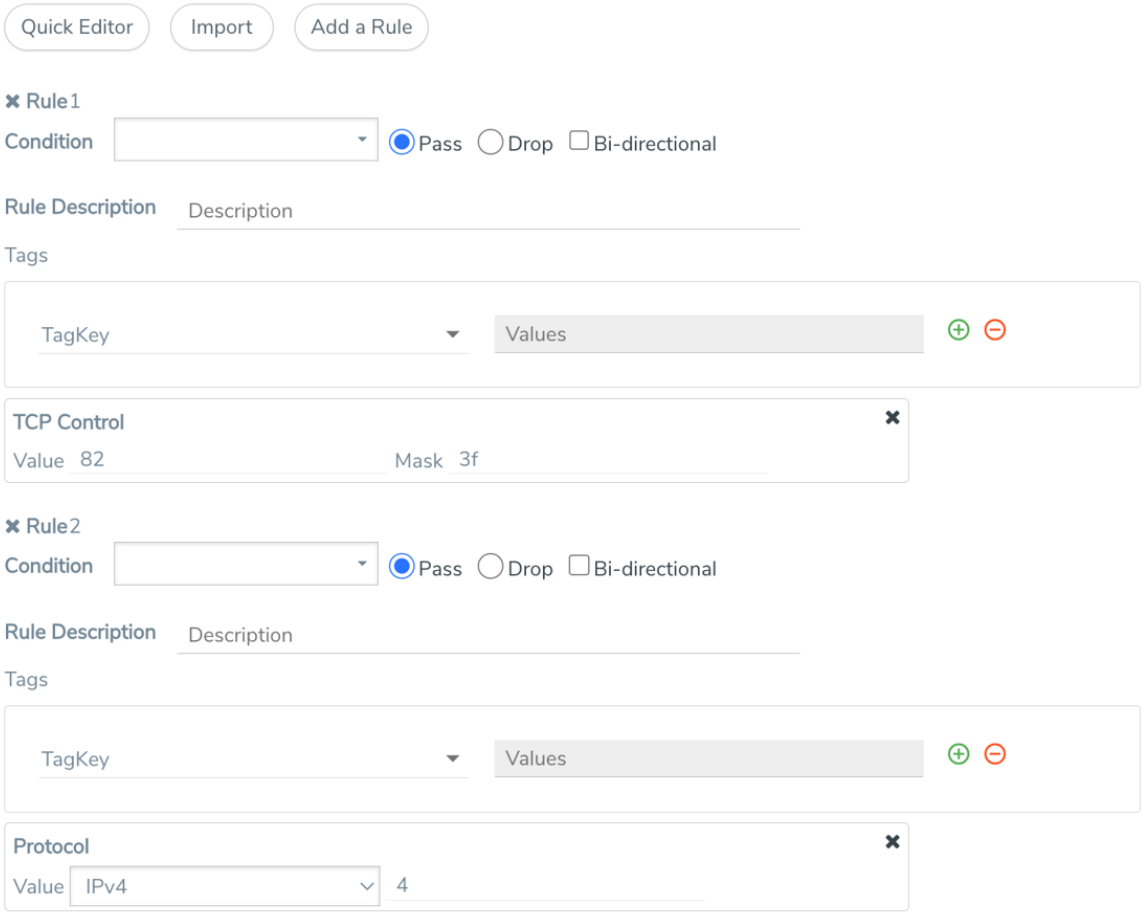

Figure 3 *Map Rule with SYN Bit Set*

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.

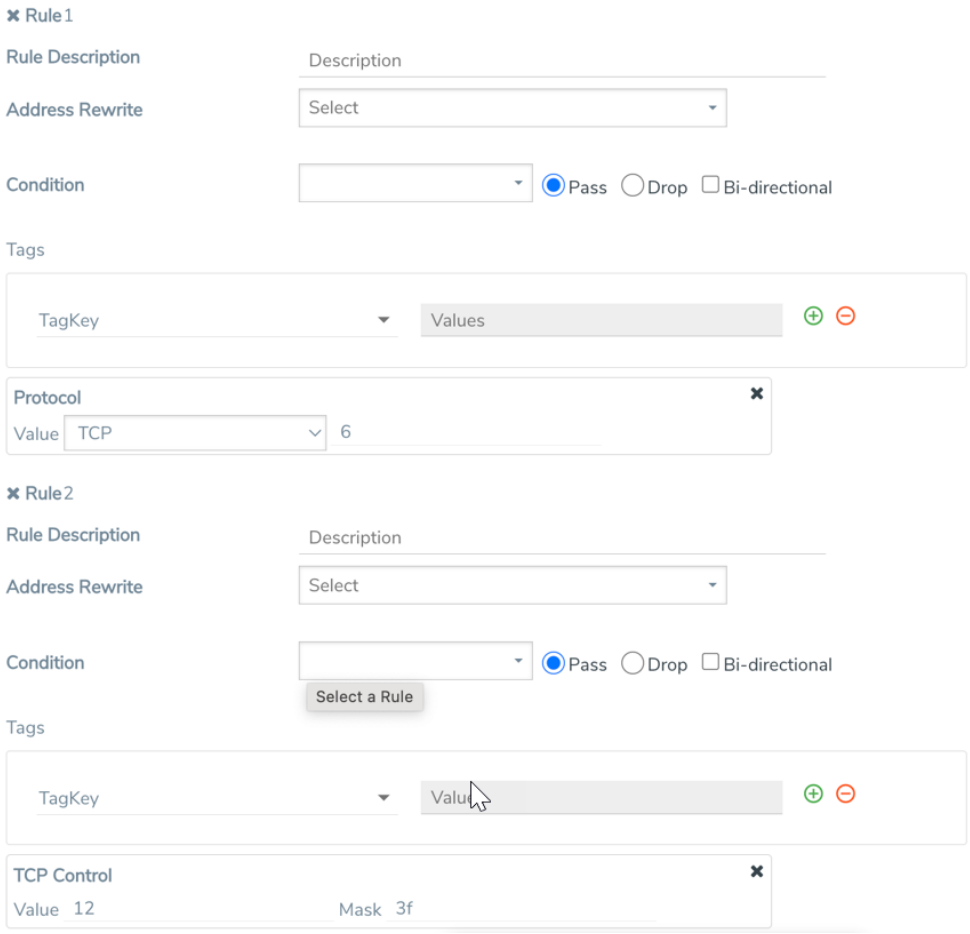

Figure 4 *Map Rule Matching All Packets with Both ACK and SYN Bits set*

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.

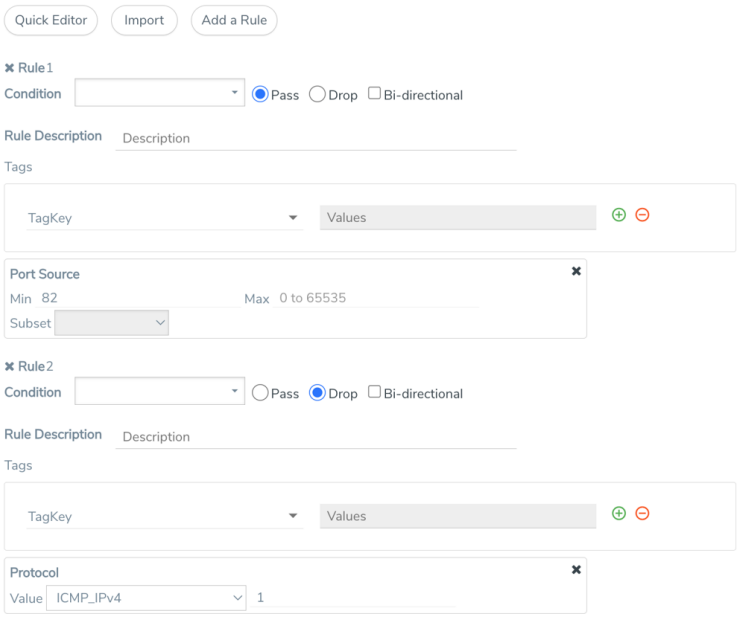

Figure 5 *Pass and Drop Rules in a Map*

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.

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.
- On GigaVUE-TA400, both UDA1 and UDA2 support 8-byte data match from offset 0 to 160. UDA1 data match is from the start of ethertype, UDA2 data match is from the start of the L3 header.
- 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

Description	UI Step
Save the map.	**13.** Click Save.
Map #2 – GTP_Map228 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.

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.

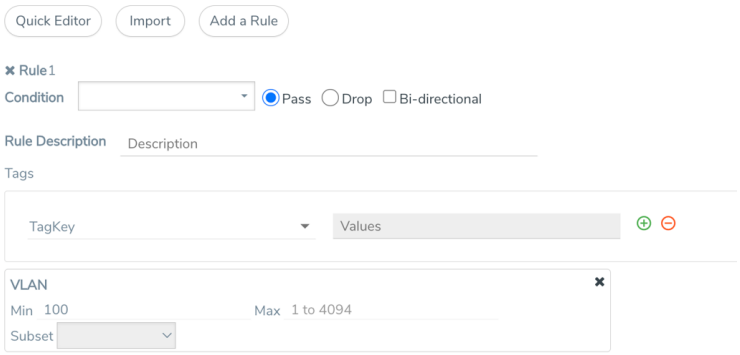

- For traffic that matches the map drop rule shown in the following figure, Q-in-Q packets of TPID ethertype 0x8100, 0x88A8, and 0x9100 are dropped.

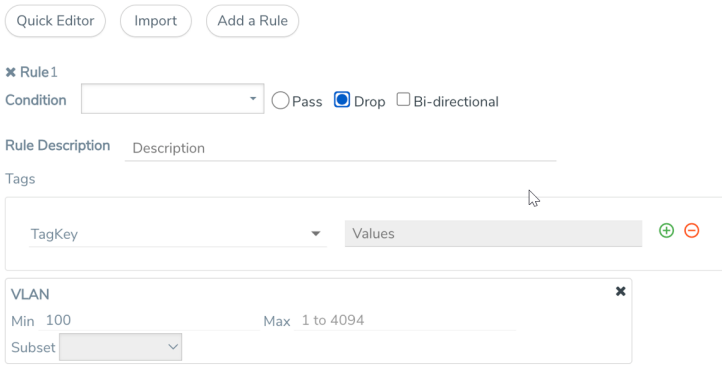

You do not specify TPID EtherTypes 0x8100, 0x88A8, and 0x9100 explicitly in a rule. If you specify these values in the **Value** field an **EtherType** rule, the map is blocked and one of the following error messages is displayed:

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
Two tags: outer TPID 8100 VID 100, inner TPID 88a8 VID 100, ethertype 0800	pass	drop	drop
Two tags: outer TPID 8100 VID 100, inner TPID 9100 VID 100, ethertype 0800	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	×	✓	✓	✓	✓

Permissions	Level 1	Level 2	Level 3	Level 4	Admin
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 6Map for VLAN 100..199.

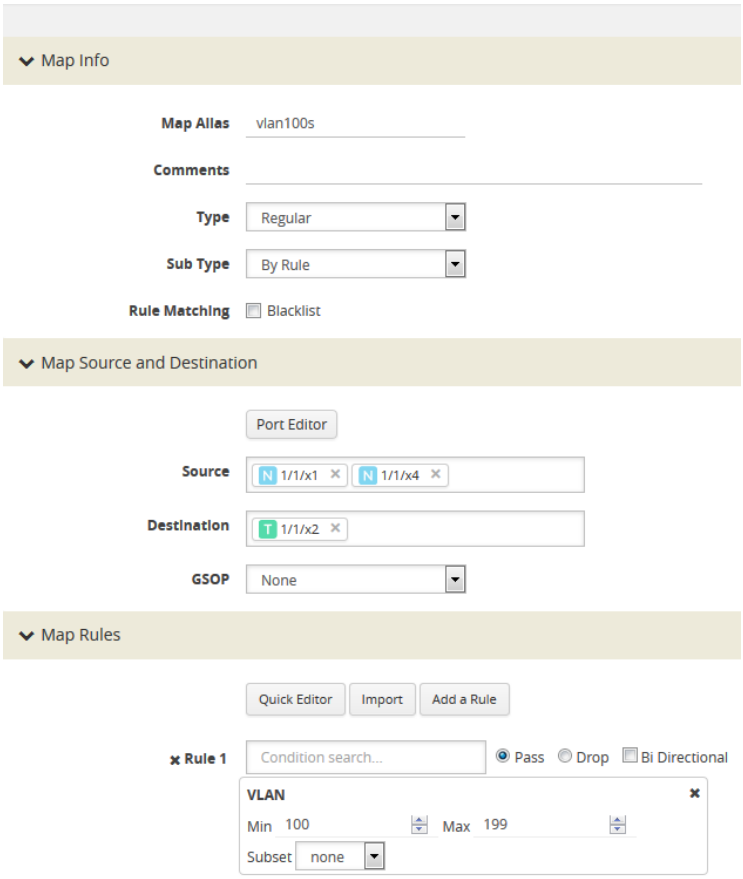

Figure 6 *Map for VLAN 100..199*

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 7Sending 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 7Sending 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 7Sending 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.

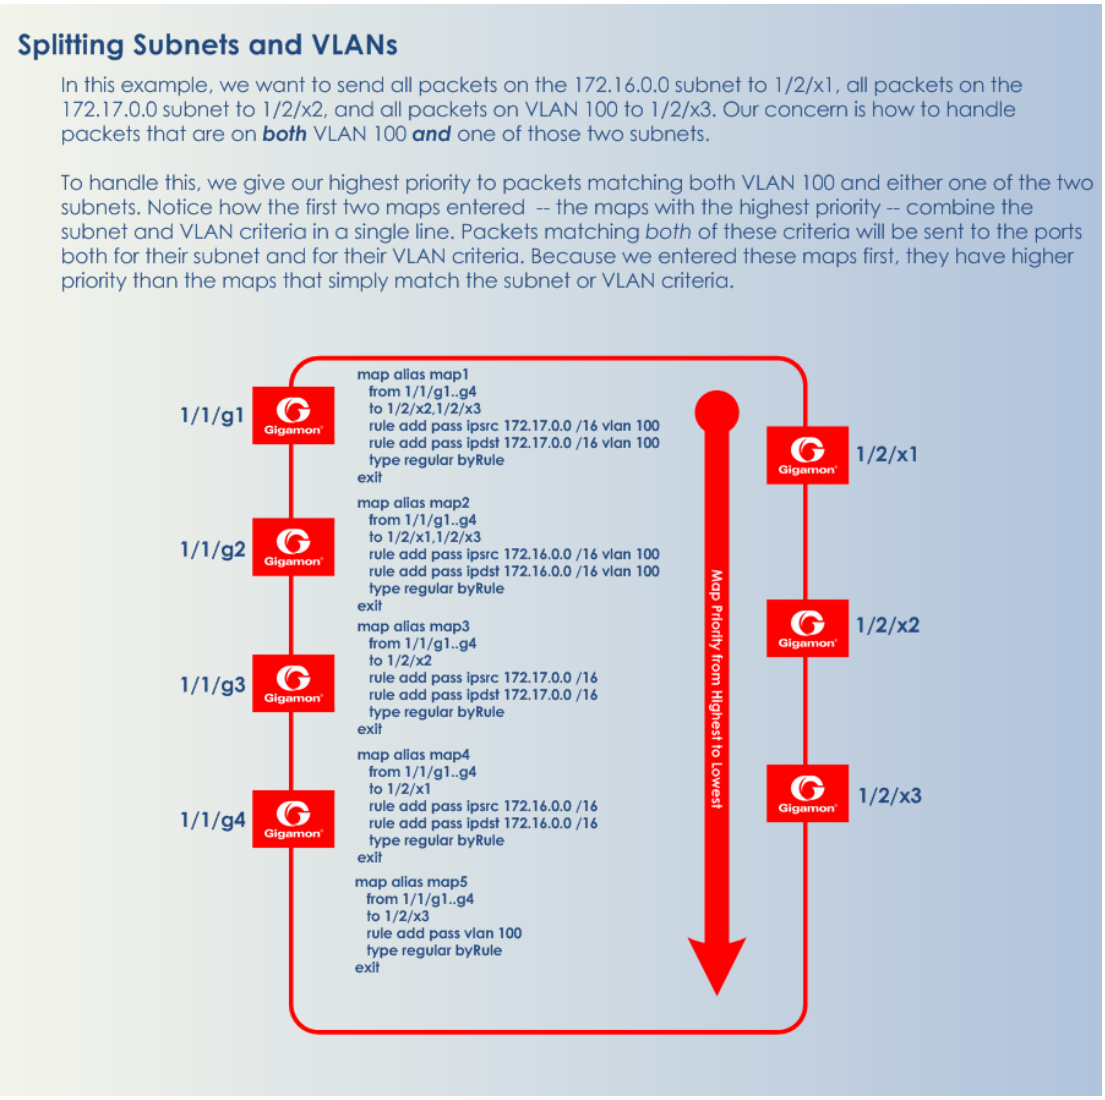

Figure 7 *Sending Subnets and VLANs to Different Ports*

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 APIs. The maps can be displayed in following types of views:

- List View
- List View

NOTE: Starting in software version 5.5.01, any change in the map health status is indicated immediately in the Maps page.

List View

The List View is the default view of the maps when the Maps page is opened after selecting **Maps > Maps**. 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 on the icon. The Topology View of the maps appear. Each map in the Maps page gets displayed as Topologies.

1. Use the zoom and scroll buttons at the bottom of the page to zoom in and out.
2. Click on a map to view the connections.

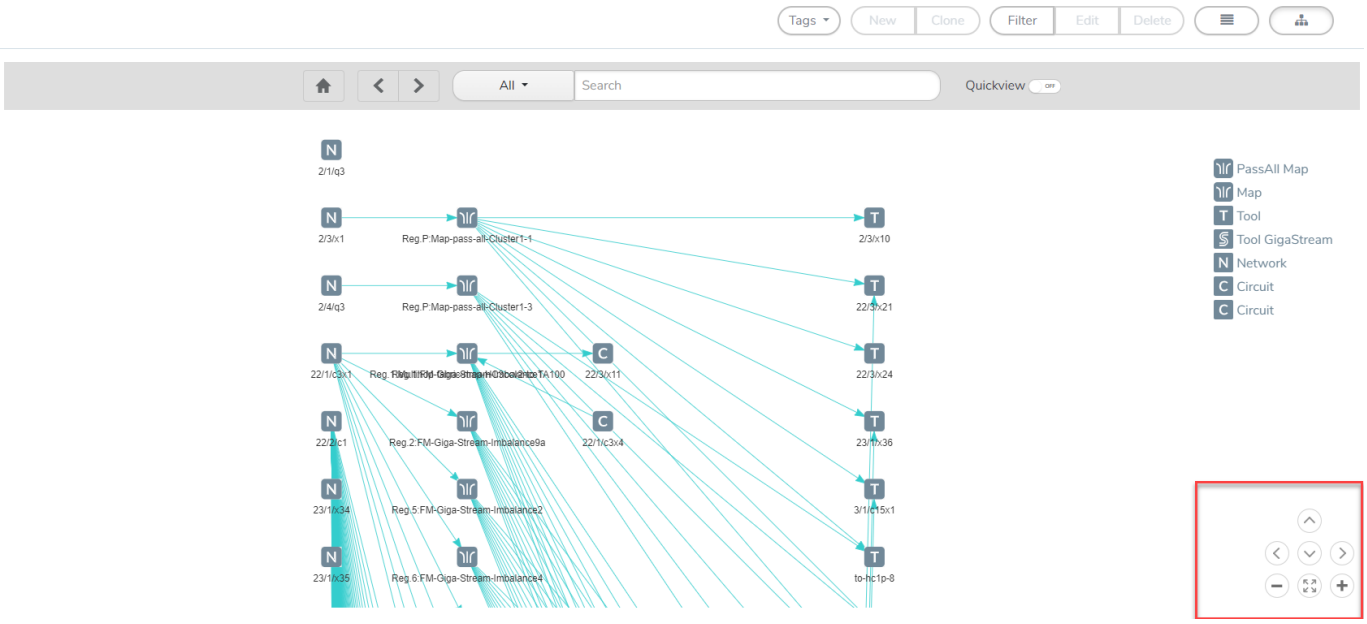

Manage Maps

This section provides the basic steps for doing the following tasks:

- Map Aliases
- 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

Map Aliases

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.

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 about 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.
- To include double quotation marks (") inside the quotation marks, precede it with a backslash (\).

  See also Description to Map Rules.

   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 are matching.

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 operation from the **GigaSMART Operations (GSOP)** drop-down list.

   e. If you select 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 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.

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.

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.

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.

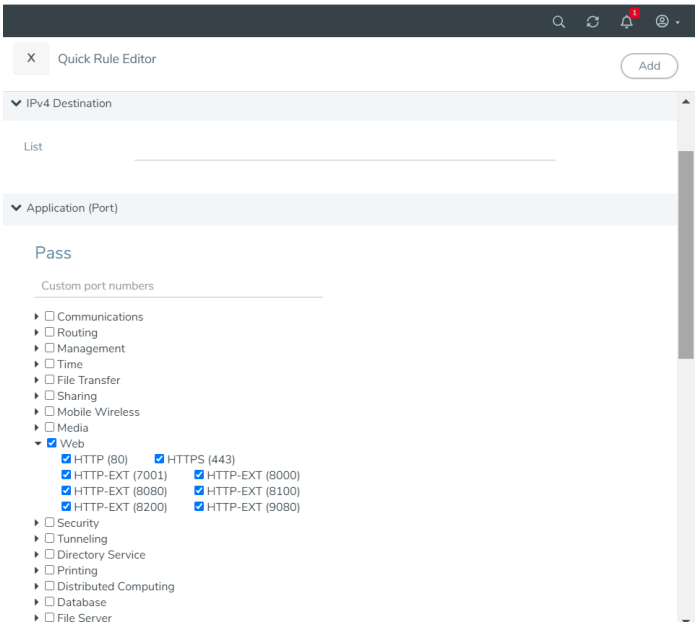

3. Click **Add**.

   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.

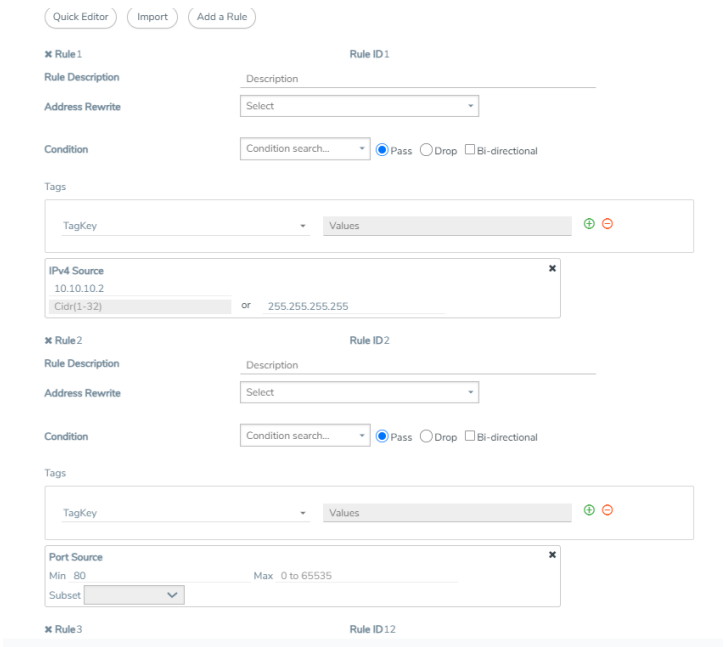

How to Use the Quick Port Editor to Add IP Address

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	✓	✓	✓	✓	✓
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

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.

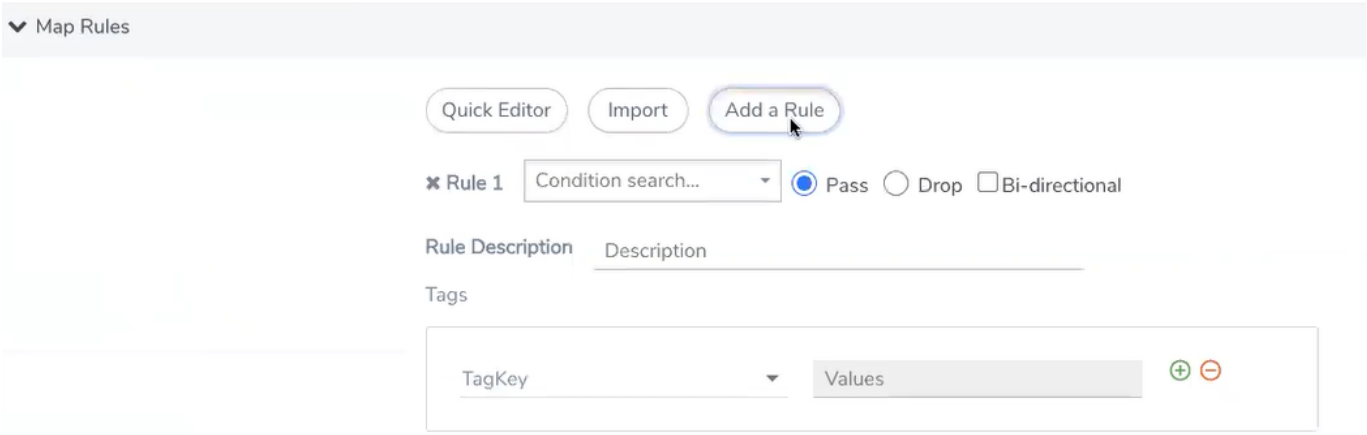

Notes

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

Visualizations	Displays...
	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 and GigaVUE-TA400.

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 or GigaVUE-TA200. 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.

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

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, and 36k rules per node on GigaVUE-TA25E node.

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.

- 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.

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 slot** 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. The total cost of qualifiers for a map rule must not exceed 10.

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 6: 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
portsrc	16	1
portdst	16	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

Qualifier	Bits	Cost
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.

The qualifier cost is the cost of all qualifiers + 54 bits.

- If the cost is less than or equal to 80 bits, 6K rules/slot are supported.
- If the cost is greater than 80 bits but less than 160 bits, 3K rules/slot are supported.
- If the cost is greater than or equal to 160 bits, 1K rules/slot are supported.

Examples:

- For the ip6src and vlan qualifiers—ip6src is 128 bits, vlan is 16 bits, so the total is 128+16+54 bits, which is a cost greater than 160 bits, so 1K rules per slot are supported.
- For the portdst qualifier only—portdst is 16 bits, so the total is 16+54 bits, which is a cost less than 80 bits, so 6K rules per slot are supported.

The maximum cost supported is 480 bits/template.

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 7: 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.

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-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 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-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 8: 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 9: 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

Node Type	Rules per map
GigaVUE-TA100	4K
GigaVUE-TA200	4K
GigaVUE-TA200E	4K
GigaVUE-TA400	4K

How Many Maps Can Run at Once?

The maximum number of maps that can run is only limited by the total number of rules used by the maps.

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?

There is no limit to the number of maps that can share a network 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 8Tool Mirror Page.



Figure 8 *Tool Mirror Page*

Does Flow Mapping® Support port-pairs?

Yes.

Select **Ports > Port Pairs** to go to the Port Pairs page shown in Figure 9Creating Port Pairs. For more information about port pairs and details on configuring two ports as an inline TAP, refer to Port Pairs.

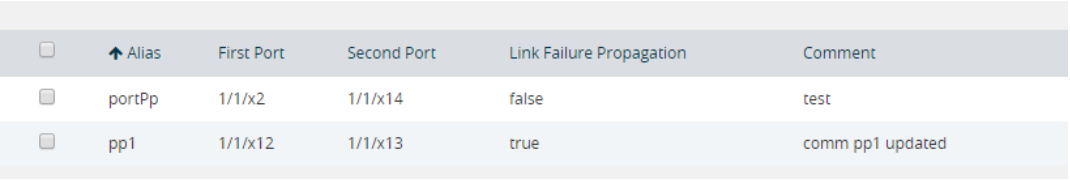

Alias	First Port	Second Port	Link Failure Propagation	Comment
portPp	1/1/x2	1/1/x14	false	test
pp1	1/1/x12	1/1/x13	true	comm pp1 updated

Figure 9 *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. On the left navigation pane, click on 
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**.

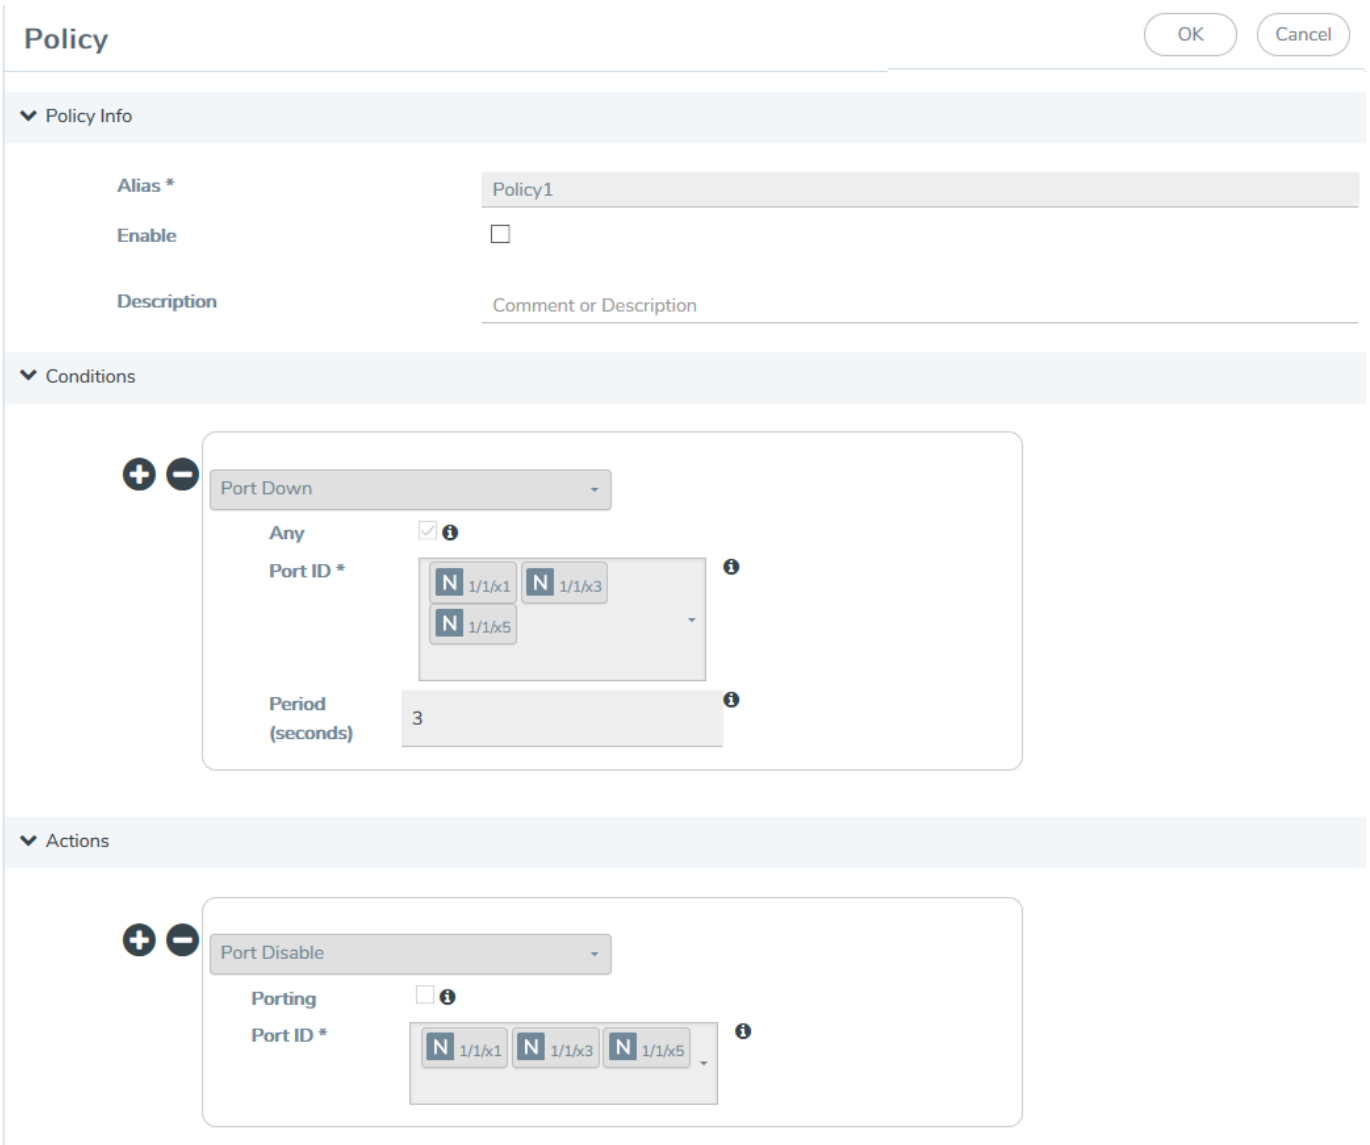

The newly configured policy is displayed in the Policies page. When a policy has not been executed, it will show a status of NOT EXECUTED. When a policy has been executed, it will show a status of SUCCESS or FAILURE.

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. On the left navigation pane, click on 
2. On the left navigation pane, go to **Physical > Active Visibility,** and then select the required cluster or node ID.
3. On the left navigation pane, click **Conditions** to view the list of pre-defined conditions or click **Actions** to view the list of pre-defined actions.

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 10: 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 11: 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

The VLAN manipulation adds a new VLAN tag to the outgoing traffic with the user-configured VLAN value. The modified packets are then delivered as per Flow Mapping® configurations. This allows the user to maintain the confidentiality of the outgoing traffic.

VLAN manipulation can be enabled in two ways:

- **Rule based**- The VLAN manipulation functionality is enabled for traffic that qualifies as a specific rule on a map. This can be enabled only for pass rules. Rule based VLAN manipulation adds a new VLAN tag to the incoming traffic that matches the rule.

- **Map Based**- The VLAN manipulation 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 VLAN manipulation adds a new VLAN tag to the incoming traffic that matches the rule and can also be applied to a deployed map. Refer to Map VLAN manipulation Source and Destination Compatibility Matrix for more information.

Table 12: 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
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	No
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 VLAN manipulation functionality in the same map, then rule-based configuration takes priority.

Configuring VLAN manipulation

1. To enable VLAN manipulation functionality through GigaVUE-FM:
   a. **Map-based Configuration**- To configure VLAN manipulation 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 If you want to configure the MAC address.
      - Select the **Add** option from the VLAN Action field.
      - Enter a **VLAN ID** value between 1 and 4095 in the VLAN ID field.

- Select the **TPID** for the VLAN Tag. The default value of TPID is 0x8100. Select the other supported values 0x9100 and 0x88a8 from the drop-down list.
- Click on **OK** to complete the configuration.

b. **Rule based Configuration**- To configure VLAN manipulation 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**.
- Enter the description of the rule in the Rule Description field.
- Select the **Add** option from the VLAN Action field.
- Enter a **VLAN ID** value between 1 and 4095 in the VLAN ID field.
- Select the **TPID** for the VLAN Tag. The default value of TPID is 0x8100. Select the other supported values 0x9100 and 0x88a8 from the drop-down list.
- Click on **OK** to complete the configuration.

2. To enable VLAN manipulation 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:

```
vlan-op add <vlan id> tpid <value>
no vlan-op
```

Refer to the GigaVUE-OS CLI Reference Guide for more information.

License

You do not need a license to enable this feature for the GigaVUE HC Series. To enable this feature for the GigaVUE TA Series ensure you have an Advanced Features License.

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 Tunnel Encapsulation, 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.
- The cluster 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.

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.

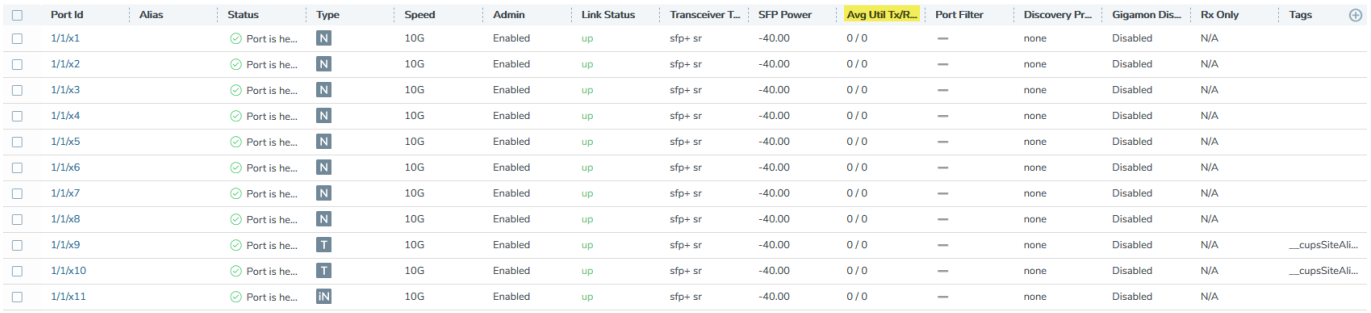

Port Id	Alias	Status	Type	Speed	Admin	Link Status	Transceiver T...	SFP Power	Avg Util Tx/R...	Port Filter	Discovery Pr...	Gigamon Dis...	Rx Only	Tags
1/1/x1		Port is he...	N	10G	Enabled	up	sfp+ sr	-40.00	0 / 0	—	none	Disabled	N/A	
1/1/x2		Port is he...	N	10G	Enabled	up	sfp+ sr	-40.00	0 / 0	—	none	Disabled	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.

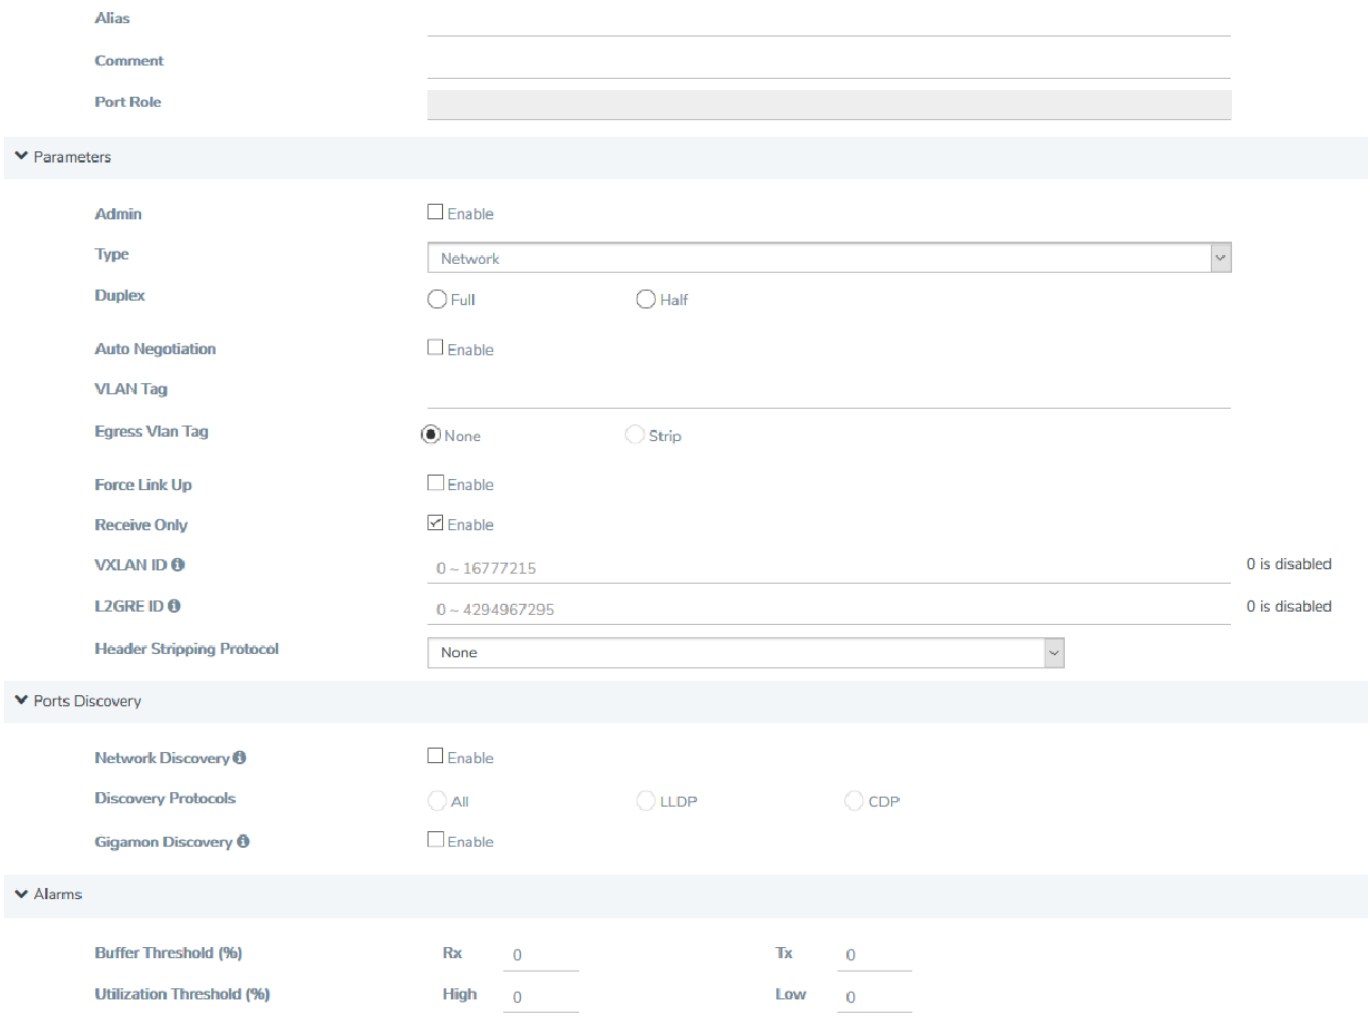

Utilization Alarm/SNMP Trap Generation

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**.

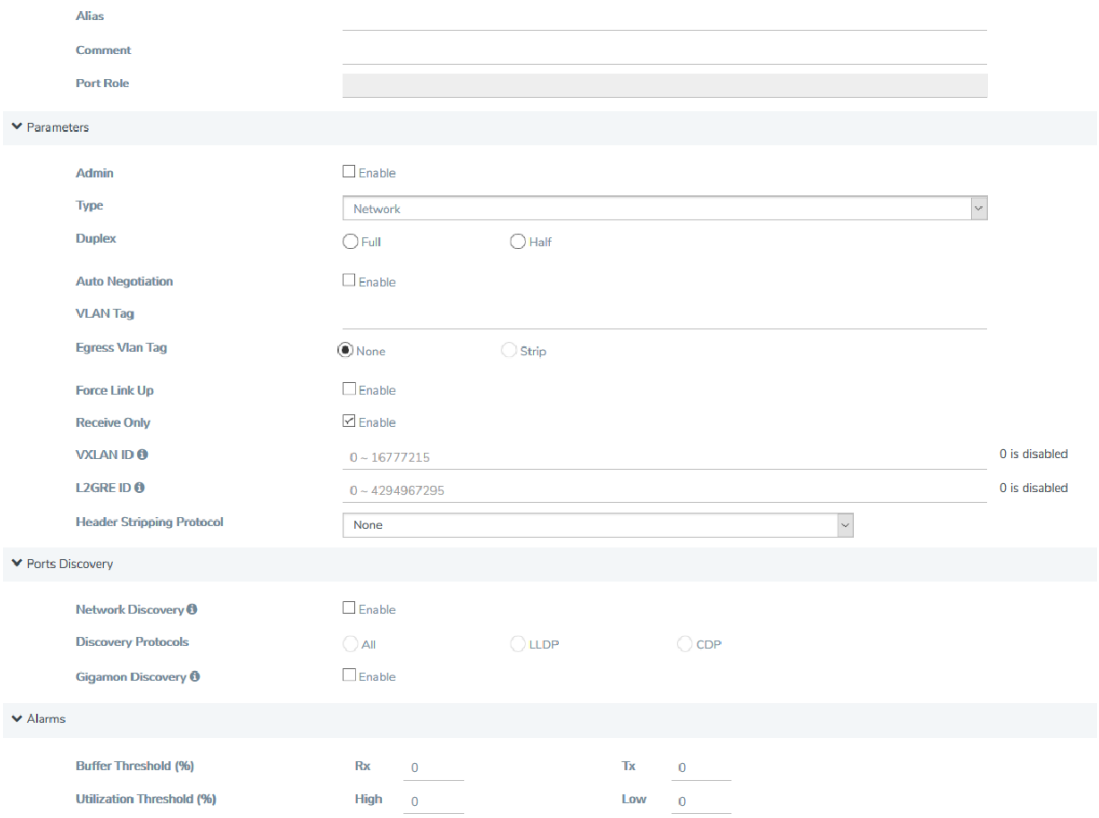

Microburst

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.

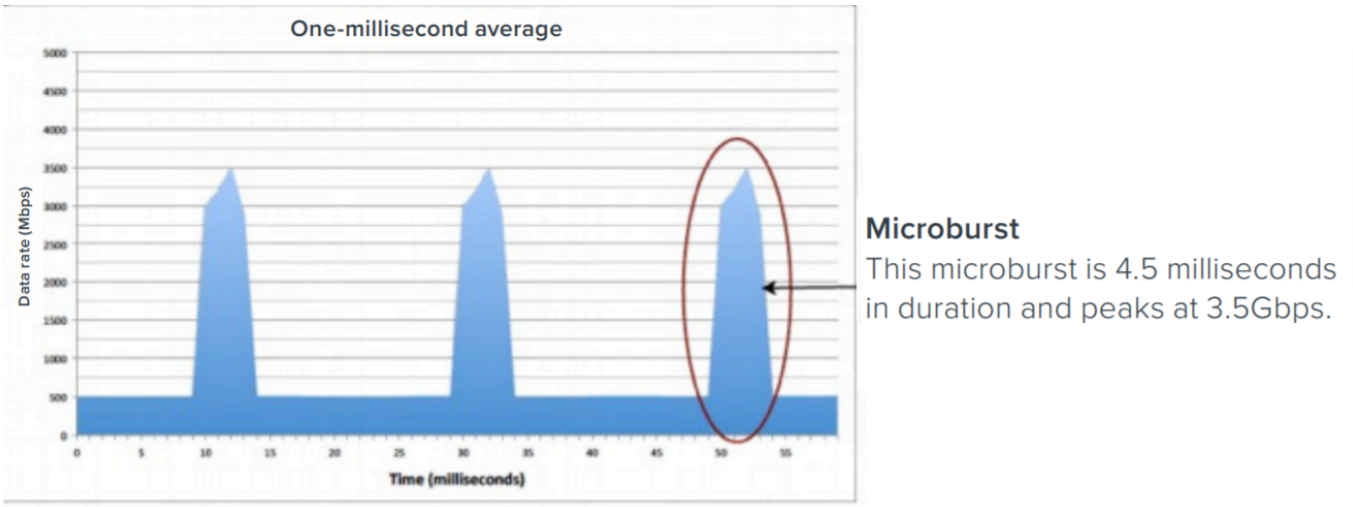

When the traffic is aggregated, packets are received at the egress port, usually from multiple ingress ports, in parallel rather than one after another. This results in queuing of packets. GigaVUE-OS devices have some amount of buffer to queue these packets. Each port accesses this buffer dynamically, 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:

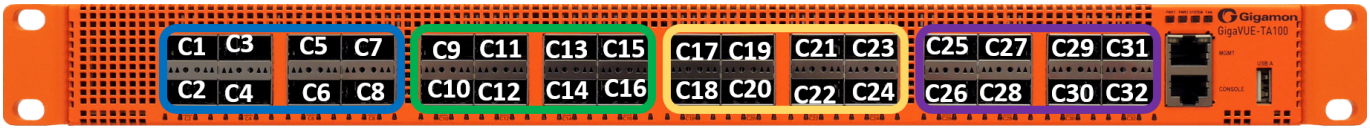

- GigaVUE-TA200—Following diagram illustrates the logical grouping in GigaVUE-TA200:

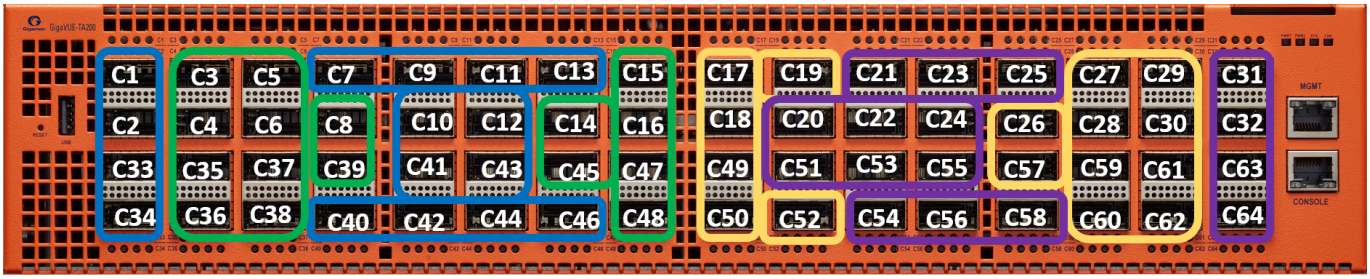

Legend

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 in the *GigaVUE Fabric Management Guide v5.13* or earlier.

NOTE: Classic Inline Bypass Solution is not supported in GigaVUE-HCT and GigaVUE-HC1-Plus devices.

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 10Flexible Inline Arrangements Scenario illustrates a complex inspection scenario that is enabled by flexible inline arrangements.

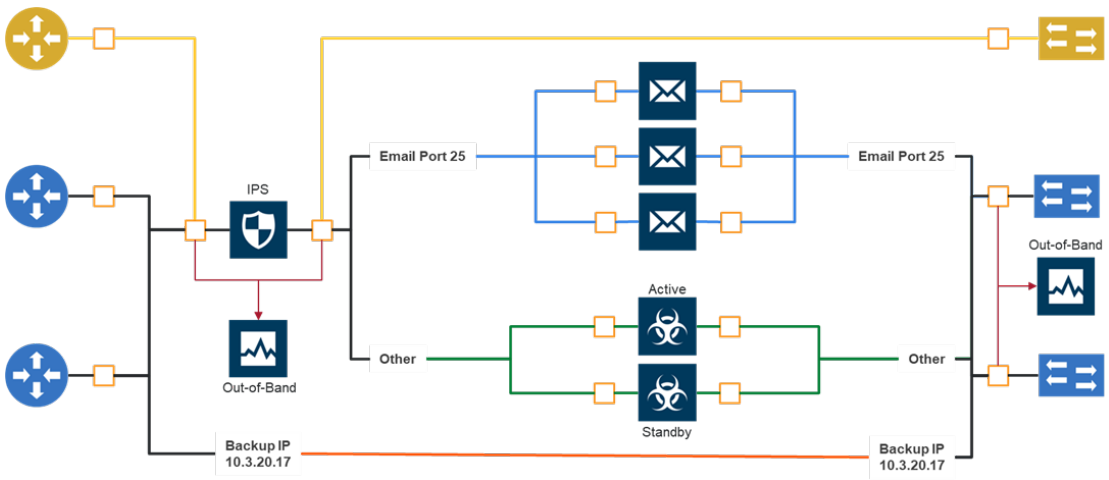

Figure 10 *Flexible Inline Arrangements Scenario*

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 10Flexible 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

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

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.

GRIP Supported by Flexible Inline Arrangements

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. Flexible inline arrangements support GRIP.

NOTE: Gigamon Resiliency for Inline Protection (GRIP™) is not supported in GigaVUE-HCT devices.

Flexible Inline Solution Supported in Clustered Nodes

The GigaVUE® HC Series, GigaVUE-TA25 and GigaVUE-TA200 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 10Flexible 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

Keep in mind the following when working with flexible inline arrangement:

- If an inline tool is associated with a flexible inline map, it cannot be used in a classic inline map or inline decryption map. All inline networks and inline tools must participate exclusively in either flexible inline maps or classic inline maps.
- You cannot create multiple unidirectional collector maps for the same inline network using the flexible inline canvas. For example, consider that you want to have different VLANs in each direction on the collector map, then you must create additional unidirectional flexible inline maps and associate unique VLAN tags (or it is automatically assigned by GigaVUE-FM).
- The following functionalities are not supported in GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA25E,GigaVUE-TA25 and GigaVUE-TA400 such as:
  - Physical Bypass Functionality is not supported due to the absence of BPS card.
  - Flexible and Resilient Inline SSL functionality is not supported due to the absence of the GigaSMART card.
  - GRIP functionality is not supported due to the absence of a BPS card.
  - ICAP functionality is not supported due to the absence of the GigaSMART card.
  - Inline Bypass Solution (Classic)
- In GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-HC1-Plus, when a Flexible Inline Single VLAN Tag is enabled in the map, below are the limitations:
  - inline-network traffic path BYPASS WITH MONITORING and inline-tool/inline-tool-group flex-traffic-path MONITORING cannot be configured.
  - oob-copy from inline-network cannot be configured
- For GigaVUE-HC1, GigaVUE-HC3, GigaVUE-HC3 (CCv1 & CCv2), GigaVUE-TA200E, GigaVUE-TA200 , GigaVUE-TA25E, GigaVUE-TA400 and GigaVUE-HC1-Plus nodes, the scalable number of bidirectional flex inline maps per device is 256, and the scalable number of unidirectional flex inline maps per device is 512.
- For GigaVUE-TA25, the limit is 126 scalable number of bidirectional maps and 252 for unidirectional maps.
- The following combinations are not supported:
  - Flexible Inline SSL Decryption
  - Inline Network LAG

    When you attempt to add an Inline SSL App to an Inline Network LAG Flexible Map you get the following error message: ***"An Inline SSL APP cannot be used when the traffic source is an inline network LAG"*** as shown in below figure.

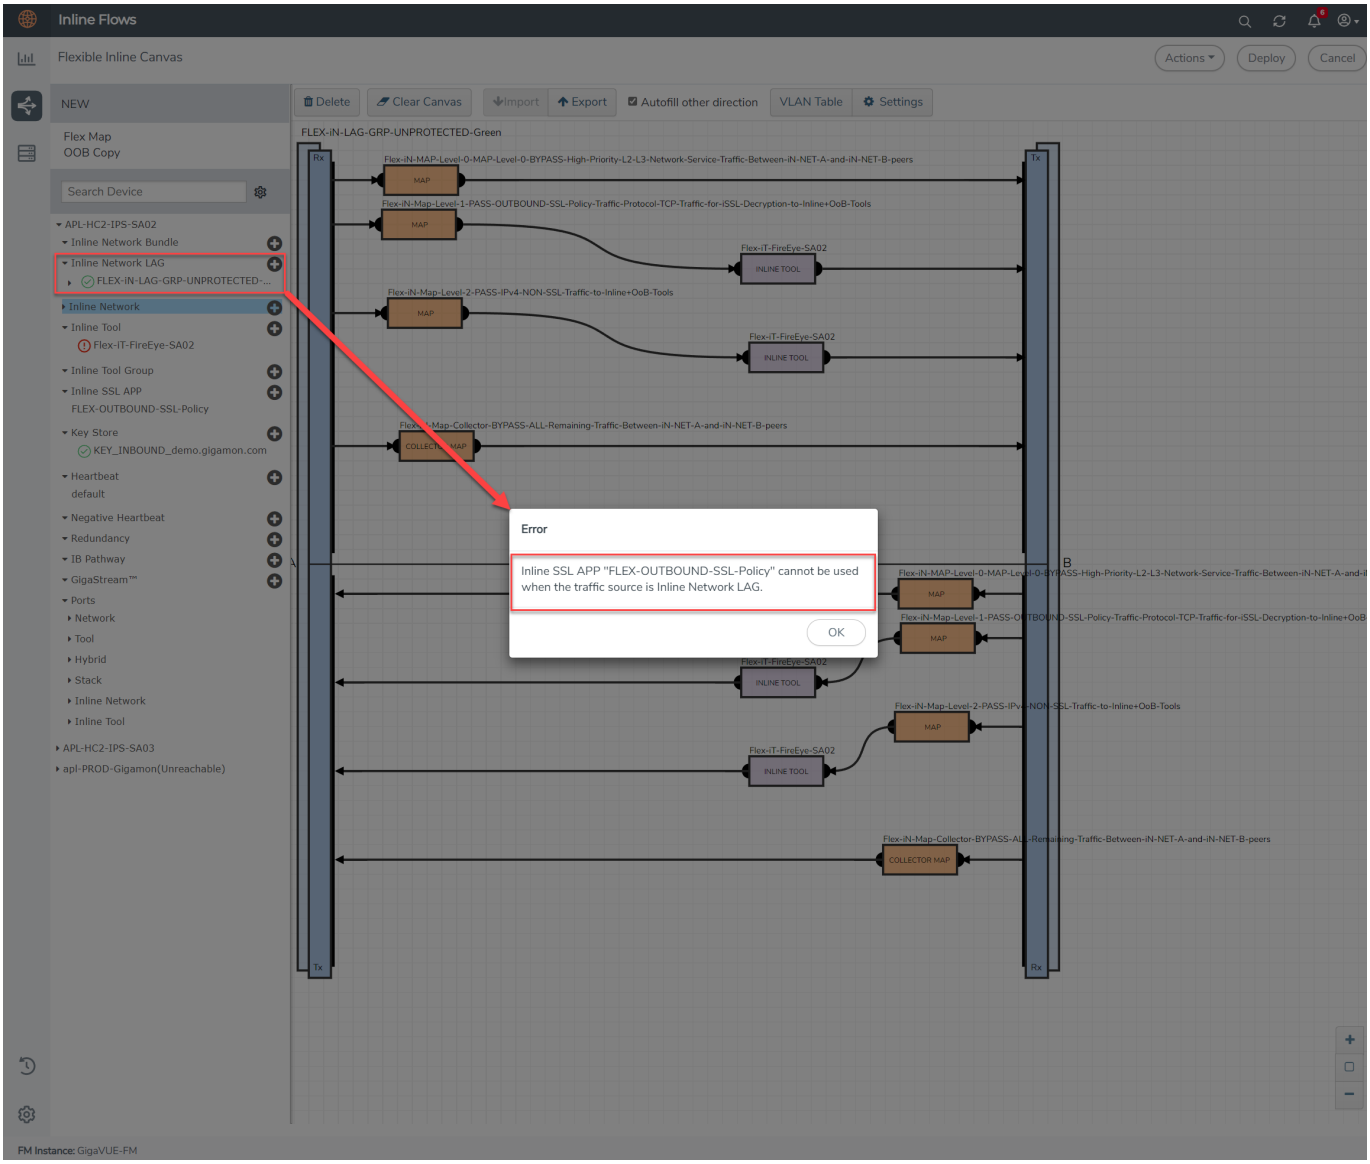

- Setting the Flex Traffic Path of inner chain Inline-tools as “Drop” does not drop the Inline SSL traffic and continues to reach the inline network egress.
- The Egress Port Filter on inline-network/tool/hybrid does not support VLAN-based filtering in GigaVUE-TA400.
- If an inline tool in a flex inline map is set to monitoring mode inGigaVUE-TA400, then all inline tools of the said map should be set to the same shared mode (true or false).
- Asymmetric hashing options: a-srcip-b-dstip and b-srcip-a-dstip on the inline-tool-group are not supported in GigaVUE-TA400.
- When an Inline Network Link Aggregation is configured as a source in GigaVUE-TA400, the CDP: pass-through option will not be supported, whereas the Bypass LACP/CDP/LLDP will be supported.

- For GigaVUE-TA400 the maximum Inline Network and Inline Tools is 48. The maximum Inline Network Link Aggregation List size is 24. The maximum Inline tool or Inline tool group per direction is 16. The maximum number of oob-copy entries per direction is 17. The maximum number of oob-copy ports per entry is 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

This section describes about the different flexible inline flows and provides step-by-step instructions on how to configure them using GigaVUE-FM. It also provides information about the forwarding states of the inline network.

Refer to the following sections for details:

- Configure Inline Network Ports and Inline Network
- Configure IP Interface
- Configure Inline Network Link Aggregation Group (LAG)
- Configure Inline Network Bundle
- Configure Inline Tool Ports and Inline Tools
- Configure Inline Tool Group
- Configure Inline Single Tag
- Configure Resilient Inline Arrangement
- Configure Hardware Security Model (HSM)
- Configure Flexible Inline TLS/SSL Decryption Solution
- Configure Internet Content Adaptation Protocol (ICAP)
- Configure Gigamon Resiliency for Inline Protection

Configure Inline Network Ports and Inline Network

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-HC series node.

Following are the two types of inline network:

- Unprotected inline network—It is an arrangement of two ports of the inline network type. The arrangement facilitates access to a bidirectional link between two networks (two far-end network devices) that need to be linked through an inline tool.
- Protected inline network—It is implemented using bypass combo modules. It is based on the pairs of ports associated with physical protection switches on the bypass combo modules. For a protected inline network, the ports are created automatically when the bypass combo modules are recognized by the GigaVUE® HC Series node.

To configure inline network ports and an inline network:

1. On the left navigation pane, go to **Physical > Orchestrated Flows > 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 inline network.
3. Click the '+' icon next to the **Inline Network** option to create a new inline network.

   Figure 11 *Inline Network Configuration*
4. In the **Alias** and **Description** fields, enter a name and description for the inline network, and then click **Port Editor**.
5. In the **Quick Port Editor**, scroll down to the inline network ports that you wish to configure. Select **Enable** to administratively enable inline network ports, and then click **OK**.
6. From the **Port A** and **Port B** drop-down lists, select the ports that you want to configure as the inline network pair.
7. From the **Traffic Path** drop-down list, select one of the following options:
   - **Bypass**—all traffic that originates from the inline network bypasses the sequence of inline tools and inline tool groups and is redirected to the opposite-side inline network port.
   - **Drop**—all traffic originating from the inline network is dropped.

- **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 redirected 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.

NOTE: If the Inline network is set to bypass with monitoring, GigaSMART Inline TLS/SSL decryption will not function.

8. Select the **Link Failure Propagation** check box to ensure that the inline network link failure on one side of the inline network is propagated to the other side. For details, refer to Network Port Link Status Propagation Parameter.
9. Select the **Accept Regular Heartbeat** check box to ensure that the inline network port pair accepts the heartbeat packets that are sent from the inline tool port pair. For details, refer to Heartbeat Support Between GigaVUE Nodes.
10. Click **OK** to save the configuration.
11. Drag the **Inline Network** object to the canvas and click **Deploy**.

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.

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.

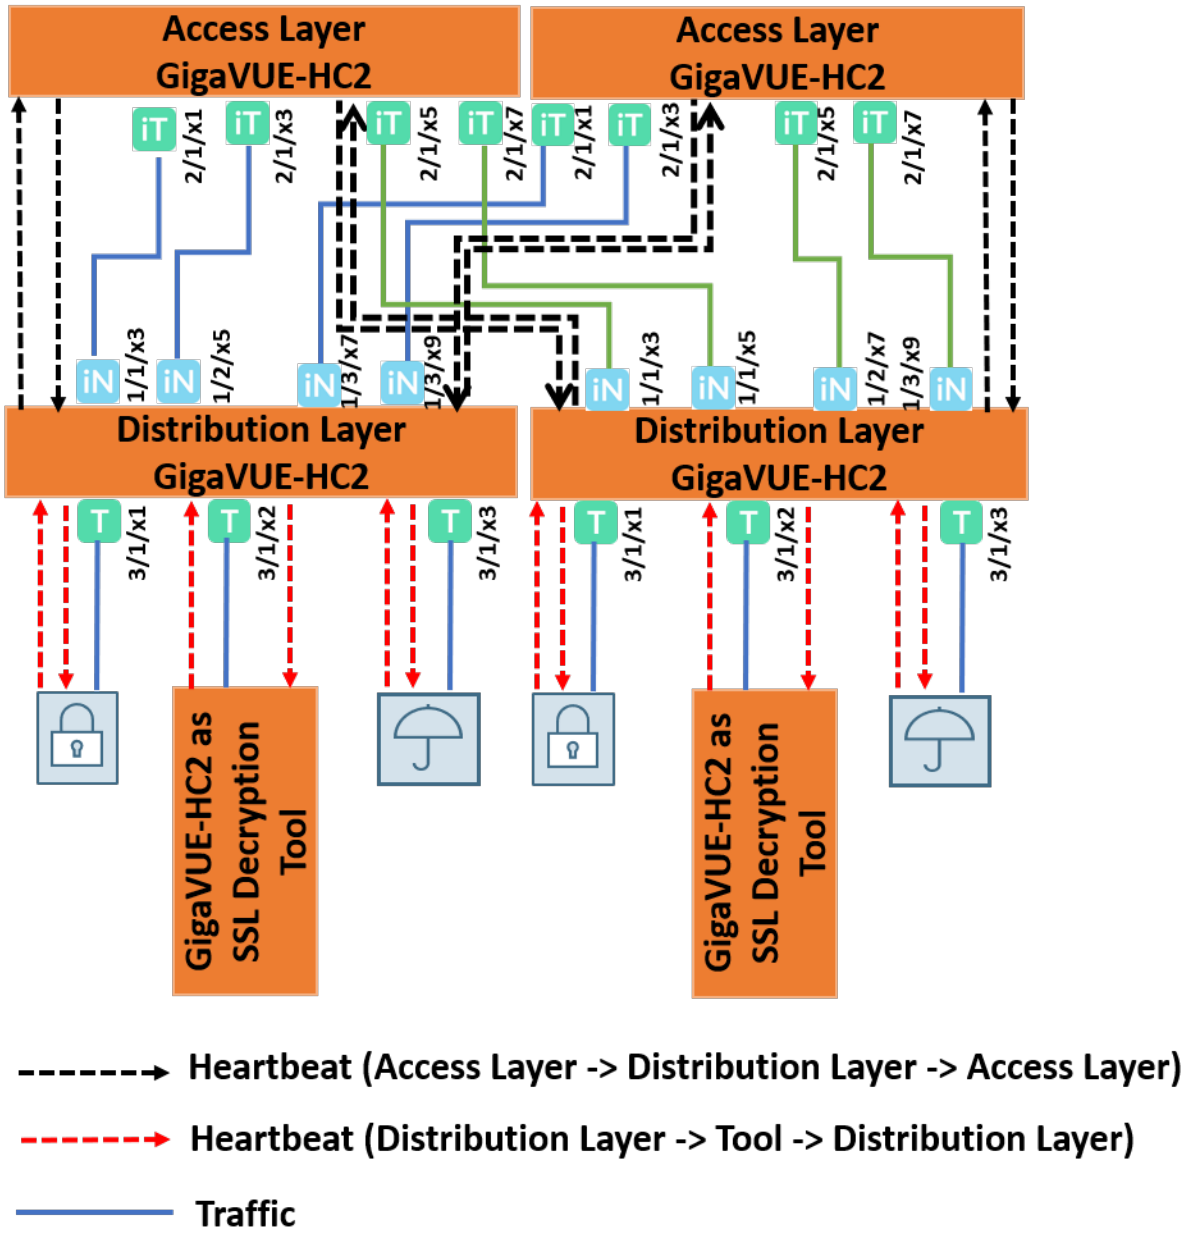

The GigaVUE node at the access layer accesses the network traffic, gets the traffic processed by the tools at the tool layer, and transmits the processed traffic back to the network.

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 TLS/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.

Configure IP Interface

To configure IP Interface:

1. On the left navigation pane, go to **Physical > Orchestrated Flows > 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 IP Interface.
3. Click the '+' icon next to the **IP Interface** option to create a new IP Interface.

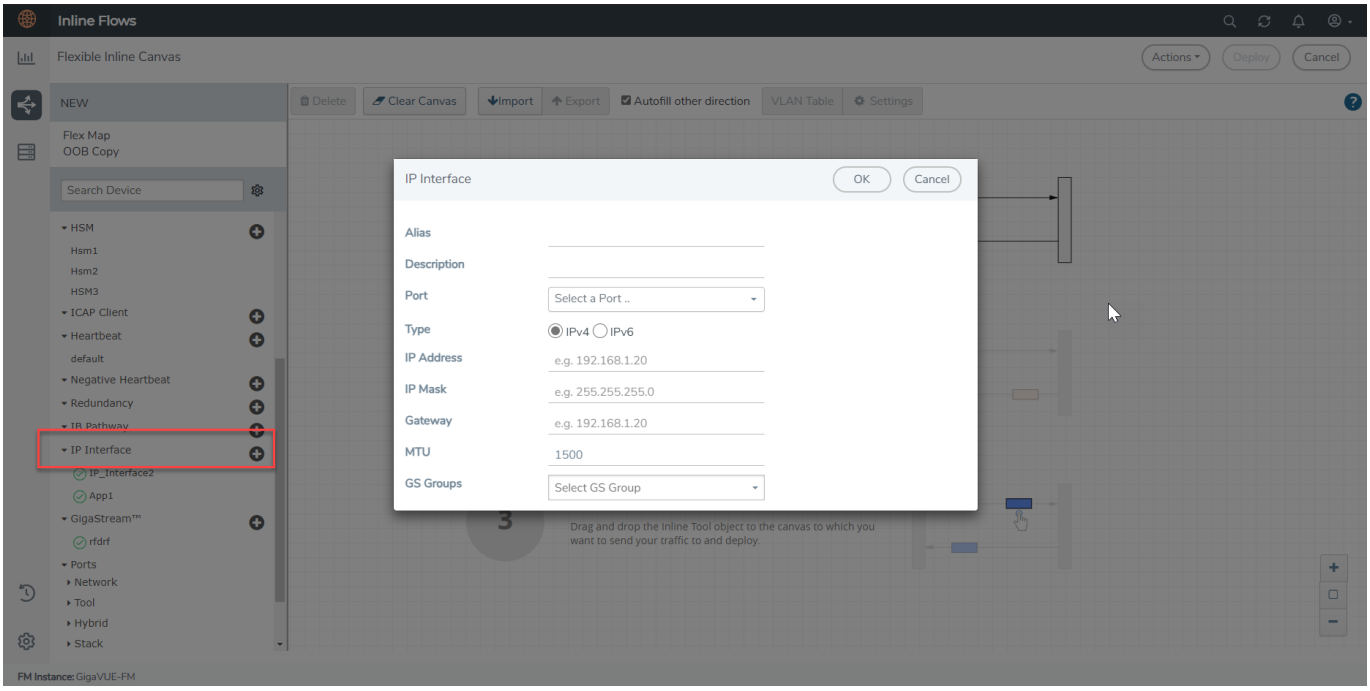

4. In the **Alias** and **Description** fields, enter a name and description for the IP Interface.
5. From the **Port** drop down list, select the port that you want to configure as the IP Interface.
6. Select the **Type** of the IP Interface that you want to configure.
7. Enter an **IP Address**. For example, 192.168.1.20.
8. Enter an **IP Mask**. For example, 255.255.255.0.
9. Enter a **Gateway**. For example, 192.168.1.20.
10. .Enter the Maximum Transmission Unit (MTU) for this port in the **MTU** field. For example, 1500.
11. Select the required GigaSMART Group you created from the **GS Groups** drop down list.
12. Click **OK** to save the configuration.

Configure Inline Network Link Aggregation Group (LAG)

Refer to the following sections that provide details about the inline network LAG, its limitations, and instructions on how to configure the inline network LAG:

- About Inline Network LAG
- Inline Network LAG—Rules and Notes
- Configure Inline Network LAG

About Inline Network LAG

A Link Aggregation Group (LAG) is a method of combining a number of physical ports together to make a single high-bandwidth data path, and thereby implement the traffic load sharing among the member ports in the group and to enhance the connection reliability. If you have a LAG in your network that must be inspected inline, the Flexible inline network LAG feature allows you to group the inline networks as one logical entity, instead of creating separate inline networks for each link in the LAG. Moreover, you can configure a flexible inline map with the inline network LAG as the source.

Traffic from the inline network LAG is grouped and sent to the inline tools with the same VLAN ID. The return traffic from the inline tools is hashed to the other side of the inline network LAG so that the incoming and the outgoing inline networks are different.

Each inline networks in an inline network LAG has their own specific forwarding states and traffic path settings. When one member link in the LAG goes down, the traffic is sent to the other member links.

Inline Network LAG—Rules and Notes

Keep in mind the following when working with inline network LAG:

- You cannot combine protected and unprotected inline networks, or inline networks with different speed in an inline network LAG.
- Inline TLS/SSL Decryption is not supported using Inline Network LAG. Use an Inline Network Bundle instead

  NOTE: It is highly recommended that the first flexible inline map bypasses the LACP/PAgP protocols between the inline network peers.

Refer to the example below:

```
# Flex Bypass Map configuration for Inbound and Outbound TLS/SSL Decryption:
-------------------------------------------------------------------------------------------------------------------------------
Flex-iN-MAP-L0-BYPASS-High-Priority-L2-L3-Network-Service-Traffic-Between-iN-NET-A-and-iN-NET-B-peers
rule add pass macdst 01:80:C2:00:00:00 ff:ff:ff:ff:ff:ff comment STP RST
rule add pass macdst 01:00:0c:cc:cc:cd ff:ff:ff:ff:ff:ff comment PVST+ RPVST+
rule add pass ethertype 8809 macdst 01:80:c2:00:00:02 FF:FF:FF:FF:FF:FF comment LACP
rule add pass macdst 01:00:0c:cc:cc:cc ff:ff:ff:ff:ff:ff comment CDP_DTP_PAgP_UDLD_VTP
rule add pass macdst 01:80:c2:00:00:0e ff:ff:ff:ff:ff:ff comment LLDP
rule add pass ipdst 224.0.0.2 255.255.255.255 portdst 1985 bidir comment HSRP
rule add pass ipdst 224.0.0.102 255.255.255.255 portdst 1985 bidir comment HSRPv2
rule add pass protocol 70 ipdst 224.0.0.18 255.255.255.255 comment VRRP
```

rule add pass protocol 59 comment OSPF-hex-59=decimal-89

rule add pass protocol 58 comment EIGRP-hex-58=decimal-88

rule add pass protocol tcp portdst 179 bidir comment BGP

rule add pass protocol udp portdst 520 bidir comment RIPv1

rule add pass protocol udp portdst 521 bidir comment RIPv2

 # Other Flex Inline Maps for Inbound and Outbound TLS/SSL Decryption:

Flex-iN-Map-L1-DROP-QUIC-Traffic-protocol-udp-portdst-80+443

rule add drop protocol udp prtdst 80 bidir

rule add drop protocol udp prtdst 443 bidir

Flex-iN-Map-L1-PASS-INBOUND-Decrypt-SSL-Traffic-protocol-tcp+ip+prtdst-80+443-bidir-to-Inline+OoB-Tools

rule add pass protocol tcp ipdst 192.0.2.100 255.255.255.255 prtdst 80 bidir comment "Support StartTLS for HTTP (RFC 2817) midstream SSL/TLS decryption"

rule add pass protocol tcp ipdst 192.0.2.100 255.255.255.255 prtdst 443 bidir

 # Important Note:

 # PATH: FM: Traffic > Orchestration > Inline Flow > Flexible Inline Canvas > Inline SSL App >

 # Inline SSL App Name: Flex-INBOUND-SSL-Policy > starttls add l4port 80 (see flex-image-1 below)

Flex-iN-Map-L1-PASS-OUTBOUND-Decrypt-SSL-Traffic-protocol-tcp-to-Inline+OoB-Tools

rule add pass protocol tcp comment "Gigamon does not support QUIC UDP 443 + UDP 80 protocol decryption and should be dropped to force TLS/SSL traffic to the TCP protocol"

Flex-iN-Map-L2-PASS-IPv4-NON-SSL-Traffic-to-Inline+OoB-Tools rule add pass ipver 4

Flex-iN-Map-Collector-BYPASS-ALL-Remaining-Traffic-Between-iN-NET-A-and-iN-NET-B-peers

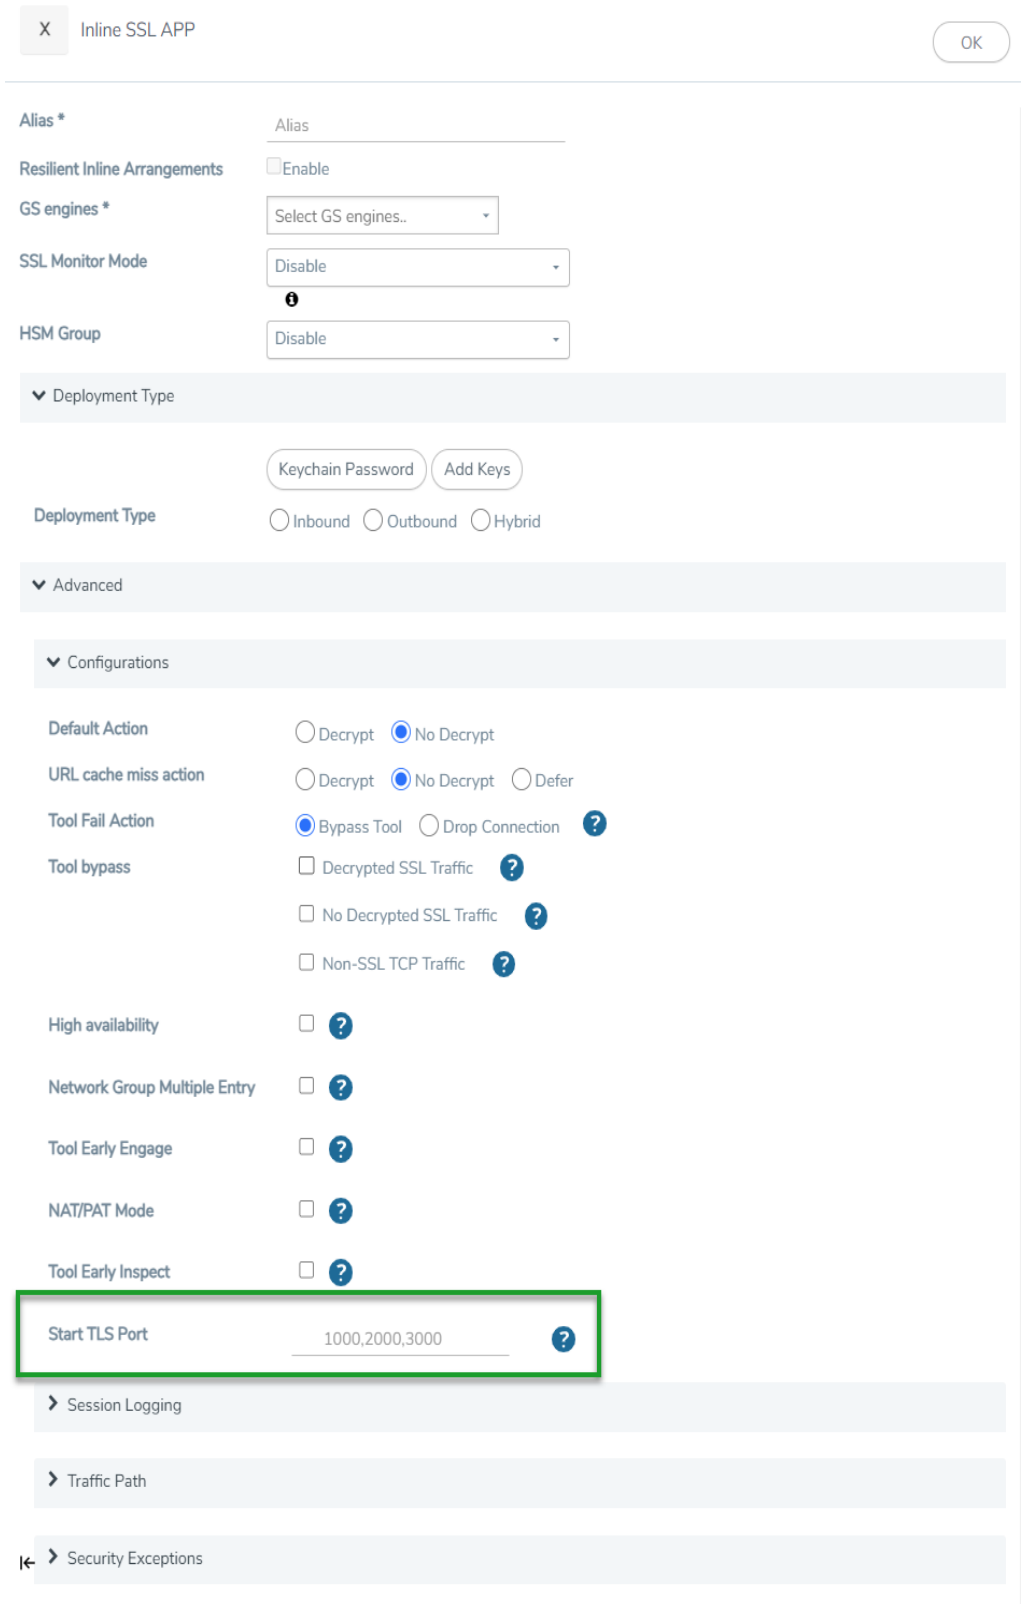
X
Inline SSL APP
OK
Alias *
Alias
Resilient Inline Arrangements
Enable
GS engines *
Select GS engines..
SSL Monitor Mode
Disable
HSM Group
Disable
Deployment Type
Keychain Password
Add Keys
Deployment Type
Inbound
Outbound
Hybrid
Advanced
Configurations
Default Action
Decrypt
No Decrypt
URL cache miss action
Decrypt
No Decrypt
Defer
Tool Fail Action
Bypass Tool
Drop Connection
Tool bypass
Decrypted SSL Traffic
No Decrypted SSL Traffic
Non-SSL TCP Traffic
High availability
Network Group Multiple Entry
Tool Early Engage
NAT/PAT Mode
Tool Early Inspect
Start TLS Port
1000,2000,3000
Session Logging
Traffic Path
Security Exceptions

Configure Inline Network LAG

Before you configure an inline network LAG, ensure that you configure the required inline network ports and inline networks. Refer to Configure Inline Network Ports and Inline Network.

To configure an inline network LAG:

1. On the left navigation pane, go to **Physical > Orchestrated Flows >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 inline network LAG.
3. Click the '+' icon next to the **Inline Network LAG** option to create a new inline network LAG.

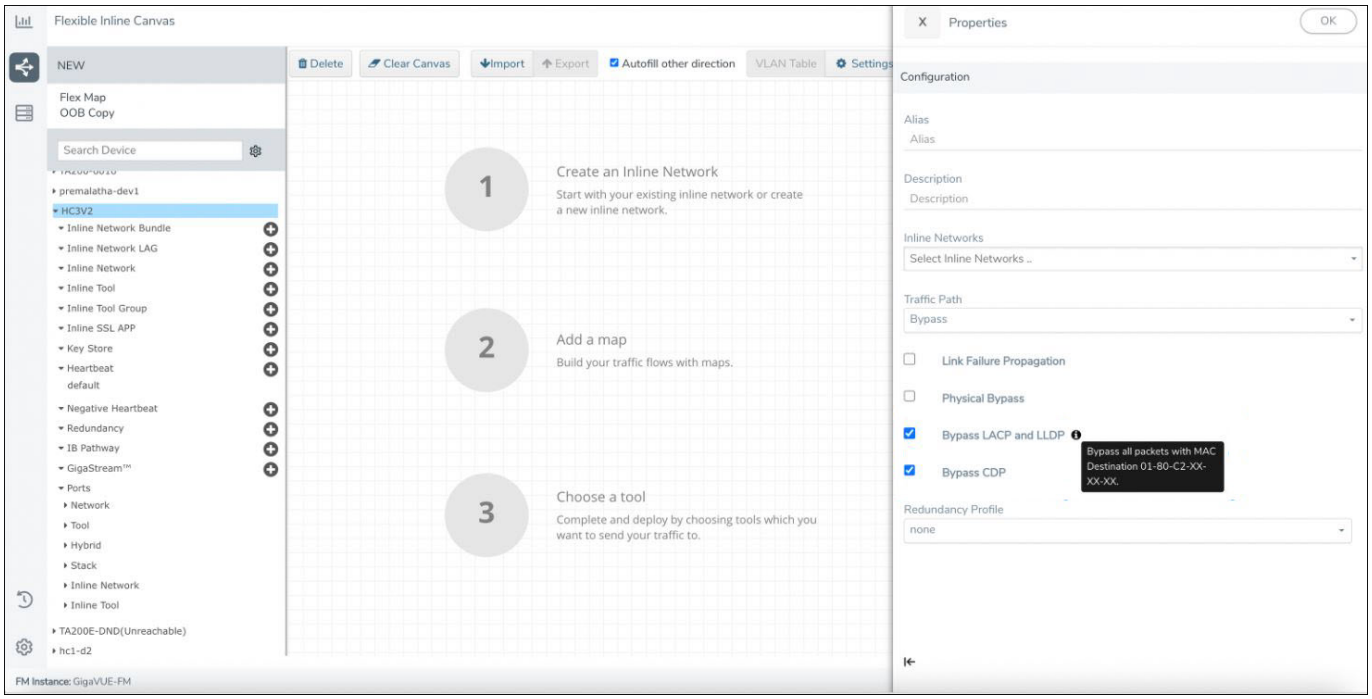

4. In the Properties pane that appears on the right, enter the name and description of the inline network LAG in the **Alias** and **Description** fields.
5. From the **Inline Networks** drop-down list, select the required inline networks that need to be part of the inline network LAG.
6. From the **Traffic Path** drop-down list, select one of the following options:
   - **Bypass**—All traffic arriving at Port A of the inline network is directly forwarded to Port B of the inline network. Similarly, all traffic arriving at Port B of the inline network is directly forwarded to Port A of the inline network.

- **Drop**—Traffic is not exchanged between 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**—Traffic is forwarded to the sequence of inline tools.

7. Select the **Link Failure Propagation** check box if you want to bring down a port when its pair goes down.
8. Select the **Physical Bypass** check box if you want the traffic to flow directly between Port A and Port B of the inline network pair when a device or a module is powered down.
9. If there is a group of links, which are part of a port channel that use LACP, select the **Bypass Link Aggregation Control Protocol and Link Layer Discovery Protocol** check box to maintain the port channel functionality on the links that are connected to inline network LAG ports.

   NOTE: Inline Network LAG needs a bypass map to handle LACP bypass.

   NOTE: When the Bypass LACP and LLDP checkbox is enabled, all protocol packets with MAC Destination 01-80-C2-XX-XX-XX are bypassed.

10. If there is a group of links, which are part of a port channel that supports CDP, select the **Bypass Cisco Discovery Protocol** check box to maintain the CDP discovery functionality on the links that are connected to inline network LAG ports.

    NOTE: Inline Network LAG needs a bypass map to handle CDP bypass.

    NOTE: When CDP/LLDP bypass is enabled, the CDP/LLDP neighborship discovery will not be established on the respective inline networks.

11. Click **OK** to save the configuration.
12. Drag the **Inline Network LAG** object to the canvas.
13. Configure the required flexible inline maps and then, click **Deploy**.

Configure Inline Network Bundle

The Flexible inline network bundle feature allows you to group multiple inline networks into an inline network bundle. You can configure flexible inline maps with the inline network bundle as the source. GigaVUE-FM configures separate inline maps for each inline networks

that are grouped in the inline network bundle. The inline maps are configured based on the Tool Side VLAN tags for the multiple inline networks and the rules that you specified when configuring the inline map for the network bundle.

Before you configure an inline network bundle, ensure that you configure the required inline network ports and inline networks. Refer to Configure Inline Network Ports and Inline Network.

To configure an inline network bundle:

1. On the left navigation pane, go to **Physical > Orchestrated Flows > 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 inline network bundle.
3. Click the '+' icon next to the **Inline Network Bundle** option to create a new inline network bundle.
4. In the **Alias** field, enter the name of the inline network bundle.
5. From the **Inline Networks** drop-down list, select the required inline networks that you want to add to the inline network bundle.
6. Click **OK** to save the configuration.
7. Drag and drop the inline network bundle into the canvas, and then configure the required inline map.
8. Enter the **Tool Side VLAN Tag** for each inline network added in the inline network bundle.
9. 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.
10. Add the required rules for the inline map, and then click **OK** to save the configuration.
11. Click **Deploy**. GigaVUE-FM configures separate inline maps for each inline networks that are grouped in the inline network bundle.

Configure Inline Tool Ports and Inline Tools

An inline tool consists of a pair of inline tool ports that run at the same speed, on the same medium (fiber or copper). Both the inline tool ports must be on the same GigaVUE-HC series node. Moreover, the inline tool ports must be on the same GigaVUE-HC series node in which the inline network ports reside. The inline tools are attached to the inline tool ports.

An inline tool can also be a pass-through device that performs packet inspection and selective forwarding, such as Intrusion Protection System (IPS). This is a physical device, external to the GigaVUE HC series node.

To configure the inline tool ports and the inline tools:

1. On the left navigation pane, go to **Physical > Orchestrated Flows > 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 inline tool.
3. Click the '+' icon next to the **Inline Tool** option to create a new inline tool.
4. In the **Properties** pane, in the **Alias** and **Description** fields, enter a name and description for the inline tool.
5. From the **Type** drop-down list, select one of the following options:
   - **External**—To configure a third-party tool.
   - **GigaVUE Node**—To configure a GigaVUE node as a tool.
6. Click **Port Editor**, and in the **Quick Port Editor**, scroll down to the inline tool ports that you wish to configure. Select **Enable** to administratively enable the inline tool ports, and then click **OK**.
7. From the **Port A** and **Port B** drop-down lists, select the inline tool ports according to the direction the inline tool expects traffic from the network.
8. Verify that the **Enabled** check box is selected.
9. From the **Failover action** drop-down list, select one of the following options:
   - **Tool Bypass**—For every map involving the inline tool or inline tool group that triggered this failover action, the traffic coming to such an inline tool or inline tool group is redirected to the next inline tool or inline tool group in the ordered list defined in Port A and Port B or to the respective inline network port.
   - **Network Bypass**—All inline networks configured as the source of any map involving the inline tool or inline tool group that triggered this failover action, will be put in the bypass mode, that is, all traffic coming to side A will be directed to side B and vice versa.
   - **Tool Drop**—For every map involving the inline tool or inline tool group that triggered this failover action, the traffic coming to such an inline tool or inline tool group is dropped and the traffic is redirected to a dummy VLAN with no members.

     NOTE: When failover-action 'drop' is triggered for an inline-tool present in the Flex iSSL solution, all the traffic entering GigaSMART is dropped at vport.

   - **Network Drop**—All inline networks configured as the source of any map involving the inline tool or inline tool group that triggered this failover action, will be put in the drop mode, that is, all traffic coming to side A or side B will be dropped.

- **Network Port Forced Down**—For all inline networks configured as the source of any map involving the inline tool or inline tool group that triggered this failover action, the inline network ports will be brought down.

10. Select the **Recovery Mode** as **manual** or **automatic**.
11. Select the **Enable** check box for the **Inline tool Sharing mode** if you want to define additional tags on the tool side.

    NOTE: If you choose to disable the **Inline tool Sharing mode,** the inline tool can be used only in one flexible inline map.

12. From the **Flex Traffic Path** drop-down list, select one of the following options:
    - **Drop**—Traffic is dropped at the inline tool.

    NOTE: When failover-action 'drop' is triggered for an inline-tool present in the Flex iSSL solution, all the traffic entering GigaSMART is dropped at vport.

    - **Bypass**—Traffic bypasses the inline tool. Use this option for performing maintenance on an inline tool.
    - **Monitoring**—Traffic is fed to the inline tool and absorbed, while a copy of the traffic is sent to the next inline tool in the sequence. Traffic returned from side B of the network is also absorbed at the inline tool in the monitoring mode.
    - **To Inline Tool**—Traffic is forwarded to the inline tool.
13. Select the **Enable** check box below the **Regular Heartbeat**, if required, and then from the **Regular Heartbeat Profile** drop-down list, select a suitable profile.
14. In the HB IP Address A and HB IP Address B fields, enter the IP address of side A and side B defined in the Heartbeat profile.
15. Select the **Enable** check box below the **Negative Heartbeat**, if required, and then from the **Negative Heartbeat Profile** drop-down list, select a suitable profile.
16. Click **OK** to save the configuration.
17. Drag the **Inline Tool** object to the canvas.
18. Configure the required flexible inline maps and then, click **Deploy**.

Configure Inline Tool Group

An inline tool group is an arrangement of multiple inline tools. Traffic is distributed to the inline tools that are part of an inline tool group based on hardware-calculated hash values. For example, if one tool in a group goes down, traffic is redistributed to other tools in the group using hashing. You can also configure redundancy, such as 1+1 and N+1.

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-Plus, node, but can be on different modules on the node. On the GigaVUE-HC1, all the inline tool ports of the inline group must be on either the base module or the bypass combo module. The inline tool ports must also be on the same GigaVUE-HC3, GigaVUE-HC1-Plus, or GigaVUE-HC1 node as the inline network ports.

When an inline tool is removed from an inline tool group , removed inline tool must be added as an Inline Spare tool . This would require you to first delete the inline tool from the tool group and save the changes .Then add the removed tool as a spare tool from the canvas.

In a cluster environment, you can configure the inline tool group on GigaVUE® HC Series nodes through the cluster leader. The inline constructs must all be configured on one GigaVUE-HC3, GigaVUE-HC1-Plus,, or GigaVUE-HC1 node, not across nodes, even if the nodes are in a cluster.

Resilient weighted hashing provides you the ability to distribute traffic to the inline tools by assigning either an equal weight or a custom weight 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 weight or the custom weight 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.

Before you configure an inline tool group, ensure that you configure the required inline tools. Refer to Configure Inline Tool Ports and Inline Tools.

To configure an inline tool group:

1. On the left navigation pane, go to **Physical > Orchestrated Flows > 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 inline tool group.
3. Click the '+' icon next to the **Inline Tool Group** option to create a new inline tool group.
4. In the **Properties** pane, in the **Alias** and **Description** fields, enter a name and description for the inline tool group.
5. From the **Inline Tools** drop-down list, select the required inline tools.
6. From the **Weighting** drop-down list, select one of the following options:

- **Equal**—Traffic is distributed equally to all the inline tools in the inline tool group. This is the default option.
- **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. However, asymmetrical hashing is not supported for **Relative** or **Percentage** options.

7. From the **Inline Spare Tool** drop-down list, select the inline tool to which the traffic will be forwarded when the first failure occurs in the set of primary inline tools. To update the Inline Spare Tool list by adding a removed inline tool, delete the inline tool from the inline tool group. Then from the canvas under the Inline Tool group option update the released tool as a spare tool.

   NOTE: You cannot select an inline spare tool if you have selected a **Weighting** option.

8. Select the **Enabled** check box to make the inline tool group available for deployment.
9. Select the **Release Spare if Possible** check box to ensure that the inline spare tool is released from the active set of tools to become the spare again when the primary inline tool recovers from the failure.
10. From the **Failover Action** drop-down list, select one of the following options:
    - **Tool Bypass**—For every map involving the inline tool or inline tool group that triggered this failover action, the traffic coming to such an inline tool or inline tool group is redirected to the next inline tool or inline tool group in the ordered list defined in Port A and Port B or to the respective inline network port.
    - **Network Bypass**—All inline networks configured as the source of any map involving the inline tool or inline tool group that triggered this failover action, will be put in the bypass mode, that is, all traffic coming to side A will be directed to side B and vice versa.
    - **Tool Drop**—For every map involving the inline tool or inline tool group that triggered this failover action, the traffic coming to such an inline tool or inline tool group is dropped.
    - **Network Drop**—All inline networks configured as the source of any map involving the inline tool or inline tool group that triggered this failover action, will be put in the drop mode, that is, all traffic coming to side A or side B will be dropped.
    - **Network Port Forced Down**—For all inline networks configured as the source of any map involving the inline tool or inline tool group that triggered this failover action, the inline network ports will be brought down.

11. From the **Failover Mode** drop-down list, select **Spread** to redistribute all the traffic coming from the inline network (or inline network group) to the active inline tools (excluding the failed inline tool or tools).

    NOTE: This field is not applicable when there is only one inline tool in the tool list.

12. From the **Minimum Healthy Group Size** drop-down list, select the minimum number of inline tools that must be up so that the entire inline tool group is considered to be up. The minimum number must include the inline spare tool as well.
13. From the **Hash** drop-down list, select one of the following options to distribute packets across a number of inline tools that belong to the inline tool group:
    - **Advanced**—Specifies symmetrical hashing, which is derived from the combination of packet fields based on the criteria selected for the advanced-hash algorithm. 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.
    - **SideA as sourceIP & SideB as destinationIP**—Specifies asymmetrical hashing, which is derived from the source IP address for side A of the inline network and the destination IP address for side B of the inline 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.
    - **SideB as sourceIP & SideA as destinationIP**—Specifies asymmetrical hashing, which is derived from the destination IP address for side A of the inline network and the source IP address for side B of the inline 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.

    NOTE: This field is not available for selection if you have selected the **Relative** or **Percentage** options in the **Weighting** drop-down list.

14. From the **Flex Traffic Path** drop-down list, select one of the following options:
    - **Drop**—Traffic is dropped at the inline tool group.
    - **Bypass**—Traffic bypasses the inline tool group. Use this option for performing maintenance on an inline tool group.
    - **Monitoring**—Traffic is fed to the inline tool group and absorbed, while a copy of the traffic is sent to the next inline tool in the sequence. Traffic returned from side B of the network is also absorbed at the inline tool group in the monitoring mode.
    - **To Inline Tool**—Traffic is forwarded to the inline tool group.
15. Click **OK** to save the configuration.
16. Drag the **Inline Tool Group** object to the canvas.
17. Configure the required flexible inline maps and then, click **Deploy**.

Configure Inline Single Tag

During the flexible Inline Bypass operations, network traffic sent to Inline-tools contains an extra VLAN tag. The extra VLAN tag is used to help distinguish the traffic coming from inline-tools and to make sure traffic is routed to the right inline networks. Using extra VLAN tags often presents problems for various inline-tools. The flexible Inline single tag can be used to replace the extra VLAN tag on incoming traffic. Using flexible inline single tags, you can map incoming VLANs on the network side to the outgoing VLANs on the tool side.

NOTE: The **OOB Copy** tag attribute **none** is invalid for single tag maps. The attribute **original** should be used.

To configure an inline single tag:

1. On the left navigation pane, go to **Physical > Orchestrated Flows > 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 inline network from the list of devices.
3. Drag and drop the inline network into the canvas.
4. Click **Settings** to open the **Settings** pane.
5. Select the **Enable** check box for the **Show Single Tag Options**, and then enter the VLANS expected on the inline network.
6. Drag and drop a flexible inline map object into the canvas, and then click the map to open the **Properties** pane.
7. Select the **Enable** check box for the **Single Tag Mode**, and then enter the tool side VLAN tags. 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. The VLAN qualifier is added to the rules by GigaVUE-FM. If you do not specify any rules, GigaVUE-FM adds a rule with the VLAN qualifier to the map. You can use the check box 'Enable Network side VLAN Tag' to bulk enable or disable the VLAN tags added for flexinline maps.

   NOTE: You can choose to enable or disable the **Single Tag Mode** for collector maps, if required.

8. Drag and drop the required inline tools into the canvas.
9. Drag and drop the OOB Copy into the canvas, if required.

10. Click **Port Editor**, and then in the **Quick Port Editor**, scroll down to the hybrid or tool ports that you wish to configure. Select **Enable** to administratively enable the ports, and then click **OK**.
11. From the **Destination Ports** drop-down list, select the required hybrid or tool ports that you want to configure as destination ports. You can also select a hybrid or tool GigaStream. For information about GigaStream, refer to *"How to Use GigaStream" section*.
12. From the **VLAN Tag** drop-down list, select the required tags.
13. Click **OK** to save the configuration.
14. Click **Deploy**.

After deploying the solution , any unused VLAN Tags will be preserved and displayed. The VLAN tags are displayed in the same order that they were configured. Similarly, the VLAN Table window would list the maps as displayed in the configuration canvas.

NOTE: In GigaVUE-HC1-Plus, if the inline networks traffic path is in monitoring mode, you cannot enable the Single Tag Mode.

NOTE: If you enable the Single VLAN tag option in Flexible iSSL solution, you should also enable the Single VLAN Tag configuration in the inline-ssl app profile deployed in the solution

Configure Resilient Inline Arrangement

Refer to the following sections that provide details about the resilient inline arrangement feature and instructions on how to configure it:

- Resilient Inline Arrangement
  - Resilient Inline Arrangement With Single VLAN Tag
  - Resilient Inline Arrangement—Classic
  - Inter-broker Pathway (IB-P)
- Resilient Inline Arrangement—Rules and Notes
- Deploy Resilient Inline Arrangement

Resilient Inline Arrangement

Resilient inline arrangement is a method of configuring and deploying inline threat prevention tools for dual-path, redundant network architectures. A successful deployment of resilient inline arrangements provides traffic management for dual-path high availability environments.

NOTE: Resilient inline arrangement is not supported in GigaVUE-HCT devices.

The following figure illustrates the resilient inline arrangement.

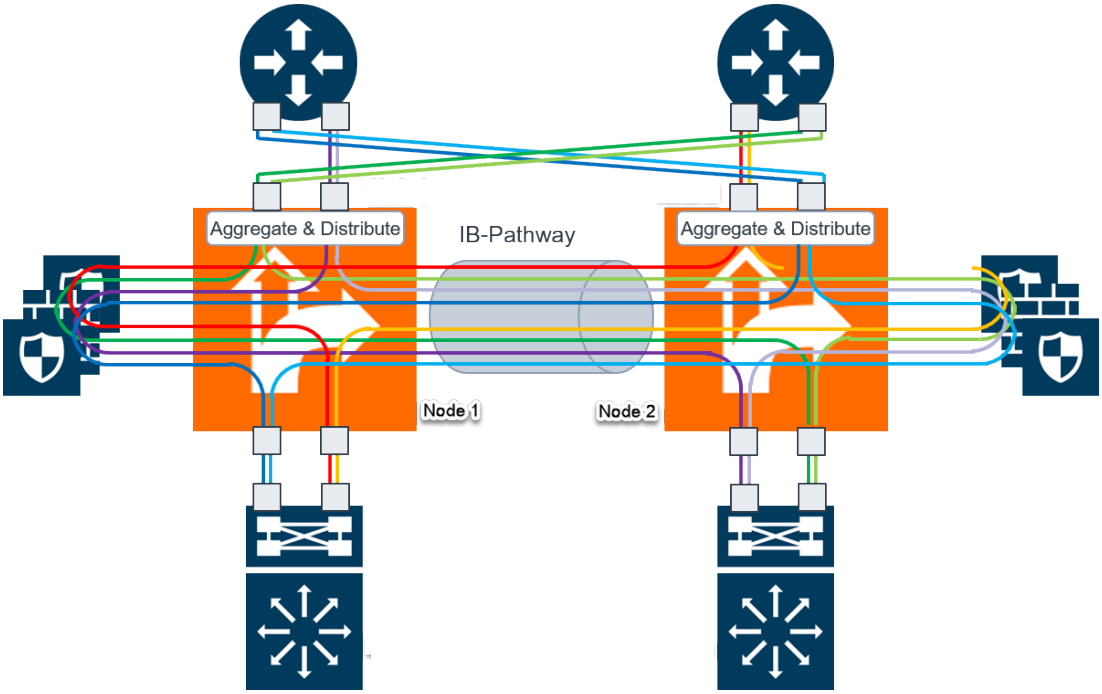

Figure 12 *Resilient Inline Arrangement*

The resilient inline arrangement shows the Gigamon devices, which consolidate the traffic from multiple intercepted links before routing the traffic to inline tools. To protect such an inspection arrangement from any failure of the Gigamon devices, a redundant arrangement of inline packet broker is shown. Both the inline packet brokers are interconnected by an Inter-broker Pathway (IB-P). For details, refer to Inter-broker Pathway (IB-P).

Each inline packet broker is attached to a set of inline tools that are identical to each other, that is, both inline packet brokers must have equal number of inline tools. Moreover, the inline tools on both sides must be of the same type, port speed, and processing capacity.

Resilient inline arrangement is based on an aggregation and distribution principle that divides the packets received by an inline packet broker, between Node 1 and Node 2. The inline packet broker on the left, guides the Node 1 class of packets through its local tools and

Node 2 class of packets through the remote tools that are reachable by a resilient inter-broker pathway. Similarly, the inline packet broker on the right, guides the Node 2 class of packets through its local tools and Node 1 class of packets through the remote tools.

Each link intercepted by the inline packet broker must be configured with the following component maps:

- either a bidirectional original component map or two unidirectional original component maps,
- two unidirectional export component maps, and
- two unidirectional import component maps.

GigaVUE-FM configures the required export and import component maps for all the links that are intercepted by both the inline packet brokers. GigaVUE-FM configures the maps based on the tool side VLAN tags and the rules that you specified when configuring the flexible inline map.

The component maps use VLAN tags to transfer the traffic from inline network to inline tools and back through the inter-broker pathway. Refer to the following sections:

- Resilient Inline Arrangement—Classic
- Resilient Inline Arrangement With Single VLAN Tag

Resilient Inline Arrangement—Classic

When a packet is received from an inline network, an additional VLAN tag is added to the packet before guiding it to the inline tools. The additional VLAN tag is useful when the inline tools are shared by multiple traffic flows. It helps to distinguish the traffic coming from inline-tools and to make sure the traffic is routed to the right inline networks. You can configure the additional VLAN tags when you create the flexible inline maps.

Resilient Inline Arrangement With Single VLAN Tag

You can choose to deploy a resilient inline arrangement with a single VLAN tag in which a packet received from an inline network is guided to the inline tool using a single VLAN tag, which you can configure when creating a flexible inline map. You must configure the packet's original VLAN tag as the network side VLAN tag and provide the required tool side VLAN tag when you create the flexible inline maps. The single VLAN tag is useful when your inline tools do not support Q-in-Q VLAN tags.

You can configure a Flexible Inline SSL and RIA iSSL solution with Single VLAN Tagging (SVT).

The following table explains the compatibility matrix between single VLAN tag enabled and disabled maps. Symbol (√) denotes the engine ports that are supported, and symbol (X) denotes the engine ports that are not supported.

Maps	SVT enabled RIA iSSL map	
	same gs_engine	different gs_engine in different maps
RIA	√	X
RIA + SVT	√	X
RIA + iSSL	√	X

Inter-broker Pathway (IB-P)

The inter-broker pathway provides link aggregation and distribution and is responsible for moving traffic between Node 1 and Node 2. You must configure tool ports in the inter-broker pathway. Following are the IB-P states:

- inter-broker pathway-up—the traffic is handled as follows:
  - If the traffic is governed by the original component maps in which the traffic path is set to Bypass, the traffic bypasses the sequence of inline tools and inline tool groups and is re-directed to the inline network port that is configured on the opposite-side.
  - If the traffic is governed by the export component maps in which the traffic path is set to any value other than Bypass, the traffic is routed through the inter-broker pathway based on the tag value defined in the map. If the tag value matches the VLAN attribute configured in the import component map, the traffic is sent to the inline packet broker on the opposite side. The traffic is then routed through the inline tools or inline tool groups based on the sequence defined in the import component map. After inspection, the traffic is sent back to the inter-broker pathway with the same tag value. Finally, the traffic is intercepted by the export component map and is guided to the respective exit inline network port.
- inter-broker pathway-down—the traffic is handled based on the failover action selected for the inline map configured, as follows:
  - If the failover is set to 'bypass', the traffic is passed directly between the respective inline network ports.
  - If the failover is set to 'original-map', the traffic is passed through the path that is defined by the respective original map.

  NOTE: Traffic can be moved from 'bypass' to 'original-map' and vice-versa, when the inter-broker pathway is in 'down' state.

The failover-action set for an inline tool or an inline tool group that is configured on Node 2 will affect the inter-broker pathway as follows:

- If the failover-action for the inline tools on Node 2 is set to 'network-bypass', all traffic received on the Node 2 will be by-passed and referred back to Node 1.
- If the failover-action is set to 'network-drop', all traffic received on Node 2 of the inter-broker pathway will be dropped.
- If the failover-action is set to 'network-port-forced-down', all ports on Node 2 of the inter-broker pathway will be brought down.

Resilient Inline Arrangement—Rules and Notes

Keep in mind the following rules and notes when working with Resilient Inline Arrangement:

- Ensure that the names on both GigaVUE devices are identical, that is, the inline networks, inline tools, inline tool groups, out-of-band tools, and out-of-band tool GigaStreams must all have the same alias names on both the devices.
- If you choose to use the inline network bundle, the alias of the inline network bundle on both the devices must be identical. However, the inline networks that are grouped into the bundle can have different aliases.
- In GEN2 GigaSMART card, a maximum of 14 VLANs will be supported for a single inline-network per GS Group. In the case of multiple inline-network ports (number of inline-network ports x number of VLANs), the number should not exceed 14 per GS Group.
- In GEN3 GigaSMART card, a maximum of 16 VLANs will be supported for a single inline-network per GS Group. In the case of multiple inline-network ports (number of inline-network ports x number of VLANs), the number should not exceed 16 per GS Group.

Deploy Resilient Inline Arrangement

Following are the prerequisites that you must complete before you configure Resilient inline arrangement:

- Configure the required inline networks. Refer to Configure Inline Network Ports and Inline Network.
- Configure the required inline network LAG. Refer to Configure Inline Network Link Aggregation Group (LAG).
- Configure the required inline tools. Refer to Configure Inline Tool Ports and Inline Tools.
- Configure the required inline tool group. Refer to Configure Inline Tool Group.

Complete the following tasks to successfully deploy resilient inline arrangement:

Create Inter-broker Pathway

To create a new inter-broker pathway:

1. On the left navigation pane, go to **Physical > Orchestrated Flows > 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 create the inter-broker pathway.
3. Click the '**+**' icon next to the **IB Pathway** option to create a new inter-broker pathway.
4. In the **Properties** pane, in the **Alias** and **Description** fields, enter a name and description for the inter-broker pathway.
5. From the **Ports** drop-down lists, select the required tool ports to attach to the inter-broker pathway.

   NOTE: If the required tool ports are not available, you can choose to administratively enable the tool ports. Click **Port Editor**, and in the **Quick Port Editor** page, scroll down to the tool ports that you wish to configure. Select **Enable**, and then click **OK**.
6. In the **Minimum Ports Up** field, enter the minimum number of tool ports that must be operationally up so that the status of the inter-broker pathway will be up.
7. From the **Traffic Path** drop-down list, select one of the following options:
   - **Bypass**—Traffic bypasses the inter-broker pathway and is redirected to the next inline network port.
   - **Monitoring**—Traffic is forwarded to the sequence of inline tools in the monitoring mode.
   - **To Inline Tool**—Traffic is forwarded to the sequence of inline tools you have configured.
8. Click **OK** to save the configurations.

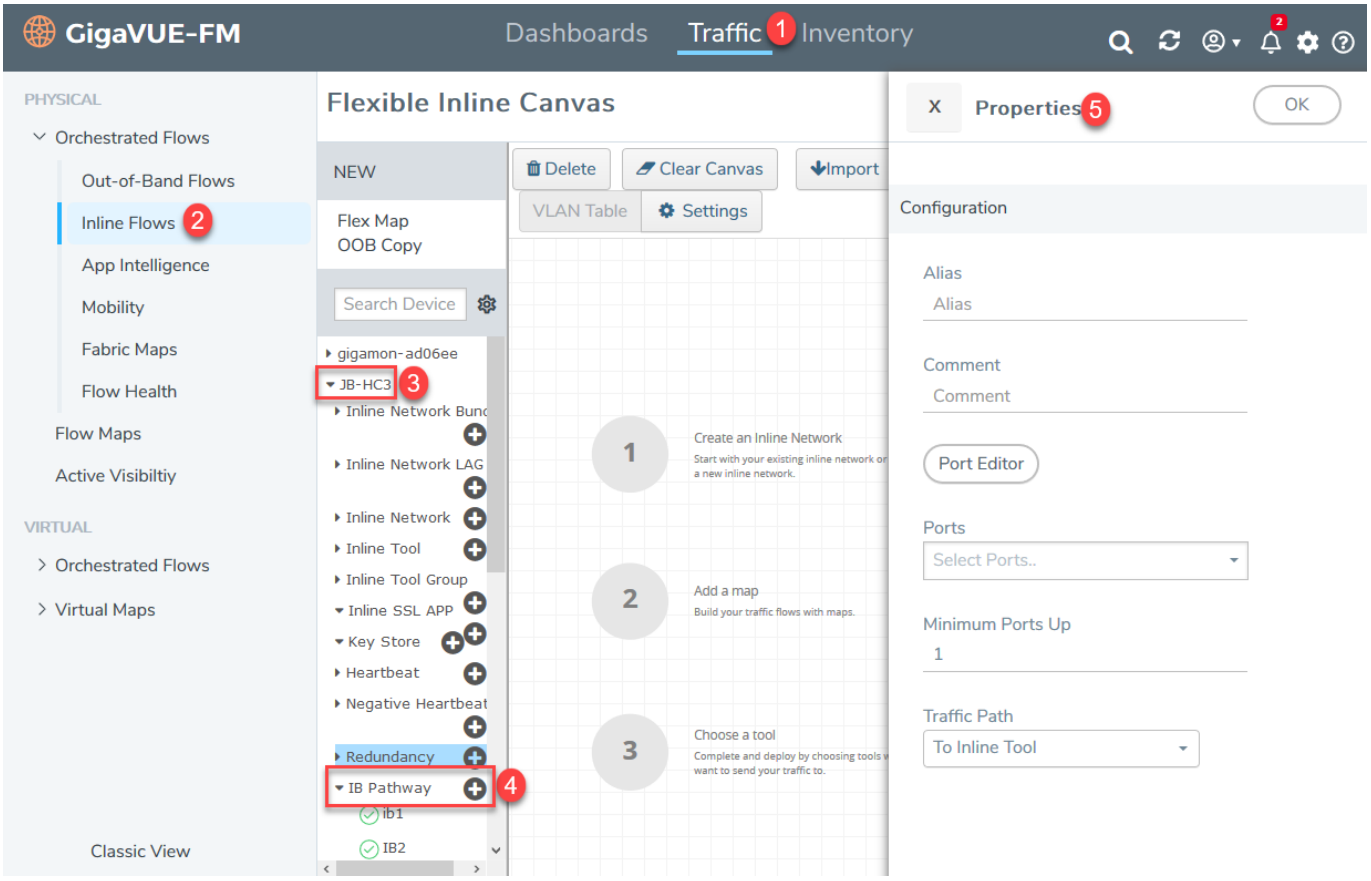

Configure Resilient Inline Arrangement

To configure a resilient inline arrangement:

1. Drag and drop the required inline network or inline network LAG into the flexible inline canvas, and then click **Settings**.
2. In the **Settings** pane, select the **Enable** check box next to **Show Single Tag Options** to configure resilient inline arrangement with a single VLAN tag.

   > **NOTE:** Enable **Show Single Tag Options** only when your inline tools does not support Q-in-Q VLAN tags.

3. Select the **Enable** check box next to **Show Resilient Inline Menu**.
4. Select the required **Node 1**, **Node 2**, **IB Pathway1**, and **IB Pathway2** for the resilient inline arrangement.
5. From the **Hashing Source** drop-down list, select one of the following options:

- **Side A**—Hashing is done based on either the source IP address or the source port from side A. On side B, hashing is done based on either the destination IP address or the destination port.
- **Side B**—Hashing is done based on either the source IP address or the source port from side B. On side A, hashing is done based on either the destination IP address or the destination port.

6. From the **Hashing Type** drop-down list, select one of the following options:
   - **L3 (IP Based)**—Hashing is done based on the IP address.
   - **L4 (Port Based)**—Hashing is done based on the transport layer port number.
7. From the **Hashing LSB Node** drop-down list, select one of the following options:
   - **Node1 as 0**—All traffic from IP addresses ending with 0 will be hashed to node 2.
   - **Node2 as 0**—All traffic from IP addresses ending with 0 will be hashed to node 1.

   NOTE: This field is available only if you select the **L3 (IP Based)** option in the **Hashing Type** field.
8. From the **Hashing Port** drop-down list, select one of the following options:
   - **Node1 as odd**—All traffic with odd port numbers will be hashed to node 2 and traffic with even port numbers will be hashed to node 1.
   - **Node2 as odd**—All traffic with odd port numbers will be hashed to node 1 and traffic with even port numbers will be hashed to node 2.

   NOTE: This field is available only if you select the **L4 (Port Based)** option in the **Hashing Type** field.
9. Click **OK** to save the settings.
10. Drag and drop the flexible inline map into the canvas, and then click the map to open the **Properties** pane.
11. In the **Alias** and **Description** fields, enter the name and description of the inline map.
12. Select the **Enable** check box next to **Single Tag Mode** if you want to deploy resilient inline arrangement with single VLAN tag. Refer to Resilient Inline Arrangement With Single VLAN Tag.

    NOTE: You can choose to disable the **Single Tag Mode** for collector maps, if required.
13. Enter the **Tool Side VLAN Tag** for the inline network for which you are configuring the map.
14. 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.
15. From the **FlexInline Failover** drop-down list, select one of the following options:
    - **Bypass**—the traffic is passed directly between the respective inline network ports.

- **Original Map**—the traffic is passed through the path that is defined in this flexible inline map.

16. Add the required rules for the inline map, and then click **OK** to save the configuration.
17. Drag and drop the required inline tools or inline tool group into the canvas.
18. Drag and drop the **OOB Copy** into the canvas, if required.
19. From the **Destination Ports** drop-down list, select the required hybrid or tool ports.
20. From the **VLAN Tag** drop-down list, select one of the following options:
    - **None**—No VLAN tag is used and the traffic is routed to a different destination.
    - **Original**—Uses the original VLAN tag of the packet received from the inline network.
    - **As Inline**—Uses the same VLAN tag that was configured for the flexible inline map.

    NOTE: The **As Inline** is the only option that is available when you configure Resilient Inline Arrangement with single VLAN tag.
21. Click **Deploy**. The Deploy pop-up window appears.
22. In the Deploy pop-up window, select a traffic path and click **OK**.

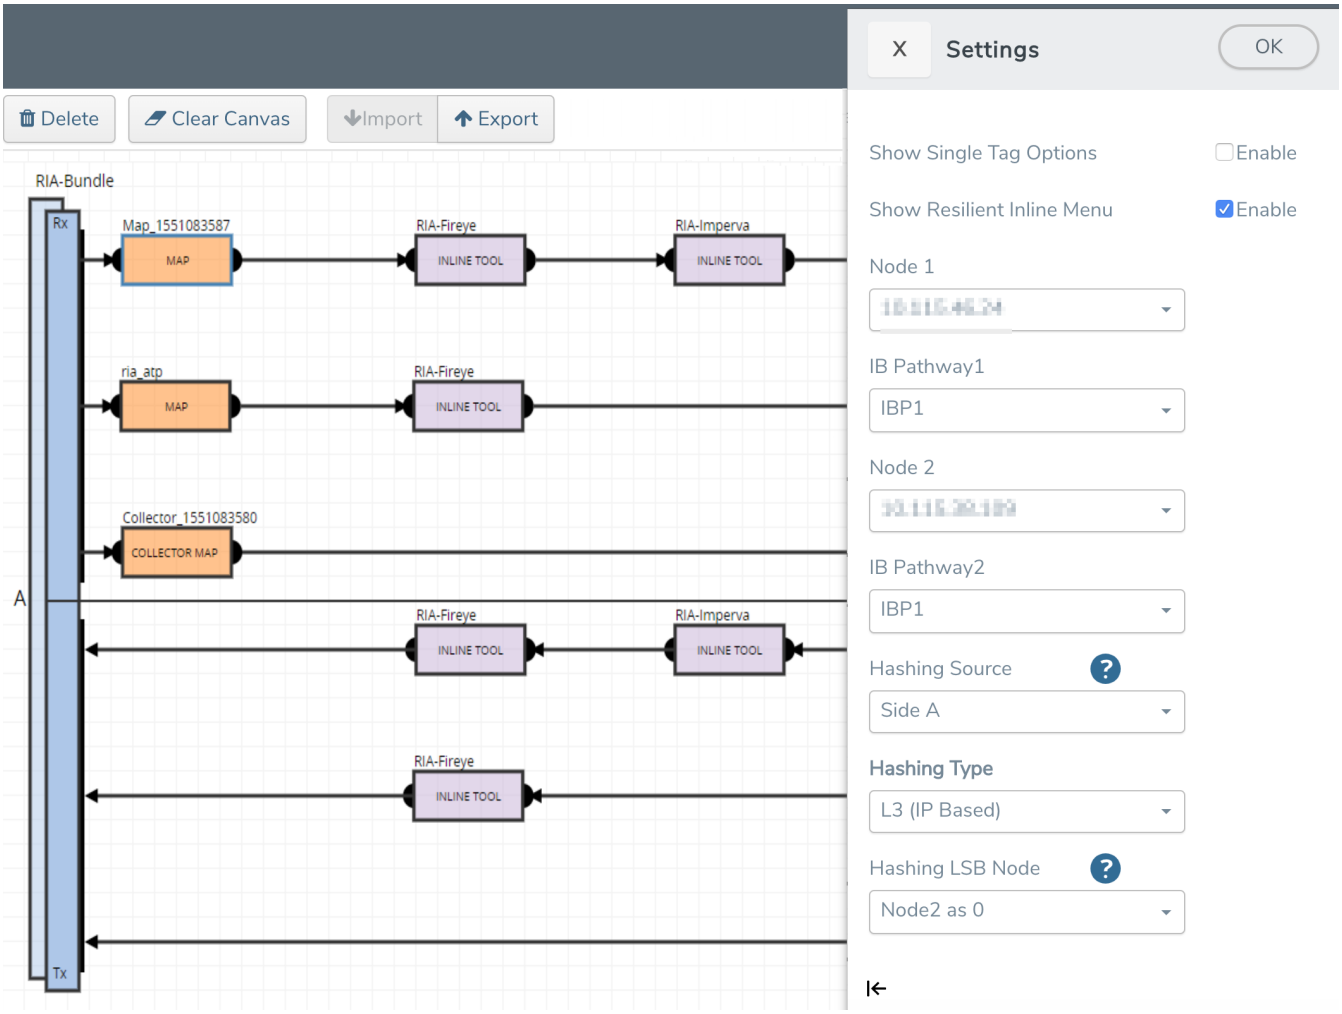

Configure Hardware Security Model (HSM)

Refer to the following sections that provide details about HSM Group, its limitations, and instructions on how to configure HSM Group in Flexible Inline Canvas:

- About HSM
- HSM Group - Limitations
- Supported Platforms
- Configure HSM Group

About HSM

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.

Some enterprises where security is paramount use Entrust nShield HSM to keep sensitive information such as private keys safe.

Starting in software version 6.4, the current Inline TLS/SSL is enhanced to include Thales-Luna Network HSM support in addition to the already supported Entrust nShield HSM solution.

.

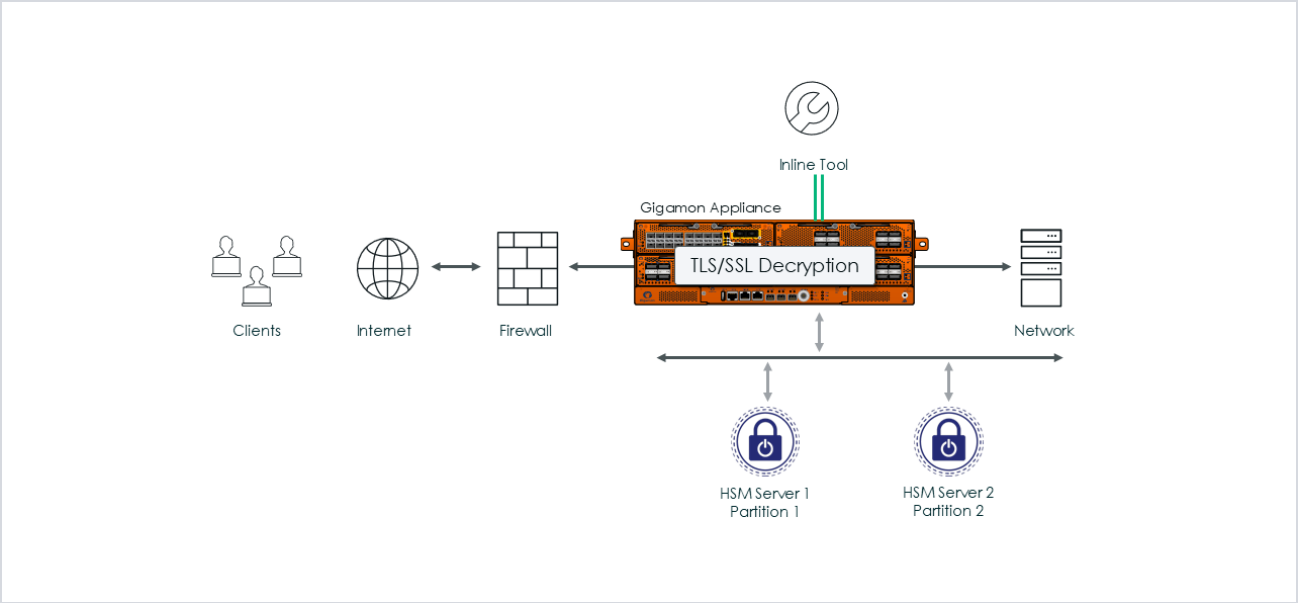

Figure 13 *ISSL with Thales Luna - Inbound support*

Thales -Luna Network HSM allows a single physical HSM to be divided into logical HSM partitions, each with independent data, access controls, and administrative policies. HSM partitions allow separate data storage and administration policies that would be maintained by multiple applications sharing one HSM without compromising other partitions residing on it. Each HSM partition has its own access control.

HSM Group - Limitations

Keep in mind the following limitations when configuring an HSM Group:

- Entrust nShield HSM and Thales-Luna Network HSM cannot be configured together in a HSM Group configuration.
- Thales-Luna Network HSM features like cluster, standby, and non HA are not supported.
- IPv6 support for Thales-Luna Network HSM server configuration will not be done along with IPv6 stack port support.
- When uploading RSA and ECDSA keys, validity check for protocol mismatch cannot be performed since the private keys are available on the HSM server.
- 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 application and deploy the solution.
  2. Additionally, the user needs to register the new DHCP IP with the HSM server for client use.

Supported Platforms

HSM Group is supported in the following platforms:

- GigaVUE-HC1 Gen3
- GigaVUE-HC3 Gen3
- SMT-HC1-S
- GigaVUE-HC1-Plus
- GigaVUE-HCT (Out-Band tools /Passive SSL decryption)

Configure HSM Group

To configure HSM Group:

1. On the left navigation pane, go to **Physical > Orchestrated Flows > Inline Flow**, and then click **Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the Flexible Inline Canvas that appears, select the required device for which you want to configure the HSM Group.
3. Click the '+' icon next to the **HSM Group** option to create a new HSM.

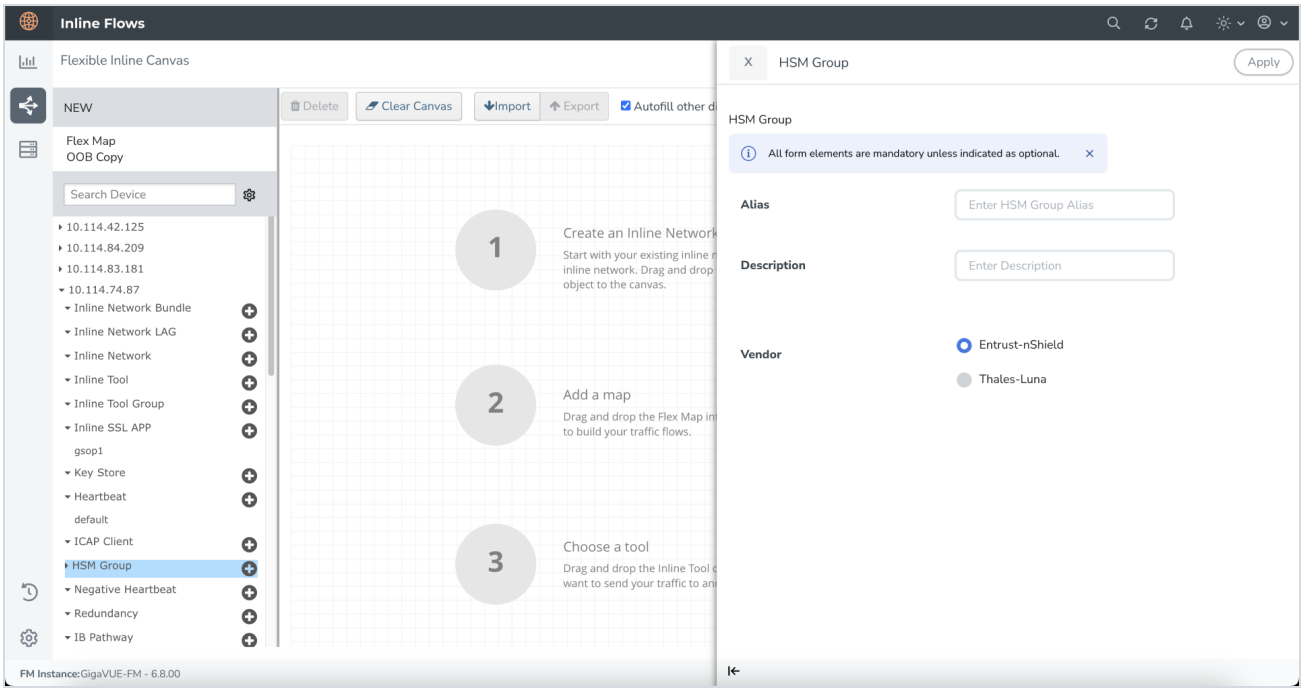

4. In the **HSM Group** properties pane that appears on the right, enter the name and description of the HSM Group alias in the **Alias** and **Description** fields.
5. Select the required vendor type from the options (**Entrust-nShield** or **Thales-Luna**) to create the respective HSM Group.

   > **NOTE:** You can create a maximum of 16 HSM units per device for Thales-Luna Network HSM.

6. Click **Apply** to save the configurations. All the individual HSM units that you create will be listed under the configured HSM Group. Refer below for detailed information on the configuration details for Entrust nShield and Thales-Luna.

   > **NOTE:** If the operational status of the HSM Group is 'Registration pending'. Please execute the below steps to register the HSM client (GigaSMART engine) in the HSM Server. 1) Register the GigaSMART engine IP address in the HSM Server 2) Assign the HSM partition to the GigaSMART engine IP address.

Configure Entrust nShield HSM:

1. Click the '+' icon next to the configured nShield HSM.

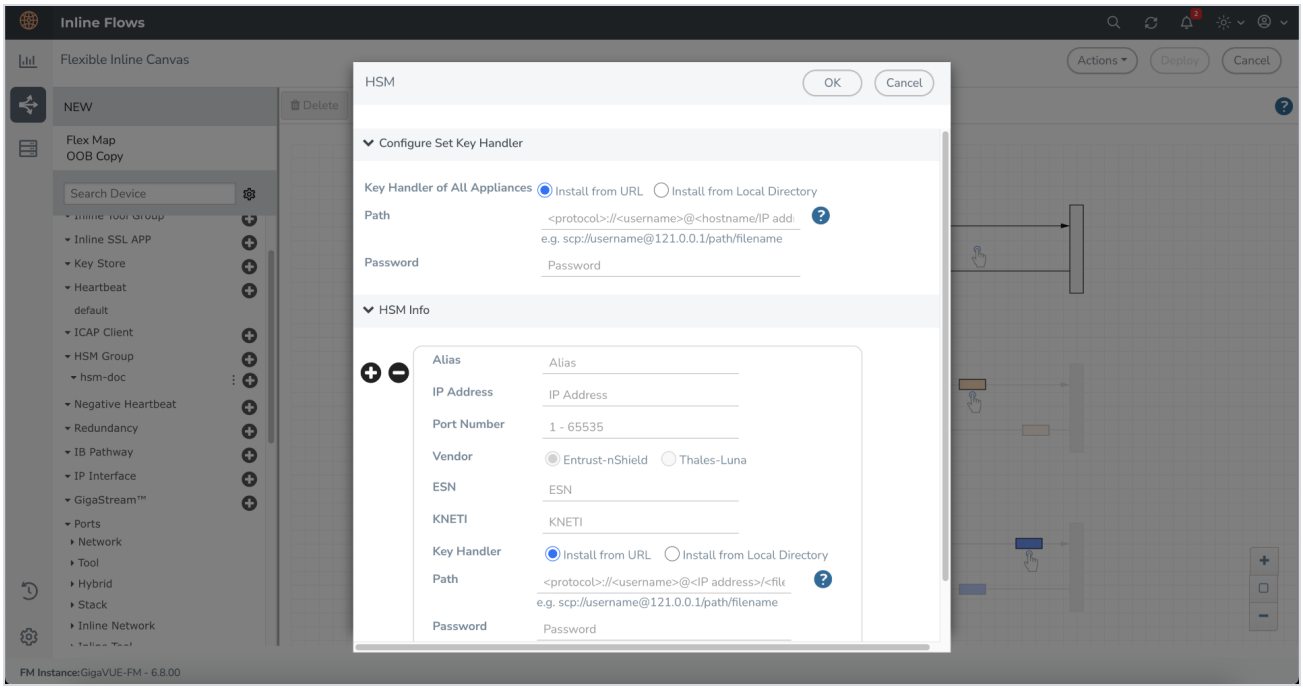

2. In the HSM pop-up pane, choose one of the following methods to install the key handler file:
   - **Install from URL**—Enter a valid directory path including the file name and enter the password to access the server.

   > **NOTE:** SCP, SFTP, HTTP, and FTP 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".

3. In the **Alias** field, enter a name for the HSM appliance.
4. Enter a valid **IP address** and **Port Number** details.

   > **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. By default, Entrust nShield is selected and Thales-Luna is disabled in the **Vendor** type when configuring Entrust nShield.
6. Enter the **ESN** and **KNETI** that you obtained from the HSM administrator.
7. Choose one of the following methods to select the required key handler file:
   - **Install from URL**—Enter a valid directory path including the file name and enter the password to access the server.

- **Install from Local Directory**—Browse and select the key handler file from your local directory.

8. Click **OK** to save the configuration.

To configure through GigaVUE-OS CLI refer to Entrust nShield HSM for SSL Decryption for Out-of-Band Tools.

Configure Thales-Luna HSM:

1. Click the '+' icon next to the configured Thales-Luna Network HSM.

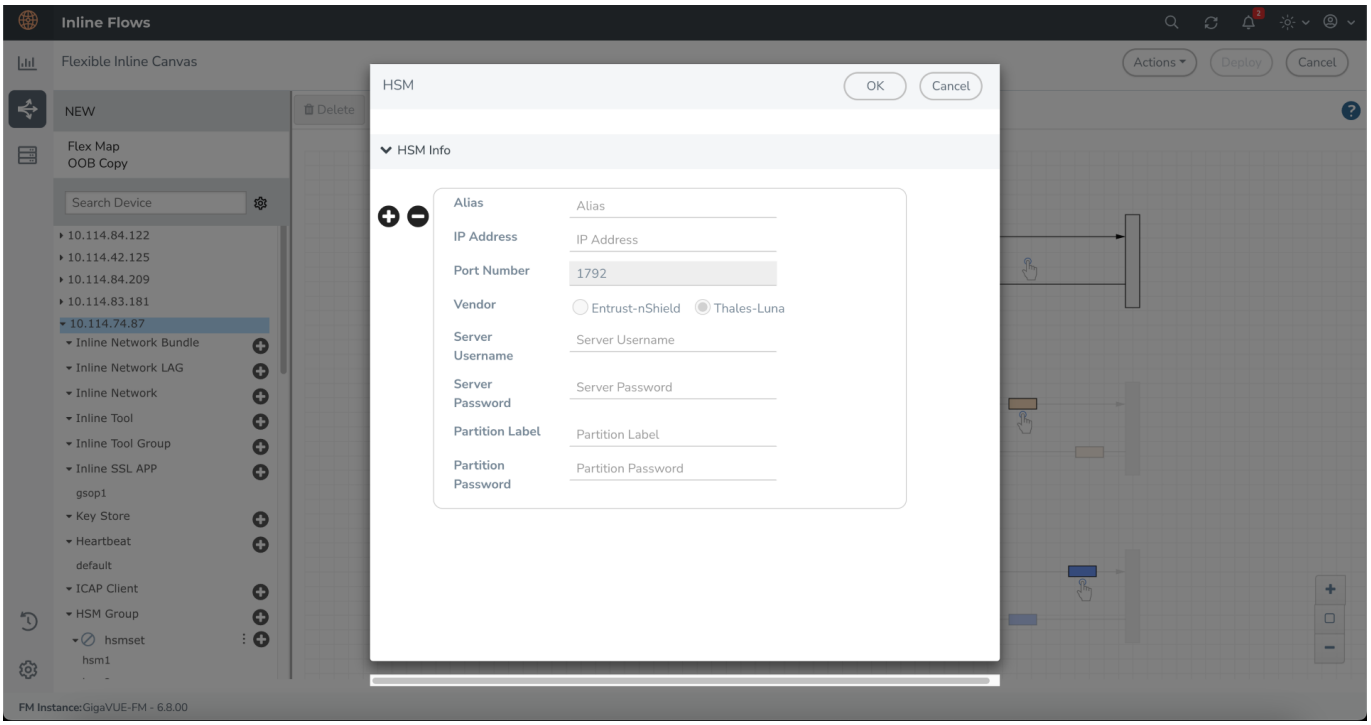

2. In the HSM pop-up that appears, enter a name for the HSM appliance in the **Alias** field.
3. Enter a valid **IP address** and **Port Number** details.

> **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.

4. By default, Thales-Luna is selected and Entrust nShield is disabled in the **Vendor** type when configuring Thales-Luna.
5. Enter the valid username and password in the **Server Username** and **Server Password** fields.
6. Enter the valid details in the **Partition Label** and **Partition Password** fields.

> **NOTE:** When adding multiple HSM appliances, make sure to keep the Partition Password same for all the partitions.

7. Click **OK** to save the configuration.
8. Drag the Inline Network object to the canvas and click **Deploy**.
9. Once you configure and deploy the HSM solution, you should register the configured GigaSMART client in the Luna server. Refer to Client Register - Luna Command Reference for more details.

> **NOTE:** Monitor the HSM Group status by utilizing the tooltip provided on the HSM Group in GigaVUE-FM or the **show apps hsm-groups** status command. Once the status turns to '**RegisterPending**', register your client in the Luna server within 10 minutes. In case of a reload, wait up to 90 seconds before you proceed with the registration.

To configure through GigaVUE-OS CLI refer to Entrust nShield HSM for SSL Decryption for iSSL.

Modifying a HSM Decryption Deployment

If an HSM Decryption deployment is modified follow the below steps:

1. Move the Inline Network traffic path to bypass mode.
2. Make the desired deployment change such as:
   a. A non-HSM based decryption to HSM Luna based decryption.
   b. A non-HSM based decryption to HSM Entrust nShield based decryption.
   c. An HSM Entrust nShield based decryption to HSM Thales-Luna Network based decryption
3. . Reboot the GigaSMART card.
4. Move the Inline Network out of bypass mode to 'To Inline Tool' mode.

Configure Flexible Inline TLS/SSL Decryption Solution

Refer to the following sections that provide details about the flexible inline decryption solution and instructions on how to configure it:

- Flexible Inline TLS/SSL Decryption Solution

- Benefits of Flexible Inline TLS/SSL Decryption Solution
- Flexible Inline TLS/SSL Decryption Solution—Rules and Notes
- Configure Flexible Inline TLS/SSL Decryption Solution

Flexible Inline TLS/SSL Decryption Solution

The flexible inline TLS/SSL decryption solution combines the flexible inline arrangements feature with the inline TLS/SSL decryption solution. It includes the GigaSMART-based packet processing, which is the inline TLS/SSL decryption functionality into the flexible inline arrangements framework. In the flexible inline TLS/SSL decryption solution, the outer maps guide the inline traffic to GigaSMART for preprocessing. The inner map guides the traffic processed by GigaSMART through a user-defined sequence of inline tools and inline tool groups.

Figure 14Flexible Inline Decryption Solution illustrates an example of how the inline TLS/SSL decryption functionality is incorporated in the flexible inline arrangements framework to form the flexible inline decryption solution.

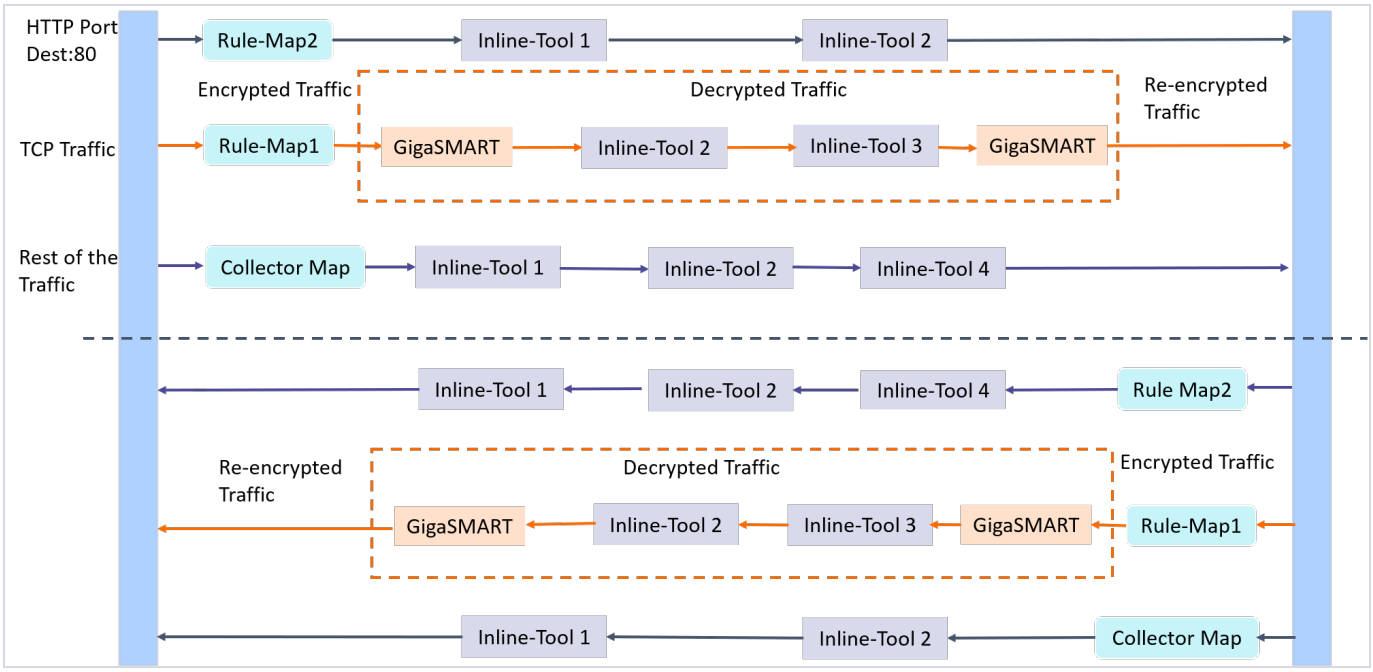

Figure 14 *Flexible Inline Decryption Solution*

In this example, the HTTP Port Destination:80 traffic is guided through the Flexible inline map, Rule Map2 to the sequence of inline tools, Inline Tool 1 and Inline Tool 2. The TCP traffic is encrypted and guided through Rule Map1 to GigaSMART, where it is decrypted and guided to Inline Tool 2 and Inline Tool 3. The decrypted traffic is guided back to GigaSMART,

and then it is re-encrypted and routed to the network. Here, the Rule Map 1 is the outer map, which guides the traffic to GigaSMART for preprocessing. The inner map guides the traffic processed by GigaSMART through a series of inline tools or inline tool groups.

The rest of the traffic is guided through the Collector Map to a series of inline tools, and then to the network.

Figure 15Flexible Inline Decryption Maps illustrates the different maps that guide the traffic in the flexible inline TLS/SSL decryption solution.

The outer maps are similar to the flexible inline maps but include virtual port alias along with inline tools and inline tool groups as the destination port. The same virtual port alias can be used in multiple outer maps.

Depending on the flexible inline TLS/SSL decryption solution that you create, there may be two types of inner maps:

- Inner Proxy Maps, which guides the traffic from GigaSMART to the inline tools or inline tool groups.
- Inner Non-proxy Maps or Network-end Maps, which guide the traffic that is bypassed from GigaSMART to the inline tools or inline tool groups.

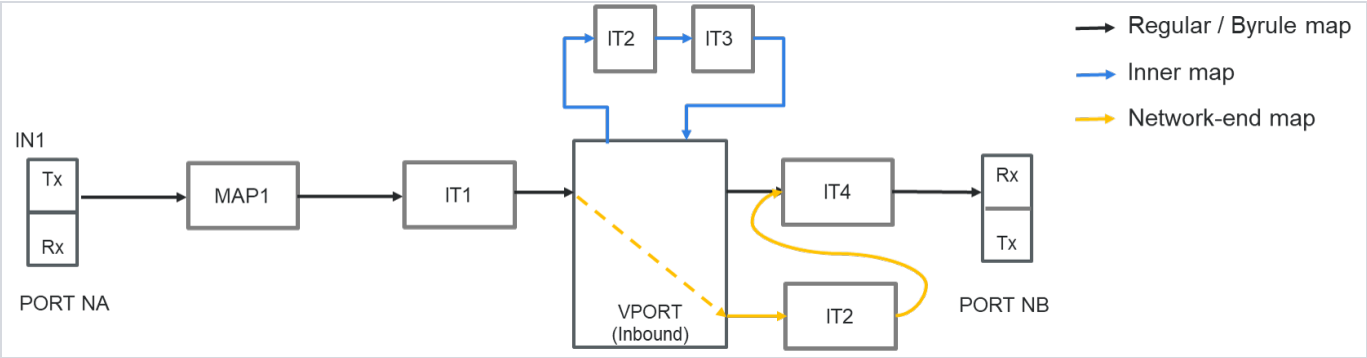

Figure 15 *Flexible Inline Decryption Maps*

Benefits of Flexible Inline TLS/SSL Decryption Solution

The flexible inline TLS/SSL decryption solution incorporates the inline TLS/SSL decryption with the flexible inline arrangement and offers the following benefits:

- Enables you to perform the inline decryption configuration, required map deployments, and flexible inline flow configurations, all in one canvas.
- Shares the same inline tool or inline tool group across multiple inline network links and across multiple inline maps. Refer to Figure 16Flexible Inline TLS/SSL APP—Deployed.

- Allows you to tap OOB copies of decrypted traffic from GigaSMART, either before or after the inspection of security tools. Refer to Figure 16Flexible Inline TLS/SSL APP—Deployed.
- Allows you to selectively decrypt and guide traffic to the attached inline tools or inline tool groups.

Flexible Inline TLS/SSL Decryption Solution—Rules and Notes

Keep in mind the following rules and notes when working with flexible inline decryption solution:

- When you want to migrate from the inline decryption to the flexible inline decryption solution,
  - ensure that you delete all the inline decryption virtual ports, GigaSMART, and maps configurations, and then reconfigure them using the flexible inline canvas. However, if there are OOB maps from the inline network ports, before you delete the OOB maps, ensure that the **Traffic Path** for the inline network is not set to 'Bypass'.
  - It is recommended to delete the OOB map from vPort before deleting other maps. If the OOB map from the vPort is not deleted while deleting all the inline TLS/SSL maps, then GigaVUE-FM throws an error on first time. You need to click **Delete All** again to delete all the maps.
  - Modifying the VLAN settings on an out-of-band map is not allowed if another out-of-band map has the same port as destination.
- When there is a multiple tool failover across the inner and outer maps; and if any of the tool comes up, the traffic in the inline network does not flow as expected.
- If you want to switchover from the flexible inline decryption solution to the inline decryption solution, you must delete the flexible inline SSL APP, and then reconfigure the ports, GigaSMART, and maps using the inline decryption workflow.
- When you configure the flexible inline decryption solution using GigaVUE-FM, the keychain password will be unlocked automatically when the device participating in the solution reboots.
- If an inline TLS/SSL profile is used across multiple map configurations with different inline network pairs, the tool set used across the maps is also the same. Consider the following example:
  - **Flexible Inline Map 1** with inline network pair 1 uses inline TLS/SSL Profile 1 with tools A and B.
  - **Flexible Inline Map 2** with inline network pair 2 also uses inline TLS/SSL Profile 1. This map also has tools A and B (filled in automatically). You cannot configure this as A or C or A, C.
  - The following combinations are not supported:
    - Flexible Inline TLS/SSL Decryption

- Inline Network LAG

  When you attempt to add an Inline SSL App to an Inline Network LAG Flexible Map you get the following error message: ***"An Inline SSL APP cannot be used when the traffic source is an inline network LAG"*** as shown in below figure.

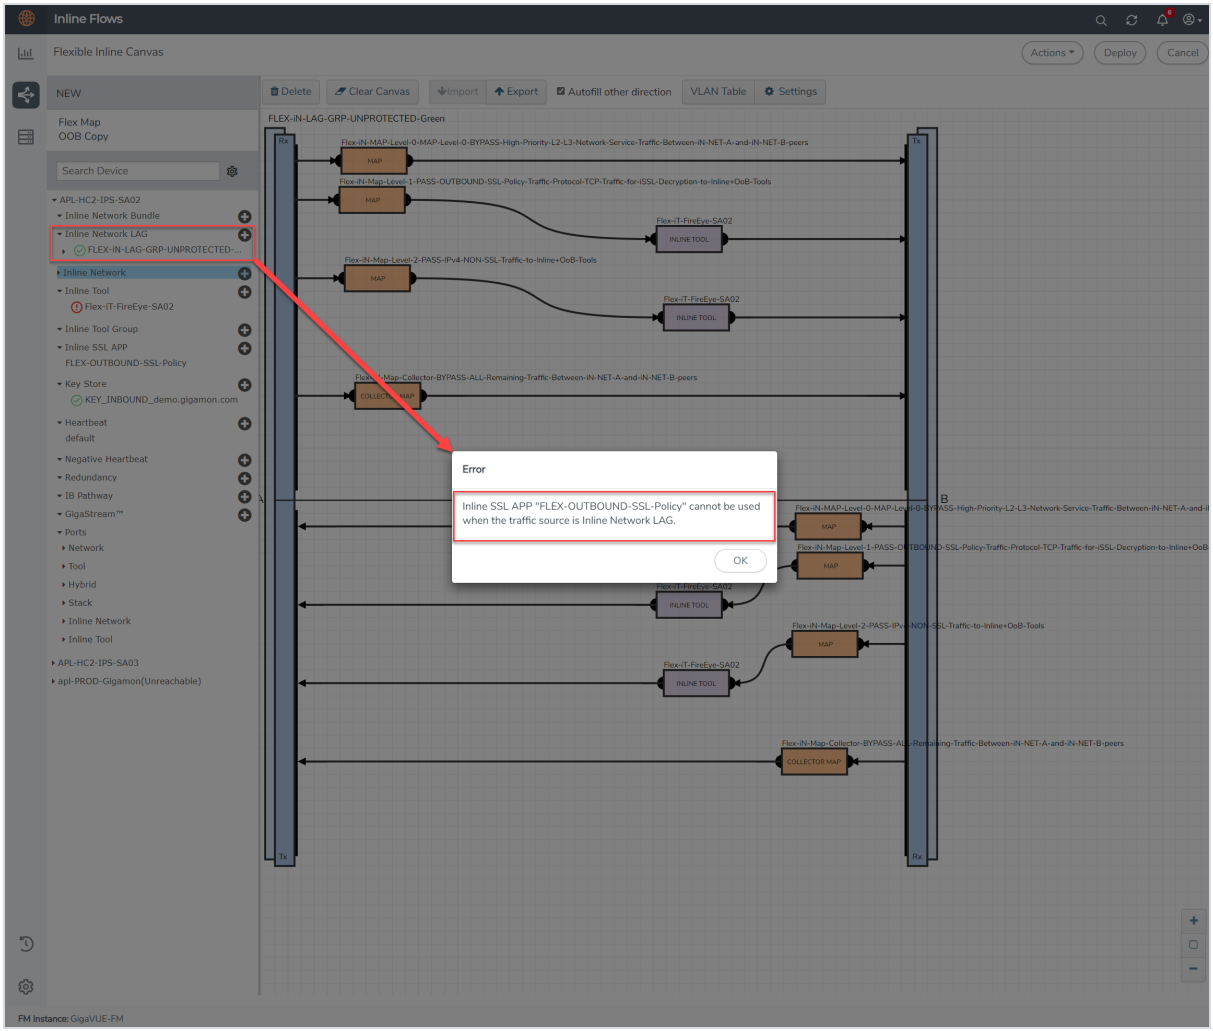

- For the Flexible Inline TLS/SSL Maps, tag of the outer map cannot be edited in the configuration canvas. To change the tag, follow these steps:
  a. delete the map
  b. deploy the solution
  c. re-add the map with the updated tag and deploy the solution again.
- Traffic is not decrypted when inline-network traffic path is in monitor mode.

- Setting the Flex Traffic Path of inner chain Inline-tools as “Drop” does not drop the Inline TLS/SSL traffic and continues to reach the Inline network egress.

Configure Flexible Inline TLS/SSL Decryption Solution

Following are the prerequisites that you must complete before you configure the flexible inline decryption solution:

- Configure the required inline networks or inline network bundle. Refer to Configure Inline Network Ports and Inline Network or Configure Inline Network Bundle.
- Configure the required inline tools. Refer to Configure Inline Tool Ports and Inline Tools.
- Configure the required inline tool group. Refer to Configure Inline Tool Group.
- Ensure that there are no inline decryption configurations such as inline decryption policy, inline decryption virtual port, GigaSMART group, GigaSMART operations, or inline decryption map configured on the device.

To configure a flexible inline decryption solution:

1. On the left navigation pane, go to **Physical > Orchestrated Flows > Inline Flows**, and then click **Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the Flexible Inline Canvas that appears, select the required device for which you want to configure the flexible inline decryption solution.
3. Click the ‘+’ icon next to the **Inline SSL APP** option to create a new flexible inline decryption solution.
4. In the **Inline SSL APP** page that appears, enter a name for the Inline SSL APP, and then complete the required fields. Refer Inline SSL App—Field References for details.
5. Click **OK** to save the configurations.
6. Drag and drop the required inline network or inline network bundle in to the flexible inline canvas.
7. Drag and drop the flexible inline map into the canvas.
8. In the **Properties** pane, in the **Alias** and **Description** fields, enter the name and description of the inline map.
9. Enter the **Tool Side VLAN Tag** for the inline network for which you are configuring the map.
10. 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.
11. Add the required rules for the inline map, and then click **OK** to save the configuration.
12. Drag and drop the Inline SSL APP into the canvas.

13. Drag and drop the required inline tools or inline tool group into the canvas.
14. Drag and drop the **OOB Copy** into the canvas, if required.
15. Click **Deploy**.

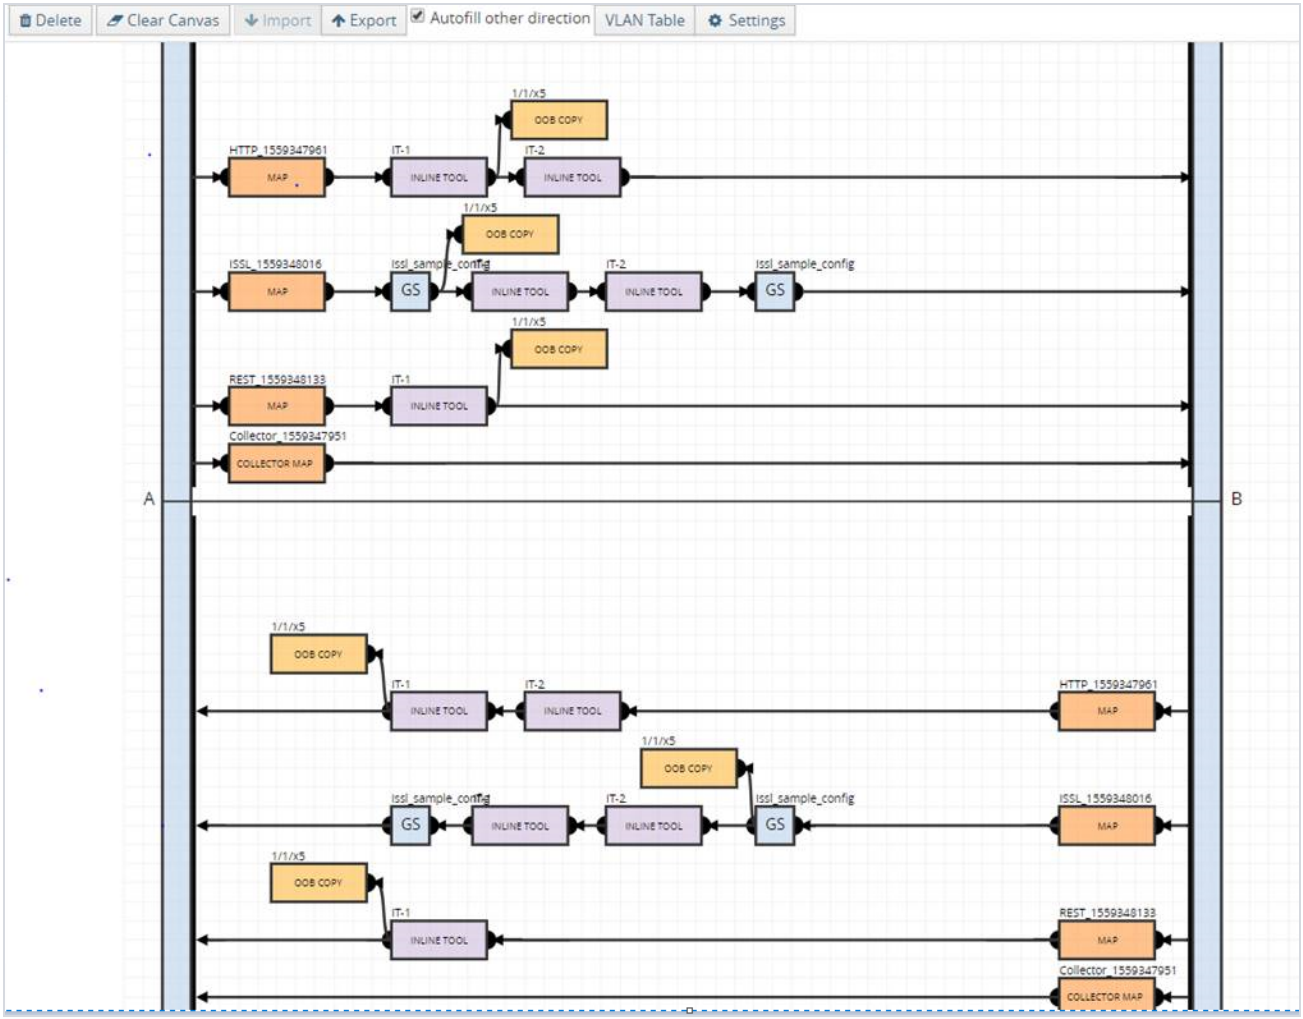

Figure 16 *Flexible Inline TLS/SSL APP—Deployed*

Inline SSL App—Field References

The following table lists and describes the attributes that define the flexible inline decryption solution.

Field	Description
Alias	Enter a unique name for the flexible inline SSL APP.
GS engines	Select the required GigaSMART engines.

Field	Description
TLS/SSL Monitor Mode	Select an TLS/SSL Monitor Mode from one of the following options: ▪ Enable—When the monitor mode is enabled, the TLS/SSL decryption or encryption is off. The monitor application collects information such as the TCP ports that are in use and VLAN information about the incoming traffic, and forwards the packets to the tool port or network port based on the non-TLS/SSL TCP bypass action. ▪ Disable—This is the default value. When the monitor mode is disabled, the TLS/SSL decryption or encryption is on. Use this mode during the deployment stage. ▪ Inline—Both monitor mode and TLS/SSL decryption or encryption is on. Use this mode to debug issues. Refer to Inline TLS/SSL Monitor Mode for details.
HSM Group	Select a HSM Group alias that you have configured from the drop-down list. Select **Disable** from the drop-down list to disable the HSM Group. Refer to Configure Hardware Security Model (HSM) for details. Notes: • Thales-Luna Network HSM configuration is supported in Inbound, Outbound, and Hybrid deployment types. • Entrust nShield HSM configuration is supported in Inbound, Outbound, and Hybrid deployment types.
Advanced Session Statistics	Enable this option to visualize advanced Inline SSL Session dashboards, such as Session Insights and Session Table, in the Fabric Health Analytics dashboard. The basic dashboards are available by default as you configure an Inline SSL session. Refer Default Dashboards and GigaSMART Inline TLS/SSL Dashboards to know more.
Keychain Password	The keychain password must be configured before installing certificates and keys. Refer to Configure Keychain Password for details. To add or reset the Keychain Password: **a.** Click **Keychain Password**, and then choose either **Add** or **Reset**. **b.** If you choose to reset the Keychain Password, enter a password that is 8 to 30 characters long and contains at least one numerical character, one uppercase character, one lowercase character, and one special character. **c.** Select the **Auto Login** check box to enable GigaVUE-FM to unlock the keystore when the device reboots. Refer to Support for unattended restart of TLS/SSL decryption in managed nodes for details. **d.** Click **OK** to save the Keychain Password.
Add new keys	To configure a certificate-key pair: **a.** Click **Add new keys** to open the Key page. **b.** Enter a name and description for the key. **c.** Select the required **Key Type** and **File Type**.

Field	Description
	d. You can choose to include a Passphrase for the key when you select PEM or PKCS12 as File type if required. **e.** When you choose Luna-HSM, enter the Key label for the key. **f.** Add the required **Private Key** and **Certificate**. **g.** Click **OK** to save the configuration.
Deployment Type	Select one of the following deployment types: ▪ Inbound—For inbound deployments, add a new **Server Key Mapping**. Enter the domain name or IP address of the server, and then select the required **Key Pair Alias**. Refer to TLS/SSL Session, Inbound Deployment for details. ▪ Outbound—For outbound deployments, add a primary and a secondary signing Certificate Authorities (CA). Refer to TLS/SSL Session, Outbound Deployment for details. ▪ Hybrid—For hybrid deployments, add a new **Server Key Mapping**, and a primary and a secondary signing CA. Refer to TLS/SSL Keys and Certificates and Generate and Add a Certificate to Key Store for details.
Configurations	
Default Action	Select one of the following options : ▪ Decrypt—Decrypt all the traffic that is guided into the Inline SSL APP. ▪ No Decrypt—Do not decrypt the traffic that is guided into the Inline SSL APP.
URL Cache Miss Action	Select one of the following options: ▪ Decrypt—Decrypt all the traffic that is guided into the Inline SSL APP. ▪ No Decrypt—Do not decrypt the traffic that is guided into the Inline SSL APP. ▪ Defer—Delay the decryption until the Defer Timeout seconds provided.
Tool Fail Action	The failover action taken in response to a failure of an inline tool. Select one of the following options: ▪ Bypass Tool—The traffic bypasses the failed inline tool. ▪ Drop Connection—The traffic is dropped.
Tool Bypass	Select the required options: ▪ Decrypted TLS/SSL Traffic—Bypasses the decrypted SSL traffic. ▪ No Decrypted TLS/SSL Traffic—Bypasses the non-decrypted SSL traffic. ▪ Non-TLS/SSL TCP Traffic—Bypasses the non-TLS/SSL, that is the TCP intercepted traffic.
High Availability	Select the check box to detect the link switchover by upstream device that is in active or standby mode. **NOTE:** Do not select this check box if the inline network links are in active state. Refer to High Availability Active Standby for details.

Field	Description
Network Group Multiple Entry	Select this check box to allow the traffic from different inline network to reenter GigaSMART. Refer to Inline Network Group Multiple Entry for details.
Tool Early Engage	Select this check box to allow the inline tools to change the MAC address or VLAN IDs. When a connection request is received from the client, GigaSMART establishes the connection with the inline tool first, before connecting with the server. This helps the inline tools to modify the MAC address or VLAN IDs when sending the traffic back to the server. Refer to Tool Early Engage and One-Arm Mode for additional information and limitations.
Tool Early Inspect	Select this check box to allow the inline tool to inspect the decrypted data first before connecting to the server. This will allow the inline tool to validate the data and ensure that only valid connections are sent to the server. Notes: • You can access **Tool Early Inspect** feature from the flex Inline SSL APP only. **Tool Early Inspect** cannot co-exist with features such as RIA, NAT/PAT mode, Tool Early Engage, One-Arm, and Decryption Port Mapping. • If **Tool Early Inspect** is enabled, you can configure the connections timeout value. Connection timeout represents the time by which the tool should respond after receiving the first decrypted data. If no response is received within the configured time interval, the connections will be reset. Refer to Tool Early Inspect for details.
StartTLS Port	Enter the required SSL/TLS ports. Refer to StartTLS and HTTP CONNECT for details.
Session Logging	
Session Logging	Select the Enable checkbox to log the Inline-TLS/SSL session related information to a remote server.
IP Version	Select IPV4 or IPV6 as the IP Version for the Session Logging server. You can select one session logging configuration per GigaSMART group.
Remote Syslog Server IP	Enter the IP address of the remote syslog server.
Associated IP Interface	In the Associated IP interface drop-down list, select the IP interface that you assigned to the GigaSMART group. You can create the IP interface by clicking the **Create IP Interface** button and the IP Interface window will open. Complete the fields to create the IP Interface: • In the **Alias** and **Description** fields, enter the name and description for the IP interface. • Select the **Port**.

Field	Description
	• Select **IPV4** or **IPV6** as the IP Version. • Enter an **IP Address**. For example, 192.168.1.20. • Enter an **IP Mask**. For example, 255.255.255.0. • Enter a **Gateway**. For example, 192.168.1.20. • Enter the Maximum Transmission Unit (MTU) for this port in the **MTU** field. For example, 1500. • Select the GigaSMART Group you created from the **GS Groups** field.
Remote Syslog Port Number	Enter the port number of the remote syslog server.
Log Level	In the Log Level drop-down list, select the severity log level of the events that you want to send to the inline TLS/SSL session logging server.
Traffic Path	
Single VLAN Tag	Enable the check box to deploy flexible inline TLS/SSL solution with a single VLAN tag. If an inline tool is involved in an inline TLS/SSL map, the inline tool can be supported across multiple maps with different single VLAN tags. **NOTE:** Deploying a flexible iSSL solution with SVT is optional, and you can choose to enable or disable the Single VLAN Tag option. If you choose to enable the Single VLAN Tag option in the iSSL solution, you should also enable the Single VLAN Tag configuration in the flex map deployed in that solution. **NOTE:** If you enable the Single VLAN tag option in the Flexible iSSL solution, you should also enable the Single VLAN Tag configuration in the inline-ssl app profile deployed in the solution Refer to Single VLAN Tagging (SVT) in iSSL for more details.
Tool Side VLAN Tag	Enter the required tool side VLAN tag for the inline network.
TPID	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.
Traffic Path	Select one of the following options: ▪ Drop—Traffic is dropped at the virtual port. ▪ Bypass—Traffic bypasses the virtual port. ▪ Monitoring—Traffic is fed to the virtual port and absorbed, while a copy of the traffic is sent to the next inline tool in the sequence. Traffic returned from side B of the network is also absorbed at the virtual port in the monitoring mode. **NOTE:** You can select the **Monitoring** option only if you have set the **SSL Monitor Mode** to either **Enable** or **Inline**. ▪ To Inline Tool—Traffic is forwarded to the inline tool. This is the default value.

Field	Description
Inline Failover Action	Select one of the following options: ▪ Virtual port bypass—All virtual ports configured as the source of any map that triggered this failover action, will be put in the bypass mode, that is all traffic will bypass the virtual port and will be guided to the inline tool or inline tool group. ▪ Virtual port drop—All virtual ports configured as the source of any map that triggered this failover action, will be put in the drop mode, that is all traffic will be dropped at the virtual port. ▪ Network bypass—All inline networks configured as the source of any map involving the inline tool or inline tool group that triggered this failover action, will be put in the bypass mode, that is, all traffic coming to side A will be directed to side B and vice versa. ▪ Network drop—All inline networks configured as the source of any map involving the inline tool or inline tool group that triggered this failover action, will be put in the drop mode, that is, all traffic coming to side A or side B will be dropped. ▪ Network port forced down—For all inline networks configured as the source of any map involving the inline tool or inline tool group that triggered this failover action, the inline network ports will be brought down.
Security Exceptions	You can choose to either decrypt or drop the traffic for the following certificates: ▪ Self-signed certificate ▪ Unknown CA certificate ▪ Invalid certificate ▪ Expired certificate You can also choose to configure the security exceptions for certificate revocation validation based on OCSP or CRL on inline decryption profile. Select one of the following options: ▪ Soft Fail—If you select this option, the client browser displays the secondary MitM certificate and the inline decryption session stats in GigaVUE-FM displays as Decrypt. ▪ Hard Fail—If you select this option, the client browser displays the certificate from DigiCert and the inline decryption session stats in GigaVUE-FM displays as Bypass: Unknown Revocation. Refer to Certificate Revocation List (CRL), Online Certificate Status Protocol (OCSP), CRL and OCSP, and Checking Certificate Revocation Status for details.
No-decrypt list/Decrypt list	Select the following check boxes: • No-decrypt list—Allows traffic from certain classes such as sites, domains, host-based IP address and IP subnets (decision based on LPM) to bypass decryption. • Decrypt list—Allows traffic from certain sites, domains, host-based IP address and IP subnets (decision based on LPM) to always be decrypted. Select from the below operations that can be performed on an uploaded list: • Append _ This would add to the uploaded list. • Replace- This would remove the previously added list and add a new list. This option is supported only on Generation 3 cards.

Field	Description
	• Clear- This would completely clear the list. • Download - This would download the list that has been uploaded. If you select Append/Replace, you can enter the list using any of the following options: · Copy and Paste · Install from URL · Install from Local Directory Refer to No-decrypt Listing Policy and Decrypt Listing Policy for details.
Policy Rules	Add the required policy rules for the inline decryption profile. Click **Add a Rule**. In the **Condition** field, Select one of the following options from the drop-down list: ▪ Category ▪ Domain ▪ IPv4 Destination ▪ IPv4 Source ▪ IPv6 Destination ▪ IPv6 Source ▪ L4 Port Destination ▪ L4 Port Source ▪ VLAN ▪ X509 Certificate Issuer Name Select one of the following options: ▪ Decrypt—Decrypt all the traffic that is guided into the Inline SSL APP. ▪ Decrypt—Do not decrypt the traffic that is guided into the Inline SSL APP.
Network Access	Network access configuration is used to get URL categorization updates. Refer to URL Categorization for details. To configure the network access for the GigaSMART engine ports: a. Select either DHCP or IP Address as the network access configuration mode. b. If you select IP Address as the mode, enter the IP Address, Netmask, Gateway, DNS, MTU, and VLAN. c. Select either Eth2 or Eth3 as the Interface. d. If you want to attach a Proxy profileto your Inlie SSL deployment select a **Proxy Server Profile** from the drop-down. To configure a new Proxy Server Profile, click on **Create new Proxy**. Refer Proxy Server Configuration to know more. **Notes:** • The Eth3 option is available only for GigaVUE-HC3 devices. • IP Address configuration mode details should be entered when you select Luna HSM configuration from the **HSM Group** drop down.

Field	Description
	• If your Proxy Server profile is associated with an Inline SSL application, choose **'None'** in the Proxy Server profile field on the Inline SSL configuration page to disconnect the proxy server profile prior to deleting the profile.
Decryption Port Mapping	The TCP destination port for decrypted traffic sent to inline tools can be configured as part of the inline decryption profile. Configure the required Priority 1 map, which is user configurable and Priority 2 map, which is the default out port. Refer to Inline TLS/SSL Decryption Port Map for details.
Trust Store	The trust store contains a trusted certificate authority (CA) for server validation. You can choose to either append or replace the trust store. Refer to Trust Store for details.
TCP Settings	Configure the required TCP settings as follows: • **TCP Inactive Timeout**— TCP Inactive session timeout in minutes. • **TCP Delayed ACK**—GigaSMART Inline TLS/SSL decryption ACKs every TCP packet by default. If TCP Delayed ACK is enabled, then GigaSMART decryption will wait for 100ms or ACK every third packet – whichever comes first. • **TCP SYN Retries**—number of retries made by the MitM to initiate a session with the destination server. If a SYN/ACK response isn't received from the destination server on initial TCP SYN, GigaSMART attempts for additional number of TCP SYN Retries as defined by the user. • **TCP TIMEWAIT Timeout**— Configure the 'TCP TIMEWAIT' timeout value from 0-300 seconds. The default value is 30 seconds. The TCP connection in the TIME_WAIT state gets deleted after the timeout period.
Split-Proxy Settings	
Split-Proxy	Select the check box to enable the split proxy settings for the inline decryption solution. The TLS connection between the server and client is divided into two independent connections, and the security parameters are kept separate.
Non-PFS Ciphers (Server)	Select the check box to enable the non-PFS ciphers settings for the inline decryption solution that has the split proxy settings enabled. This setting is to indirectly force the server to use protocols that are lower than TLS1.3 with non-PFS ciphers. This means that the ciphers with DHE/ECDHE key-exchange will not be used on the server side.
Miscellaneous (Global Settings)	

Field	Description
SSL/TLS Version	Select the minimum and maximum SSL/TLS version.
Connection Reset Action	Select one of the following options for the minimum SSL/TLS version: ▪ Drop—Closes all sessions that are below the minimum SSL/TLS version specified. This ensures that the network is safe from the weak TLS/SSL connections. This is the default option. ▪ No Decrypt—Bypasses all sessions that are below the minimum SSL/TLS version specified. Select one of the following options for the maximum SSL/TLS version: ▪ No Decrypt—Bypasses all sessions that are above the maximum SSL/TLS version specified. This is the default option. ▪ Drop—Closes all sessions that are above the maximum SSL/TLS version specified.
Caching persistence	Select this check box to allow the information to be saved on the GigaVUE node in the control card's persistent storage so that it can be retrieved in case of reboots. Refer to Cache Persistence for details.

Support for unattended restart of TLS/SSL decryption in managed nodes

The keychain is an encrypted database of certificates and private keys. On individual nodes, the keychain is stored in flash memory until reboot. The user needs to enter a keychain password to access the keychain. The keychain password is cached in the RAM of the control plane processor to allow decryption of the keychain file, but the keychain password is not cached across reboots. The SSL processing is not possible without the keychain password.

The keychain password is stored in GigaVUE-FM to automatically unlock the keychain during reboots and processing the TLS/SSL decryption without any intervention. The keychain password is stored in an encrypted database for key protection and risk management. Enable the **Auto Login** option when you set up or reset the keychain password to automatically unlock the keychain during reboots.

Single VLAN Tagging (SVT) in iSSL

You can configure to deploy flexible iSSL solution with a single VLAN tag. With shared mode enabled, through this support, the inline tool that is shared with a single VLAN tag in a packet will be supported in the flexible inline TLS/SSL arrangements across multiple maps with different single VLAN tags. You can configure the VLAN tags when you create the flexible inline maps.

NOTE: If you enable the Single VLAN tag option in Flexible iSSL solution, you should also enable the Single VLAN Tag configuration in the inline-ssl app profile deployed in the solution

NOTE: The PCP and CFI fields in the VLAN header cannot be retained when SVT is enabled in Flexible Inline maps and it will always be zero when sent out of the inline network ports.

You can configure a Flexible Inline SSL and RIA iSSL solution with Single VLAN Tagging (SVT).

The following table explains the compatibility matrix between single VLAN tag enabled and disabled maps. Symbol (√) denotes the engine ports that are supported, and symbol (X) denotes the engine ports that are not supported.

Maps	SVT enabled iSSL map	
	same gs_engine	different gs_engine in different maps
Flex	√	√
Flex + SVT	√	√
Flex + iSSL	√	√

Rules and Notes

Keep in mind the following rules and notes when deploying a flexible iSSL solution with single VLAN tag:

1. An untagged iSSL session map can also use a SVT enabled GigaSMART application.
2. An untagged or non-SVT solution should be of the least priority.
3. With SVT enabled in the TLS/SSL app, VLAN based policy rule is not applicable hence, the option will not be shown once SVT is enabled in TLS/SSL app.

Supported Platforms

The following devices support single VLAN tag feature in flexible inline TLS/SSL solution in both Gen2 and Gen3:

- GigaVUE-HC1-Plus
- GigaVUE-HC1
- GigaVUE-HC3V1

- GigaVUE-HC3V2

Limitations

- In all H-series devices, the same GSApp profile can be part of both SVT enabled and non-SVT iSSL solution, but if SVT is disabled in the gs_app profile then it will not be part of the SVT enabled solution.
- OOB-copy from vport with tag as original will not be supported for both proxy and non-proxy map because even if multiple VLAN tags are configured, only one proxy map will be created.
- For inline-tools that are part of a proxy(inner) map, oob-copy from inline-tool with tag as original will not supported.
- Flexible Inline single vlan tag configuration in GigaVUE-HC1-Plus fails to send traffic with a proper VLAN tag with monitoring mode configured in either the inline-network, inline-tool, or vport.
- In **GEN2** GigaSMART card, a maximum of 28 VLANs will be supported for a single inline-network per GS Group. In the case of multiple inline-network ports (number of inline-network ports x number of VLANs), the number should not exceed 28 per GS Group.
- In **GEN3** GigaSMART card, a maximum of 32 VLANs will be supported for a single inline-network per GS Group. In the case of multiple inline-network ports (number of inline-network ports x number of VLANs), the number should not exceed 32 per GS Group.
- With SVT enabled in the TLS/SSL app, VLAN based policy rule is not applicable hence, the option will not be shown once SVT is enabled in TLS/SSL app.

NOTE: If you edit the Single VLAN Tag in the inline SSL app present in a solution, you should explicitly edit the **Single Tag Mode** on flex maps manually in the same or other solutions that use the same SSL app.

Configure Internet Content Adaptation Protocol (ICAP)

Refer to the following sections that provide details about ICAP, its limitations, and instructions on how to configure the ICAP Client:

- About ICAP
- Supported Platforms
- Configure ICAP Client
- ICAP - Rules, Notes, and Limitations

About ICAP

The ICAP protocol serves as a communication interface for security tools like Data Loss Prevention (DLP) systems. Until software version 6.3, you cannot integrate inline SSL deployments with DLP-ICAP server as an inline tool. DLP tools are deployed outside of Gigamon® decryption zone as DLP tools need clear text traffic to inspect HTTP request/response.

Starting in software version 6.4, the ICAP Client app enables integration with the DLP ICAP server by functioning as an inline tool within the GigaSMART engine. Decrypted traffic from inline TLS/SSL is sent to the ICAP client and then forwarded to the ICAP server for inspection, providing enhanced security and visibility. If the ICAP application receives the decrypted traffic from a source other than Gigamon, it will function even without the iSSL deployment. ICAP app is part of TLS/SSL license SKUs.

The following figures illustrates the deployments in Inbound and Outbound with ICAP client support.

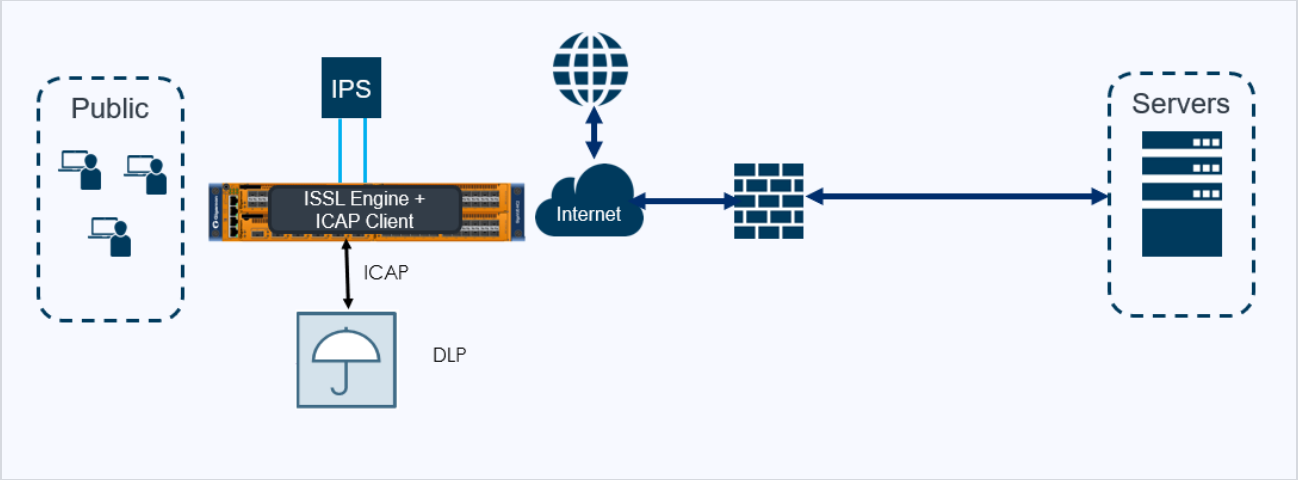

Figure 17 *Inline SSL Solution with ICAP Client support – Outbound Deployment*

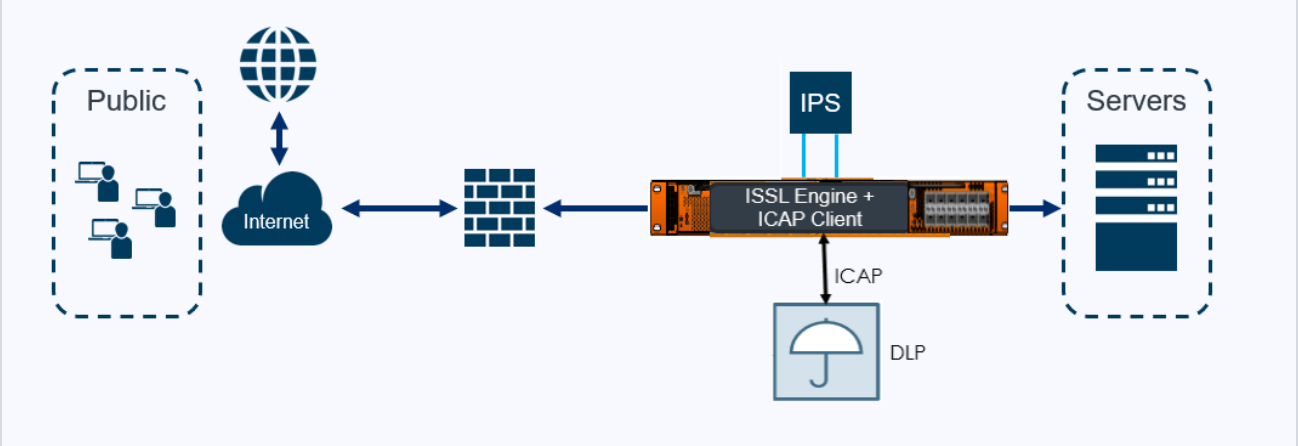

Figure 18 *Inline SSL Solution with ICAP Client support – Inbound Deployment*

Supported Platforms

ICAP Client app is supported in the following platforms and cards:

Platform	Card
GigaVUE-HC1 Gen3	SMT-HC1-S
GigaVUE-HC1P Gen3	SMT-HC1-S
GigaVUE-HC3 Gen3	SMT-HC3-C08

Configure ICAP Client

Following are the prerequisites that you must complete before you configure the ICAP Client:

- Configure the required inline networks. Refer to Configure Inline Network Ports and Inline Network.
- Configure the required IP Interface. Refer to Configure IP Interface.

  NOTE: For ICAP, it is not necessary to add GS Groups when configuring IP interface. It will be added automatically when the port is added to ICAP Client.

To configure the ICAP Client:

1. On the left navigation pane, go to **Physical > Orchestrated Flows > Inline Flows**, and then click **Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the Flexible Inline Canvas that appears, select the required device for which you want to configure the ICAP Client.
3. Click the '+' icon next to the **ICAP Client** option to create the new ICAP Client.

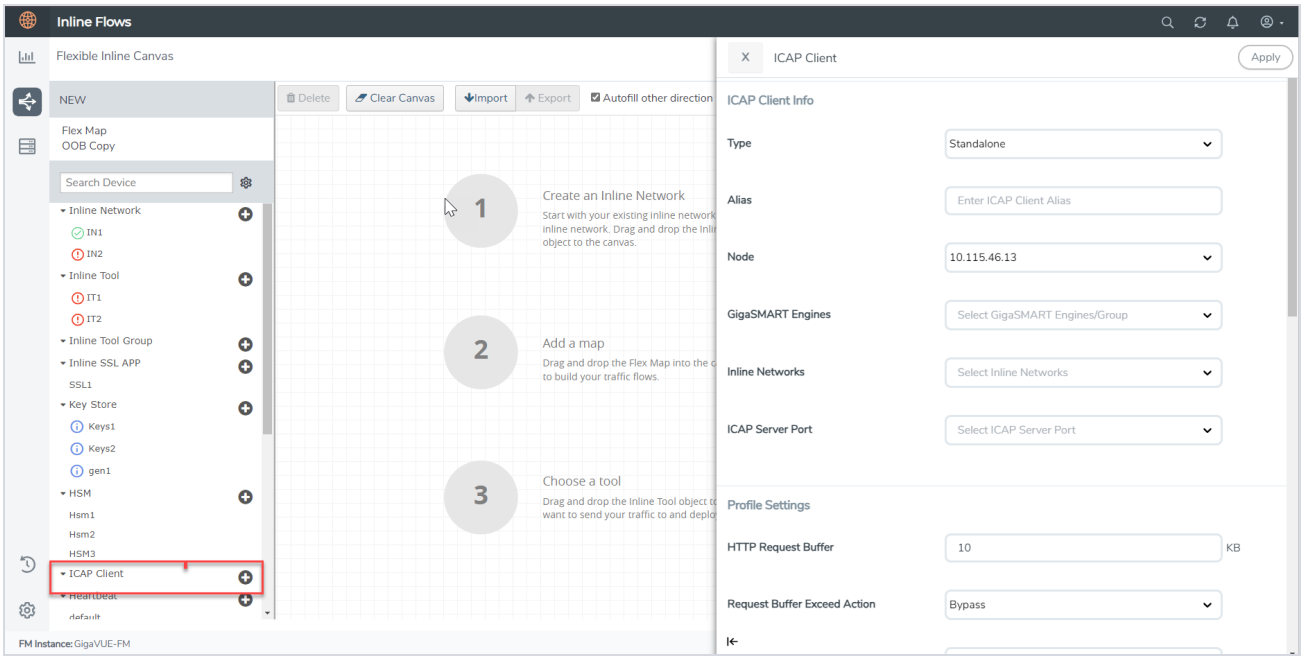

4. In the **ICAP Client** properties pane that appears on the right, complete the required fields in the ICAP Client Info, Profile Settings, and the Server sections. Refer to ICAP Client—Field References for more details.
5. Click **Apply** to save the configurations.
6. Drag and drop the required inline network and inline tool into the flexible inline canvas.
7. Drag and drop the Inline SSL APP into the canvas.
8. Click **Deploy**.

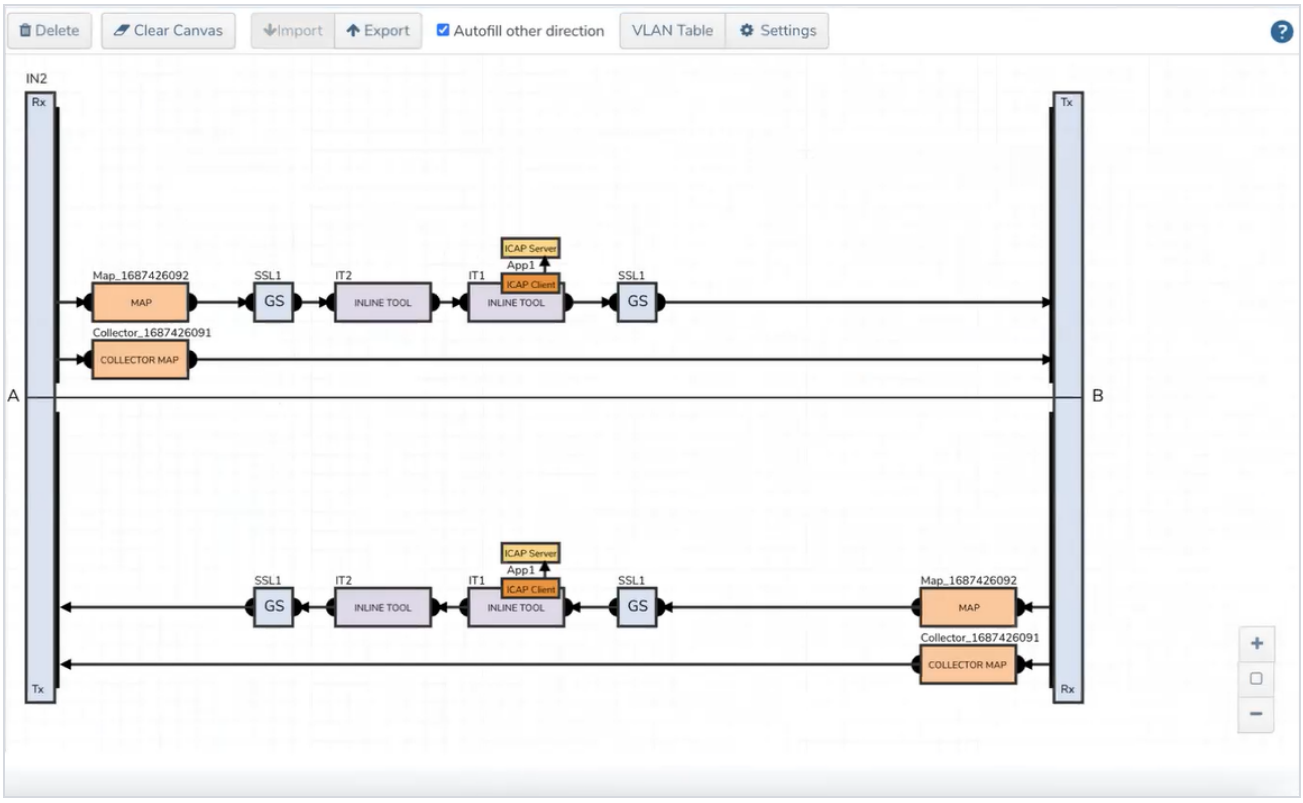

Figure 19 *ICAP Client APP—Deployed*

ICAP Client—Field References

The following table lists and describes the attributes that define the ICAP Client.

Field	Description
ICAP Client Info	
Type	Select the required type from the following: • If ICAP is deployed without iSSL and the traffic from devices other than Gigamon is sent to ICAP, select the **Standalone** option. • If ICAP and iSSL are integrated and deployed on the same node, select the **Same Node** option. • If ICAP and iSSL are integrated and deployed on different node, select the **Different Node** option.
Alias	Enter a unique name for the ICAP Client.
Node	Select the required node for which you want to configure the ICAP Client.
GigaSMART engines	Select the required GigaSMART engine.
Inline Networks	Select the required inline networks, which are the source for the ICAP app.

Field	Description
Source Ports	Select the required tool port of iSSL, which is connected to the inline network. (This option will not appear if you select standalone as type.)
ICAP Server Port	Select the required IP interface physical port.
Profile Settings	
HTTP Request Buffer	Enter the HTTP request buffer size.
Request Buffer Exceed Action	Select the buffer action on exceeding the size: • Drop - The traffic is dropped. • Bypass - The traffic bypasses on exceeding the buffer size.
Preview Size	Enter the required preview size.
REQMOD	Enable or disable the Request Modification Mode.
RESPMOD	Enable or disable the Response Modification Mode.
Response Timeout	Enter the response timeout in seconds.
Response Timeout Action	Select the response timeout action on exceeding the timeout: • Drop - The traffic is dropped • Bypass - Traffic bypasses on exceeding the response timeout limit.
Inactivity Timeout	Enter the session inactivity timeout value in seconds.
Port Range	Enter the ICAP client source port range for connecting to ICAP server.
Server	
Server Alias	Enter a unique name for the ICAP Server.
Description	Enter a description for the ICAP Server.
Port	Enter the L4 port number on which the ICAP server is listed.
Address	Enter the L3 IPV4/IPV6 address of the ICAP server.
Request Modification URL	Enter the Request Modification service URL.
Response Modification URL	Enter the Response Modification service URL.
Options URL	Enter the ICAP options URL (if necessary).

ICAP - Rules, Notes, and Limitations

Keep in mind the following rules, notes, and limitations when configuring the ICAP Client app:

- ICAP app gsop is not supported in the cluster environment.

- It is not possible to chain ICAP app gsop with other gsops.
- For a gsgroup, only one instance of ICAP gsop is allowed.
- **HTTP2** traffic ICAP inspection will not be supported in the current implementation for software version 6.4. As an alternative, you can enable HTTP2 downgrade in iSSL to enforce traffic in HTTP 1.1 version.
- If HTTP messages are within preview length and if any Trailer is present, they will not be sent to the ICAP server. This is a limitation as part of the RFC 3507 design itself.
- Any Trailer present in HTTP messages will be sent to the ICAP server as it is. No modifications will be made. There is no specific definition in ICAP RFC for HTTP trailers.
- HTTP 1.1 pipeline will not be supported. That is multiple HTTP requests that are sent one after another. Only one outstanding HTTP request/response is supported in our ICAP Client.
- The chunk extensions (a note that is added along with the data chunk during data transfer in http/1.1) in the http communication between client and server will not be sent to the ICAP server for inspection.
- ICAP app cannot be deployed as inline tool with One-Arm iSSL, L3 NAT iSSL, and RIA iSSL.
- In HC1P, the inline-network ports must be in the same slot.
- ICAP server traffic will be disrupted randomly on the HC1P device when an inline pair map is deployed with multiple GigaSMART engine ports. TCP SYN packet from inline network IN1 side A reaches the engine port1 whereas SYN/ACK packet from inline network IN2 side B reaches the other engine port2. Issues in traffic will be observed due to this asymmetric pattern of traffic.
- ICAP is not supported in GigaVUE-HCT chassis.
- Refer to the table below for information on the number of GS engine ports required for each deployment module:

Type	Number of GS Engine	Comments
Standalone	1	1 Gen3 for ICAP
Inline SSL: Same Node	2	1 Gen3 for ICAP and 1 Gen3 or Gen2 for iSSL **NOTE:** Inline SSL does not support HTTP2 Downgrade on Gen2. If HTTP2 Downgrade is required for iSSL, use 1 Gen3 for ICAP and 1 Gen3 for iSSL.
Inline SSL: Different Node	2	1 Gen3 for ICAP in the same node and 1 Gen3/Gen2 for ISSL in the other node

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, and GigaVUE-HC1P 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, or GigaVUE-HC1P
- 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.

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 modules on GigaVUE-HC1. The capacity will be 10Gb.
- two 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.

Between the two GigaVUE nodes, a 10Gb fiber signaling link is cabled using stack ports. Also, two inline tools are needed for the GRIP solution.

Refer to Figure 20Traffic Flows Through Node with Primary Role, Figure 21Traffic Flows Through Node with Secondary Role after Primary is Lost, and Figure 22Both Nodes Fail; No Traffic Monitoring.

Figure 20Traffic Flows Through Node with Primary Role shows traffic coming from a network (for example, the Internet) through an edge router at the top of the figure. Two GigaVUE nodes with an inline monitoring tool attached to each node are shown in the middle of the figure. Traffic to end devices on a private network are shown at the bottom of the figure.

The GigaVUE node on the left of the figure is acting in the primary role, while the GigaVUE node on the right is acting in the secondary role. The nodes are synchronized through a signaling link using a stack link between two stack ports.

As shown in Figure 20Traffic Flows Through Node with Primary Role, traffic only flows through the node with the primary role. On the primary node, the bypass protection switch relays are open. Traffic is directed to the inline tool attached to the primary node. The node with the secondary role watches the state of the signaling link. If the primary node is up, the link is up, and the secondary node takes no action. The bypass protection switch relays on the secondary node are in a closed state. In Figure 20Traffic Flows Through Node with Primary Role, the dotted lines depict the inactive traffic path.

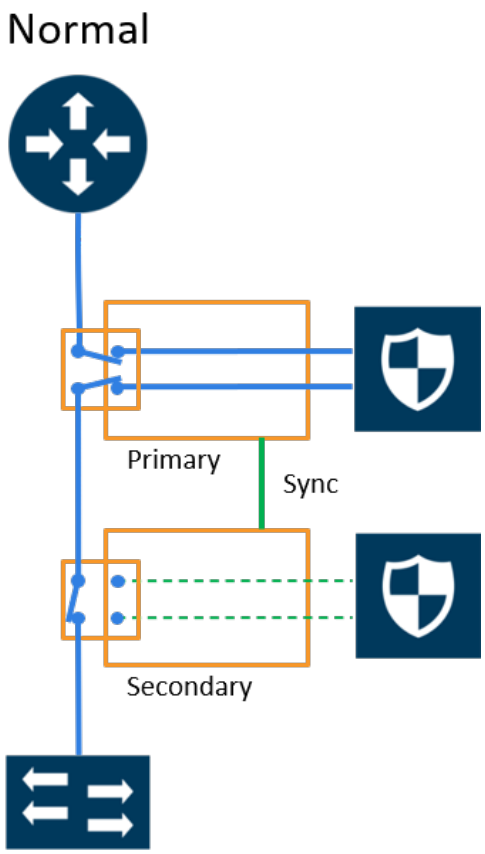

Figure 20 *Traffic Flows Through Node with Primary Role*

In Figure 21Traffic Flows Through Node with Secondary Role after Primary is Lost, power is lost to the GigaVUE node in the primary role. The bypass protection switch relays on the primary node close automatically when the node is down. The secondary node receives a signal through the signaling link that the primary node is down. The secondary node opens its bypass protection switch relays. Now traffic flows through the secondary node and traffic is directed to the inline tool attached to the secondary node.

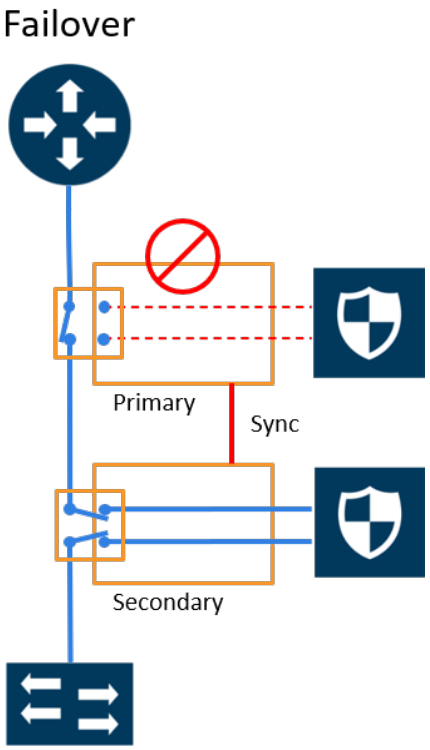

Figure 21 *Traffic Flows Through Node with Secondary Role after Primary is Lost*

In Figure 22Both Nodes Fail; No Traffic Monitoring, both nodes have lost power. The bypass protection switch relays are closed on both GigaVUE nodes. Traffic flows between the networks, but without going through the inline tools, which are both bypassed.

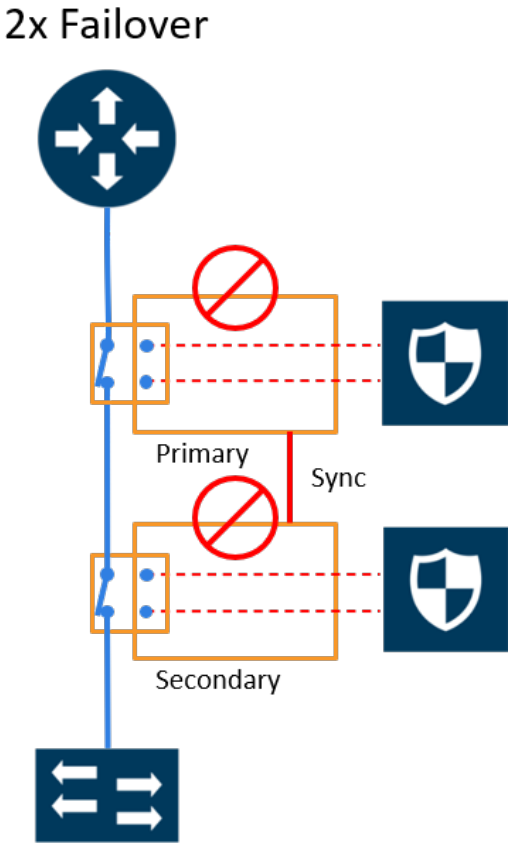

Figure 22 *Both Nodes Fail; No Traffic Monitoring*

How to Handle Recovery

In the scenario in Figure 21Traffic Flows Through Node with Secondary Role after Primary is Lost, 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 Figure 20Traffic Flows Through Node with Primary Role with the primary on the left and the secondary on the right:

- Connect the network shown at the top of Figure 20Traffic Flows Through Node with Primary Role 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 Figure 20Traffic Flows Through Node with Primary Role.
- 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, click on **> Nodes.**Click on the Left Node in which the configuration needs to be done.
2. From the left navigation pane, 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. On the left navigation pane, go to **Physical > Orchestrated Flows > 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 Figure 21Traffic Flows Through Node with Secondary Role after Primary is Lost, 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 13: Redundancy Control States

Table 13: 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**".

GRIP for Mixed Topologies

GRIP supports mixed topologies. The two nodes in the GRIP configuration do not both have to be the same GigaVUE® HC Series nodes but the speed of both nodes must be matching. For example, one node can be a GigaVUE-HC1 and the other node can be a GigaVUE-HC1-Plus. However the port speed on both nodes (10Gb) must match.

Upgrade Procedure

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.

Pros and Cons

Upgrade Secondary Node followed by Primary Node

Pros:

- Traffic will always be inspected, except for the delay during the primary to secondary node switch.
- Even when the primary node falls under some unreachable state or gets a continuous reboot during the upgrade, the traffic will still be inspected in the secondary.

Cons:

- Time-consuming as upgrades need to be done one by one.

Upgrade both nodes together

Pros:

- Time taken for upgrade will be minimal as both the nodes will get upgraded together.

Cons:

- During the upgrade, all the traffic will get bypassed, and no inspections will be done.
- If the primary node falls into any issuing state and only the secondary comes up, then traffic will still not get inspected in the tool and bypassed.

Rules and Notes

Keep in mind the following rules and notes when you work with the Gigamon Resiliency for Inline Protection feature:

1. Physical Bypass should not be enabled when the redundancy profile is configured in inline network ports.

2. The signaling port type should be a stack port, and only one port should be used.

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. 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.

5. Link Failure Propagation is not recommended when copper ports (TAP card ports) are involved in Inline Network ports, and only one path (a-to-b alone) is available. If the Inline Network has both paths, then it is a best practice to leave LFP enabled.

6. Gigamon Resiliency for Inline Protection (GRIP™) is not supported in GigaVUE-HCT devices.

7. GRIP is not supported in other GigaVUE TA Series devices due to the absence of BPS modules.

Refer to Flexible Inline Arrangements—Rules and Notes, which also apply to GRIP.

Limitations

- All the inline components should be located in the same box within the cluster..
- Although GRIP is supported in a cluster, there is a limitation when the suspended protection role is used on the standby node in the cluster. It is recommended to either switch the standby to the leader or apply the suspended role in the redundancy profile to the leader.

Troubleshoot

If any of the below issues occur then kindly follow the given steps to isolate 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.

How many members can we have in the GRIP configuration?

It is recommended to have two nodes to do the GRIP configuration where one node acts as the Primary and the other acts as the Secondary.

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.

Is GRIP Supported in Flexible Inline Maps?

Yes, GRIP is supported in both Classic and Flexible Inline Maps.

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.

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.

Example: Gigamon Resiliency for Inline Protection

Example is an Inline Bypass solution for GRIP using TAP-HC1-G10040 modules on the GigaVUE-HC1, with copper ports.

First, configure the GigaVUE-HC1 with the primary role, then configure the GigaVUE-HC1 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-HC1

Task	Description	UI Steps
1.	Configure ports on the TAP-HC0-G100C0 module as passive (in	**1.** From the left navigation pane, go to **System > Ports > Ports > All Ports**.

Task	Description	UI Steps
	passive mode, relays are closed). Also configure ports, port type (inline-network).	**2.** Select a port of the TAP-HC1-G10040 module. **3.** Click to open the Ports page. **4.** Select Passive for TapTX **5.** Click Save. **6.** Repeat steps 2 through 6 for each port on the TAP-HC1-G10040 module **7.** Configure Inline Network ports **a.** Select the port. **b.** Click Quick Port Editor. **c.** 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. **2.** Select Enable for Admin. **3.** Select Stack for Type. **4.** 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 Physical > Orchestrated Flows > Inline Flows > Configuration Canvas > Redundancy. **2.** Click '+' icon. **3.** Enter a name for the profile in the Alias field. For example, RP_001. **4.** Click in the Signaling Port field and select the stack port configured in Task 2. **5.** Select Primary for Protection Role. **6.** Click OK.
4.	Configure inline network. This step associates the redundancy profile to the inline network and also disables link fail propagation on the inline network.	Refer to the Configure Inline Network Ports and Inline Networksection for configuration details.
5.	Configure inline tool ports, port type (inline-tool), and administratively enable them.	**1.** Select Physical > Orchestrated Flows > Inline Flows > Configuration Canvas > Inline Tool. **2.** Select the first port (for example, 1/4/x1) to configure as an inline-tool port. **3.** Select Inline Network for Type and select Enable for Admin. **4.** Select the second port (for example, 1/4/x2) and repeat steps 3 through 5.

Task	Description	UI Steps
6.	Configure inline tool and failover action. then enable inline tool.	Refer to the Configure Inline Tool Ports and Inline Toolssection for configuration details.
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.	Refer to the Configure Inline Network Bundle section for configuration details.

Configure Secondary Role GigaVUE-HC1

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.** From the left navigation pane, go to **System > Ports > Ports> All Ports**. **2.** Select a port of the TAP-HC1-G10040 module. **3.** Click to open the Ports page. **4.** Select Passive for TapTX **5.** Click Save. **6.** Repeat steps 2 through 6 for each port on the TAP-HC1-G10040 module **7.** Configure Inline Network ports **a.** Select the port. **b.** Click Quick Port Editor. **c.** 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. **2.** Select Enable for Admin. **3.** Select Stack for Type. **4.** 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	**1.** Select Physical > Orchestrated Flows > Inline Flows > Configuration Canvas > Redundancy. **2.** Click '+' icon.

Task	Description	UI Steps
	protection role (secondary).	**3.** Enter a name for the profile in the Alias field. For example, RP_001. **4.** Click in the Signaling Port field and select the stack port configured in Task 2. **5.** Select Secondary for Protection Role. **6.** Click OK.
4	Configure inline network. This step associates the redundancy profile to the inline network and also disables link fail propagation on the inline network.	Refer to the Configure Inline Network Ports and Inline Networksection for configuration details.
5	Configure inline tool ports, port type (inline-tool), and administratively enable them.	**1.** Select Physical > Orchestrated Flows > Inline Flows > Configuration Canvas > Inline Tool. **2.** Select the first port (for example, 1/4/x1) to configure as an inline-tool port. **3.** Select Inline Network for Type and select Enable for Admin. **4.** 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.	Refer to the Configure Inline Tool Ports and Inline Tools section for configuration details.
7	Configure map passall, from inline network to inline tool.	Refer to the Configure Inline Network Bundlesection for configuration details.
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.	Refer to the Configure Inline Network Ports and Inline Networksection for configuration details.

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.

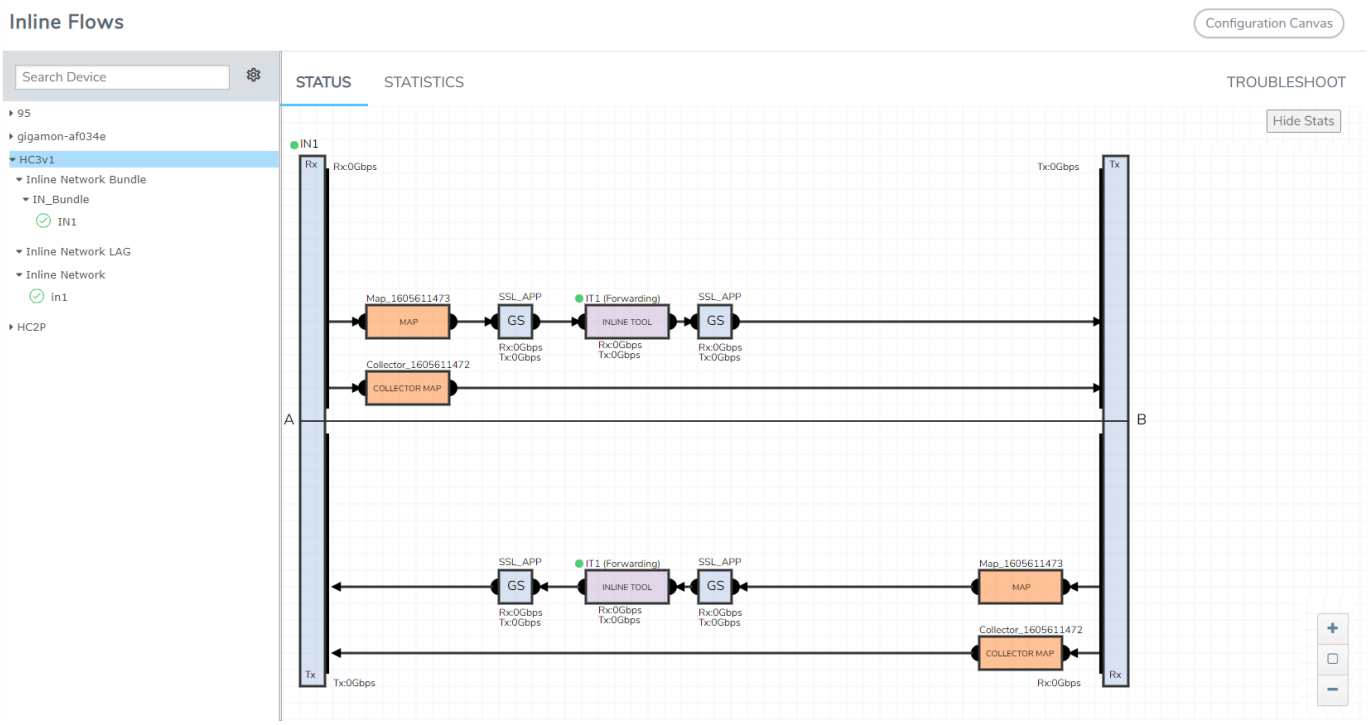

Rules and Notes

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.

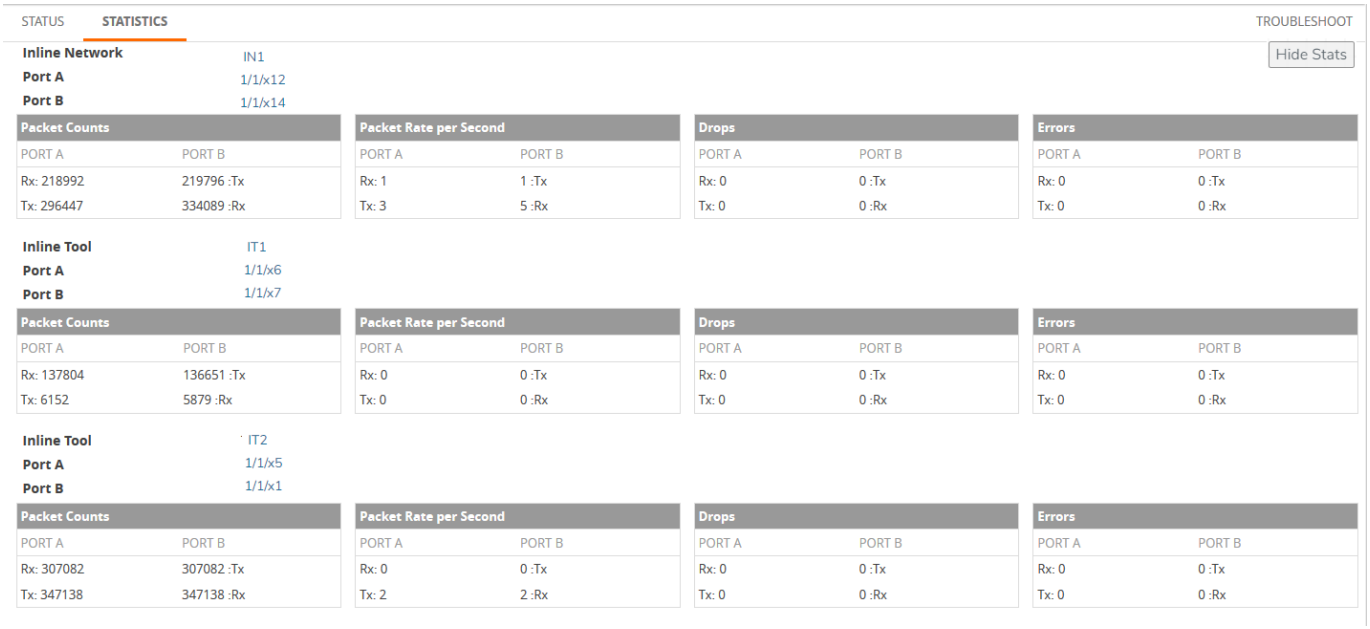

HSM Statistics

To view HSM Statistics, click the **View HSM Stats** button on the Statistics tab.

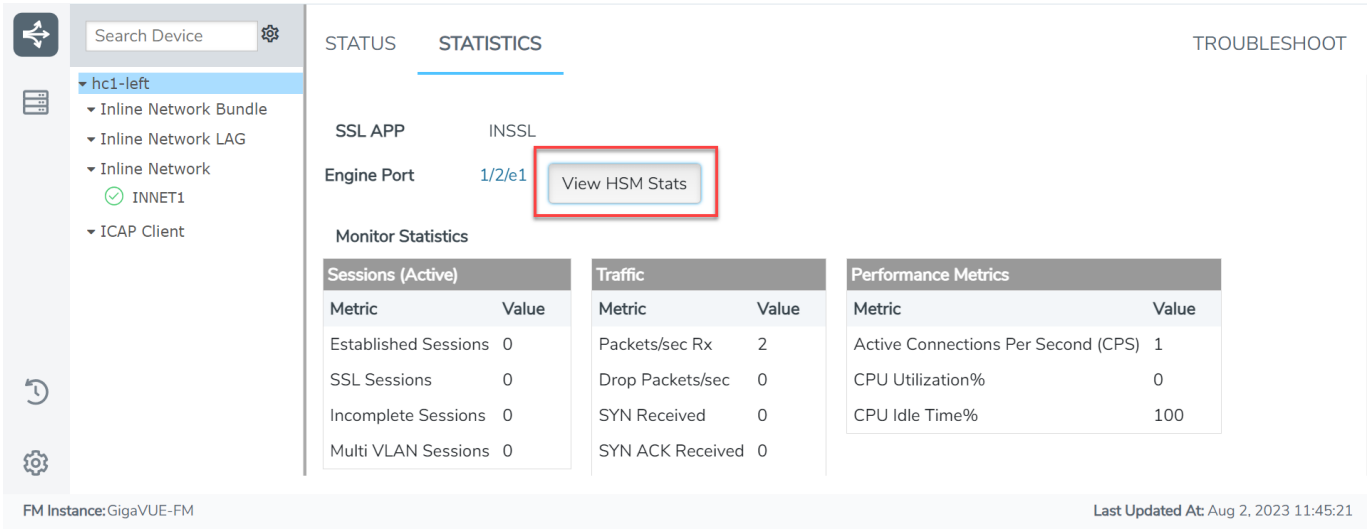

The HSM Statistics window includes the following tabs:

- **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.

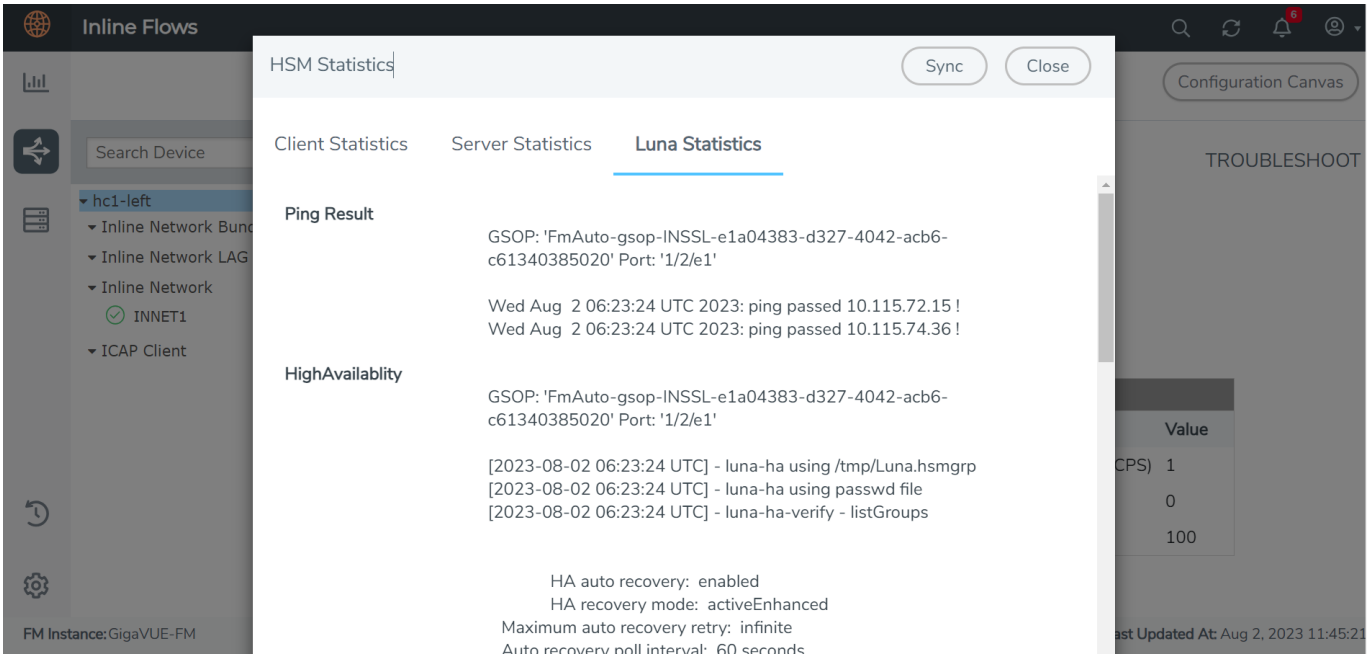

ICAP Show Stats

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.

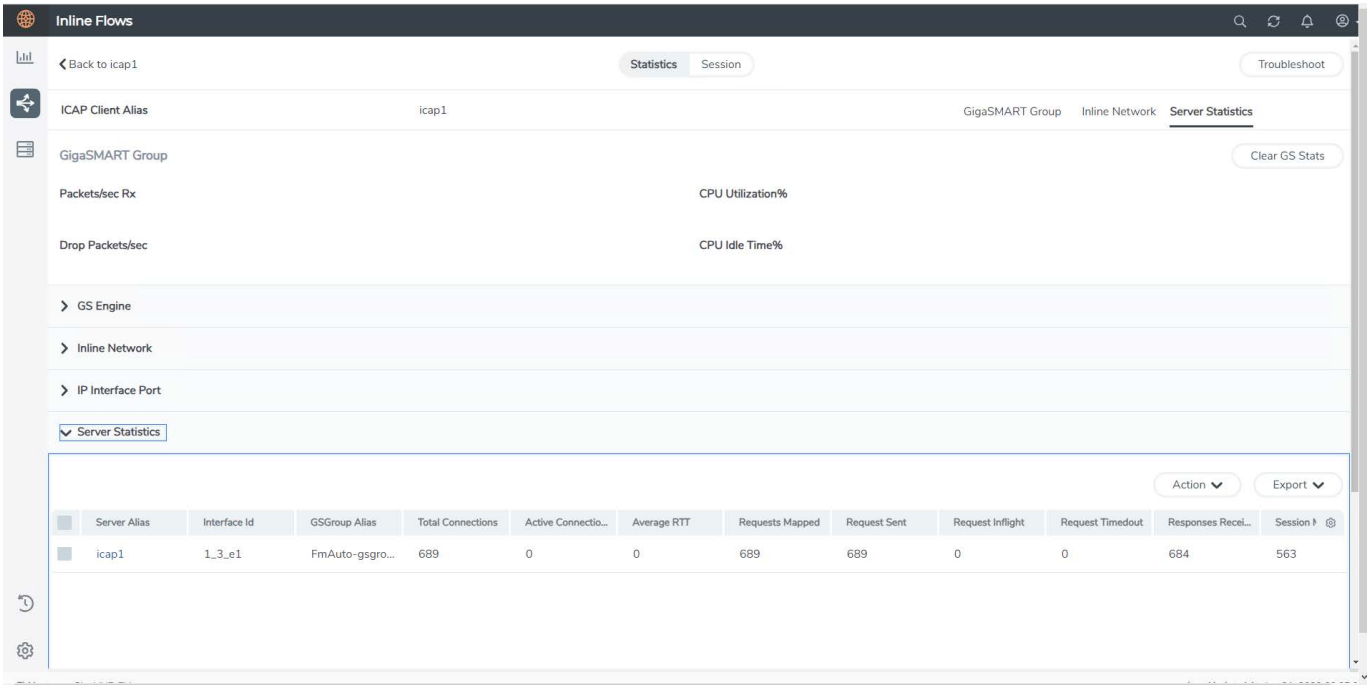

To view ICAP session statistics details, click the **Sessions** tab.

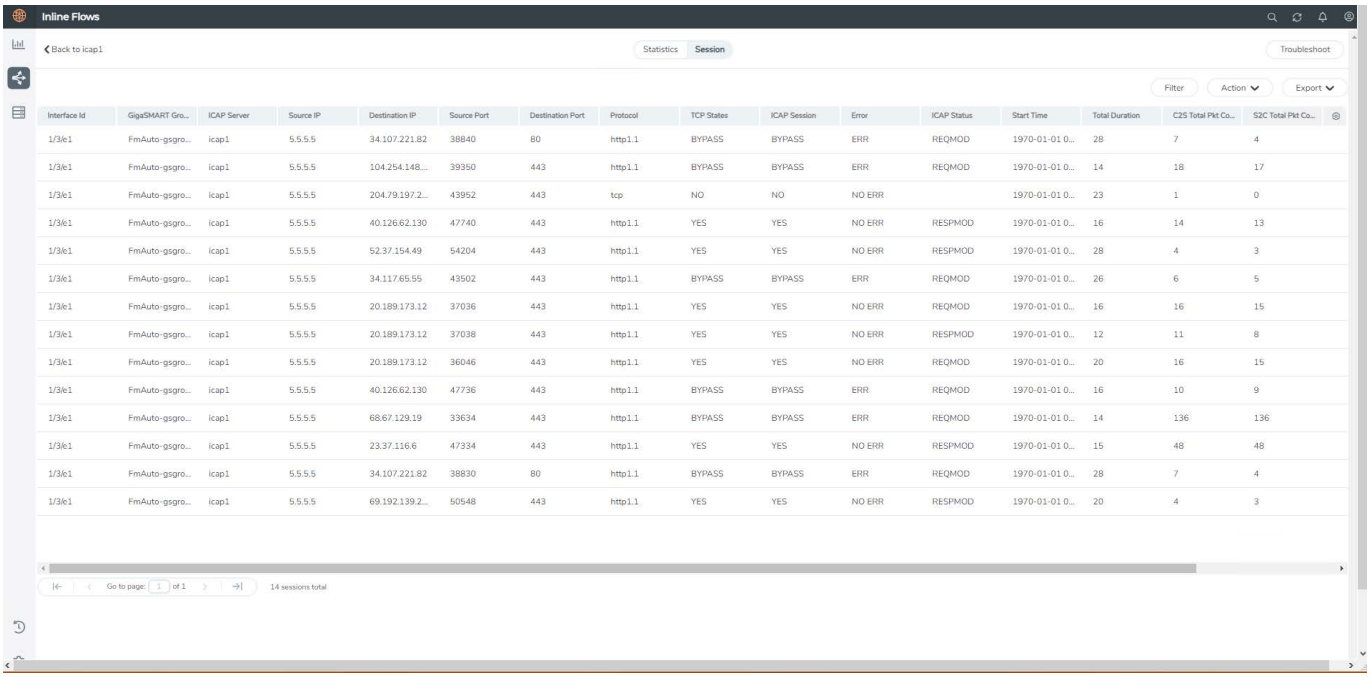

Refer to the below table for more details on the sessions tab.

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.

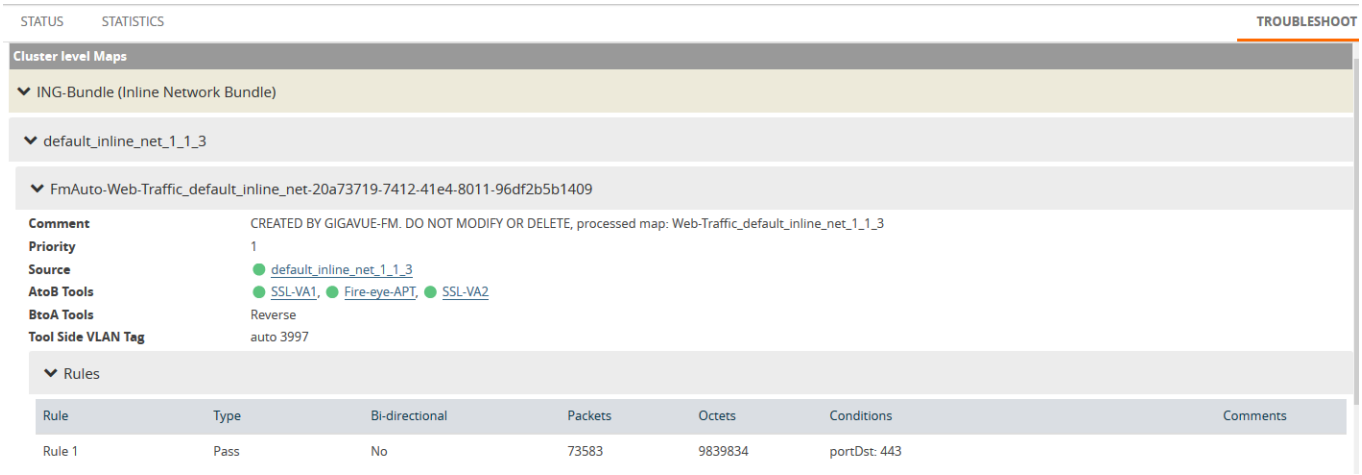

To troubleshoot any issue or failure in your flexible inline flow:

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.

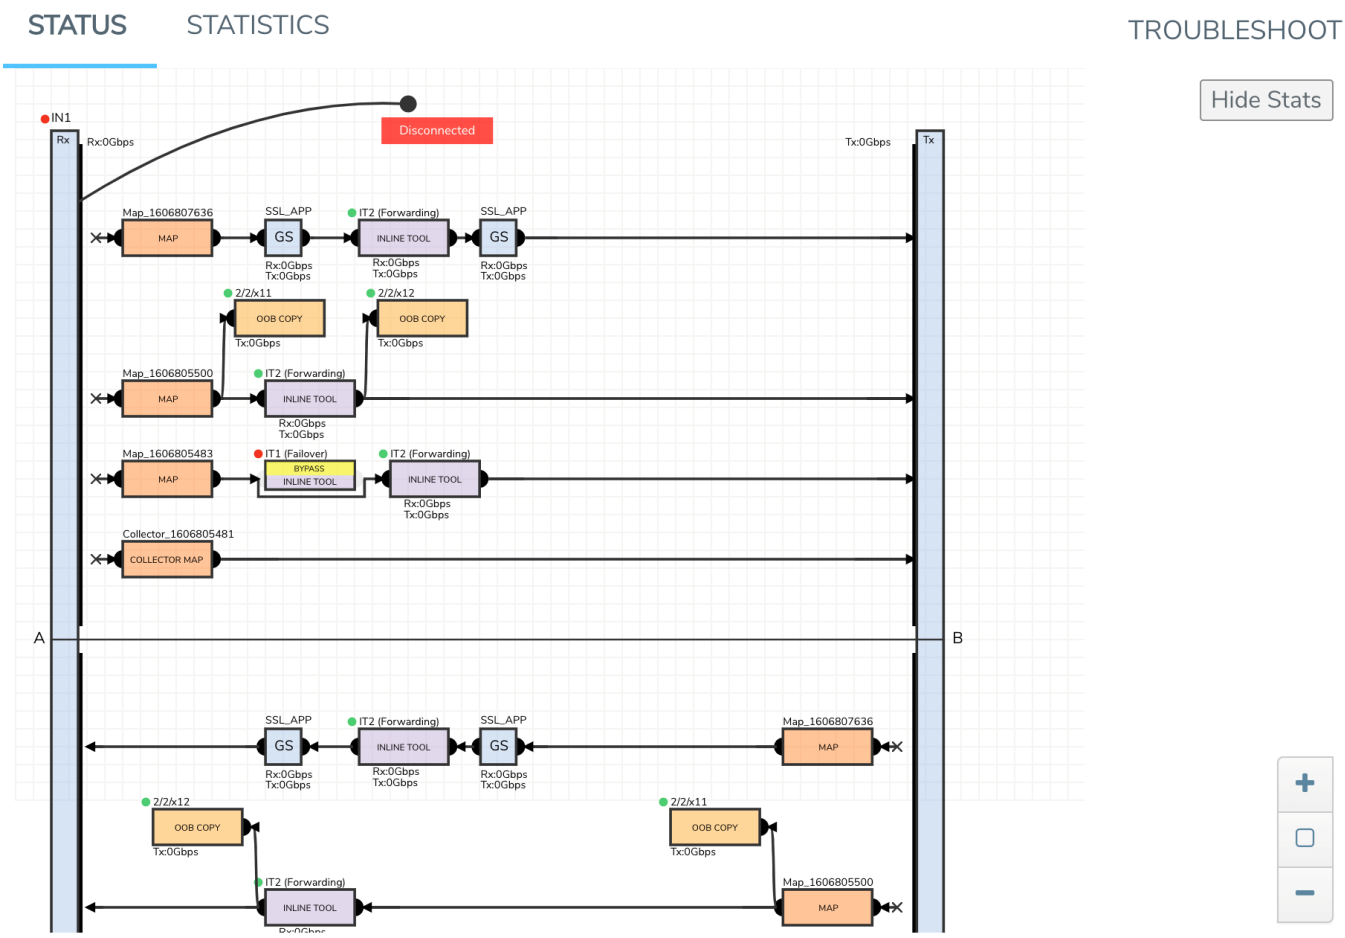

Figure 23 *Inline Network Forwarding States*

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 14: 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 the inline network	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
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.
disable	traffic path set to to-inline-tool	both far-end ports are up	all inline tools involved (directly	NORMAL	The traffic is guided between

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
			or indirectly as members of inline tool groups) in the maps originating from the inline network are in the *up* operational state		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 inline tool groups involved in the maps originating from the inline	NETWORK PORTS FORCED DOWN	No traffic is exchanged between the inline network ports, and the inline network

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
			network configured with the traffic path parameter to-inline-tool and failover action of network-port-forced-down is in the *down* operational state		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 from the inline network configured with to-inline-tool	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
			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 to-inline-tool and failover action of network-bypass	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
			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.

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.

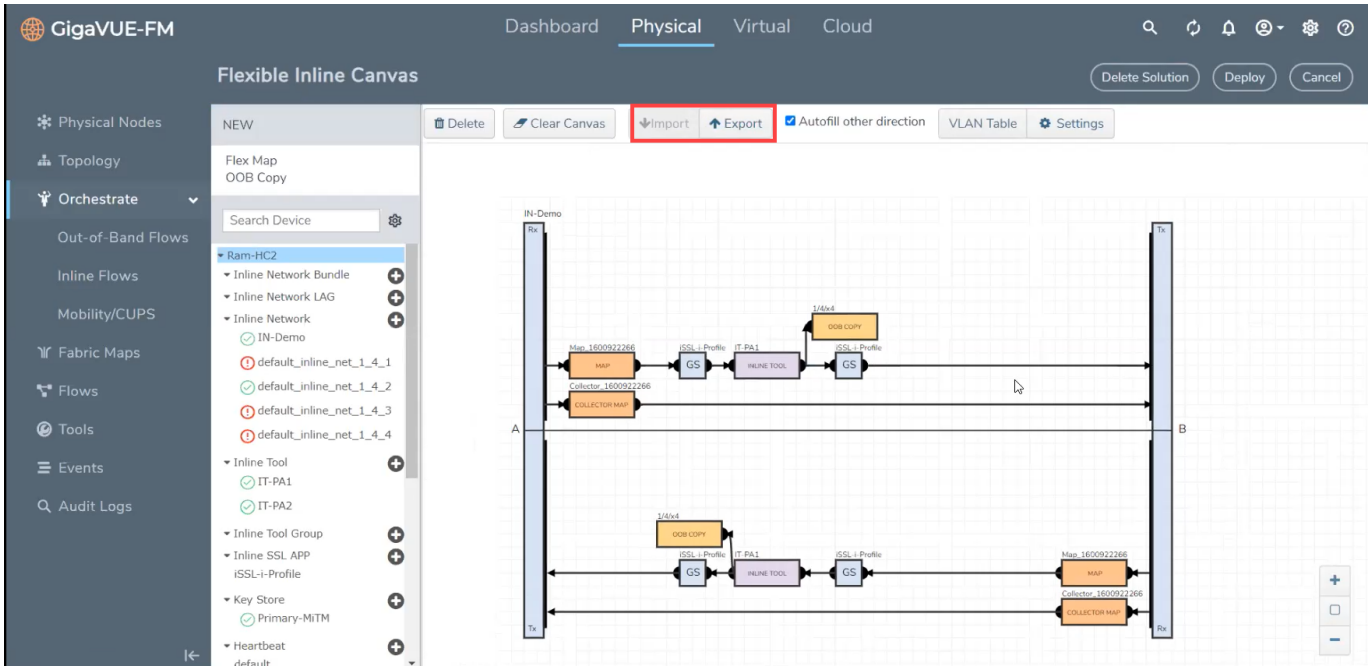

Rules and Notes

- 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"*.

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 24Timestamping illustrates how timestamps are added to ingress and egress packets.

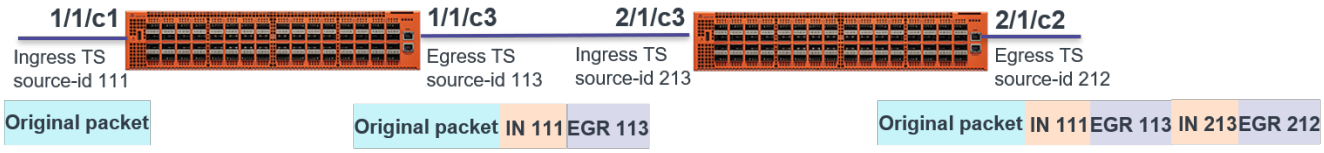

Figure 24 *Timestamping*

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.

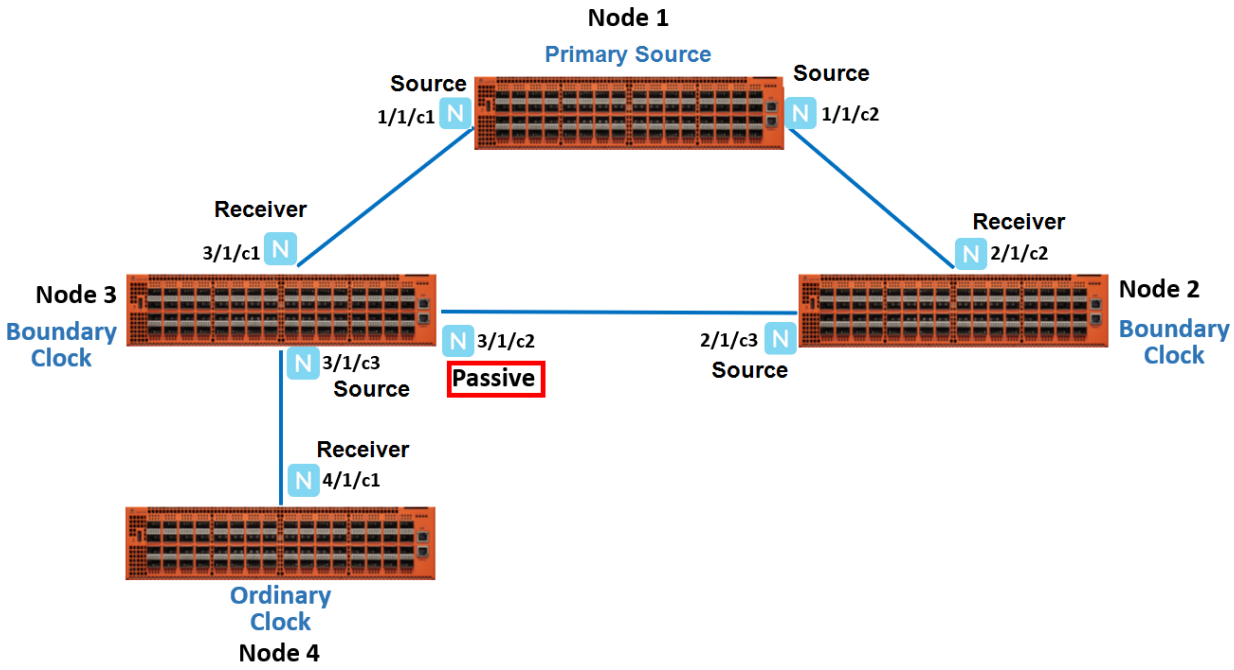

Figure 25 *PTP Deployment*

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 26PTP State Transition).

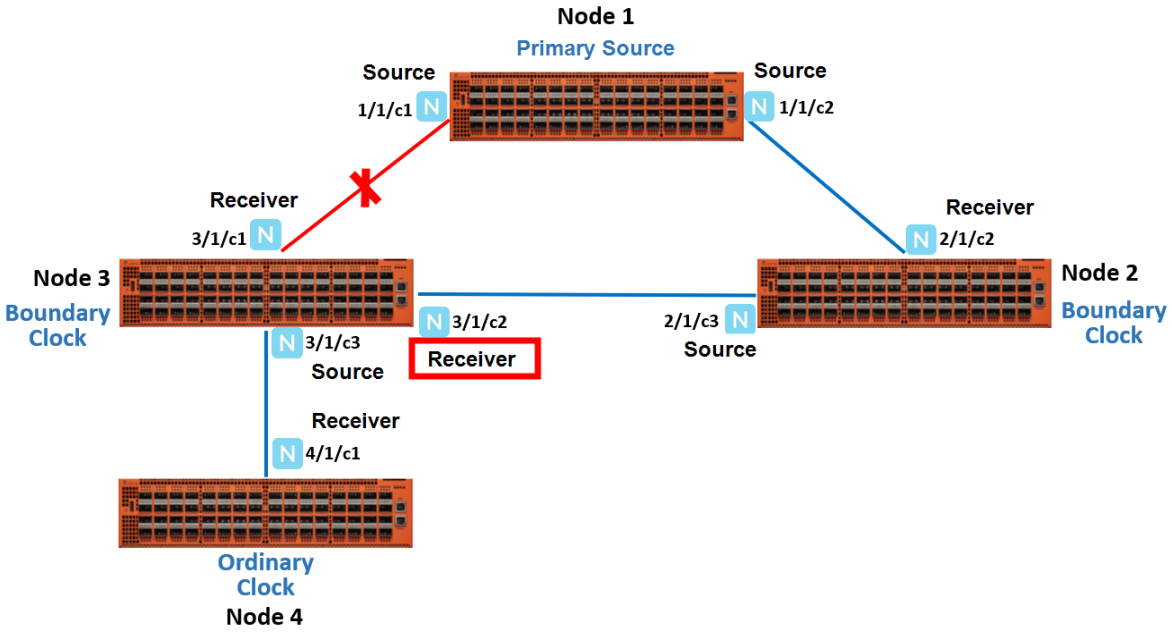

Figure 26 *PTP State Transition*

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. On the left navigation pane, click  , and then select **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. On the left navigation pane, click  , and then select **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. On the left navigation pane, click , and then select **Physical > Nodes**..
2. In the Physical Nodes page, click the required cluster ID for which you want to configure timestamp.
3. From the left navigation pane, 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 27 *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

To view:	Refer to:
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 28PTP Port State shows the PTP configuration details.

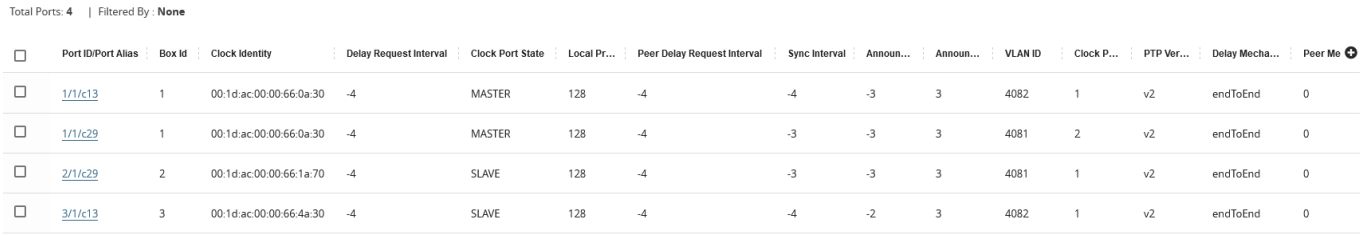

Total Ports: 4 | Filtered By : None

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 28 *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.

PTP Port State	Description
	• 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 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	The minimum interval allowed between PTP delay request

PTP Port State	Description
Interval	messages between peer ports.
Sync Interval	The time interval between PTP synchronization messages on the port.
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 29PTP Clock Details shows the PTP clock details.

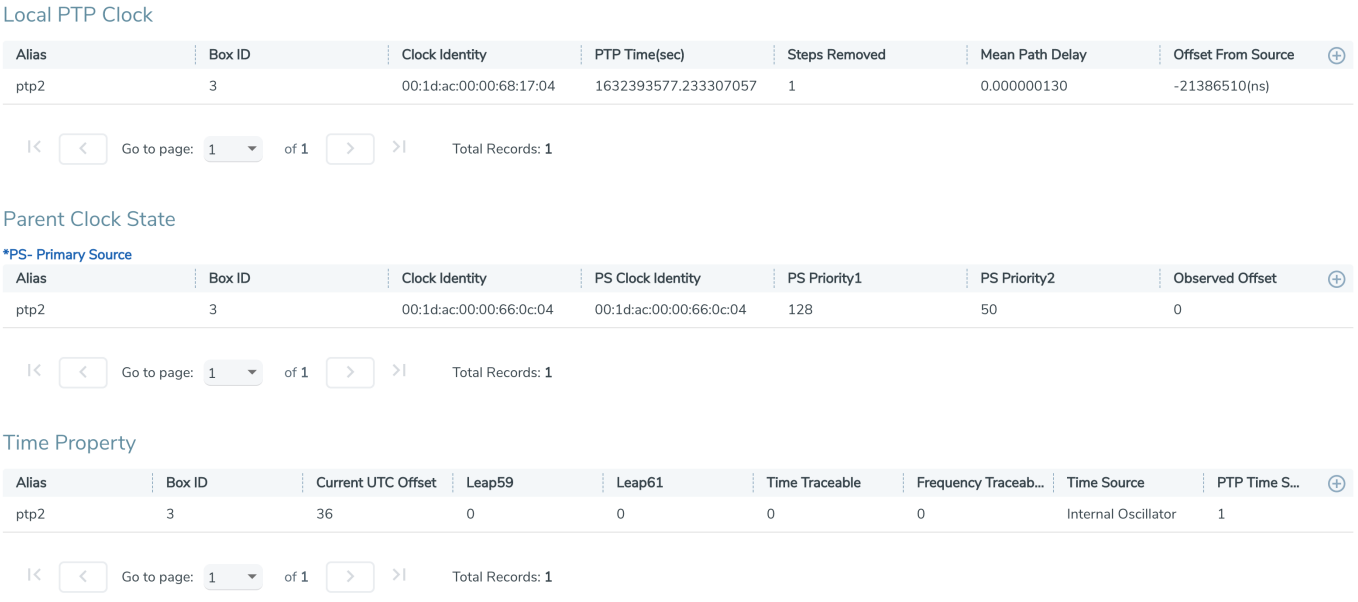

Figure 29 *PTP Clock Details*

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 29PTP 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.

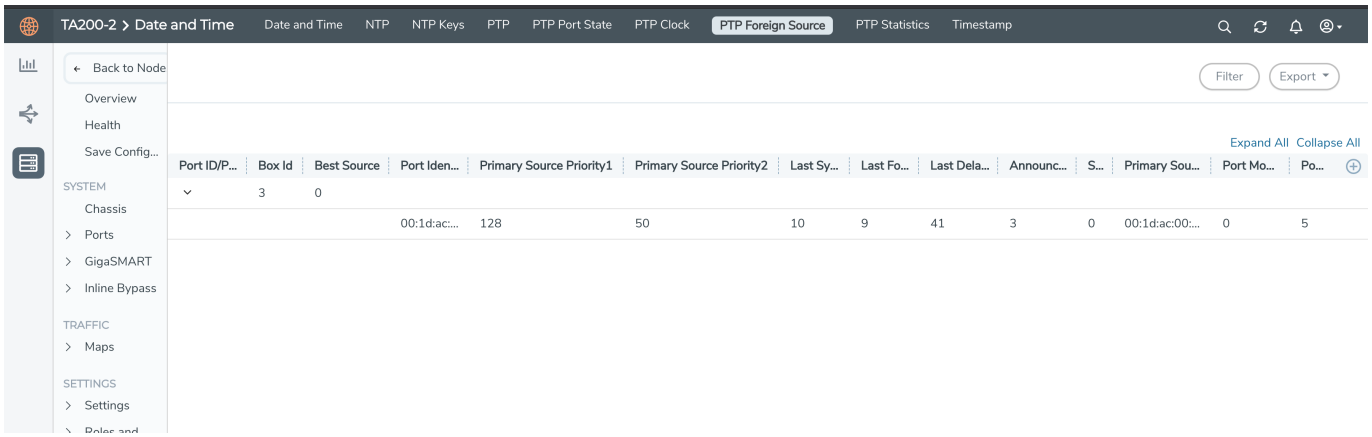

Figure 30 *PTP Foreign Source Details*

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 31PTP Statistics shows the PTP statistics details.

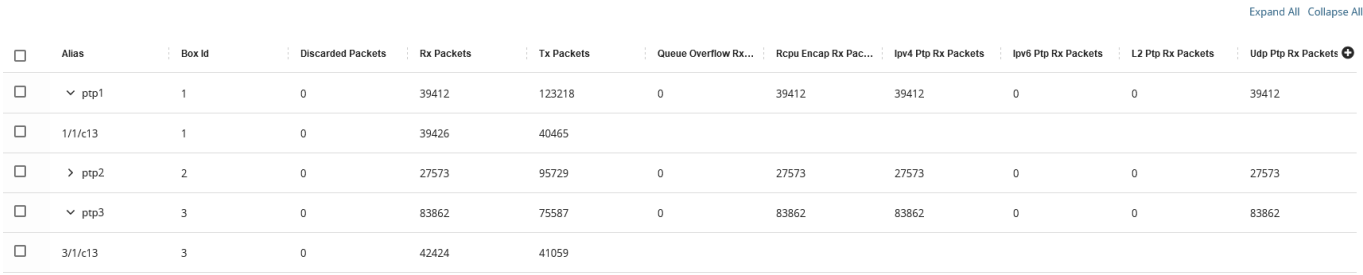

Expand All Collapse All

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 31 *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 and timestamp is supported only on GigaVUE-TA200 and GigaVUE-TA200Edevice.
- 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.
- 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.
  - Do not connect standalone nodes or clusters to the spine nodes.
  - 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.

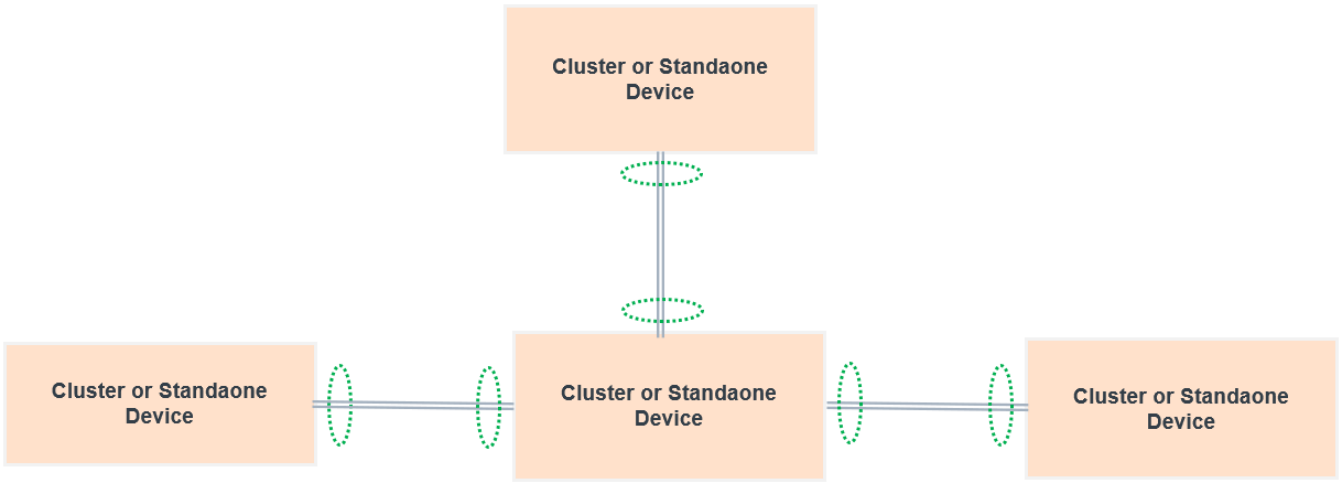

Supported Multihop Topology

The following multihop topology has only one path from source to destination.

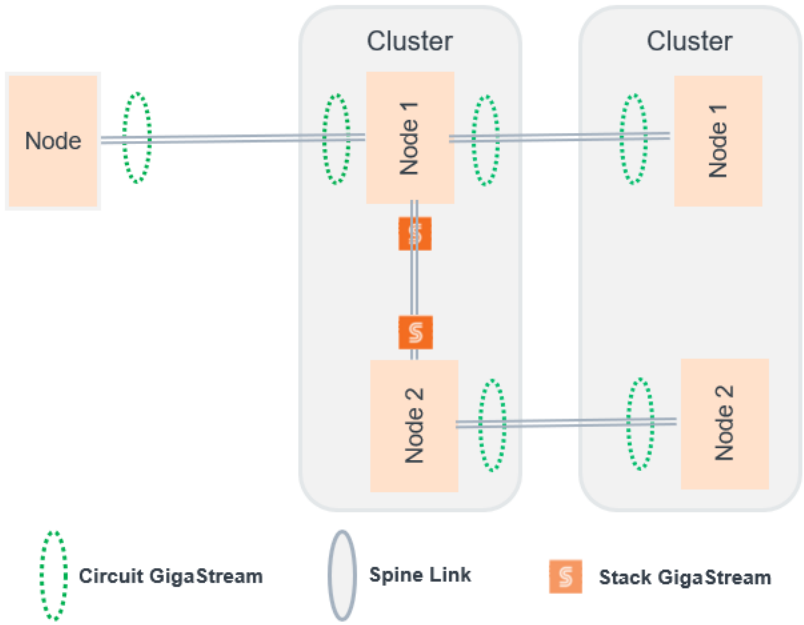

Unsupported Multihop Topology

The following multihop topology has more than one path from source to destination and therefore it is not supported.

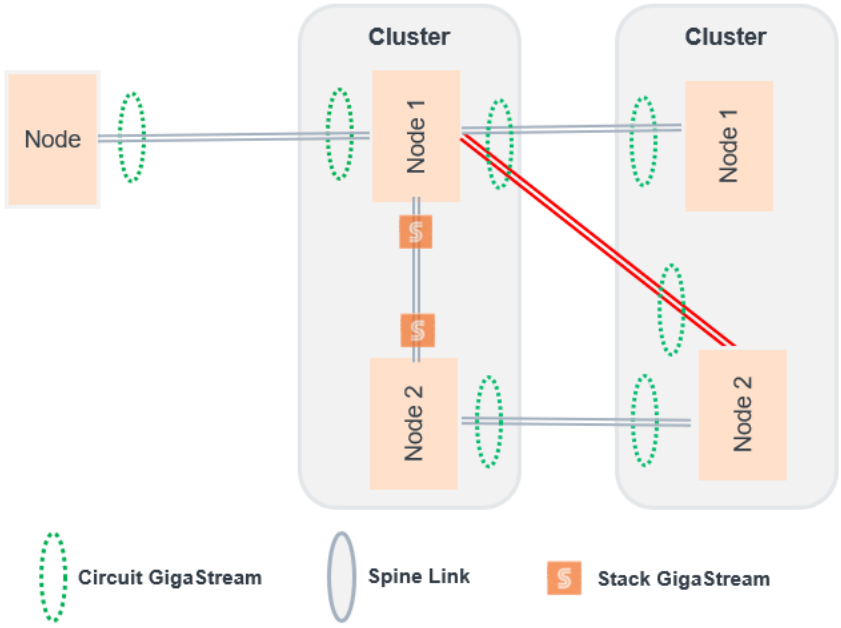

Leaf and Spine Topology

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.

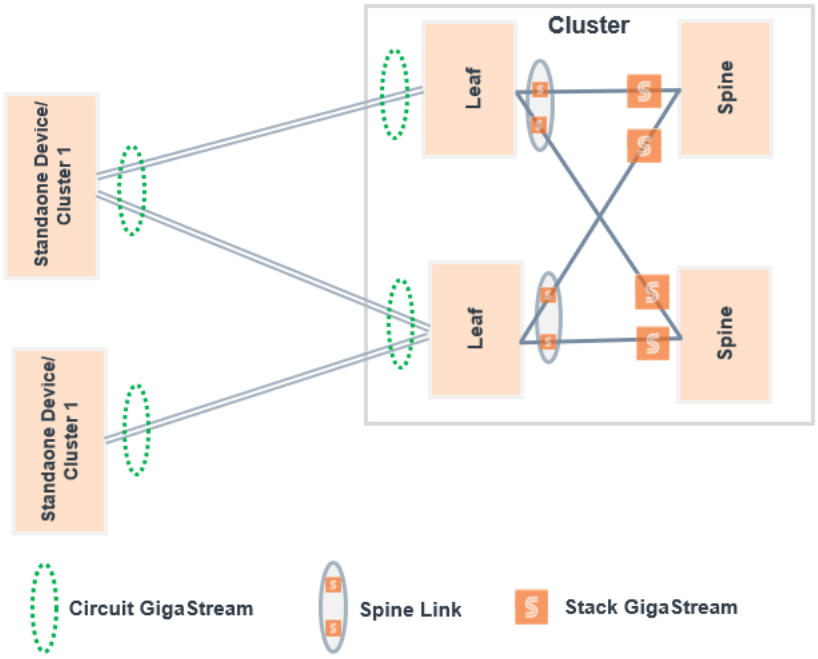

Dual Multipath Leaf and Spine Topology

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.

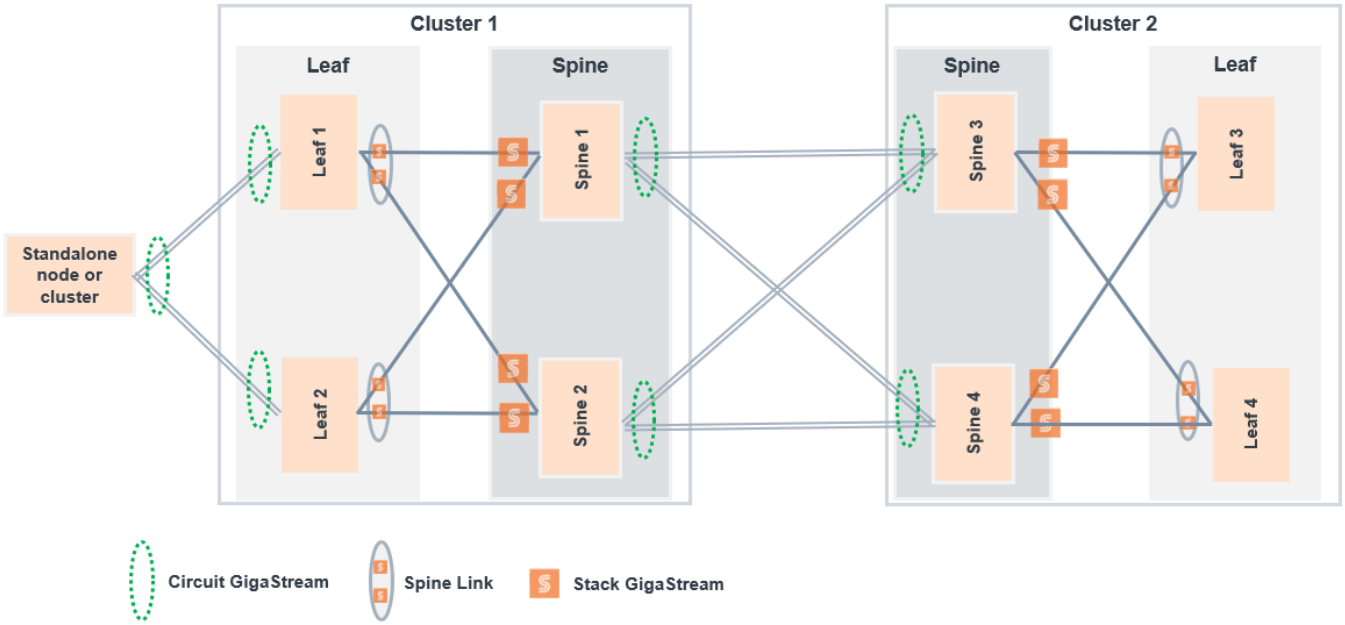

Supported Failovers

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

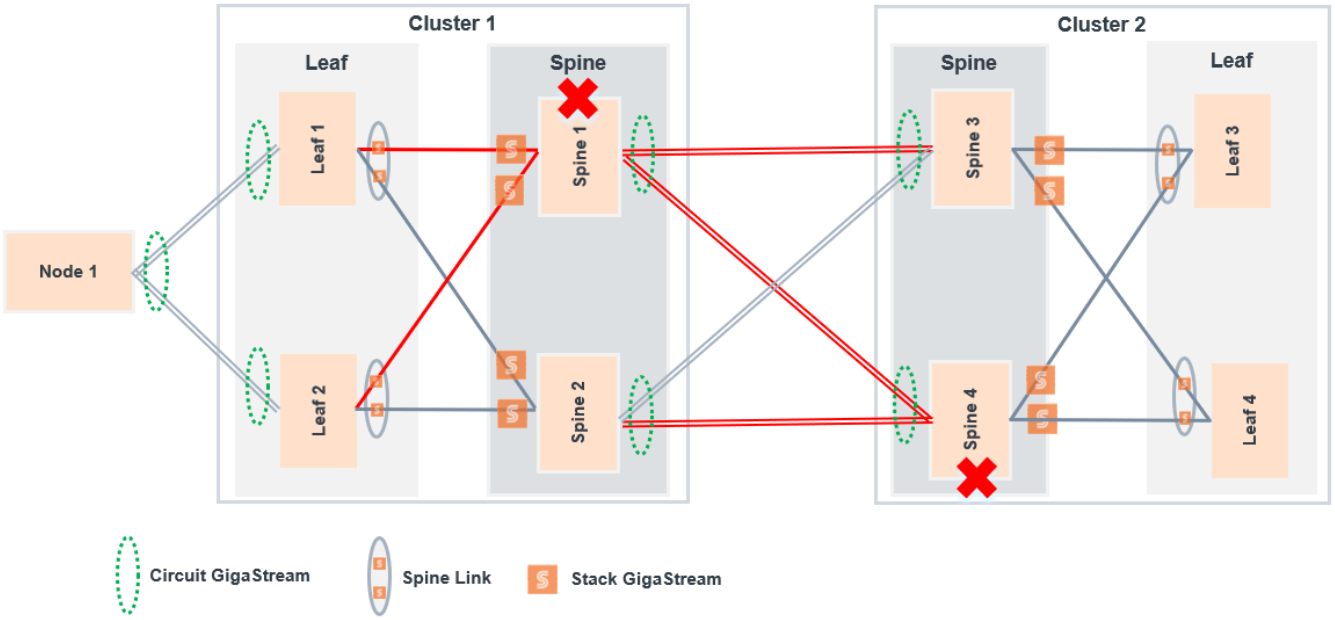

In the above topology:

- 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

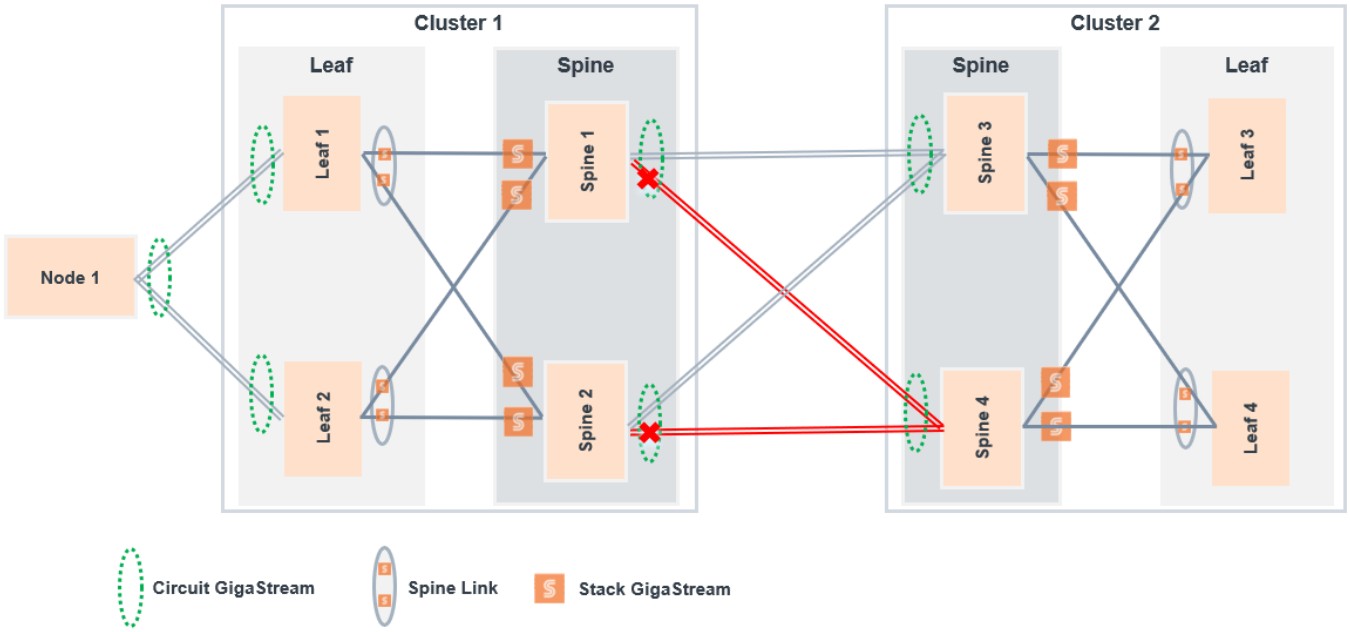

In the above topology:

- 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

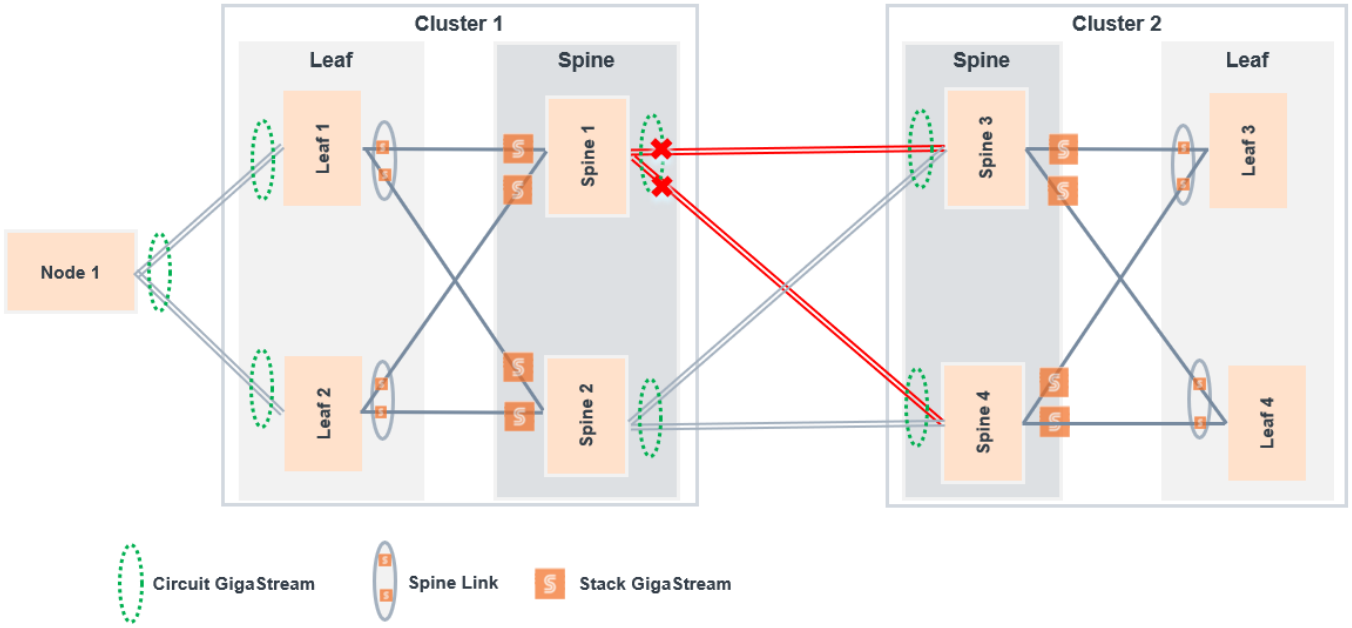

In the above topology:

- 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

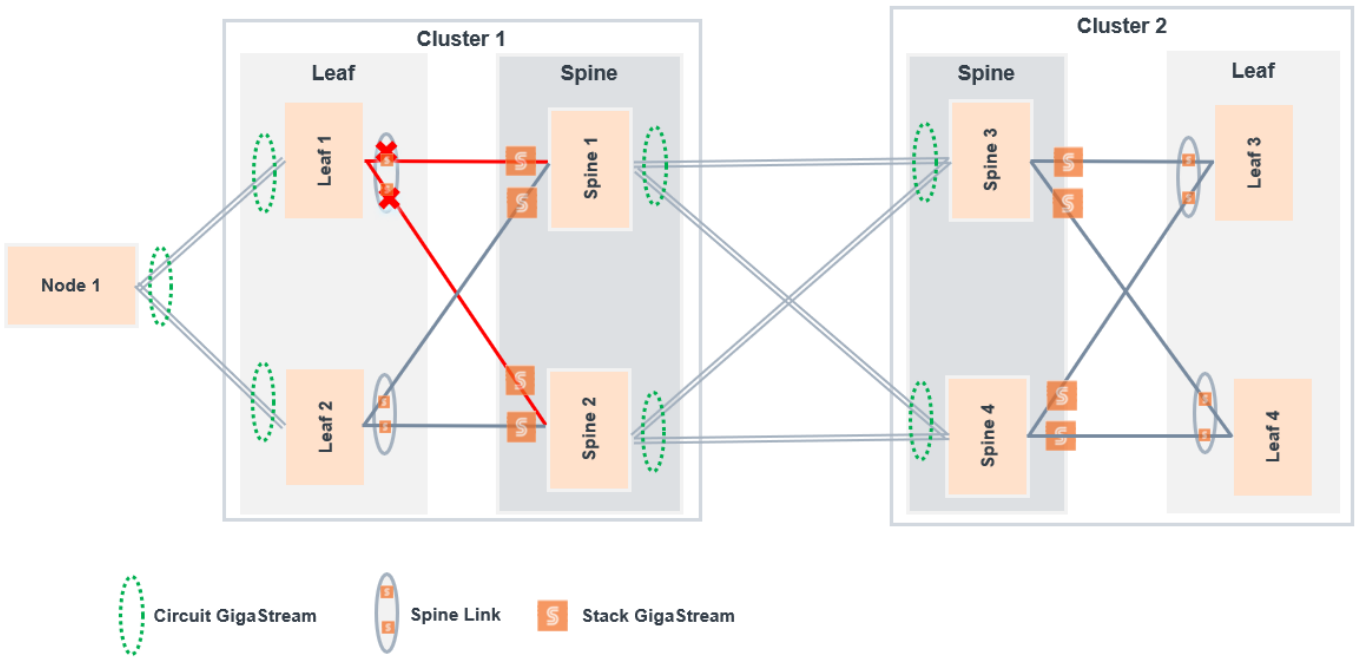

In the above topology, traffic received on leaf 1 will be discarded as the following spine links are down:

- 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 Share Circuit ID Resource.
- 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.

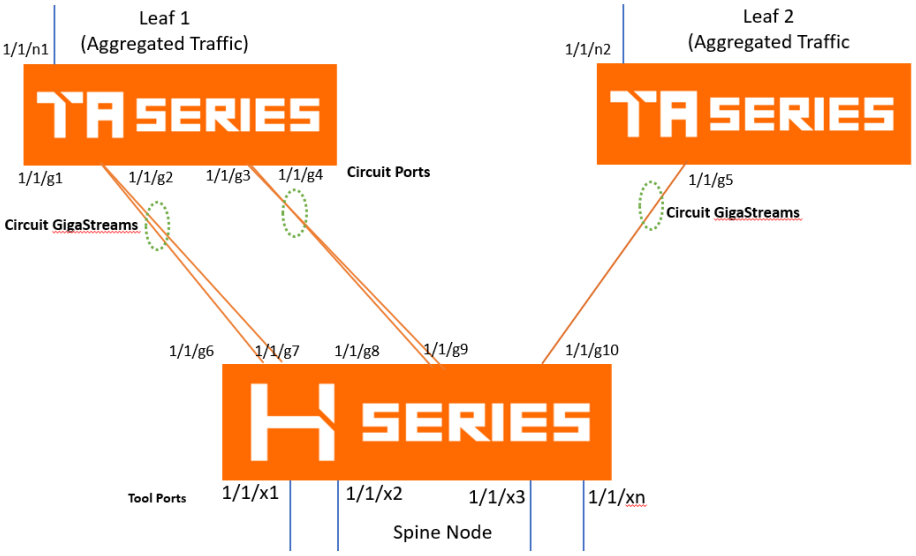

Figure 32 *Circuit GigaStreams in Fabric Maps*

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, click on 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.

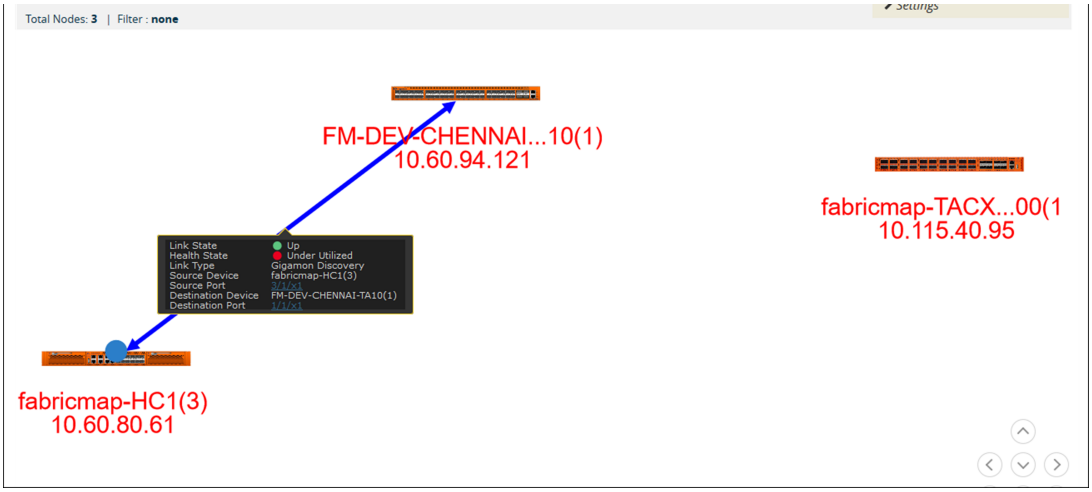

Figure 33 Connecting Devices

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. On the left navigation pane, click  , and then select **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.

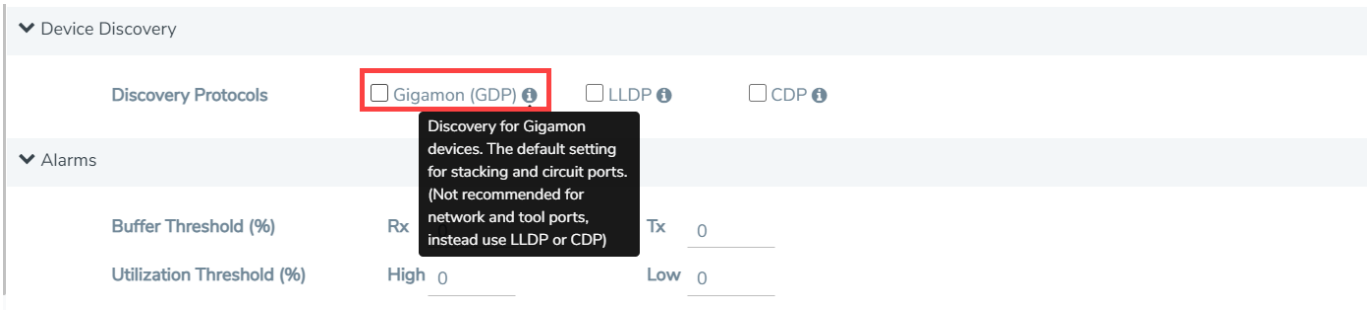

Figure 34 Gigamon Discovery Protocol

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.

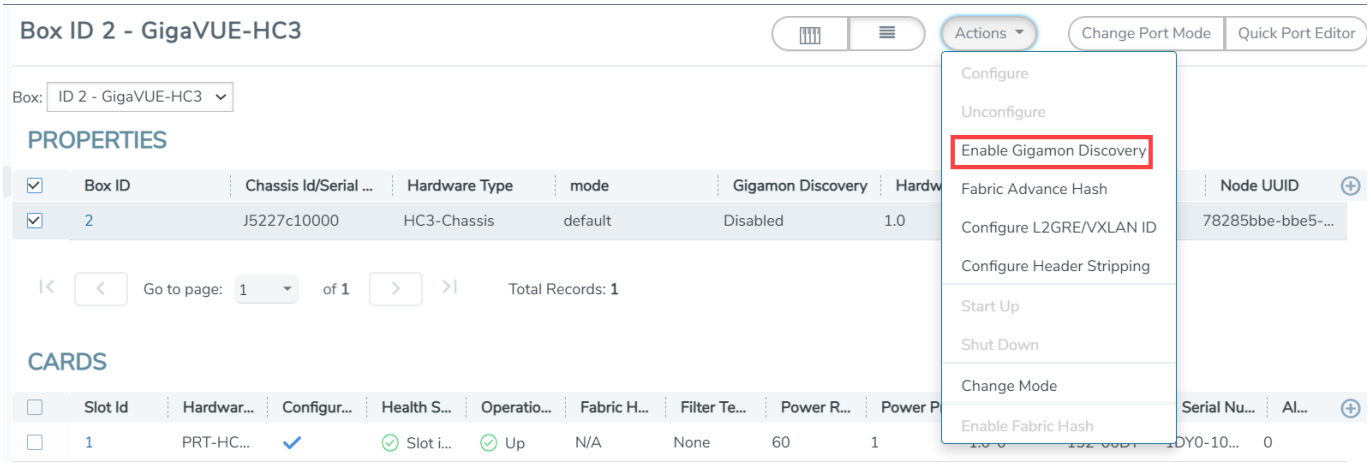

NOTE: This action must be done for both the source and the destination devices.

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. On the left navigation pane, click on  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
Type	Port type
Status	Health status of the fabric map

1. Click **Create**.The create map screen is displayed.

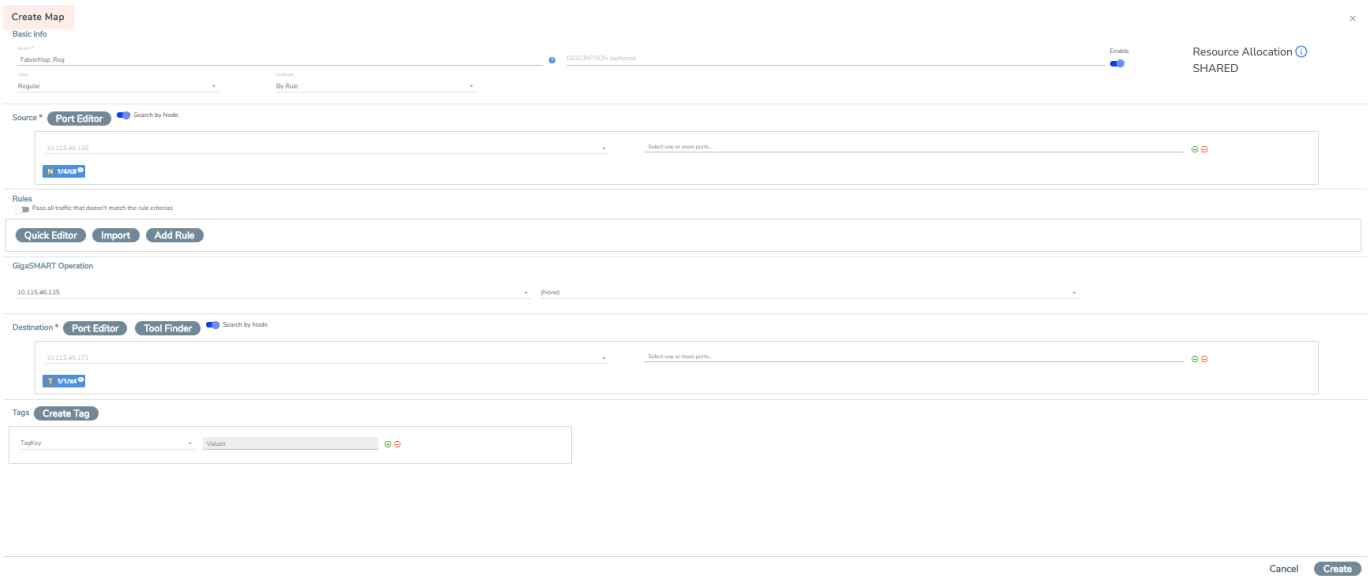

2. Enter map parameters as described in the following table:

Field	Description
Basic Info	

Field	Description
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 ▪ 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.

Field	Description
	For detailed information about map rules, refer to Map Rules. **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. • 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.

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.

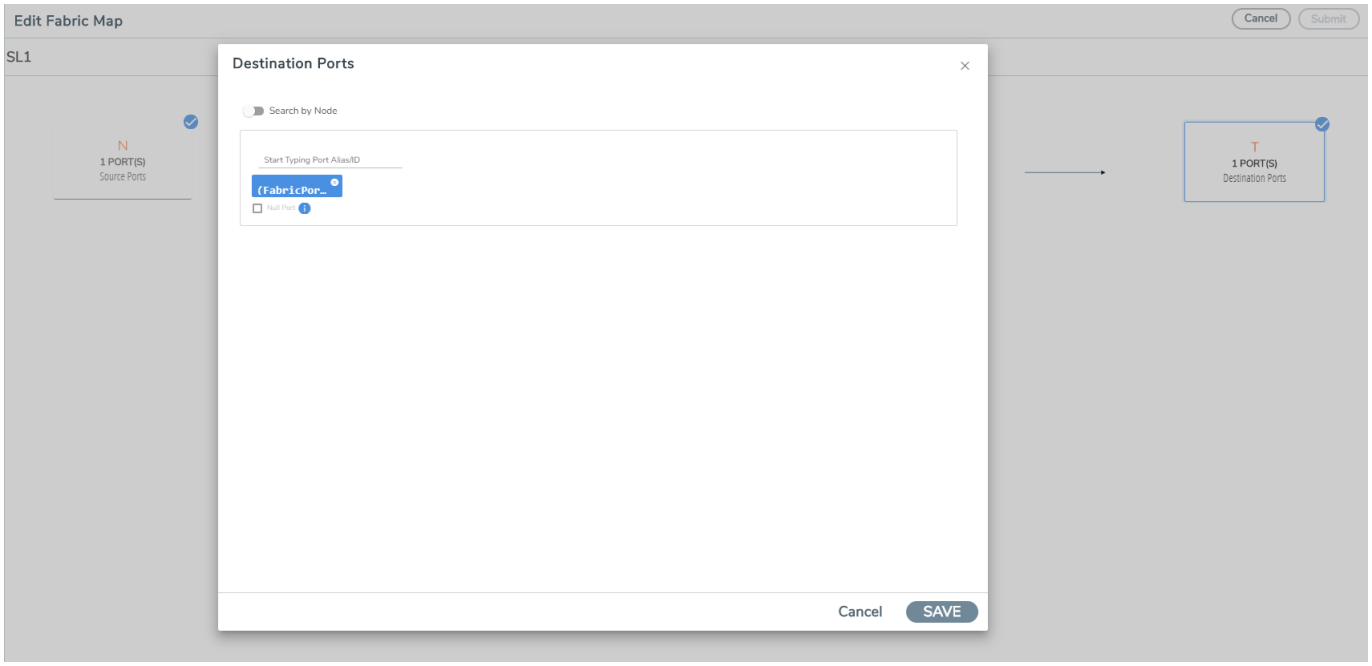

Delete Fabric Maps

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, click on 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.

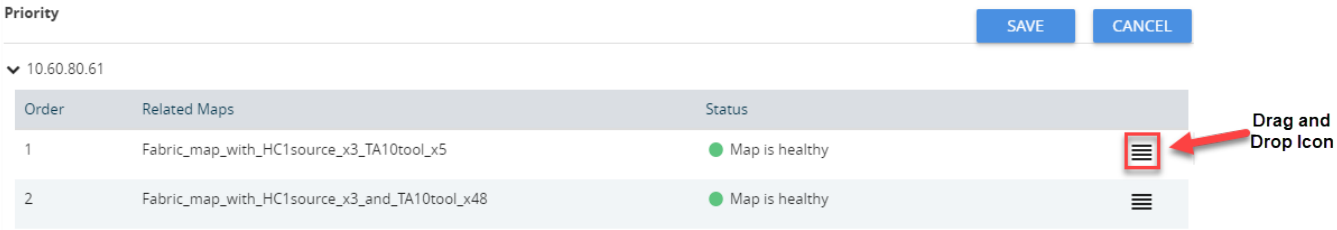

6. Click **Save** to save the changed priority.

NOTE: The priority list includes both flow maps and cluster maps, which are grouped and prioritized based on the cluster ID.

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.

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, click on  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.

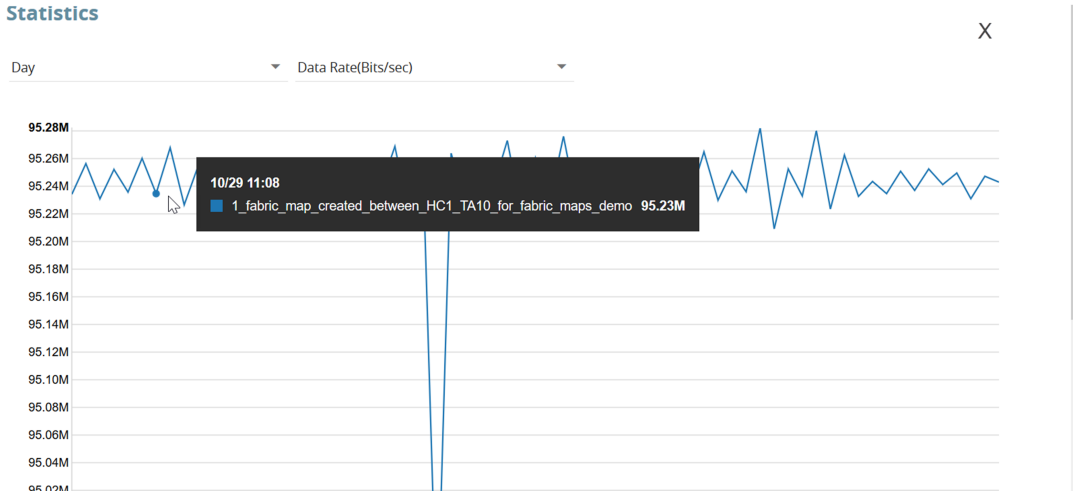

Figure 35 Fabric Maps Statistics

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)

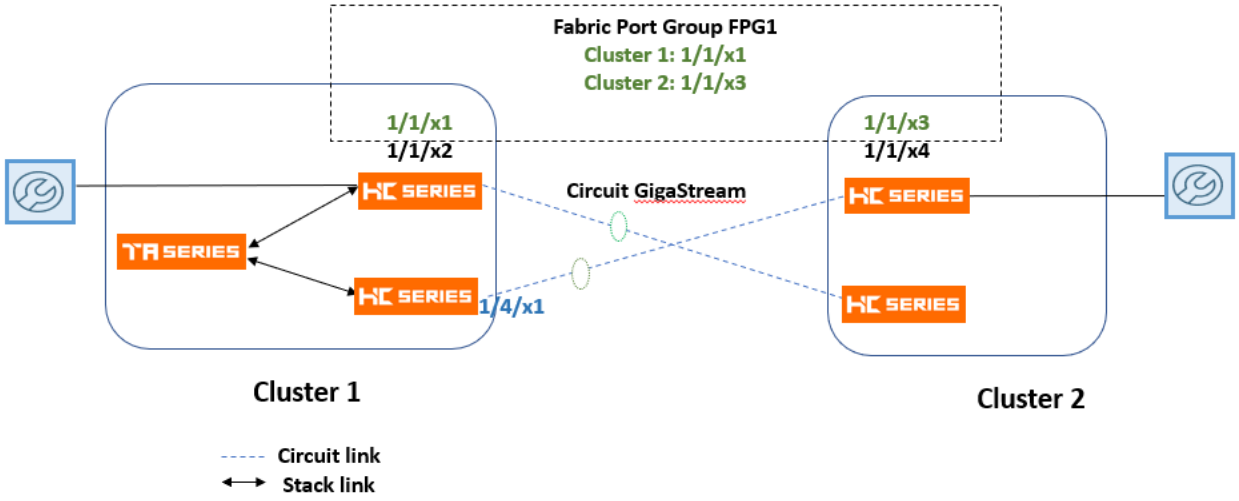

Create a Fabric Port Group FPG1 using port 1/1/x1 of cluster 1 and port 1/1/x3 of cluster 2. Define the port weight for each of the ports. For the Fabric Port Group that is created, GigaVUE-FM:

- 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. On the left navigation pane, click on  from**Physical**select**Orchestrated Flows >Fabric Maps> Fabric Port Group.**
2. Click **Create**. The **Create Fabric Port Group** page appears.

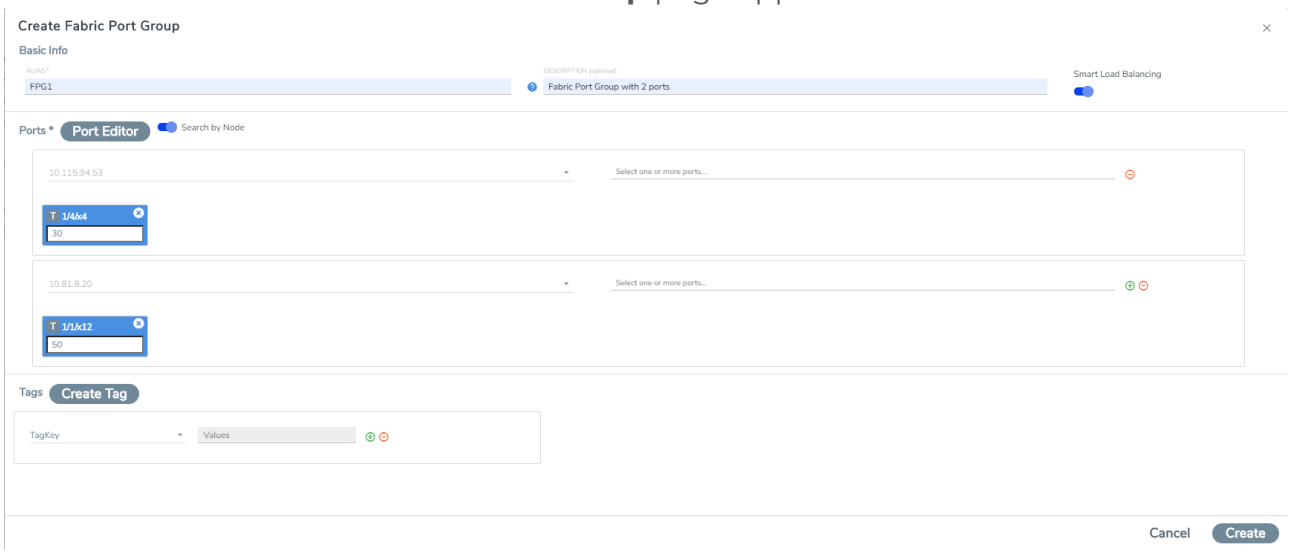

.

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**.

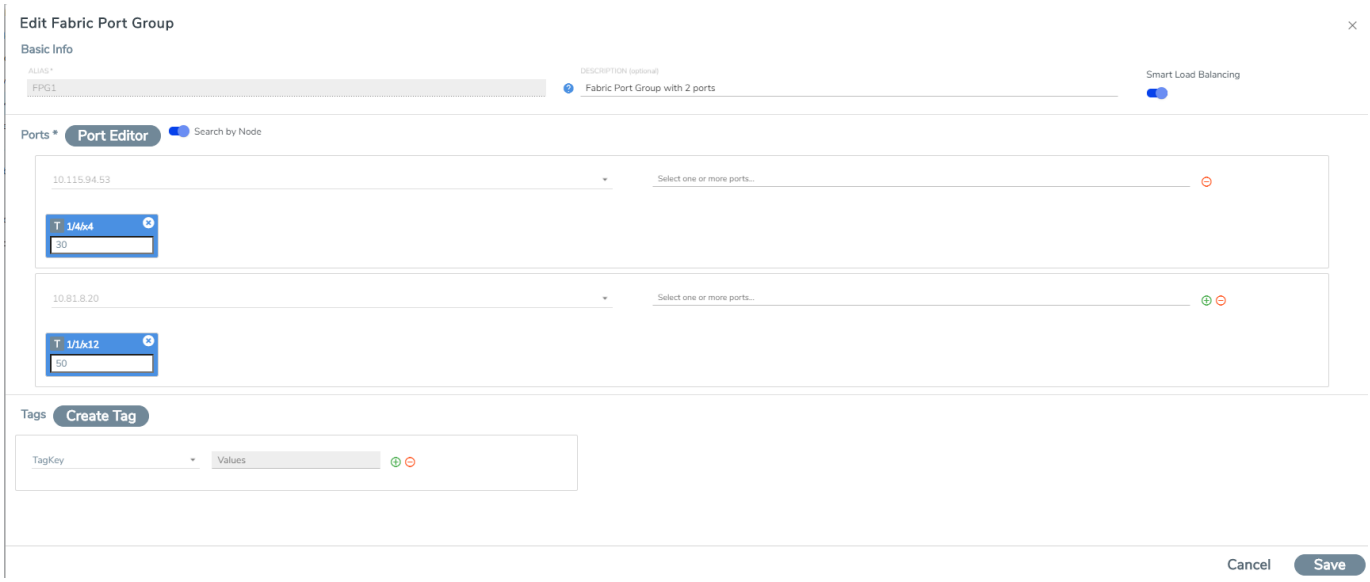

Configure Fabric Map Using Fabric Port Group

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

There may be 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, click on  under **Physical > Fabric Maps**.
2. Select a **Fabric Map** to troubleshoot.
3. Select **View Details** from the **Task** drop-down menu.

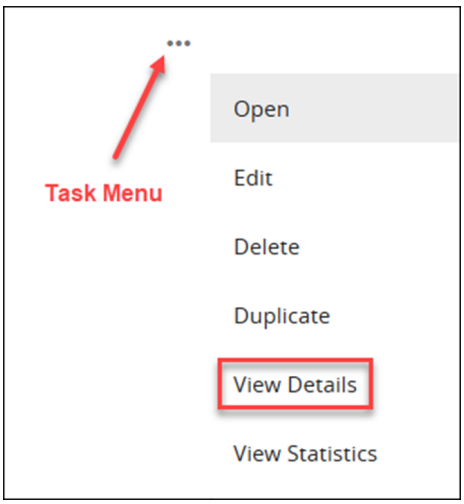

Figure 36 View Details.

The Details screen displays fabric map component details.

4. Click the **More Info** link to display the fabric map details.
5. Click the **Troubleshoot** link to display troubleshooting options.

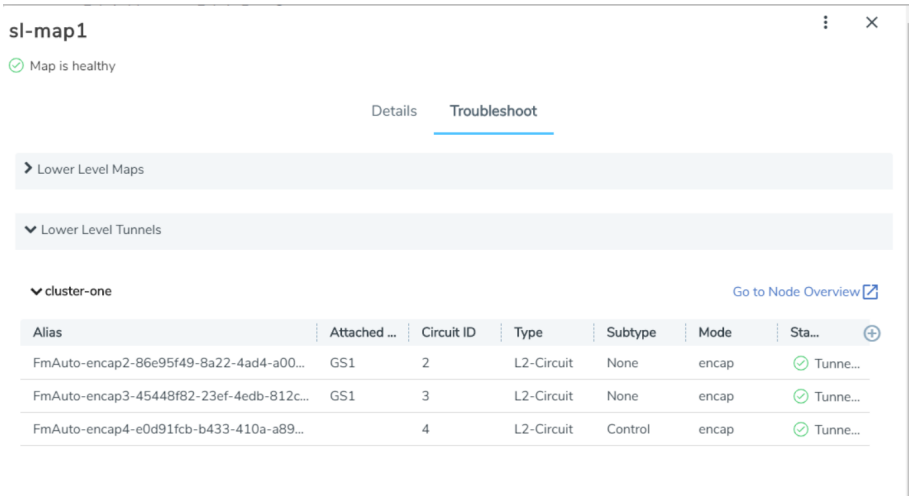

Figure 37 Troubleshooting

6. Click 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 • There is 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.	• See How to Troubleshoot Partial Success Errors for more details.

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 means there could be a port that 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 38 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.

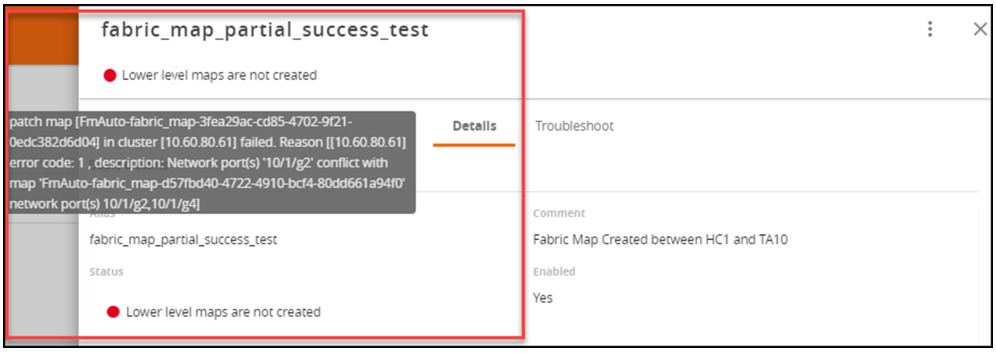

Figure 39 Fabric Maps Details

3. Click the **Troubleshooting** link. Scroll down the screen to investigate the Source and Destination ports configurations.

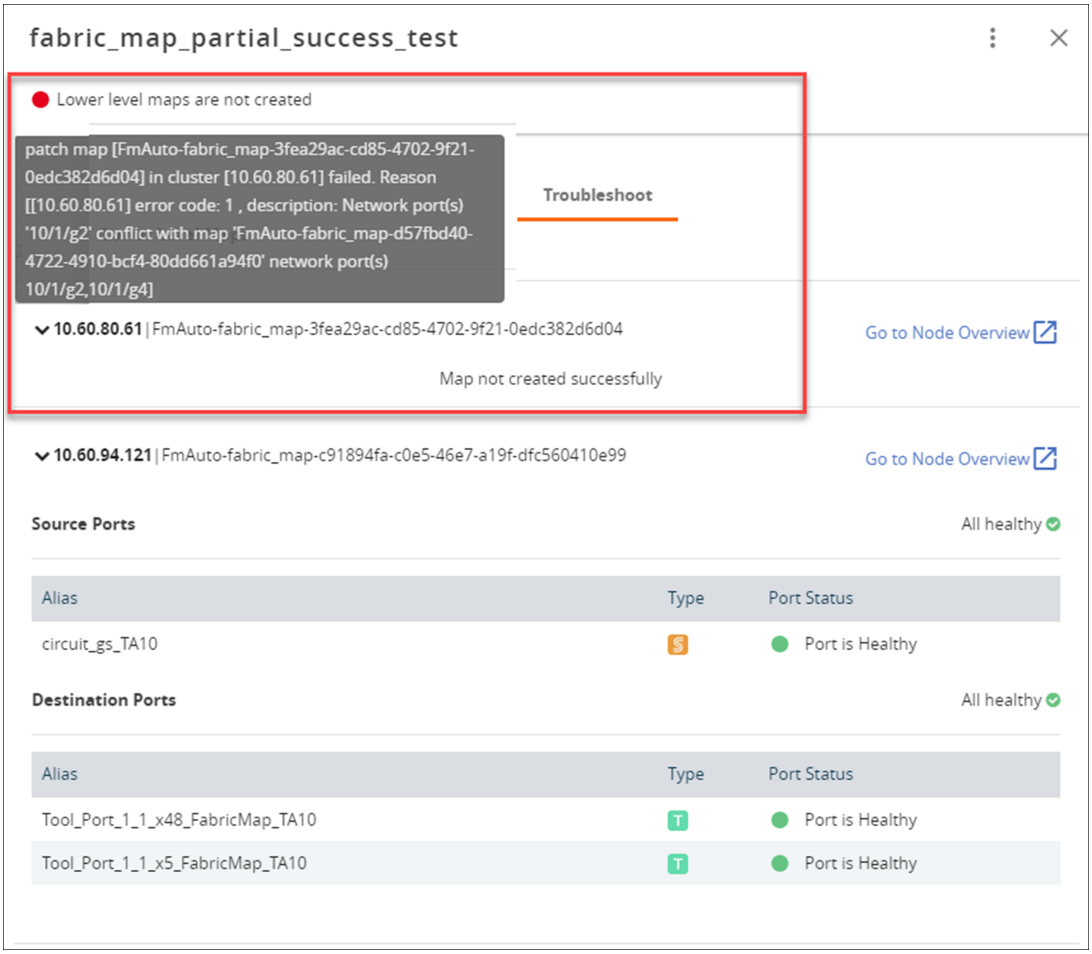

Figure 40

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 41 Editing Fabric Maps

5. Select the **Source Ports** component.

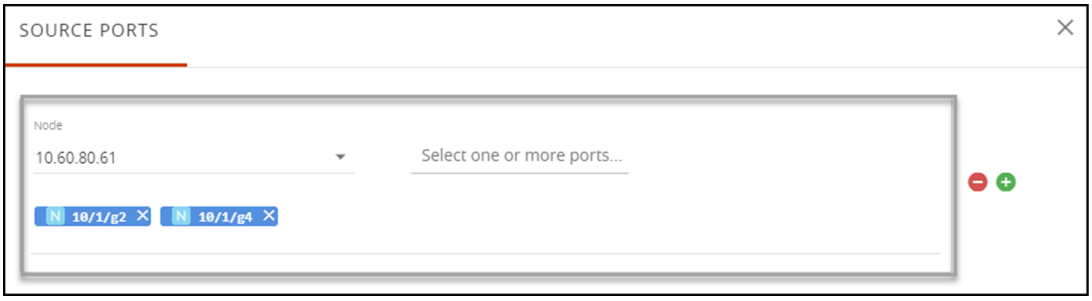

Figure 42 Source Port Component

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, click on  under **Physical > Fabric Maps.**
2. Click the **Config Audit** link at the top of the screen.

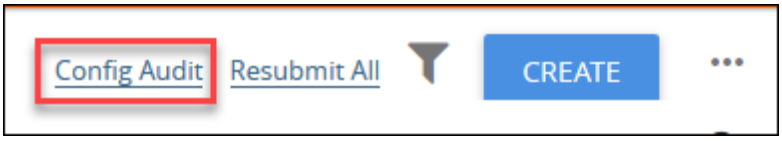

After the config audit is complete, a notification appears with the audit results.

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.**

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.

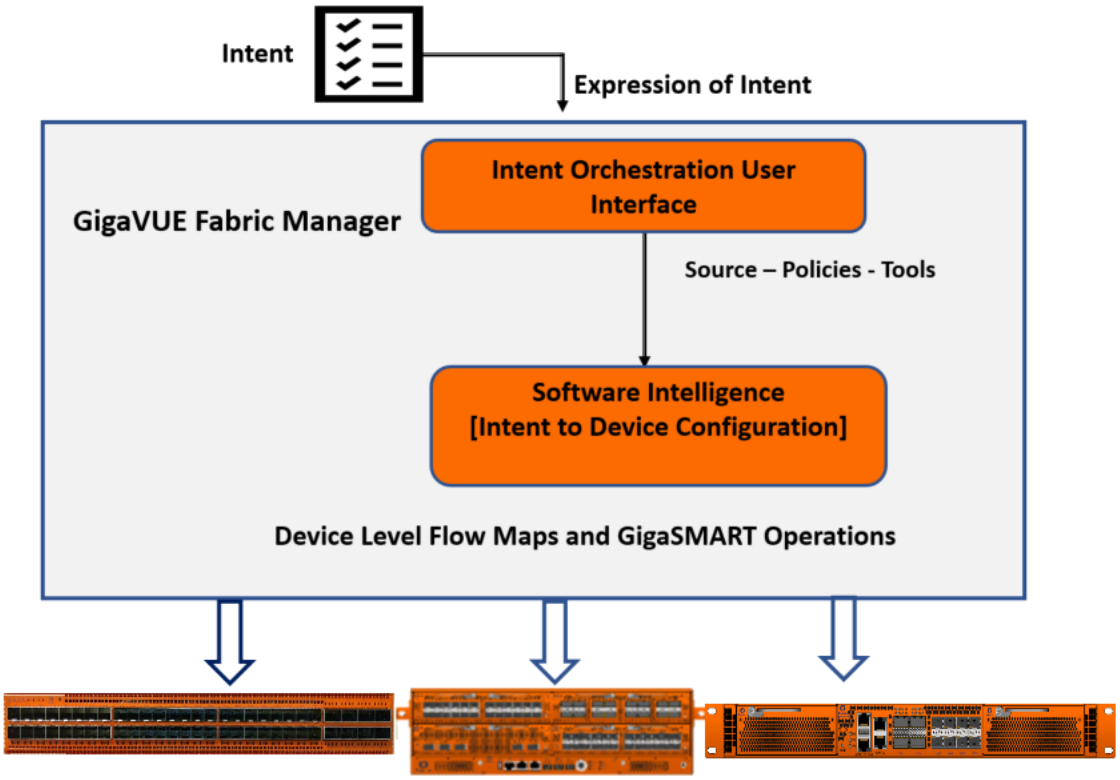

With Orchestrated Configurations, you can:

- 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.

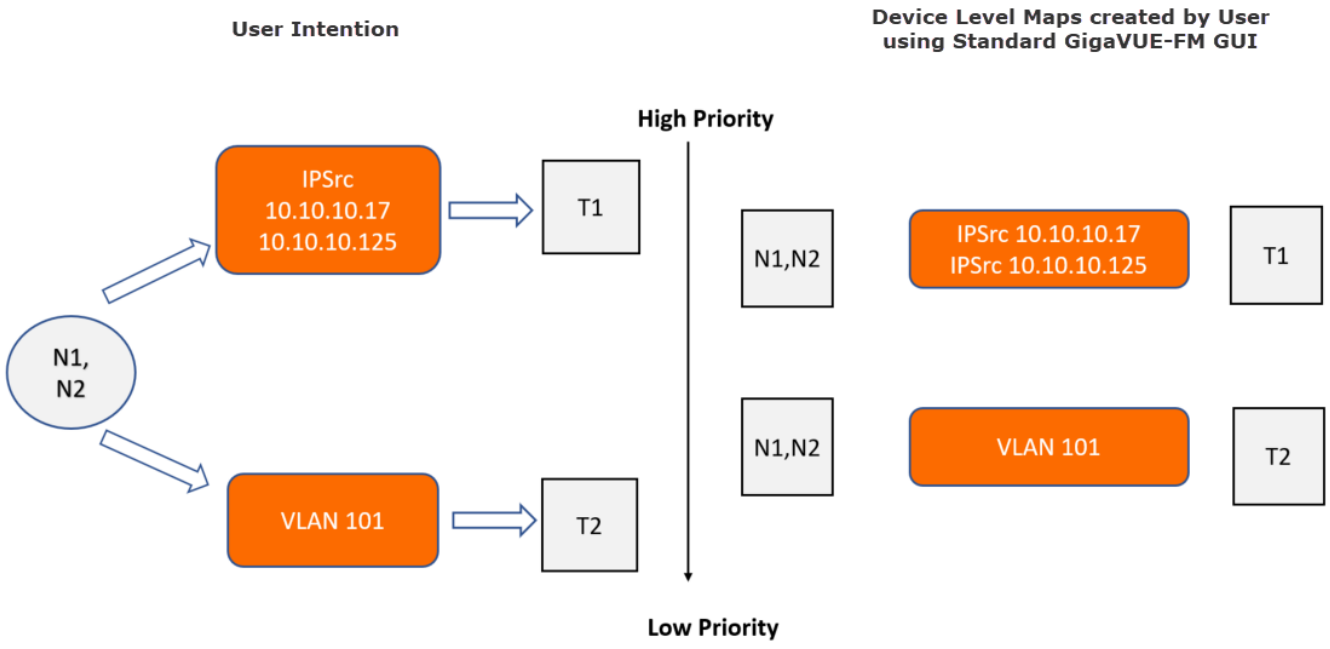

By default, the first map configured always has the highest priority. If a traffic has both IP and VLAN sources, IP traffic is sent to tool T1 but VLAN traffic is not delivered to tool T2.

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.

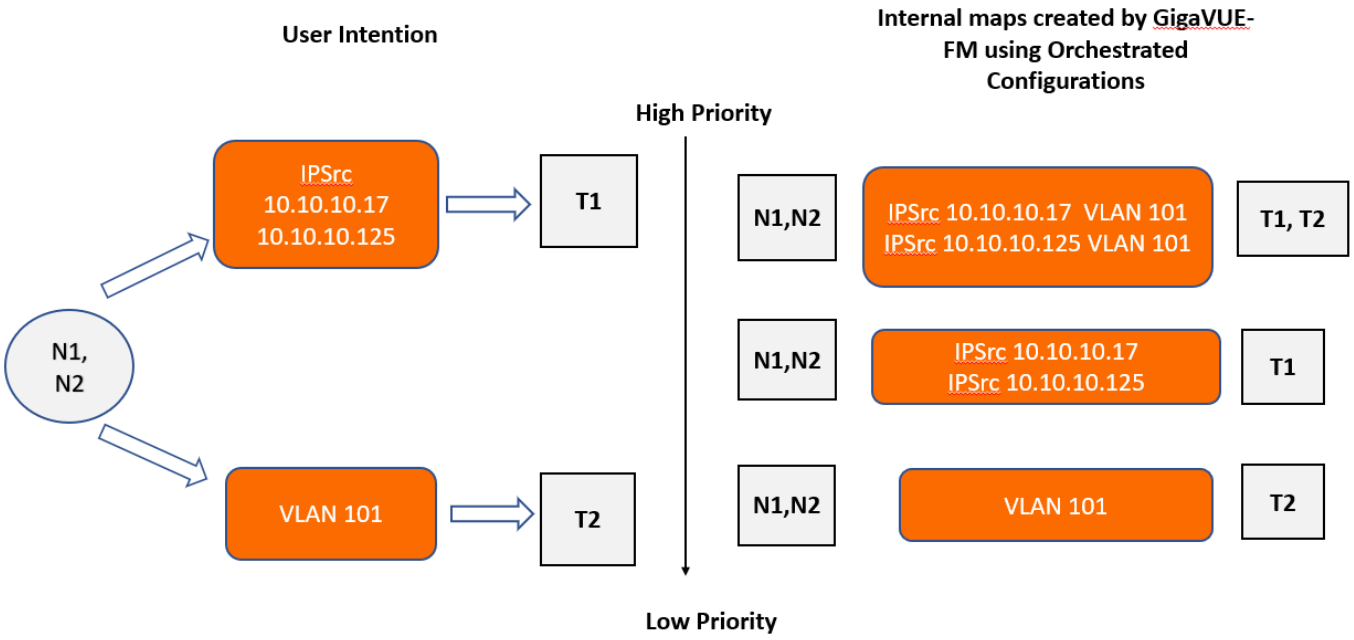

NOTE: You can create priority-based policies by enabling the Prioritize Rules checkbox in the Orchestrated Configuration GUI. Refer to How to Create a Policy for details.

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.

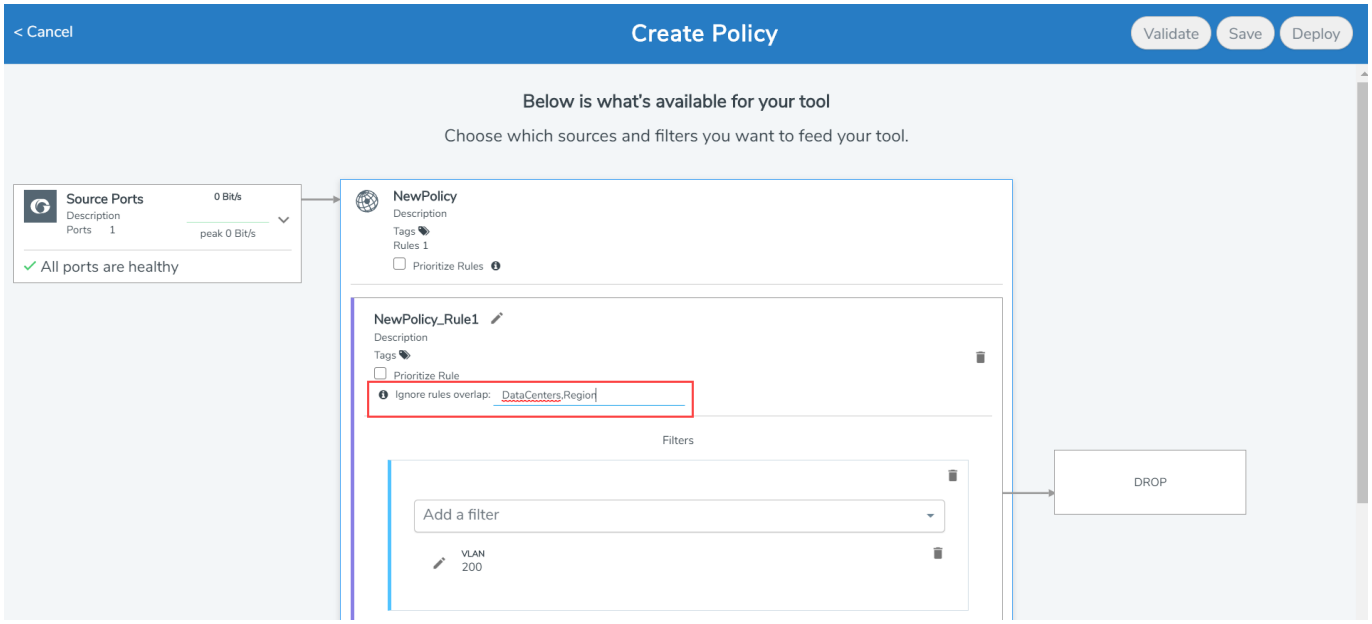

Overlapping Sources in Policies

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.

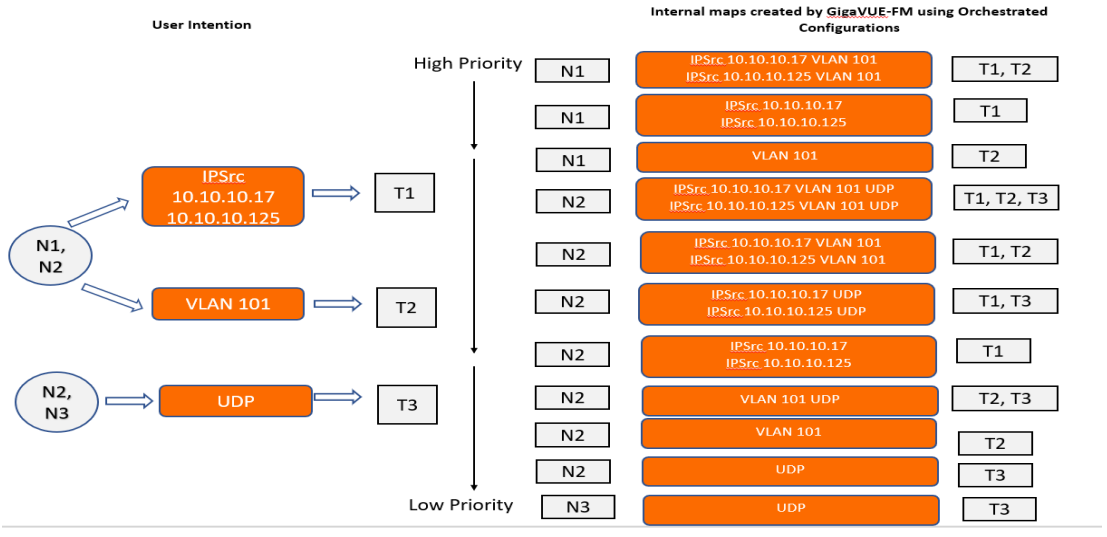

Supported Topologies

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

On the left navigation pane, click on  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. On the left navigation pane, click on  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.

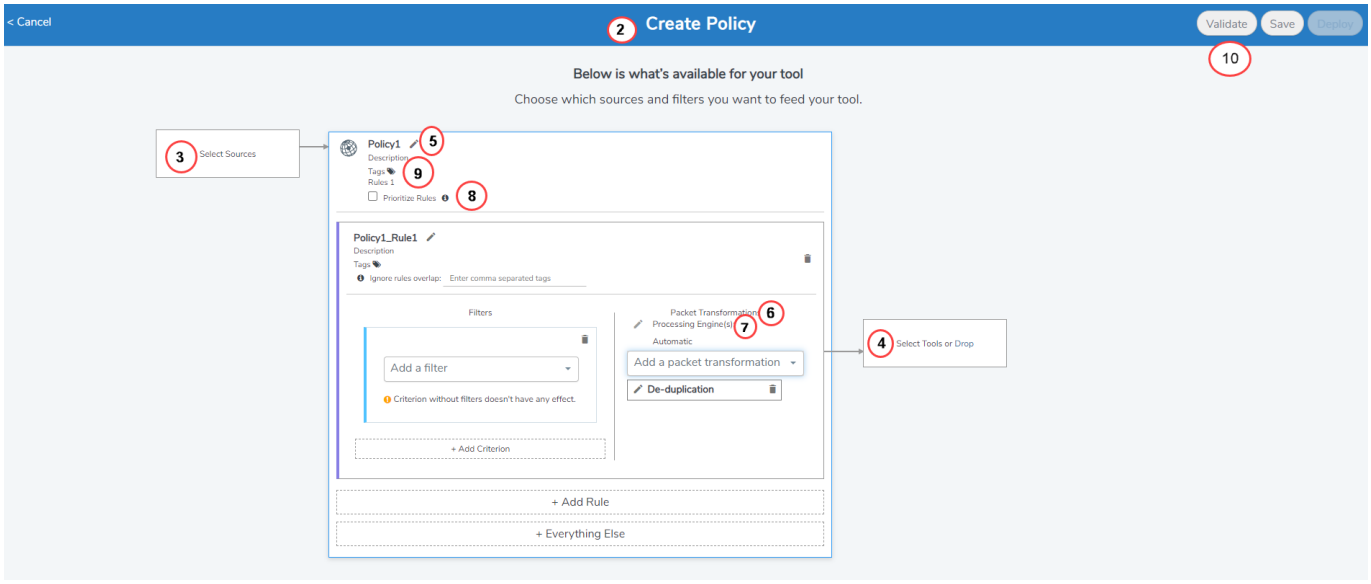

Tagging and RBAC Support

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.

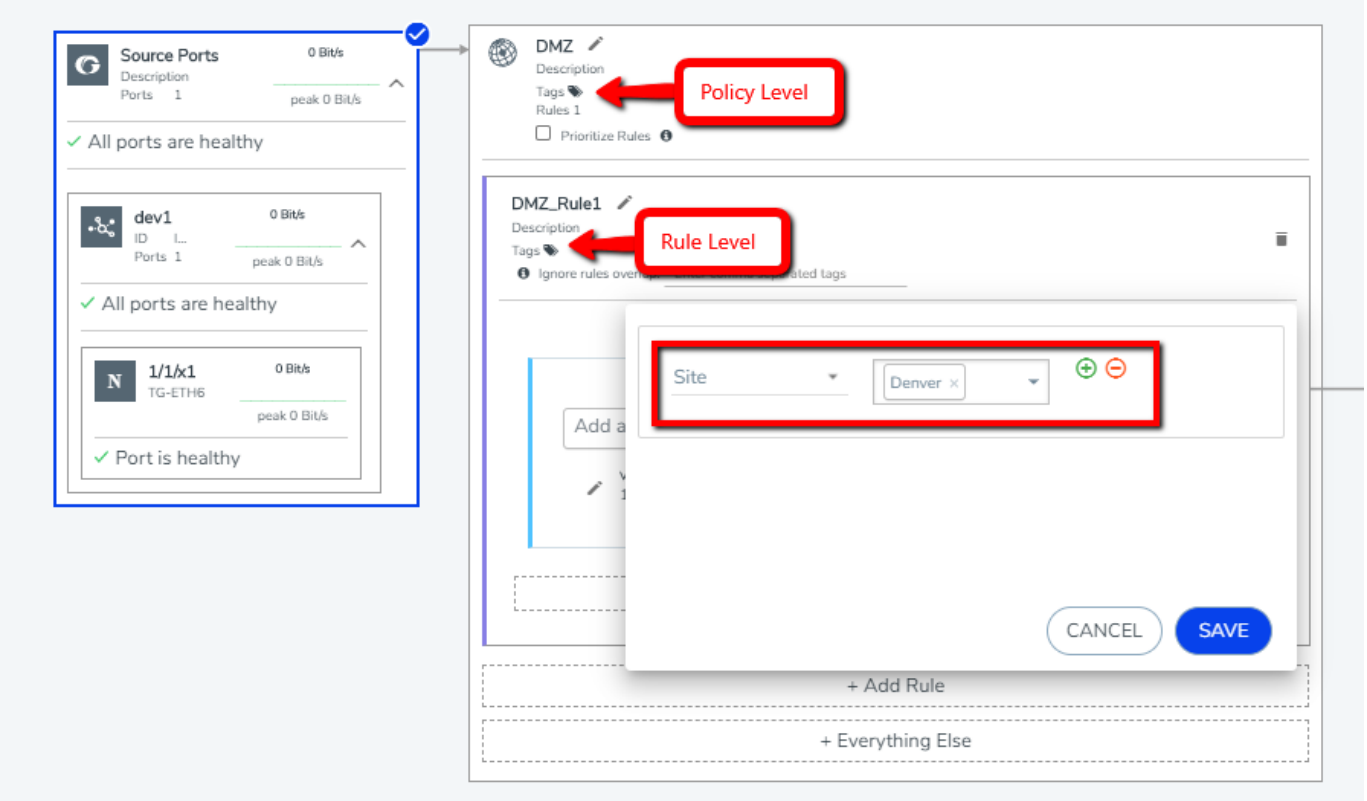

Drop Rules

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.

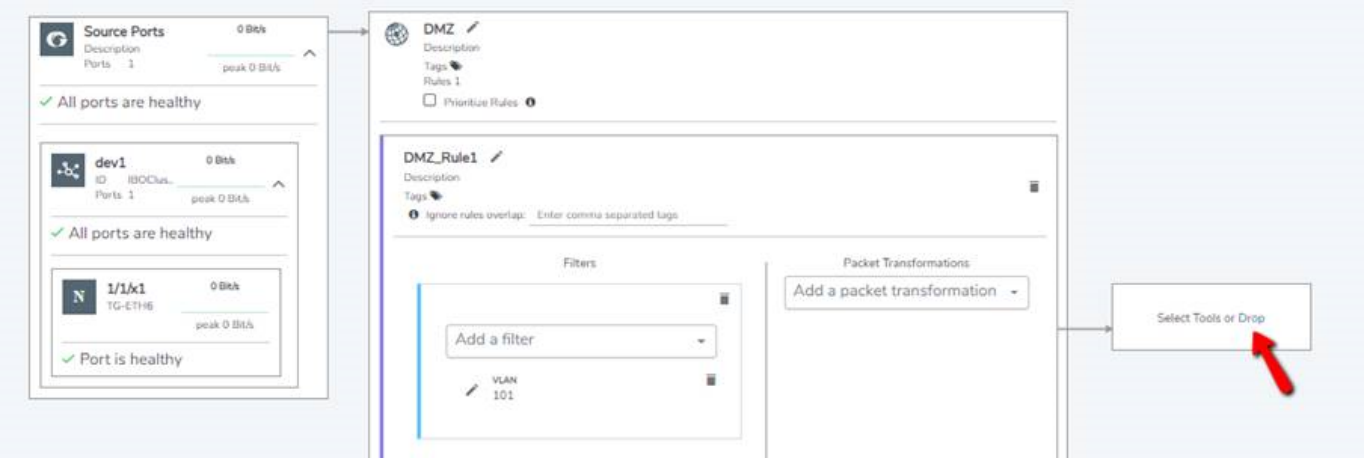

In orchestrated configurations, there are no rule priorities for pass rules. However, the following is the order of priority for prioritized and unprioritized drop rules together with the pass rules:

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.

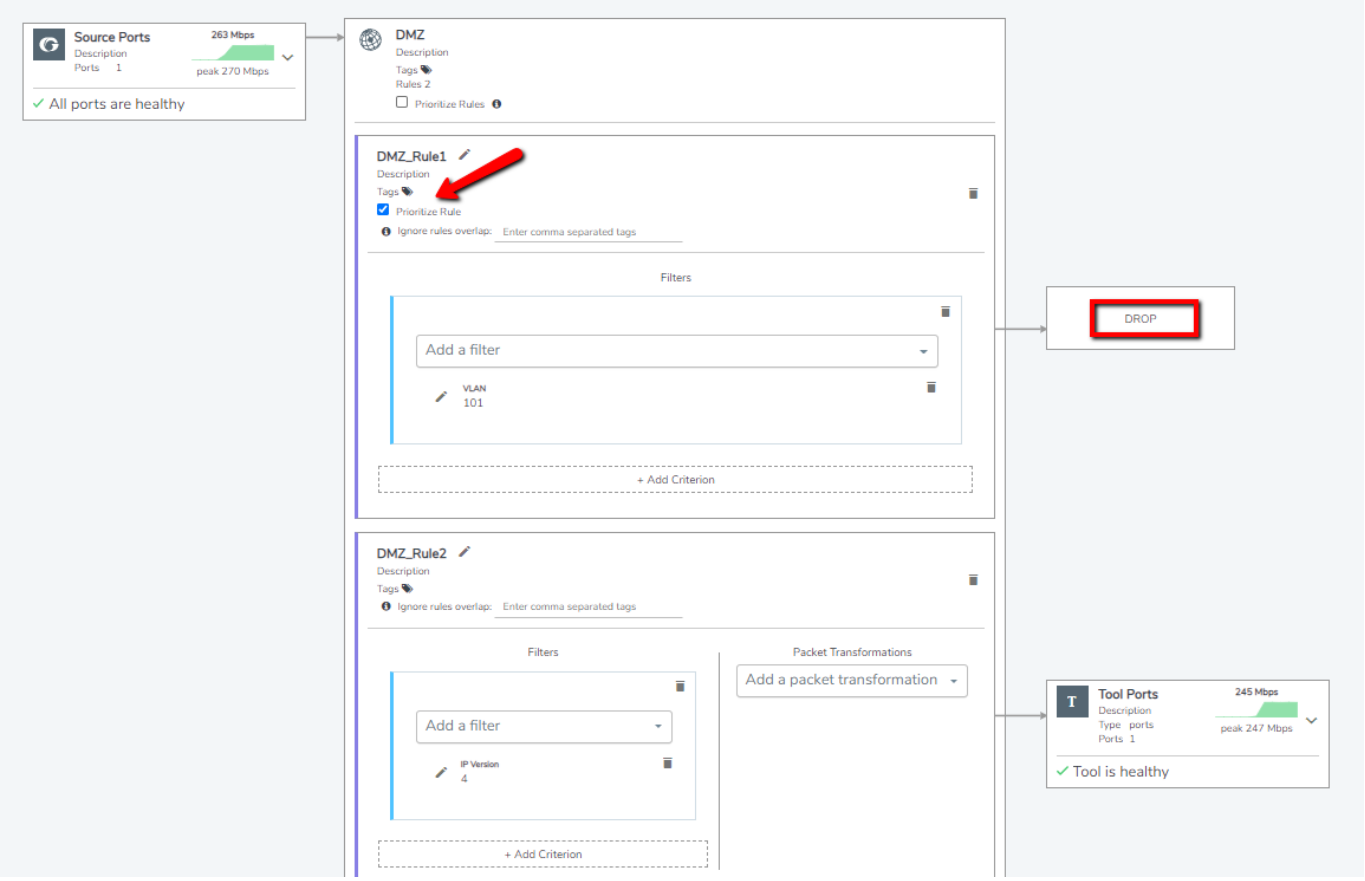

NOTE: To prioritize the drop rules in deployed policies, you must first undeploy the policy.

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.

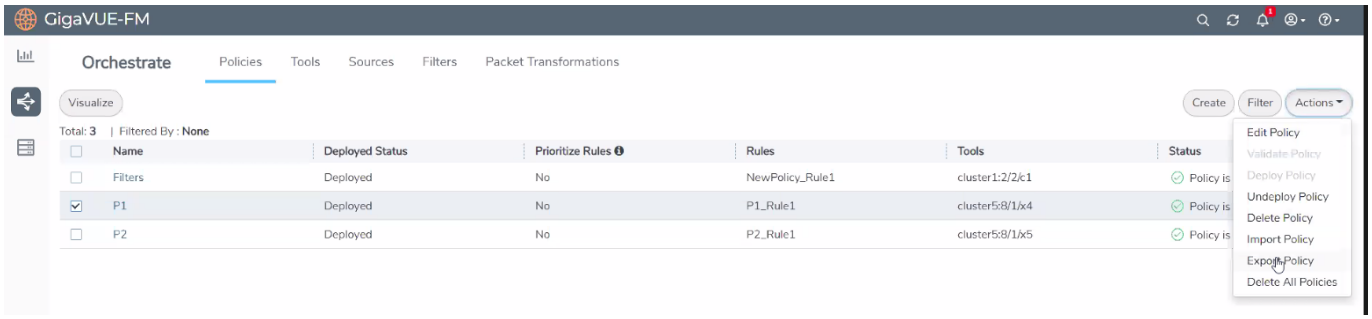

To import a policy:

- 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.

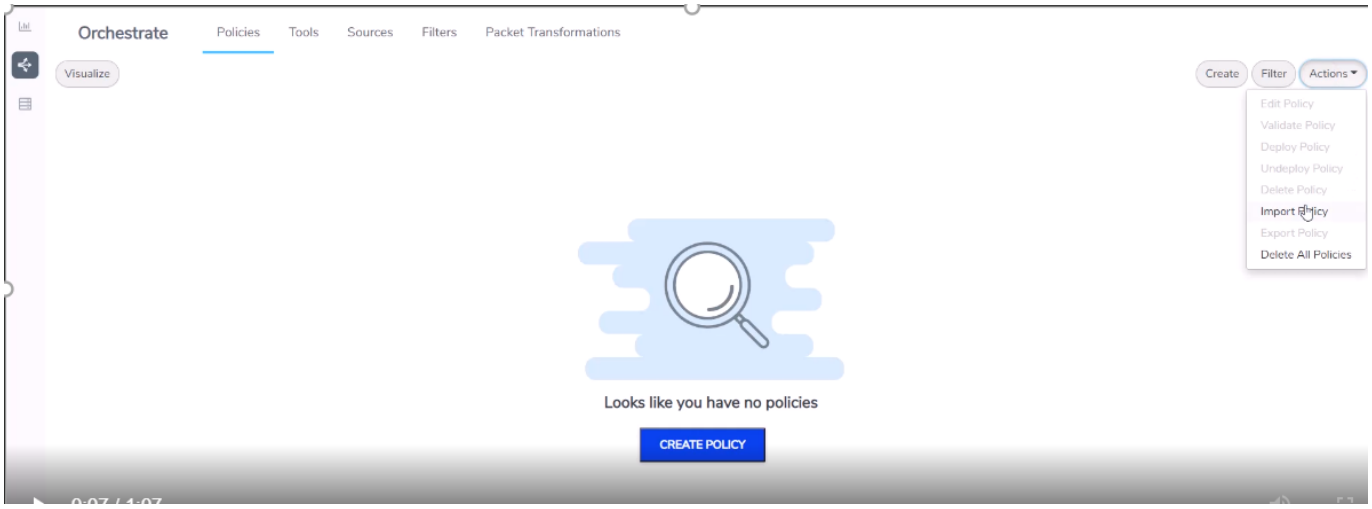

Egress Filters for Additional Filtering Capabilities

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.

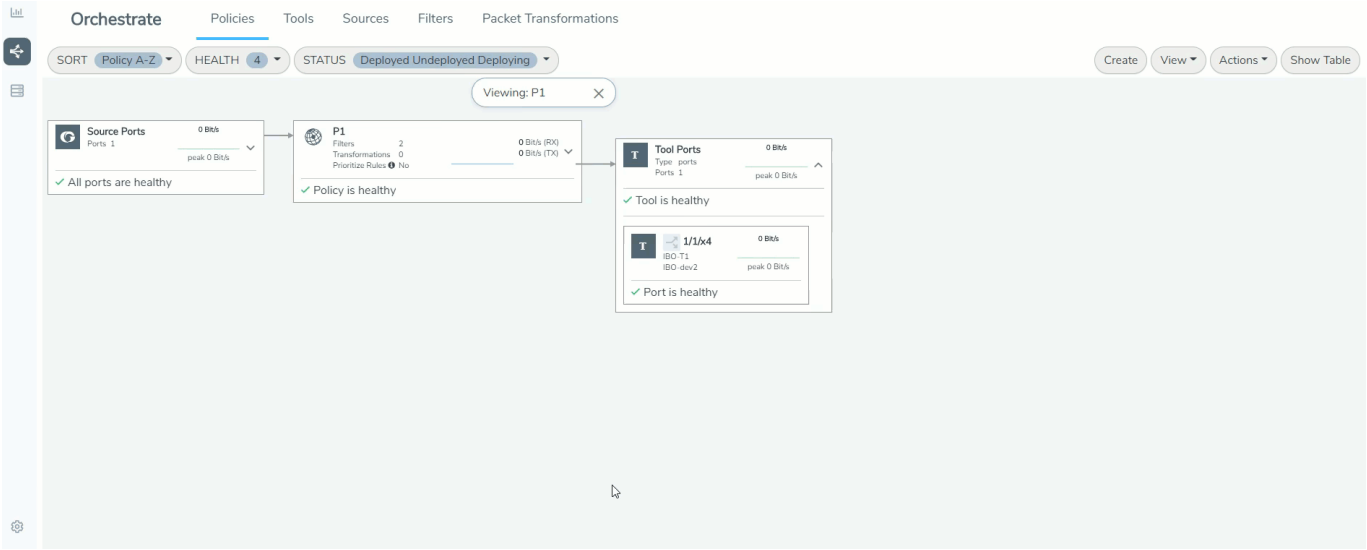

Glossary

- ***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

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 Inline and 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-C05 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	• Configure Inline TLS/SSL Decryption Using GigaVUE-FM • 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

Designed to meet the needs of security operations teams, these GigaSMART operations provide complete network and application visibility to increase the efficiency of security and analytic tools.

Application Intelligence	
Application Intelligence	▪ Application Intelligence ▪ Application Filtering Intelligence ▪ Application Metadata Intelligence ▪ Application Metadata Exporter
Application Session Filtering	▪ GigaSMART Application Session Filtering (ASF) and Buffer ASF

Application Intelligence

This chapter describes about the Application Intelligence solution and its operations. Refer to the following sections for details:

- About Application Intelligence
- How Application Intelligence Works
- Application Intelligence Session

- View the Details of an Application Intelligence Session
- Application Filtering Intelligence
- View the Application Intelligence Dashboard
- About De-duplication
- Application Metadata Intelligence
- View the Health Status of a Solution

Required License: Application Filtering Intelligence, and Application Metadata Intelligence. The Application Visualization capability is included as part of any one of this license. The license for Application Visualization is included in both the licenses; not any one of theese licenses.

Licensing

- 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.

The following table provides details about the features available with the license:

License	Application Visualization	Application Filtering Intelligence	Application Metadata Intelligence	NetFlow/IPFIX
Application Filtering Intelligence	✓	Default Filtering Features available.[2]	×	×
Application Metadata Intelligence	✓	Default Filtering Features available.[1]	✓	✓
NetFlow/IPFIX (NF1)	×	×	✓	✓

License	Application Visualization	Application Filtering Intelligence	Application Metadata Intelligence	NetFlow/IPFIX
(GigaVUE HC SeriesGen3)				**NOTE:** NetFlow/IPFIX supports exporting standard information elements in NetFlow v5/v9, IPFIX and CEF formats.
NetVUE PLUS (GigaVUE HC Series Gen2)	✓	✓	×	×
SecureVUE (GigaVUE HC Series Gen2)	✓	×	×	×
SecureVUE PLUS (GigaVUE HC Series Gen2)	✓	✓	✓	✓
Zero Trust Architecture (GigaVUE HC Series Gen2) & (GigaVUE HC Series Gen3)	✓	✓	✓	✓

[1.] By default, pass-all map is created in Application Filtering Intelligence session with null port. You can also define the physical port, if required.

[2.] By default, pass-all map is created in Application Filtering Intelligence session with null port. You can also define the pass and/or drop application rules and advanced GigaSMART rules, if required.

About Application Intelligence

GigaVUE-FM Application Intelligence provides a comprehensive solution that:

- identifies the applications contributing to the network traffic.
- isolates preferred application-specific traffic and directs it to the appropriate tools.
- exports relevant application metadata for further analytics and analysis.

Application Intelligence Capabilities

Application Intelligence provides the following capabilities for both physical devices and virtual nodes:

- **Application Visualization (earlier known as Application Monitoring)** - Identifies and monitors all applications contributing to the network traffic, and reports on all applications, and their application families and tags and the total bandwidth they consume over a select period. Able to identify over 4000 applications. It displays the traffic statistics in bytes, packet and flows.
- **Application Filtering Intelligence**- Enables traffic filtering by layer 7 applications, which means you can filter out high-volume, low-risk traffic from reaching the tools and direct high-risk network traffic of interest to the right tool at the right time. You can select applications based on the application name, family or tags.
- **Application Metadata Intelligence** - Supports exporting up to 6000 attributes of metadata that provide relevant usage context for over 4000 applications, thus enabling you to rapidly identify indicators of compromise (IoC) for security analytics and forensics tools.
- **Application Metadata Explorer**-Enables ingesting application metadata in cloud-native tools. It converts CEF records from Application Metadata Intelligence (AMI) into JSON or OCSF format. JSON records can be exported over HTTP, HTTPS or Kafka. It can also enrich metadata from user-plane traffic with control-plane metadata for mobile networks or enrich application metadata with key host environment details for cloud workloads. Refer to Application Metadata Exporter to know more.

NOTE: Application Intelligence requires a separate GigaSMART engine or engines in group.

Supported Platforms

Application Intelligence is supported on the following platforms :

- GigaVUE-HC1
- GigaVUE-HC3
- GigaVUE-HC1-Plus
- GigaVUE-HCT
- VMware ESXi
- VMware NSX-T
- GigaVUE Cloud Suite for AWS
- GigaVUE Cloud Suite for Azure
- GigaVUE Cloud Suite for Third Party Orchestration
- GigaVUE Cloud Suite for OpenStack

Refer to Supported GigaSMART Operations for more details on the devices that support GigaSMART operations.

Application Intelligence Feature Compatibility

The table below shows the compatibility between Application Intelligence features and devices.

Application Intelligence Features	GigaVUE-HC1		GigaVUE-HC3		GigaVUE-HC1 Plus		GigaVUE-HCT
	Gen 2 On-board	Gen 3 SMT-HC1-S	Gen 2 (C05)	Gen 3 (C08)	Rear Gen3 SMT-HC1A-R	Front Gen 3 SMT-HC1-S	Gen 3 SMT-HC1-S
User Defined Application Signatures	x	√	x	√	√	√	√
Fast Mode	x	√	x	√	√	√	√
Aggregate Mode	x	√	x	√	√	√	√
Application Family and Tags	x	√	x	√	√	√	√
Max Packet Size	x	√	x	√	√	√	√
DPI Packet Limit	x	√	x	√	√	√	√
AMI Cache Monitoring	x	√	x	√	√	√	√

User Defined Application

This feature allows you to identify unclassified TCP, UDP, HTTP, and HTTPS applications, and extract their application name and the lower layer protocol attributes.

To configure User Defined Application signatures :

Step Number	Task	Refer the following
1	Create rules under User Defined Application Section	Create rules under User Defined Application
2	Configure Application Intelligence Session	For Physical: Application Intelligence Session
3	Monitor User Defined Application	View the Application Intelligence Dashboard

Create Rules under User Defined Application

1. Click **Inventory**.

2. Click **User Defined Applications** to create rules based on a set of **Supported Protocols and Attributes**. For information on **Supported protocols and Attributes** refer **User Defined Application** topic. This helps the physical or virtual node to classify the traffic based on the protocols and attributes selected in the created rule.

3. Click **New** in the **User Defined Applications** screen to create a new rule.

4. Enter **Application Name**.

5. Enter **Priority**. The value must be between 1 and 120.

Note: The least value will have the highest priority.

6. In the created rule:

   a. Choose the **Protocol** from the list of protocols.

> **NOTE:** The **HTTP2** protocol is not listed under the list of protocols. You need to select **HTTP** for configuring the signatures for HTTP2 applications.

   b. Choose the **Attributes** from the list of attributes.

   c. Choose the **Values** from the list of values.

7. Click **Apply**. The rule is now created. For information on the limitations for creating rules refer Configuration Limitations section.

8. Click the application listed under the **Applications** column.

9. Click the **Rule** tab.

10. Select a rule to view its protocol details.

Supported Protocols and Attributes

The DPI engine will match the rules defined based on the following protocols and attributes within the first 500 bytes of a packet payload.

For supported Regex patterns, refer Supported RegExp Syntax

Protocol	Attributes	Attribute Labels	Description	Direction	Supported Data Type	Example Value
http	cts-uri	Request URI	Partially Normalized URL (path + request)	Client to Server Only	REGEXP	\/fupload\/(create_file|new_slice|upload_slice)\?.*upload_token=.*
	cts-server	Server Name	Web Server Name from URI or Host	Client to Server Only	REGEXP	(.*\.)?gigamon\.com

	mime_type	MIME Type	Content type of Request or the Web page	Both, Client to Server or Server to Client	REGEXP	http
	cts-user_agent	User Agent	Software / Browser used for request	Client to Server Only	REGEXP	mozilla
	cts-referer	Referer URI	Source address where client got the URI	Client to Server Only	REGEXP	http:\/\/gigamon.com\/
	stc-server_agent	Server Agent	Software used for the server	Server to Client Only	REGEXP	NWS_TCloud_PX
	stc-location	Redirect Location	Destination address where the client is redirected to	Server to Client Only	REGEXP	.*\/football\/.*
	cts-cookie	Cookie (Raw)	Raw value of the HTTP Cookie header line	Client to Server Only	REGEXP	.*tEstCoOkie.*
	content	Content	Message body content	Both, Client to Server or	REGEXP	.*GIGAMON.* mindata = 206 Refer Mindata

				Server to Client		
http2	cts-uri	Request URI	Partially Normalized URL (path + request)	Client to Server Only	REGEXP	\/fupload\/(create_file|new_slice|upload_slice)\?.*upload_token=.*
	cts-server	Server Name	Web Server Name from URI or Host	Client to Server Only	REGEXP	(.*\.)?gigamon\.com
	cts-user_agent	User Agent	Software / Browser used for request	Client to Server Only	REGEXP	mozilla
	cts-referer	Referer URI	Source address where client got the URI	Client to Server Only	REGEXP	http:\/\/gigamon.com\/
ssl	common_name	Domain Name	Domain name from Client Hello message or the certificate		REGEXP	(.*\.)?gigamon\.com
	stc-subject_alt_name	Subject Alt Name (s)	List of host names which belong to the same certificate	Server to Client Only	REGEXP	(.*\.)?gigamon\.com

rtmp	cts-page_url	Page URL	URL of the webpage where the audio/video content is streamed	Client to Server Only	REGEXP	http:\/\/www.music.tv\/recorded\/1234567
tcp	stream	Payload Data	Data payload for a packet, excluding the header.		REGEXP	.*GIGAMON.* mindata = 70 Refer Mindata
	port	Server Port	Server (listen) port number		UINT16 RANGE as REGEXP String	80-4350
udp	stream	Payload Data	Data payload for a packet, excluding the header		REGEXP	.*GIGAMON.* mindata = 100 Refer Mindata
	port	Server Port	Server (listen) port number		UINT16 RANGE as REGEXP String	80-4350
sip	user_agent	User Agent	Software used	Both, Client to Server or	REGEXP	GVUE-release 6.2.0

				Server to Client		
icmp	code	Message Code	Code of the ICMP message	Both, Client to Server or Server to Client	UINT8 as REGEXP String	200
	typeval	Message Type	Type of ICMP message	Both, Client to Server or Server to Client	UINT8 as REGEXP String	10
ip	address	Server IP Address	IP address of the server		IPV4 as REGEXP String	62.132.12.30\/24
	dscp	DSCP Value	DSCP from Differentia ted Service (DS) Field in IP header		UINT8 as REGEXP String	33

	resolv_ name	DNS Name	Server's DNS name		REGEXP	gigamon.com
ipv6	address	Server IP Addres s	IP address of the server		IPV6 as REGEXP String	2001:0:9d38:6ab8:307b:16a 4:9c66:5f4 2001:0:9d38::9c66:5f4/64
	dscp	DSCP Value	DSCP from Different ia ted Service (DS) Field in IP header		UINT8 as REGEXP String	43

Mindata

The mindata value is the number of payload bytes to buffer and match a given pattern. You can configure mindata value for HTTP content, TCP stream, and UDP stream. The buffer size is calculated from the start of the payload and the default buffer size is different for each protocol (HTTP - 206, TCP - 67, and UDP - 48.)

For example, for pattern ".*TEST.*" that may be present within the first 67 bytes of TCP payload, you can specify the mindata value as 4 (which is the length of the input string) or as 67 (which is the default buffer size of TCP payload). In case, the pattern is present in between 65 to 68 bytes of the payload and the mindata is specified as 4 or 67, it will not match. For this case, you must specify the mindata value as 68.

Supported RegExp Syntax

Pattern	Description
.	Matches any symbol
*	Searches for 0 or more occurrences of the symbol or character set that precedes it
+	Searches for 1 or more occurrences of the symbol or character set that precedes it
?	Searches for 0 or 1 occurrence of the symbol or character set that precedes it
( )	Groups a series of expressions together

[ ]	Matches any value included within the bracket at its current position Example: [Dd]ay matches Day and day
| [<start>-<end>]	Separates values contained in ( ). Searches for any one of the values that it separates. Example: The following expression matches dog or cat: (dog | cat). Matches any value contained within the defined range (a hyphen indicates the range). You can mix character class and a hexadecimal range Example: [AaBbCcDdEeFf0-9]
\0 <octal_ number>	Matches for a direct binary with octal input
\x<hexadecimal-number>\x	Matches for a direct binary with hexadecimal input
\[<character-set>\]	Matches a character set while ignoring case. WARNING: Not performance friendly

Limitations

- The maximum number of user defined application that can be configured is 120 per GigaVUE-FM. These applications can be spread across one or more application intelligence sessions.
- The maximum number of rules that can be created per application is 8.
- The maximum number of protocols that can be configured per rule is 3.

How Application Intelligence Works

Figure 1Application Intelligence Solution illustrates how the Application Intelligence solution works.

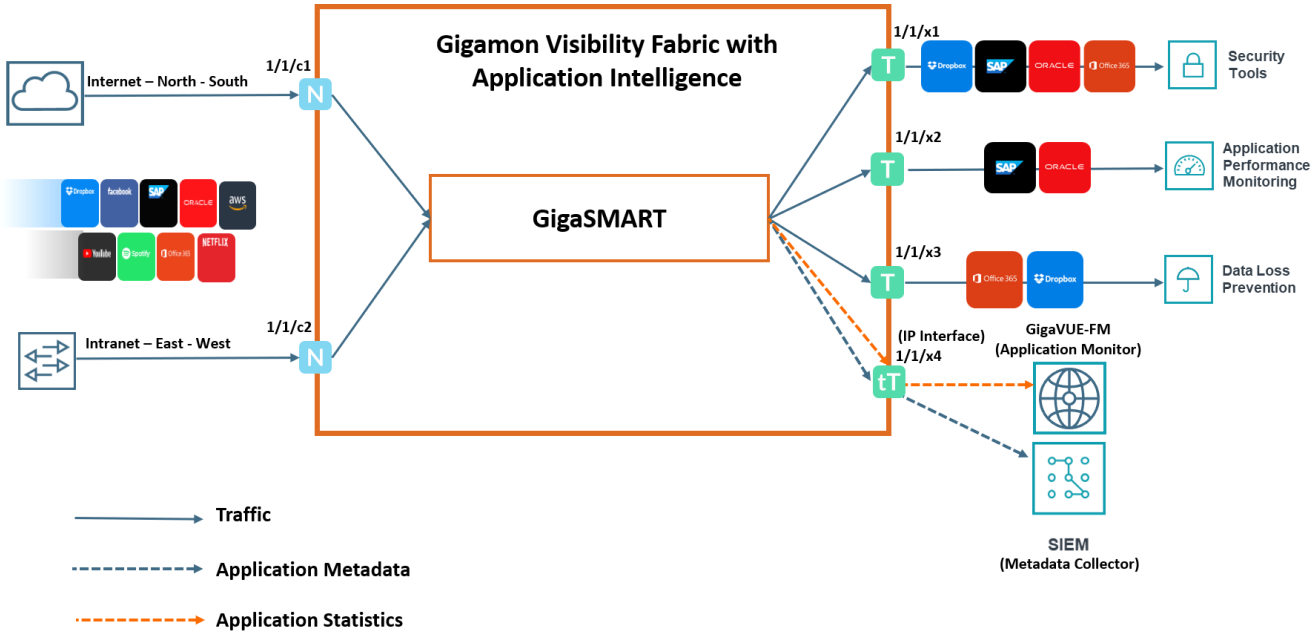

Figure 1 *Application Intelligence Solution*

The Gigamon device that is configured with the Application Intelligence capability receives the traffic through the network ports. Application Visualization (earlier it was known as Application Monitoring) exports applications related information to GigaVUE-FM, which renders the information on the Application Intelligence Dashboard. Based on the pass or drop rules configured in the maps, Application Filtering Intelligence lets relevant applications be guided to the tools. In this example, you can see that except some of the audio or video streaming applications such as Spotify®, YouTube, and Netflix, other applications are sent to the security tool. This is because the audio or video streaming applications are high-volume, low-risk traffic. The threat detection tool need not inspect such traffic. Hence, these applications are dropped based on the configured drop rules.

Application Metadata Intelligence allows you to export metadata from applications that are detected in the network traffic. The records can be exported to a collector either in IPFIX or CEF format through the IP interface or the management interface. You can also use the application metadata attributes for purposes other than security, such as to determine the network or application health, to track the long-lived sessions seen in the network, and so on.

GigaVUE-FM Application Protobook

Protobook provides a complete list of supported protocols and their attributes of metadata. These protocols can also be viewed as groups by Tags, Family and Classification method.

You can access the Application Protobook from the GigaVUE-FM in the following ways:

- On the left navigation pane, click then click the "App Protobook" link under the **Support** section OR
- Go to **admin > App Protobook**

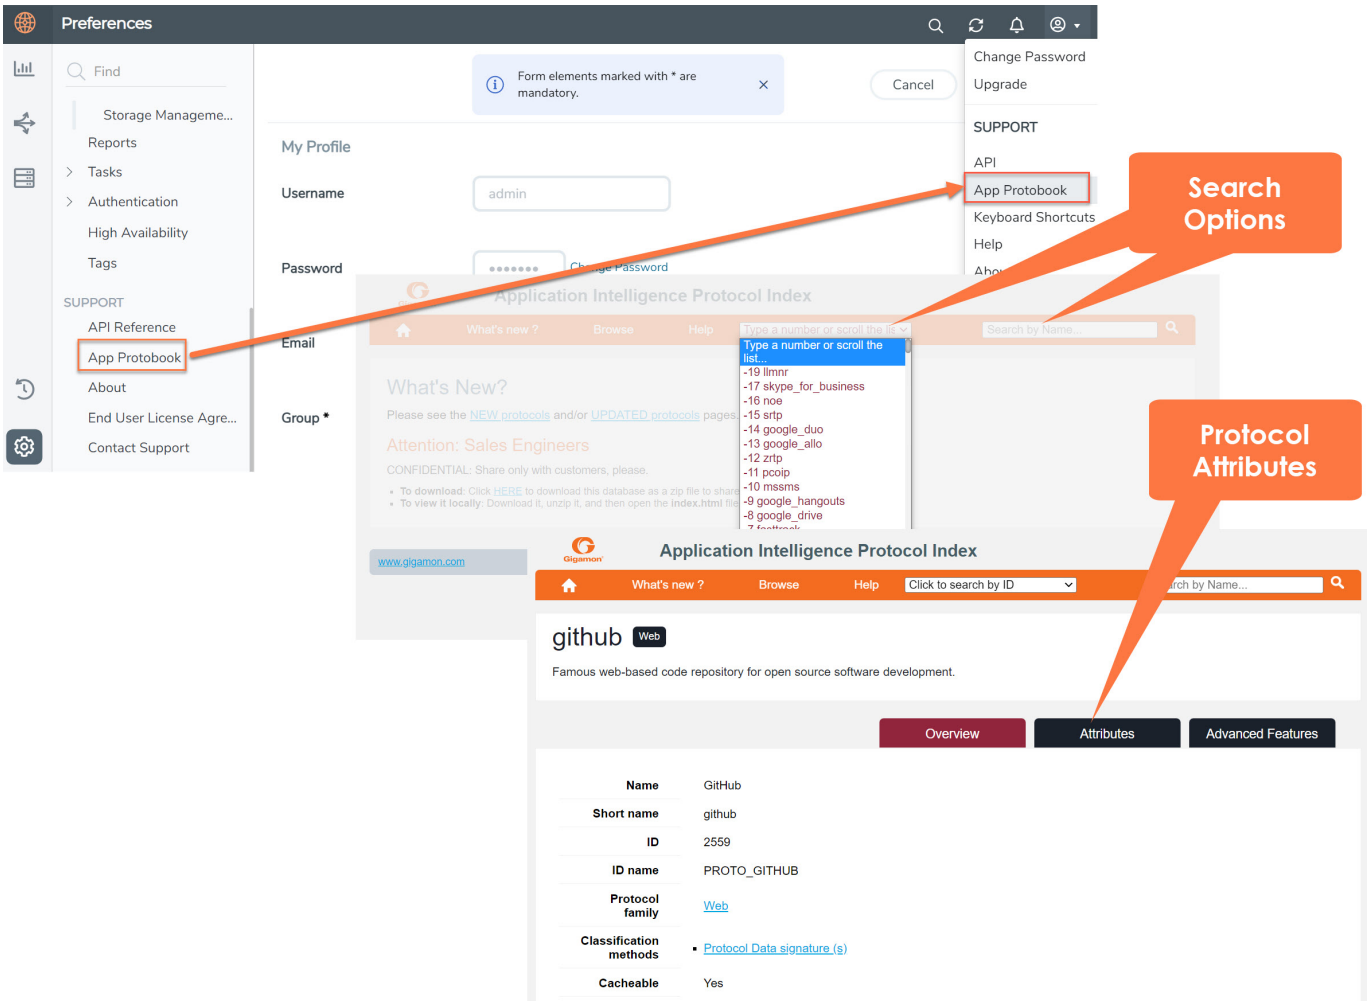

Application Intelligence—Rules and Notes

Keep in mind the following rules and notes when working with the Application Intelligence solution:

- When both the AMI and AFI licenses are installed on Gen 2 cards, the non-5-tuple packets will be dropped, and the packets will not pass to the AMI application for attribute extraction.
- Whenever you perform a backup and restore operation, you must create a backup of both the device and GigaVUE-FM and then restore the backed-up data on both as well.

- The best results for Application Metadata Intelligence can be obtained only if the traffic is bi-directional.
- In Gen 2 cards, if you want to filter traffic based on IP address, ensure you have configured both IP source and destination in GS rules when bidirectional Application filtering settings are enabled.
- It takes 5 minutes for the Application Filtering Intelligence and Application Metadata Intelligence statistics to be populated in the GigaVUE-FM GUI.

Application Intelligence Session

Required License: The Application Monitoring(Application Visualization) capability is included as part of any one of these licenses. Application Filtering Intelligence, and Application Metadata Intelligence.

Supported Devices: GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus, GigaVUE-HCT.

Refer to Supported GigaSMART Operations for more details on the devices that support GigaSMART operations.

Application Visualization (formerly known as Application Monitoring) collects application statistics and sends this information to GigaVUE-FM, which acts as an application monitor. The monitoring reports are sent to GigaVUE-FM through destination port 2056. The application statistics are displayed as a set of monitoring reports, providing easy-to-read graphical representations of application usage data. This gives you better insight and control over how your network is utilized and which applications consume the most resources. To enable Application Monitoring, you need to create the necessary application intelligence sessions on the nodes managed by GigaVUE-FM.

Application Visualization uses various classification methods to identify applications. The application statistics exported to GigaVUE-FM may take up to seven minutes to populate in the following dashboards:

- Top 10 Applications

- Total Traffic

- Total Applications

Users can view the statistics based on Bits, Packets, and Flows. Historical data can be accessed for up to seven days, and users can export the application visualization records using the IP interface. However, it's not recommended to use the IP interface, as disruptions to the data plane may prevent GigaVUE-FM from collecting and displaying statistics.

NOTE: Users may want to filter traffic upfront based on VLAN, subnet, and host IP address. Filtering traffic upfront can also reduce the load on the DPI engine. Application Intelligence supports configuring L2-L4-based rules to filter such traffic.

Create an Application Intelligence Session in Physical Environment

To create an Application Intelligence Session:

1. On the left navigation pane, select **Traffic > Solutions > Application Intelligence**.
2. Click **Create New**. The **Create Application Intelligence Session** page appears.

NOTE: If the Create button is disabled, check whether a valid license for Application Metadata Intelligence or Application Filtering Intelligence is available.

3. In the **Basic Info** section complete the following:

- Enter the name and description (optional).
- Select **Physical** in the Environment field.
- Select the node from the list of nodes.

4. In the **Configurations** section, view the following:
   a. The export interval is fixed at five minutes, at which the Application Intelligence session will generate the reports for application monitoring.

   NOTE: When you upgrade to 6.4, the prior configured export interval will be updated to five minutes.

   NOTE: The Monitoring option is enabled by default. To increase the session performance it is recommended to disable this option.

   b. Select a GigaSMART Group. You can also choose to create a new GigaSMART Group.
      - Provide a name in the **Alias** field.
      - Select a port or multiple ports from the **Port List**.
      - Click **Save**.

      For more information about creating a new GigaSMART group, refer to the Configure a GigaSMART Group section.

   c. Select the required interface. By default, Management Interface is enabled. To export the data through IP interface, uncheck the Management Interface check box. You can view the **IP Interface** field only after you select a GigaSMART Group.

d. Select an IP interface that is used to create a dedicated channel to communicate all application statistics to GigaVUE-FM. You can also choose to create a new IP interface. For more information, refer to the IP Interfaces section.

e. Enable the **Fast Mode** option for performance (less CPU cycles and less memory utilization) improvement. When the **Fast Mode** is enabled, some or all of the attributes of the applications will be disabled. You can view the list of attributes/applications available in the fast mode by navigating to the app editor under AMI feature in theGigaVUE-FM. If all the attributes of the application are disabled then the application itself is disabled. Refer Fast Mode for more information.

NOTE: The Fast Mode feature is supported in GigaVUE HC Series Gen 3 GigaSMART and GigaVUE V Series modules. The **Fast Mode** option can be enabled or disabled only when creating a new application intelligence session.

5. In the **User Defined Applications** section, select the template from the inventory where it was created.
6. If you are unable to view the required port in the **Port** field, perform these steps:
   - Click **Port Editor.** Select the **Type** as **Tool** from the drop-down list for the required **Port Id**. Select **OK**.

     The selected Port appears in the list.
   - Select the **Type** as:
   - IPv4 - to allow the traffic in IPv4 interface.
   - IPv6 - to allow the traffic in IPv6 interface.
   - Provide the **IP Address**, **IP Mask, Gateway**, and **MTU**. Provide the IP address corresponding to the IP interface selected.
   - Click **Save**.
7. In the **DestinationSettings**, enter the destination IP address. The version of IP address in the Destination field should be same as the version of IP address defined in the IP-interface (applicable only when IP interface is selected). By default, the IP address of the GigaVUE-FM interface is displayed.
8. In the **Source Traffic** section, select a source port that require application monitoring in the **Source ports** field. Source port can be a single port, multiple ports, and port groups.

   NOTE: Ports already used as source ports in the intent-based orchestrated solution will not be listed in the drop-down.
9. Click **Save**. The session created is added in the list view.
10. In the created session, click **Edit** to perform operations related to **Application Filtering**, **De-duplication** and **Application Metadata**.

11. Configure the rules for filtering the required traffic in the **L2-L4 Rules** fields. To configure a rule:
    a. Click **Select Conditions**. Select the required parameters from the drop-down list.
    b. Select the value for the parameters from the drop-down.
    c. Select the required options:
        - Pass or Drop - Based on the parameter selected in the Conditions fields, the traffic that matches the conditions will either be passed or dropped.
        - Bidirectional - Allows the traffic in both directions of the flow.

    NOTE: Click "**+**" to create multiple rules for filtering the required traffic, and click "**+ New Source Traffic**" to create multiple sources with filtering options.

The total applications participating in the network traffic are displayed in the Application Intelligence Dashboard. For more information about the dashboard, refer to the View the Application Intelligence Dashboard.

If the session configuration is unsuccessful, troubleshoot the error notified (refer to View the Health Status of a Solution). Click the **Reapply all pending solutions** button in the dashboard to redeploy the configuration.

NOTE: GigaVUE-FM takes few minutes to display the application statistics.

You can also filter the traffic based on the applications. For more information, see Application Filtering Intelligence.

Upgrading the Protocol Signature

You can upgrade the protocol signature by upgrading the image file on the GigaSMART card, and by uploading the GigaSMART card image to GigaVUE-FM from the Internal Server or External Server. To select the image, follow the steps:

- Internal Server—Upload the GigaSMART card image to GigaVUE-FM from the internal server and select the image that you need to upgrade from the selected GigaSMART card.
- External Server—Provide the location of the image in the external server that you need to upgrade from the GigaSMART card.

For more information on upgrading the image, refer to the following topics in the GigaVUE-OS Upgrade guide:

- Upgrade from an External Image Server
- Upgrade with GigaVUE-FM as the Image Server

Disable Application Visibility in Application Intelligence Session

To disable Application Visibility, do the following:

1. Edit the Application Intelligence Session.
2. Disable the checkbox next to **Monitoring Enabled**.
3. Select **Save**.

After disabling the monitoring session, the Application Intelligence dashboard does not receive the updates and a message " **Monitoring has been disabled. Please enable monitoring to get latest data** " is shown in red at the top of the dashboard.

Fast Mode

The DPI engine provides a performance optimization functionality called “fast mode”. In this mode the performance is increased by the use of light parsers for processing the HTTP and DNS traffic. This increases the performance of HTTP and DNS based protocols and reduces the memory utilization. This is supported only in GigaVUE HC Series.

Limitations of Fast Mode

- This affects the classification of applications over HTTP. Only limited applications based on HTTP can be classified.
- This affects the attribute extraction. For example, only some attributes are extracted for the HTTP traffic and no attributes are extracted for the DNS traffic. When the fast mode is enabled, the GigaVUE-FM automatically displays only the attributes that are supported.

NOTE: In the protobook, you can check if a specific application supports the fast mode or not, by navigating to the **Attributes** tab of the specific application. If the **Basic DPI Support** field is yes, then the application supports the fast mode option.

Application Intelligence can also be configured for virtual environment, refer to section in *GigaVUE V Series Applications Guide* for more detailed information.

View the Details of an Application Intelligence Session

To view the details and the statistics of a session, do the following steps:

1. Select the session from the **Application Intelligence Sessions** pane for which you need to view the sessions details, health status and statistics, and click the ellipsis as shown:

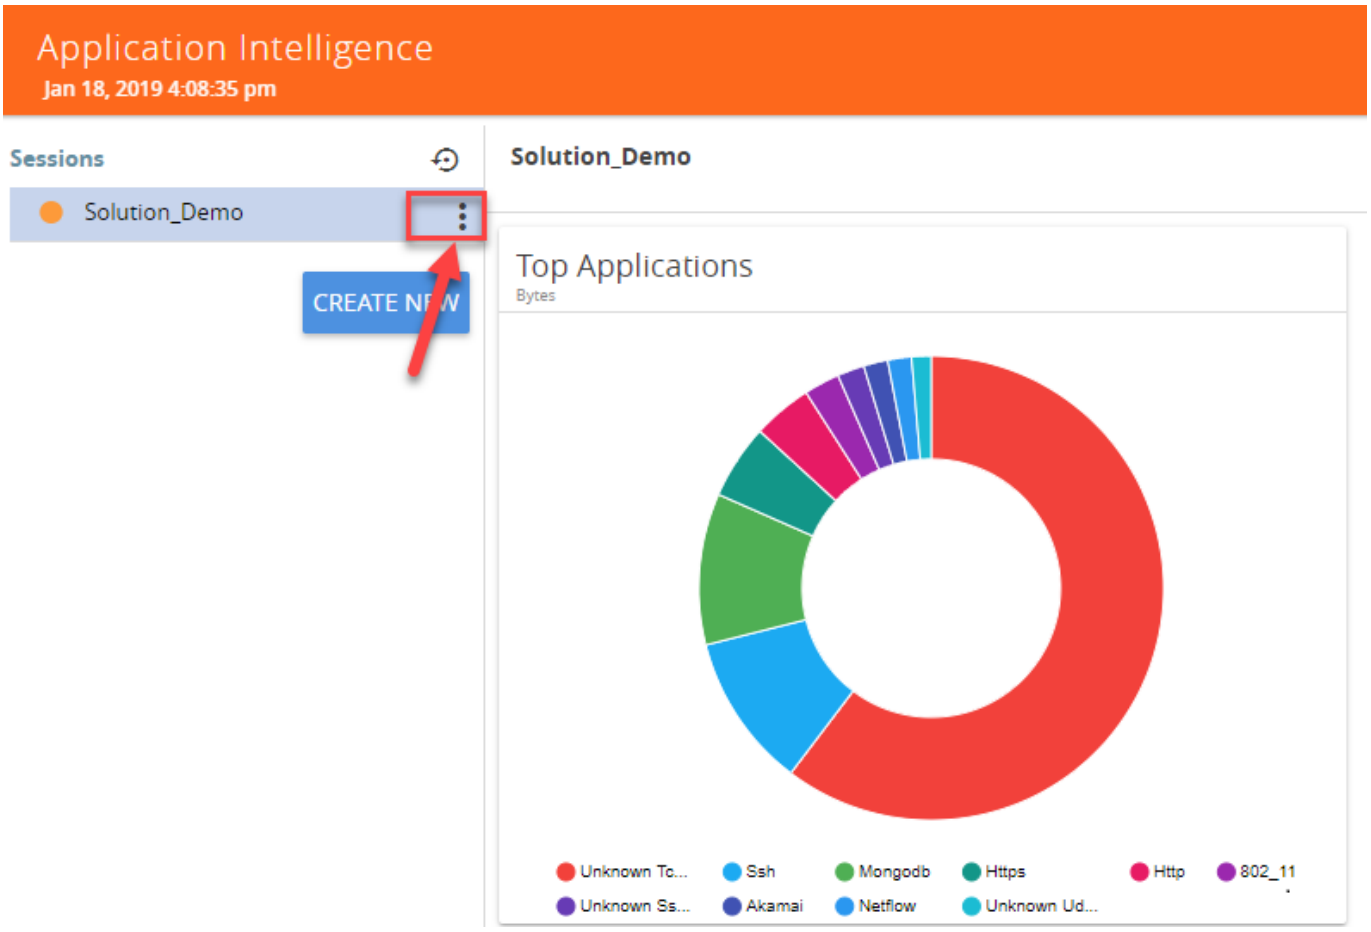

2. Select **View Details** from the drop-down list.

You can view the Monitoring Details, Destination Settings, Source Traffic, and Application Filtering-Destination Traffic, and Application Metadata-Destination Traffic in the Details page.

NOTE: You can edit the Source Traffic, Application Filtering-Destination traffic, and Application Metadata-Destination Traffic from the view details window.

3. Click **Troubleshoot** to view the current statistics and health status of components associated with the solution.

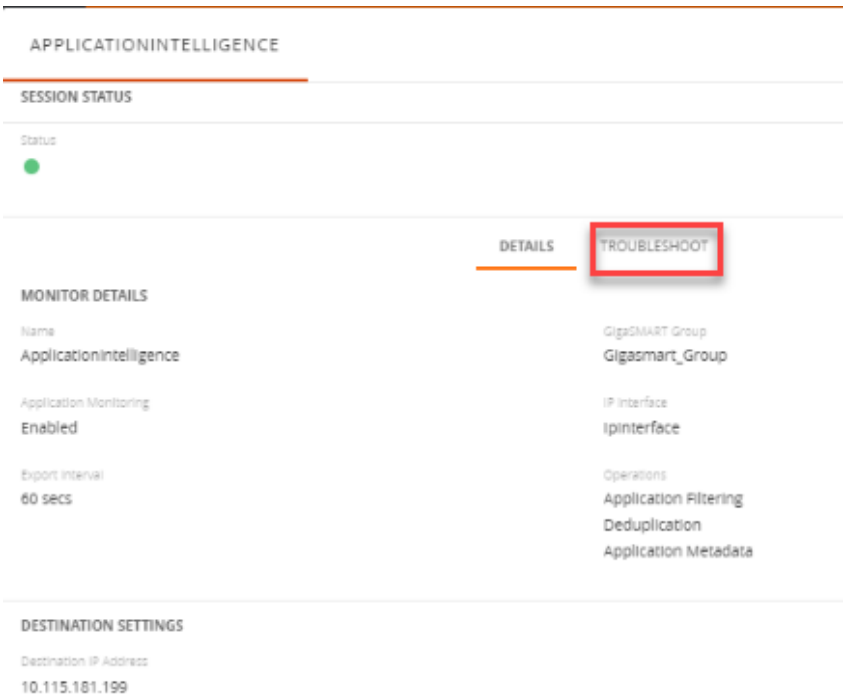

The details of the components that can be viewed for a solution are shown in the following table:

Components	Details of Components
Source Traffic	◦ Network Ports ◦ First Level Maps ◦ L2-L4 Rules
GigaSMART	◦ GigaSMART Port ◦ GigaSMART Group ◦ Virtual Port
Application Monitoring	◦ IP interface ◦ Exporter

Components	Details of Components
Application Filtering - Destination Traffic	◦ Second Level Map ◦ Applications and Advanced Rules ◦ Tool Ports ◦ GigaSMART Operation (GSOP)
Application Filtering	◦ Application Session Filtering
Application Metadata	◦ IP Interface ◦ Cache ◦ Metadata Tools

The troubleshooting page has a flow diagram representing the components associated to the solution. You can also click on the blocks in the flow diagrams to view the details of the corresponding components.

To learn more about the color indication and the health status of a solution refer to View the Health Status of a Solution.

Application Filtering Intelligence

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC1 Gen 3, GigaVUE-HC3 Gen 2, GigaVUE-HC3 Gen 3, GigaVUE-HC1-Plus, GigaVUE-HCT.

Refer to Supported GigaSMART Operations for more details on the devices that support GigaSMART operations.

Application Filtering Intelligence (AFI) functionality on GigaSMART allows filtering of traffic based on the application (such as YouTube, NetFlix, Sophos, or Facebook) or application family (such as antivirus, web, erp, or instant-messaging) or application tag (such as Multimedia Streaming, Gaming, Cryptocurrency).

Organizations can utilize AFI to effectively filter and route crucial applications to one or multiple tools or to a Null Port. This empowers organizations to oversee each flow independently. They have the flexibility to filter in only the packet headers by adjusting the Packet Count and discarding the rest.

NOTE: Application Filtering Intelligence(AFI) and Application Metadata Intelligence (AMI) licenses are available for individual purchase or as a bundle on GigaVUE HC Series. When obtained together, all applications passed by AFI are directed to packet

monitoring tools and AMI. In certain scenarios, users may prefer to export NetFlow/IPFIX or application metadata for the filtered applications instead of monitoring raw packets. In such cases, users can opt for the Null Port (dummy tool port) as the tool destination for AFI. Traffic sent to a Null Port is internally discarded.

In diverse environments, organizations may need to monitor different types of traffic separately and block specific applications from being monitored. AFI allows for the establishment of distinct maps to either forward or block applications to the relevant tools, and these maps are processed using logical OR operation.

You can configure up to five maps with priorities. Higher priority maps take precedence over lower ones. It's best to prioritize maps with specific rules. Advanced rules can be set within each map to optimize traffic further, using a logical AND operation for multiple rules.

NOTE: The **Application Families and Tag** feature is supported only on Gen 3 GigaSMART module.

Application Filtering Intelligence (AFI) supports filtering for 4300 applications. When filtering encrypted applications, GigaVUE-FM will add eight applications (SSH, SSL, IPSEC, ISAKMP, TOR, TOR2WEB etc.) by default as a filtering criterion. These application works as base of top-level applications like HTTPS. Application Filtering Intelligence will filter all valid applications which have either of the eight applications in their protocol path.

To upgrade the protocol signature, refer to Upgrading the Protocol Signature

Elephant Flows in Application Filtering Intelligence

An elephant data flow is a single session (TCP Session) with a relatively long-running network connection that consumes a large or disproportionate amount of bandwidth, buffers, and queues. Because of this nature, elephant flows can cause packet drops in other traffic and significantly increase the mean-time-to-completion (mttc) of smaller flows (mouse flows)[1].

Elephant flows are considered to affect the traffic in the following ways:

- Disproportionately affects mouse data flows mean-time-to-completion (mttc).
- Causes significant issues to tools, detecting problems with applications and next-generation firewall (NGFW), as it causes high CPU spikes and bandwidth consumption.

[1]Mouse data flows are emails, web pages, data requests, or other short-lived data flow

- Elephant flows are often related to high use low inspection traffic, for example, backups, database replication, VM migrations, data migrations, etc., inside the data centers that impact network bandwidth for minutes or hours or more.

Handling Elephant Flows

Application Filtering Intelligence detects and handles the elephant flows in the traffic. This feature helps to optimize the performance of the following GigaSMART cards when elephant flows are present in the traffic:

- HC1-X12G4
- SMT-HC3-C05
- SMT-HC0-Q02X08

In tunneled traffic, this feature detects the elephant flows, but it doesn't involve in optimizing the performance of the GigaSMART engine.

To detect the elephant flows in the traffic, do the following in the GigaVUE-FM:

1. On the left navigation pane, click on , go to **Physical** > **Nodes**.
2. Click on the required Cluster ID.
3. From the device view, go to **System** > **GigaSMART** > **GigaSMART Groups**.
4. Click **New** to create a new GigaSMART Group for detecting the traffic with elephant flow.
   a. Enter the name of the group in the **Alias** field.
   b. Select the ports in the **Port List**.

   You can also include the detection of elephant flow in a existing GigaSMART group.
5. In the **GigaSMART Parameters** > **Eflow** section, enable the **Eflow** checkbox to enable the detection of elephant flow.
6. Enable the **Log** check box to print the parameters of the elephant flow including the 5-tuple information into the GigaSMART logs.

NOTE: It is recommended to disable the check box after collecting the required parameters.

7. Enter the following parameters to identify the elephant flow:
   a. **Interval** — The interval within which packet-count and packet-ratio for a traffic flow are examined. The interval should be specified in seconds. The range lies between 0 to 3600. Specify the interval as 0 to ignore this parameter. The default value is **2 secs**.

b. **Packet Count**— Enter the maximum number of packets to be received by the flow within the given interval to categorize the flow as an elephant flow. The default value is **10,000**.

c. **Packet Ratio** — Enter the packet ratio, which is the percentage concentration of the packets in the flow to the packets seen overall by the gsgroup. Specify 0 to ignore this parameter. The default value is **0**.

You can handle the elephant flows in Application Filtering Intelligence Solution by using the gsgroup created to detect the elephant flow.

Refer to the GigaVUE-OS CLI Reference Guide to learn about the commands that must be configured to detect and handle the elephant flow of traffic.

Create Application Filtering Intelligence for Physical Environment

GigaVUE-FM allows you to create Application Filtering Intelligence by selecting the applications available from the **Total Applications** displayed on the Application Intelligence (AFI) dashboard. To create Application Filtering Intelligence, follow these steps:

1. On the left navigation pane, click Traffic  , select **Solutions** > **App Intelligence**.

NOTE: If you are creating Application Filtering Intelligence immediately after creating Application Monitoring, then you can proceed from Step 2.

2. Select the required application from the **Total Applications** in the right pane of the Application Intelligence dashboard. You can also select multiple applications for creating the Application Filtering Intelligence.
3. Click **Operations,** and select **App Filtering** from the drop-down list.

4. 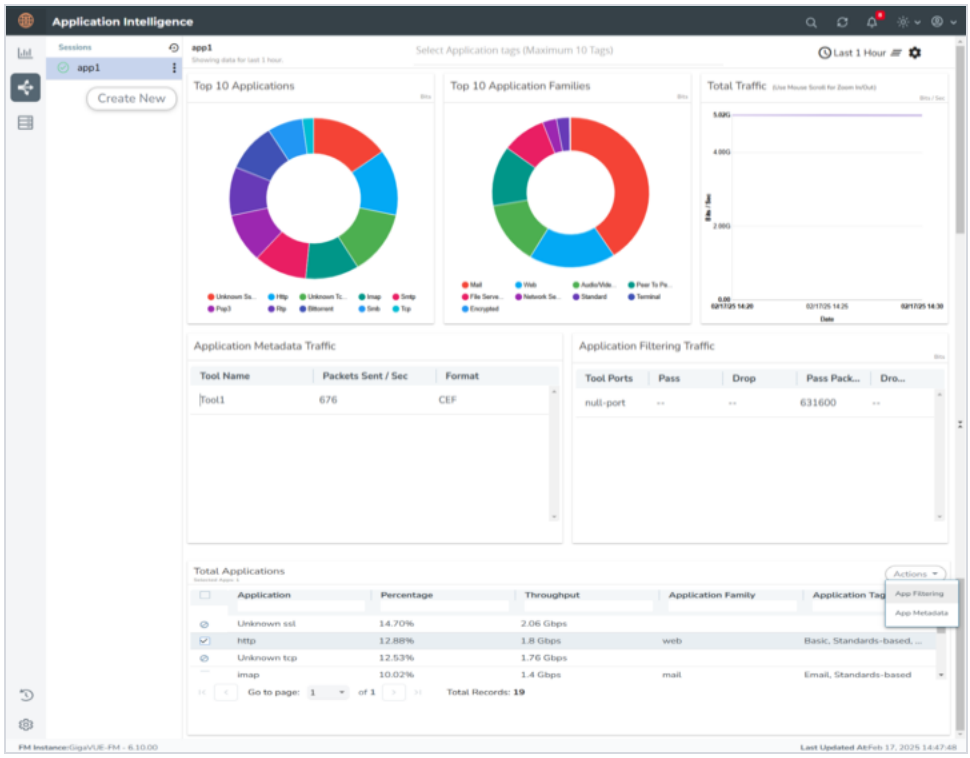

You can view the list of applications selected in the **Selected Applications** section.

5. Select either the **Pass** or **Drop** check box for an application to allow it to either pass through or get dropped off in the tool port present in the **Destination Traffic Priority**. You can also perform a search operation to filter the required application from the list of applications.
6. Use the **Destination Traffic Priority** section**,** to either choose the available tool port or add a new port for creating a traffic priority. In the **Select ports...** field, select the tool ports for sending the filtered applications traffic to the external tool. If you are unable to view the required port in the **Port** field, perform these steps:
   - Click **Port Editor**. Select the **Type** as **Tool** from the drop-down list for the required **Port Id**. Select **OK**.

     The selected Port appears in the list.
   - Click **Save**.

NOTE: You can select the Null Port option, when you need to collect only the application metadata without sending the filtered traffic to any physical port.

7. In the **Priority** section, you can perform the following actions:
   - Enable the **Pass All** check box to pass all the applications when there are no matching rules.
   - Click **Advanced Rules > Add a rule** to add new rules to perform advanced filtering on the application. For more details on adding a rule, refer to Adding Rules section.

> **NOTE:** Pattern match is supported in GigaSMART Generation 3 nodes.

> **NOTE:** To view the statistics of packets that are sent due to no rule match pass, view the Map Rule Counters. Refer Review Map Statistics with Map Rule Counters

8. In the **Destination Traffic Priority** section, click **+ Add New** to create additional **Destination Traffic Priority** (second level maps). In Application Filtering Intelligence, you can create a maximum of five Destination Traffic Priorities.

> **NOTE:** You can click and drag the icon ≡ to reorder the map priority when there are multiple priorities.

9. Click **Filter to** button for the corresponding **Priority** in a **Destination Traffic Priority** section for passing and dropping the applications to the required tool ports.

> **NOTE:** You cannot filter the traffic using applications when you select pattern match in the rules configuration.

In the **Application Filtering Intelligence Settings,** you can edit the following options while creating the Application Filtering Intelligence**:**

- Bidirectional (Default option).
- Buffer/Buffer count - The option is enabled default.
- Protocol- The default value is TCP-UDP. You can also select any of the following options:
  - TCP only
  - UDP only
  - TCP-UDP
  - TCP, UDP and SCTP
  - SCTP only
- Packet Count
- Timeout in seconds - The default value is 15 seconds.
- Sessions Field

10. Click **Save**.

You can view the **Application Filtering Intelligence** Statistics from the Application Intelligence Dashboard page.

Adding Rules

You can use **Advanced Rules** option to add more rules in Application Filtering. To add rules, do the following:

1. Click **Priority** > **Advanced Rules**> **Add a Rule**
2. Click the field **Select Options** in **Rule 1**. and select any of the following options:

- DSCP
- ERSPAN ID
- EtherType
- GRE Key
- GTP-U TEID
- IP Fragmentation
- IP Version
- IPv4 Destination
- IPv4 Source
- IPv4 protocol
- IPv4 TOS
- IPv4 TTL
- IPv6 Destination
- IPv6 Flow Label
- IPv6 Next Header
- IPv6 Source
- MAC Source
- MAC Destination
- MPLS Label
- Pattern Match — You can select the pattern type as either as follows and provide the respective values.
    - **String**
    - **Regex**- For example, you can use it if you want to filter HTTP sessions that include JSON and API traffic. This is supported in Gen2 and Gen3 nodes.

NOTE: For Regex examples, refer the **Supported Protocols and Attributes** section in the User Defined Application topic.

NOTE: The pattern matching rule and the application filter rule cannot be used together within the same second-level map.

- Port Destination
- Port Source
- TCP Control

- VLAN
- VN-Tag Destination VIF ID
- VN-Tag Source VIF ID
- VN-Tag VIF List ID
- VXLAN ID

3. Click **Pass** or **Drop** check box to allow it to either pass through or get dropped off in the tool port present in the **DestinationTrafficPriority**.
4. Click **Save**.

Application Filtering Intelligence can also be configured for virtual environment, refer to Application Filtering Intelligence section in *GigaVUE V Series Applications Guide* for more detailed information.

AFI with Slicing and Masking

The steps to configure AFI with slicing are as follows:

1. Create **GigaSMART group** for GigaSMART Engine.
2. Enable **Application Session Filtering**.
3. Set the **Buffer Size** to 2.
4. Create **GigaSMART Operation**.
5. Select **GigaSMART Group**.
6. Select **ASF**, **APF** and **Slicing** for **GigaSMART Operations (GSOP)**.
7. For Slicing, select protocol and desired offset.
8. Create **Application Session Filtering** profile from GigaVUE-FM using the following steps.
    a. Click **GigaSMART**.
    b. Click **App Idenfication**.
    c. Click **Application Session Filtering**.
    d. Enable **Bidirectional**
    e. Enable **Buffer**.
    f. Set the **Buffer Count Before Match** to 20.
    g Select **TCP-UDP** for **PROTOCOL**.
    h. Select **5tuple** for **Session Fields**.
9. Create **New Virtual Port**.
10. Select the GigaSMARTS group and **ASF Profile** created in the previous steps.
11. Create **First Level** Map using the following steps.

a. Click **Traffic**.

b. Click **Maps**.

c. Click **New**.

d. Select **Type** as **First Level**.

e. Select **Source** as Network Ports and **Destination** as Vport.

12. Create **Second Level** Map using the following steps.

a. Click **Traffic**.

b. Click **Maps**.

c. Click **New**.

d. Select **Type** as **Second Level**.

e. Select **Source** as **Vport** and To Port as **Tool Port**

13. Attach the GSOP.
14. Create the Application profile using the following steps.

a. Click Create New Application Profile drop down in the PASS/DROP application profile option.

b. Search and select applications for Pass/Drop.

15. Save the map.

The solution is deployed and you can verify using the following commands:

- Show map stats
- Show vport stats
- Show gsop stats
- Show gsgroup stats and port stats of Source port, GS engine and Tool port

NOTE: The same procedure can be used for Masking and Advanced Flow Slicing GSOPs by selecting Masking or Advanced Flow Slicing instead of Slicing in the 6th step.

View the Application Intelligence Dashboard

After creating the Application Intelligence Session, you can monitor the applications in the network by the reports displayed in the Dashboard as shown in the following figure:

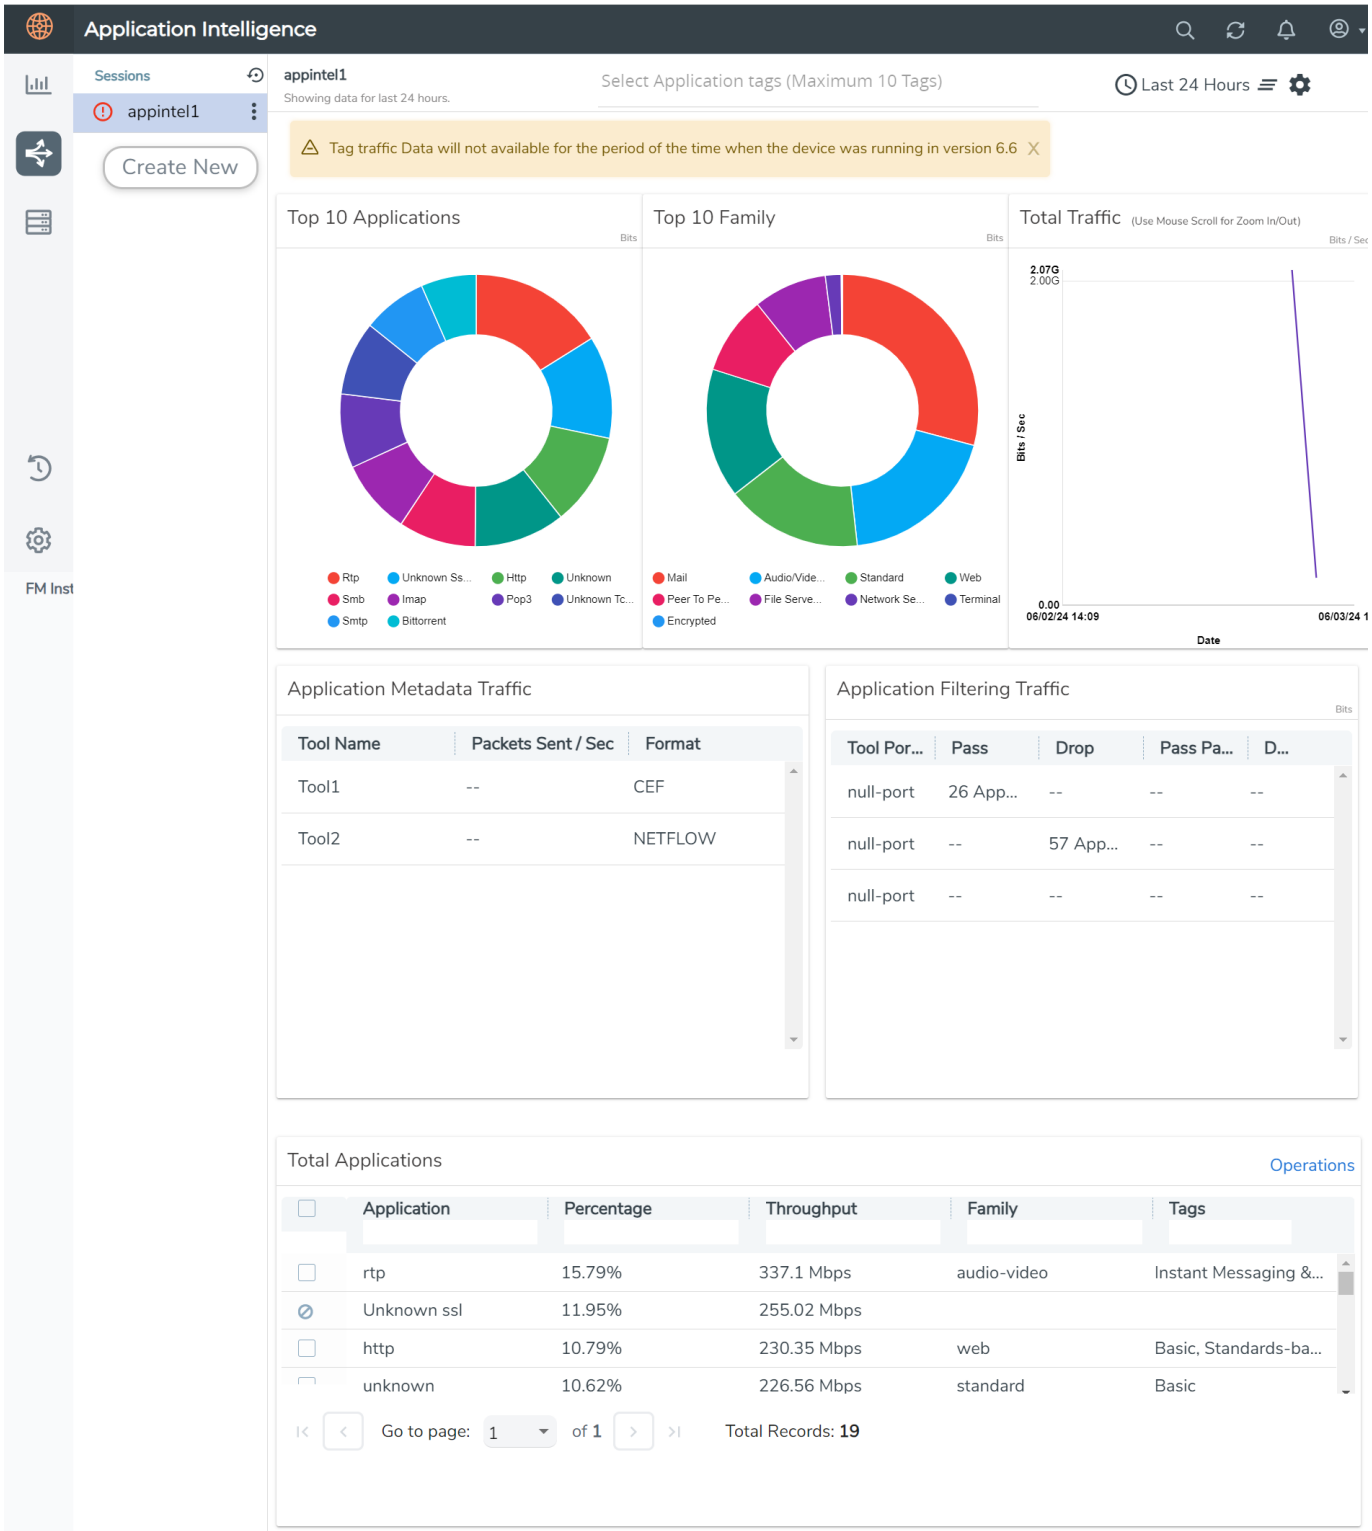

Application Intelligence Dashboard displays the following metrics:

NOTE: The **Top 10 Application Families** metric and the **Select Application Tags** option are supported only on Gen 3 GigaSMART module and available from software release version 6.7.00 only.

- **Top 10 Applications**: You can view a graphical representation of top 10 applications running in the network based on the metrics. When you hover over the Pie-chart, GigaVUE-FM shows the application name in the network. The legend for the Pie-chart appears at the bottom. When you select a pie, you can view the corresponding data highlighted in the Total Applications table.
- **Top 10 Application Families**: You can view a graphical representation of top 10 application families running in the network based on the metrics. When you hover over the Pie-chart, GigaVUE-FM shows the application family name in the network.

NOTE: The **Top 10 Applications** Pie-chart may include user defined applications, but their corresponding Application Family will not be seen in the **Top 10 Application Families** Pie-chart. The user defined applications are not categorized under any family. So, if user defined applications are involved, there may be a variation in the **Top 10 Applications** and **Top 10 Application Families** Pie-charts.

- **Total Traffic**: You can view the total traffic of the network represented in the linear form of a graph.
- **Total Applications**: You can view the applications and their bandwidth in the network. You can also select the required application for filtering and exporting metadata by using the Operations field. When the following application family names are displayed in the Application list, it implies that the application could not be successfully identified:
  - unknown — unknown application name and protocol
  - unknown-TCP — unknown application name that belongs to TCP flow.
  - unknown-HTTP —unknown application name that belongs to HTTP flow.
  - unknown-UDP —unknown application name that belongs to UDP .
  - unknown-SSL —unknown application name that belongs to SSL flow.

NOTE: You can filter the applications listed under **Total Applications** using the parameters such as **Application**, **Percentage**, **Throughput**, **Family**, and **Tags**.

NOTE: You can forward the unknown application family traffic to a packet capture tool, analyze the packet capture and configure User Defined Application signatures to identify them.

- **Application Filtering Traffic**: You can view the statistics of the applications that are filtered in the tool ports in the dashboard after creating an Application Filtering Intelligence session for a device. You can view the number of applications that are dropped and passed through for the Application Filtering Traffic.
- **Application Metadata Traffic:** You can view the details of the Tool Name and the Format in which the metadata of the application is exported.

You can use the **Select Application Tags** option to filter the traffic specific to the selected application tags. On selecting the tags for filtering, the following widgets will show only the data specific to the selected tag:

- **Total Traffic**
- **Top 10 Applications**
- **Top 10 Application Families**
- **Total Applications**

GigaVUE-FM enables you to view the above metrics for a particular period by selecting the date and time from the dashboard.

GigaVUE-FM takes more than five minutes to display the application statistics since the export interval is fixed at five minutes. For the first fifteen minutes after creating the solution, if GigaVUE-FM receives traffic, it will show real-time data. If there is no traffic during this time, it will take at least eleven minutes to display the statistics once traffic is received.

You can also choose to view the graphs in the dashboard for the metrics in bytes, packets or flows. To view the metrics in bytes, packets or flows, click the gear ⚙ button in the right corner of the Application Intelligence dashboard.

NOTE: By default, you can view the metrics in the dashboard for the last one hour, if the period is not selected.

NOTE: After upgrading to 6.4.00, you will lose the complete historical application visualization data except for the last twenty-four hours.

About De-duplication

Required License: De-duplication

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. For more information, refer to the GigaSMART De-Duplication.

To enable De-duplication in the Application Filtering Intelligence or Application Metadata Intelligence, click **De-duplication.**

Requirements and guidelines:

- You must install licenses for De-duplication and Application Filtering Intelligence on your GigaSMARTModules to enable the features in GigaVUE-FM.
- Application Filtering Intelligence and De-duplication can be combined on the same GigaSMART engine.
- When combining Application Filtering Intelligence (AFI) and De-duplication, AFI has the higher priority and it is executed before De-duplication.

Application Metadata Intelligence

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 for more details on the devices that support GigaSMART operations.

Application Metadata Intelligence generates more than 6000 attributes for more than 4300 applications without impacting the users, devices, applications, or the network appliances. The feature identifies applications even when the traffic is encrypted.

Application Metadata Intelligence (AMI) is enabled to multi-collect protocols with more than one metadata attribute of the same type. The multi-collect feature supports additional protocols such as DNS, GTP,GTPV2, DHCP, HTTP, HTTPS, SSL, HTTP_PROXY, HTTP2, KERBEROS5, and DHCP6.

NOTE: The DPI attributes of certain applications will be extracted and exported only if the traffic is decrypted. Please refer to the list of such applications under "Decrypted Traffic Support" in the GigaVUE-FM Application Protobook.

The generated metadata is exported in IPFIX (IP Flow Information Export) format and CEF (Common Even Format) to security analytics and forensics tools thereby providing greater visibility to enforce corporate compliance.

NOTE: Work with your Gigamon representative to get the latest IPFIX elements.

The output from the Application Metadata Intelligence in CEF format can also be converted to JSON format using Observability Gateway Application (AMX) application. To learn more about OGW application refer to Application Metadata Exporter.

Application Metadata Intelligence generates metadata only if the application is allowed to be passed in Application Filtering Intelligence. For example, Application Metadata Intelligence has the capability to generate metadata for HTTP traffic only if Application Filtering Intelligence filters in the HTTP traffic.

Create Application Metadata Intelligence for Physical Environment

You can create an Application Metadata Intelligence session in GigaVUE-FM by selecting the applications available from the Total Applications displayed on the Application Intelligence (AMI) dashboard.

To create an Application Metadata Intelligence session, follow these steps:

1. On the left navigation pane, click Traffic  , select **Solutions** > **App Intelligence**. .
2. In the **Application Intelligence Session** , click **Application Metadata**.

   You must configure Application Intelligence session to monitor the application on the network and to display them on the **Total Applications**. To create Application Intelligence session refer to Application Intelligence Session.
3. From the navigation pane, click **App Intelligence**. Select the applications from the **Total Applications** in the right pane of the Application Intelligence dashboard.
4. Click **Operations** and select **App Metadata** from the drop-down list.

   You can view the list of applications selected in the **Selected Applications** section.

   Application Metadata Intelligence generates more than 6000 attributes for more than 4300 applications without impacting the users, devices, applications, or the network appliances. The feature identifies applications even when the traffic is encrypted.
5. Expand the application and select the attributes to be extracted.

NOTE: You can select the attributes only if the application has attributes. If Application Filtering Intelligence License is available, you must create Application Filtering to create Application Metadata Intelligence.

NOTE: The attributes IP source and IP destination cannot be configured to be extracted from the **App Editor** section. To export, them utilize the **Advanced Settings** > **Collects** section.

6. In the **Destination Traffic** section, you can attach five exporters to a GigaSMART group. To add an exporter, click **+ Add New** and enter the following information:

   a. **Tool Name** - Enter the name of the tool to which you want to export the application-specific metadata.

b. **Tool IP Address** - Enter the IP address of the tool to which you want to export the application-specific metadata. You can provide either IPv4 or IPv6 address.

c. **IP Interface -** Select the IP interface through which the GigaVUE-FM needs to export the application-specific metadata to the tools. You can select either IPv4 or IPv6 interface. Provide IPv4 address as the destination address for IPv4 interface and provide IPv6 address as the destination address for IPv6 interface. You can also choose to create a new IP Interface.

d. **Tool Template** - Select the tool template containing the list of applications and attributes that needs to be exported. Click the Save option to save the template as a new template.

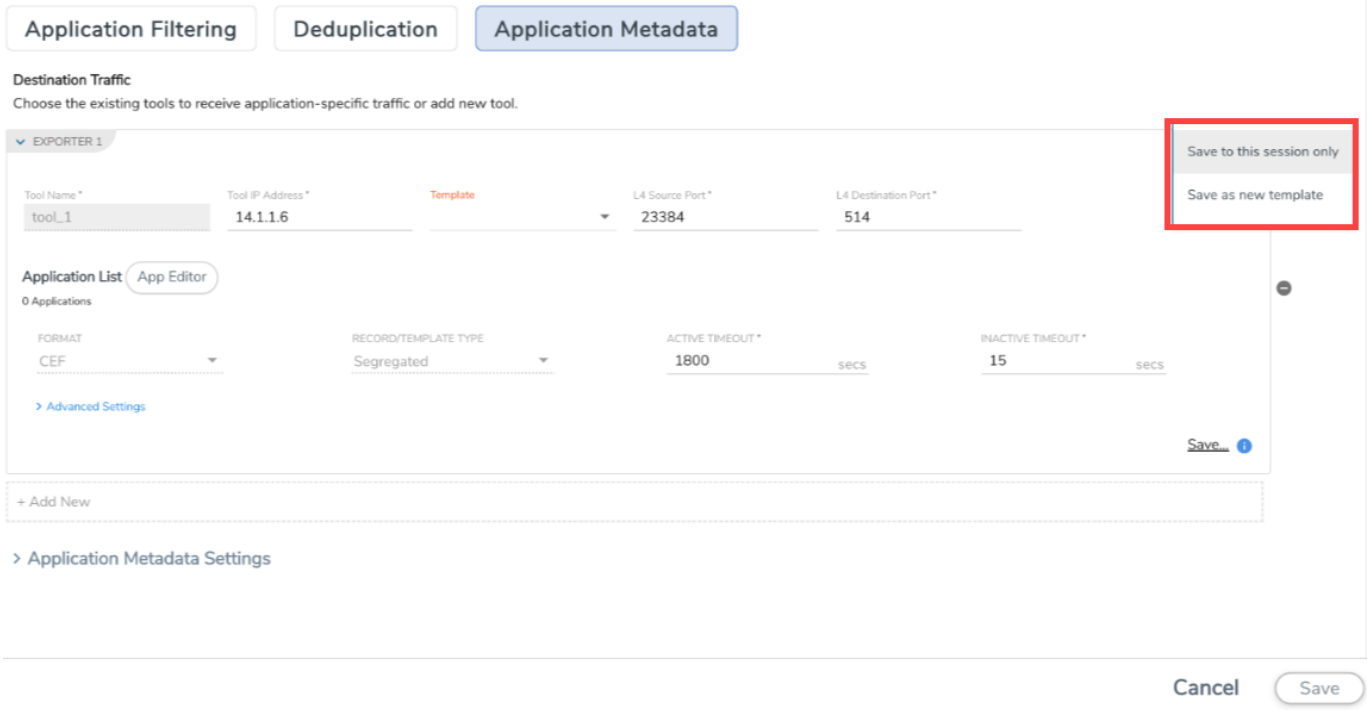

When editing the exporter template, if you change any of the non-editable fields (Format, Record Type, NetFlow Version), the solution fails.

NOTE: When you create a session with flow-behaviour as bi-directional, GigaVUE-FM allows you to select Netflow v5 and v9 templates. When you edit the same session, you cannot select the Netflow v5, and v9 templates.

e. Enter the L4 Source Port, and the L4 Destination Port details.

NOTE: If the export format is CEF, the default value for L4 destination port is 514. If the export format is NetFlow, the default value for L4 destination port is 2055.

f. Select the Format as CEF or NetFlow from the **Format** drop-down list.

g. Select the Record/Template type as either of the following:

- Segregated - The application-specific attributes and the generic attributes will be exported as an individual record 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. Enter the Active Timeout, InActive Timeout and Template Refresh time interval.

NOTE: The format and the record/template type get selected automatically, after selecting the Tool Template.

7. In the **Advanced Settings** > **Collects** section, you can select the following packet attributes:
   - Counter - Select the Bytes, and Packets.
   - IPv4 - Select the required attributes. By default, Source Address, Destination Address, and Protocol are enabled.
   - IPv6 - Select the required attributes. By default, Source Address, Destination Address, and Next Header are enabled.
   - Transport -Select the required attributes. By default, Source Port, Destination Port are enabled.

   By default, the above collect types are displayed. Click ⊕ to add the following collect types:
   - Data Link - Select any one of the parameters such as Source Mac, Destination Mac and VLAN.
   - Timestamp - Select the required timestamp such as System Uptime First, Flow Start, System Uptime Last, and Flow End.
   - Flow - Select the parameter as End Reason if required.
   - Interface - Select any one of the parameter such as Input Physical, Output Physical and Input Name.
8. In the **Application Metadata Settings** section:
   a. Select the Flow Behavior as any one of the following:
      - Uni-Directional
      - Bi-Directional. The default value is Bi-Directional.
   b. Enter the Timeout and Cache Size. The cache size range allowed for different platform is given in the following table:

Platform	Range in million
GigaVUE-HC1	1
GigaVUE-HC2	1
GigaVUE-HC3	1-2

   c. You can enable or disable the **Multi-Collect** option to perform the following:

- **Enable** — Enables the multi-collect of attributes within a given Metadata Store cache which means that if a configured attributes is seen in multiple packets within the same flow, each of these information is collected. By default, when a new cache is created, multi-collect is enabled. When upgraded from an older release, the multi-collect option is enabled.
- **Disable** — Disables the multi-collect of attributes within a given Metadata Store cache.

d. You can use the toggle button to enable or disable the **Aggregate Mode**, which is disabled by default. You need to delete the existing solution and recreate the solution to enable the **Aggregate Mode**. The **Aggregate Mode** option is applicable only for Gen 3 devices. Only one exporter is supported with the **Aggregate Mode** enabled.

NOTE: You need to enable the **Aggregate Mode** option to export the minimum, maximum, and mean of RTT values for the following list of supported protocols and attributes and also the aggregate of TCP Lost byte values collected per export time interval.

Protocol Name	Attribute
http	rtt
icmp	rtt
icmp6	rtt
ssh	rtt
tcp	rtt
tcp	rtt_app
telnet	rtt
wsp	connect_rtt
wsp	query_rtt

e. You can enable or disable the **Advance Hash** option to perform the following:

- **Enable** — Configures metadata cache advance-hash for encapsulated flows . This feature improves the efficiency of scheduling the distribution of encapsulated flows. It also improves the distribution of flows in service provider deployment cases. By default, when a new cache is created, advance hash is enabled. When upgraded from an older release, the advance hash is enabled.
- **Disable** — Disables the metadata cache advance-hash for flows.

f. If you want to include the VLAN ID along with the 5-tuple to identify the traffic flow, select the **Data Link** and enable the

g. **VLAN** option.

h. In the **Observation Domain ID** field, enter a value to identify the source from where the metadata is collected. The range is from 0 to 255.

For example: If you enter 5 in this field, then the observation domain ID is calculated as follows:

Observation Domain ID (4-Bytes)	
Byte 1	0
Byte 2	1
Byte 3	GS engine slot (for e.g. 2 if 1/2/e1)
Byte 4	User defined (for e.g. 5). Default : 0.

The calculated value of Observation Domain Id in Hexadecimal is **00 01 02 05**, and in Decimal is **66053**.

9. In the **Selected Applications** section, select **Export** and click **Export To** for the applications that needs to be exported to the destination tool.
10. Click **Save**.

Create NetFlow Session for Physical Environment

Note: This configuration is applicable only when using NetFlow License.

1. On the left navigation pane, select **Traffic > Solutions > Application Intelligence**.
2. Click **Create New**. The **Create Application Intelligence Session** page appears.

NOTE: If the Create button is disabled, check whether a valid license for Application Metadata Intelligence or Application Filtering Intelligence is available.

3. In the **Basic Info** section complete the following:

- Enter the name and description (optional).
- Select **Physical** in the Environment field.
- Select the node from the list of nodes.

4. In the **Configurations** section, view the following:

a. The export interval is fixed at 5 minutes, at which the Application Intelligence session will generate the reports for application monitoring.

> **NOTE:** When you upgrade to 6.4, the prior configured export interval will be updated to 5 minutes.

b. Select a GigaSMART Group. You can also choose to create a new GigaSMART Group.

- Provide a name in the **Alias** field.
- Select a port or multiple ports from the **Port List**.
- Click **Save**.

For more information about creating a new GigaSMART group, refer to the Configure a GigaSMART Group section.

5. If you are unable to view the required port in the **Port** field, perform these steps:
   - Click **Port Editor.** Select the **Type** as **Tool** from the drop-down list for the required **Port Id**. Select **OK**.

     The selected Port appears in the list.
   - Select the **Type** as:
   - IPv4 - to allow the traffic in IPv4 interface.
   - IPv6 - to allow the traffic in IPv6 interface.
   - Provide the **IP Address**, **IP Mask, Gateway**, and **MTU**. Provide the IP address corresponding to the IP interface selected.
   - Click **Save**.
6. In the **Source Traffic** section, select a source port that require application monitoring in the **Source ports** field. Source port can be a single port, multiple ports, and port groups.

   > **NOTE:** Ports already used as source ports in the intent-based orchestrated solution will not be listed in the drop-down.

7. Configure the rules for filtering the required traffic in the **L2-L4 Rules** fields. To configure a rule:

   a. Click **Select Conditions**. Select the required parameters from the drop-down list.

   b. Select the value for the parameters from the drop-down.

   c. Select the required options:

   - Pass or Drop - Based on the parameter selected in the Conditions fields, the traffic that matches the conditions will either be passed or dropped.
   - Bidirectional - Allows the traffic in both directions of the flow.

   > **NOTE:** Click "**+**" to create multiple rules for filtering the required traffic, and click "**+ New Source Traffic**" to create multiple sources with filtering options.

8. Click on the **NetFlow** tab.

9. In the **Destination Traffic** section, click **+ Add New** to create an exporter to receive application-specific traffic. You can only create a maximum of 5 exporters. Enter the following details:

Field	Description
Tool Name	Enter the tool name.
Tool IP Address	Enter the tool IP address.
Template	Select the tool template. Refer to Tool Templates for more details on what tool templates are and to create custom tool templates.
L4 Source Port	Port from which the connection will be established to the target. For Example, if A is the source and B is the destination, this port value belongs to A.
L4 Destination Port	Port to which the connection will be established from the source. For Example, if A is the source and B is the destination, this port value belongs to B.
APPLICATION ID	Enable to export the data with Application Id.
Format	NetFlow
Record / Template 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.
Active Timeout	Enter the active timeout value in seconds.
Inactive Timeout	Enter the inactive timeout in seconds.
Version	Select the NetFlow version.
Template Refresh Interval	Enter the time interval at which the template must be refreshed in seconds.

10. In the **Advanced Settings** > **Collects** section, the following details are already configured.

> **NOTE:** When the template is NetFlow v5 or when the format is NetFlow and the version as V5 you cannot modify the **Collects**.

> **NOTE:** The GTP -U collects are disabled if the template is the Netflow v9 .

- TimeStamp
- Counter
- Interface
- IPv4
- Transport

11. In the **ApplicationMetadata Settings** section:

    a. Select the Flow Behavior as any one of the following:
       - Uni-Directional
       - Bi-Directional. The default value is Bi-Directional.

    b. Enter the Timeout and Cache Size.

    c. You can enable or disable the **Multi-Collect** option to perform the following:
       - **Enable** — Enables the multi-collect of attributes within a given Metadata Store cache which means that if a configured attributes is seen in multiple packets within the same flow, each of these information is collected. By default, when a new cache is created, multi-collect is enabled. When upgraded from an older release, the multi-collect option is enabled.
       - **Disable** — Disables the multi-collect of attributes within a given Metadata Store cache.

    d. You can use the toggle button to enable or disable the **Aggregate Mode**, which is disabled by default. You need to delete the existing solution and recreate the solution to enable the **Aggregate Mode**. The **Aggregate Mode** option is applicable only for Gen 3 devices. Only one exporter is supported with the **Aggregate Mode** enabled.

Protocol Name	Attribute
http	rtt
icmp	rtt
icmp6	rtt
ssh	rtt
tcp	rtt
tcp	rtt_app
telnet	rtt
wsp	connect_rtt
wsp	query_rtt

NOTE: You need to enable the **Aggregate Mode** option to export the minimum, maximum, and mean of RTT values for the following list of supported protocols and attributes and also the aggregate of TCP Lost byte values collected per export time interval.

e. You can enable or disable the **Advance Hash** option to perform the following:
   - **Enable** — Configures metadata cache advance-hash for encapsulated flows . This feature improves the efficiency of scheduling the distribution of encapsulated flows. It also improves the distribution of flows in service provider deployment cases. By default, when a new cache is created, advance hash is enabled. When upgraded from an older release, the advance hash is enabled.
   - **Disable** — Disables the metadata cache advance-hash for flows.

f. If you want to include the VLAN ID along with the 5-tuple to identify the traffic flow, select the **Data Link** and enable the **VLAN** option.

g. In the **Observation Domain ID** field, enter a value to identify the source from where the metadata is collected. The range is from 0 to 255. The calculated value of Observation Domain Id in Hexadecimal is **00 01 02 05**, and in Decimal is **66053**.

12. Click **Save**.

NetFlow Dashboard

In Appviz, only the traffic statistics are displayed as applications cannot be configured and used in the NetFlow configuration.

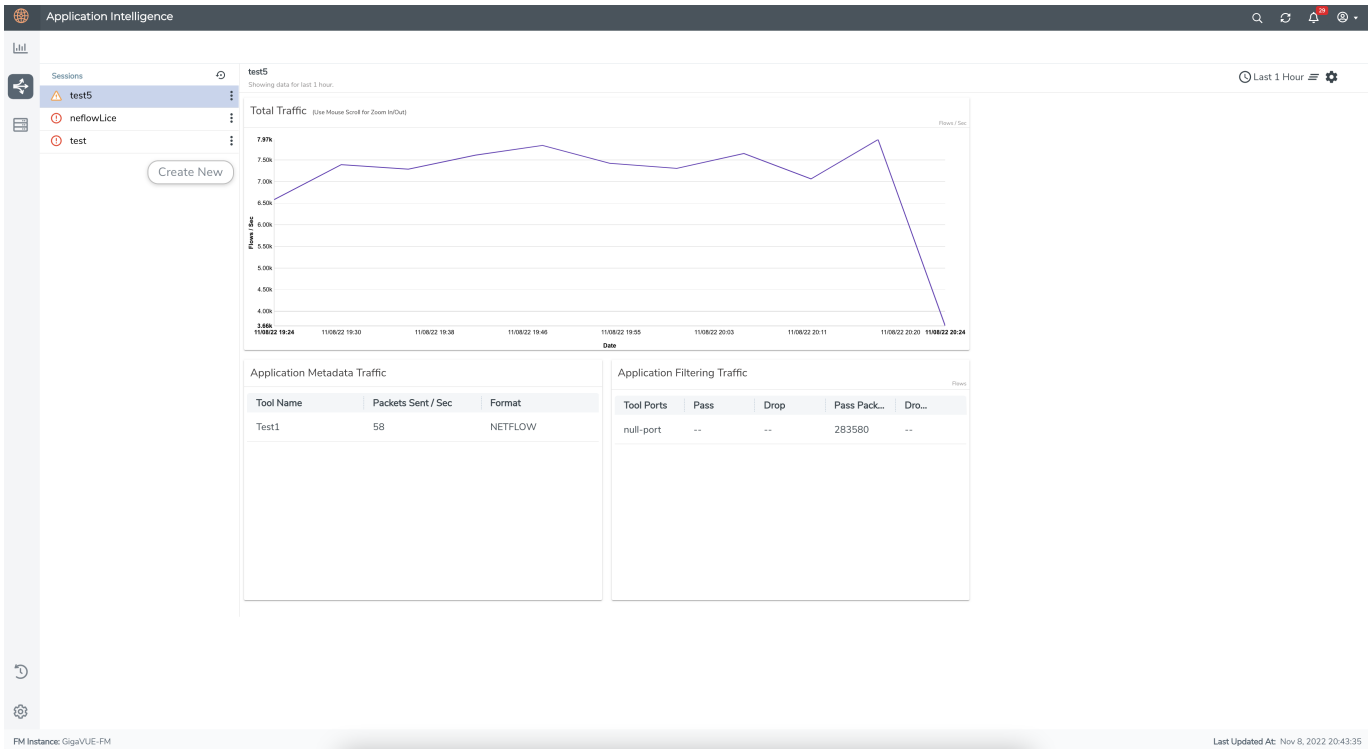

Tool Templates

Tool template in Application Metadata Intelligence predefines a list of applications and its attributes, which you can choose as per your requirements while configuring 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. By default, you can find the following tool templates:

- BroMetadata Template
- Netflow V5 Template
- SplunkMetadata Template

Starting from software version 6.2.00, the following tool templates are supported:

- SecurityPosture
- RogueActivity
- SuspiciousActivities
- AnomalousTraffic
- Troubleshooting
- M2131Logging
- UnmanagedAssets

The following table provides the purpose of each of the tool templates when used in Application Metadata Intelligence :

Tool Template	Purpose
BroMetadata Template	For selecting applications and attributes that can be detected by Bro sensor
Netflow V5 Template	For emulating NetFlow V5 behavior
SplunkMetadata Template	For providing a quick insight into the network traffic generated by various applications and protocols
SecurityPosture	For detecting flaws in securing the applications in the network
RogueActivity	For detecting unsanctioned applications that can pose challenges to network security
SuspiciousActivities	For detecting issues related to unmanaged devices, suspicious connections, and traffic outside normal limits in the network
AnomalousTraffic	For detecting challenges with HTTP, HTTPS, and DNS traffic in the network

Troubleshooting	For detecting latency, connectivity, and protocol errors in the network
M2131Logging	For U.S. Office of Management and Budget M-21-31 logging requirements
UnmanagedAssets	For providing visibility into unmanaged hosts and devices in the network

NOTE: You cannot edit the above templates. Hover over the Description column to view the description of the default tool templates.

You can create new tool templates according to your requirements. You can also edit and clone the templates. Refer to Create Custom Tool Templates for more information.

Create Custom Tool Templates

To create a customized Tool Template, do the following:

1. On the left navigation pane, click  , and then select **Resources > Tool Templates**.

   You can view two system defined templates by default.
2. Click **Create**. The **New Tool Template** page appears.

a. Enter the **Template Name**.
b. Select the **GS Version**.
c. Enable the **Application Family** or **Application Tag** option. Only if you choose the **GS Version** as **1.680.0.24**, the **Application Family** and **Application Tag** options appear.
d. Enable the **fastmode** check box. When the **fastmode** option is enabled, only the **fastmode** supported attributes will be listed. Only if you choose the **GS Version** as **1.680.0-20** or **1.680.0.24**, the **fastmode** option appears.
e. Enter the **Description**.

3. Click the **App Editor** button. The **App Editor** page appears.
4. Click the **Application Family** field and select an Application Family such as antivirus, webmail that needs to be filtered from the traffic. You can also select multiple application families.

   If you choose to add or delete all the applications in a family, click **Add All Application in Families** or **Delete All Application in Families**. or

Click the **Application Tags** field and select an **Application Tags** that needs to be filtered from the traffic.

You can also select multiple application tags. If you choose to add or delete all the applications in **Application Tags**, click **Add All Application in Tags** or **Delete All application in Tags**.

5. Select the format (CEF/NetFlow) in which the application details need to be exported to the tool.
6. Select the record or template 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.
7. Enter the **Active timeout** and **Inactive timeout** in seconds.

   NOTE: Select the Version and the Template Refresh Interval for the Netflow format.

8. In the **AdvancedSettings** > **Collects** section, you can view and configure the following collect types:
   - **Counter** - Select the Bytes and Packets.
   - **IPv4** - Select and enter the required information in the fields such as Source Address, Destination Address, Fragmentation, and Header Size and Payload Size.
   - **IPv6** - Select and enter the required information in the fields such as Source Address, Destination Address, Fragmentation, and Header Size and Payload Size.
   - **Transport** -Select and enter the required information in the fields such as Source Address, Destination Address, and Header Size and Payload Size.

   By default, the above collect types are displayed. Click ⊕ to add the following collect types:
   - **Data Link** - Select any one of the parameters such as Source Mac, Destination Mac and VLAN.
   - **Timestamp** - Select the required timestamp such as System Uptime First, Flow Start, System Uptime Last, and Flow End.
   - **Flow** - Select the parameter as End Reason if required.
   - **Interface** - Select any one of the parameter such as Input Physical, Output Physical and Input Name.
9. Click **Save**. The new tool template is added to the list view.
10. Select a tool template and click the ellipsis to perform the following:
    - **View Details** - To view the details in the template.
    - **Edit** - To edit the parameters and fields in the template.
    - **Delete** - To delete the template from the list.
    - **Duplicate** - To duplicate the template in the list.
    - **Export** - To export a tool template.

NOTE: When a solution is **fastmode** enabled, you can only select the tool templates that has all the attributes supported in the **fastmode**. The other tool templates that are listed will be greyed out. Refer "Create an Application Intelligence Session" topic for information on **fastmode**.

NOTE: You can edit an existing template and the attributes associated with it as required, and save the updates as a new template while creating an Application Metadata Intelligence session. Refer to Create Application Metadata Intelligence for Physical Environment for details.

NOTE: You can select the required applications without selecting the application family.

a. Click the **Application** field and select an application or multiple applications that needs to be exported from the traffic.

b. Click **Add** to add the Application details.

11. Select the format (CEF/NetFlow) in which the application details need to be exported to the tool.
12. Select the record or template 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.
13. Enter the **Active timeout** and **Inactive timeout** in seconds.

    NOTE: Select the Version and the Template Refresh Interval for the Netflow format.

14. In the **AdvancedSettings** > **Collects** section, you can view and configure the following collect types:
    - **Counter** - Select the Bytes and Packets.
    - **IPv4** - Select and enter the required information in the fields such as Source Address, Destination Address, Fragmentation, and Header Size and Payload Size.
    - **IPv6** - Select and enter the required information in the fields such as Source Address, Destination Address, Fragmentation, and Header Size and Payload Size.
    - **Transport** -Select and enter the required information in the fields such as Source Address, Destination Address, and Header Size and Payload Size.

By default, the above collect types are displayed. Click ⊕ to add the following collect types:

- **Data Link** - Select any one of the parameters such as Source Mac, Destination Mac and VLAN.
- **Timestamp** - Select the required timestamp such as System Uptime First, Flow Start, System Uptime Last, and Flow End.
- **Flow** - Select the parameter as End Reason if required.
- **Interface** - Select any one of the parameter such as Input Physical, Output Physical and Input Name.

15. Click **Save**. The new tool template is added to the list view.
16. Select a tool template and click the ellipsis to perform the following:
    - **View Details** - To view the details in the template.
    - **Edit** - To edit the parameters and fields in the template.
    - **Delete** - To delete the template from the list.
    - **Duplicate** - To duplicate the template in the list.
    - **Export** - To export a tool template.

NOTE: When a solution is **fastmode** enabled, you can only select the tool templates that has all the attributes supported in the **fastmode**. The other tool templates that are listed will be greyed out. Refer "Create an Application Intelligence Session" topic for information on **fastmode**.

NOTE: You can edit an existing template and the attributes associated with it as required, and save the updates as a new template while creating an Application Metadata Intelligence session. Refer to Create Application Metadata Intelligence for Physical Environment for details.

Import and Export Tool Templates

You can export and import tool templates. The following are the advantages of exporting and importing tool templates:

- Retrieve a tool template that was deleted unintentionally.
- Deploy the template in another device.
- Re-deploy the template in the device after GigaVUE-FM is upgraded to a new version (in case of issues in the existing solution).

To export a tool template:

1. In the Tool Templates page, select a template that must be exported.
2. Click the ellipsis and select **Export**.
3. Save the file to the required location.

To import a tool template:

1. Click **Import**.
2. Select the required file from the local folder.
3. Click **Deploy** to deploy the template.

RTT Attribute in AMI

The RTT attribute supported in AMI export is in the struct timeval format. The size of the attribute is 16 bytes. The first 8 bytes represent the value in seconds and the next 8 bytes represent the value in microseconds.

The values are exported in the network byte order.

In the following example:

- The PEN id 3365 has a RTT value of 0 seconds and 0xb827 microseconds
- The PEN id 3384 has a RTT value of 0 seconds and 0xbfb8 microseconds

```
Flow 1
    Enterprise Private entry: (Gigamon Systems LLC) Type 3365: Value (hex bytes): 00 00 00 00 00 00 00 00 00 00 00 00 00 00 b8 27
    Enterprise Private entry: (Gigamon Systems LLC) Type 3384: Value (hex bytes): 00 00 00 00 00 00 00 00 00 00 00 00 00 00 bf b8
```

View the Health Status of a Solution

The health of an Application Intelligence solution is determined by the health status of the following components, and the configuration status during deployment of the solution in a device:

- IP interface
- Source Port
- Destination Port
- GigaSMART Engine

You can view the health status of the solution by a color indication next to the name of the solution as shown in the following figure:

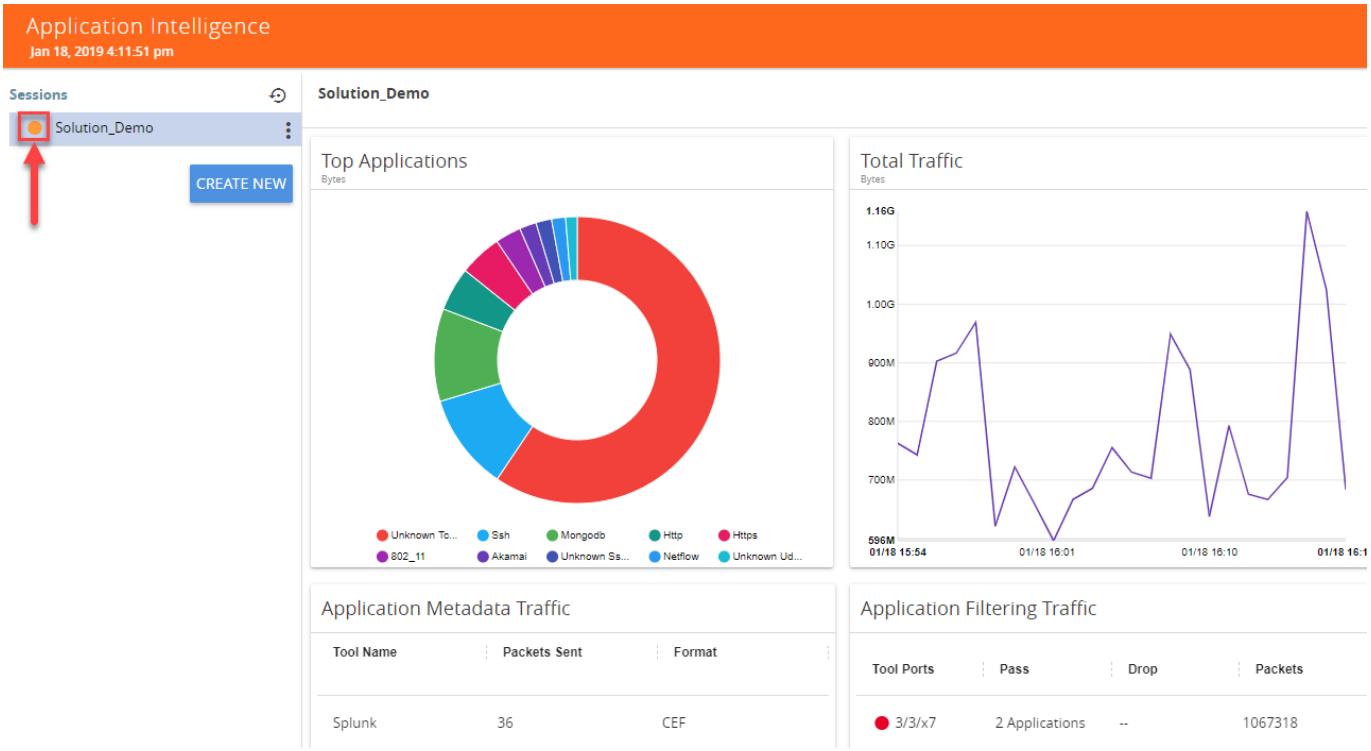

The health status of a solution is indicated by the following colors:

Color	Health Status of a Solution
Green	Healthy - All the components in a solution are functioning properly.
Red	Unhealthy - Any of the components in a solution is not functioning properly.
Amber	Partially Healthy

You can also view the reason for a failover when you hover your mouse over the color indicator next to the name of the solution. To avoid this scenario:

1. On the right side of the top navigation bar, click  .
2. On the left navigation pane, select **System** > **Thresholds**.
3. Set the threshold value of GigaSMART engine port packet drops to zero.

The following table provides the health state of the Application Intelligence solution corresponding to the health state of its associated components:

Table 1: Health status of Solution Vs Health status of Associated Components

Health Status of a Solution	GSGroup	IP Interface	Network Port and Tool Port	Metadata Exporter	Configuration Deployment Status
Red	Unhealthy	Healthy	Healthy	Healthy	Success
Amber	Partially Healthy	Healthy	Healthy	Healthy	Success
Red	Healthy	Unhealthy	Healthy	Healthy	Success
Amber	Healthy	Partially Healthy	Healthy	Healthy	Success
Amber	Healthy	Healthy	Some Maps are Unhealthy/Partially Healthy	Healthy	Success
Red	Healthy	Healthy	All Maps are Unhealthy	Healthy	Success
Red	Healthy	Healthy	Healthy	All Metadata Exporters are Unhealthy	Success
Amber	Healthy	Healthy	Partially Unhealthy	Healthy	Success
Amber	Healthy	Healthy	Healthy	Some metadata exporters are Unhealthy	Success
Red	Healthy	Healthy	Healthy	Healthy	Failed
Amber	Healthy	Healthy	Healthy	Healthy	Partial Success
Red	Healthy	Healthy	Healthy	Healthy	Failed
Green	Healthy	Healthy	Healthy	Healthy	Success

To view the details of a solution, click **View Details** > **Troubleshoot**. For more details, refer to View the Details of an Application Intelligence Session.

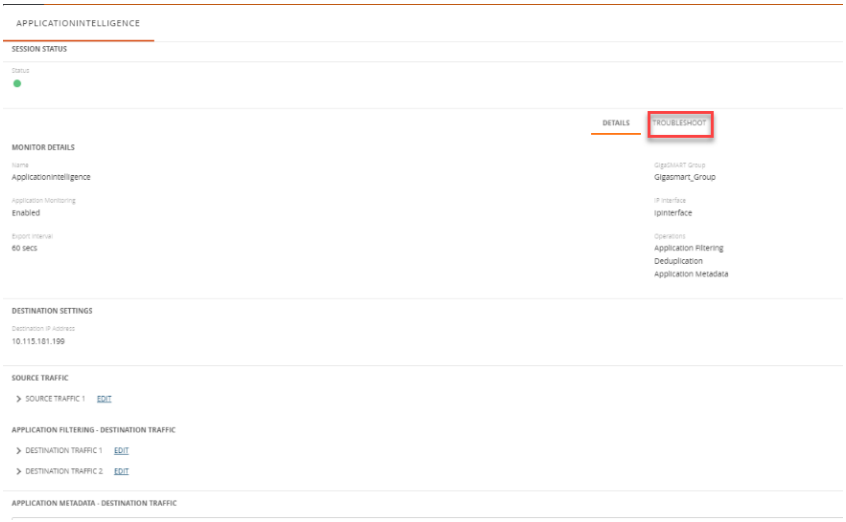

The solution page provides you the details of the health status of its associated components.

GigaSMART Application Session Filtering (ASF) and Buffer ASF

Required Licenses: Adaptive Packet Filtering (APF) and Application Session Filtering (ASF)

> **NOTE:** The ASF license requires the APF license to be installed as a prerequisite.

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 for more details on the devices that support GigaSMART operations.

DEPRECATION ANNOUNCEMENT: GigaSMART Application Session Filtering is end-of-sale in 5.6.00. Refer to the Application Intelligence to learn about the feature set and licensing for the improved application filtering and application metadata functionality.

Application Session Filtering (ASF) provides additional filtering on top of Adaptive Packet Filtering (APF). With APF, you can filter on any data patterns within a packet. With ASF, you apply the pattern matching and then send all the packet flows associated with the matched packet to monitoring or security tools.

ASF allows you to filter all traffic corresponding to a session. Use ASF to create a flow session and send the packets associated with the flow session to one or more tools. A flow session consists of one or more fields that you select to define the session. Either the packets for the whole session can be captured or only the packets following a pattern match.

A flow session is a session defined by protocol fields in the packet. For example, you can define a flow session with source and destination IP (two tuple), source and destination IP plus source and destination port (four tuple), or any combinations with inner or outer IP/port and protocol.

For example, use APF to filter TCP packets to capture the SYN packet. Then use ASF with GigaSMART Load Balancing to send all subsequent packets associated with the session to multiple tool ports. This example is illustrated in Example 1: ASF, Forward TCP Traffic. For information on capturing a whole session by buffering packets, refer to Application Session Filtering with Buffering.

Or use APF to create pattern-matching filters in which the pattern is a sequence of data bytes at a variable or fixed offset within a packet. Then use ASF with a specified session definition to capture subsequent packets belonging to the session. When an incoming packet matches an APF rule, a flow session is created. The subsequent incoming packets that match the value of the fields in the flow session will be forwarded to the same tool port as the matching packet.

For example, use APF to pattern match the string *www.gigamon.com*. Use the 5tuple field to identify the flow session. This creates the signature of the session. All the packets associated with the session will be forwarded to a tool port, hence APF becomes flow-aware or session-aware.

ASF provides the following session capabilities:

- filter on one, two, or both MPLS labels and/or VLAN IDs
- filter on both inner and outer IP addresses, Layer 4 ports, and protocols

Pattern matching examples are illustrated in Pattern Match with Type String to Example 4: ASF, Forward GTP Traffic

For information on load balancing, refer to stateful load balancing in the section GigaSMART Load Balancing.

ASF operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to Groups of GigaSMART Engine Ports for details.

In ASF and buffer ASF second level maps, a maximum of five (5) maps can be attached to a virtual port (vport). Each map can contain up to 25 gsrules.

Application Session Filtering (with or without buffering) is a pillar of the GigaSECURE Security Delivery Platform. Refer to .

Session-Aware APF (SAPF) Renaming and Licensing Change to ASF

The GigaSMART feature named Session-Aware Adaptive Packet Filtering (SAPF) in GigaVUE-OS 4.3 was renamed to Application Session Filtering (ASF) in GigaVUE-OS 4.4. It is now encompassed within ASF as the non-buffering equivalent to Application Session Filtering with buffering.

In addition, the license has moved from the APF license to the ASF license.

Application Session Filtering with Buffering

ASF captures packets of a session after an APF rule match. When the APF match occurs in the middle of a session, packets in the session prior to the match are not captured. With some tools needing all the packets of a flow session in order to perform data analysis, GigaSMART uses buffering to ensure that all packets belonging to a flow session are captured and forwarded to the tools. This is referred to as Application Session Filtering with buffering, or buffer ASF.

Buffer ASF uses the pattern-matching and regular expression engine in APF to select packet flows based on matching criteria with one or more packets in the flow session. Buffering ensures that the entire session, from start to finish, is either dropped or forwarded to the security tools or the performance monitoring tools.

To capture all packets belonging to a flow session prior to the APF rule match, ASF needs to know the first packet of a flow session. For this, GigaSMART supports both TCP and UDP connections.

For TCP connections, the TCP SYN packet is used to indicate the start of a session. GigaSMART captures and stores (or buffers) all packets of a flow session until an APF match occurs. After that, GigaSMART will either forward or drop all stored packets belonging to that session based on the APF pass or drop rule that is configured. Subsequent packets after the APF match will be forwarded or dropped as they arrive.

For UDP connections, there is no special packet that indicates the start of a UDP flow session from a Layer 4 perspective. GigaSMART will take the first UDP packet of a session it

encounters as the start of a session flow. This may result in incomplete capture at the beginning of configuration or at system boot up, but as new UDP sessions arrive, GigaSMART will capture the first packet of the flows.

ASF and Buffer ASF Session Definitions

Use the **Session Field** of the **ASF** page to define ASF and buffer ASF sessions by specifying session field attributes to add or delete. A session field is a group of one or more fields that define a flow session. (From the device view, select **GigaSMART > ASF** and click **New** to open the ASF page.)

A flow session consists of field names and attributes. Some field names include multiple attributes, which provide a quick way to define sessions.

The field names and attributes are as follows:

- gtpu-teid
- ipv4 (ipv4-src, ipv4-dst)
- ipv4-5tuple (ipv4-src, ipv4-dst, l4port-src, l4port-dst, ipv4-protocol)
- ipv4-dst
- ipv4-l4port-dst (ipv4-src, ipv4-dst, l4port-dst)
- ipv4-protocol
- ipv4-src
- ipv4-src-l4port-dst (ipv4-src, l4port-dst)
- ipv6 (ipv6-src, ipv6-dst)
- ipv6-5tuple (ipv6-src, ipv6-dst, l4port-src, l4port-dst, ipv6-protocol)
- ipv6-dst
- ipv6-l4port-dst (ipv6-src, ipv6-dst, l4port-dst)
- ipv6-protocol
- ipv6-src
- ipv6-src-l4port-dst (ipv6-src, l4port-dst)
- l4port (l4port-src, l4port-dst)
- l4port-dst
- l4port-src
- mpls-label
- vlan-id

An ASF session definition consists of combinations of the fields and attributes in the list above. In addition, for all IP and L4 port fields in the packet, each ASF session field must specify **outer** or **inner** for location. Outer specifies the first IP or L4 port in the packet. Inner specifies the second IP or L4 port in the packet (usually inside tunneling). For VLAN ID and MPLS label fields, a position (1 or 2) must be specified. Position 1 is the first occurrence of the protocol header or field in the packet. Position 2 is the second occurrence of the protocol header or field in the packet.

A buffer ASF session definition consists of combinations of the fields and attributes in the list above. **One restriction is that ipv4-src or ipv6-src needs to be defined, as a minimum.** In addition, the following restrictions apply to buffer ASF session definitions:

- the gtpu-teid field name is not supported
- the IP and L4 port fields only support location **outer**
- the VLAN ID field only supports position 1

All packets belonging to the same source and destination IP will be considered as the same flow session. This is useful if you want to capture all packets belonging to separate TCP/UDP connections that have the same IPs, such as control or data flows.

Define ASF Session

When defining an ASF session, enter the fields, attributes, and options, then save. The changes only take effect when after you save. For example, in Figure 2An ASF Configuration, the settings are:

- Alias is asf2, which is name of the ASF Profile for the GigaSMART Operation
- buffer enabled
- Buffer Count before Match is 5
- ipv4-5tuple outer
- vlan-id position 1

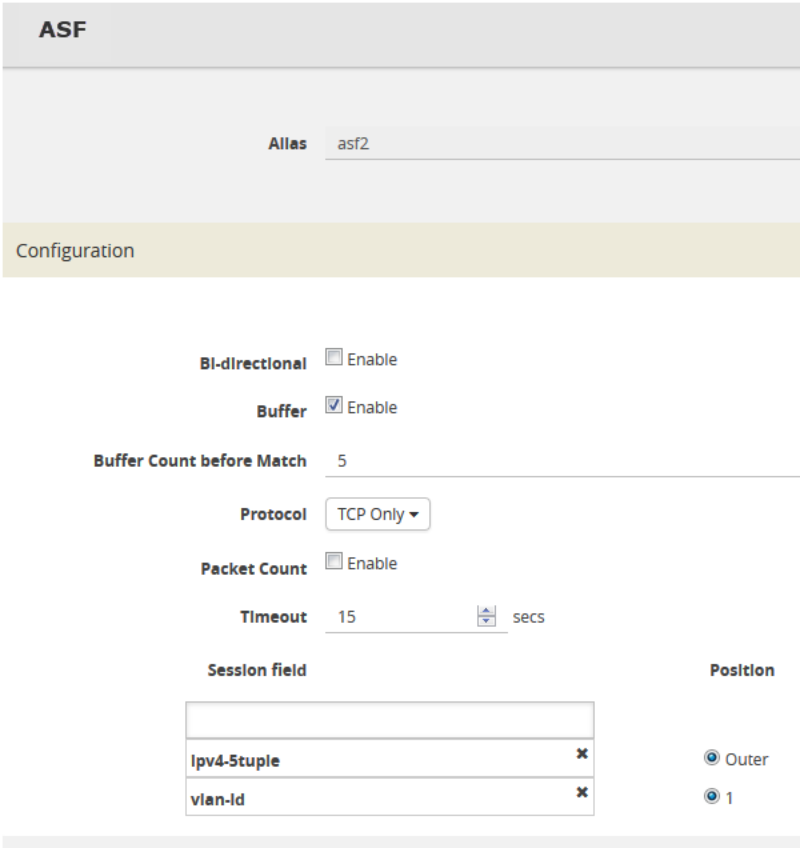

Figure 2 *An ASF Configuration*

Quick Session Delete for Buffer ASF

For a buffer ASF session defined with ipv4-5tuple or ipv6-5tuple, there is a quick session delete for TCP connections. The session is deleted 4 seconds after RST or both FIN packets are detected, signaling the end of the flow.

Specify Resources for Buffer ASF

On GigaVUE-HC3, and GigaVUE-HC1 nodes buffer ASF supports from 2 to 5 million sessions.

The large number of sessions can use a lot of memory resources on the 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.

The GigaSMART line card or module does not need to be reloaded the first time buffer ASF resources are allocated. But subsequent changes, such as increasing from 3 million to 4 million sessions, or decreasing from 3 million to 2 million sessions will require a reload for the changes to take effect.

You can reload the card from the UI, by doing the following:

1. On the left navigation pane, click  , and then select **Physical > Nodes**.
2. From the left navigation pane, go to **System > Chassis**.
3. Clicking the Table View button to open the Table View of the Chassis.
4. Under Cards, select the card to reload.
5. From the **Actions** menu, select **Shut Down**.
6. Now select **Start Up** from **Actions** menu.

Alternatively, you can use the following GigaVUE-FM API to reload the card:

PATCH /inventory/chassis/cards/{slotId}

For more information about the GigaVUE-FM APIs, refer to the *GigaVUE-FM REST API Reference in GigaVUE-FM User's Guide*.

You can also reload the card with the following two CLI commands:

(config) # card slot <slot ID> down

(config) # no card slot <slot ID> down

For more information about the CLI commands, refer the *GigaVUE-OS CLI Reference Guide*.

ASF and Buffer ASF Session Notes

The following table summarizes notes relating to ASF and buffer ASF:

Notes
A session field can only be modified or deleted if it is *not* configured in any GigaSMART operation.
A session field can only contain the same session attribute and position pair once.
A session field cannot contain overlapped session attribute and position. For example, the following is not valid: ipv4-5tuple outer and ipv4-src outer.
Up to a maximum of 25 flow session aliases are supported for ASF.
A total of 4 session tables per GigaSMART engine are supported for ASF. Each table has its own session definition.
Up to 2 million session entries are supported for ASF. The entries are shared by all session aliases.

Notes
Each session table (session alias) can only be used once within a gsgroup.
The number of buffer ASF sessions supported is configurable from 2 to 5 million.
The number of packet buffers supported for buffer ASF is from 2 to 5 million.

Buffer ASF Packet Processing Special Cases

The following are special cases of packet processing for buffer ASF:

- Non-TCP SYN packet received and no session matched—When a non-TCP SYN packet is received and there is no session matched, the packet will be considered as *no match* and will be passed to other maps. If there is no match after all maps have been processed, the packet will be forwarded to a shared collector, if one is configured.
- Out of session—When a TCP SYN packet is received and no free session is available, the packet will be considered as *no match*. Other packets belonging to this session will also be considered as *no match*, as for the special case described above.
- Out of packet storage buffer—When the buffer is full for the first packet of a session, the session will not be created and the packet will be considered as *no match*. When the buffer is full for an existing session, and the APF match has not yet occurred, a flag will be set and the current packet and all buffered packets will be considered as *no match*. Subsequent packets will also be considered as *no match*.
- Exceeded configured buffering limit—When there is no APF match after the configured number of packets have been buffered, all buffered packets and all subsequent packets belonging to this session will be considered as *no match*.

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 ASF or buffer ASF map, all matching rules will have that packet counted against them and the overall map status pass counter will show 1.

ASF and Buffer ASF Examples

Refer to the following ASF examples (non-buffered):

- Example 1: ASF, Forward TCP Traffic
- Pattern Match with Type String
- Example 3: ASF, Forward Traffic Matching a Pattern

- Example 4: ASF, Forward GTP Traffic

Refer to the following buffer ASF examples:

- Example 1: Buffer ASF, Drop Netflix Traffic
- Example 2: Buffer ASF, Drop YouTube Traffic
- Example 3: Buffer ASF, Drop Windows Update Traffic
- Example 4: Buffer ASF, Forward VNC Traffic
- Example 5: Buffer ASF, Forward HTTPS Traffic on Non-Standard Port

Example 1: ASF, Forward TCP Traffic

In Example 1, ASF is used with GigaSMART Load Balancing and Adaptive Packet Filtering to load balance TCP traffic among multiple tool ports. TCP SYN indicates the start of a connection. Once the TCP SYN packet is detected, subsequent packets belonging to the same TCP connection will be forwarded to a configured tool port. Packets belonging to the same connection will be sent to the same tool port, regardless of the number of connections.

NOTE: This example uses APF to filter TCP packets to capture the SYN packet. Alternatively, use buffer ASF to capture a whole session by buffering packets.

Task	Description	UI Steps
1	Create a flow session.	**1.** From the device view, select GigaSMART > ASF. **2.** Click New. **3.** Type asf4 in the Alias field. **4.** Select ipv4-tuple from the Session field list. **5.** Select outer. **6.** Click Save.
2	Create a port group and specify the tool ports for load balancing.	**1.** Select Ports > Port Groups > All Port Groups. **2.** Click New. **3.** Type portgrp1 in the Alias field. **4.** Select Tool. **5.** Select SMART Load Balancing. **6.** Click in the Ports field and select the tool ports. For example, 1/1/x6,1/1/x7,1/2/x3, and 1/2/x4.
3	Configure a GigaSMART group and associate it with GigaSMART engine ports.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsgrp1 in the Alias field.

Task	Description	UI Steps
		4. Select two engine ports from the Port List field. For example, 1/3/e1 and 1/3/e2 **5.** Click Save.
4	Configure the combined GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation and create two GigaSMART Operations. **2.** Click New. **3.** Type gsop1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Select the operations. o APF o ASF with asf4 for the ASF profile o Load Balancing with Stateful Type ASF, and Round Robin **6.** Click Save.
5	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.
6	Create a first level map.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. **a.** Type map11 in the Alias field. **b.** Select First Level for Type. **c.** Select By Rule for Subtype. **d.** Select the network port 1/1/x1 for the Source. **e.** Select the virtual port vp1 for the Destination. **6.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Version and set version to 4 **4.** Click Save.
7	Create a second level map. The gsrule captures the first packet of a session.	**1.** Click New. **2.** Configure the map. **a.** Type map22 in the Alias field. **b.** Select Second Level for Type. **c.** Select By Rule for Subtype.

Task	Description	UI Steps
		d. Select the virtual port vp1 for the Source. **e.** Select the port group portgrp1 for the Destination. **f.** Select gsop1 form the GSOP list. **7.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select TCP Control. **d.** Enter 2 for Value. **e.** Enter 0 for Mask. **f.** Set Position to 1. **7.** Click Save.

Example 2: ASF, Forward VNC Traffic

In Example 2, traffic from a Virtual Network Computing (VNC) application is forwarded from network port 1/1/x1 to tool port 1/1/x6. Packets will be matched with a VNC signature. Once a packet is matched, subsequent packets with the same IPv4 5tuple will be forwarded to the same destination as the matching packet. By default, both the forward and the reverse traffic of the same session will be captured and forwarded.

Step	Description	Command
1	Create a flow session.	**1.** From the device view, select GigaSMART > ASF. **2.** Click New. **3.** Type asf1 in the Alias field. **4.** Select ipv4-tuple from the Session field list. **5.** Select outer. **6.** Click Save.
2	Configure a GigaSMART group and associate it with GigaSMART engine ports.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsgrp1 in the Alias field. **4.** Select two engine ports from the Port List field. For example, 1/3/e1 and 1/3/e2 **5.** Click Save.
3	Configure the combined GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations and create

Step	Description	Command
		two GigaSMART Operations. **2.** Click New. **3.** Type gsop1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Select the operations. o APF o ASF with asf1 for the ASF profile **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 vp1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. **a.** Type map11 in the Alias field. **b.** Select First Level for Type. **c.** Select By Rule for Subtype. **d.** Select the network port 1/1/x1 for the Source. **e.** Select the virtual port vp1 for the Destination. **6.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Version and set Version to 4 **4.** Click Save.
6	Create a second level egress map. The gsrule contains the VNC signature.	**1.** Click New. **2.** Configure the map. **a.** Type map22 in the Alias field. **b.** Select Second Level for Type. **c.** Select By Rule for Subtype. **d.** Select the virtual port vp1 for the Source. **e.** Select the port group portgrp1 for the Destination. **f.** Select gsop1 form the GSOP list. **7.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass.

Step	Description	Command
		c. Select Pattern Match. d. Select regex for Type and enter ^rfb 00[1-9]\.00[0-9]\x0a$ e. Set Offset from 16 to 1000 6. Click Save.

Example 3: ASF, Forward Traffic Matching a Pattern

In Example 3, the traffic that matches a particular pattern (ymsg|ypns|yhoo) is forwarded from network port 1/1/x1 to tool port 1/1/x6 after adding a VLAN tag. Packets will be matched with the special signature. Once a packet is matched, subsequent packets with the same source IP, source port, and VLAN ID will be forwarded to the same destination as the matching packet (after the VLAN header is inserted). By default, both the forward and the reverse traffic of the same session will be captured and forwarded.

Task	Description	UI Steps
1	Create a flow session and other parameters.	1. From the device view, select GigaSMART > ASF. 2. Click New. 3. Type asf2 in the Alias field. 4. Enable Packet Count. 5. Set Number of packets to 50. 6. Set the session field. o Select ipv4-src outer o Select vlan-id position 1 7. Select outer. 8. Click Save.
2	Configure a GigaSMART group and associate it with GigaSMART engine ports.	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type gsgrp1 in the Alias field. 4. Select two engine ports from the Port List field. For example, 1/3/e1 and 1/3/e2 5. Click Save.
3	Configure the GigaSMART operation.	1. From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation and create two GigaSMART Operations. 2. Click New.

Task	Description	UI Steps
		3. Type gsop1 in the Alias field. **4.** Select gsgrp1 from the GS Groups list. **5.** Select the operations. o Adaptive Packet Filtering o Add Header and set VLAN to 1000 o ASF with asf2 for the ASF profile **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 vp1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. **a.** Type map11 in the Alias field. **b.** Select First Level for Type. **c.** Select By Rule for Subtype. **d.** Select the network port 1/1/x1 for the Source. **e.** Select the virtual port vp1 for the Destination. **6.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Version and set Version to 4 **4.** Click Save.
6	Create a second level map. The gsrule contains the special signature.	**1.** Click New. **2.** Configure the map. **a.** Type map22 in the Alias field. **b.** Select Second Level for Type. **c.** Select By Rule for Subtype. **d.** Select the virtual port vp1 for the Source. **e.** Select the too port 1/1/x6 for the Destination. **f.** Select gsop1 form the GSOP list. **7.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Pattern Match.

Task	Description	UI Steps
		d. Select regex for Type and enter (ymsg|ypns|yhoo) **e.** Set Offsett from 16 to 1000 **6.** Click Save.

Example 4: ASF, Forward GTP Traffic

In Example 4, GTP traffic from network port 1/1/x1 is load balanced based on inner IP and tunnel ID to four tool ports: 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4. APF filters GTP-u packets. Once a packet is matched, subsequent packets in the same direction with the same gtpu-teid and inner IP will be forwarded to the same destination as the matching packet. In Example 4, both the outer and inner IP are IPv4.

Task	Description	UI Step
1	Create a flow session and other parameters.	**1.** From the device view, select GigaSMART > ASF. **2.** Click New. **3.** Type asf3 in the Alias field. **4.** Set timeout to 90. **5.** Set the session field. o Select gtpu-teid o Select Ipv4 inner **6.** Select outer. **7.** Click Save.
2	Create a port group and specify the tool ports for load balancing.	**1.** From the left navigation pane, go to **System > Ports > Ports > Port Groups > All Port Groups.** **2.** Click New. **3.** Type portgrp1 in the Alias field. **4.** Select Tool. **5.** Select SMART Load Balancing. **6.** Click in the Ports field and select the tool ports. For example, 1/1/x6,1/1/x7,1/2/x3, and 1/2/x4.
3	Configure a GigaSMART group and associate it with GigaSMART engine ports.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsgrp1 in the Alias field. **4.** Select two engine ports from the Port List field. For example, 1/3/e1 and 1/3/e2

Task	Description	UI Step
		5. Click Save.
4	Configure the combined GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOPS) > GigaSMART Operation and create two GigaSMART Operations. **2.** Click New. **3.** Type gsop1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Select the operations. o Adaptive Packet Filtering o ASF with asf3 for the ASF profile o Load Balancing with Stateful, Type ASF, and Least Conn **6.** Click Save.
5	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.
6	Create a first level map.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. **a.** Type map11 in the Alias field. **b.** Select First Level for Type. **c.** Select By Rule for Subtype. **d.** Select the network port 1/1/x1 for the Source. **e.** Select the virtual port vp1 for the Destination. **6.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Protocol and set Value to UDP. **d.** Select Port Destination and set the port value to 2152 **5.** Click Save.
7	Create a second level map.	**1.** Click New. **2.** Configure the map. **a.** Type map22 in the Alias field. **b.** Select Second Level for Type. **c.** Select By Rule for Subtype. **d.** Select the virtual port vp1 for the Source.

Task	Description	UI Step
		e. Select the port group **portgrp1** for the Destination. f. Select gsop1 form the GSOP list. **7.** Add a rule. a. Click Add a Rule. b. Select Pass. c. Select Pattern Match. d. Select IPv4 Protocol and enter the IPv4 address. Set Position to 1. e. Select Ipv4 Destination and set the port value to 2152. Set Position to 1. **6.** Click Save.

Example 1: Buffer ASF, Drop Netflix Traffic

In Example 1, the goal is to drop all Netflix traffic. The flow session is defined by the 5tuple field and the first occurrence of VLAN ID. The Netflix traffic is expected to be identified in the first 6 packets of a session. (Configure the maximum number of packets buffered before the match to 5.) A maximum of 3 million sessions is specified.

Task	Description	UI Steps
1	Configure a GigaSMART group and associate it with GigaSMART engine ports and Define the maximum number of sessions, in millions.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsgrp1 in the Alias field. **4.** Select two engine ports from the Port List field. For example, 1/3/e1 and 1/3/e2 **5.** Under Params Resource Buffer select ASF and set the Buffer size to 3. **6.** Click Save.
2	If needed, reload the GigaSMART line card or module to allocate the resources for buffer ASF.	If you reset the buffer size of an ASF profile in Task 1, go to the Chassis page and select Table View. Under Cards, select the card to reload. From the Actions menu, select Shut Down and then Start Up. You can also issue the following CLI commands to reboot the card (the card is in slot 3 in this example): **(config) # card slot 3 down** **(config) # no card slot 3 down**

Task	Description	UI Steps
3	Create a flow session, specify the buffer count before the match, and enable buffering. **NOTE:** The default protocol is TCP, so it does not need to be specified.	**1.** From the device view, select GigaSMART > ASF. **2.** Click New or select an existing ASF profile then click Edit. **3.** Type asf2 in the Alias field if this is a new ASF profile. **4.** Enable Buffer. **5.** Set Buffer Count before Match to 5. **6.** Set the session field. o Select ipv4-5tuple outer o Select vlan-id position 1 **7.** Click Save.
4	Configure the combined GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation and create two GigaSMART Operations. **2.** Click New. **3.** Select gsgrp1 from the GigaSMART Groups list. **4.** Type gsop1 in the Alias field. **5.** Select the operations. o APF o ASF with asf2 for the ASF profile
5	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.
6	Create a first level map.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. **a.** Type map11 in the Alias field. **b.** Select First Level for Type. **c.** Select By Rule for Subtype. **d.** Select the network port 1/1/x1 for the Source. **e.** Select the virtual port vp1 for the Destination. **6.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Protocol and set Value to UDP. **d.** Select Port Destination and set the port value to 2152 **5.** Click Save.

Task	Description	UI Steps
7	Create a second level map. The gsrule specifies the traffic to drop, using keywords. Buffered packets and all subsequent packets will be dropped.	**1.** Click New. **2.** Configure the map. **a.** Type map22 in the Alias field. **b.** Select Second Level for Type. **c.** Select By Rule for Subtype. **d.** Select the virtual port vp1 for the Source. **e.** Select the tool port 1/1/x6 for the Destination. **f.** Select gsop1 form the GSOP list. **7.** Add a rule. **a.** Click Add a Rule. **b.** Select Drop. **c.** Select Pattern Match. **d.** Select regex and enter netflix|nflxvideo|nflximg|Netflix|nflxext. **e.** Set the offset from 0 to 1460 **f.** Set Protocol to tcp and set Position to 1. **7.** Click Save.

Example 2: Buffer ASF, Drop YouTube Traffic

In Example 2, the goal is to drop all YouTube traffic. The YouTube traffic is expected to be identified in the first 7 packets of a session. (Configure the maximum number of packets buffered before the match to 6.) A maximum of 4 million sessions is specified.

Step	Description	Command
1	Configure a GigaSMART group and associate it with GigaSMART engine ports and define the maximum number of sessions, in millions	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsgrp1 in the Alias field. **4.** Select two engine ports from the Port List field. For example, 1/3/e1 and 1/3/e2 **5.** Under Params Resource Buffer, select ASF and set the Buffer Size to 4. **6.** Click Save.
2	If needed, reload the GigaSMART line card or module to allocate the resources for buffer ASF.	If you reset the buffer size of an ASF profile in Task 1, go to the Chassis page and select Table View. Under Cards, select the card to reload. From the Actions menu select Shut Down and

Step	Description	Command
		then Start Up. You can also issue the following CLI commands to reboot the card (the card is in slot 3 in this example): **(config) # card slot 3 down** **(config) # no card slot 3 down**
3	Create a flow session, specify the buffer count before the match, and enable buffering. **NOTE:** The default protocol is TCP, so it does not need to be specified.	**1.** From the device view, select GigaSMART > ASF. **2.** Click New or select an existing ASF profile then click Edit. **3.** Type asf2 in the Alias field if this is a new ASF profile. **4.** Enable Buffer. **5.** Set Buffer Count before Match to 6. **6.** Set the session field to ipv4-5tuple outer **7.** Click Save.
4	Configure the combined GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations and create two GigaSMART Operations. **2.** Click New. **3.** Type gsop1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Select the operations. o Adaptive Packet Filtering o ASF with asf2 for the ASF profile **6.** Click Save.
5	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.
6	Create a first level map.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. **a.** Type map11 in the Alias field. **b.** Select First Level for Type. **c.** Select By Rule for Subtype. **d.** Select the network port 1/1/x1 for the Source. **e.** Select the virtual port vp1 for the Destination. **6.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass.

Step	Description	Command
		c. Select IPv4 Version and set Version to v4. 4. Click Save.
7	Create a second level map. The gsrule specifies the traffic to drop, using keywords. Buffered packets and all subsequent packets will be dropped.	1. Click New. 2. Configure the map. a. Type map22 in the Alias field. b. Select Second Level for Type. c. Select By Rule for Subtype. d. Select the virtual port vp1 for the Source. e. Select the tool port 1/1/x6 for the Destination. f. Select gsop1 form the GSOP list. 7. Add a rule. a. Click Add a Rule. b. Select Drop. c. Select Pattern Match. d. Select regex and enter youtube|ytimg|yt3.ggpht|tubeMogul|tmogul. e. Set the offset from 0 to 1460 f. Set Protocol to tcp and set Position to 1. 7. Click Save.

Example 3: Buffer ASF, Drop Windows Update Traffic

In Example 3, the goal is to drop all Windows update traffic. The Windows update traffic is expected to be identified on the HTTP request packet of a session. A maximum of 2 million sessions is specified.

Task	Description	UI Steps
1	Configure a GigaSMART group and associate it with GigaSMART engine ports and define the maximum number of sessions, in millions.	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type gsgrp1 in the Alias field. 4. Select two engine ports from the Port List field. For example, 1/3/e1 and 1/3/e2 5. Under Params Resource Buffer, select ASF and set the Buffer Size to 2. 6. Click Save.

Task	Description	UI Steps
2	If needed, reload the GigaSMART line card or module to allocate the resources for buffer ASF.	If you reset the buffer size of an ASF profile in Task 1, go to the Chassis page and select Table View. Under Cards, select the card to reload. From the Actions menu select Shut Down and then Start Up. You can also issue the following CLI commands to reboot the card (the card is in slot 3 in this example): **(config) # card slot 3 down** **(config) # no card slot 3 down**
3	Create a flow session, specify the buffer count before the match, and enable buffering. **NOTE:** The default protocol is TCP, so it does not need to be specified.	**1.** From the device view, select GigaSMART > ASF. **2.** Click New or select an existing ASF profile then click Edit. **3.** Type asf2 in the Alias field if this is a new ASF profile. **4.** Enable Buffer. **5.** Set Buffer Count before Match to 3. **6.** Set the session field to ipv4-5tuple outer **7.** Click Save.
4	Configure the combined GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations and create two GigaSMART Operations. **2.** Click New. **3.** Type gsop1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Select the operations. o Adaptive Packet Filtering o ASF with asf2 for the ASF profile **6.** Click Save.
5	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.
6	Create a first level map.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. **a.** Type map11 in the Alias field. **b.** Select First Level for Type. **c.** Select By Rule for Subtype. **d.** Select the network port 1/1/x1 for the Source. **e.** Select the virtual port vp1 for the Destination.

Task	Description	UI Steps
		6. Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Version and set Version to v4. **4.** Click Save.
7	Create a second level map. The gsrule specifies the traffic to drop. Buffered packets and all subsequent packets will be dropped.	**1.** Click New. **2.** Configure the map. **a.** Type map22 in the Alias field. **b.** Select Second Level for Type. **c.** Select By Rule for Subtype. **d.** Select the virtual port vp1 for the Source. **e.** Select the tool port 1/1/x6 for the Destination. **f.** Select gsop1 form the GSOP list. **7.** Add a rule. **a.** Click Add a Rule. **b.** Select Drop. **c.** Select Pattern Match. **d.** Select regex and enter msdownload/update/software. **e.** Set the offset from 0 to 1460 **f.** Set Protocol to tcp and set Position to 1. **7.** Click Save.

Example 4: Buffer ASF, Forward VNC Traffic

In Example 4, the goal is to forward VNC traffic from network port 1/1/x1 to tool port 1/1/x6. All packets belonging to the TCP connection need to be sent to the tool port. The first data packet after the TCP handshake is expected to contain the VNC pattern match. A maximum of 2 million sessions is specified.

Task	Description	UI Steps
1	Configure a GigaSMART group and associate it with GigaSMART engine ports and define the maximum number of sessions, in millions.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsgrp1 in the Alias field. **4.** Select two engine ports from the Port List field. For example, 1/3/e1 and 1/3/e2

Task	Description	UI Steps
		5. Under Params Resource Buffer, select ASF and set the Buffer Size to 2. **6.** Click Save.
2	If needed, reload the GigaSMART line card or module to allocate the resources for buffer ASF.	If you reset the buffer size of an ASF profile in Task 1, go to the Chassis page and select Table View. Under Cards, select the card to reload. From the Actions menu select Shut Down and then Start Up. You can also issue the following CLI commands to reboot the card (the card is in slot 3 in this example): **(config) # card slot 3 down** **(config) # no card slot 3 down**
3	Create a flow session, specify the buffer count before the match, and enable buffering. **NOTE:** The default protocol is TCP, so it does not need to be specified.	**1.** From the device view, select GigaSMART > Application Session Filtering. **2.** Click New or select an existing ASF profile then click Edit. **3.** Type asf1 in the Alias field if this is a new ASF profile. **4.** Enable Buffer. **5.** Set Buffer Count before Match to 3. **6.** Set the session field to ipv4-5tuple outer **7.** Click Save.
4	Configure the combined GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation and create two GigaSMART Operations. **2.** Click New. **3.** Type gsop1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Select the operations. o Adaptive Packet Filtering o ASF with asf1 for the ASF profile **6.** Click Save.
5	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.
6	Create a first level map.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. **a.** Type map11 in the Alias field.

Task	Description	UI Steps
		b. Select First Level for Type. **c.** Select By Rule for Subtype. **d.** Select the network port 1/1/x1 for the Source. **e.** Select the virtual port vp1 for the Destination. **6.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Version and set Version to v4. **4.** Click Save.
7	Create a second level map. The gsrule specifies the traffic to pass. Buffered packets and all subsequent packets will be passed.	**1.** Click New. **2.** Configure the map. **a.** Type map22 in the Alias field. **b.** Select Second Level for Type. **c.** Select By Rule for Subtype. **d.** Select the virtual port vp1 for the Source. **e.** Select the tool port 1/1/x6 for the Destination. **f.** Select gsop1 form the GigaSMART Operations (GSOP) list. **7.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Pattern Match. **d.** Select regex and enter ^rfb 00[1-9]\.00[0-9]\x0a$. **e.** Set Protocol to tcp and set Position to 1. **6.** Click Save.

Example 5: Buffer ASF, Forward HTTPS Traffic on Non-Standard Port

In Example 5, the goal is to forward HTTPS traffic that uses a non-standard Layer 4 port. All packets belonging to the TCP connection need to be sent to the tool port. A maximum of 5 million sessions is specified.

Task	Description	UI Steps
1	Configure a GigaSMART group and associate it with GigaSMART engine ports and define the maximum number of sessions, in millions.	**1.** GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsgrp1 in the Alias field. **4.** Select two engine ports from the Port List field. For

Task	Description	UI Steps
		example, 1/3/e1 and 1/3/e2 **5.** Under Params Resource Buffer, select ASF and set the Buffer Size to 2. **6.** Click Save.
2	If needed, reload the GigaSMART line card or module to allocate the resources for buffer ASF.	If you reset the buffer size of an ASF profile in Task 1, go to the Chassis page and select Table View. Select the card in the table. From the Actions menu select Shut Down and then Start Up. You can also issue the following commands to reboot the card (the card is in slot 3 in this example): **(config) # card slot 3 down** **(config) # no card slot 3 down**
3	Create a flow session, specify the buffer count before the match, and enable buffering. **NOTE:** The default protocol is TCP, so it does not need to be specified.	**1.** From the device view, select GigaSMART > ASF. **2.** Click New or select an existing ASF profile then click Edit. **3.** Type asf2 in the Alias field if this is a new ASF profile. **4.** Enable Buffer. **5.** Set Buffer Count before Match to 3. **6.** Set the session field to ipv4-5tuple outer **7.** Click Save.
4	Configure the combined GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations and create two GigaSMART Operations. **2.** Click New. **3.** Type gsop1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Select the operations. o APF o ASF with asf2 for the ASF profile **6.** Click Save.
5	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.
6	Create a first level map.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map.

Task	Description	UI Steps
		a. Type map11 in the Alias field. **b.** Select First Level for Type. **c.** Select By Rule for Subtype. **d.** Select the network port 1/1/x1 for the Source. **e.** Select the virtual port vp1 for the Destination. **6.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Version and set Version to v4. **4.** Click Save.
7	Create a second level map. The gsrule specifies the traffic to pass. The RegEx expression identifies the traffic as SSL. Buffered packets and all subsequent packets will be passed.	**1.** Click New. **2.** Configure the map. **a.** Type map22 in the Alias field. **b.** Select Second Level for Type. **c.** Select By Rule for Subtype. **d.** Select the virtual port vp1 for the Source. **e.** Select the tool port 1/1/x6 for the Destination. **f.** Select gsop1 form the GSOP list. **7.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Pattern Match. **d.** Select regex and enter x16\x03.{3}\x01. **e.** Set Protocol to tcp and set Position to 1. **6.** Click Save.

Display ASF Statistics

To display ASF statistics on the GigaSMART operation, select **GigaSMART > Statistics**.

Refer to ASF Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

Enhanced Application Session Filtering

GigaSMART Enhanced Application Session Filtering

Enhanced Application Session Filtering (ASF) allows you to filter a specific application field from the incoming traffic. Enhanced Application Session Filtering supports the following session capabilities:

- Applications SSL —Supports SNI field in a hello packet.
- HTTP—Supports HOST and User-agent field in a request packet

You can use the matching pattern in a regrex profile using the regular expression format. Rules are applied based on the order in which it was created.

When a rule matches, the corresponding pass or drop action is taken on the packets. In a session, when a packet matches a pass rule, the packet and its subsequent packets belonging to the same session are forwarded to and processed by the next GS application belong to the same GSOP.

When a packet of a session matches a pass rule, the packet and the subsequent packets belonging to the same session are forwarded to and processed by the next GigaSMART application that belongs to the same GSOP. If there is no GigaSMART application in the downstream application chain in the GSOP, the packet and the subsequent packets belonging to the same session are forwarded to the tool port defined in the Map.

When a packet of a session matches a drop rule, the packet and the subsequent packets belonging to the same session are dropped.

When a packet of a session does not match a rule, it will be examined by the next rule. When a packet does not match any rules defined in the Enhanced ASF profile, the packet will be forwarded to and processed by the GSOP configured in the next Map belonging to the same VPORT.

GigaVUE Enriched Metadata for Mobile Networks

The metadata enrichment enhances service provider analytics by generating metadata on 5G/4G/3G network traffic. The AMX correlates the user plane metadata produced by AMI with the control plane metadata produced by the GTP/5G correlation mobility application to produce an enriched metadata feed for the mobile networks. This data feed helps the customers with use cases like service personalization, planning, and many others by containing information about the

- Subscriber Session
- Over the Top Application
- Handset Type
- Location
- Flow throughput calculation attributes - DL, UL bytes and time stamp.
- Application Protocol
- Core Network Information
- User Tunnel Information

For information on Control Plane Metadata, refer Control Plane Metadata

For information on AMX, refer Configure Application Metadata Exporter Application.

The architecture of GigaVUE Enriched Metadata for Mobile Networks is as follows:

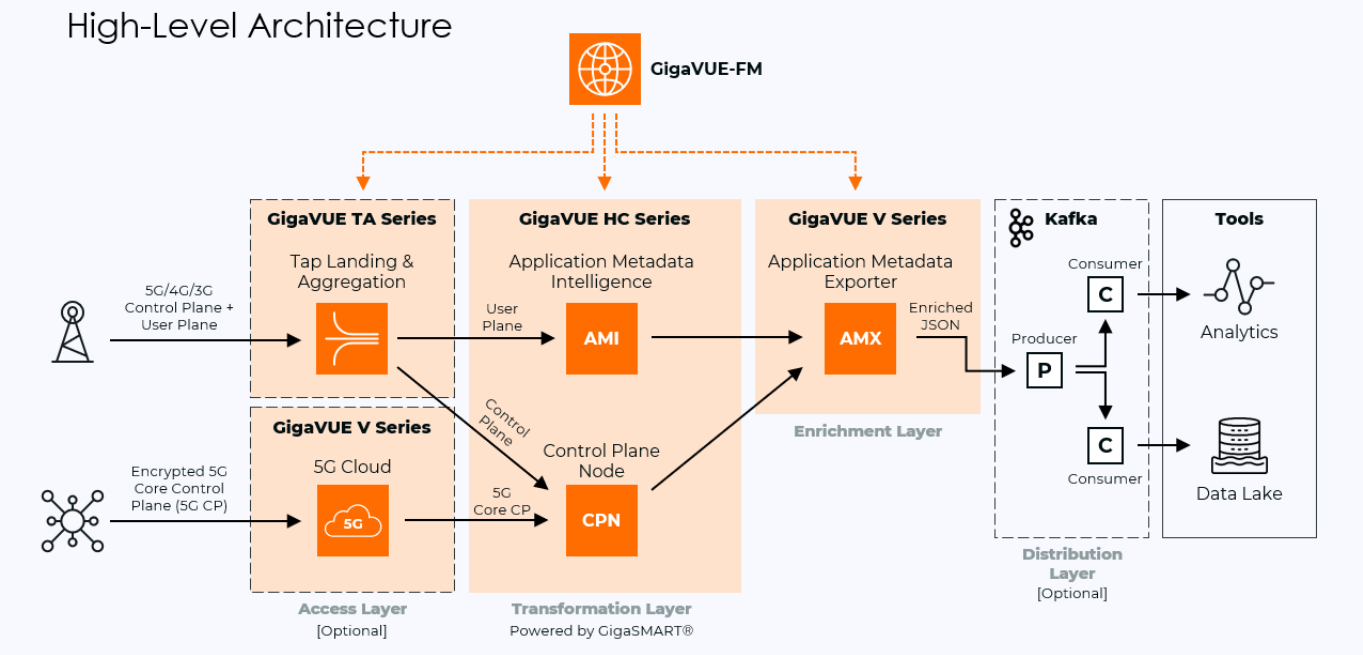

AMI Configurations

The following CLI commands show the configurations to be done before the attributes are exported:

```
gsgroup alias GSG_S1U_AMI_01 port-list 1/2/e2 hash advanced
apps metadata app-profile alias AMI_APP_PROFILE
application add id
counter add bytes size 4
counter add inner-bytes size 4
counter add packets size 4
flow add end-reason
ipv4 add destination prefix minimum-mask /32
ipv4 add protocol
ipv4 add source prefix minimum-mask /32
ipv6 add destination prefix minimum-mask /128
ipv6 add next-header
ipv6 add source prefix minimum-mask /128
timestamp add flow-end-msec
timestamp add flow-start-msec
transport add destination-port
transport add source-port
type export
exit
```

The following CLI commands show how frequently the attributes are exported and to which destination (destination IP address, destination port number) they are exported:

```
apps metadata exporter alias AMI_METADATA_EXPORTER
app-profile attach alias AMI_APP_PROFILE
cef timeout active 30
cef timeout inactive 15
destination dscp 0
destination ip ver4 172.31.251.1
destination l4 protocol udp
l4 port dest 30100
l4 port src 35100
max-packet-size disable
monitor timeout 60
netflow template-refresh-interval 60
netflow timeout active 60
netflow timeout inactive 15
netflow version ipfix
record-type cohesive
```

```
source ip-interface attach HC3-TO-AMX-ETH0
ttl 64
type cef
exit
```

The following CLI commands explain about the cache details and the flow behavior of AMI Metadata Exporter:

```
apps metadata cache alias AMI_APP_CACHE
advance-hash disable
aggregate-mode disable
dpi-inject-limit disable
events enable transaction-end
exporter attach alias AMI_METADATA_EXPORTER
flow-behaviour bi-direction
match ipv4 add destination prefix minimum-mask /32
match ipv4 add protocol
match ipv4 add source prefix minimum-mask /32
match ipv6 add destination prefix minimum-mask /128
match ipv6 add next-header
match ipv6 add source prefix minimum-mask /128
match transport add destination-port
match transport add source-port
multi-collect disable
observ-domain-id 0
size flow 10
timeout idle 300
exit
gsparams gsgroup GSG_S1U_AMI_01
resource metadata 10
metadata fastMode enable
exit
```

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.
```

S.No	Groups-Input
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 for 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

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 3GTP Version) or EPC interface Gn/GP, S5/S8, or S10, S2B, or S11/S1U (refer to Figure 4GTP 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.

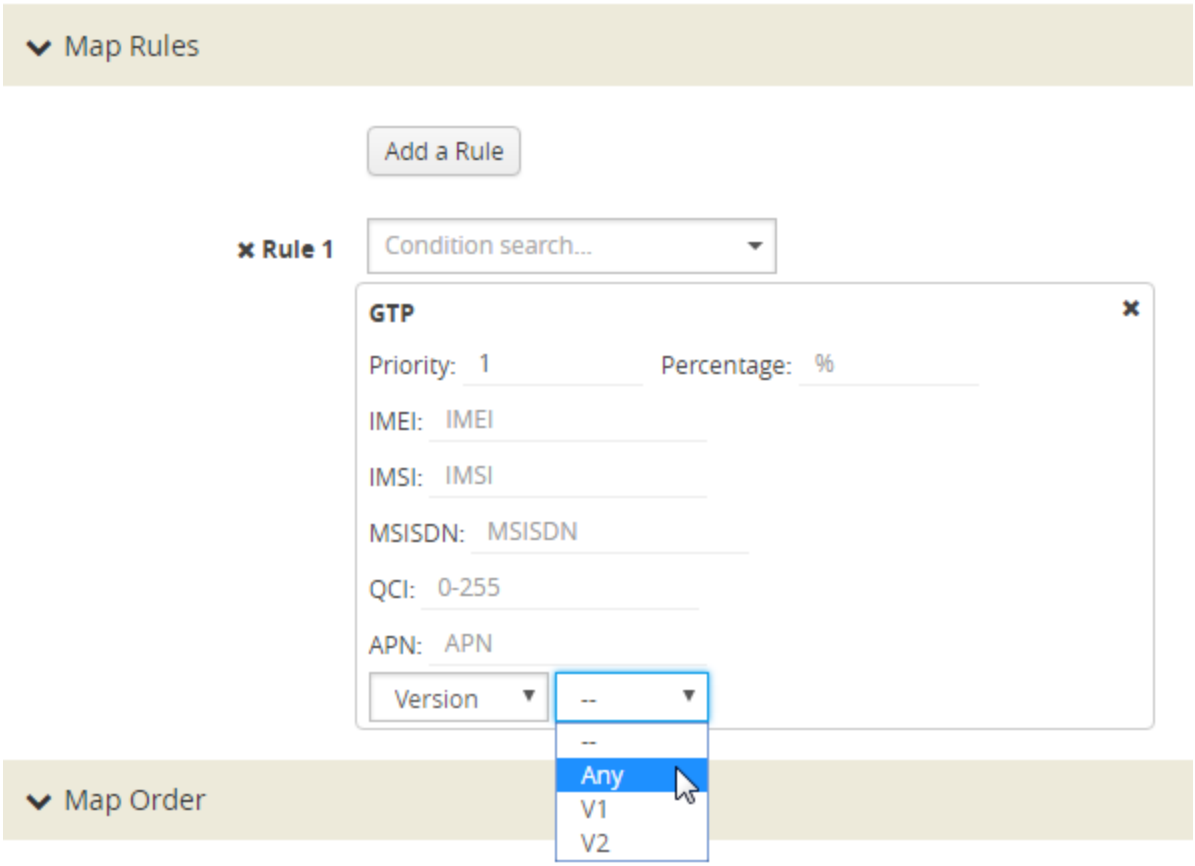

Figure 3 *GTP Version*

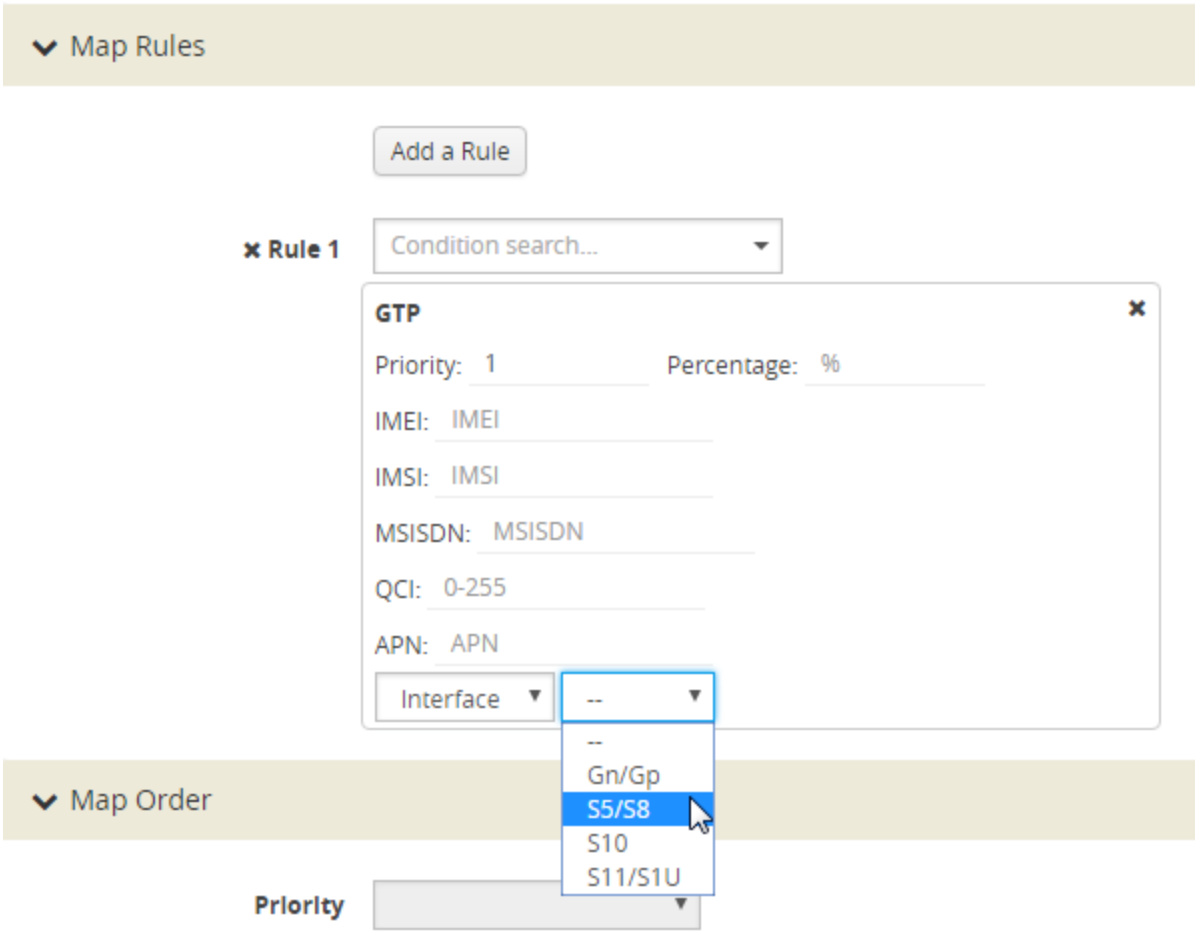

Figure 4 *GTP EPC Interface*

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.

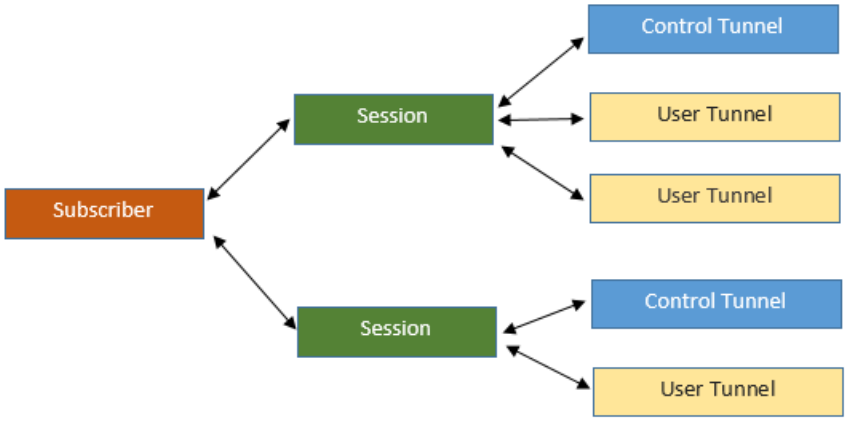

In a second level map, the following can be specified:

- 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.

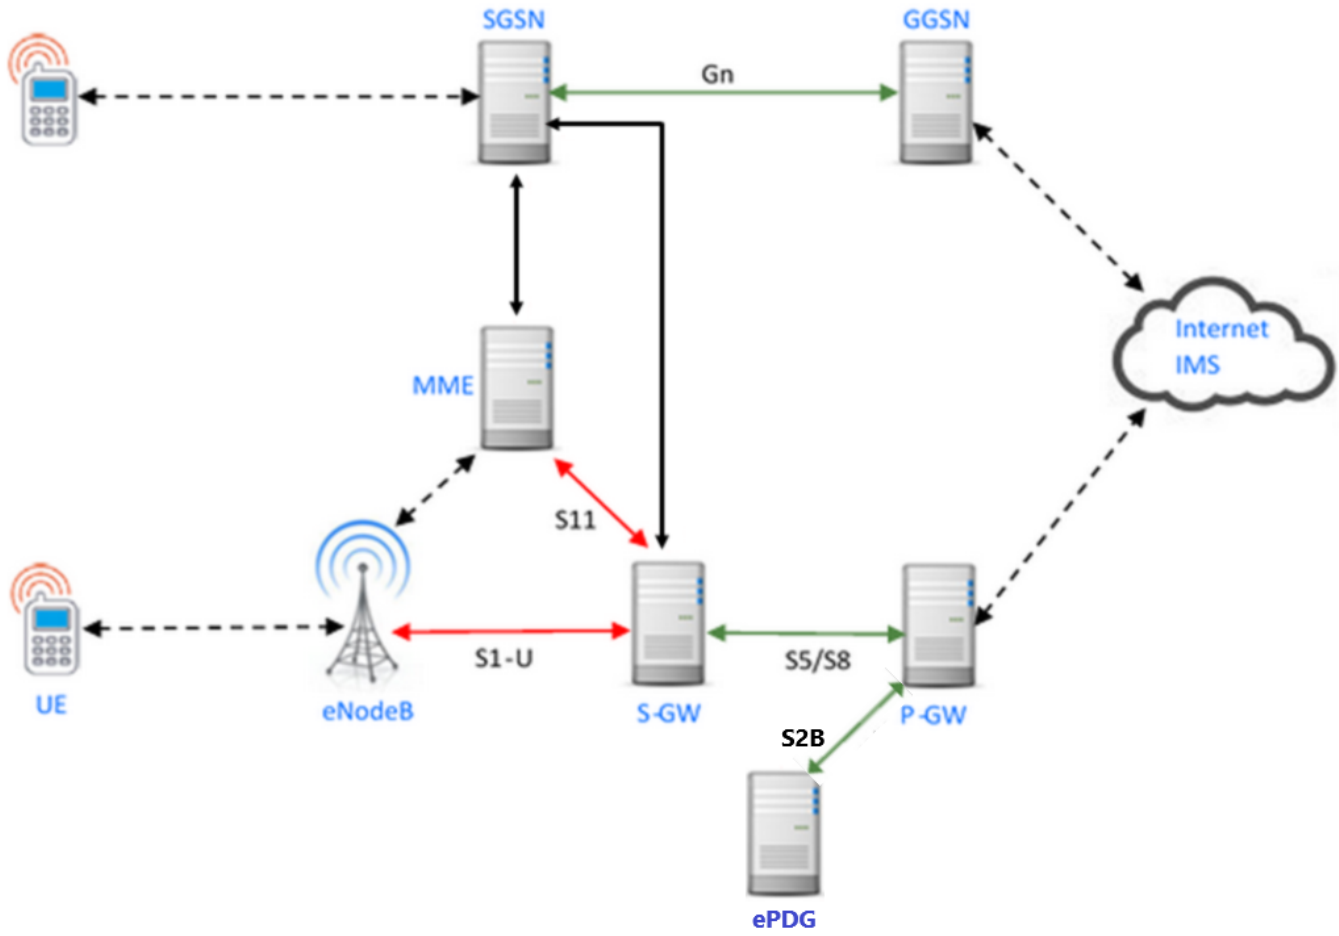

For LTE networks, the following GTP traffic will be correlated to the specific mobile subscriber and routed to the same tool port:

- 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.

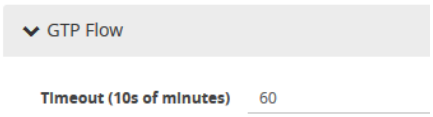

5. Click **Save**.

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:

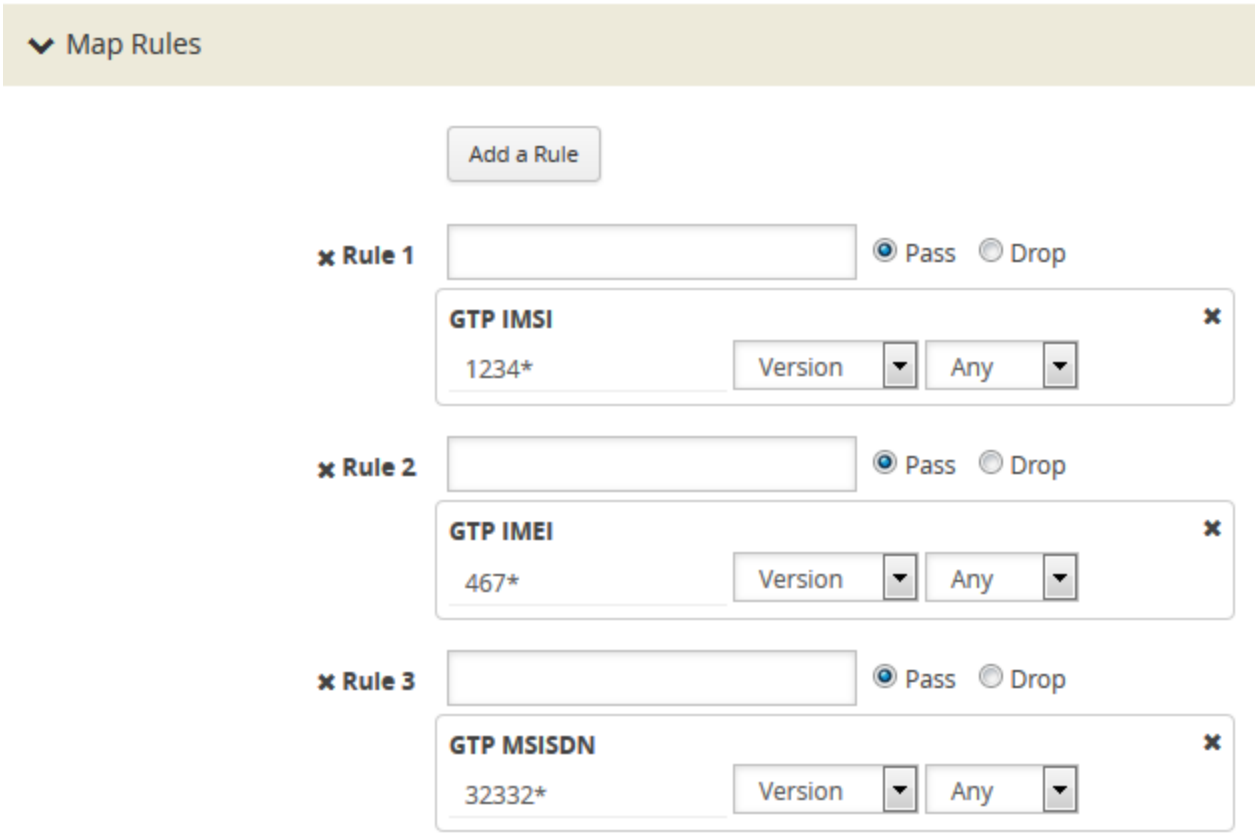

In addition, in one map, all drop rules are matched first and all pass rules are matched next.

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:

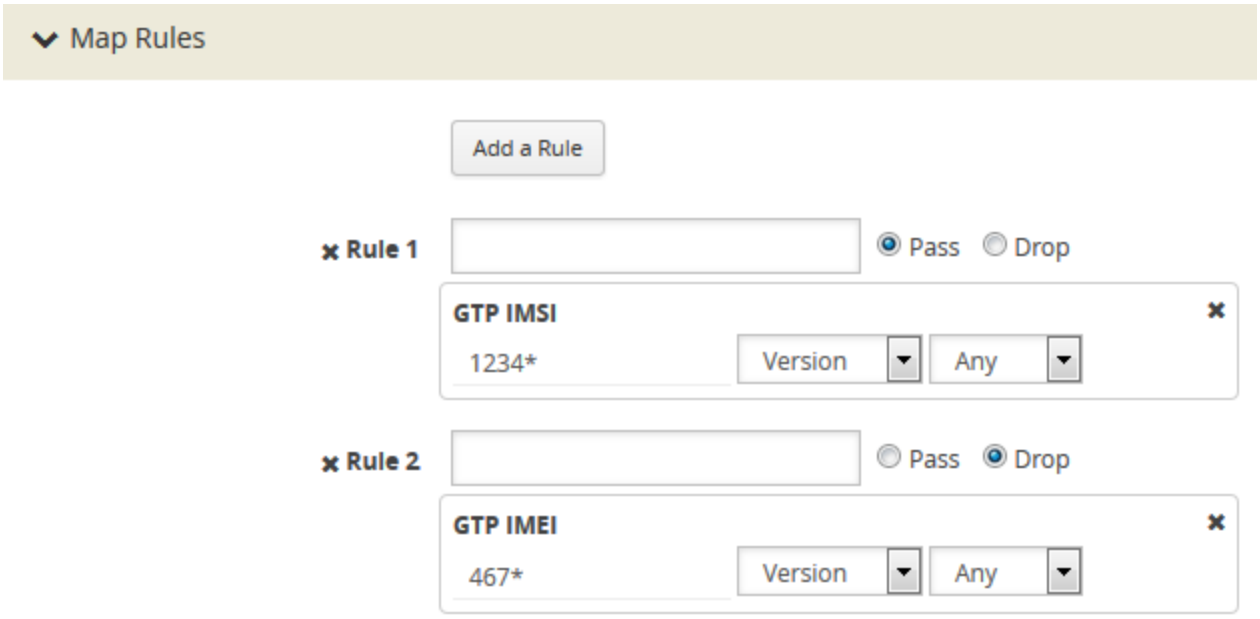

If multiple maps are matched, the map with the highest priority will be considered for further processing. For example, the Figure 5Rule in Map1 shows the rule for map1 while the rule Figure 117NetFlow Generation Gigamon Solution shows the rule for map2. In a GTP session, if an IMSI matches the first rule in map1 and an IMEI matches the first rule in map2, because map1 has higher priority, the packet will be passed.

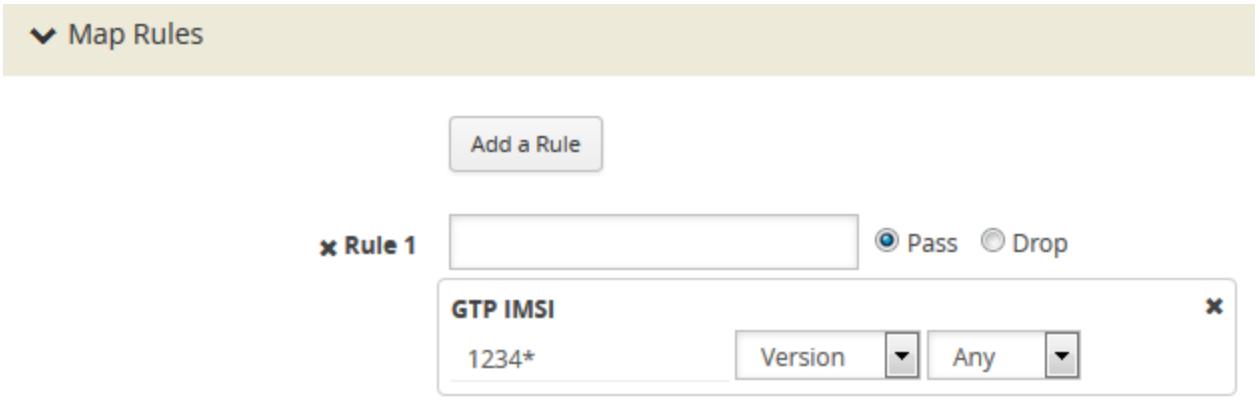

Figure 5 *Rule in Map1*



Figure 6 *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.

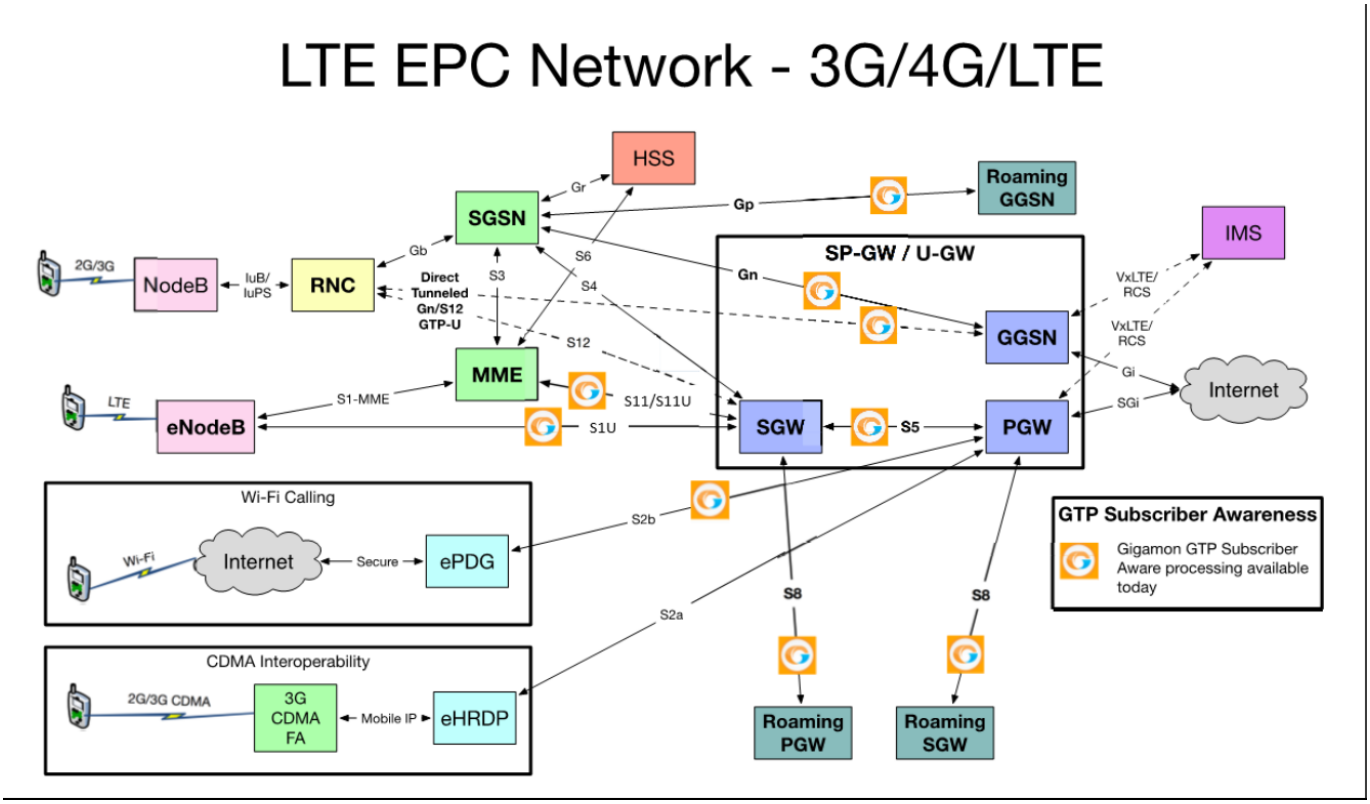

In Example 1, filter rules are configured to identify and forward all the traffic related to subscribers identified by an IMSI prefix. All traffic specific to the filtered IMSIs 22222222222223*, including GTP-c and GTP-u, is forwarded to a monitoring tool. A shared collector is configured to which traffic 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.

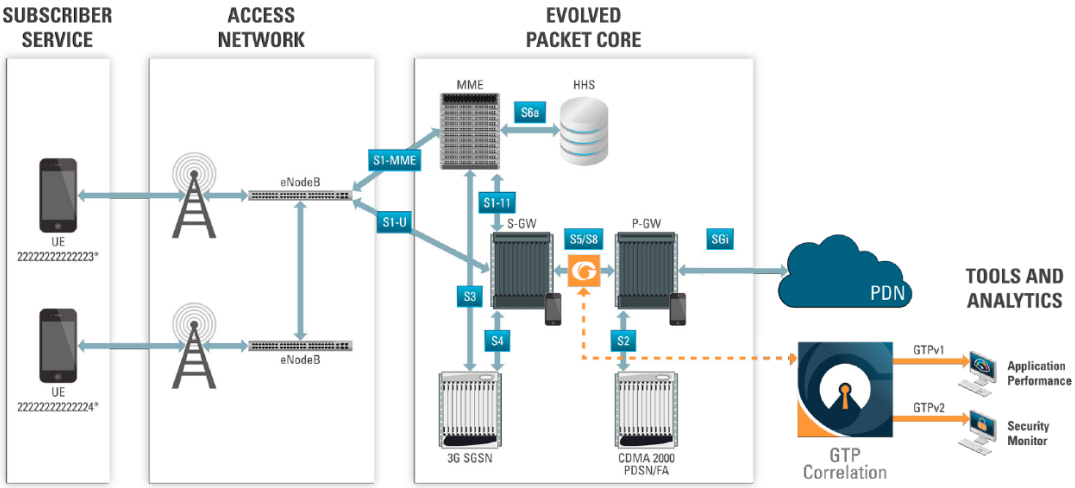

In Example 2, EMEI traffic is distributed based on GTP versions as follows:

- 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 GTP Rotational Sampling

Supported Devices: GigaVUE-HC3 Gen 2

Refer to Supported GigaSMART Operations for more details on the devices that support GigaSMART operations.

GTP Rotational Sampling samples GTP 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 the amount it shifts or slides every period is also configurable.

This capability is used in network security and network performance, where sampling is desired, and sampling needs to be deterministic so that it is always known whether a subscriber is sampled in or out.

In classic Deterministic Sampling, the pool of IMSIs remains constant, but in Rotational Sampling, the pool of IMSIs is updated at regular intervals, such that all subscribers get an opportunity to participate in the sampling.

The rotations do not disrupt treatment of existing sessions. Entire sessions are sampled in or out based upon the occurrence of the "Create Session" event. 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 diagram illustrates an example of GTP Rotational Sampling:

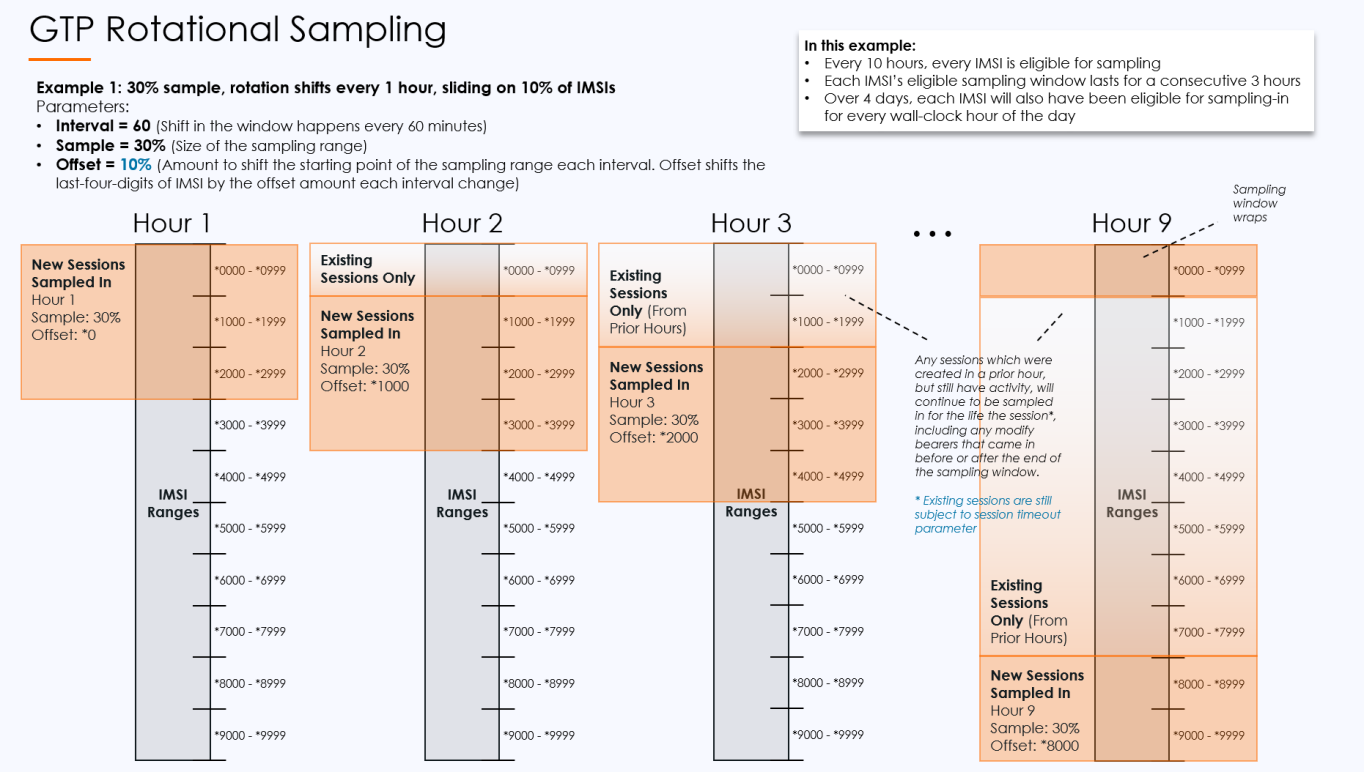

Supported Platforms

The feature is supported in the following cards:

- SMT-HC3-C05
- SMT-HC0-Q02X08

Configure Rotational Sampling

GigaVUE-FM allows you to configure Rotational Sampling in the second level map of GTP Flow Sample and GTP Flow Overlap.

To configure Rotational Sampling in second level maps, do the following:

1. On the left navigation pane, click on .
2. Go to **Physical** > **Nodes** > **Cluster ID**.
3. Go to **Traffic** > **Maps**.
4. Click **New**.
5. 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 Correlation
      - Flow SampleOverlap
   e. Enable the **Rotational Flow Sampling** check-box.
   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 10000 minutes.
6. Follow the steps from 5 to 9 in the section Create a New Mapto complete the configuration of rotational sampling.

Viewing the Configuration of Rotational Sampling for non-CUPS LTE

After configuring rotational sampling through Ansible, GigaVUE-FM allows you to view the configuration for non-CUPS LTE. 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**, **Current Offset Amount**, **Remaining Sampling**.

In GigaVUE-FM, Rotational Flow Sampling support provides you the ability to select the flow sampling type as "default" or "rotational" and configure the timer and offset values during the creation and edition of Mobility solution. When the user selects fstype as default the existing behaviour is retained (random flow sampling will be applied on maps). The user can edit all the "fstype" properties (type ,timer, and offset) from the Mobility solution. The "fstype" is only applicable for 'secondLevel/flowSample', 'secondLevel/flowSampleOverlap' map types.

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.

Rotational Sampling support in Mobility Solution

You can enable the rotational flow sampling type in the sampling section for Non Cups LTE only. When you specify **enable= true** in the rotational flow sampling , solution will set the **fstype** of second level maps as **rotational** and sets the **offset** and **timer** values. When you set **enable= false** then fstype is not set in second level maps.

In the UI, a column by name **Rotational Sampling** is added which can be either **Enabled** or **Disabled**. If the **Rotational Sampling** is **Enabled**, it shows as a link to quick view.

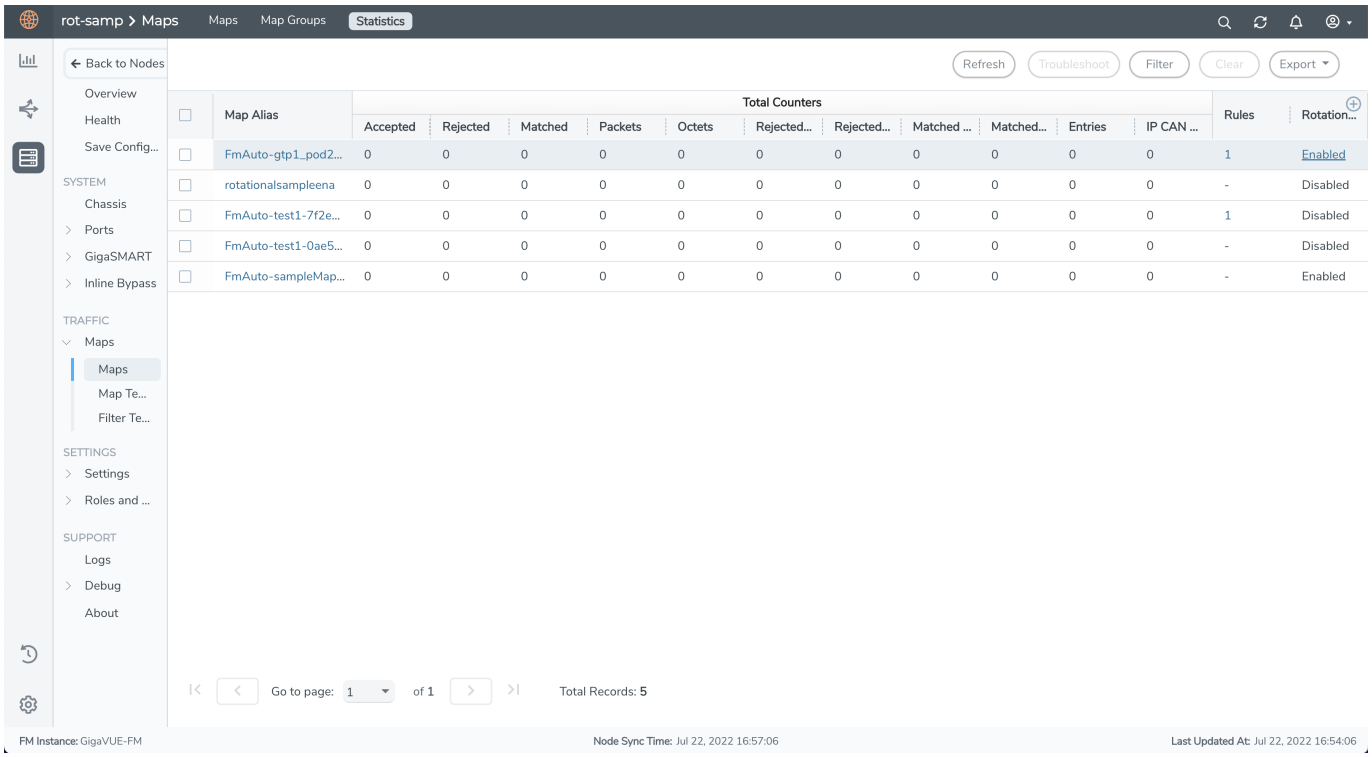

The quick view contains the following columns :

- Current Offset Amount
- Remaining Sampling Time
- Current Interval Start
- Next Interval Start

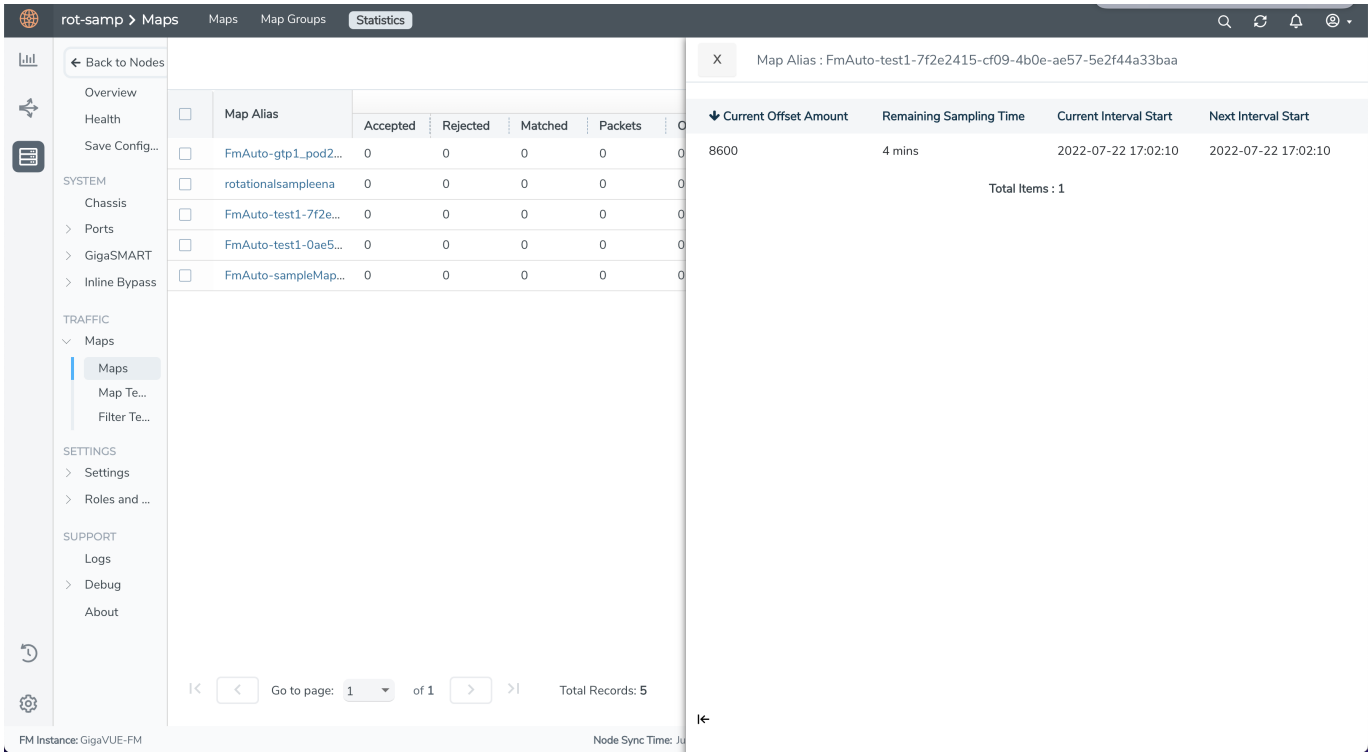
rot-samp > Maps
Maps
Map Groups
Statistics
Back to Nodes
Overview
Health
Save Config...
SYSTEM
Chassis
Ports
GigaSMART
Inline Bypass
TRAFFIC
Maps
Maps
Map Te...
Filter Te...
SETTINGS
Settings
Roles and ...
SUPPORT
Logs
Debug
About
Map Alias
Accepted
Rejected
Matched
Packets
FmAuto-gtp1_pod2...
rotationalsampleena
FmAuto-test1-7f2e...
FmAuto-test1-0ae5...
FmAuto-sampleMap...
Go to page: 1 of 1
Total Records: 5
FM Instance: GigaVUE-FM
Map Alias : FmAuto-test1-7f2e2415-cf09-4b0e-ae57-5e2f44a33baa
Current Offset Amount
Remaining Sampling Time
Current Interval Start
Next Interval Start
8600
4 mins
2022-07-22 17:02:10
2022-07-22 17:02:10
Total Items : 1

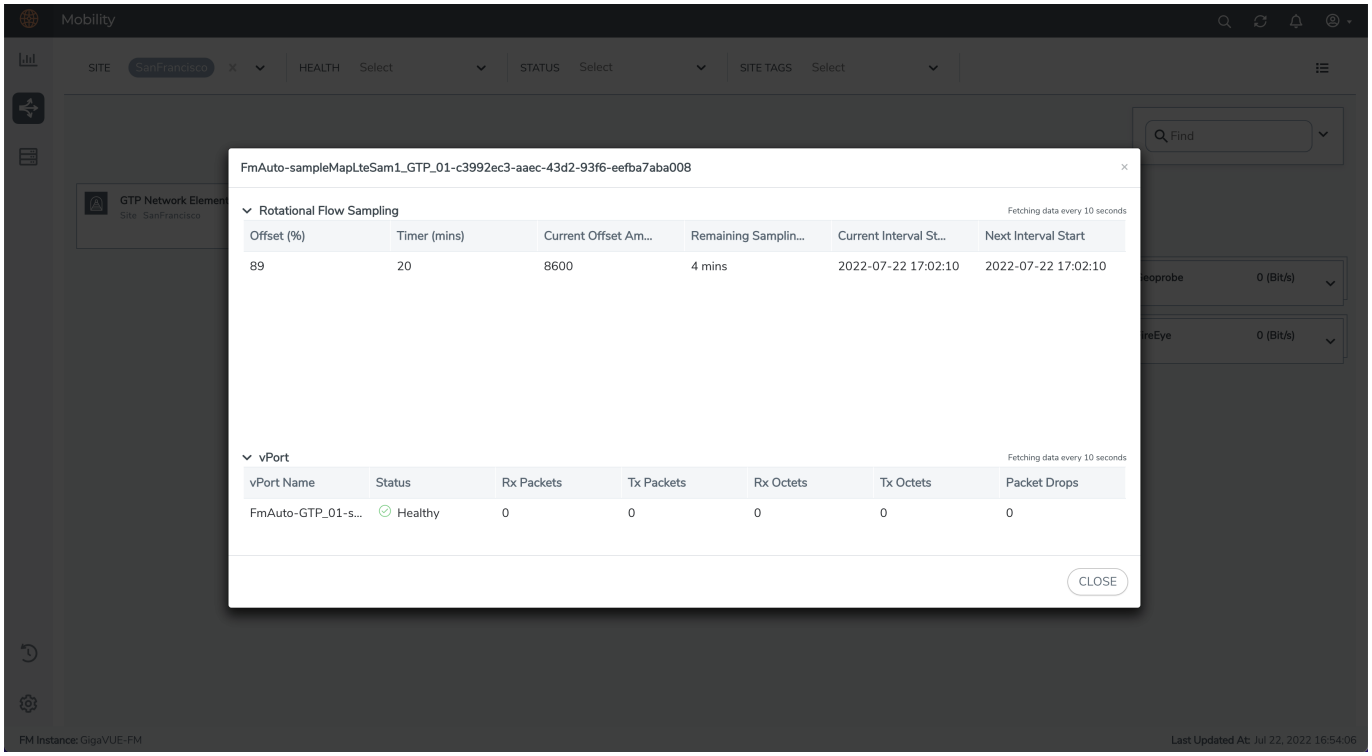
Mobility
SITE
SanFrancisco
HEALTH
Select
STATUS
Select
SITE TAGS
Select
Find
FmAuto-sampleMapLteSam1_GTP_01-c3992ec3-aaec-43d2-93f6-eefba7aba008
Rotational Flow Sampling
Fetching data every 10 seconds
Offset (%)
Timer (mins)
Current Offset Am...
Remaining Samplin...
Current Interval St...
Next Interval Start
89
20
8600
4 mins
2022-07-22 17:02:10
2022-07-22 17:02:10
vPort
Fetching data every 10 seconds
vPort Name
Status
Rx Packets
Tx Packets
Rx Octets
Tx Octets
Packet Drops
FmAuto-GTP_01-s...
Healthy
0
0
0
0
0
CLOSE
FM Instance: GigaVUE-FM
Last Updated At: Jul 22, 2022 16:54:06

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 for 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.

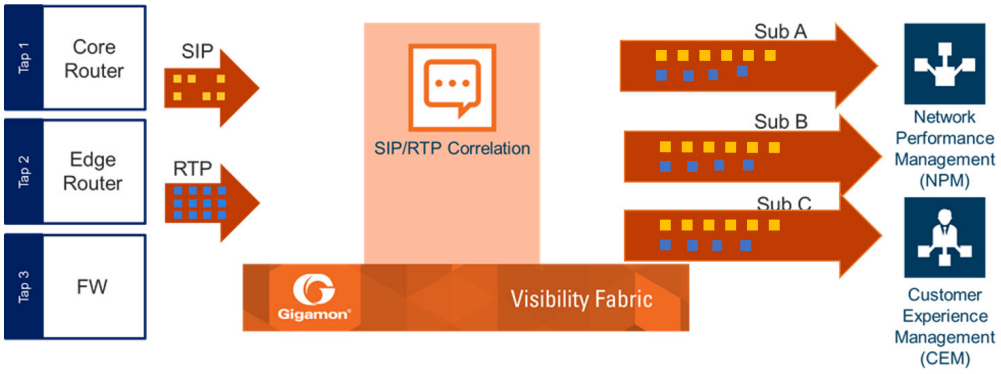

Figure 7 *SIP/RTP Correlation*

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 8 *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.

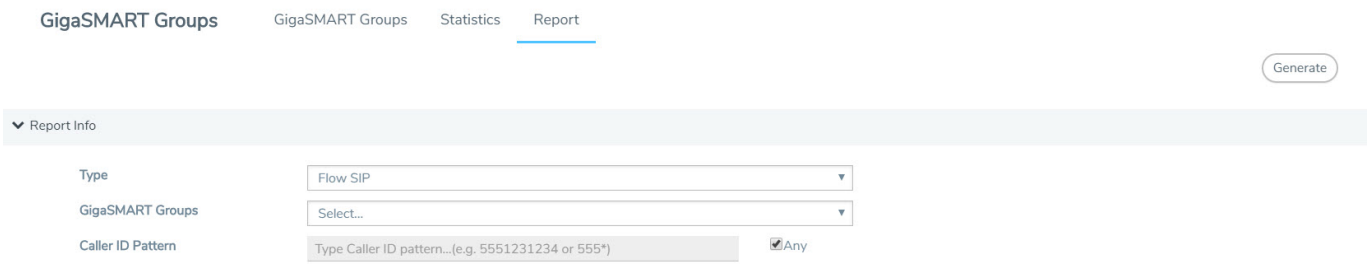

Figure 9 *Generate SIP Report*

6. Click the Generate button. The SIP Messages Report displays.



Figure 10 *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, click on from **Traffic** select **Maps> Statistics.** The Statistics page displays a count of the rules that actually matched in a map.



Figure 11 *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.

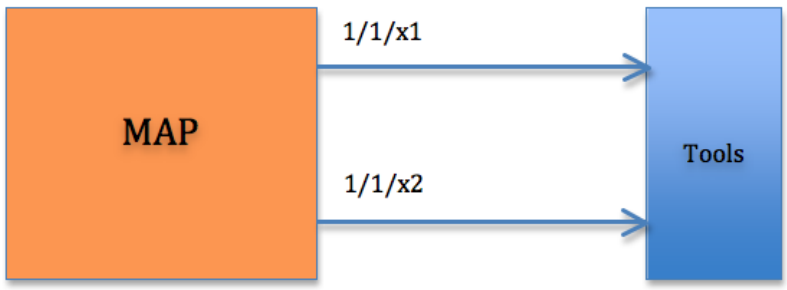

Admission Control

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.

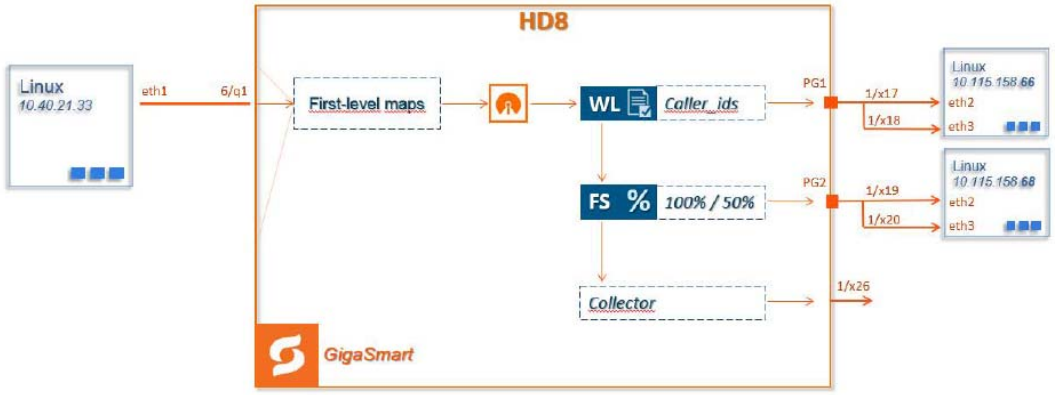

Figure 12 *SIP/RTP UDP/TCP Support*

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.

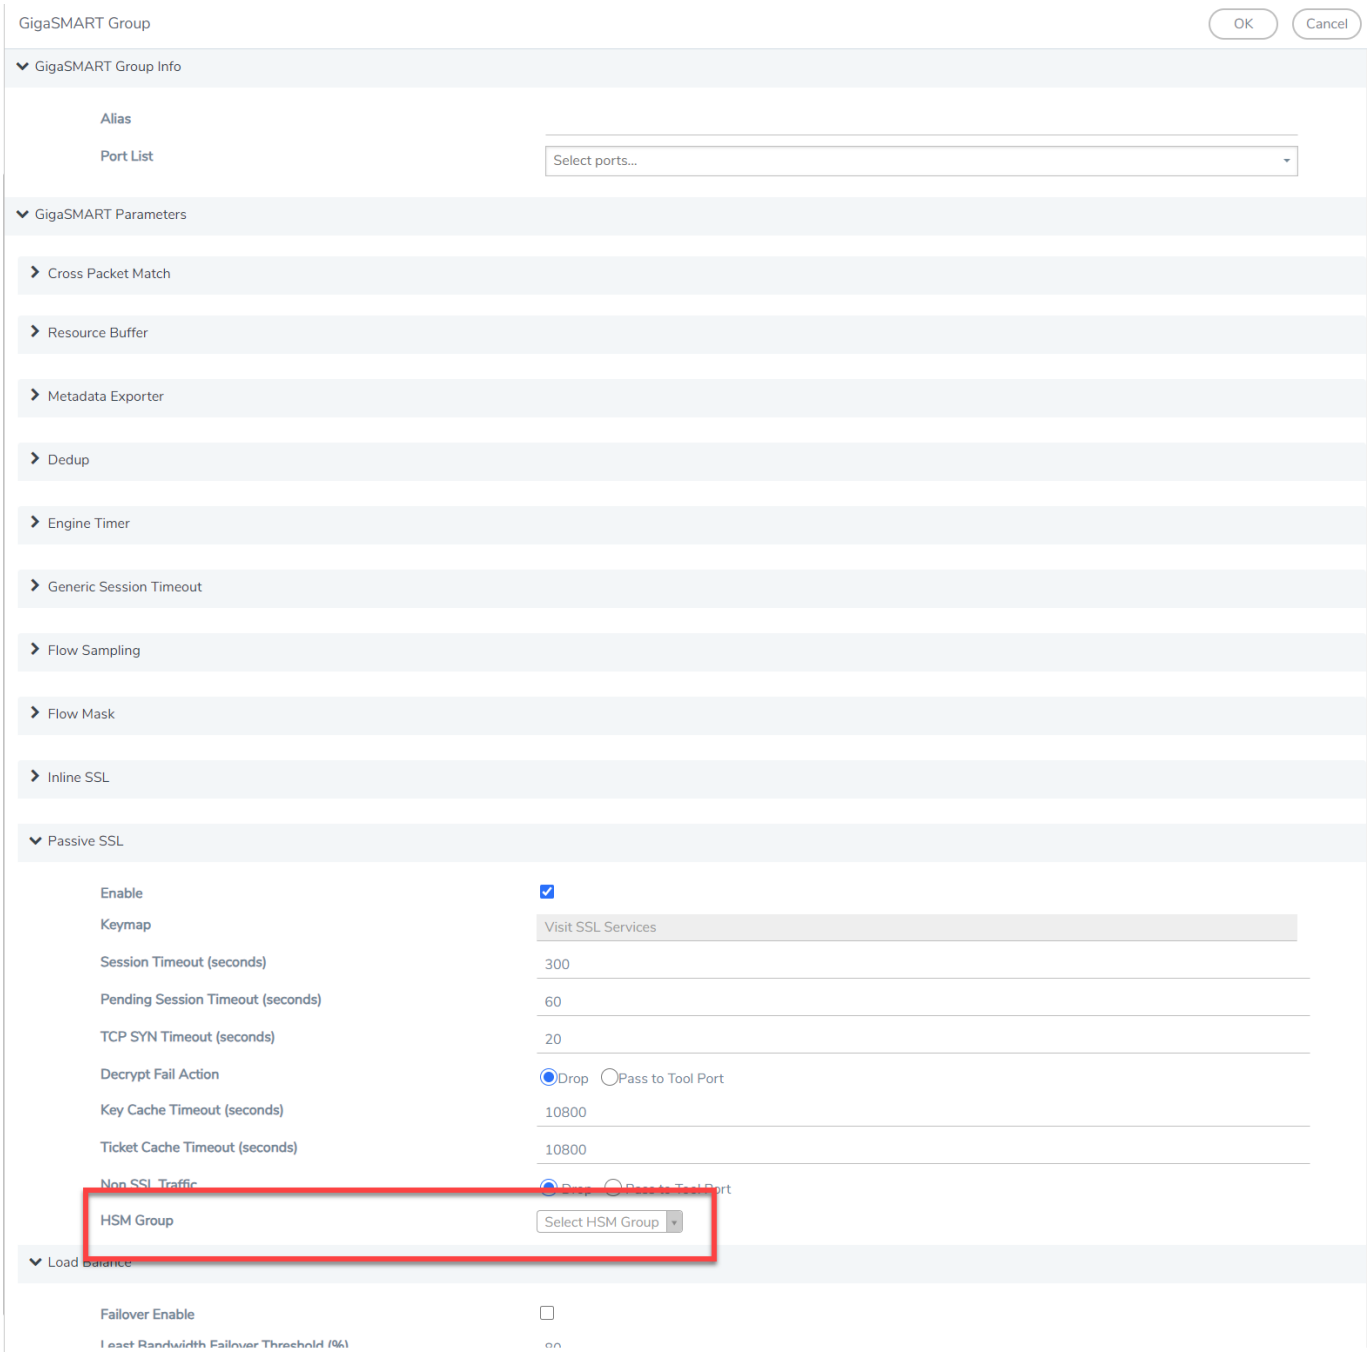

Figure 13 *GigaSMART Group Port Info*

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**.

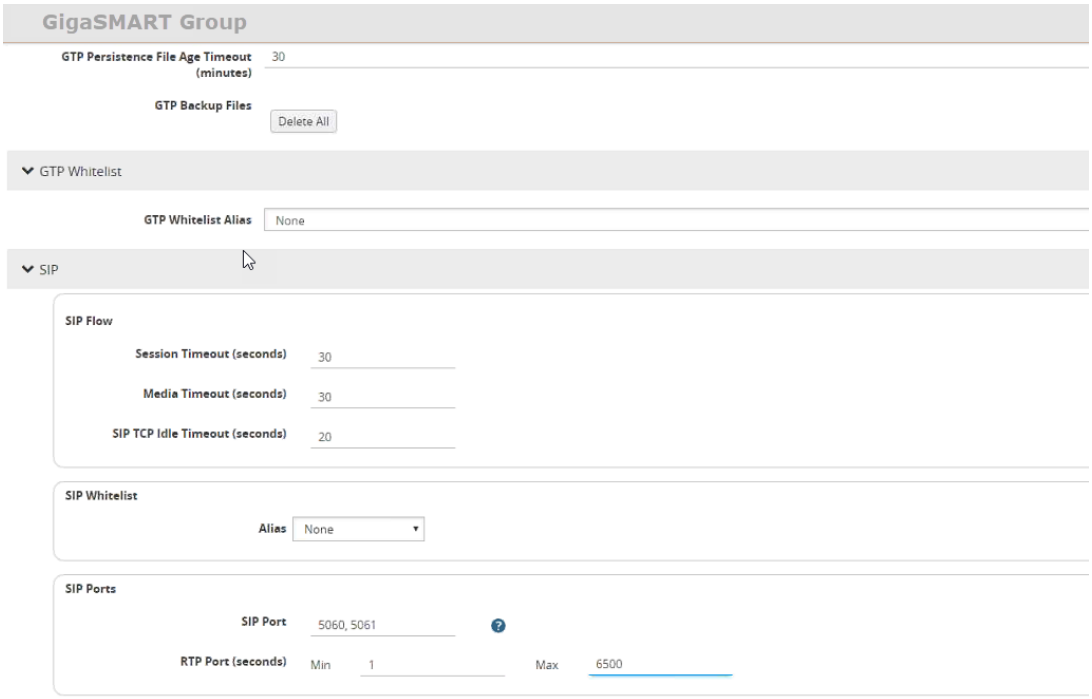

Figure 14 *SIP Port parameters*

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**.

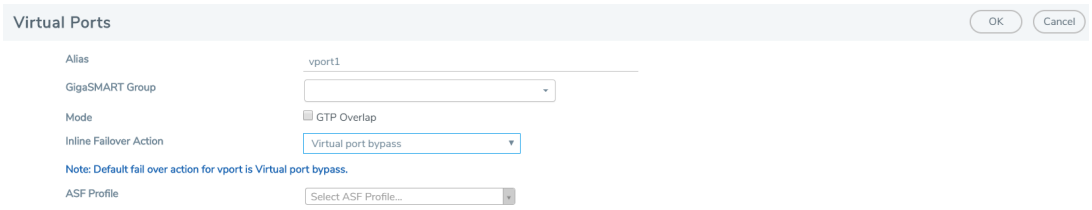

Figure 15 *Virtual Ports*

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:

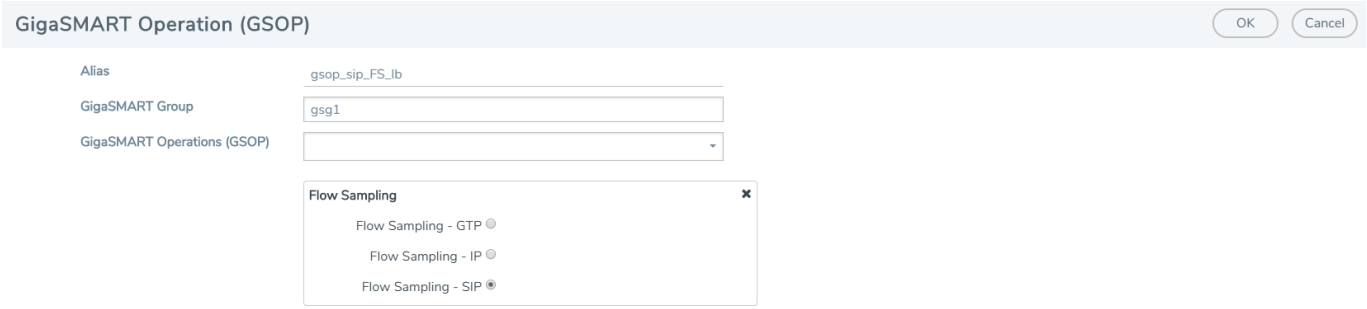

Figure 16 *GigaSMART Operations page*

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.

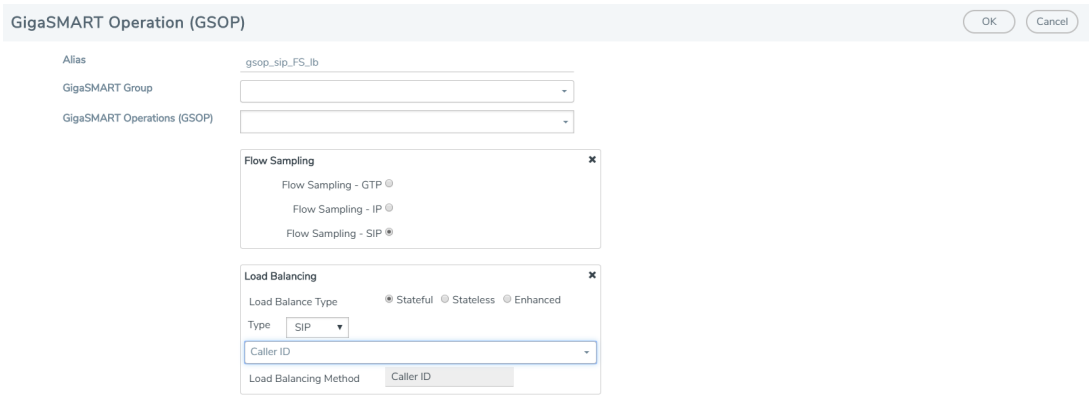

Figure 17 GigaSMART Operations - Load Balancing

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, click on  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.

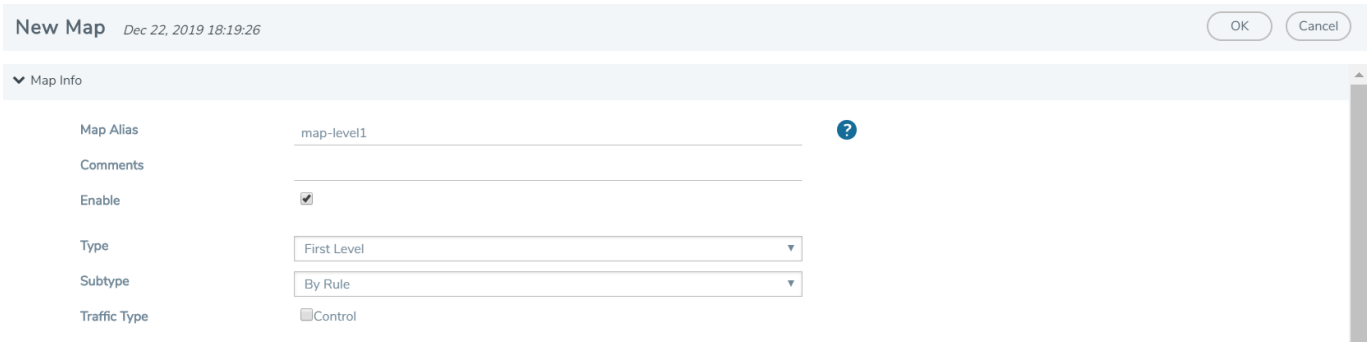

Figure 18 *Create New Map*

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, click on 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

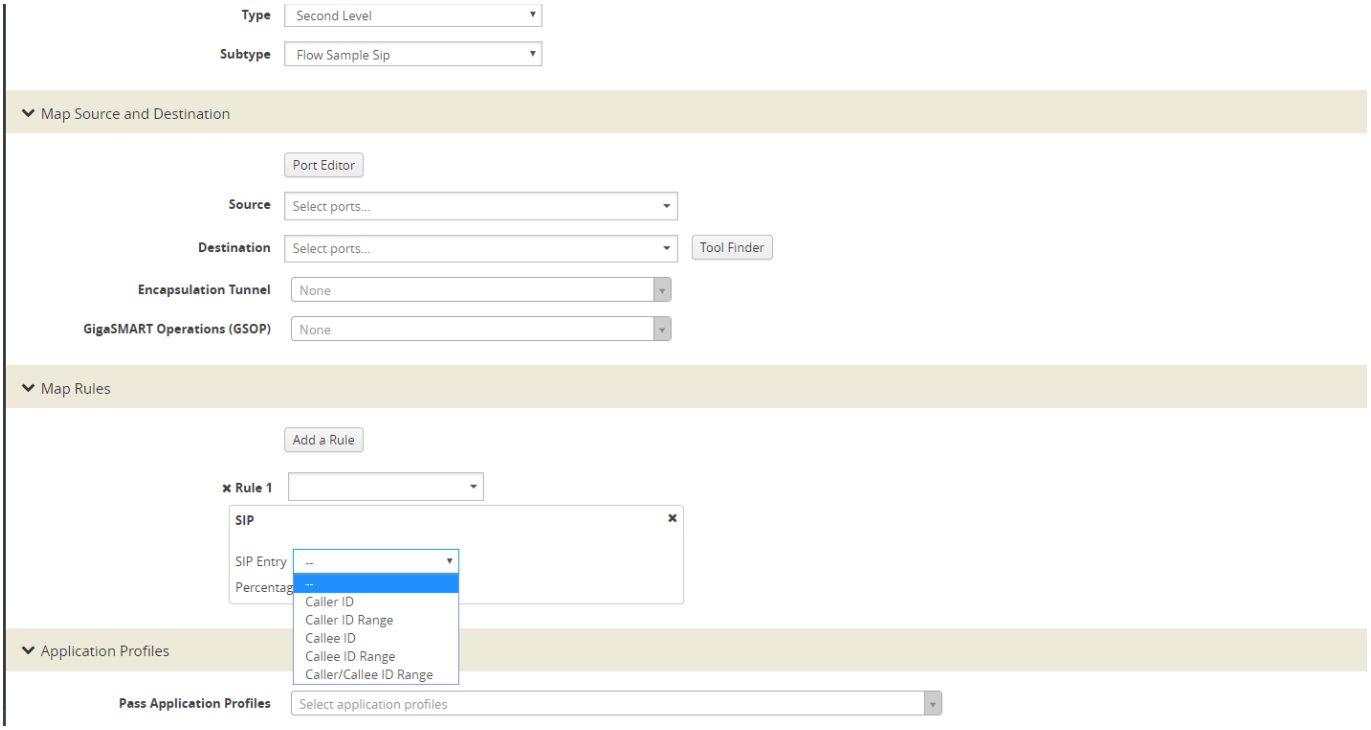

Figure 19 *Create Second Level Map*

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, click on  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 for 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.** From the left navigation pane, 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 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, click on  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.** From the left navigation pane, 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, click on  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.** From the left navigation pane, 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.
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.

Task	Description	UI Steps
		5. Click Save.
5	Create a first level map and direct traffic to the virtual port.	**1.** On the left navigation pane, click on  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, click on  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. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Port Source. **d.** Enter 80 for the port value. **e.** Select 2 for Position.

Task	Description	UI Steps
		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 for 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, click on  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.

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, click on  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, click on  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, click on  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, click on  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.** From the left navigation pane, 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, click on  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, click on  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.** From the left navigation pane, 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, click on  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, click on  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, click on 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 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* **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, click on 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, click on  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, click on  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.** From the left navigation pane, 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, click on  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, click on  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, click on  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.** From the left navigation pane, 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.

Task	Description	UI Steps
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.** 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

Task	Description	UI Steps
		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.
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, click on  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.

Task	Description	UI Steps
9.	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 whitelist map.	**a.** On the left navigation pane, click on and 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, click on 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:

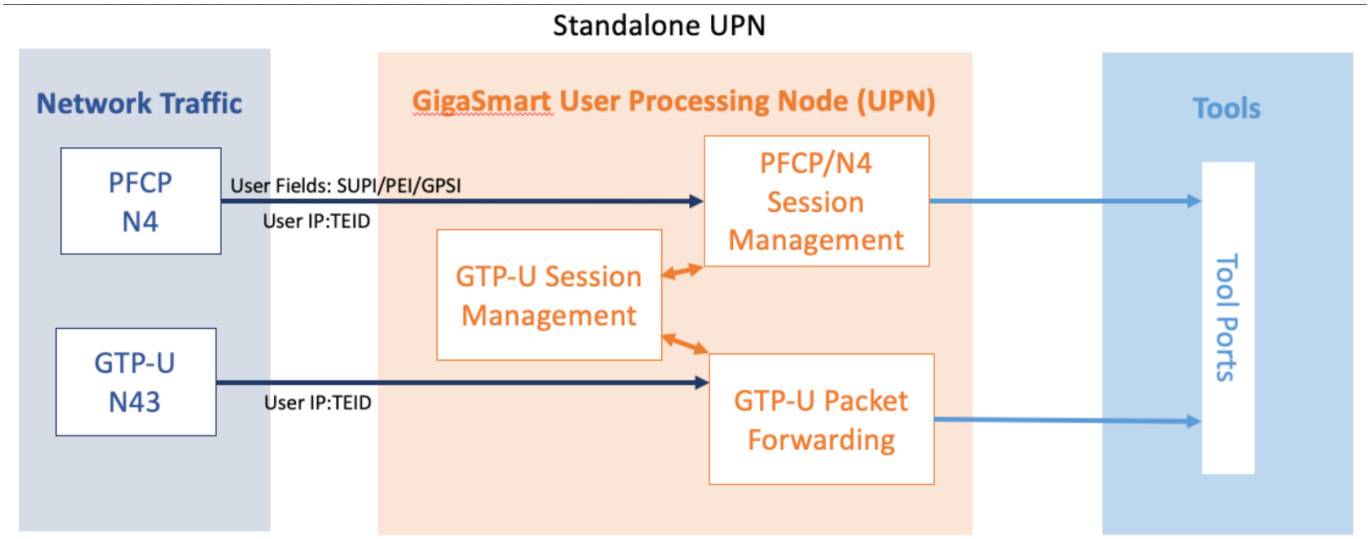

Rules and Notes

- 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 20 *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.

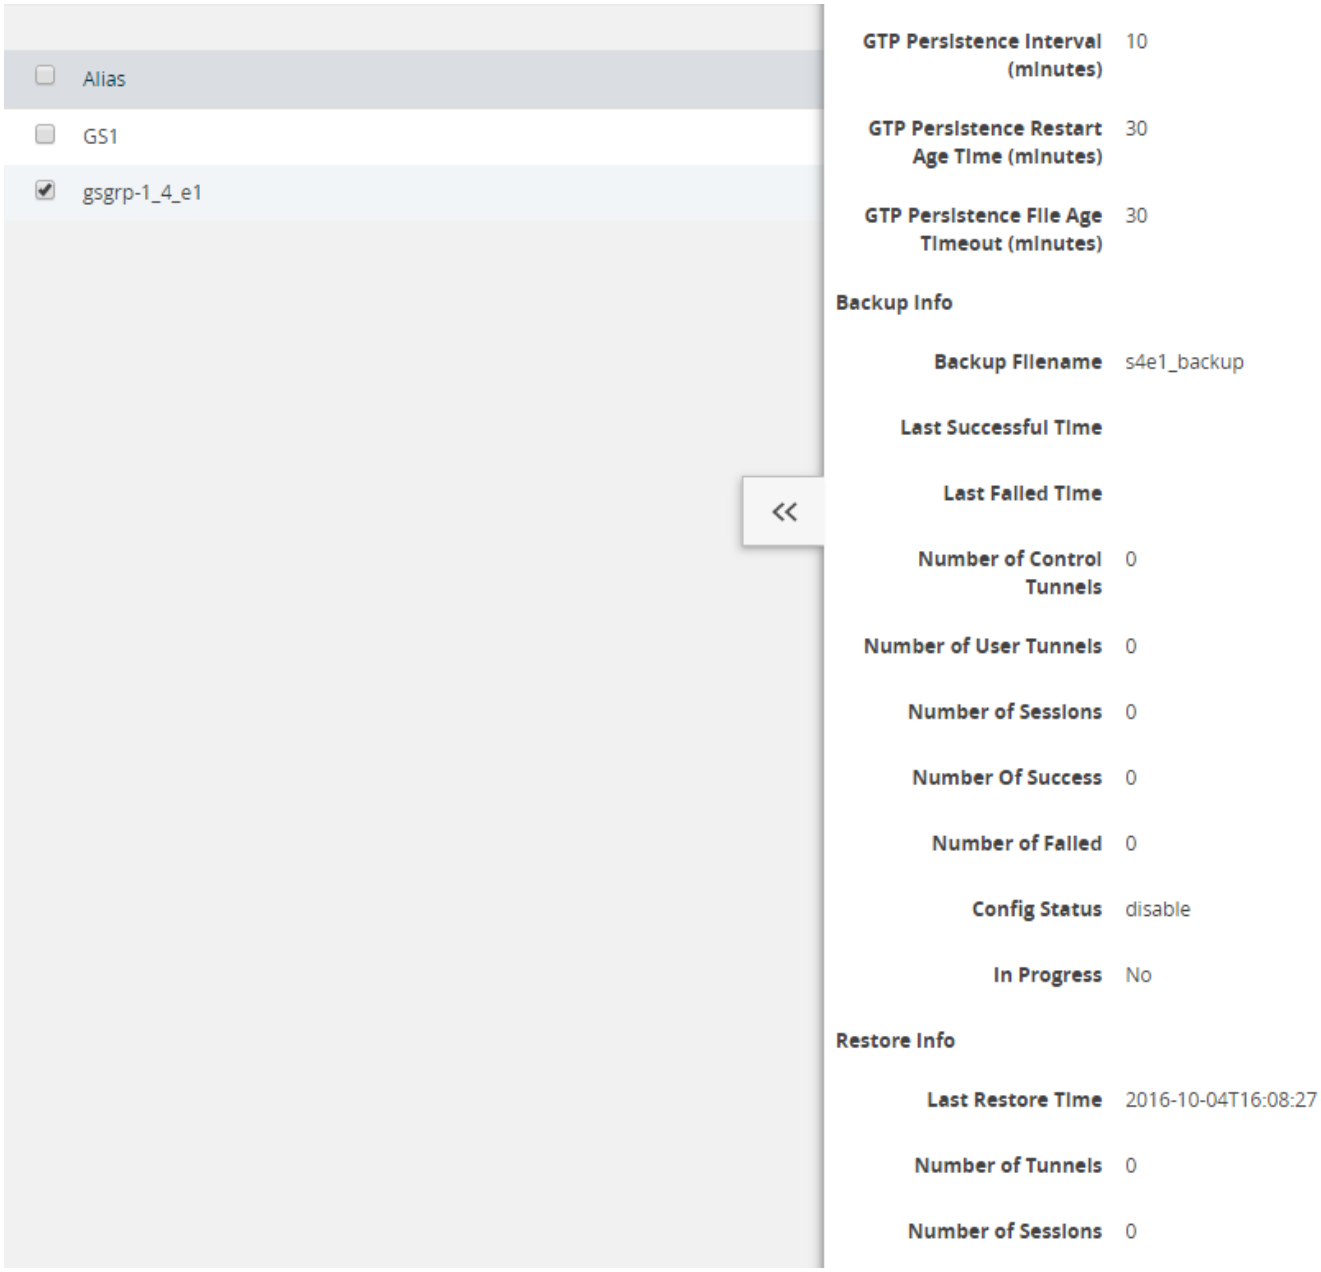

Figure 2 GTP Persistence Information

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

Use the following sequence to help you determine the maximum number of port group members:

1. Determine the number of members per port group and add 1 to the number.

2. Multiply each port group result times each other.

3. The total multiplication should not exceed 1024.

For instance, assume the following configuration in a GTP overlap mapping group:

- Port Group 1—2 load balanced GigaStream
- Port Group 2—3 load balanced GigaStream

- Port Group 3—1 load balanced tool port
- Port Group 4—1 load balanced GigaStream
- Port Group 5—4 load balanced tool ports The total number becomes:

(2+1)*(3+1)*(1+1)*(1+1)*(4+1) = 240

Since this does not exceed the maximum number of multicast IDs (1024), the tool configuration shown is accepted.

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.

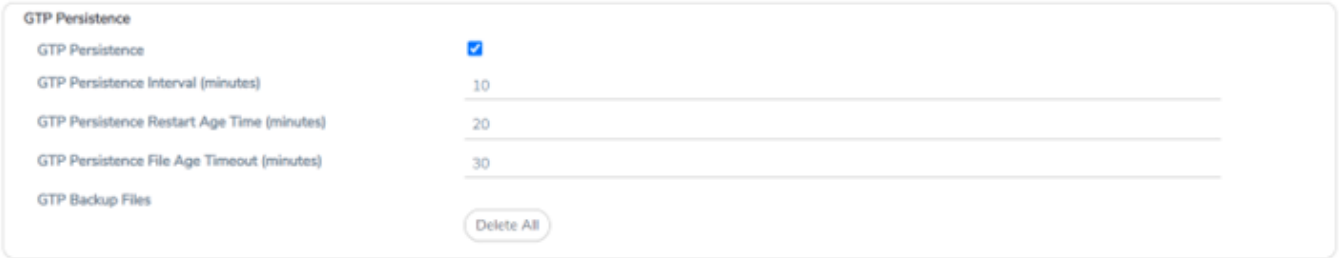

5G Stateful Session Recovery

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.

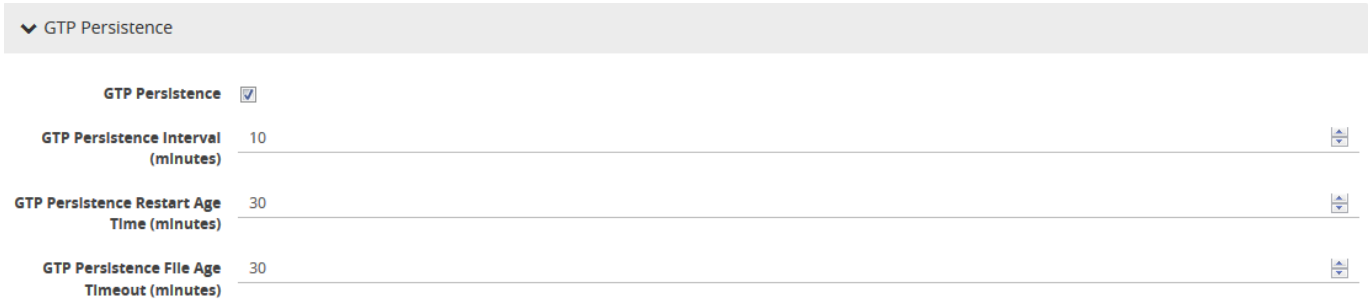

Figure 21 *GTP Persistence GigaSMART Parameters*

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.

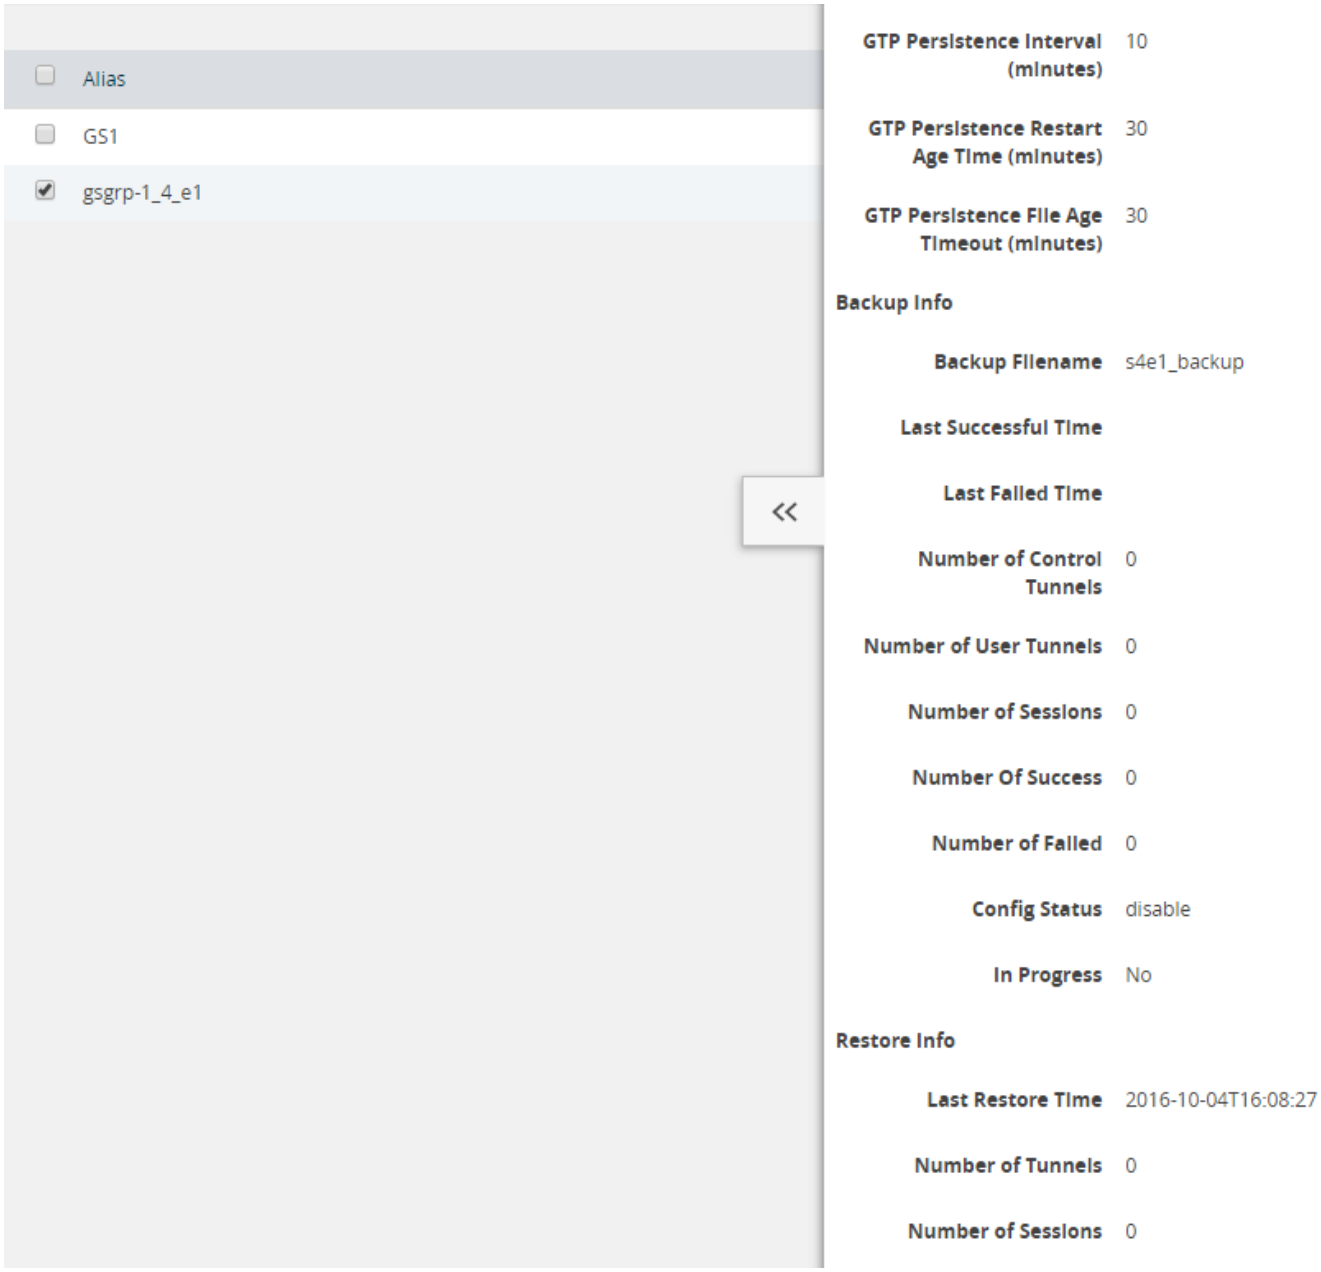

Figure 2 GTP Persistence Information

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

Use the following sequence to help you determine the maximum number of port group members:

1. Determine the number of members per port group and add 1 to the number.

2. Multiply each port group result times each other.

3. The total multiplication should not exceed 1024.

For instance, assume the following configuration in a GTP overlap mapping group:

- Port Group 1—2 load balanced GigaStream
- Port Group 2—3 load balanced GigaStream

- Port Group 3—1 load balanced tool port
- Port Group 4—1 load balanced GigaStream
- Port Group 5—4 load balanced tool ports The total number becomes:

(2+1)*(3+1)*(1+1)*(1+1)*(4+1) = 240

Since this does not exceed the maximum number of multicast IDs (1024), the tool configuration shown is accepted.

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).

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.

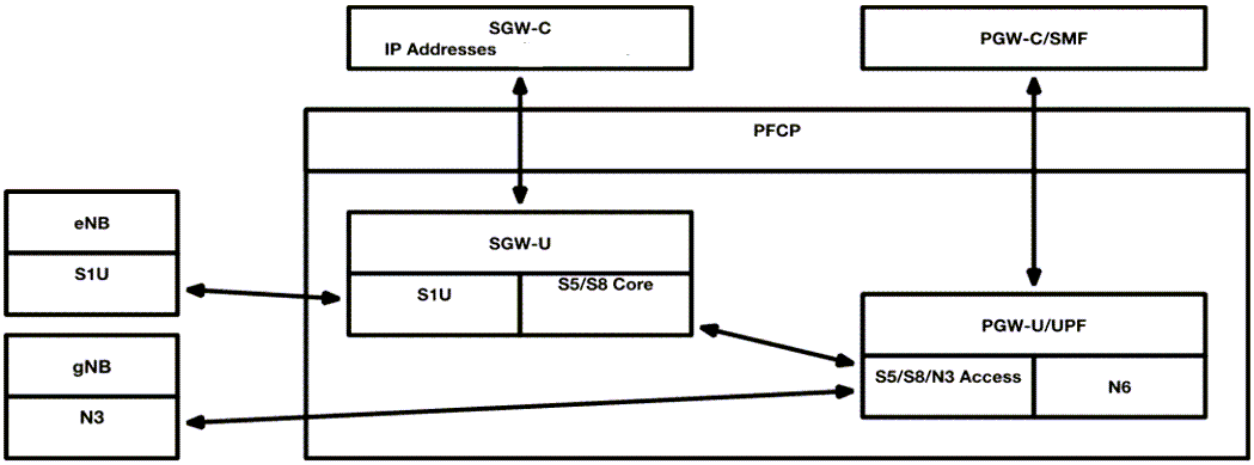

Supported Platforms:

- 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 for 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:

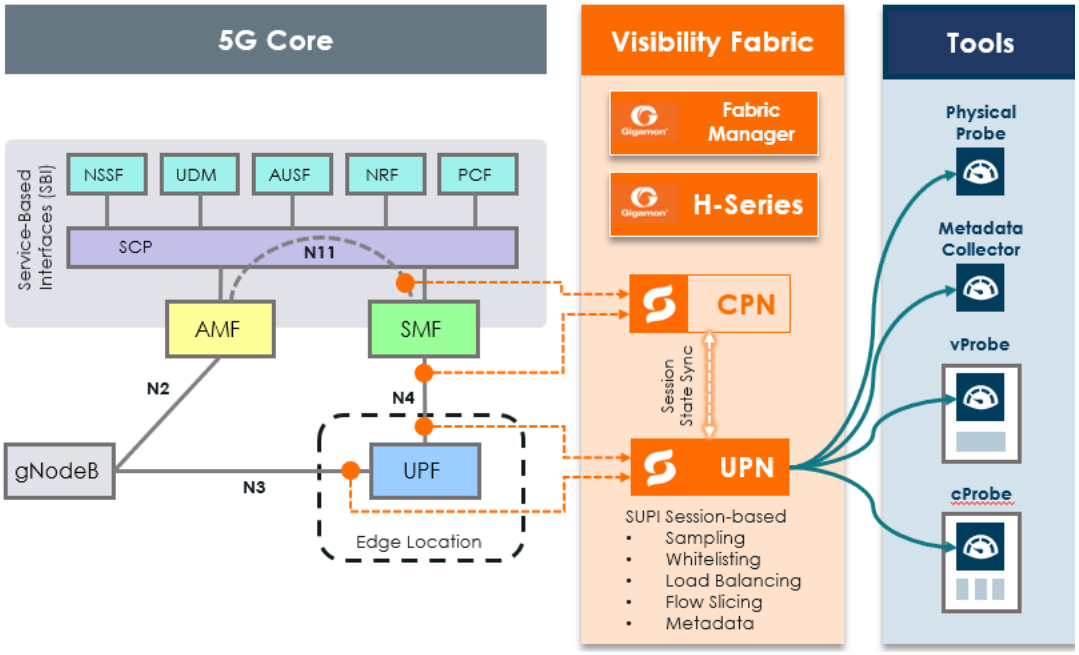

Figure 22 *5G Stand Alone Packet Core Network*

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:

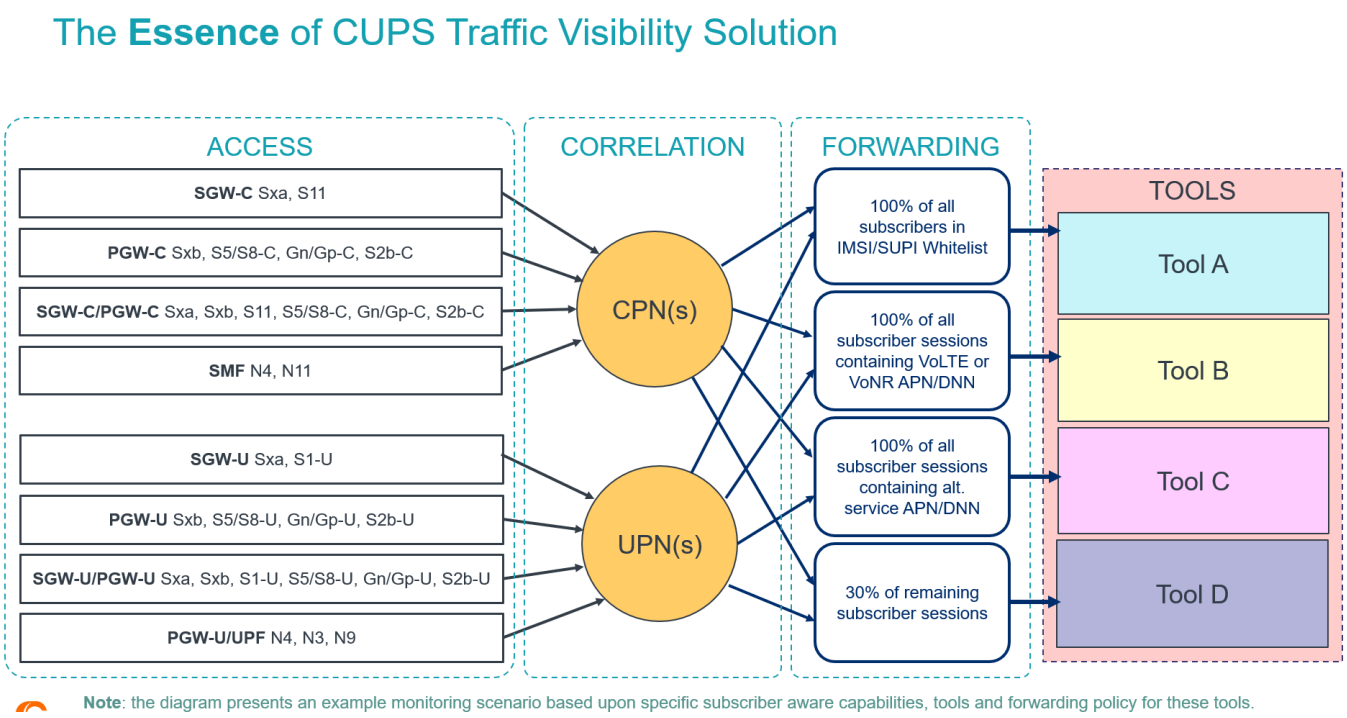

Figure 23 *CUPS Traffic Visibility Solution*

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.

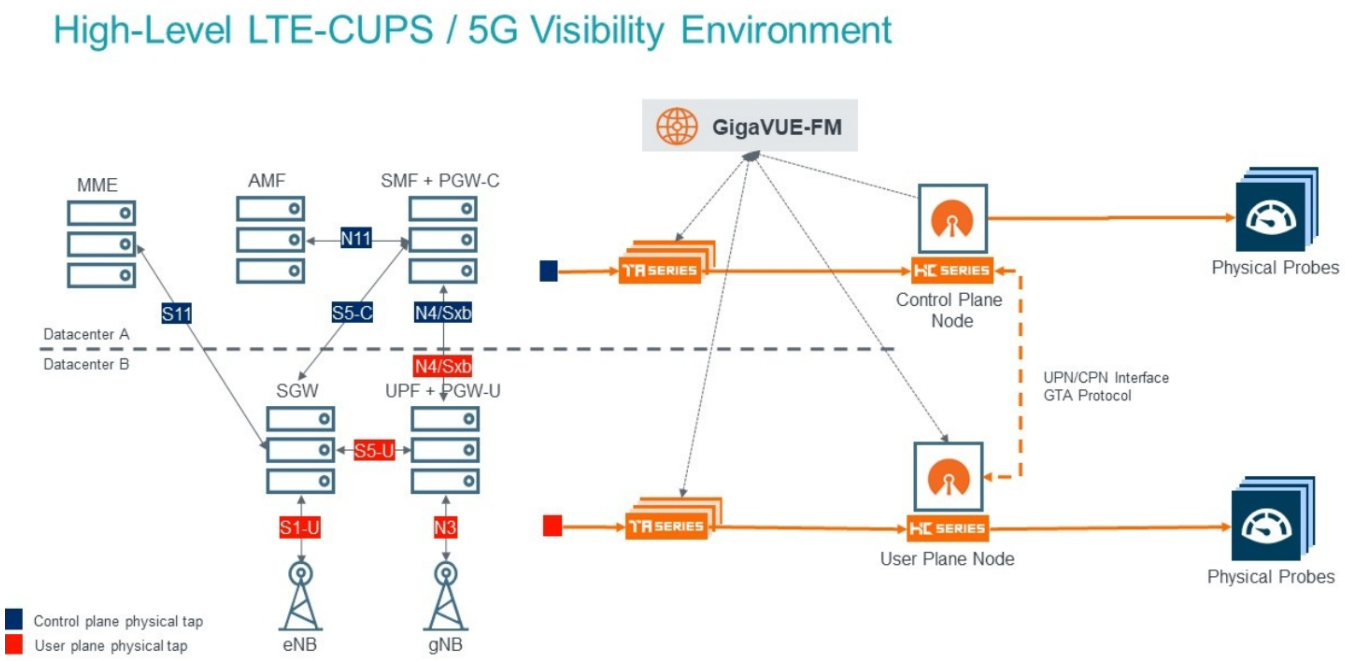

5G Load Balancing

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.

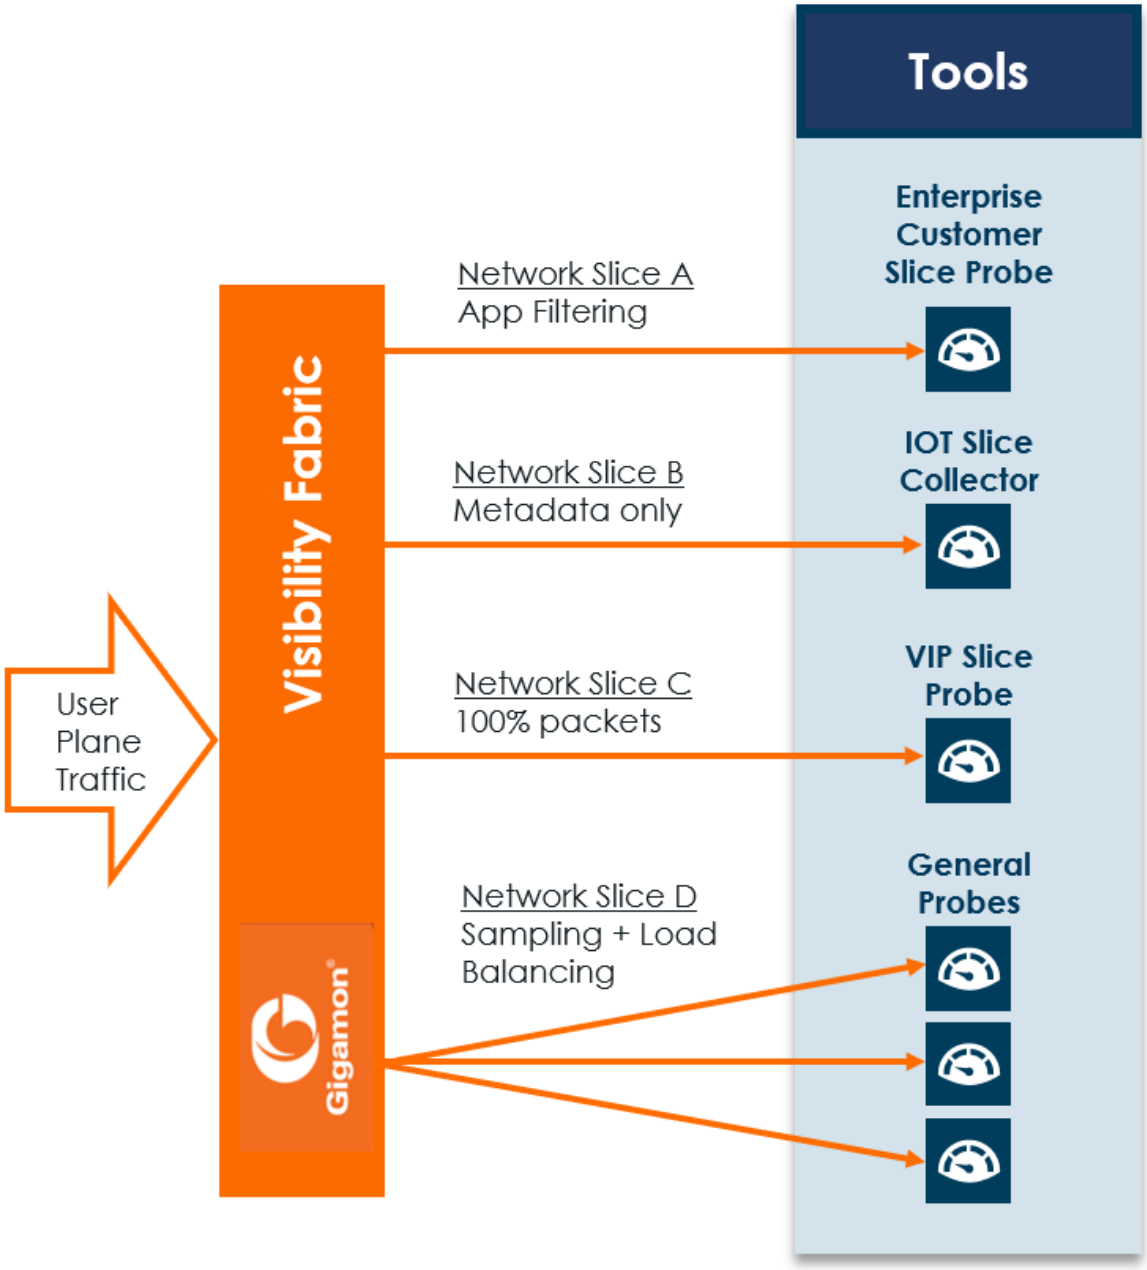

5G Network Slice Correlation uses 5G Correlation. It exposes Network Slice instance ID (NSI). NSI is available for filtering through flow sampling maps. 5G subscriber aware filter includes the Network Slice Selection Assistance Information (NSSAI) value for 5G correlation.

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.

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

Use the following sequence to help you determine the maximum number of port group members:

1. Determine the number of members per port group and add 1 to the number.
2. Multiply each port group result times each other.
3. The total multiplication should not exceed 4096.

For instance, assume the following configuration in a GTP overlap mapping group:

- Port Group 1—2 load balanced GigaStreams
- Port Group 2—3 load balanced GigaStreams
- Port Group 3—1 load balanced tool port
- Port Group 4—1 load balanced GigaStream
- Port Group 5—4 load balanced tool ports

The total number becomes:

(2+1)*(3+1)*(1+1)*(1+1)*(4+1) = 240

Since this does not exceed the maximum number of multicast IDs (4096), the tool configuration shown is accepted.

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 24GTP Persistence GigaSMART Parameters. The timers are pre-configured with default values.

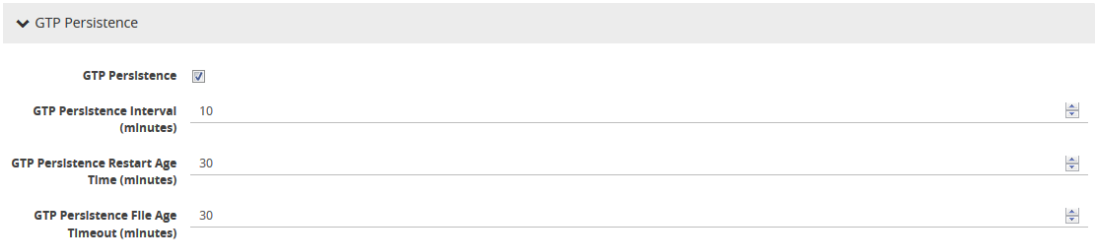

Figure 24 *GTP Persistence GigaSMART Parameters*

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 25GTP Persistence Information, GigaSMART Group gsgrp-1_4_e1 is selected and the Quick View is displayed.

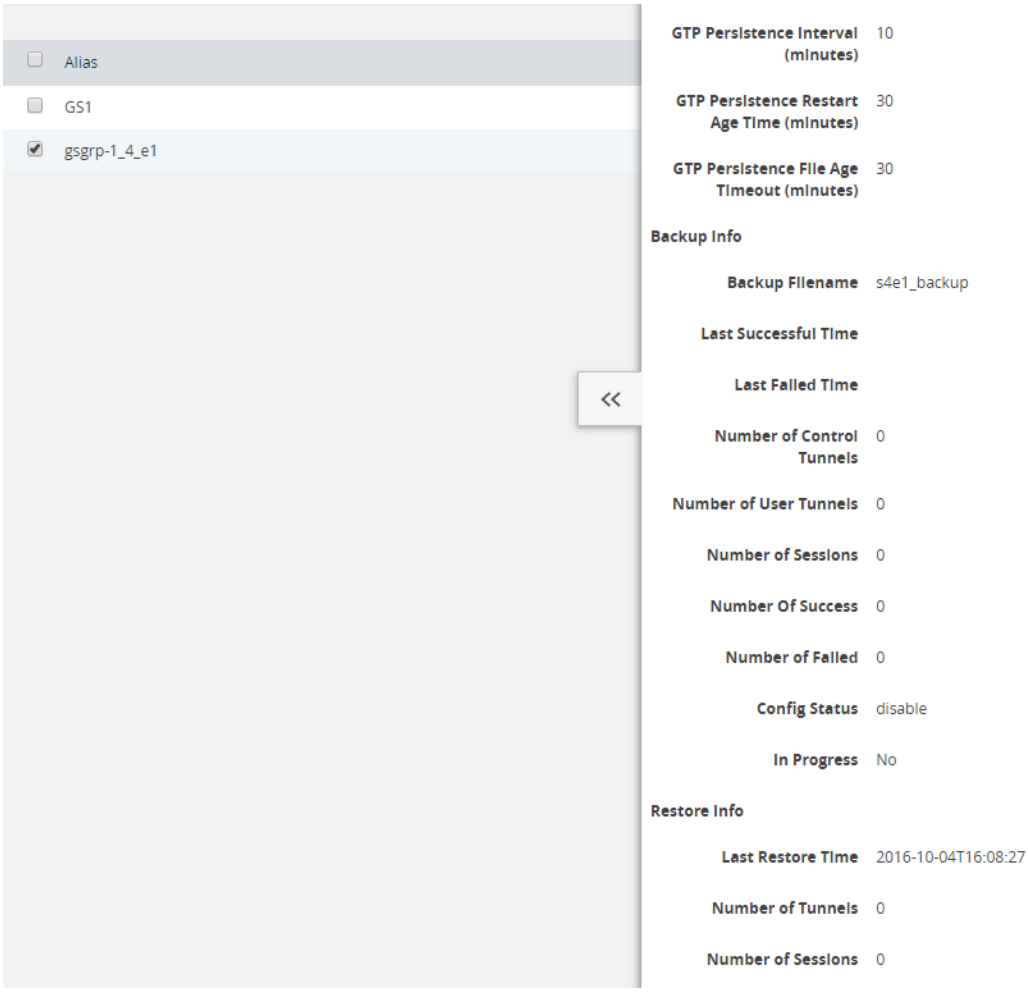

Figure 25 *GTP Persistence Information*

The following table describes persistence information.

Table 2: 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 26GTP Backup Files Delete)
4. Click **Delete All** under **GTP Backup Files**.

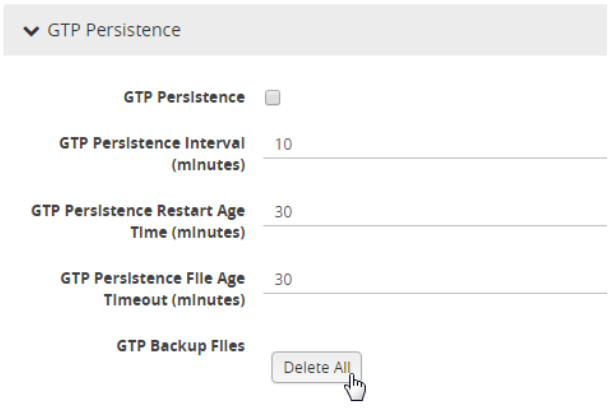

Figure 26 *GTP Backup Files Delete*

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:

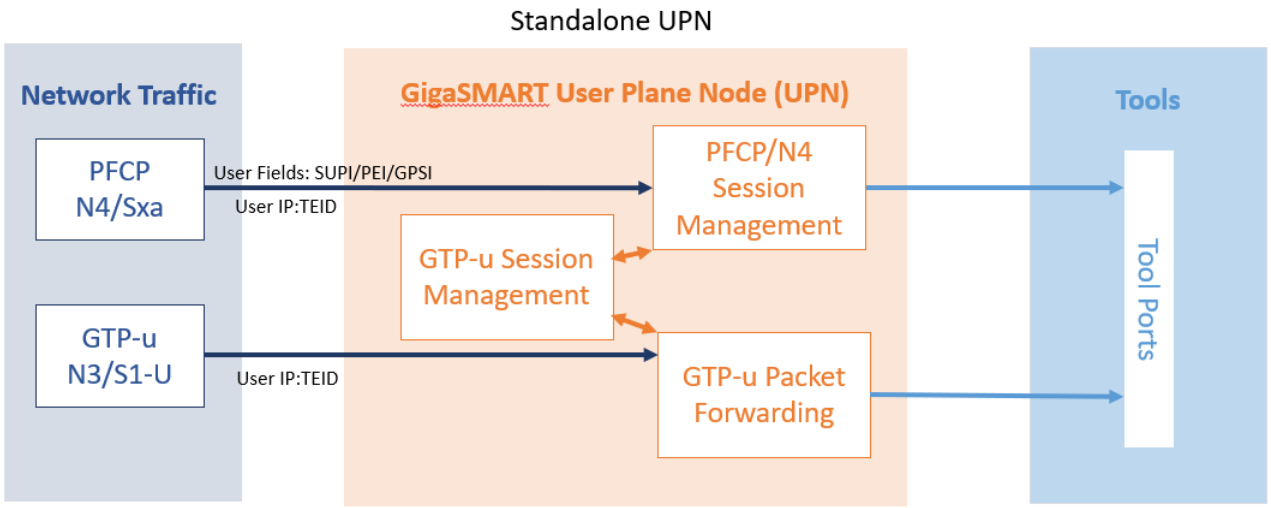

Rules and Notes

- 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.

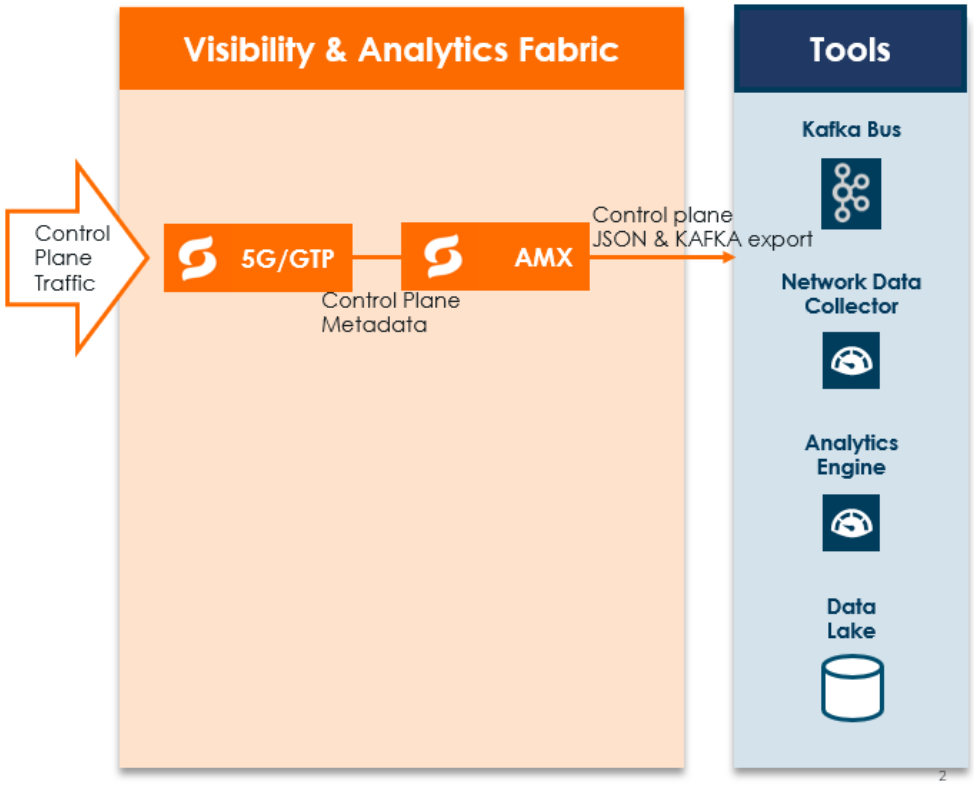

Configuration of 3G/4G Control Plane Metadata using Ansible

To configure 3G/4G control plane metadata, perform the following steps:

S.No	Steps	Refer to..
1.	Creating Inventory Directory	Configuration of 3G/4G Control Plane Metadata using Ansible
2.	Creating ***fmInfo.yml***	Configuration of 3G/4G Control Plane Metadata using Ansible
3.	Creating ***ansible_inputs.json***	Configuration of 3G/4G Control Plane Metadata using Ansible
4.	Creating 3G/4G inventory file	Configuration of 3G/4G Control Plane Metadata using Ansible
5.	Creating **host_vars** directory	Configuration of 3G/4G Control Plane Metadata using Ansible
6.	Creating **host_vars** files	Configuration of 3G/4G Control Plane Metadata using Ansible

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.

S.No	Groups-Input
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	Configuration of 5G Control Plane Metadata using Ansible
2.	Creating ***fmInfo.yml***	Configuration of 5G Control Plane Metadata using Ansible
3.	Creating ***ansible_inputs.json***	Configuration of 5G Control Plane Metadata using Ansible
4.	Creating 3G/4G inventory file	Configuration of 5G Control Plane Metadata using Ansible
5.	Creating **host_vars** directory	Configuration of 5G Control Plane Metadata using Ansible
6.	Creating **host_vars** files	Configuration of 5G Control Plane Metadata using Ansible

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.

S.No	Groups-Input
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]
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

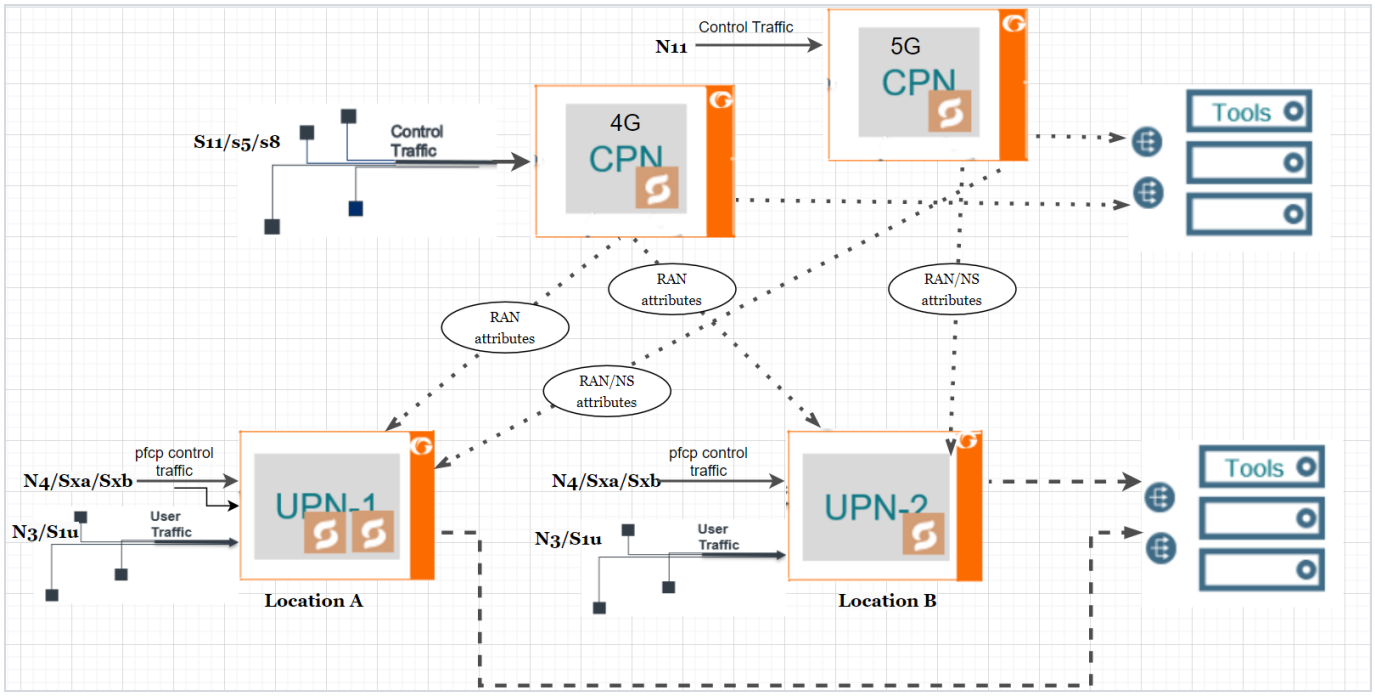

The above topology explains how CPN and UPN communicate with each other.

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:

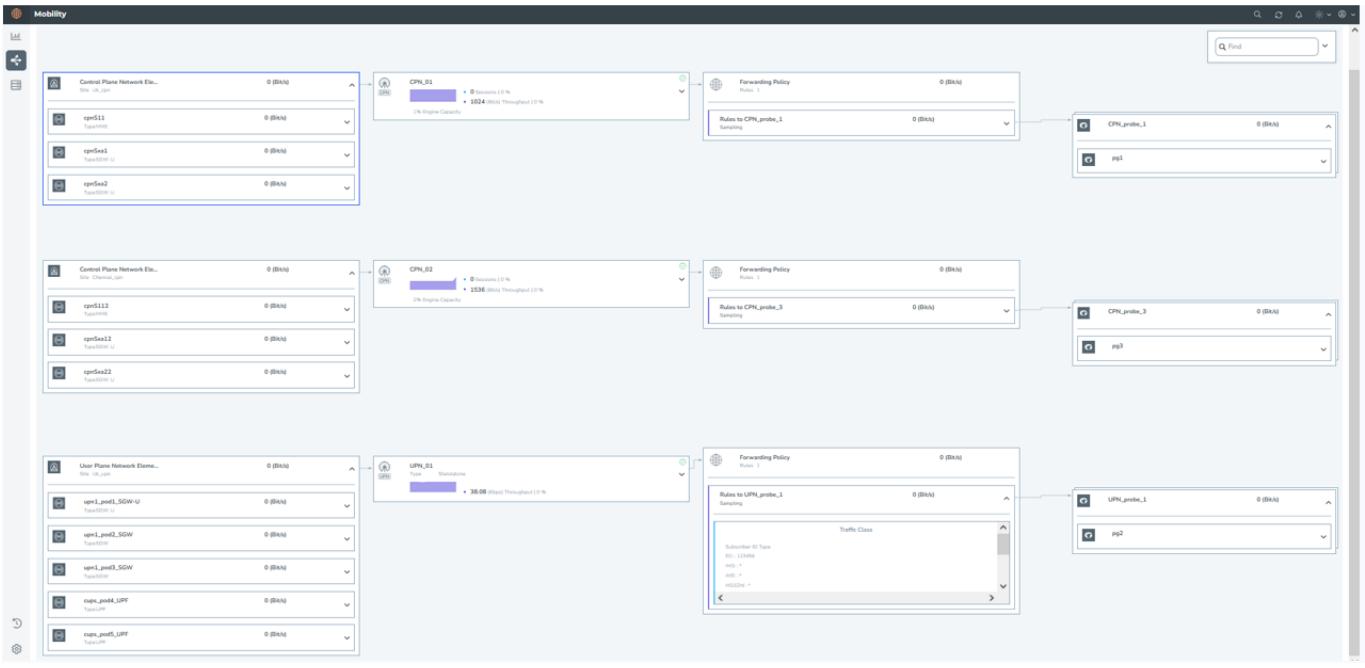

After configuring CPN-UPN Communication solution, you can view the analytics and statistics in the CUPS Communication dashboard. Refer to View CUPS Communication Dashboard for more detailed information.

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.

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** • PFCP packets matching IMSI are forward listed • RAN forward list is not applicable	**IMSI** • 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	**RAN + APN + Interface** • GTP-u packets matching the

Database Contains	Forward list Type	Forward list Rule - APN	Forward list Rule - Interface	PFCP Packets Filtered for Matching	GTP-U Packets Filtered Matching
				listed • RAN forward list is not applicable	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:

Table 3: Packet Statistics Dashboard

Dashboard	Description	Visualizations	Details
Packet Statistics	Displays the visualization details of Packet Statistics.	**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.

Table 4: GPFCP Statistics Dashboard

Dashboard	Description	Visualizations	Details
GPFCP Statistics	Displays the visualization details of GPFCP Statistics.	**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.

Table 5: Enrichment Statistics Dashboard

Dashboard	Description	Visualizations	Details
Enrichment Statistics	Displays the visualization details of Enrichment Statistics.	**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

Dashboard	Description	Visualizations	Details
			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.
		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
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-FMafter 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. On the left navigation pane, click on  .
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:

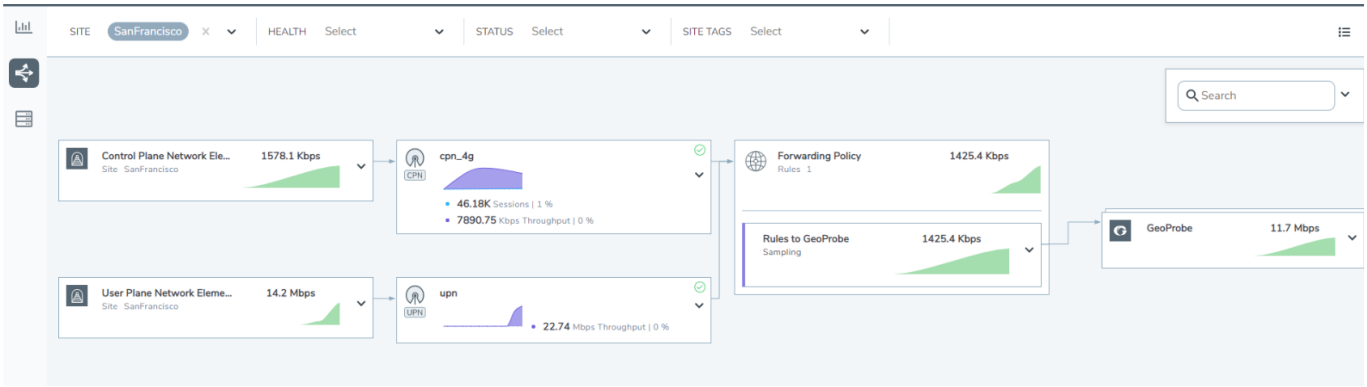

GigaVUE-FM also helps you to view the statistics and health information of the following:

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 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.

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

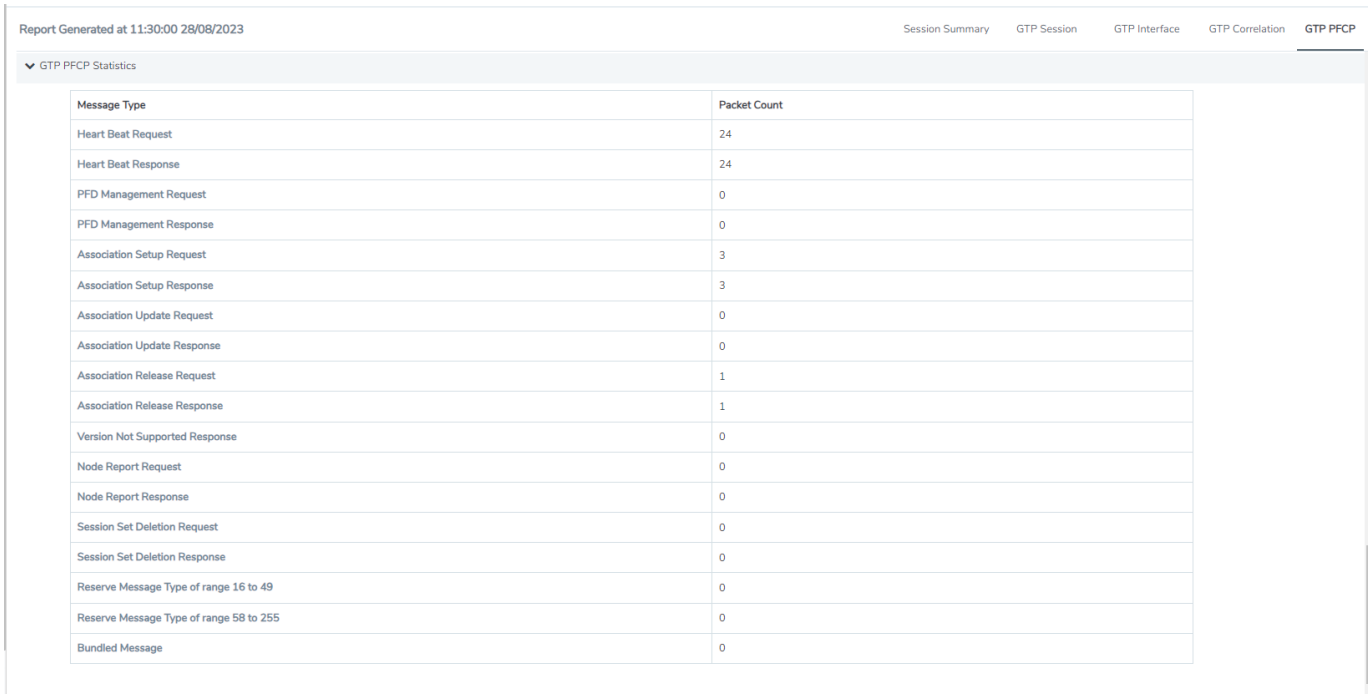

Message Type	Packet Count
Heart Beat Request	24
Heart Beat Response	24
PFD Management Request	0
PFD Management Response	0
Association Setup Request	3
Association Setup Response	3
Association Update Request	0
Association Update Response	0
Association Release Request	1
Association Release Response	1
Version 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 engines. The packets that are sent to the older tunnels are not taken into account in the flow-ops report.

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
			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.

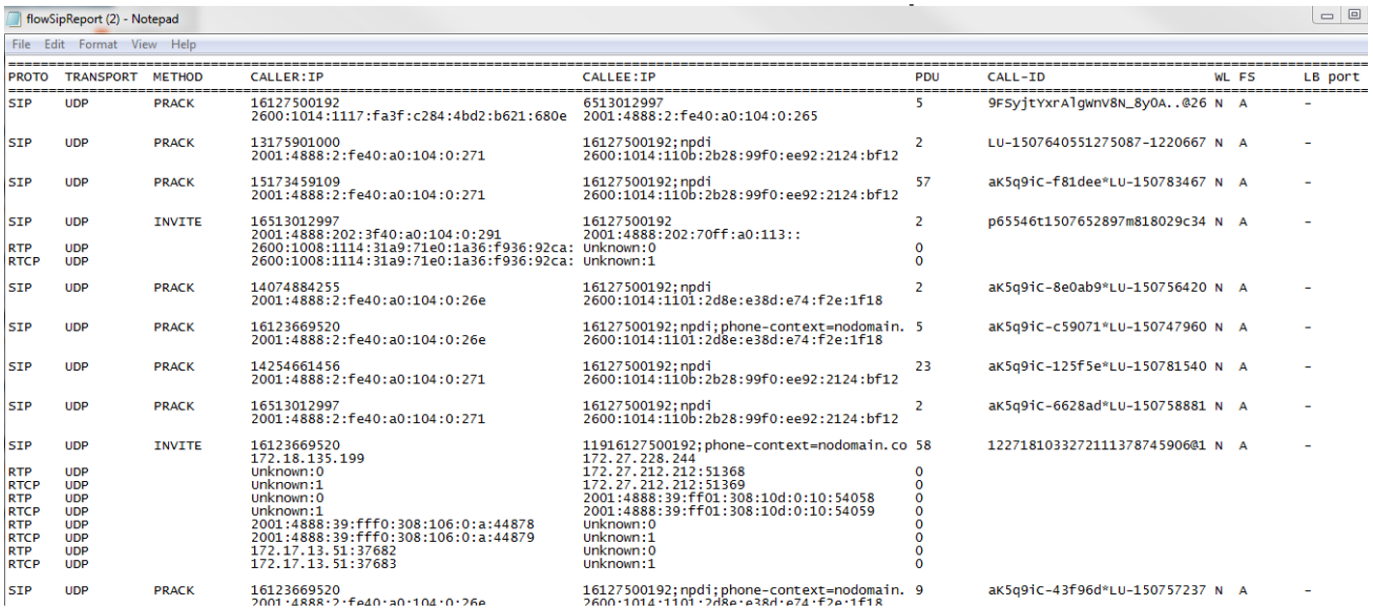

flowSipReport (2) - Notepad

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_8y0A..@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.

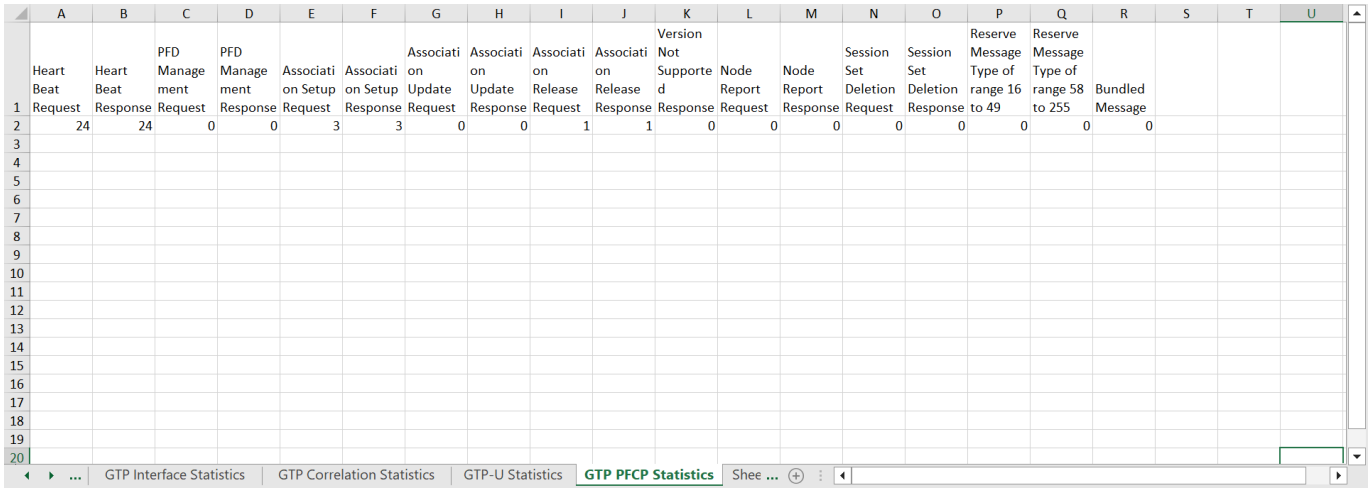

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U
1	Heart Beat Request	Heart Beat Response	PFD Management Request	PFD Management Response	Association Setup Request	Association Setup Response	Association Update Request	Association Update Response	Association Release Request	Association Release Response	Version 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			
3																					
4																					
5																					
6																					
7																					
8																					
9																					
10																					
11																					
12																					
13																					
14																					
15																					
16																					
17																					
18																					
19																					
20																					

... 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.

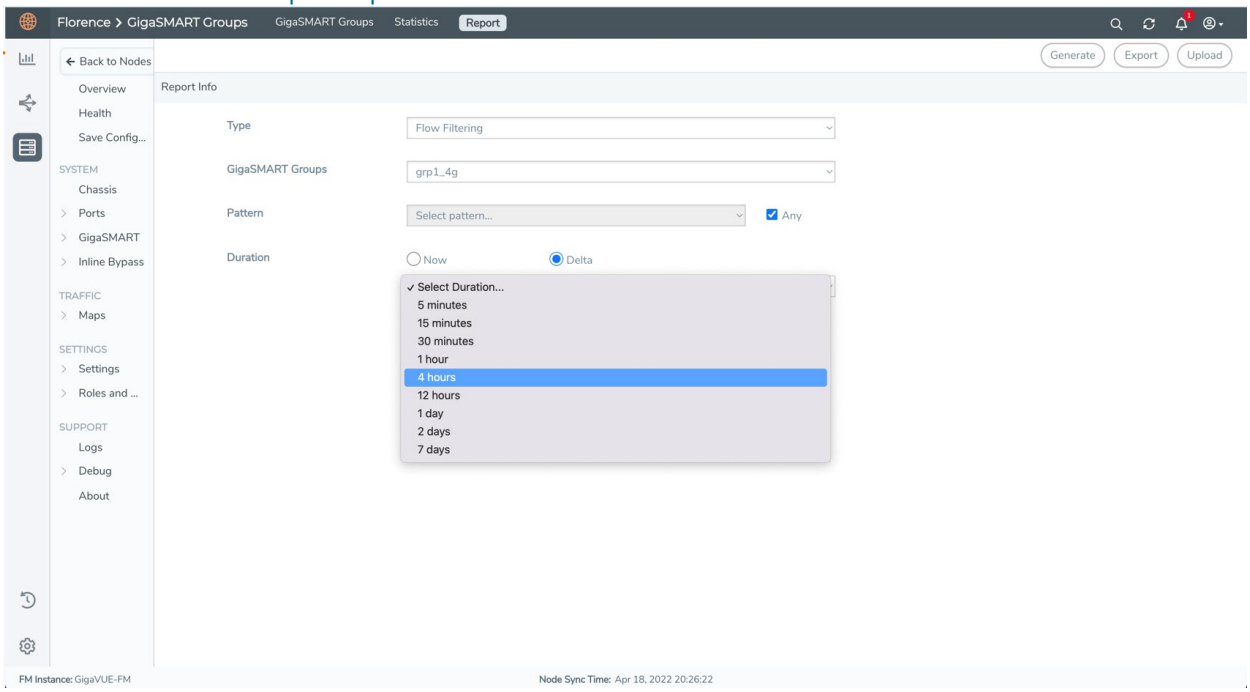

3. Click **Generate**. The delta report for the selected duration is displayed. The following details are displayed as part of the report:
   - Duration of the delta report in the page header
   - Detailed timestamps as part of the report header
   - Within the report, the following details are displayed:
     - **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:

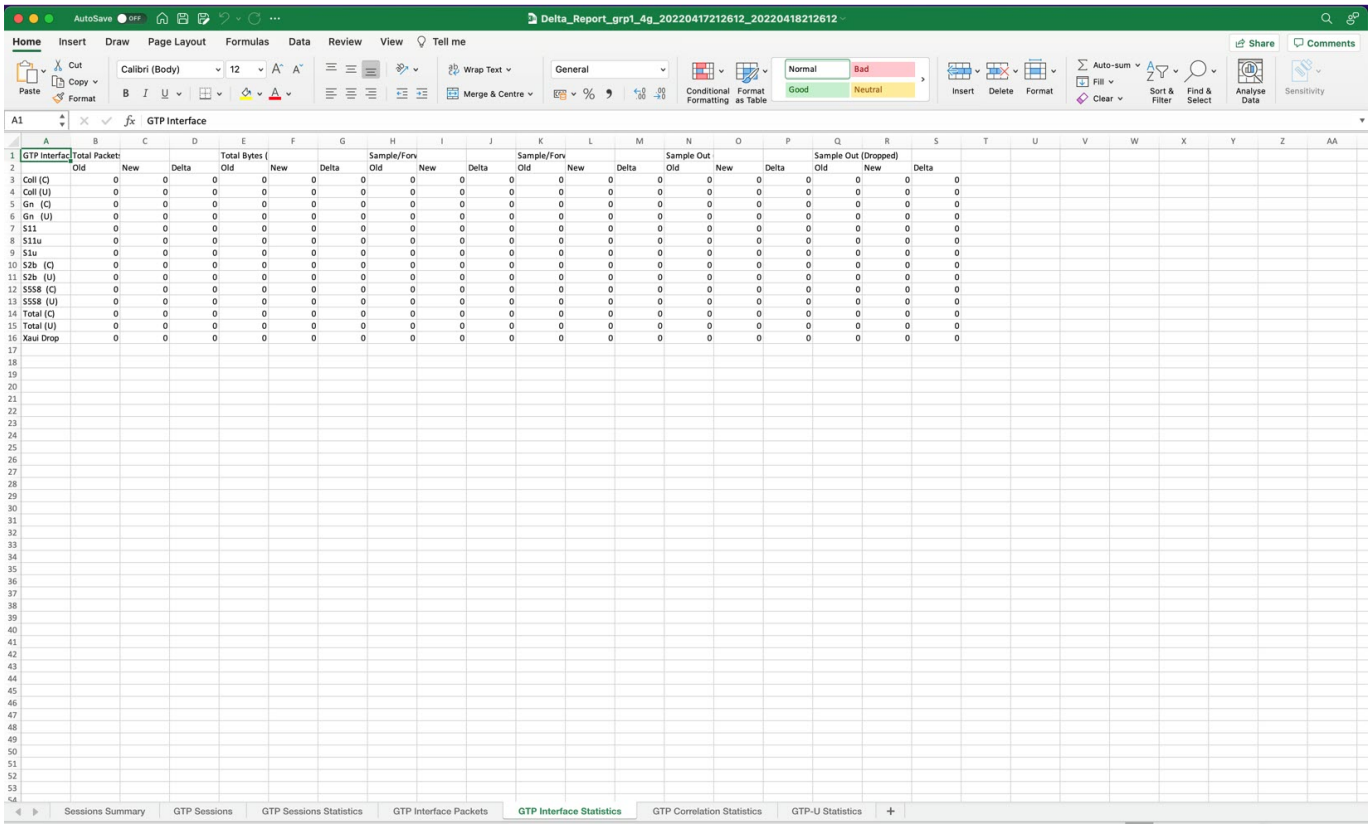

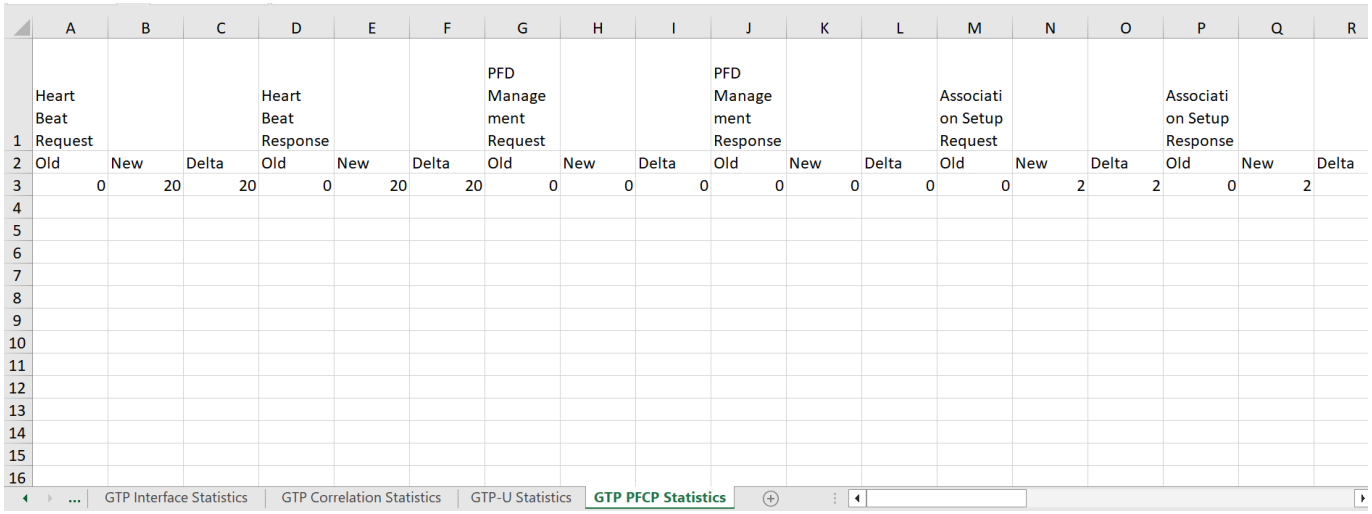

GTP Overlap Flow Sampling Maps

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 28GTP 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.

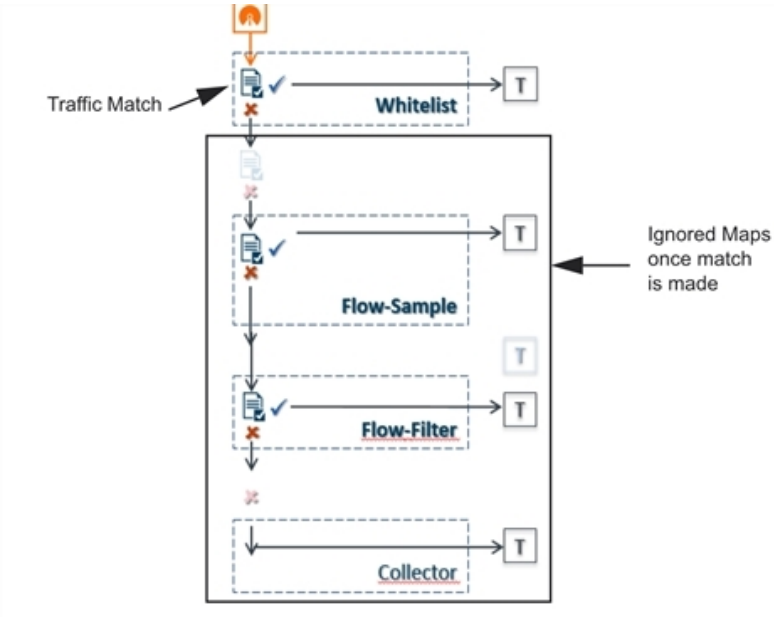

Figure 27 *Non-Overlap GTP Mapping Mode*

This contrasts with GTP overlap flow sampling maps. In Figure 27Non-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.

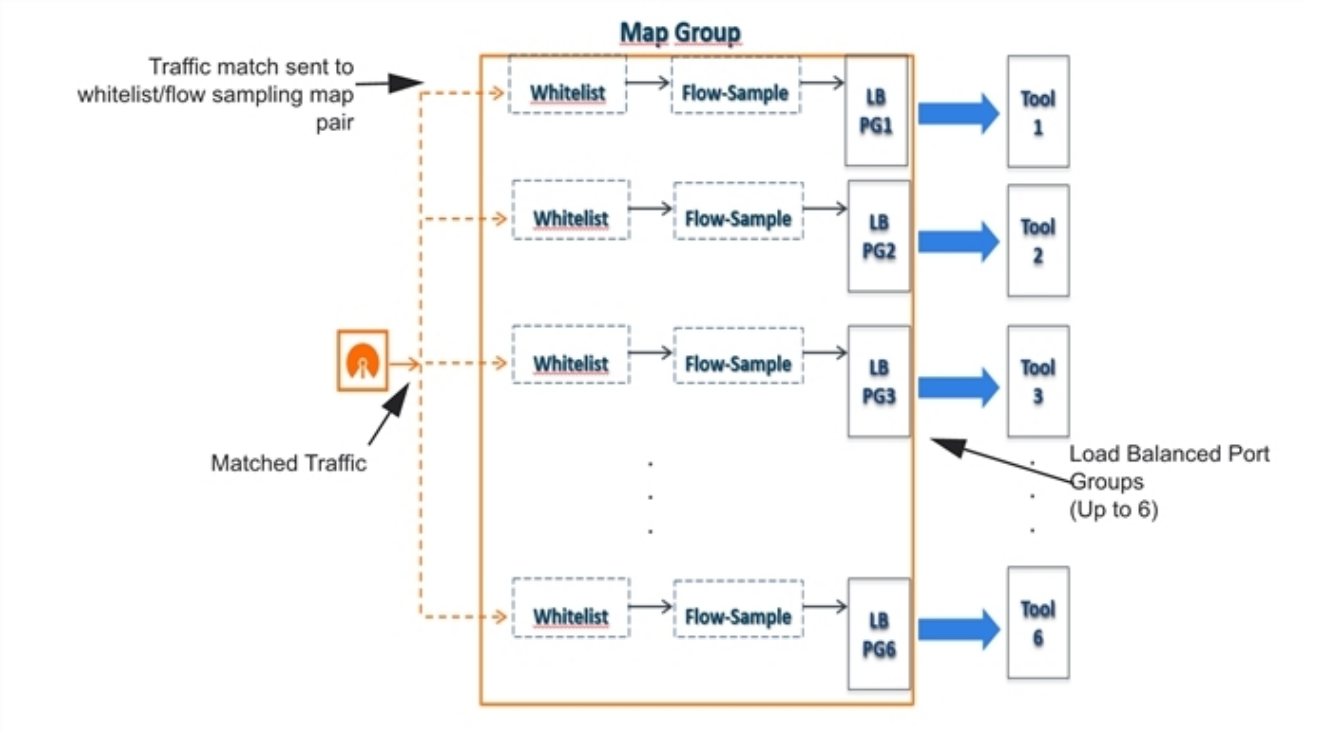

Figure 28 *GTP Overlap Flow Sampling Map Groups*

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

Use the following sequence to help you determine the maximum number of port group members:

1. Determine the number of members per port group and add 1 to the number.
2. Multiply each port group result times each other.
3. The total multiplication should not exceed 512.

For instance, assume the following configuration in a GTP overlap mapping group:

- Port Group 1—2 load balanced GigaStream
- Port Group 2—3 load balanced GigaStream
- Port Group 3—1 load balanced tool port
- Port Group 4—1 load balanced GigaStream
- Port Group 5—4 load balanced tool ports

The total number becomes:

(2+1)*(3+1)*(1+1)*(1+1)*(4+1) = 240

Since this does not exceed the maximum number of multicast IDs (512), the tool configuration shown is accepted.

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 29GTP Overlap Mode Example 1).

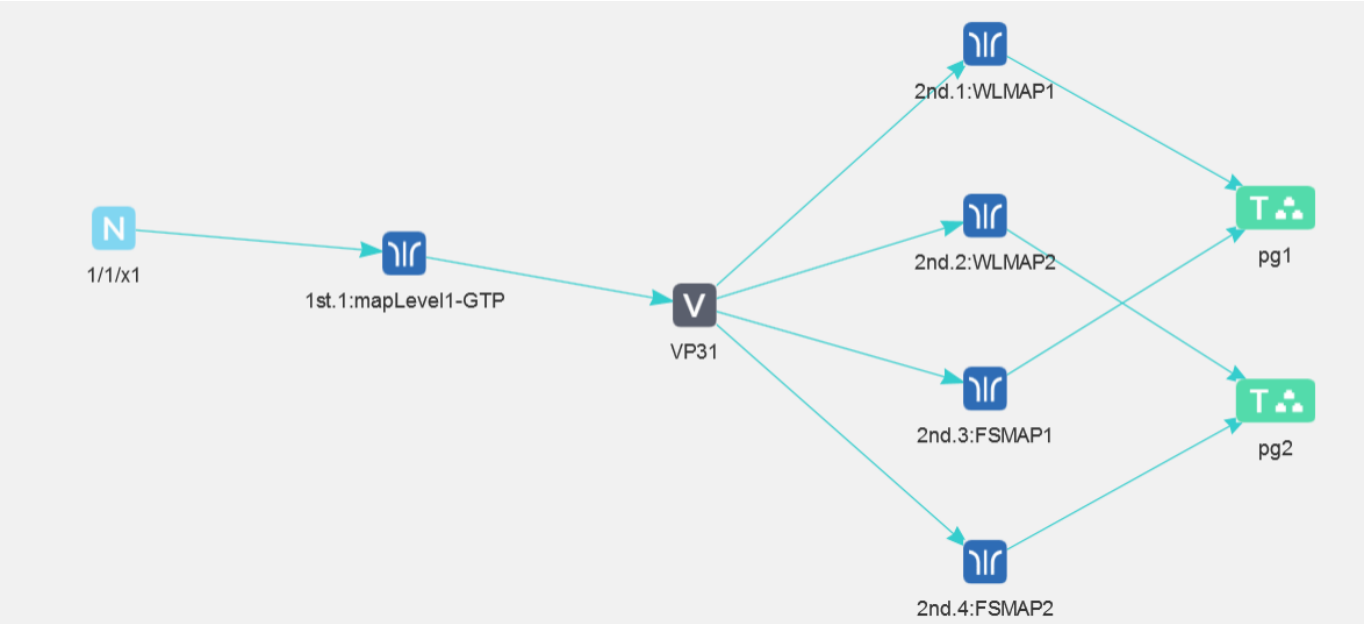

Figure 29 *GTP Overlap Mode Example 1*

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 persistence, as well as to configure timers, use to the GTP Persistence fields under GigaSMART Parameters on the GigaSMART Group configuration page shown in Figure 30GTP Persistence GigaSMART Parameters and select **GTP Persistence** The timers are preconfigured with default values.

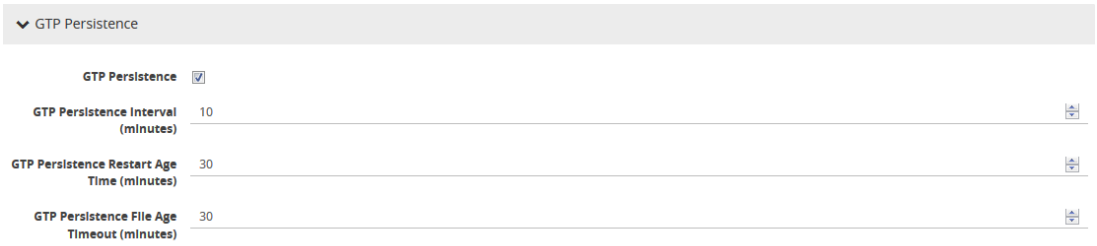

Figure 30 *GTP Persistence GigaSMART Parameters*

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 31GTP Persistence Information, GigaSMART Group gsgrp-1_4_e1 is selected and the Quick View displayed.

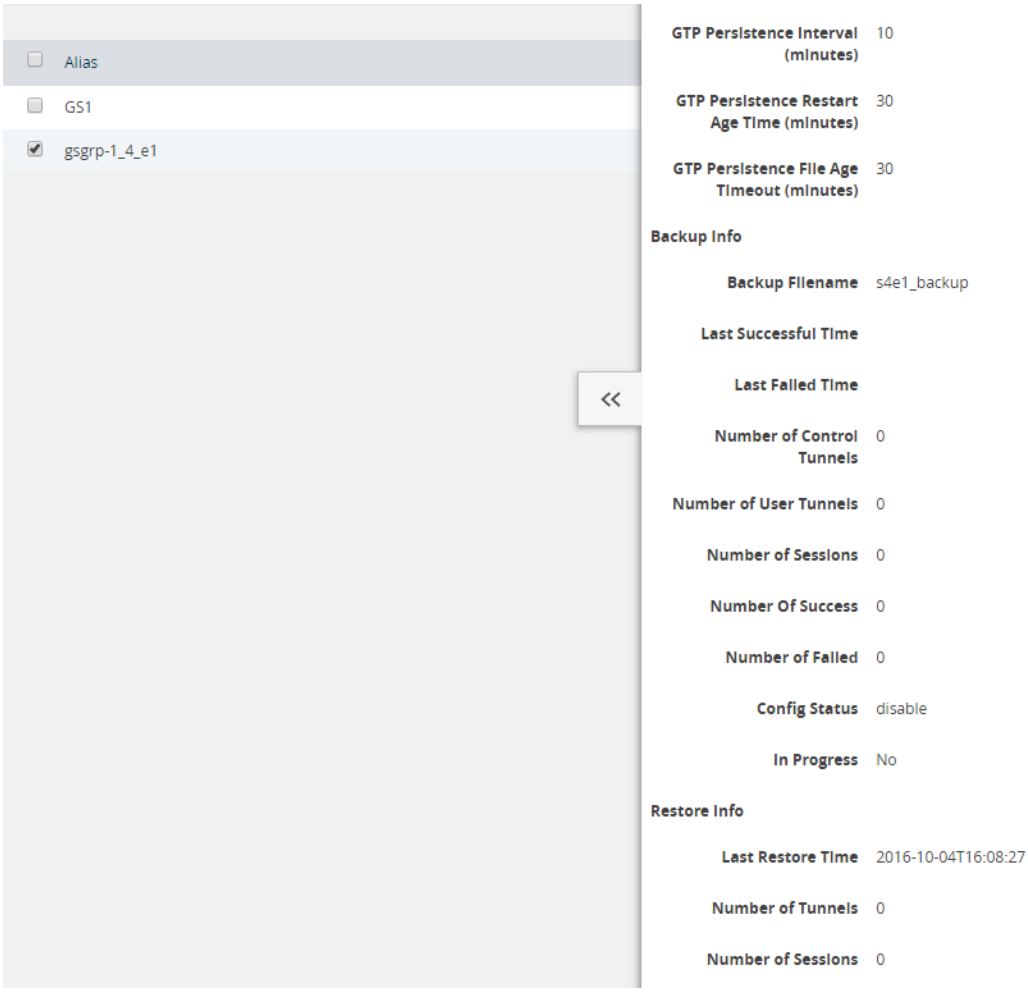

Figure 31 *GTP Persistence Information*

The following table describes persistence information.

Table 6: 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.
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 32GTP Backup Files Delete) and click **Delete All** under **GTP Backup Files**.

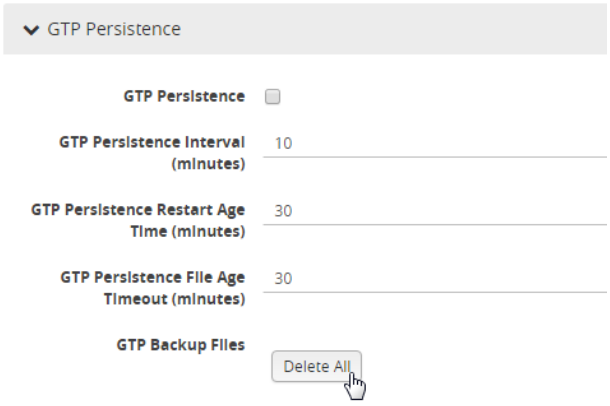

Figure 32 *GTP Backup Files Delete*

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.

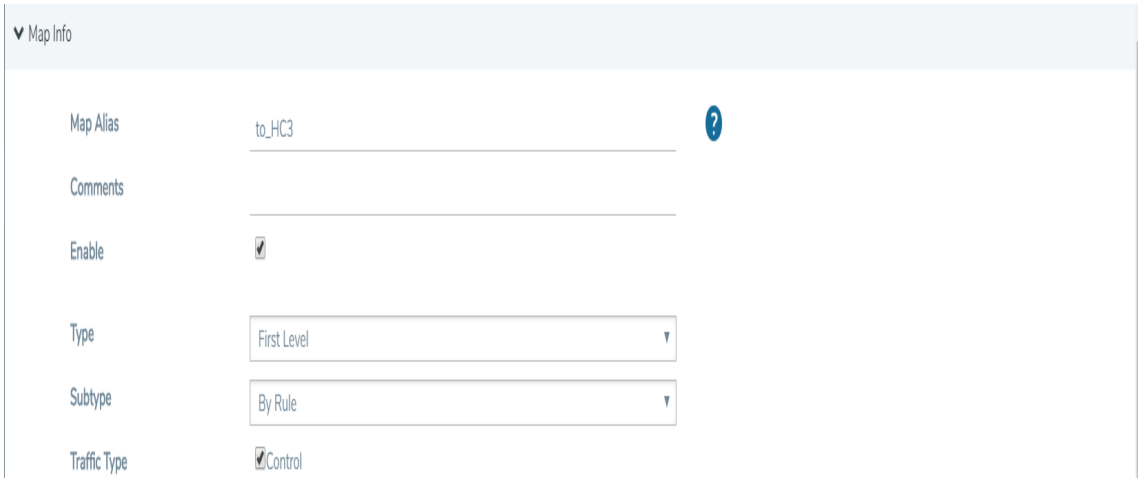

Upgrade from Earlier Release

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.

Task	Description	UI Steps
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 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.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ 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.

Task	Description	UI Steps
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. **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

Task	Description	UI Steps
		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 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.	**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. **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.

Task	Description	UI Steps
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**. ▪ Four tool type of ports for a GigaStream that will be created in **Task 4**. Then administratively enable the ports.	**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. **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-	**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.

Task	Description	UI Steps
	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**	**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. **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	**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.

Task	Description	UI Steps
	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 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**.	**f.** Click Save. **7.** Configure a Flow Sampling operation. **a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **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	**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.

Task	Description	UI Steps
	grouping. In addition to the virtual port, traffic is also sent to a GigaStream and a tool port.	▪ Select the GigaStream port hc3-gs-1, the virtual port p-hc3scale-4engines-slots1and3, and a tool port (for example, 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.

Task	Description	UI Steps
		▪ Select hc3-scale-wl from the GSOP list. **3.** Click Save.
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.

Task	Description	UI Steps
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-q2x32-1-4.	**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 ▪ 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.

Task	Description	UI Steps
		3. Click Save.

GigaSMART TLS/SSL Decryption for Inline and Out-of-Band Tools

Required License: TLS/SSL Decryption for Inline and 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 for more details on the devices that support GigaSMART operations.

TLS/SSL decryption for inline and out-of-band tools is described in Inline TLS/SSL Decryption and GigaSMART Passive TLS/SSL Decryption. It is supported on all GigaVUE-HC1, GigaVUE-HC1-Plus, and GigaVUE-HC3 devices.

Limitations

Flexible Inline SSL solution does not support the following:

- Inline-Netlag as source
- Iboss and resilient hashing
- Cluster

Inline TLS/SSL Decryption

This chapter describes TLS/SSL decryption for inline and out-of-band tools, referred to as inline TLS/SSL decryption. It provides introductory material as well as configuration examples for GigaVUE-FM.

In this section:

- About Inline TLS/SSL Decryption
- Inline TLS/SSL Decryption Deployments
- Get Started with Inline TLS/SSL Decryption

- Configure Inline TLS/SSL Decryption
- Resilient Inline Arrangement with GigaSMART Flex Inline Solution

Refer Active / Standby Resilient Inline SSL Solution using GRIP for more detailed information.

About Inline TLS/SSL Decryption

This section introduces inline TLS/SSL decryption.

What Inline TLS/SSL Decryption Provides

Inline TLS/SSL decryption provides the following:

- Identifies/detects encrypted traffic flows (TLS/SSL traffic) in a network across any port.
- Intercepts encrypted traffic flows between a client and a server.
- Filters encrypted traffic flows based on policy. For example, if the encrypted traffic flows contain health care or financial information, let those flows bypass decryption.
- Decrypts packets. Inline TLS/SSL decryption decrypts packets once at a single decryption point.
- Delivers decrypted traffic flows to multiple security tools. The tools can be inline or out-of-band. The tools can detect threats such as malware in the decrypted traffic flows.
- Re-encrypts traffic flows after receiving them back from the inline tools.
- If a tool acts on traffic flows based on the threats it finds, when malware is found in the decrypted traffic flows, the tool can:
  - modify the traffic flows
  - terminate the connection
- If the tool modifies the packets, GigaSMART will re-encrypt them. If the tool terminates the connection, GigaSMART will terminate the connection between the client and the server.

When TLS/SSL traffic is decrypted, sensitive data will be exposed in the connected tools. For example, if email traffic is decrypted, user passwords might be exposed or if financial data is decrypted, social security numbers might be exposed in the decrypted traffic.

Because TLS/SSL connections might carry sensitive data, not all connections should be inspected. Some of the TLS/SSL connections carrying user data such as financial or medical information should be bypassed without inspection, based on a configured policy.

Inline TLS/SSL decryption addresses acceptable use policies and adheres to privacy and compliance requirements. It offers advanced controls to select the traffic to decrypt.

Inline TLS/SSL supports the following applications:

- HTTPS
- FTPS
- StartTLS can be used to decrypt SMTP, IMAP, and POP3 (refer to StartTLS and HTTP CONNECT.

Important Inline TLS/SSL Rules and Notes

This following list describes important caveats and limitations of working with Inline TLS/SSL:

- Clustering is not supported with inline TLS/SSL.
- IPV6 traffic decryption is supported only for GEN 3 cards. Refer to the GigaVUE-HC1 Hardware Installation Guide and GigaVUE-HC3 Hardware Installation Guide for the list of GEN 3 card numbers.
- IPV6 traffic decryption is not supported for the One-Arm mode.
- Gigamon Resiliency for Inline Protection (GRIP) support with inline TLS/SSL can be referred through the validated design Active Standby Resilient Inline SSL Solution using GRIP (5.10.01).
- Resilient Inline Arrangements (RIA) configurations is supported with inline SSL for asymmetric traffic visibility with TLS/SSL Decryption.

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.

Supported Ciphers

Inline TLS/SSL decryption supports modern cryptographic algorithms. It supports the commonly-supported TLS 1.2 and TLS 1.3 ciphers.

Combining the following ciphers, MACs, and Key Exchange Algorithms results in many cipher suites:

- Ciphers: AES_128_CBC, AES_128_GCM, AES_256_GCM, AES_256_CBC, Camellia, Chacha20
- MAC: SHA, SHA256, SHA384, Poly1305
- Key Exchange Algorithms: RSA, DHE_RSA, ECDHE_RSA, ECDHE_ECDSA.

Inline TLS/SSL Supported TLS 1.3 Ciphers

Cipher Name	Encryption (Enc)	MAC
TLS_AES_256_GCM_SHA384	AES_ 256_GCM	SHA384
TLS_CHACHA20_POLY1305_SHA256	CHACHA20_ POLY1305	SHA256
TLS_AES_128_GCM_SHA256	AES_128_GCM	SHA256

Inline TLS/SSL Supported TLS 1.2 Ciphers

Cipher Name	Key Exchange (Kx)	Authentication (Au)	Encryption (Enc)	MAC
TLS_DHE_RSA_WITH_AES_128_CBC_SHA	DHE	RSA	AES128_CBC	SHA
TLS_DHE_RSA_WITH_AES_256_CBC_SHA	DHE	RSA	AES256_CBC	SHA
TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA	DHE	RSA	CAMELLIA128	SHA
TLS_RSA_WITH_CAMELLIA_128_CBC_SHA	RSA	RSA	CAMELLIA128	SHA
TLS_RSA_WITH_CAMELLIA_256_CBC_SHA	RSA	RSA	CAMELLIA256	SHA
TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA	DHE	RSA	CAMELLIA256	SHA
TLS_DHE_RSA_WITH_AES_128_CBC_SHA256	DHE	RSA	AES128_CBC	SHA256
TLS_DHE_RSA_WITH_AES_256_CBC_SHA256	DHE	RSA	AES256_CBC	SHA256
TLS_DHE_RSA_WITH_AES_128_GCM_SHA256	DHE	RSA	AES128_GCM	SHA256
TLS_DHE_RSA_WITH_AES_256_GCM_SHA384	DHE	RSA	AES256_GCM	SHA384
TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305	ECDHE	RSA	CHACAH20	POLY1305
TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305	ECDHE	ECDSA	CHACAH20	POLY1305
TLS_DHE_RSA_WITH_CHACHA20_POLY1305	DHE	RSA	CHACAH20	POLY1305
TLS_RSA_WITH_AES_128_CBC_SHA	RSA	RSA	AES128_CBC	SHA
TLS_RSA_WITH_AES_256_CBC_SHA	RSA	RSA	AES256_CBC	SHA
TLS_RSA_WITH_AES_128_CBC_SHA256	RSA	RSA	AES128_CBC	SHA256
TLS_RSA_WITH_AES_256_CBC_SHA256	RSA	RSA	AES256_CBC	SHA256
TLS_RSA_WITH_AES_128_GCM_SHA256	RSA	RSA	AES128_GCM	SHA256
TLS_RSA_WITH_AES_256_GCM_SHA384	RSA	RSA	AES256_GCM	SHA384

Cipher Name	Key Exchange (Kx)	Authentication (Au)	Encryption (Enc)	MAC
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA	ECDHE	ECDSA	AES128_CBC	SHA
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA	ECDHE	ECDSA	AES256_CBC	SHA
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	ECDHE	RSA	AES128_CBC	SHA
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	ECDHE	RSA	AES256_CBC	SHA
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	ECDHE	ECDSA	AES128_CBC	SHA256
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384	ECDHE	ECDSA	AES256_CBC	SHA384
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	ECDHE	RSA	AES128_CBC	SHA256
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	ECDHE	RSA	AES256_CBC	SHA384
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	ECDHE	ECDSA	AES128_GCM	SHA256
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	ECDHE	ECDSA	AES256_GCM	SHA384
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	ECDHE	RSA	AES128_GCM	SHA256
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	ECDHE	RSA	AES256_GCM	SHA384

Diffie Hellman Ephemeral (DHE) is a key exchange protocol.

Inline TLS/SSL decryption supports key cipher suites and exchanges without downgrading cryptography levels of the organization.

Ciphersuites are a standard combination of the following:

- **bulk encryption algorithm**—Specifies how to encrypt communications, including the algorithm, key size, and the cryptographic mode used. For example, AES_128_CBC is AES with 128-bit keys in Cipher Block Chaining mode.
- **key exchange algorithm**—Specifies how both sides authenticate each other during the TLS/SSL handshake. For example, RSA.
- **message authentication code (MAC)**—Specifies the hash algorithm used to verify that communications have not been tampered with. For example, SHA.

- **pseudorandom function**—Specifies how a 384-bit master secret, which is used as a source of randomness for session keys, is generated.

NOTE

- TLS/SSL transactions with unsupported ciphers will be bypassed/TCP proxied.
- The new TLS1.3 cipher suites are defined differently and do not specify the certificate types (RSA/DSA/ECDSA) or the key exchange mechanism (DHE/ECHDE).
- The Inline TLS/SSL session is now equipped to receive a client hello with the key exchange X25519Kyber768 and now fall back to using just X25519. This ensures the system maintains secure and functional connections, even if it cannot use the newer, quantum-resistant algorithm now.

The following key sizes are supported:

- **RSA**—2048, 3072, 4096, 8192
- **DH**—1024, 2048, 4096
- **ECC**—prime256v1, ecsecp256r1, ecsecp384r1, ecsecp521r1, secp160k1, secp160r1, secp160r2, secp192k1, secp192r1, secp224k1, secp224r1, secp256k1, secp256r1, secp384r1, secp521r1, brainpoolP256r1, brainpoolP384r1, brainpool512r1, X25519, X448

The following TLS extension is supported:

- **RFC7301**—Application-Layer Protocol Negotiation (ALPN)

More About Inline TLS/SSL Decryption

Refer to the following sections for additional details:

- TLS/SSL Decryption for Inline Tools
- Example Inline TLS/SSL Decryption
- TLS/SSL Sessions
- TLS/SSL Terminology and Acronyms

TLS/SSL Decryption for Inline Tools

TLS/SSL decryption for inline tools provides visibility into encrypted traffic. Inline TLS/SSL decryption delivers decrypted packets to tools that can be placed inline or out-of-band. The tools look into decrypted packets for threats, such as viruses or other malware.

The amount of Internet traffic that is encrypted is increasing, and much of it is encrypted with Secure Sockets Layer (SSL) or Transport Layer Security (TLS) protocols.

> **NOTE:** Throughout this document, the terms Secure Sockets Layer (SSL) and Transport Layer Security (TLS) are used interchangeably.

Malware increasingly uses encrypted TLS/SSL traffic, thus a significant percentage of attacks hide in TLS/SSL. Inline TLS/SSL decryption offers visibility into encrypted applications and hidden threats in your organization.

Many applications, such as email, also use TLS/SSL. Encryption protects data from being viewed in transit over the Internet such as in an exchange of emails. Encryption also keeps the data private. But when data is encrypted, packets are not inspected, which can create blind spots in your network.

Providing visibility into encrypted traffic eliminates this blind spot. SSL/TLS blind spots in your network can be eliminated across any port or application, for example, port 443, or email, Web, or VoIP applications.

Inline TLS/SSL decryption differs from the existing GigaSMART SSL/TLS decryption application, which is passive. Passive TLS/SSL decryption delivers decrypted traffic to out-of-band tools that can then detect threats entering the network. When a threat is detected, the tools can send a notification to the user.

Inline TLS/SSL decryption offloads the decryption task so that tools can inspect traffic easily and effectively. The advantage of operating inline is that tools can act when a threat is detected.

Inline TLS/SSL decryption supports TLS/SSL version 3.0 and TLS versions 1.0, 1.1, 1.2, and 1.3.

Also, the inline TLS/SSL decryption solution is able to decrypt Perfect Forward Secrecy (PFS) ciphers, for example, ECDHE-RSA-AES256-SHA384 and DHE-RSA-AES128-SHA256.

Example Inline TLS/SSL Decryption

Refer to Figure 33Inline TLS/SSL Decryption Example, Outbound for an example of inline TLS/SSL decryption.

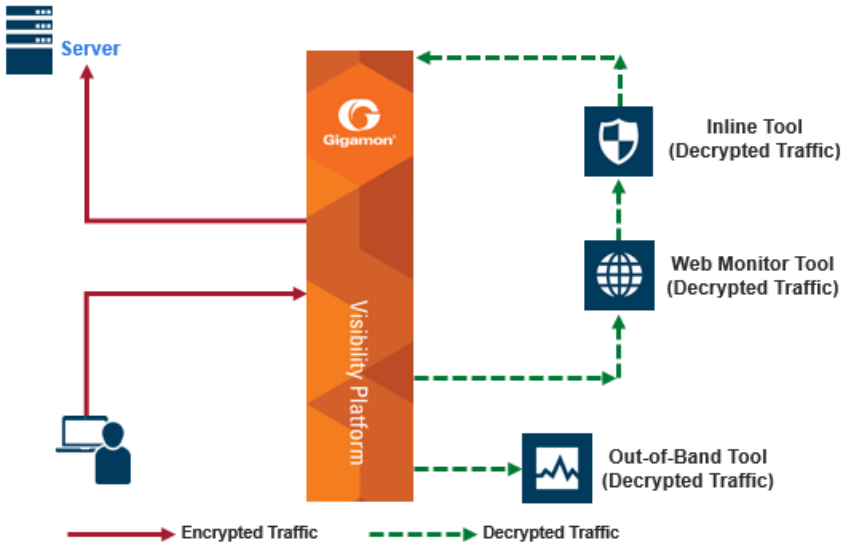

Figure 33 *Inline TLS/SSL Decryption Example, Outbound*

Use case for inline TLS/SSL decryption:

- Clients in internal network
- Servers on the Internet
- Organization does not have the private key of the server
- Diffie-Hellman and Perfect Forward Secrecy is being used

Figure 33Inline TLS/SSL Decryption Example, Outbound shows the inline TLS/SSL decryption solution with the client and GigaVUE node within an enterprise. The server, in the top left of the figure, resides on the Internet. This is an example of an outbound deployment.

The client is on the lower left of the figure. The client is a user who is, for example, using a browser to go to a website on the Internet, such as a bank or a search engine. The traffic from the user could be encrypted or it might not be encrypted. For example, the user might be going to a bank website using the HTTPS protocol or going to a search engine website using the HTTP protocol. The solid line from the user to the GigaVUE node represents encrypted traffic, but there might also be traffic from the user that is not encrypted. Traffic that is not encrypted can either be bypassed or it can go to tools for inspection.

In Figure 33Inline TLS/SSL Decryption Example, Outbound, instead of the user interacting directly with the server at the top left, the GigaVUE node is placed in the middle. Thus, the GigaVUE node intercepts the client/server session.

In the GigaVUE node, encrypted packets are identified, then filtered. Selected packets are decrypted and sent to tools for inspection. The dotted lines represent decrypted packets. Packets are decrypted once, then the same decrypted packets can be sent to inline tools and/or out-of-band tools connected to the GigaVUE node.

The traffic from the inline tools is returned to the GigaVUE node to be re-encrypted and then sent to the destination on the Internet, for example, the website that the user is visiting. The solid line to the server represents traffic that has been re-encrypted in the GigaVUE node.

TLS/SSL Terminology and Acronyms

Table 7: TLS/SSL Terminology provides definitions of TLS/SSL terminology:

Table 7: TLS/SSL Terminology

Term	Definition
Plaintext	The original, unencrypted data.
Ciphertext	The encrypted data.
Cryptography	The practice of secure communications.
Encryption	The process of turning plaintext into ciphertext.
Decryption	The process of turning ciphertext into plaintext.
Encryption algorithm	The algorithm used to perform encryption and decryption. It is also called the cipher.
Encryption key	The key used for encryption.
Decryption key	The key used for decryption.
Symmetrical encryption algorithm	The algorithm used for encryption in which the encryption key and the decryption key are identical.
Asymmetrical encryption algorithm	The algorithm used for encryption in which the encryption key and the decryption key are different.
Public key	The key used for encryption.
Private key	The key used for decryption.

Table 8: TLS/SSL Acronyms lists TLS/SSL acronyms:

Table 8: TLS/SSL Acronyms

Acronym	Definition
AES	Advanced Encryption Standard

Acronym	Definition
CA	Certificate Authority
CBC	Cipher Block Chaining
CN	Common Name
CRL	Certificate Revocation List
DES	Data Encryption Standard
DH, D-H	Diffie-Hellman
DSA	Digital Signature Algorithm
ECDH	Elliptic Curve Diffie-Hellman
ECDSA	Elliptic Curve Digital Signature Algorithm
FQDN	Fully Qualified Domain Name
IDS	Intrusion Detection System
IPS	Intrusion Prevention System
MAC	Message Authentication Code
MD	Message Digest
MitM	Man-in-the-Middle
OCSP	Online Certificate Status Protocol
OoB	Out-of-Band
PEM	Privacy Enhanced Mail
PFS	Perfect Forward Secrecy
PKCS12	Public Key Cryptography Standard #12
PKI	Public Key Infrastructure
RSA	Rivest-Shamir-Adleman
SHA	Secure Hash Algorithm
SSL	Secure Sockets Layer
TCP	Transmission Control Protocol
TLS	Transport Layer Security
URL	Uniform Resource Locator

InlineTLS/SSL show command Field Descriptions

The Inline SSL show command fields have the below descriptions and values. Refer to apps inline-ssl to know more about GigaVUE-OS CLI configurations.

Fields	Description	Values
Source IP	Source IP of the Host/Client from where connection was initiated.	IPv4 or IPv6 address.
Destination IP	Destination IP of the Host/Server at which connection is terminated.	IPv4 or IPv6 address.
Source Port	Source Port at client side.	TCP port range (0 to 65535).
Destination Port	Destination Port at the server side.	TCP port range (0 to 65535).
Protocol	Inline SSL Session's protocol details.	TLS/SSL Inbound -Inbound TLS/SSL session TLS/SSL Outbound -Outbound TLS/SSLsession. Non-SSL- Not an SSL session, need not be decrypted.
TCP State	TCP state of all the legs/ports of ISSL such as: • Na - Inline Network Port A. • Nb - Inline Network Port B. • Ta - Inline Tool Port A. • Tb - Inline Tool Port B..	All Possible TCP protocol state for each leg: • INIT – TCP connection is yet to be started from GigaSMART engine. • SYN-SENT – GigaSMART engine has initiated the TCP connection to the server. • SYN-RECEIVED – GigaSMART engine has received a TCP connection request from a client. • ESTABLISHED – GigaSMART engine has established an ongoing TCP connection. • FIN-WAIT-1 – GigaSMART engine has initiated connection termination request • FIN-WAIT-2 – GigaSMART engine has received an acknowledgement for its connection termination request, awaiting for connection termination from other end. • CLOSE-WAIT – GigaSMART engine has received connection

Fields	Description	Values
		termination request. • CLOSING – GigaSMART engine is awaiting acknowledgment for its connection termination request. • LAST-ACK – GigaSMART engine has acknowledged connection termination request and sent the connection termination request, and is awaiting for the last acknowledgement before closing. • TIME-WAIT – GigaSMART engine waiting for the other host to receive the last acknowledgement sent. • CLOSED – TCP connection is closed at GigaSMART engine. N/A – TCP not expected to be performed at this leg for this session
Decryption	Session traffic is being decrypted.	Yes, No
Error	Specifies if this session processing has any error.	ERR, NO_ERR
SSL State	State of SSL processing.	• N/A, • Handshake, • Decrypting, • Bypass:no_config, • Bypass:clientHello_err, • Bypass:no_sni, • Bypass:no_cert, • Drop:cert_mismatch, • Drop:tool_block, • Bypass:policy, • Bypass:handshake_fail, • Bypass:unknown_revocation, • Bypass:client_auth, • Bypass:version_mismatch, • Bypass:src_ssl_init_error, • Bypass:ssl_key_mismatch, • Drop:version_mismatch, • Bypass:unsupported_ciphers, • SSL_Proxy:non_http
Cert Subject Name	Client's Certificate subject name.	
C2S Status	Client to Server TCP state details.	All Possible TCP protocol state: (For detailed explanation, refer previous explanation on TCP states on the

Fields	Description	Values
		table) INIT, CLOSED, SYN SENT, SYN RECD, EST, CLOSE WAIT, FIN WAIT 1, CLOSING, LAST ACK, FIN WAIT 2, TIME WAIT, N/A
S2C Status	Server to Client TCP state details.	All Possible TCP protocol state: (For detailed explanation, refer previous explanation on TCP states on the table) INIT, CLOSED, SYN SENT, SYN RECD, EST, CLOSE WAIT, FIN WAIT 1, CLOSING, LAST ACK, FIN WAIT 2, TIME WAIT, N/A
Tool Status	Status of Tool connection status.	• TOOL_NOT_BYPASS – Traffic will flow through tools. • TOOL_BYPASS – Traffic will skip tools.
PolicyVerdict	ISSL policy decision determined for this session.	• NO_DECRYPT • PENDING • DECRYPT
PolicyMatchFields	The fields through which policy matching is derived.	• DEFAULTSRC_IP • DST_IP • SRC_PORT • DST_PORT • VLAN • DOMAIN • CATEGORY • ISSUER • URL_CACHE_MISS • PENDING
URLCategory	Category of the URL based on the certificate of the server.	• Uncategorized • RealEstate • ComputerandInternetSecurity • FinancialServices • BusinessandEconomy • ComputerandInternetInfo • Auctions • Shopping • CultandOccult • Travel • AbusedDrugs • AdultandPornography • Home

Fields	Description	Values
		• Military • SocialNetwork • DeadSites • IndividualStockAdviceandTools • TrainingandTools • Dating • SexEducation • Religion • EntertainmentandArts • PersonalSitesandBlogs • Legal • LocalInformation • StreamingMedia • JobSearch • Gambling • Translation • ReferenceandResearch • SharewareandFreeware • P2P • Marijuana • Hacking • Games • PhilosophyandPoliticalAdvocacy • Weapons • PaytoSurf • HuntingandFishing • Society • EducationalInstitutions • OnlineGreetingcards • Sports • Swimsuits&IntimateApparel • Questionable • Kids • HateandRacism • OnlinePersonalStorage • Violence • KeyloggersandMonitoring • SearchEngines • InternetPortals • WebAdvertisements • Cheating • Gross

Fields	Description	Values
		• WebbasedEmail • MalwareSites • PhishingandOtherFrauds • ProxyAvoidandAnonymizers • SpywareandAdware • Music • Government • Nudity • NewsandMedia • Illegal • CDNs • InternetCommunications • BotNets • Abortion • Health&Medicine • ConfirmedSPAMSources • SPAMURLs • UnconfirmedSPAMSources • OpenHTTPProxies • DynamicContent • ParkedSites • AlcoholandTobacco • PrivateIPAddresses • ImageandVideoSearch • FashionandBeauty • RecreationandHobbies • MotorVehicles • WebHosting • N/A • Unknown
URLFilterResult	Based on URL filter what decision was made for this session	
InterfacePair	Interface network pair used for this session	Na/Nb, Nb/Na
ToolInterfacePair	Interface tool pair used for this session	Ta/Tb, Tb/Ta
StartTime	Start time of the session	
EndTime	End time of the session	
DurationSeconds	Seconds for which session was served	
SSLVersion	SSL Version used	SSLv1.2, TLSv1, TLSv1.2, TLSv1.3.
SSLCipher	SSL Cipher used for encryption/decryption	All Possible standard ciphers.

Fields	Description	Values
CertIssuer	Certificate Issuer of client certificate	
CertValidation	Certificate validation status.	• VALID • INVALID • SELF_SIGNED • HOST_MISMATCH • EXPIRED • REVOKED • UNKNOWN_CA • N/A
ClitoServTotalPktcount	Total Packet count from client to server.	
ServtoCliTotalPktcount	Total Packet count from server to client.	
ClitoServClearTextByte	Total Decrypted bytes from client to server.	
ServtoCliClearTextByte	Total Decrypted bytes from server to client.	
ClitoServOriginalByte	Total Bytes transferred from client to server.	
ServtoCliOriginalByte	Total Bytes transferred from server to client.	
SSLflags(devuse)	This is for debugging purpose, can only be determined with the knowledge of code of that particular version.	

TLS/SSL Sessions

Secure Sockets Layer (SSL) is a protocol that allows the transmission of secure data between a server and client. Transport Layer Security (TLS) is a cryptographic protocol that adds security to TCP/IP communication.

Inline TLS/SSL decryption supports SSL version 3.0 and TLS versions 1.0, 1.1, 1.2, and 1.3.

TLS and SSL are used in communications such as Web browsing, email, instant messaging, and voice over IP (VoIP). TLS and SSL encrypt these communications.

The client initiates the TLS/SSL session. The GigaVUE node intercepts the connection and negotiates an TLS/SSL session with the client.

The GigaVUE node monitors all TCP connections, then intercepts the TLS/SSL session. Non-TCP traffic is passed transparently without any changes.

All the incoming TLS/SSL traffic terminates on the GigaVUE node. The TLS/SSL connections are decrypted in inbound or outbound deployments, passed to the inline tools, and eventually to the server.

The session to the client is terminated on the GigaVUE node, but information about the session, such as the initiator's IP address is maintained, so that the GigaVUE node can "reconnect" the client and server.

The GigaVUE node performs TLS/SSL decryption and feeds tools, either inline or out-of-band.

The session to the server is from the GigaVUE node to the server. The GigaVUE node negotiates a new TLS/SSL session with the server.

TLS/SSL Handshake

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 34Basic RSA TLS/SSL Handshake shows the basic TLS/SSL handshake between a client and server to establish a session.

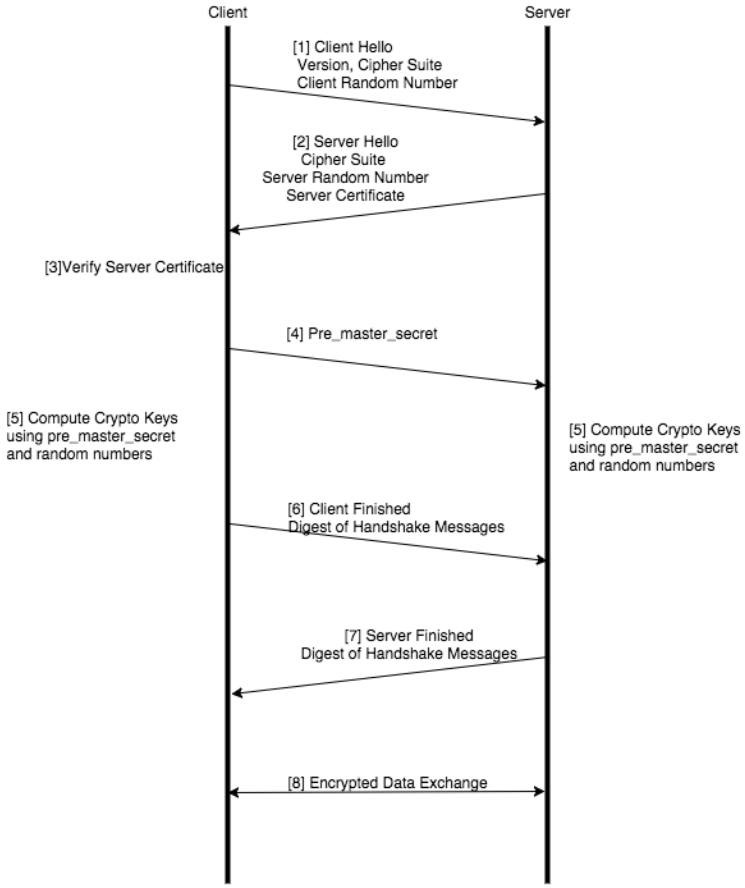

Figure 34 *Basic RSA TLS/SSL Handshake*

TLS/SSL Handshake Steps

The TLS/SSL handshake steps in Figure 34Basic RSA TLS/SSL Handshake are as follows:

1. The SSL or TLS client sends a Client Hello message that contains information such as the SSL or TLS version and the list of cipher suites supported by the client. The message also contains a client random number.
2. The SSL or TLS server responds with a Server Hello message that contains the SSL or TLS version it supports and the cipher suite chosen by the server from the list provided by the client, and a server random number. The server also sends its digital certificate.
3. The SSL or TLS client verifies the server's digital certificate.
4. Using the random numbers from the Hello messages, the SSL or TLS client computes the pre_master_secret that enables both the client and the server to compute the secret key to be used for encrypting subsequent message data. The pre_master_secret is encrypted with the server's public key and sent.

5. The SSL or TLS server verifies the client's certificate.
6. The SSL or TLS client sends the server a Client Finished message containing a digest (MAC) of the messages in the handshake, which is encrypted with the secret key, indicating that the client part of the handshake is complete.
7. The SSL or TLS server sends the client a Server Finished message containing a digest (MAC) of the messages in the handshake, which is encrypted with the secret key, indicating that the server part of the handshake is complete.
8. For the duration of the SSL or TLS session, the server and client can now exchange messages that are symmetrically encrypted with the shared secret key.

Figure 35Man-in-the-Middle TLS/SSL Handshake shows the TLS/SSL handshake when a Man-in-the-Middle sits between the client and server.

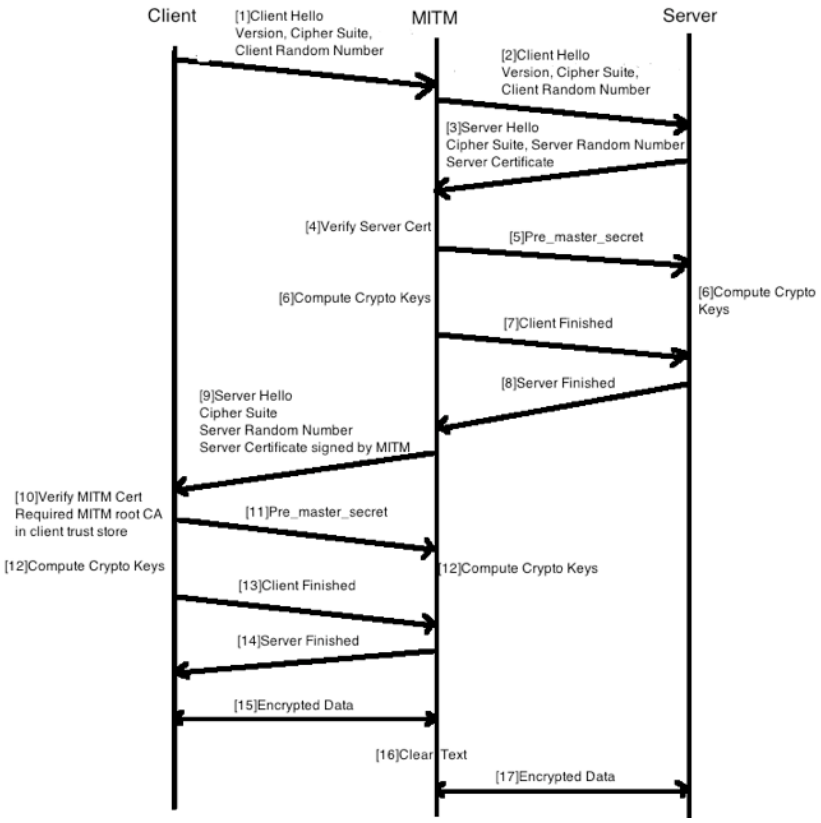

Figure 35 *Man-in-the-Middle TLS/SSL Handshake*

TLS/SSL Session, Inbound Deployment

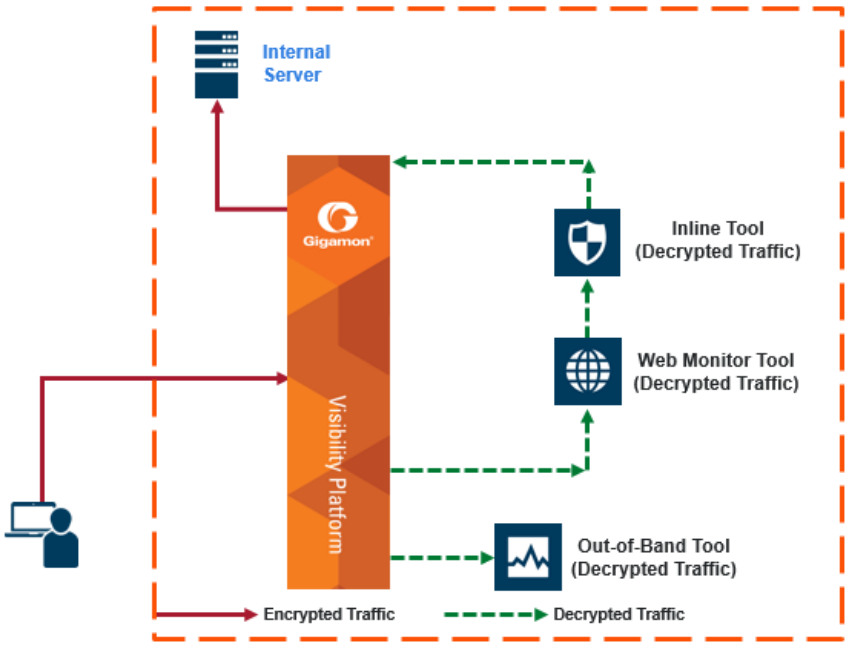

Figure 36 *Inbound Deployment of Inline TLS/SSL Decryption*

Figure 36Inbound Deployment of Inline TLS/SSL Decryption shows an inline TLS/SSL decryption inbound deployment. The client is on the Internet. The server and the GigaVUE node are located within the same enterprise network, with the GigaVUE node deployed on the server side. The GigaVUE node needs access to the private keys of the server.

The TLS/SSL session is created as follows:

1. Client traffic, such as from the Internet, arrives on an inline-network port on the GigaVUE node and establishes a TCP connection.
2. The GigaVUE node initiates a TCP connection to the server.
3. Gigamon establishes the TLS/SSL handshake with the server parameters from the TLS/SSL client's Hello request.
4. Gigamon establishes the TLS/SSL handshake with the client using the server's certificate and key, and using the appropriate parameters.

TLS/SSL Session, Outbound Deployment

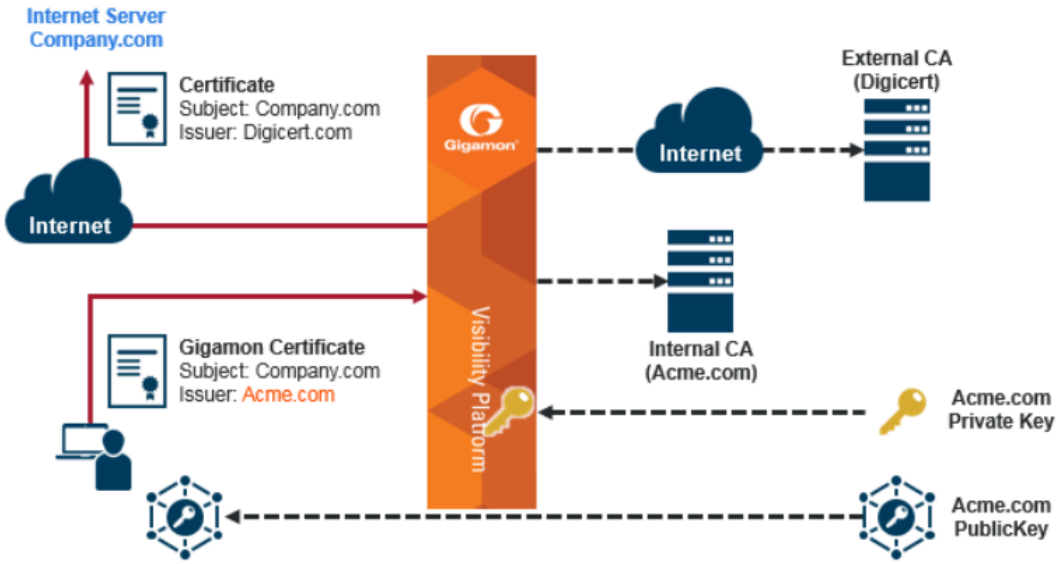

Figure 37 *Outbound Deployment of Inline TLS/SSL Decryption*

Figure 37Outbound Deployment of Inline TLS/SSL Decryption shows an inline TLS/SSL decryption outbound deployment. The client is in your network, and it connects to a server outside your network on the Internet. Traffic is destined to servers where you do not have access to the private keys.

The TLS/SSL session is created as follows:

1. Client traffic, such as from an employee on the enterprise's Intranet is destined to a server on the Internet. The traffic arrives on a network port on the GigaVUE node.
2. Gigamon initiates a new connection to the server as the client.
3. The server responds with its server certificate and presents it to Gigamon.
4. Gigamon spoofs the server certificate and presents it to the client, but now the certificate is signed by Gigamon.
5. The end client verifies that the certificate is valid as it belongs to the server and has been signed by Gigamon, which has been listed as one of the valid CAs.
6. Gigamon now maintains two TCP connections, one with the client and one with the server.

TCP Transition States during TLS/SSL Session

The TCP transition starts with INIT state during a session.

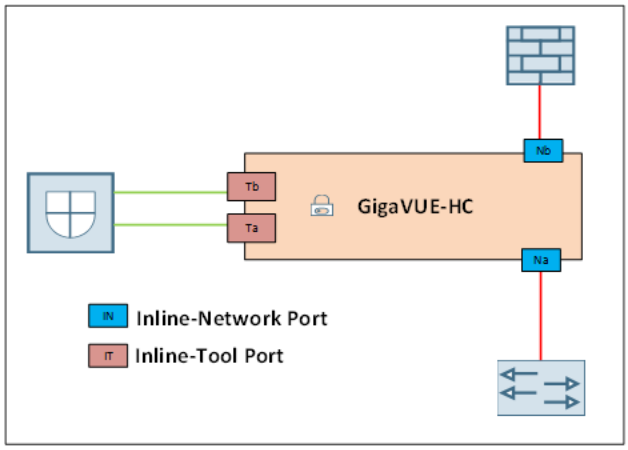

The following are the different states and its what it indicates:

- **(Na:SYN RCV:Nb:SYN_SENT:INIT:INIT)** - This status indicates that the client side has received the SYN and due to that the TCP session handshake has started from the server side of GigaVUE device. If the TCP handshake on the server is not successful, that is if the server is unreachable/busy, this connection is reset.
- **(Nb:SYN_RCV: EST:INIT:INIT)**- This status indicates that TCP handshake has been established at the server side of GigaVUE device.
- **(Na:EST:Nb:EST:Na:INIT:Nb:INIT)**- This status indicates that the TCP handshake has been established between the Client and Server side. The session will either be decrypted or not decrypted based on the policy.
- **(Na:EST:Nb:EST:Ta:SYN_SENT:Tb:INIT)** - This status indicate that either decryption or no decryption decision has been made and now the tool side TCP handshake would be initiated.
- **(EST:EST:Ta:SYN_SENT:Tb:SYN_RCV)** - This status indicates that either decryption or no decryption decision has been made, so the tool side TCP session is initiated, with SYN received on the tool server side.
- **(EST:EST:EST:EST)** - This indicates that a successful TCP handshake has been established.

TLS/SSL Session Resumption

TLS/SSL sessions can be resumed to improve performance. TLS/SSL session resumption speeds up the TLS/SSL handshake.

Once a session has been established, the keys are saved so a session can be resumed efficiently later. The resumed TLS/SSL handshake has fewer steps.

Session identifier-based resumption is supported. The GigaVUE node maintains the session identifier data in the cache. Session ticket-based resumption is not supported.

By default, resumption is enabled.

TLS/SSL Session Search

Starting in software version 5.2, you can search an existing session based on a hostname. The input is matched against the Server Name Indication (SNI) or the certificate subject name of the current sessions.

StartTLS and HTTP CONNECT

The Inline TLS/SSL Decryption solution looks for the CLIENT HELLO packet and, when found, it switches to TLS/SSL mode. The StartTLS command is initiated from the client or the server when switching from non-SSL to SSL mode.

Using the StartTLS mechanism, protocols such as SMTP, IMAP, and POP3 can be decrypted. These protocols start in plaintext mode and then upgrade to TLS/SSL mode on the existing port instead of using another port.

HTTP CONNECT provides a mechanism for explicit proxies where an HTTP session is established between a client and a server, and then upgraded to TLS. HTTP CONNECT is automatically detected.

Both StartTLS and HTTP CONNECT are upgrade related protocols that adds security to an existing insecure protocol and included in iSSL StartTLS command.

When enabling StartTLS, the specific ports to monitor StartTLS traffic must be specified. Up to 20 ports can be monitored.

Inline TLS/SSL Decryption Behavior with StartTLS

For connections that use StartTLS to upgrade from non-TLS mode to secure mode, the inline TLS/SSL decryption solution decrypts correctly if the decision to decrypt or not is made in the certificate phase.

If the CLIENT HELLO packet does not have SNI information, the inline TLS/SSL decryption solution will apply policy rules in certificate phase of the policy evaluation.

For explicit proxy connections policy, rules are applied in certificate phase of policy evaluation. For information on the certificate phase of policy evaluation, refer to Policy Evaluation.

NOTE: You need to enable StartTLS for decryption sessions established through explicit proxy.

TLS/SSL Keys and Certificates

The TLS/SSL 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.

An TLS/SSL certificate is a digital document containing a public key, host information, and a digital signature from the certificate issuer, known as a Certificate Authorities (CAs). The certificate allows trust to be established between two communicating endpoints.

The inline TLS/SSL decryption solution has a trust store, which is a collection of certificates of CAs. Gigamon only trusts server certificates that have a trust anchor in the configured trust store; in other words, the certificate chain must be built with one of the root CAs in the trust store. Gigamon ships with a default trust store, which you can replace if needed.

The inline TLS/SSL decryption solution acts as a Break-and-Inspect. In the outbound deployment case, the MitM generates server certificates on-the-fly signed by the installed Signing CA. In the inbound deployment case, server certificate generation is not needed but the server's private key and certificate chain need to be made available to the MitM.

The inline TLS/SSL decryption solution also has a key store, which is a collection of TLS/SSL private keys (for inbound deployments) and TLS/SSL certificates and corresponding private keys that are used to digitally sign the emulated server certificates (for outbound deployments).

Figure 38Sample Certificate shows a sample certificate and its relevant parts.

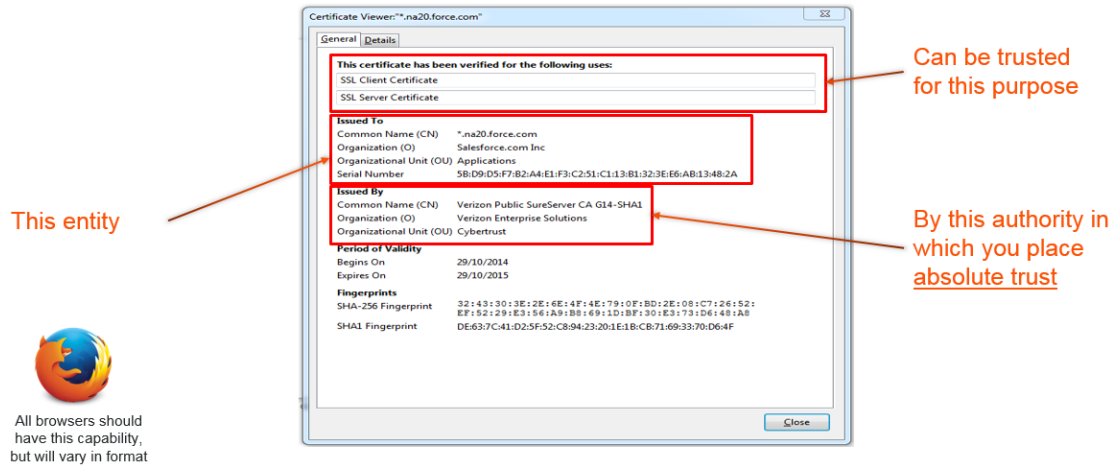

Figure 38 *Sample Certificate*

Key Store

The key store contains keys and certificate-key pairs. The key store can contain a maximum of 1000 key pairs.

A particular key in the key store can be selected only for decryption (inbound deployment) or for re-signing and re-encryption (outbound deployment).

Starting in software version 5.2, encrypted or password protected PEM (Privacy Enhanced Mail) is supported for fetching or downloading a private key.

Starting in software version 5.6.00, ECDSA keys are supported for both inbound and outbound deployments.

Set Up Key Store Certificate Management

A Key Store certificate can be setup to be auto- enabled, auto-deleted and auto-retained to an inline-SSL profile. The configuration can be done as follows:

1. On the left navigation pane, click  , and then select **Physical > Nodes**.
2. In the Physical Nodes page, select the node for which you want to configure Key Store settings.
3. Go to **GigaSMART**>**Inline SSL** >**Key Store** .
4. From the **Actions** drop-down, select **Settings**. Configure the following settings:
    - **Auto Enable New Certificates** – Associates new certificates to an inline-TLS/SSL profile, if the new certificate shares a common name with an existing certificate(s) of the inline-TLS/SSL profile. The setting will be triggered as and when a new certificate is uploaded or pushed by a third-party tool.
    - **Auto Delete Expired Certificates** – Deletes 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 inline-TLS/SSL profile. 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.

Generate and Add a Certificate to Key Store

To generate and add an inline-TLS/SSL signing certificate to key store for outbound deployment, perform the following steps:

1. Create an internal root CA (for example, using Microsoft AD, or OpenSSL, or any other CA implementation).
2. Push this CA root certificate to all devices.
3. Issue a sub-CA certificate for Gigamon.
4. Upload the sub-CA certificate to Gigamon with the private key of only the sub-CA.

Display Key Store Certificates

The key store certificates added would be displayed in Key Store page.

To access the Key Store page:

1. On the left navigation pane, click  , and then select **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 > Inline 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 an inline-SSL profile and the profile is part of the inline-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 .

Component	Description
	Blue - The key certificate is not actively participating in any traffic flow. Red - The key certificate has been 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.

Trust Store

The trust store contains a trusted certificate authority (CA) for server validation. A default trust store from Mozilla is included with this GigaVUE-OS software version. The trust store is updated periodically. You can also fetch a trust store file containing certificates in PEM format if you want to replace the default trust store.

Starting in software version 5.2, CA certificates can be appended to the trust store or specific certificates can be deleted from the trust store. Also, the trust store can be queried for a specific certificate by its fingerprint. The query includes the hex representation of the first four octets of the certificate's SHA1 fingerprint.

NOTE: Whenever you make changes to the trust store, you should re-configure the inline-SSL profile again in order to apply the changes.

Certificate Validation

Gigamon needs to validate the server certificate so that an incoming untrusted certificate is not made legal by Gigamon re-signing the certificate.

The certificate validation process includes several steps as follows:

1. **Certificate expiration date and validity period**: The GigaVUE node compares the current date to the validity period listed in the certificate. If the expiration date has not passed and the current date is within the period, the certificate is good.
2. **Certificate issued by trusted CA**: The GigaVUE node maintains a list of trusted root CAs. This list determines the certificates that the client will accept. The trust store acts as a trust anchor during certificate validation. The GigaVUE node validates that each incoming certificate chain is trusted by one of the certificates in the trust store.
3. **Server name**: The GigaVUE node validates that the server certificate is valid for the hostname mentioned in the SNI. This validation is not performed if the client does not send SNI.
4. **Certificate revocation check**: The GigaVUE node validates the server certificate status using OCSP and CRL lists downloaded from the concerned CAs. Internet connectivity is required for this functionality. The certificate revocation check determines the revocation status of the server certificate.

Certificates that pass the validation are accepted. The primary MitM CA signs the forged certificate. Certificates that fail validation can be accepted if security exceptions are configured and the secondary MitM CA signs the corresponding forged certificate. For self-signed certificates, the forged certificate will also be self-signed.

Client applications will typically add the primary MitM CA to their trust store and not to the secondary. This will act as a mechanism to bubble up certificate validation errors on the client applications and provide the end users an opportunity to reject the connection.

If there are certificate validation errors, the TLS/SSL connection is dropped unless explicitly permitted by the security exceptions in the policy profile. The certificate validation errors are grouped into the following four categories:

- **Expired**: The validity period of the certificate is in the past.
- **Self-signed**: The certificate is self-signed, meaning that the subject and the issuer are the same. The validity period is current and the certificate signature is valid.
- **Unknown CA**: The CA is not valid. The issuer certificate cannot be obtained from the certificate chain or is not in the trust store. You can download the root certificates for these sites and import them into the trust store if the certificate is valid and trusted.
- **Invalid**: The certificate is not valid for the given SNI (for the Client Hello containing the SNI). There might have been a failure to decrypt or decode the certificate signature or the fields, not and before/not, were not read.

For security exceptions for expired, unknown CA, and invalid certificates, the resulting certificate is signed (re-signed) by the secondary MitM CA, so the user can accept or reject the connection.

Client Authentication

A server can challenge the client by requesting for a client certificate after the server responds with its Server Hello message. The client then respond with its certificate and the server validates the client certificate.

Gigamon does not support client authentication for outbound and inbound connections. If client authentication is detected during server handshake, then the connection is bypassed.

Re-Signed Certificates

As a MitM, Gigamon re-signs certificates. The following fields are copied from the original certificate:

- subject name
- certificate validity
- subject alternative name

The following fields are set in the removed list:

- authority Information Access
- certificate Policies
- CRL distribution points
- SCT List

The following fields are set in the re-signed certificate:

- certificate type—v3
- issuer
- version
- public key
- serial number—randomly generated
- signature algorithm / hash
- thumbprint
- v3 extensions:
  - basicConstraints CA—True

    NOTE: As this is a CA certificate, basicConstraints is set to True. For leaf certificates, basicConstraints is set to False.

  - keyUsage—digitalSignature and keyEncipherment
  - extendedkeyUsage—serverAuth

- subjectKeyIdentifier—hash
- authorityKeyIdentifier—keyid,issuer:always

Checking Certificate Revocation Status

All server certificates for decrypted outbound connections are issued by the GigaVUE node. The issuing CA is imported into the client's browsers as a trusted CA. Thus, the clients will trust all certificates signed by the GigaVUE node. This interferes with the ability of the clients to check for the revocation status of the certificates. Thus the burden is upon the node to perform the revocation checks on the original server certificates before regenerating the certificates to the clients.

If revocation check is enabled with soft fail, decryption will continue even if the revocation status is not already known, whereas with hard fail, traffic will not be decrypted unless the revocation status is determined for certain.

If revocation check is enabled, once the GigaVUE node determines that the server certificate is revoked, further TLS/SSL connections to the server are dropped.

By default, revocation check is disabled.

There are two methods to check the revocation status of the certificates as follows:

- using a Certificate Revocation List (CRL) from the issuing Certificate Authorities (CAs). Refer to Certificate Revocation List (CRL).
- using the Online Certificate Status Protocol (OCSP). Refer to Online Certificate Status Protocol (OCSP).

Certificate Revocation List (CRL)

A Certificate Revocation List (CRL) is an online database of certificates that have been revoked.

Each issuing CA in the PKI infrastructure maintains a list of revoked certificates that they had issued earlier. The list contains the serial number of the revoked certificates and the reasons for the revocation. Any revoked certificate should not be trusted even if the signatures are valid. CAs publish the revocation list periodically.

Each server certificate will contain the CRL location in the "CRL distribution points" X.509 extension.

Online Certificate Status Protocol (OCSP)

The Online Certificate Status Protocol (OCSP) is an Internet protocol used for obtaining the revocation status of an X.509 certificate.

Certificate status can be verified near real-time by querying the OCSP link in the certificate. The link will be present in the “Authority Information Access” X.509 extension.

CRL and OCSP

If both CRL and OCSP are enabled, OCSP is performed first, followed by CRL.

Both CRL and OCSP require Internet connectivity. Refer to Set up the Stack Port Interface for more information.

Policy Profile

The policy profile consists of multiple rules, with each rule having a decrypt or no-decrypt action for the match condition. For example, there might be a policy to decrypt all but financial-related traffic.

In addition to the rules, the profile also consists of various configuration options that affect the decryption decision as follows:

- The default action to take if none of the rules match.
- The URL cache miss action to take if the URL category-based rules are configured, but GigaSMART does not have the category information.
- For decrypted traffic, options to override expired, invalid, self-signed, and unknown CA certificates and to enable or disable the certificate revocation check.
- Whether or not to send decrypted/non-decrypted traffic through the tools.

Each policy rule consists of a match condition and the decrypt or no-decrypt action for the match. The following rule types are available:

- URL category
- hostname/domain name
- server certificate issuer
- source and destination IP address
- source and destination port numbers
- VLAN identifier

NOTE: You can configure up to 2048 policy rules under a policy profile.

Policy Evaluation

Policies are evaluated by GigaSMART at various phases as shown in Figure 39Policy Validation Flow.

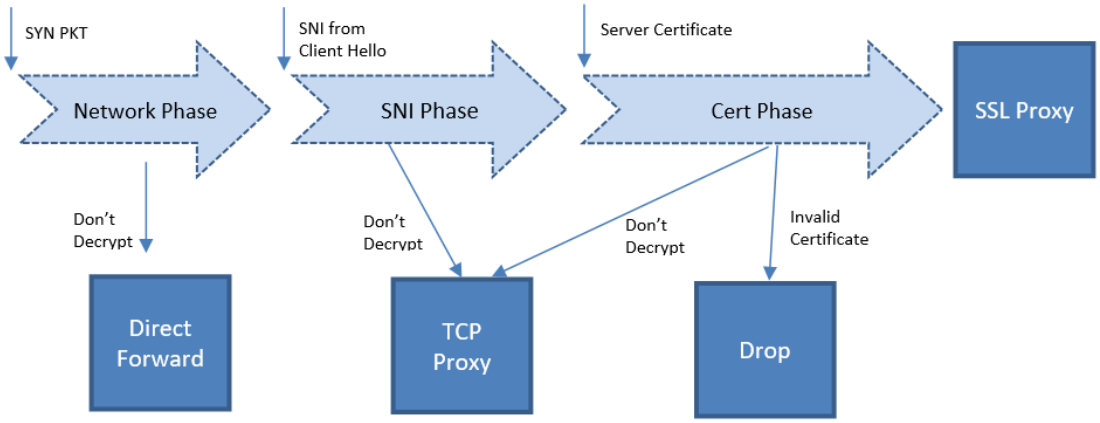

Figure 39 *Policy Validation Flow*

Network Phase

The Network Phase of the policy is done on the SYN packet. In this phase:

- The policy engine inputs are source and destination IP addresses, source and destination ports, and VLAN identifiers.
- The no-decrypt rules are evaluated before the decrypt rules, which are ordered by the following match conditions: source IP, destination IP, source port, destination port and VLAN.
- If traffic is not to be decrypted, packets are processed as non-proxy traffic. The bypass VLAN is used. Refer to Direct Forward in Figure 39Policy Validation Flow.

  NOTE: The no-decrypt verdict is always final for the lifetime of the given TCP connection.

- Traffic that is to be decrypted continues to the next phase. The decrypt verdict from the Network Phase can be overridden by more specific rules in the next phase.

Following conditions are considered for traffic that is to be decrypted:

- If the Client Hello packet contains SNI, policy is evaluated in the SNI phase and then in the Certificate phase.
- If the Client Hello packet does not contain SNI, policy is evaluated only in the Certificate phase.

- For HTTPS proxy connections, policy rules are evaluated based on the hostname from the HTTPS proxy CONNECT request in addition to the SNI.

SNI Phase

The SNI Phase of the policy is done on the Server Name Indication (SNI) from the Client Hello. In this phase:

- The policy engine input is the hostname from SNI.
- Rules are evaluated in the following order:
  - no-decrypt list/no-decrypt domain
  - decrypt list/decrypt domain
  - no-decrypt category
  - decrypt category

  NOTE: On the no-decrypt list/decrypt list wildcard entries, you can specify explicit wildcard domain rules like ***.domain.com** instead of **domain.com**. For example, ***.gigamon.com** evaluates all the sub-domains, and **gigamon.com** only evaluates single domain (i.e. gigamon.com).
- If traffic is not to be decrypted, packets are processed as TCP proxy. Refer to TCP Proxy in Figure 39Policy Validation Flow.
- For traffic that is to be decrypted, the plaintext version is sent through the tools based on the decrypt tool-bypass configuration.

If URL category-based rules are configured and a URL cache miss occurs, the policy verdict is based on the settings of the URL cache miss action. A cache miss action of decrypt or no-decrypt will cause the traffic to be decrypted or not decrypted immediately. The defer action will cause delays.

For compliance reasons, the cache miss action of no-decrypt is recommended.

Certificate Validation

In the case of a decrypt decision from the SNI phase or if the Client Hello does not contain SNI, GigaSMART will verify the server certificate using the configured trust store. Additional checks will be performed for certificate expiry, hostname mismatch, and self-signed certificate. If configured, revocation check will also be done on otherwise valid server certificates.

For valid server certificates, GigaSMART will issue the corresponding server certificate using the primary MitM CA. If the validation fails, the connection will be dropped unless a security exception is configured. The secondary MitM CA, if configured, will be used to issue server certificates in that case.

Starting in software version 5.2, the primary MitM CA is not mandatory for an inbound deployment.

Cert Phase

The certificate phase of policy evaluation is done for all connections after the certificate validation is completed. In this phase:

- The policy engine inputs are the certificate issuer and the certificate subject name.
- The Common Name (CN) attribute is extracted from the server certificate subject name. Policy evaluation in this phase is performed on the value of the CN attribute. This is similar to the SNI phase. If no matching rules are found, the certificate issuer-based rules are evaluated. For issuer-based rules, the CN attribute and the Domain Name (DN) attribute of the issuer are considered.
- If traffic is not to be decrypted, packets are processed as TCP proxy. Refer to TCP Proxy in Figure 39Policy Validation Flow. The server SSL session is reset and the Client Hello is resent to the server.
- GigaSMART supports certificate based policy evaluation for HTTPS proxied connections. It applies policy rules based on the hostname from HTTPS proxy CONNECT request for HTTPS proxy connections. Also, if required it applies no-decrypt in the certificate phase for HTTPS proxy connections.
- Certificate phase policies are also evaluated for HTTPS proxy connections.

Policy Profile Options

This section describes a few of the options for the policy profile. Refer to the following sections:

- Inline TLS/SSL Decryption Port Map
- Enable or Disable Tool Bypass
- High Availability Active Standby
- Inline Network Group Multiple Entry
- Tool Early Engage
- One-Arm Mode
- Tool Early Inspect

Inline TLS/SSL Decryption Port Map

The TCP destination port for decrypted traffic sent to inline tools can be configured as part of the profile. If not configured, incoming port is used directly as output by default.

Following are the two priorities that GigaSMART uses to decide on the TCP port number used for decrypted traffic:

- Priority 1—This is a port map, which is user configurable. You can specify both the In Port and the Out Port. The In Port is the TCP destination port from a client. The Out Port is the TCP port used to send traffic to inline tools.
- Priority 2—This is a default Out Port. This TCP port will be used if the incoming port does not match those specified in Priority 1.

NOTE: You cannot configure Decryption Port Mapping and Tool Early Inspect features together.

Enable or Disable Tool Bypass

Tool bypass can be enabled or disabled for the following types of traffic:

- TLS/SSL decrypted traffic
- non-decrypted SSL traffic (non-TLS/SSL TCP)
- non-TLS/SSL traffic (non-TCP)

By default, tool bypass is disabled on these traffic types, meaning that all decrypted TLS/SSL, non-decrypted TLS/SSL, and non-TLS/SSL traffic is sent to the tools. When tool bypass is enabled on a specified traffic type, that traffic is not sent to the tools.

High Availability Active Standby

Starting in software version 5.2, inline network high availability active standby is supported. When enabled, link switchover by an upstream device in active/standby scenario is detected.

For example, when there is an inline TLS/SSL network group topology with two network port pairs (Na1, Nb1 and Na2, Nb2), the incoming traffic from one network (for example, Na1) may change to another network (for example, Na2) due to upstream devices, such as firewalls performing high availability active standby failover. If an upstream device fails over, GigaSMART will forward traffic to the correct inline network.

The default is disabled.

NOTE: Do not enable this option if the inline TLS/SSL network group links are in an active/active scenario.

Inline Network Group Multiple Entry

An inline network group topology can have multiple network port pairs (for example, Na1, Nb1 and Na2, Nb2). With multiple network port pairs, traffic from a network interface might traverse GigaSMART multiple times. Intercepted traffic from GigaSMART might reenter GigaSMART through a different network interface within the same network group as shown in Figure 40Inline TLS/SSL Inline Network Group Configuration.

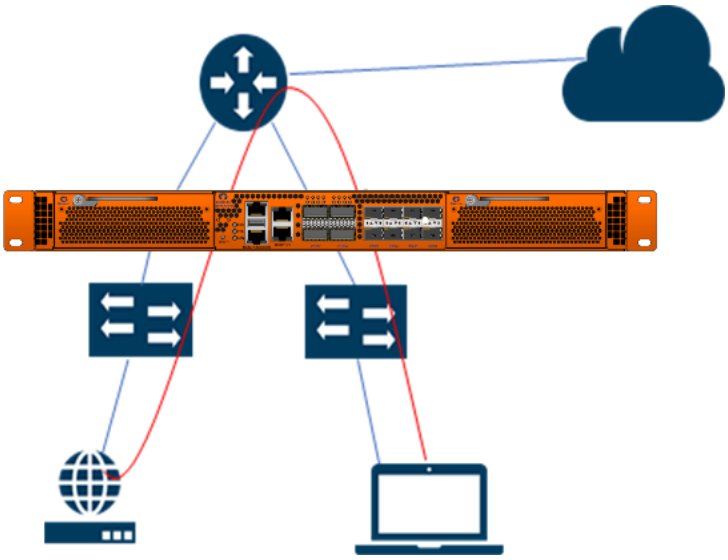

Figure 40 *Inline TLS/SSL Inline Network Group Configuration*

When the inline TLS/SSL GigaSMART sits between internal devices and the upstream router, traffic from the devices to the Internet will be intercepted by GigaSMART. When internal devices belonging to different network port pairs within the same inline network group communicate with each other, traffic initiated from a device will be intercepted by GigaSMART and sent to the upstream router. This traffic will be routed back to GigaSMART from a different network port pair to reach the destination device.

Starting in software version 5.3, the same traffic sent from GigaSMART can reenter GigaSMART.

GigaSMART remembers the inline incoming inline network interface (for example, Na1) for each connection. When traffic from the same connection reaches GigaSMART with a different inline network interface within the same network group (for example, Na2), GigaSMART will forward the traffic to the corresponding opposite network interface (for

example, Nb2), without further processing. This allows traffic from the same connection to reenter GigaSMART. GigaSMART will detect it and start forwarding traffic to the new network port pair.

However, the same traffic sent by GigaSMART reentering through the same network port pair (for example, Nb2, Na2) is not supported.

Other than the use case described above, any connection with traffic passing through GigaSMART involving more than the original network pair is not supported. If the first packet of a connection comes in through Na1, all traffic has to enter GigaSMART through the network port pair, Na1, Nb1.

You can enable or disable the inline network group multiple entry for the profile. The default is disabled.

Tool Early Engage

In a layer 3 topology, the inline tools may need to change the MAC address or VLAN IDs when the client traffic is sent back to the server. The Tool Early Engage option supports the inline tools to make this change. When a connection request is received from the client, GigaSMART establishes the connection with the inline tool first, before connecting with the server. This helps the inline tools to modify the MAC address or VLAN IDs when sending the traffic back to the server.

Note

- The Tool Early Engage setting can be enabled for a policy profile as a standalone feature without the One-Arm mode enabled.
- You cannot configure Tool Early Engage and Tool Early Inspect features together.

One-Arm Mode

With the One-Arm mode enabled, you can have both the client and server traffic travel through the same physical link or logical aggregate port channel.

Configuring One-Arm Mode

One-Arm Mode can be configured from GigaVUE-FM and GigaVUE-OS CLI

In GigaVUE-FM to enable One-Arm Mode do the following:

1. Click on and select your node.
2. While configuring maps, for the type Inline Second Level Map and click on the traffic path checkbox provided.
3. Once you enable One-Arm, destination ports will become non-configurable.

In GigaVUE-OS CLI you can configure one-arm mode through the following commands:

- **(config) # apps inline-ssl profile alias <profile_alias> one-arm enable**

Configure Inner MAP

- **map alias <inner-map alias>**
  use gsop <gsop alias>
  to <one-arm alias>
  from <vport alias>
  exit

NOTE: Enabling one-arm mode in Inline second level map is applicable only if one-arm mode is enabled in inline-ssl-profile.

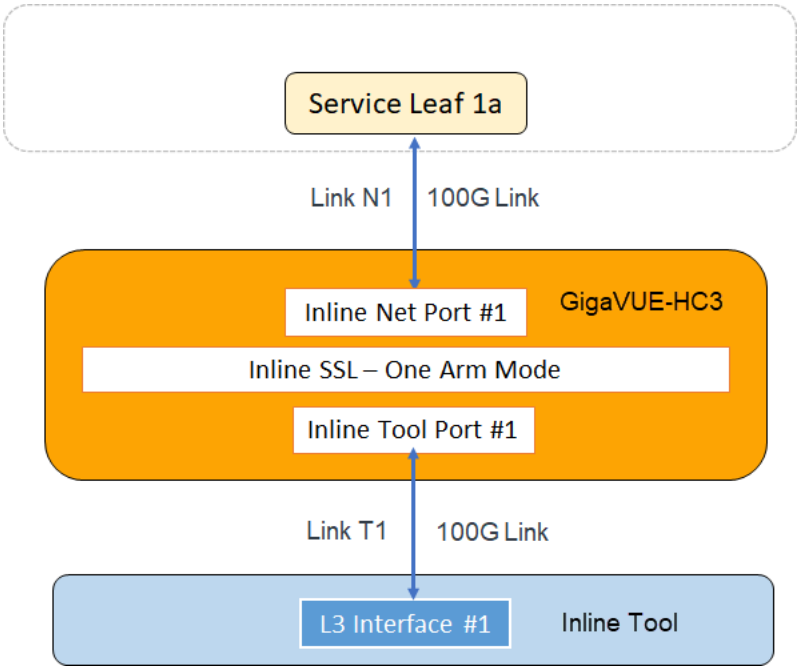

For each connection between the client and server, there are two TCP sessions established between GigaSMART and the inline network and two TCP sessions established between GigaSMART and the inline tool. In the above diagram, you can see that with the One-Arm

mode enabled, both TCP sessions from the inline network side arrive at GigaSMART on the same link N1. The TCP sessions from the inline tool side arrive at GigaSMART on the same link T1.

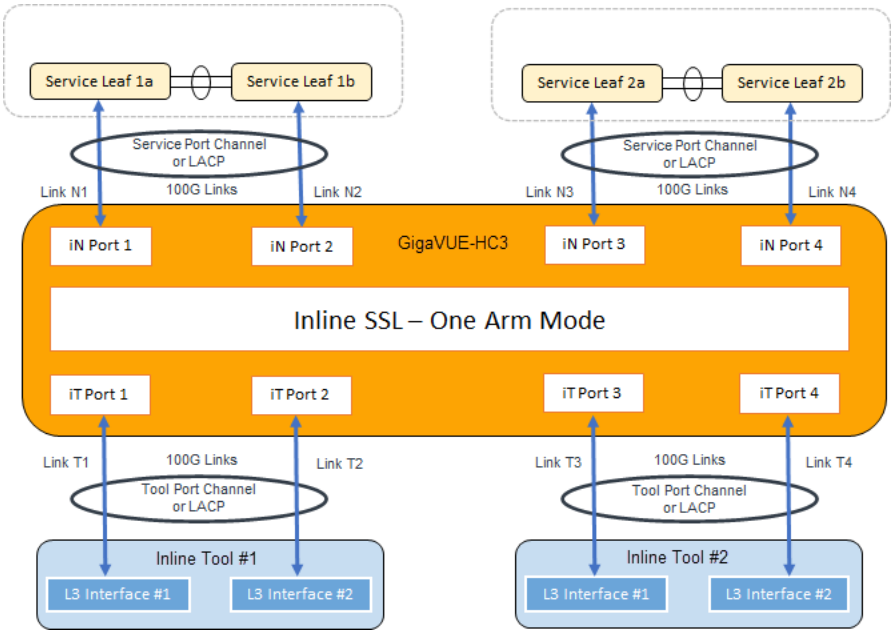

In the above figure, the inline network link and inline tool link works as a pair – (N1, T1), (N2, T2), (N3, T3), and (N4, T4). The GigaSMART sends traffic to the corresponding tool link of the received network link. Similarly, GigaSMART sends the traffic back to the server on the corresponding network link of the received inline tool link. For example, when a connection comes to GigaSMART from the inline network N1, after decryption, GigaSMART sends clear text traffic to inline tool on T1. The return traffic of the same connection arrives at GigaSMART on the inline tool link T2. GigaSMART then re-encrypts the traffic and sends it to the server on the inline network link N2.

Failover Support

When the inline network link and inline tool link work as pairs and when one of the link goes down, the corresponding link of the pair will be forced down. To achieve this behavior, you can either configure Link Aggregation Control Protocol (LACP) or enable Link Failure Propogation (LFP) between the pair of inline network link and inline tool link.

Important Rules and Notes

Keep in mind the following rules and notes before you enable the One-Arm mode for your policy profile:

- Inline network must be connected to Na side and inline tool must be connected to Nb side of the inline network.

- The **Tool Early Engage** option must be enabled.
- One-Arm mode is not supported on flexible inline decryption solution.
- To route the packet through the firewall, you must configure the client and server gateway to the router's IP address.
- Second level oob is not supported. Instead, oob can be achieved by map-passall from the tool port.
- When the Tool Early Engage and One-Arm mode is enabled, the MAC address or VLAN IDs can be changed, whereas, the IP address, ports, and protocol cannot be changed by the inline tool.
- Inline network group multiple entry and High Availability active stand-by are not supported when you enable the One-Arm mode for your policy profile.
- Ensure that LACP is enabled on both network and tool sides.
- Ensure that LFP is enabled on the inline network configuration.
- Bypass inline-tool options in the inline TLS/SSL profile (decrypt, no-decrypt, non-tls/ssl) will only work if you have another router configured on the network side to route the packets.
- The name "one-arm" is a registered key word hence do not name your Inline network ,Inline tool or Map as "one-arm" . If you have existing alias with "one-arm" as the name then modify it before upgrading to GigaVUE-OS 5.12.xx.
- Resilient Inline Arrangement will also not be supported if one-arm mode is enabled.
- If 'one-arm' is configured as a tool in inline second level map, VPort status will be "up" not "up (Normal)".
- One-Arm mode cannot co-exist with the **Tool Early Inspect** feature.

Tool Early Inspect

In the existing Inline TLS/SSL solution, the GigaVUE node intercepts the TLS/SSL connection initiated between the client and the server. The decrypted data is sent to the inline tool only after completing the TLS/SSL handshake connection between the client and the server. The downside of this method is that the TLS/SSL connection to the server is established, even for connections that the inline tool will subsequently reject.

Starting in software version 6.3.00, if Tool Early Inspect feature is enabled, the TLS/SSL handshake initiated by the client is completed first using the configured server certificate and key, and the decrypted data will be sent to the inline tool for inspection before connecting to the server. This process helps to ensure that only valid connections are sent to the server.
The following image explains the sequence of operations in the Tool Early Inspect feature.

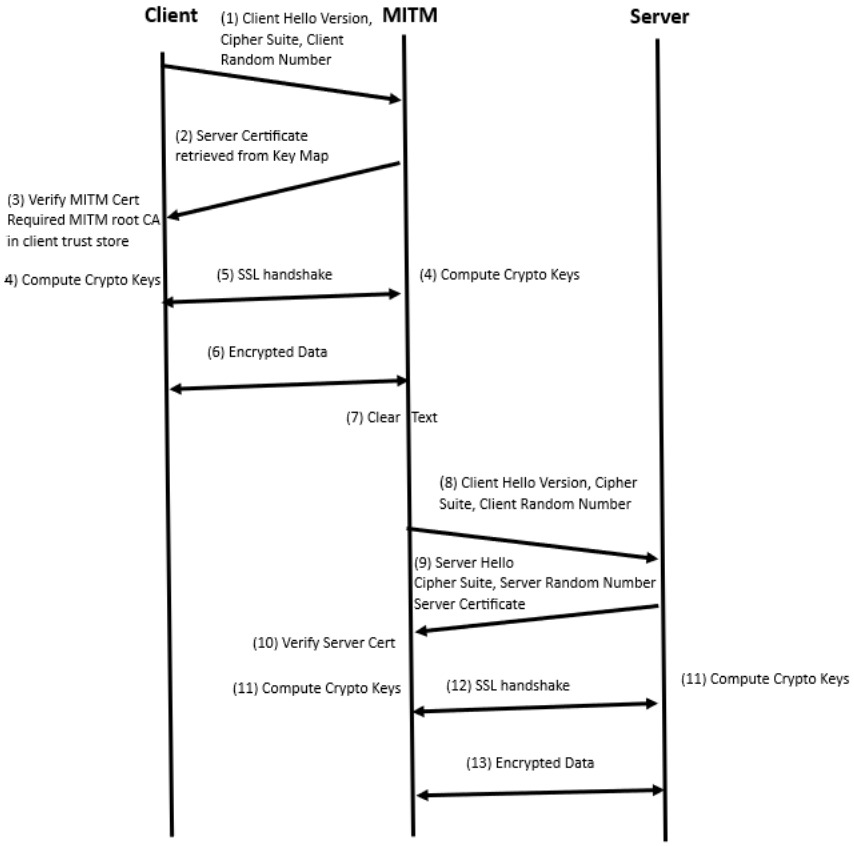

NOTE: Tool Early Inspect feature will only be supported in the inbound deployment modules; the outbound and hybrid deployment modules will not support it.

Caches

There are four in-memory caches as follows. They are not configurable.

- re-signed certificate cache
- URL category cache
- revocation certificate cache
- session resumption cache

Cache Persistence

Caches are maintained for Internet lookups such as URL categorization and certificate revocation checks using OCSP or CRL for faster subsequent lookups. The cache persistence feature allows the information to be saved on the GigaVUE node in the control card's persistent storage so that it can be retrieved in case of reboots. This allows the GigaSMART card to start with the information learned earlier. This feature is enabled by default and can be disabled if needed.

On the **Cache Persistence** page, you can:

- search for specific entries in the caches using the **Find Entries** option
- clear the caches using the **Clear Store** option from the **Actions** drop-down menu
- display a summary of the records.

GigaSMART Overload Bypass

Packet buffers, CPU, and concurrent connections are monitored for overloaded conditions. GigaSMART goes to bypass when resource usage exceeds thresholds. Existing connections will continue to be processed by GigaSMART, but any new connections will be bypassed. Refer to Overload Bypass Connections and Thresholds for information on connections and thresholds.

Table 9: Overload Bypass Connections and Thresholds

Criteria	GigaVUE-HC1	GigaVUE-HC3 (per GigaSMART Engine)
Maximum connections per second	• 1500	• 5000
Maximum connections	100000	200000
Resource Packet Buffer	Overload threshold for packet buffer resources for GigaSMART operations. Default is 80% (configurable)	
Resource CPU	Rising threshold for GigaSMART CPU statistics. Default is 90% (configurable)	
Heap exhaust	80%	

To configure the Packet Buffer and CPU threshold values, navigate to **GigaSMART** > **GigaSMART Operations (GSOP)** > **Resource Buffer** and configure the following:

- Resource Packet Buffer Overload Threshold (%)
- Resource CPU Overload Threshold (%)

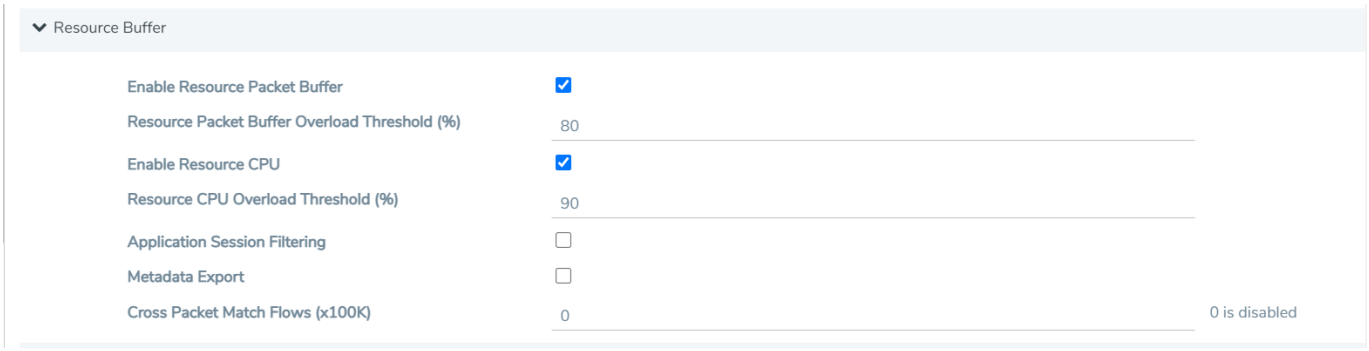

CPU Overload Threshold

Due to sudden bursts of traffic, the GigaSMART CPU can become too busy and drop packets. However, when a system or application reaches a threshold, SSL sessions can be bypassed. When a maximum CPU is reached, incoming connections will be bypassed.

When the CPU overload threshold is set to a configured value, (for example, 90%), the lower threshold is set to two-third of the CPU overload threshold configured (in this example 60%). A mean threshold is calculated, which will be the average of the CPU overload threshold and the lower threshold (in this example 75%).

The following actions will be taken:

- If the CPU hits the overload threshold, all new SSL connections will be bypassed.
- If the CPU reduces to the mean threshold, half of the new SSL connections will be bypassed.
- If the CPU reduces further to the lower threshold, all new SSL connections will be decrypted.

If you choose connectivity-over-security, the CPU overload threshold must be set to the lower threshold value.

Inline TLS/SSL Monitor Mode

Use the inline TLS/SSL monitor mode to assist in understanding your network topology. Monitor mode provides information about the traffic going to the GigaSMART card, which can help to learn about your deployment. When monitor mode is enabled, the monitor application collects information such as TCP ports used and VLAN information about the incoming traffic.

After inline TLS/SSL decryption is configured and monitor mode is enabled, the inline TLS/SSL application does not terminate the session. Instead, the monitor application collects information and forwards packets to the tool port or network port based on the configuration of the non-TLS/SSL TCP bypass action. For any Monitor mode, you can enable or disable seamlessly without any other configuration changes.

Monitor mode is disabled by default. To enable the monitor mode, refer to Configure the Inline TLS/SSL Monitor Mode.

For packets coming from the network port, the monitor application collects packet flow information.

From the information collected from monitor mode, you can analyze the following cases:

- duplicate TCP SYN—For a given session, the SYN messages with a different packet signature than 5tuple, for example, a different VLAN ID, indicates the packet is coming from multiple paths.
- asymmetric routing—For a given session, packets arriving from multiple network interfaces indicates a packet is coming from multiple paths.

Inline SSL Monitor mode only captures TCP information, not SSL information. However Inline SSL Persistent Monitor mode captures both TCP and SSL information.

NOTE: Monitor mode is supported for standalone nodes only, not for nodes in a cluster.

Configure the Inline TLS/SSL Monitor Mode

You can enable or disable the inline TLS/SSL monitor mode, or enable persistent inline TLS/SSL monitor mode using either CLI command or GigaVUE-FM.

To enable or disable the monitor mode using CLI, run the following CLI command:

(config) # apps inline-ssl profile alias sslprofile monitor enable
(config) # apps inline-ssl profile alias sslprofile monitor disable

To enable the persistent monitor mode using CLI, run the following CLI command:

(config) # apps inline-ssl profile alias sslprofile monitor inline

To enable the monitor mode using GigaVUE-FM:

1. From the device view, go to **GigaSMART > Inline SSL > SSL Profiles**.
2. From the Actions drop-down, click **Edit**.
3. From the **SSL Monitor Mode** drop-down list, select
   - **Disable** to disable SSL monitor mode, and enable SSL decryption/encryption.
   - **Enable** to enable SSL monitor mode, and disable SSL decryption/encryption.

- **Inline** to enable persistent monitor mode (both SSL monitor mode, and SSL decryption/encryption).

4. Click **Apply**.

Inline Tool Configurations

Inline tools connect to the GigaVUE node through inline bypass (BPS) modules, available on the GigaVUE-HC1, and GigaVUE-HC3.

The inline bypass arrangements supported by inline TLS/SSL decryption are as follows:

- single inline network. The inline network ports can be protected, unprotected, or a mix of protected and unprotected.
- inline network group, consisting of multiple inline networks
- single inline tool
- inline tool group, consisting of multiple inline tools over which traffic is distributed
- inline tool group with a spare inline tool, in which the failure of one tool in the inline tool group will trigger a failover to the spare
- inline series, in which traffic is guided through inline tools in a particular order

In summary, inline TLS/SSL decryption can be deployed in any combination of inline network and inline network group with any inline tool, inline tool group, or inline series.

Note: The TLS/SSL sessions will fail ,if a BPS card is shut down and brought back up with an Inline TLS/SSL configuration.

Inline tool ports can be configured in shared mode. When an inline tool is shared (true), the decrypted traffic will be VLAN tagged. The connected inline device is expected to receive VLAN tagged packets instead of untagged packets. There is an extra outer VLAN tag added to the packet, which the connected inline device needs to see. When an inline tool is not shared (false), the extra VLAN tag is not added. This allows untagged traffic to be sent to the tool ports. Use false for inline tools that are not able to handle more than one VLAN tag, such as Q-in-Q tagged packets. For tagless mode, if an inline tool is involved in an inline TLS/SSL map, the inline tool cannot be used in any other classic inline map.

Starting in software version 6.5, you can choose to deploy flexible iSSL solution with single VLAN tag in which the original tag received from the inline network will be removed and only the tool tag you create when configuring an inline map will be part of the packet. This will prevent the inclusion of additional VLAN tag in the decrypted packets that have been routed through inline tools. The inline tool that is shared with a single VLAN tag in a packet will be supported in the flexible inline TLS/SSL arrangements across multiple maps with different single VLAN tags. Shared mode (false) is not allowed when single VLAN tag is enabled in the iSSL solution.

Refer to Figure 41Simple Inline Tool Arrangement toFigure 43Serial Inline Tool Arrangement for inline tool arrangements. Encrypted traffic is shown in solid lines, decrypted traffic is shown in dotted lines.

Figure 41Simple Inline Tool Arrangement shows a simple inline tool arrangement with one inline tool connected to the GigaVUE node. Traffic is decrypted and sent to the inline tool.

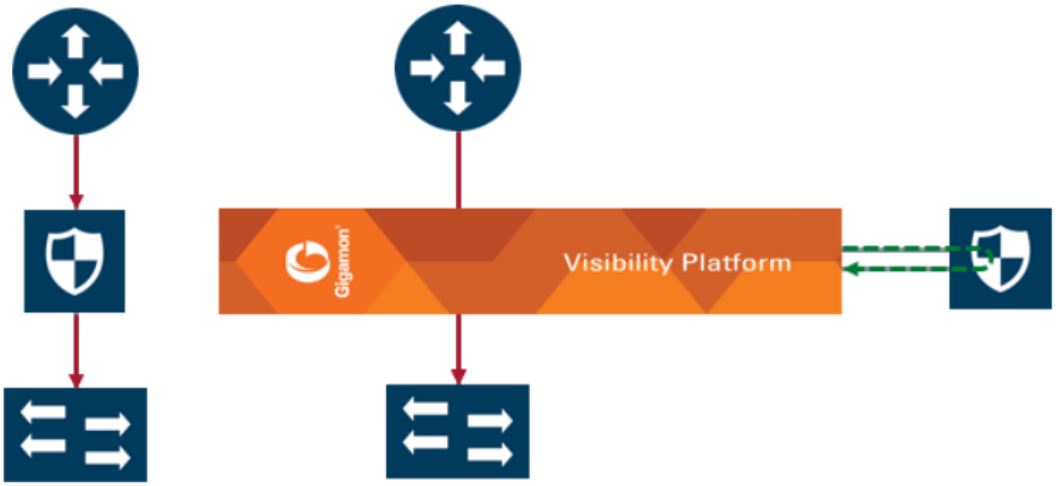

Figure 41 *Simple Inline Tool Arrangement*

Figure 42Multiple Inline Tool Arrangement shows a multiple inline tool arrangement with three inline tools connected to the GigaVUE node. Decrypted traffic is distributed across the tools.

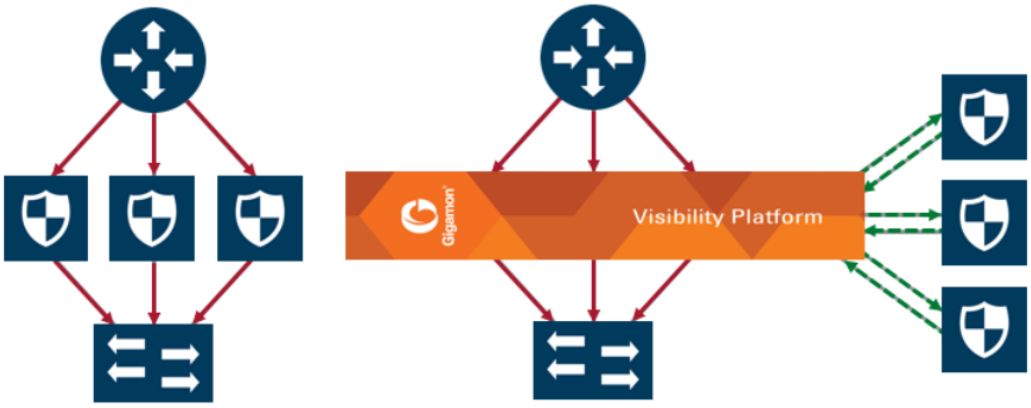

Figure 42 *Multiple Inline Tool Arrangement*

Figure 43Serial Inline Tool Arrangement shows a serial inline tool arrangement in which traffic is decrypted on the GigaVUE node, sent serially through the inline tool, and then re-encrypted on the GigaVUE node.

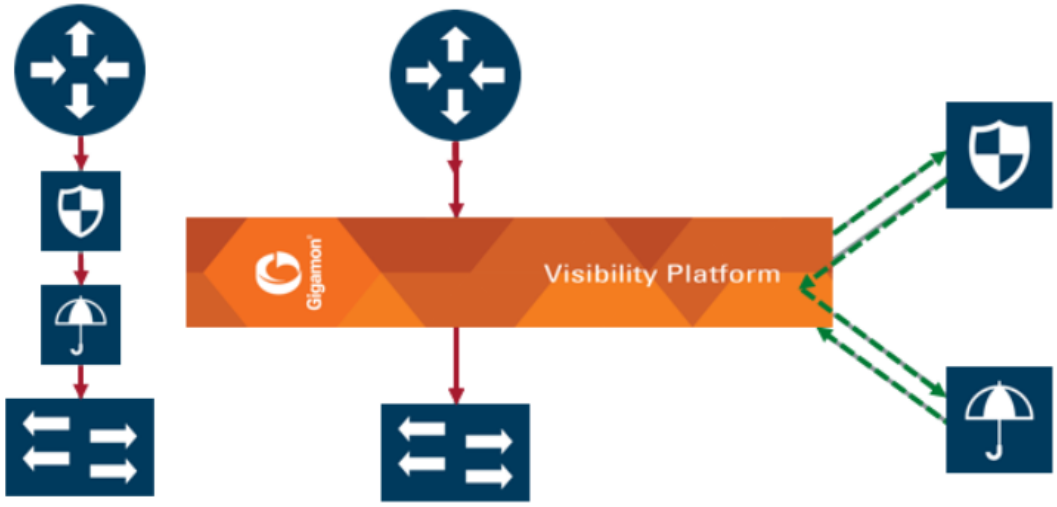

Figure 43 *Serial Inline Tool Arrangement*

Inline Bypass Restriction

The inline bypass arrangements supported by inline TLS/SSL decryption have the following restriction:

- inline network group does not support ingress VLAN tagging on the member links

Forwarding

Inline TLS/SSL decryption supports the following kinds of forwarding:

- inline forwarding—Packets can be forwarded from the inline network or inline network group to the inline tool, inline tool group, or inline series. The IPv6 traffic received from the inline network is forwarded as IPv6 traffic in tool-port. The translation from IPv6 packet in network-port to IPv4 packet in tool port will not be supported.
- inline out-of-band forwarding—Packets from inline ports can be sent to regular tool ports.
- inline bypassing—Packets can be put in loopback between two ports of an inline network.
- TLS/SSL forwarding—Packets from an inline network or inline network group can be sent to GigaSMART, then from GigaSMART to an inline tool, inline tool group, or inline series.
- GigaSMART out-of-band forwarding—Packets from GigaSMART can be copied to tool ports.

Failover

The inline bypass module detects failure either through link loss or tool heartbeat failure.

The inline bypass module supports the following failover actions:

- inline tool failover action—Specifies the failover action taken in response to a failure of an inline tool.
- inline tool group failover actions—Specifies 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.
- inline tool series failover actions—Specifies the failover action taken in response to a failure of an inline tool series as a whole. 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 for the inline tool series is respected.

The inline bypass failover actions are configurable. Refer to the *GigaVUE-OS CLI Reference Guide* for the actions and default values of the inline-tool, inline-tool-group, and inline-serial commands.

The virtual port also has configurable failover actions. Refer to the *GigaVUE-OS CLI Reference Guide* for the actions and default value of the vport command.

The GigaSMART module might also have events such as port down or card down. These failover actions are not configurable but are triggered by events. These events should not impact traffic.

Out-of-Band and Inline Tools

After decryption, traffic can be sent to multiple tools. The tools can be either inline or out-of-band.

Out-of-band tools process the decrypted packets offline. The tools are connected to the GigaVUE node through tool or hybrid ports, GigaStream, or port groups with tool or hybrid ports. The out-of-band tools receive a copy of the decrypted packets from the GigaSMART module. This is referred to as GigaSMART out-of-band forwarding.

Inline tools process the decrypted packets inline. Inline tools are connected to the GigaVUE node through inline bypass (BPS) modules.

An out-of-band tool might be an Intrusion Detection System (IDS) examining decrypted packets:

- If it detects a threat, the IDS will send a notification back, but does not have the ability to act.

An inline tool might be an Intrusion Prevention System (IPS) examining decrypted packets:

- If it does not detect a threat in the decrypted packets, the traffic comes out of the inline tool and goes back to the GigaSMART module to be re-encrypted and sent to the server.
- If it detects a threat, the IPS can act. The action depends on the tools' behavior. It can either terminate the connection or modify packets
- If the IPS terminates the connection, then GigaVUE node will terminate the connection between the client and the server.
- If the IPS modifies packets, then the modified packets will come out of the inline tool, go to the GigaSMART module to be re-encrypted, and sent to the server.

When an inline tool is shared, you must:

- configure the inline second level out-of-band map to forward proxy traffic from GigaSMART to the out-of-band tool port.

When an inline tool is not shared, you must configure only the inline second level out-of-band map to forward proxy and non-proxy traffic from GigaSMART to the out-of-band tool port.

NOTE: When OOB-copy is configured in a non-shared inline-tool with the tag set as "inline", both proxy and non-proxy out-of-band traffic will be copied with the proxy map VLAN tag.

Service Chaining of Decrypted Traffic

Service chaining of decrypted traffic may be required for compliance purposes.

This is done by directing the decrypted traffic to a hybrid port and applying the required GigaSMART operations on the traffic that is looped back from the hybrid port, before forwarding the traffic to the out-of-band tool.

The GigaSMART operations must be configured on a different GigaSMART engine than the one used for inline TLS/SSL decryption.

Inline TLS/SSL Decryption Deployments

There are two ways to deploy inline TLS/SSL decryption as follows:

- sessions are inbound
- sessions are outbound

Refer to Figure 44Inbound Deployment of Inline TLS/SSL Decryption for an example of an inbound deployment. The client is on the Internet. The server and the GigaVUE node are located within the same enterprise network, with the GigaVUE node deployed on the server side. The GigaVUE node needs access to the private keys of the server to perform Man-in-the-Middle (MitM) decryption.

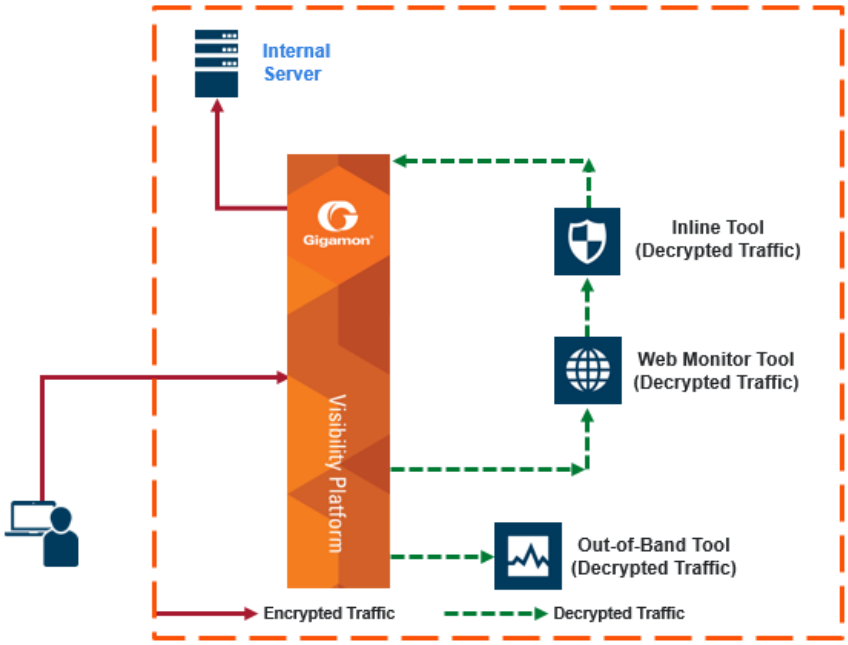

Figure 44 *Inbound Deployment of Inline TLS/SSL Decryption*

Use case for inline TLS/SSL decryption:

- Clients on the Internet
- Servers in internal network
- Organization has the private key of the server
- Diffie-Hellman and Perfect Forward Secrecy is being used

Refer to Figure 45Outbound Deployment of Inline TLS/SSL Decryption for an example of an outbound deployment. The client and the GigaVUE node are located within the same enterprise network, with the GigaVUE node deployed on the client side. The server is located in another network on the Internet. In this deployment, the role of the GigaVUE node is that of a Man-in-the-Middle (MitM). In this deployment, the GigaVUE node does not have access to the private keys of the server, but as a trusted MitM, the GigaVUE node can look at TLS/SSL traffic.

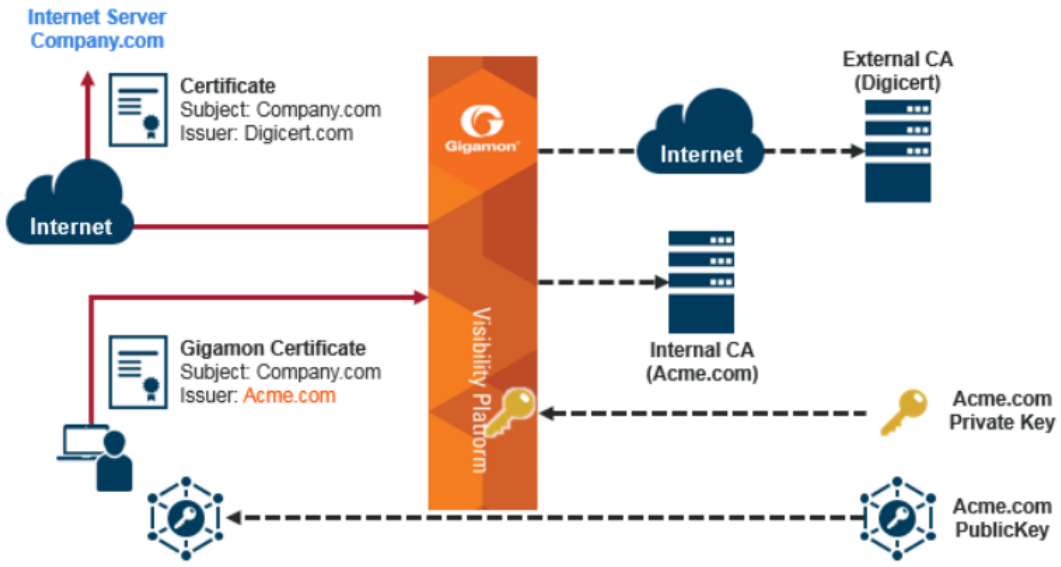

Figure 45 *Outbound Deployment of Inline TLS/SSL Decryption*

GigaVUE Modules for Inline SSL Decryption

Inline SSL decryption is supported on GigaVUE-HC1, GigaVUE-HC1-Plus, and GigaVUE-HC3 nodes, with the GigaSMART and inline bypass modules installed on the same node.

For physical inline bypass, install a fiber bypass (BPS) combo module. On GigaVUE-HC1, copper TAP also supports physical bypass.

Table 10: Inline Bypass Modules lists the inline bypass modules.

Table 10: Inline Bypass Modules

GigaVUE Node	Description
GigaVUE-HC1	Bypass Combo Module 10Gb SX/SR (50/125µm multi-mode)
	Copper TAP module
GigaVUE-HC3	Bypass Combo Module 100Gb/40Gb SR4 MPO

Figure 46GigaVUE Modules: GigaSMART and Inline Bypass shows a GigaVUE-HC3 with the GigaSMART module and the inline bypass (BPS) module. The GigaSMART module contains the SSL decryption software. The inline network ports are on the inline bypass module. The inline and out-of-band tool ports are on the same GigaVUE node.

For inline traffic, the inline network ports and the inline tool ports each need two links, called port pairs, for bidirectional traffic. For out-of-band (offline) traffic, only one link is needed, because the traffic is not bidirectional.

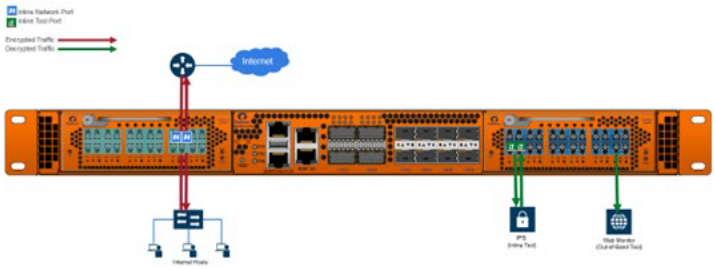

Figure 46 *GigaVUE Modules: GigaSMART and Inline Bypass*

Packet Flows

Normally, a client and server talk directly to each other, such as when you are using a browser to go to a bank website, a health care provider, or a search engine.

As shown in Figure 45Outbound Deployment of Inline TLS/SSL Decryption, the GigaVUE node is placed in the middle between the client and the server. All traffic from the Internet goes through the GigaVUE node.

Incoming traffic arrives on an ingress inline network port on the inline bypass module.

SSL traffic, which is TCP traffic, is directed to the GigaSMART module. Until the traffic is processed by the GigaSMART module, it is not known if it is SSL traffic or not.

The GigaSMART module decides what traffic is bypassed, what traffic is sent to tools without decryption, and what traffic is decrypted and then sent to tools. So, there are three types of decisions as follows:

- to bypass or not
- to decrypt or not
- to send to tools or not

Figure 47Packet Flow for Inline SSL Decryption shows the flow for a configuration consisting of a single inline network, a single inline tool, and a single out-of-band tool.

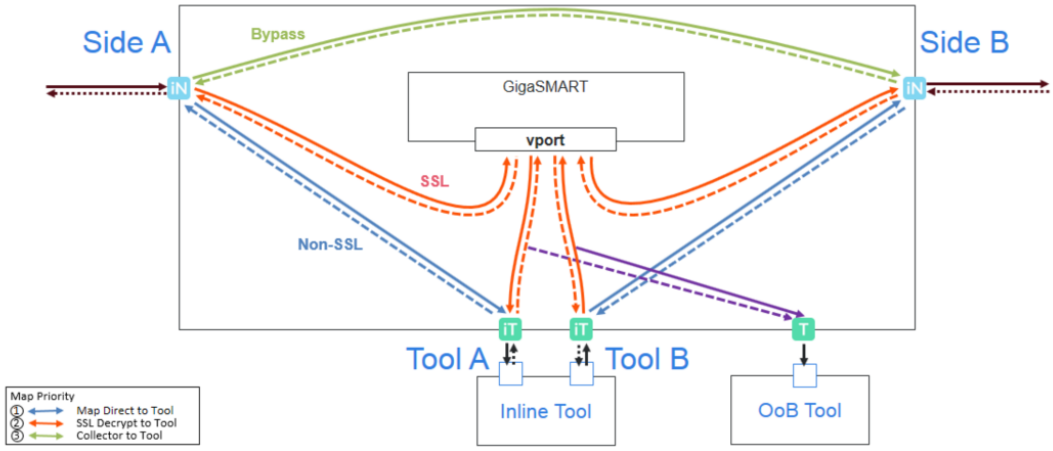

Figure 47 *Packet Flow for Inline SSL Decryption*

Traffic enters the inline SSL decryption solution at the side A inline network port on the inline bypass module.

Some traffic can be bypassed. That traffic goes from the side A inline network port to the side B inline network port on the inline bypass module as shown in the solid green line in Figure 47Packet Flow for Inline SSL Decryption. Bidirectional traffic is shown by dotted lines.

Traffic that is not encrypted (non-SSL) can be sent to tools for inspection. That traffic goes from the side A inline network port to the inline tool A port, to the inline tool, and from the inline tool to the inline tool B port and the side B inline network port on the inline bypass module without going through the GigaSMART module as shown in the solid blue lines in Figure 47Packet Flow for Inline SSL Decryption.

Traffic that is encrypted (SSL) goes from the side A inline network port on the inline bypass module to the GigaSMART module. That traffic goes to a virtual port (vport) that directs traffic to an inline SSL GigaSMART operation as shown by the red line on the left in Figure 47Packet Flow for Inline SSL Decryption. The traffic is decrypted in the GigaSMART module based on policy configuration, sent to the inline tool A port to the inline tool and from the inline tool to the inline tool B port, then back to the GigaSMART module for re-encryption as shown in the solid red lines in the center of Figure 47Packet Flow for Inline SSL Decryption. Finally, the re-encrypted traffic is sent from the GigaSMART module to the side B inline network port on the inline bypass module as shown by the red line on the right in Figure 47Packet Flow for Inline SSL Decryption.

The out-of-band tool can also receive the traffic as shown by the purple line on the right in Figure 47Packet Flow for Inline SSL Decryption.

Starting in software version 5.2, an out-of-band map from a virtual port to a single tool port is supported. Starting in software version 5.3.01, an out-of-band map from a virtual port to multiple tool ports is supported. The ports can be tool, hybrid, or GigaStream. Out-of-band maps from a vport to port groups are also supported when the ports in the group are tool or hybrid.

Modules Matrix

For traffic that is encrypted, there is a question as to whether or not it needs to be decrypted. For example, is there a policy for or against decrypting that traffic? There might be a policy, such as do not decrypt financial or health care traffic, or there might be a decrypt list that states that traffic from a particular site should always be decrypted, or there might be a no-decrypt list that states that traffic from a particular site should always be bypassed. So encrypted packets need to be filtered because some packets will not be decrypted, while others will be decrypted.

Table 11: GigaVUE Modules Used for Type of Traffic is a matrix of the GigaVUE module used for different types of traffic.

Table 11: GigaVUE Modules Used for Type of Traffic

Type of Traffic	GigaSMART	Inline Bypass	GigaSMART
Bypassing GigaSMART	No	Yes - to network	No
Not encrypted (non-SSL)	Yes	Yes - to tools then to network	Yes, depending on the configuration
Encrypted (SSL), and to be decrypted	Yes	Yes - to tools	Yes - to be re-encrypted, then to network

The GigaSMART module does the decryption as well as handling policies, no-decrypt lists, and decrypt lists. The decision to decrypt or not is made in the GigaSMART module.

Inline SSL Traffic Filtering

Because SSL/TLS connections can carry sensitive data, some organizations may require the SSL/TLS connections to avoid inspection. The SSL connections that carries user data such as financial or health care information can be bypassed without inspection, based on a configured policy.

Based on the decryption policies, some connections are not decrypted and are passed through, optionally to and through tools, without decryption. The inline SSL decryption solution respects data privacy and supports compliance.

Inline SSL decryption provides different ways to filter traffic, as follows:

- No-decrypt lists specify traffic to always pass through. A no-decrypt list policy states that traffic from certain sites should always skip decryption. Refer to No-decrypt Listing Policy.
- Decrypt lists specify traffic to always decrypt. A decrypt policy states that traffic from certain sites should always be decrypted. Refer to Decrypt Listing Policy.
- Both No-decrypt lists and Decrypt lists support comments, IP addresses, IP subnets and explicit wildcards for entries and domain rules.
- URL Web Services categorizes the URLs by their type, such as MyBank.com is a financial institution, so as a policy, do not decrypt that traffic. This is also called URL filtering. Typically, banking and health care information are not decrypted. Refer to URL Categorization.
- Policy rules based on network attributes, such as
  - Source IPv4 address
  - Destination IPv4 address
  - VLAN
  - L4 port

No-decrypt Listing Policy

No-decrypt lists are typically used in environments where the default is to decrypt, excepting for certain sites or classes of sites which cannot be decrypted for legal or compliance reasons. By default, traffic that is not to be decrypted is forwarded to the tools unless otherwise configured.

A no-decrypt list file can contain a maximum of 10,000 entries.

Decrypt Listing Policy

Decrypt listing is typically used at sites where specific classes of connections must be decrypted, although the default for other traffic is not to decrypt. Decrypt listed domains and host names will always be decrypted.

A decrypt list file can contain a maximum of 10,000 entries.

Rules and Notes while configuring a No-Decrypt/Decrypt List Policy

1. The maximum domain/hostnames support per list is 10000.
2. IP Subnets are supported from 5.13.01 version. Example, 10.10.10.0/24.
3. Special characters are not supported unless they are used to define domain names, such as * . - @ are supported for domain names and / is supported if IP subnet is defined. # is supported to comment out a line. Example of a text file format would be as follows:
   - *.google.com
   - www.gigamon.com
   - gigamon.com
   - domain-registration.com.us
   - 10.10.1.1
   - 10.10.1.0/24
4. Range of IP addresses are not supported example, 10.10.10.10-20.
5. Use a newline for each entry. Adding characters such as , ; are not supported.
6. On GigaVUE-OS pre-5.9 versions, gigamon.com as an entry matches gigamon.com and all its subdomains, that is, abc.gigamon.com, abc.xyz.gigamon.com etc.
7. Starting from GigaVUE-OS v5.9, gigamon.com as an entry matches only gigamon.com. To match all subdomains of gigamon.com on v5.9+, use *.gigamon.com.
8. If the system has large set of decrypt/no-decrypt list entries, GigaVUE-FM stats page and CLI stats command does not display any output. Wait for 5 to 10 minutes after reloading to check the inline SSL show stats command in CLI and stats page in GigaVUE-FM.

IP Address Subnet with Longest Prefix Match(LPM)

The No-decrypt and decrypt database allows the user to utilize IP subnets. This allows the user to configure overlapping IP addresses, in decrypt and no-decrypt database. The decision to decrypt or no-decrypt will be based on the longest prefix match of the IP entries available in the decrypt /no-decrypt database.

The format is as follows *subnet (no space) /prefix*.Eg: 191.1.1.0/32

URL Categorization

URL categories make it convenient to apply policies on all the possible URLs by simplifying the number of policy rules. Categorization is based on the hostname in the TLS Server Name Indication (SNI) or the hostname from the server certificate if there is no SNI. There are 83

categories including one for Uncategorized, which is a default category for URLs that do not match any of the other 82 categories. The categories are fixed meaning that categories cannot be added, deleted or modified.

GigaSMART ships with a local database of 1M entries and will also perform a cloud lookup for those hosts not found in the local database. The URL Web Service provides the URL categorization. The URL database is updated daily from the URL Web Service. Each update likely adds new entries and purges other entries, but always keeping the database at 1M entries.

NOTE: When a URL is not in the cache, for cloud look-ups the stack port interface on GigaSMART must be configured to provide Internet access. Refer to Set up the Stack Port Interface for more information.

URL Look-ups and Caching

As part of the iSSL processing, URL look-ups are performed against the database. If the URL is not found in the database, then a lookup is performed against the local cache. If the URL is not found in the local cache, then an external lookup to the URL Web Services may be performed, if configured. If the URL is found in the external look-up, then it is dynamically saved in the local cache. Future look-ups may then find the URL in the local cache instead of requiring the external look-up.

NOTE:

- For TLS connections containing SNI in the Client Hello, do not perform URL look-up in the certificate phase.
- CN based evaluation can be performed using the configuration option.

The local cache can hold up to 250k entries (in addition to the 1M entry database). The local cache works like a circular buffer – older entries are discarded to make room for newer ones if the cache is full. Each cache entry is valid for 24 hours and updated with current time stamp whenever an entry is made. If an expired entry is encountered, a new query is issued to the URL Web Services to refresh the entry in the cache. Expired entries don't get actively deleted from the cache.

While the URL Web Service is hosted on AWS, external look-ups need to occur very quickly. Gigamon provides a timeout option, up to 10 seconds for external URL look-ups via the URL cache miss defer option.

NOTE:

- URLs may get recategorized as part of updates from the URL Web Services. This is transparent to Gigamon and customers.

- The URL category classification is fixed, and a new category cannot be added. Gigamon provides the no-decrypt list/decrypt list functionality, which can achieve the same result as creating a custom category.
- If a URL belongs to multiple categories, any no-decrypt policy would take precedence over all decrypt policies.

Inline SSL URL categories

The following are the list of Inline SSL URL categories with examples.

NOTE: Gigamon does not endorse any of the following categories, descriptions, and examples, but replicated the information from the URL Web Services. Some categories are presented without examples since they are not appropriate.

Category Name	Description and Examples
Abortion	Abortion topics, either pro-abortion and anti-abortion.
Abused Drugs	Discussion or remedies for illegal, illicit, or abused drugs such as heroin, cocaine, or other street drugs. This category includes information on the misuse of non-proscribed substances (eg. "glue sniffing"), or the misuse of prescription medications.
Adult and Pornography	Sexually explicit material for the purpose of arousing a sexual or prurient interest. Online groups, including newsgroups and forums, that are sexually explicit in nature.
Alcohol and Tobacco	Sites that provide information on, promote, or support the sale of alcoholic beverages or tobacco products and associated paraphernalia.
Auctions	Sites that support the offering and purchasing of goods between individuals as their main purpose. Does not include classified advertisements. • http://ebay.co • http://quibids.com
Botnets	These are URLs, typically IP addresses, which are determined to be part of a Bot network, from which network attacks are launched. Attacks may include SPAM messages, DOS, SQL injections, proxy jacking, and other unsolicited contacts.
Business and Economy	Business firms, corporate websites, business information, economics, marketing, management, and entrepreneurship. • http://samsung.com • http://ups.com

Category Name	Description and Examples
Content Delivery Networks	Delivery of content and data for third parties, including ads, media, files, images, and video. • http://metacdn.co • http://edgestream.com
Cheating	Sites that support cheating on examinations and contain such materials, including free essays, exam copies, plagiarism, etc.
Computer and Internet Info	General computer and Internet sites, technical information. SaaS sites and other URLs that deliver internet services. • http://ranking.co • http://system.netsuite.com
Computer and Internet Security	Computer/Internet security, security discussion groups. • http://siteadvisor.co • http://webroot.com
Confirmed Spam Sources	Confirmed SPAM sources.
Cult and Occult	Internet resources which include discussion of astrology, spells, curses, magical powers, satanic rituals or supernatural beings. This includes horoscope sites.
Dating	Dating websites focused on establishing personal relationships. • http://eharmony.com
Dead Sites	These are dead sites that do not respond to http queries. Policy engines should usually treat these as “Uncategorized” sites. • http://g00gle.com • http://whitehouse.info
Dynamic Content	Domains that generate content dynamically based on arguments to their URL or other information (like geo-location) on the incoming web request. • booking.com
Education Institution	Pre-school, elementary, secondary, high school, college, university, and vocational school and other educational content and information including enrollment, tuition, and syllabus. • http://mit.edu • http://ox.ac.uk
Entertainment and Arts	Motion pictures, videos, television, music and programming guides, books, comics, movie theatres, galleries, artists or reviews on entertainment. • http://eonline.com

Category Name	Description and Examples
	• http://warnerbros.com
Fashion and Beauty	Fashion or glamour magazines, beauty, clothes, cosmetics, style. • http://visionmodels.co.uk • http://genejuarez.com
Financial Services	Banking services and other types of financial information, such as loans, accountancy, actuaries, banks, mortgages, and general insurance companies. Does not include sites that offer market information, brokerage or trading services. • http://firstpremierbankcards.com • http://paypal.com
Gambling	Gambling or lottery web sites that invite the use of real or virtual money. Information or advice for placing wagers, participating in lotteries, gambling, or running numbers. Virtual casinos and offshore gambling ventures. Sports picks and betting pools.
Games	Playing or downloading, video games, computer games, electronic games, tips, and advice on games or how to obtain cheat codes. Also includes sites dedicated to selling board games as well as journals and magazines dedicated to game playing. • http://duowan.com • http://ubi.com
Government	Information on government, government agencies and government services such as taxation, public, and emergency services. Also includes sites that discuss or explain laws of various governmental entities. Includes local, county, state, and national government sites. • http://www.nasa.gov • http://premier-ministre.gouv.fr
Gross	Sites that contain material which describe or display material which would be considered foul or disgusting. Examples would include bodily fluids, injuries, gore.
Hacking	Illegal or questionable access to or the use of communications equipment/software. Development and distribution of programs that may allow compromise of networks and systems.
Hate and Racism	Sites that contain content and language in support of hate crimes and racism.
Health and Medicine	General health, fitness, well-being, including traditional and non-traditional methods and topics. Medical information on ailments, various conditions, dentistry, psychiatry, optometry, and other

Category Name	Description and Examples
	specialties. • http://webmd.com • http://missionvalleymedical.com
Home and Garden	Home issues and products, including maintenance, home safety, decor, cooking, gardening, home electronics, design, etc. • http://homedepot.com • http://waysidegardens.com
Hunting and Fishing	Sport hunting, gun clubs, and fishing. • http://fishingworks.com • http://wildlifelicense.com
Illegal	Criminal activity, copyright and intellectual property violations, etc.
Image and Video Search	Photo and image searches, online photo albums/digital photo exchange, image hosting. • http://images.google.fr • http://gettyimages.com
Individual Stock Advice and Tools	Promotion and facilitation of securities trading and management of investment assets. Also includes information on financial investment strategies, quotes, and news. • http://stockstar.com • http://morningstar.com
Internet Communications	Internet telephony, messaging, VoIP services and related businesses. • http://skype.com • http://www.chatib.com/
Internet Portals	Web sites that aggregate a broader set of Internet content and topics, and which typically serve as the starting point for an end user. • http://yahoo.com • http://qq.com
Job Search	Assistance in finding employment, and tools for locating prospective employers, or employers looking for employees. • http://monster.com • http://51job.com
Keyloggers and Monitoring	Downloads and discussion of software agents that track a user's keystrokes or monitor their web surfing habits.
Kids	Sites designed specifically for children and teenagers. • http://www.mundogaturro.com

Category Name	Description and Examples
	• http://www.poptropica.com
Legal	Legal websites, law firms, discussions and analysis of legal issues. • http://www.pepperlaw.com • http://earlcaterlaw.com
Local Information	City guides and tourist information, including restaurants, area/regional information, and local points of interest. • http://downtownlittlerock.com • http://sandiegorestaurants.com
Malware Sites	Malicious content including executables, drive-by infection sites, malicious scripts, viruses, trojans, and code.
Marijuana	Marijuana use, cultivation, history, culture, legal issues.
Military	Information on military branches, armed services, and military history. • http://defense.gov • http://www.mod.uk
Motor Vehicles	Car reviews, vehicle purchasing or sales tips, parts catalogs. Auto trading, photos, discussion of vehicles including motorcycles, boats, cars, trucks and RVs. Journals and magazines on vehicle modifications. • http://www.carmax.com • http://carsales.com.au
Music	Music sales, distribution, streaming, information on musical groups and performances, lyrics, and the music business. • http://itunes.com • http://bandcamp.com
News and Media	Current events or contemporary issues. Also includes radio stations, magazines, online newspapers, headline news sites, newswire services, personalized news services, and weather sites. • http://abcnews.go.com • http://newsoftheworld.co.uk
Nudity	Nude or seminude depictions of the human body. These depictions are not necessarily sexual in intent or effect but may include sites containing nude paintings or photo galleries of artistic nature.
Online Greeting Cards	Online Greeting card sites. • http://123greetings.com • http://greeting-cards.com

Category Name	Description and Examples
Online Personal Storage	Online storage and posting of files, music, pictures, and other data. • http://box.net • http://freefilehosting.net
Open HTTP Proxies	The proxy servers that are accessible by any Internet user.
P2P (Peer to Peer)	Peer to peer clients and access that includes torrents, music download and programs.
Parked Sites	Parked domains are URLs which host limited content or click-through ads which may generate revenue for the hosting entities but generally do not contain content useful to the end user. Also includes Under Construction, folders, and web server default home pages. • http://000.com • http://buythisdomain.com
Pay to Surf	Sites that pay users in the form of cash or prizes, for clicking on or reading specific links, email, or web pages.
Personal Sites and Blogs	Personal websites posted by individuals or groups, as well as blogs. • http://blogger.com • http://wordpress.org
Philosophy and Political Advocacy	Politics, philosophy, discussions, promotion of a particular viewpoint or stance in order to further a cause. • http://philosophynow.org • http://political.com
Phising and Other Frauds	Phishing, pharming, and other sites that pose as a reputable site, usually to harvest personal information from a user. These sites are typically quite short-lived, so examples may not last long.
Private IP Addresses	RFC 1918, Address Allocation for Private Intranets. 10.0.0.0 - 10.255.255.255 (10/8 prefix) 192.168.0.0 - 192.168.255.255 (192.168/16 prefix)
Proxy Avoid and Anonymizers	Proxy servers and other methods to gain access to URLs in any way that bypasses URL filtering or monitoring. Web-based translation sites that circumvent filtering.
Questionable	Tasteless humor, “get rich quick” sites, and sites that manipulate the user experience or client in some unusual, unexpected, or suspicious manner.
Real Estate	Information on renting, buying, or selling real estate or properties.

Category Name	Description and Examples
	Tips on buying or selling a home. Real estate agents, rental or relocation services, and property improvement. • http://prudentialproperties.com • http://realtor.com
Recreation and Hobbies	Information, associations, forums and publications on recreational pastimes such as collecting, kit airplanes, outdoor activities such as hiking, camping, rock climbing, specific arts, craft, or techniques; animal and pet related information, including breed-specifics, training, shows and humane societies. • http://petloverspublications.com • http://craftster.org
Reference and Research	Personal, professional, or educational reference material, including online dictionaries, maps, census, almanacs, library catalogues, genealogy, and scientific information. • http://reference.com • http://wikipedia.org
Religion	Conventional or unconventional religious or quasi-religious subjects as well as churches, mosques, synagogues, or other places of worship. • http://therocksandiego.org • http://biblesociety.ca
Search Engines	Search interfaces using key words or phrases. Returned results may include text, websites, images, videos, and files. • http://google.com • http://sogou.com
Sex Education	Information on reproduction, sexual development, safe sex practices, sexually transmitted diseases, sexuality, birth control, sexual development, and contraceptives. • http://sexetc.org
Shareware and Freeware	Sites that contains softwares, screensavers, icons, wallpapers, utilities, ringtones including downloads that request a donation on open source projects. • http://download.com • http://sourceforge.net
Shopping	Department stores, retail stores, company catalogs and other sites that allow online consumer or business shopping to purchase goods and services. • http://amazon.com

Category Name	Description and Examples
	• http://groupon.com
Social Network	Social networking sites that have user communities where users interact, post messages, pictures, and otherwise communicate. • http://facebook.com • http://twitter.com
Society	A variety of topics, groups, and associations relevant to the general populace, broad issues that impact a variety of people, including safety, children, societies, and philanthropic groups. • http://dar.org • http://unicefusa.org
Spam URLs	URLs contained in SPAM.
Sports	Team or conference web sites, international, national, college, professional scores and schedules; sports-related online magazines or newsletters, fantasy sports and virtual sports leagues. • http://nba.com • http://schoenen-dunk.de
Spyware and Adware	Spyware or Adware sites that provide or promote information gathering or tracking that is unknown to, or without the explicit consent of, the end user or the organization, also unsolicited advertising popups and programs that may be installed on a user's computer.
Stream Media	Sales, delivery, or streaming of audio or video content, including sites that provide downloads for such viewers. • http://youtube.com • http://ustream.tv
Swimsuits and Intimate Apparel	Swimsuits, intimate apparel or other types of suggestive clothing.
Training and Tool	Distance education, trade schools, online courses, vocational training, software training, and skills training. • http://trainingtools.com • http://prezi.com
Translation	Language translation sites that allow users to see URL pages in other languages. • http://translate.google.com • http://microsofttranslator.com
Travel	Airlines and flight booking agencies. Travel planning, reservations, vehicle rentals, car rentals, descriptions of travel destinations,

Category Name	Description and Examples
	promotions for hotels or casinos. • http://cheapflights.com • http://expedia.com
Uncategorized	Sites that have not been categorized by URL Web Service.
Unconfirmed Spam Sources	Unconfirmed SPAM sources.
Violence	Sites that advocate violence, depictions and methods, including game/comic violence, and suicide.
Weapons	Sales, reviews, descriptions of weapons such as guns, knives, martial arts accessories.
Web Advertisements	Advertisements, media, content, and banners. • http://casalemedia.com • http://justwebads.com
Web Based Email	Sites offering web-based email and email clients. • http://google.com/mail • http://foxmail.com
Web Hosting	Free or paid hosting services for web pages and information concerning their development, publication, and promotion. • http://siteground.com • http://bluehost.com

Proxy Server Profile for URL Categorization and Certificate Revocation status

To ensure a stable security network you can now redirect URL look-ups and Certificate Revocation status checks to a Proxy Server Profile. This Proxy Server profile will be attached to your Inline SSL deployment . To learn more refer to Proxy Server Configuration.

Get Started with Inline TLS/SSL Decryption

This section describes the prerequisites needed before you begin configuring inline TLS/SSL decryption.

Topics:

- Before You Begin
  - Supported Platforms
  - GigaSMART Licensing
  - GigaSMART Compatibility
- Installation
  - Install GigaVUE Modules
  - Install Software Version
  - Install U-Boot Version on GigaVUE-HC3
  - Install MitM Certificates in Client Trust Store
- Initial Configuration
  - Configure Stack Port Interface
  - Set up GigaSMART for Inline TLS/SSL decryption on GigaVUE-HC1
  - TLS/SSL Decryption for Inline Tools
  - Configure Primary Certificate and Key

Before You Begin

Supported Platforms

Inline TLS/SSL decryption is supported on GigaVUE-HC1, and GigaVUE-HC3 nodes.

GigaSMART Licensing

The required GigaSMART license is TLS/SSL Decryption for Inline and Out-of-Band Tools.

NOTE: A GigaSMART license for Inline TLS/SSL is required to upgrade a passive TLS/SSL to inline TLS/SSL.

GigaSMART Compatibility

Inline TLS/SSL decryption is not compatible with any other GigaSMART operations, including Passive TLS/SSL decryption. Configure inline TLS/SSL decryption on a GigaSMART engine that is not shared with any other GigaSMART operation. Inbound and Outbound Inline-TLS/SSL decryption can be deployed on a single GigaSMART engine.

However, in GigaVUE-HC1 nodes, you can configure inline TLS/SSL along with other GigaSMART applications. Refer to Set up GigaSMART for Inline TLS/SSL Decryption on GigaVUE-HC1 for more information.

Installation

Install GigaVUE Modules

Install the GigaSMART and inline bypass module or copper TAP on the same GigaVUE-HC1 node or install the GigaSMART and inline bypass module on the same GigaVUE-HC3 node.

Install Software Version

Install software version 5.2.xx or higher for the GigaVUE-OS CLI, GigaVUE-OS and GigaVUE-FM.

Install U-Boot Version on GigaVUE-HC3

The U-Boot version on GigaVUE-HC3 nodes must be upgraded to version 2011.06.9 or higher. The upgrade can only be done from the CLI.

To check the U-Boot version, use the following command:

```
(config) # show version
```

For example on a GigaVUE-HC3 node, the following output is displayed:

```
U-Boot version: 2011.06.10
```

If you do not have version 2011.06.10 or higher, you will have to do a U-Boot upgrade, after the image installation. Refer to the *GigaVUE-OS Upgrade Guide* for details on installing an image.

After the image installation of the software, use the following command to upgrade the U-Boot version:

```
(config) # uboot install
```

The binary bootloader code included with the installed image is installed.

NOTE: The newer U-Boot version only goes into effect after a reload.

Install MitM Certificates in Client Trust Store

For an outbound deployment, the Man-in-the-Middle (MitM) certificates must be installed in the client trust store. Install the certificates as a Trusted Root Authority in web browsers on the client PC.

Refer to your browser's documentation for installing CA certificates to the trust store.

Initial Configuration

Set up the Stack Port Interface

A stack port is a port that is used to communicate with other nodes in the stack to expand your network capacity. GigaVUE HC Series devices provide dedicated, configurable ports for stacking.

About the Stack Port Interface

Internet connectivity is required for CRL, OCSP, and URL categorization. The stack port interface must be configured on the GigaSMART engine. You can configure one stack port for all the GigaSMART engines in a GigaVUE node. However, you must configure all the engines with unique DHCP or static IP address.

Refer to Figure 48 Stack Port Location on GigaVUE-HC3 Front for the location of the two stack ports, eth2 (default) and eth3 on the control card in the front of GigaVUE-HC3. You can use either one or both the stack ports for GigaSMART connectivity.

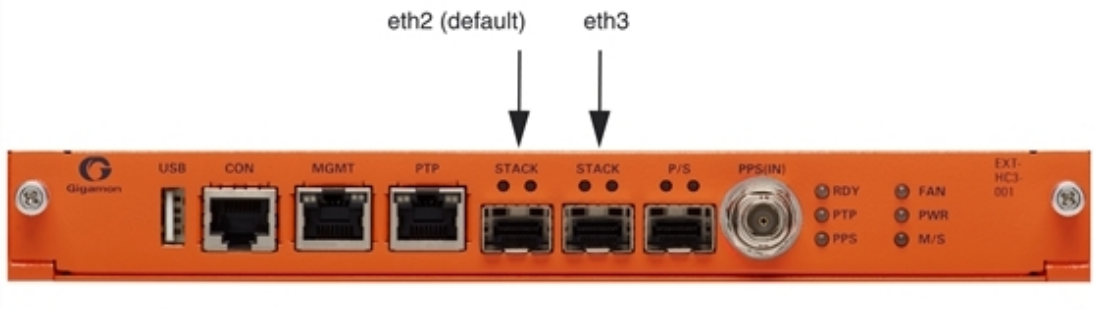

Figure 48 *Stack Port Location on GigaVUE-HC3 Front*

Refer to Figure 49 Stack Port Location on GigaVUE-HC1 Front for the location of the stack port on the front of the GigaVUE-HC1. It is the top port on the right.

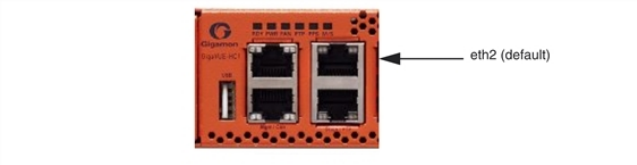

Figure 49 *Stack Port Location on GigaVUE-HC1 Front*

Configure Stack Port Interface

To configure the stack port interface:

1. Go to **Ports > All Ports**. Select a GigaSMART engine port and click **Edit**.
2. Specify either an IP Address, Netmask, Gateway, DNS IP, and optional MTU or select DHCP. Specify a VLAN ID in the range from 20 to 4094. Select the stack port interface, Eth2 or Eth3. The default is Eth2.
3. Click **OK**. The stack port interface is added.

NOTE: Proxy IP addresses cannot be used to configure for internet connectivity in the stack port. TCP Port 80 is used to connect to webroot.

Set up GigaSMART for Inline TLS/SSL Decryption on GigaVUE-HC1

You can configure GigaSMART engine resources on GigaVUE-HC1 to reduce the Inline TLS/SSL resource utilization by 50% and use the rest of the resource to configure the other GigaSMART applications. The Inline TLS/SSL application runs in the following two modes:

- **Standalone mode Enabled**—The Inline TLS/SSL feature takes the entire GigaSMART engine resource. By default standalone mode is enabled for Inline TLS/SSL.
- **Standalone mode Disabled**—The GigaSMART engine resource allocated for Inline TLS/SSL feature is reduced to 50% and the residual GigaSMART engine resource can be configured for other GigaSMART applications.

On GigaVUE-HC1, you can configure the following GigaSMART applications along with the Inline TLS/SSL feature:

  - De-duplication (SMT-HC1-DD1)
  - NetFlow Generation (SMT-HC1-NF1)
  - BSE Combo (SMT-HC1-BSE) - Masking, Slicing, and Trailer
  - Header-stripping (SMT-HC1-HS1)
  - Flow Sampling (SMT-HC1-FVU)
  - Tunneling, ERSPAN (SMT-HC1-TUN)

Limitations

- It is not recommended to configure Inline TLS/SSL feature with other GigaSMART applications, except the applications listed above.
- The Passive TLS/SSL decryption is not recommended to be configured with Inline TLS/SSL feature and combination of NetFlow and TLS/SSL decryption do not work with the Inline TLS/SSL.
- For inline TLS/SSL decryption, Internet connectivity to GigaSMART and clustering is not supported on the same interface, for example, eth2.

Enable Standalone mode for Inline TLS/SSL decryption

To enable Standalone mode for Inline TLS/SSL decryption on GigaVUE-HC1:

1. From the GigaVUE-HC1 device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.
2. Click **New** to create a new GigaSMART group or click **Edit** to modify an existing GigaSMART group.
3. Go to **Inline SSL** under **GigaSMART Parameters**, and select **Standalone**.
4. Click **OK**.

Configure Keychain Password

Configure Keychain Password

For Inbound and Outbound Inline-SSL deployments, the keychain password must be configured before installing the certificates and private keys into the keystore.

Refer to Configure Inline TLS/SSL Decryption Using GigaVUE-FM for the configuration steps.

Configure Primary Certificate and Key

For an outbound deployment, at least one of the CAs must be configured (primary or secondary). For an inbound deployment, a CA is not necessary.

The primary CA re-signs certificates for servers that present a valid certificate. The secondary CA re-signs certificates for servers that are invalid or that fail validation. If the secondary CA is not configured, the primary CA will be used for all certificates.

Refer to Configure Inline TLS/SSL Decryption Using GigaVUE-FM for the configuration steps.

Configure an Inline TLS/SSL Session Logging Server

You can configure an inline TLS/SSL session logging server to store the logged events that are generated when there are any changes made to the devices. You can specify the type of events that must be logged in to the server.

The following table provides a mapping of the severity, log level and its description:

Severity	Log Level	Description
0	Emergency	System is unusable
1	Alert	Action must be taken immediately
2	Critical	Critical condition
3	Error	Error condition
4	Warning	Warning condition
5	Notice	Normal but significant condition
6	Informational	Informational message
7	Debug	Debug message

The logged events are stored in the Common Event Format (CEF) as follows:

<SYSLOG_HEADER> <Timestamp> <hostname:engine> CEF:0|Gigamon|<Device Model>|<GigaVUE OS Version>|<Event ID>|<Event name>|<Severity>|[Extension]

Here is an example of a logged event:

Thu Jun 14 15:50:16 2018 hostname:hc2_test:1/1/e1 CEF:0|Gigamon|HC2|5.5.0|102|SESSION_DECRYPT|6|src=126.1.0.20 dst=126.1.0.10 spt=34267 dpt=443 dhost=example.com cs1Label=Certificate Subject cs1=C\=US, ST\=CA, L\=Santa Clara, CN=*.example.com cs2Label=Cipher Suite cs2=DHE-RSA-AES128-GCM-SHA256

You can view and track these logs to troubleshoot system issues, maintain audit trails, and for compliance purpose.

To configure an inline TLS/SSL session logging server using GigaVUE-FM:

Task	Description	UI Steps
1.	Configure a tool port.	a. From the device view, go to Ports > All Ports. b. Click Quick Port Editor. c. Use Quick search to find the port to configure. d. Set the type as Tool for the required port, and then select Enable. e. Click OK.
2.	Configure an IP interface with an IP address, subnet mask, default gateway, and MTU setting. Assign it to the GigaSMART group.	a. From the device view, go to Ports > Ports > IP Interfaces. b. Click New. The IP Interface page opens. c. In the Alias and Description fields, enter the name and description of the IP interface. d. From the Ports drop-down list, select the tool port that you configured in step 1. e. Select the Type of the IP interface as IPv4 or IPv6. f. Enter the IP address, subnet mask, gateway, and MTU settings in the respective fields. g. From the GS Groups drop-down list, select a GigaSMART group under which you want to configure the inline TLS/SSL session logging server. h. Click OK.
3.	Configure the inline TLS/SSL session logging server under the GigaSMART group to which you assigned the IP interface in the task 2.	a. From the device view, go to GigaSMART > GigaSMART Groups. b. Choose the GigaSMART group to which you assigned the IP interface that you configured in task 2. c. Click Edit. d. Under GigaSMART Parameters > Inline SSL Session Logging, click Add Remote Syslog Server. e. In the Remote Syslog IP and Remote Syslog Port Number fields, enter the IP address and port number of the remote syslog server.

Task	Description	UI Steps
		f. From the Associated IP Interface drop-down list, select the IP interface that you assigned to the GigaSMART group in task 2. g. From the Log Level drop-down list, select the severity log level of the events that you want to send to the inline SSL session logging server. h. Click OK.

Configure Inline TLS/SSL Decryption

This section describes the workflows for configuration inline TLS/SSL decryption using GigaVUE-FM. It also provides the details of the workflows for inline TLS/SSL map.

Introduction to Inline TLS/SSL Map Workflows

In GigaVUE-FM, workflows guide you through configuration steps. For the Inline TLS/SSL Map configuration, there are seven flows, Flow A to Flow G based on which you can perform different configurations.

Go to **Workflows** and select **Inline SSL Map** from the Inline GigaSMART Operations section as shown in Figure 50 Select Inline TLS/SSL Map Configuration.

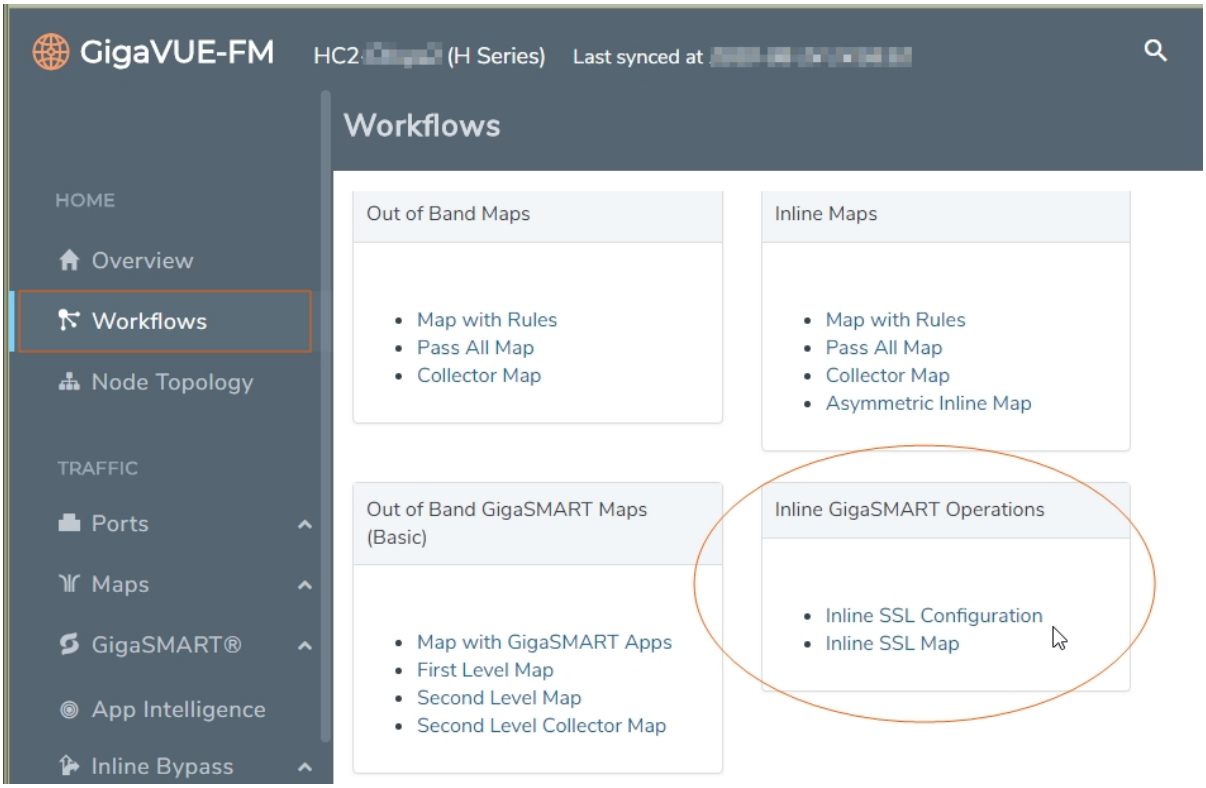

Figure 50 *Select Inline TLS/SSL Map Configuration*

Flow A to Flow G are displayed as shown in Figure 51 Flow A to Flow G.



Figure 51 *Flow A to Flow G*

The following sections describe each flow:

- Flow D
- Flow E
- Flow F

- Flow G

Flow A

Flow A is for the following use case:

- filter HTTP traffic and direct it to the tool(s)
- filter remaining TCP traffic and direct it to the GigaSMART engine for inspection before forwarding it to the tool(s)
- direct all other traffic to the tool(s)

Refer to Figure 52 FLOW A Views for a larger view of Flow A on the left and a pictorial view on the right.

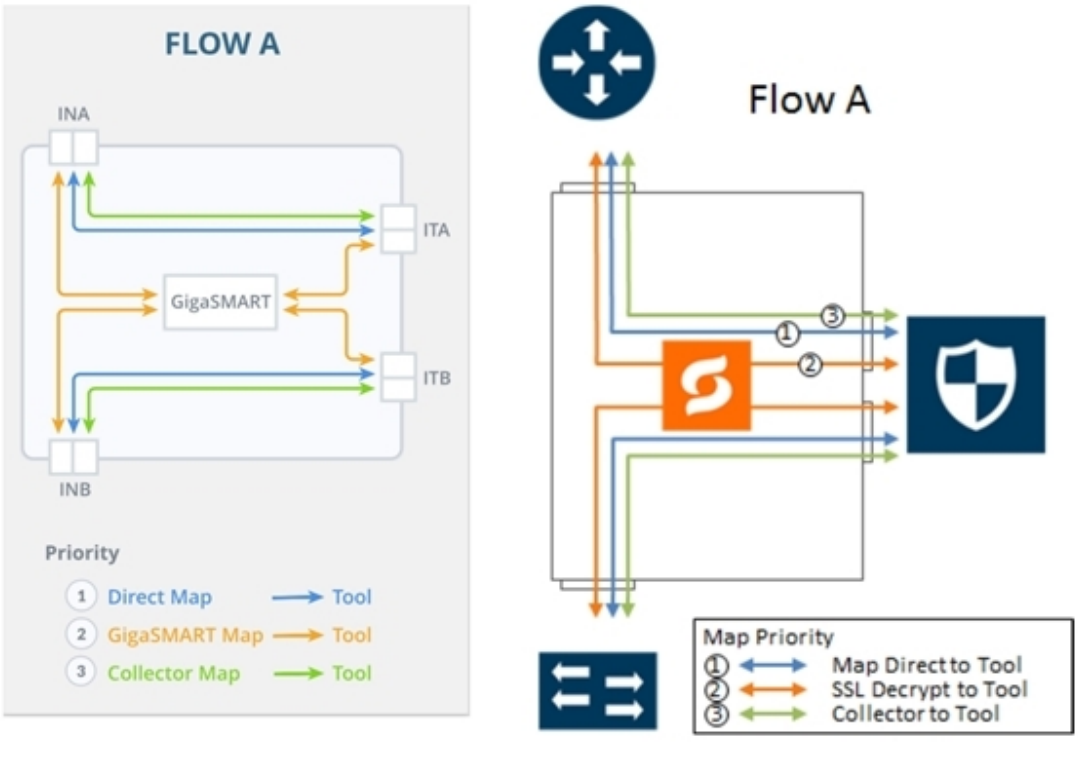

Figure 52 *FLOW A Views*

The map priorities of Flow A are as follows:

1. limit traffic going to decryption
2. selectively forward traffic for decryption
3. direct unselected traffic to a collector, which sends traffic to tool

Flow B

Flow B is for the following use case:

- filter HTTP traffic and direct it to the tool(s)
- filter remaining TCP traffic and direct it to the GigaSMART engine for inspection before forwarding it to the tool(s)
- bypass all other traffic

Refer to Figure 53 FLOW B Views for a larger view of Flow B on the left and a pictorial view on the right.

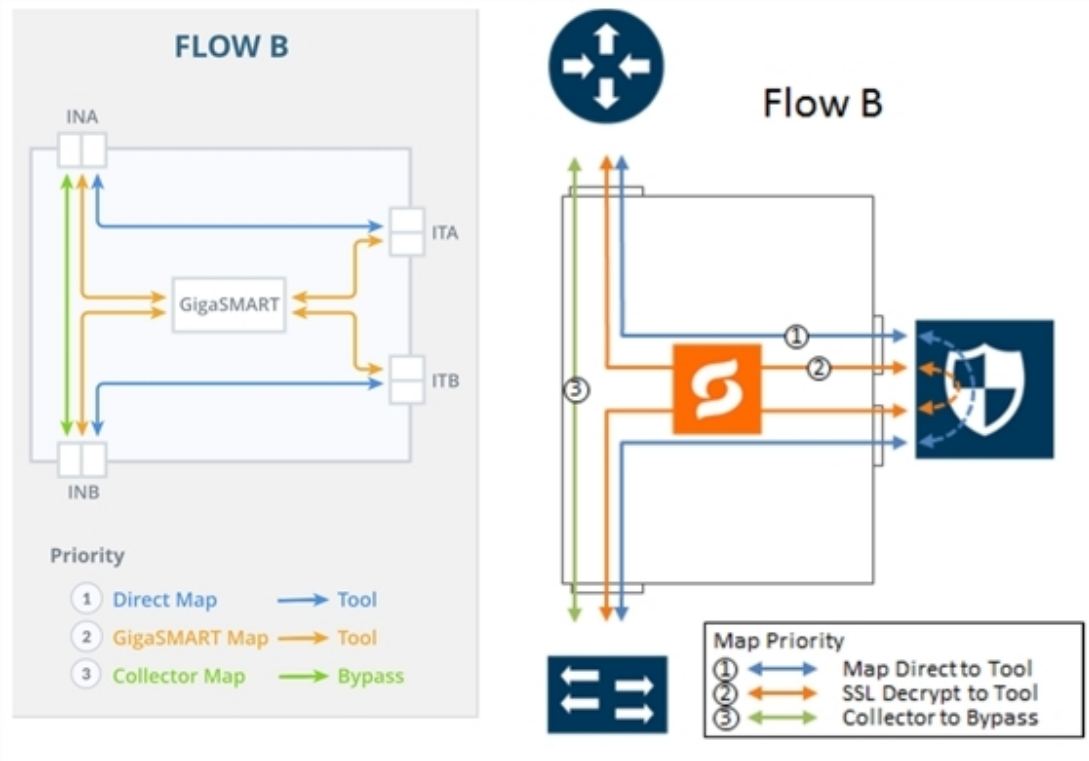

Figure 53 *FLOW B Views*

The map priorities of Flow B are as follows:

1. limit traffic going to decryption
2. selectively forward traffic for decryption
3. direct unselected traffic to a collector, which sends traffic to bypass

Flow C

Flow C is for the following use case:

- filter traffic from or to certain VLANs (for example, a guest WiFi VLAN) and direct it to bypass
- filter remaining TCP traffic and direct it to the GigaSMART engine for inspection before forwarding it to the tool(s)
- direct all other traffic to the tool(s)

Refer to Figure 54 FLOW C Views for a larger view of Flow C on the left and a pictorial view on the right.

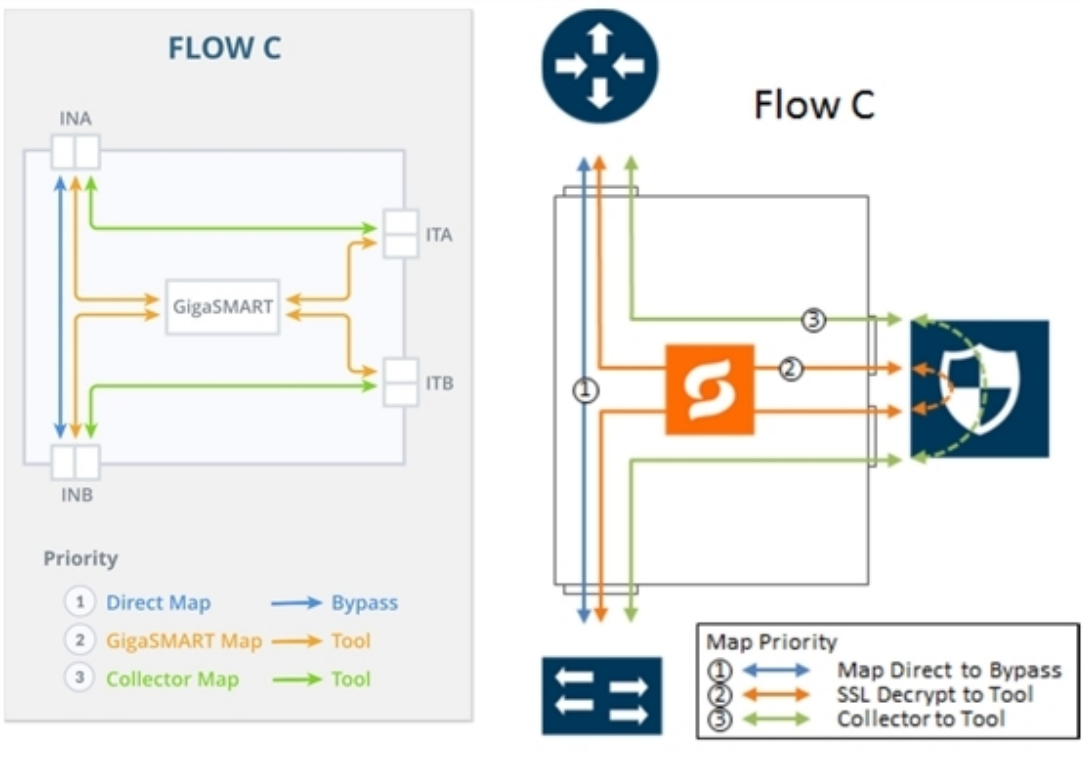

Figure 54 *FLOW C Views*

The map priorities of Flow C are as follows:

1. send trusted traffic to bypass
2. selectively forward traffic for decryption
3. direct unselected traffic to a collector, which sends traffic to tool

Flow D

Flow D is for the following use case:

- filter traffic from certain VLANs (for example, an employee WiFi VLAN) and direct it to the GigaSMART engine for inspection before forwarding it to the tool(s)
- filter traffic from or to selected internal servers and direct it to bypass
- direct all other traffic to tool(s)

Refer to Figure 55 FLOW D Views for a larger view of Flow D on the left and a pictorial view on the right.

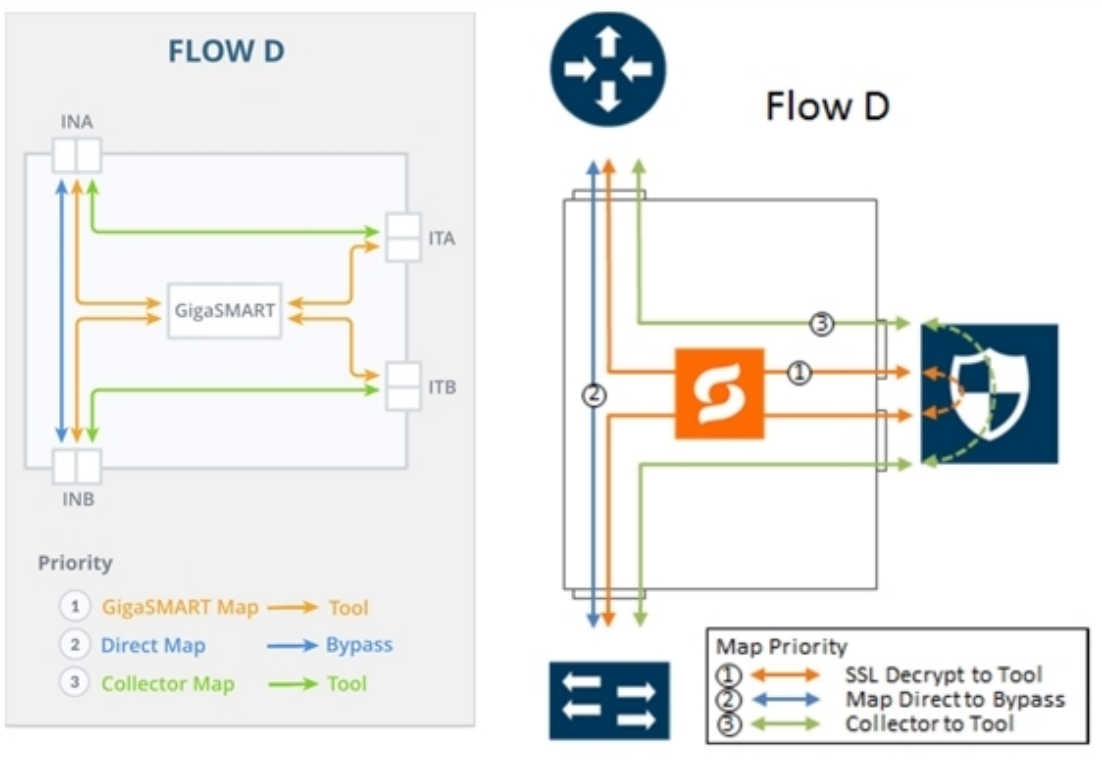

Figure 55 *FLOW D Views*

The map priorities of Flow D are as follows:

1. selectively forward traffic for decryption
2. send trusted traffic to bypass
3. direct unselected traffic to a collector, which sends traffic to tool

Flow E

Flow E is for the following use case:

- filter traffic from certain VLANs (for example, an employee WiFi VLAN) and direct it to the GigaSMART engine for inspection before forwarding it to the tool(s)
- filter traffic from or to selected internal servers and direct it to the tool(s)
- bypass all other traffic

Refer to Figure 56 FLOW E Views for a larger view of Flow E on the left and a pictorial view on the right.

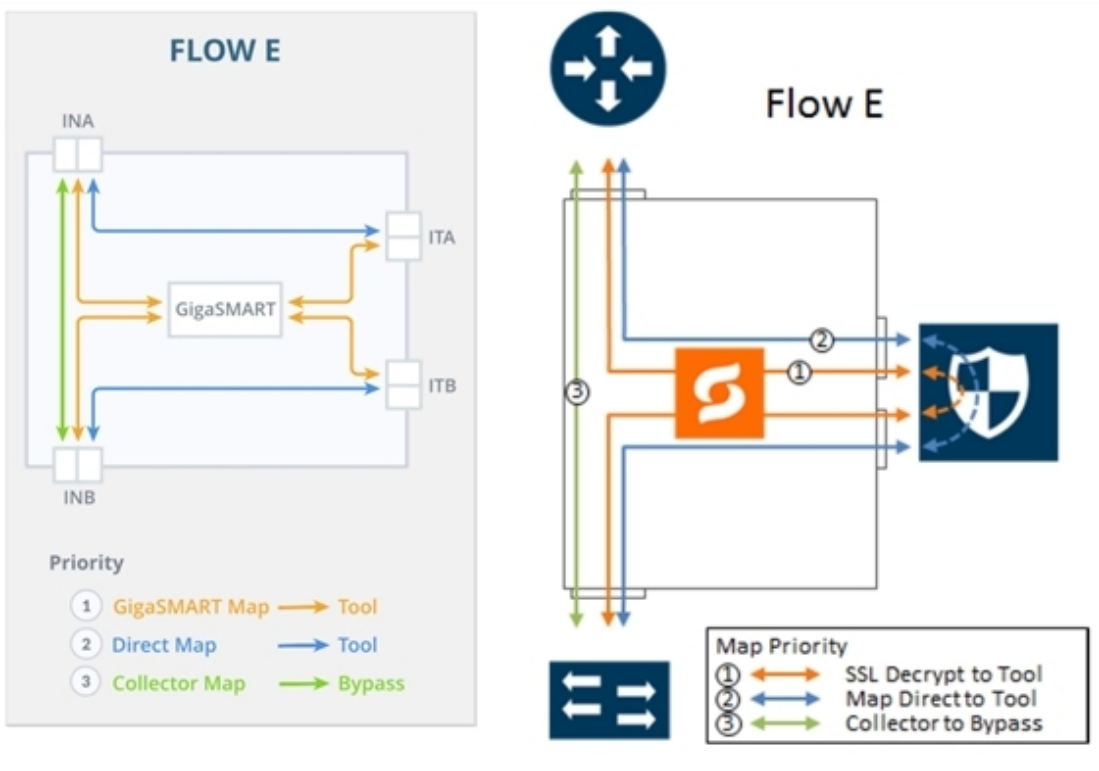

Figure 56 *FLOW E Views*

The map priorities of Flow E are as follows:

1. selectively forward traffic for decryption
2. send remaining IP traffic to tool(s)
3. direct non-IP traffic to a collector, which sends traffic to bypass

Flow F

Flow F is for the following use case:

- filter TCP traffic and direct it to the GigaSMART engine for inspection before forwarding it to the tool(s)
- direct all other traffic to the tool(s)

Refer to Figure 57 FLOW F Views for a larger view of Flow F on the left and a pictorial view on the right.

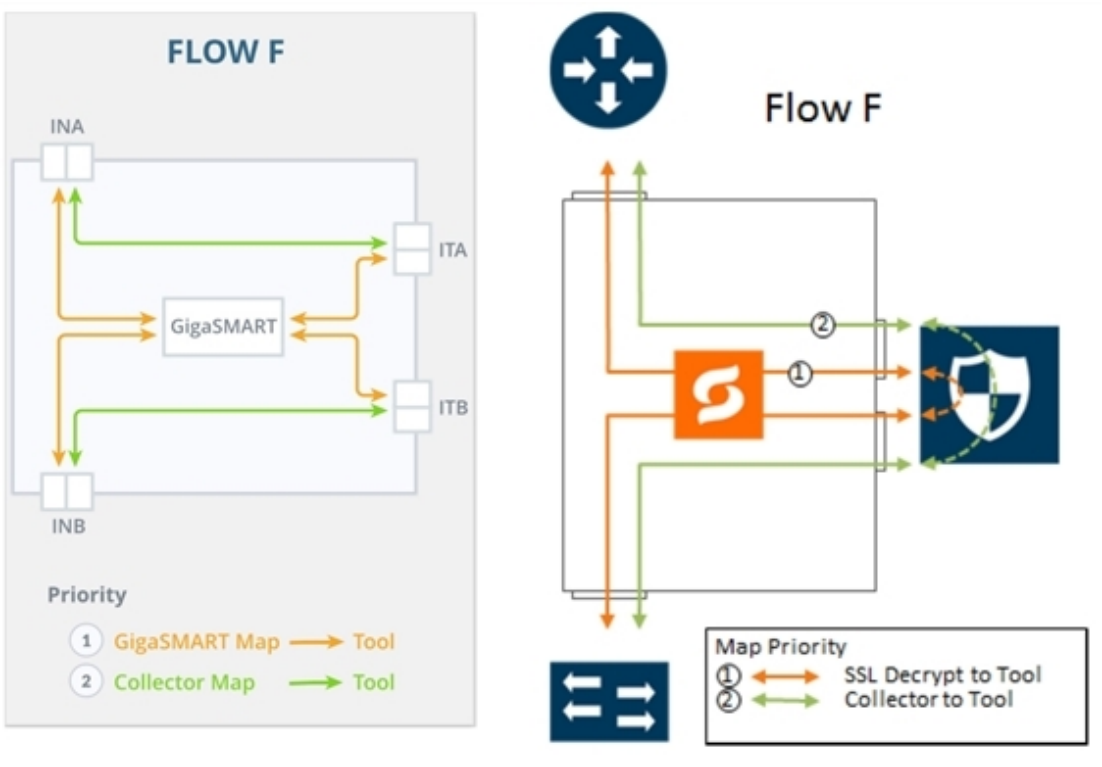

Figure 57 *FLOW F Views*

The map priorities of Flow F are as follows:

1. selectively forward traffic for decryption
2. direct unselected traffic to a collector, which sends traffic to tool

Flow G

Flow G is for the following use case:

- filter TCP traffic and direct it to the GigaSMART engine for inspection before forwarding it to the tool(s)
- bypass all other traffic

Refer to Figure 58 FLOW G Views for a larger view of Flow G on the left and a pictorial view on the right.

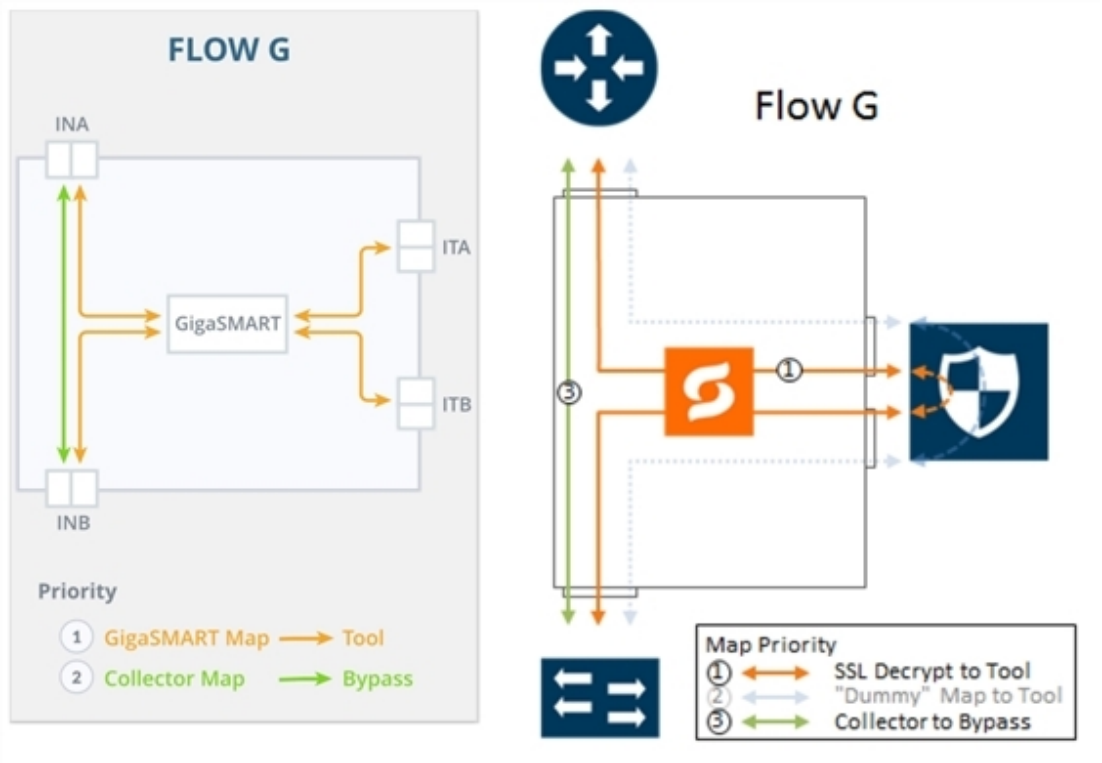

Figure 58 *FLOW G Views*

The map priorities of Flow G are as follows:

1. selectively forward traffic for decryption
2. direct unselected traffic to a collector, which sends traffic to bypass

Configure Inline TLS/SSL Decryption Using GigaVUE-FM

This section describes how to configure inline TLS/SSL decryption using GigaVUE-FM.

NOTE: Before configuring, review Get Started with Inline TLS/SSL Decryption for pre-requisites and review Introduction to Inline TLS/SSL Map Workflows.

Configure Inline TLS/SSL Decryption Using GigaVUE-FM:	Inline TLS/SSL Map Workflow Steps (for Flow B) :
• Keychain Password • Key Store • Signing CA	• Inline Network(s) • Inline Tool(s) • GS Group

• Trust Store • Policy Profile • Network Access	• Virtual Port • GS Operation • Inline Rule Based Map • Inline First Level Map • Inline Second Level Map • Collector Map (bypass)

Figure 59 *Select Inline TLS/SSL Configuration Workflow*

Workflow Overview

To configure inline TLS/SSL decryption:

1. Access the workflow
   a. Go to **Physical Nodes** and select a GigaVUE-HC1, or GigaVUE-HC3.
2. Complete the Configure Inline TLS/SSL Decryption Using GigaVUE-FM.
3. Complete the Inline TLS/SSL Map Workflow Steps (for Flow B) .
4. After completing the set-up workflows, verify Your Maps.
   a. Click **To Maps** to verify the maps created by the workflow.
   b. For inline network ports, go to **Inline Bypass > Inline Networks**. Select the inline network port and click **Edit** from the **Actions** drop-down menu. Under Configuration, select a traffic path of To Inline Tool. If using protected inline networks, disable Physical Bypass.
   c. Click **Apply**.

Inline TLS/SSL Configuration Workflow Steps

1. On the left navigation pane, click  , and then select **Physical > Nodes**.
2. In the Physical Nodes page, select the node for which you want to create the Key Store password. Go to **GigaSMART**>**Inline SSL** >**Key Store**.

3. **Set Up Keychain Password**
   a. Click **Keychain Password** from the **Actions** drop-down menu.
   b. Enter a password in the **Password** field.

      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

   c. Enable the **Auto login** check box to let GigaVUE-FM unlock the key store when the node reboots
   d. Click **Save**. The keychain password is setup.
4. **Set Up the Key Pairs**
   a. Add a New key pair:
      i. Click **New** to configure the primary signing certificate and private key. The primary CA re-signs certificates for servers that present a valid certificate.
      ii. Enter an Key Alias for the Key Pair. Click **RSA** or **ECDSA** for Key Type. Click **PEM** or **PKCS12** for Type.
      iii. (optional) For Passphrase, enter a passphrase for the key.
      iv. Select a Private Key by pasting the copied key in PEM format or installing from URL or installing from local directory.
      v. Select a Certificate by pasting the copied key in PEM format or installing from URL or installing from local directory.
      vi. Click **OK**. The Primary Key Pair is added and can be selected from the Key Store page.
   b. Still under **Key Store**, repeat the "Add a key pair" steps to configure the secondary signing certificate and private key.
5. **Set up the Signing Certificate Authority (CA)**
   a. The next step is to configure Signing CA. Click **Signing CA** from the **Inline SSL** page.
   b. Click **Add** and map each of the key pairs installed to the Primary Root CA and Secondary Root CA.
   c. Click **OK** to confirm the changes. The signing CA is configured.

6. **Set Up the Trust Store (optional)**
   a. The next step is to configure the **Trust Store**.
   b. Click **Trust Store** from the **Inline SSL** page.
   c. No configuration is required if you use the default Trust Store by selecting **Reset to Default** from the **Actions** menu.
7. **Set Up the inline SSL Policy Profile**
   a. The next step is to set up **Inline SSL Profile**.
   b. Click **SSL Profiles** from the **Inline SSL** page.
   c. Click **New** to configure an inline TLS/SSL profile. The profile specifies policy configuration, such as certificate handling and actions to take for the profile.
   d. Enter an alias for the profile. Under Policy Configuration, select a Default Action of Decrypt. Under Security Exceptions, select Decrypt or Drop.
   e. Under No-decrypt list/Decrypt list, enable the No-decrypt list and Decrypt list checkbox. Enter the paths to the files.
   f. Under Policy Rules, click **Add a Rule**.
      i. Select **Category** from the Rule drop-down (under condition) menu. Select financial_services from the Category drop-down menu. Add another rule.
      ii. Select **Category** from the Rule drop-down menu. Select health_and_medicine from the Category drop-down menu. Add another rule.
      iii. Select **Domain** from the Rule drop-down menu. Enter youtube.com in **Value** text box for the Domain. Click **Apply**.
      iv. The inline TLS/SSL profile is added. For details about the Inline TLS/SSL Policy Profile fields and their descriptions, refer to Inline TLS/SSL Policy Profile—Field References.
8. **Set Up Network Access**
   a. The next step is to set up **Network Access**.
   b. Click **Network Access** from the **Inline SSL** page.
   c. Select **DHCP** and **DHCP Enabled** for a specified GigaSMART module.
   d. Click **OK**. The network access is configured.

> **NOTE:** To verify the Network Access configuration, you can do the Ping Test as follows:
>
> a. Navigate to **GigaSMART > Inline SSL > Network Access**.
> b. In the Inline SSL Network Access window, on the Ping Test section, enter or select the GigaSMART Port and IP Address/Host Name.
> c. Click Ping to run the ping test. The test results appears in the Ping Result box.

Next, complete the Inline TLS/SSL Map Workflow Steps (for Flow B)

Inline TLS/SSL Policy Profile—Field References

The following table lists and describes the attributes that define the Inline TLS/SSL Policy Profile.

Field	Description
Alias	Enter a unique name for the flexible inline SSL APP.
Policy Configuration	
SSL Monitor Mode	Select an SSL Monitor Mode from one of the following options: ▪ Enable—When the monitor mode is enabled, the SSL decryption or encryption is off. The monitor application collects information such as the TCP ports that are in use and VLAN information about the incoming traffic, and forwards the packets to the tool port or network port based on the non-SSL TCP bypass action. ▪ Disable—This is the default value. When the monitor mode is disabled, the SSL decryption or encryption is on. Use this mode during the deployment stage. ▪ Inline—Both monitor mode and SSL decryption or encryption is on. Use this mode to debug issues. Refer to Inline TLS/SSL Monitor Mode for details.
Default Action	Select one of the following option: ▪ Decrypt—Decrypt all the traffic that is guided into the Inline SSL APP. ▪ No Decrypt—Do not decrypt the traffic that is guided into the Inline SSL APP.
URL Cache Miss Action	Select one of the following options: ▪ Decrypt—Decrypt all the traffic that is guided into the Inline SSL APP. ▪ No Decrypt—Do not decrypt the traffic that is guided into the Inline SSL APP. ▪ Defer—Delay the decryption until the Defer Timeout seconds provided.
Tool Fail Action	The failover action taken in response to a failure of an inline tool. Select one of the following options: ▪ Bypass Tool—The traffic bypasses the failed inline tool. ▪ Drop Connection—The traffic is dropped.
Decrypt Tool bypass	Bypasses the decrypted TLS/SSL traffic.
No Decrypt Tool bypass	Bypasses the non-decrypted TLS/SSL traffic.
Non-SSL TCP Traffic Tool bypass	Bypasses the non-TLS/SSL, that is the TCP intercepted traffic.
High Availability	Select the check box to detect the link switchover by upstream device that is in active or standby mode.

Field	Description
	NOTE: Do not select this check box if the inline network links are in active state. Refer to High Availability Active Standby for details.
Network Group Multiple Entry	Select this check box to allow the traffic from different inline network to reenter GigaSMART. Refer to Inline Network Group Multiple Entry for details.
Tool Early Engage	Select this check box to allow the inline tools to change the MAC address or VLAN IDs. When a connection request is received from the client, GigaSMART establishes the connection with the inline tool first, before connecting with the server. This helps the inline tools to modify the MAC address or VLAN IDs when sending the traffic back to the server. Refer to Tool Early Engage for additional information and limitations.
One-Arm Mode	Select this check box to have both the client and server traffic travel through the same physical link or logical aggregate port channel. Refer to Tool Early Engage for additional information and limitations. **NOTE:** You can enable the **One-Arm mode** only if you have enabled the **Tool Early Engage** option.
StartTLS Port	Enter the required SSL/TLS ports. Refer to StartTLS and HTTP CONNECT for details.
Security Exceptions	You can choose to either decrypt or drop the traffic for the following certificates: ▪ Self-signed certificate ▪ Unknown CA certificate ▪ Invalid certificate ▪ Expired certificate You can also choose to configure the security exceptions for certificate revocation validation based on OCSP or CRL on inline decryption profile. Select one of the following options: ▪ Soft Fail—If you select this option, the client browser displays the secondary MitM certificate and the inline decryption session stats in GigaVUE-FM displays as Decrypt. ▪ Hard Fail—If you select this option, the client browser displays the certificate from DigiCert and the inline decryption session stats in GigaVUE-FM displays as Bypass: Unknown Revocation. Refer to Certificate Revocation List (CRL), Online Certificate Status Protocol (OCSP), CRL and OCSP, and Checking Certificate Revocation Status for details.
No-decrypt list/Decrypt list	Select the following check boxes: ▪ No-decrypt list—Allows traffic from certain classes such as sites, domains, host-based IP address and IP subnets (decision based on LPM) to bypass decryption. ▪ Decrypt list—Allows traffic from certain sites, domains, host-based IP address and IP subnets (decision based on LPM) to always be decrypted.

Field	Description
	To update/append, download the existing whitelist/blacklist text file from GigaVUE-FMFM, add new entries, and upload. This will not remove the existing entries in the list. To replace the entire list, download the existing whitelist/blacklist text file from GigaVUE-FM, retain a copy, clear the list and upload a new whitelist/blacklist. Refer to No-decrypt Listing Policy and Decrypt Listing Policy for details.
Policy Rules	Add the required policy rules for the inline decryption profile.
TCP Port Map Decryption	The TCP destination port for decrypted traffic sent to inline tools can be configured as part of the inline decryption profile. Configure the required Priority 1 map, which is user configurable and Priority 2 map, which is the default out port. Refer to Inline TLS/SSL Decryption Port Map for details.
TCP Settings	Configure the required TCP settings as follows: • **TCP Inactive Timeout**—TCP Inactive session timeout in minutes • **TCP Delayed ACK**—GigaSMART Inline TLS/SSL decryption ACKs every TCP packet by default. If TCP Delayed ACK is enabled, then GigaSMART decryption will wait for 100ms or ACK every third packet – whichever comes first. • **TCP SYN Retries**—number of retries made by the MitM to initiate a session with the destination server. If a SYN/ACK response isn't received from the destination server on initial TCP SYN, GigaSMART attempts for additional number of TCP SYN Retries as defined by the user.
Server Key Map	A server key map binds a server with a key pair alias.
Split-Proxy Settings	
Split-Proxy	Select the check box to enable the split proxy settings for the inline decryption solution. The TLS connection between the server and client is divided into two independent connections and the security parameters are kept separate.
Non-PFS Ciphers (Server)	Select the check box to enable the non-PFS ciphers settings for the inline decryption solution that has the split proxy settings enabled. This setting is to indirectly force the server to use protocols that are lower than TLS1.3 with non-PFS ciphers. This means that the ciphers with DHE/ECDHE key-exchange will not be used on the server side.

Inline TLS/SSL Map Workflow Steps (for Flow B)

1. After completing the Configure Inline TLS/SSL Decryption Using GigaVUE-FM, go to **Workflows** and select **Inline SSL Map** from the Inline GigaSMART Operations section.
2. **Set Up Inline Networks**
   a. Select **FLOW B**.
   b. The first step in the Workflow is **Inline Network(s)**. Select a default inline network from the Inline Network(s) drop-down menu. This is a protected inline network.

3. **Set Up Inline Tools**
   a. Click **Next**. The next step in the Workflow is **Inline Tool(s)**.
   b. Click **Create Inline Tool**. Then click **Port Editor**. In the Quick Port Editor, locate ports and select **Type** of Inline Tool from the drop-down menu. Click **Enable** for those ports.
   c. Click **OK**. The inline tool port is added. Click **Close** to exit the Quick Port Editor.
   d. Still under **Inline Tool(s)**, enter an alias for the inline tool. Select Port A and Port B from the drop-down menus for the inline tool port pair. Under Configuration, ensure that **Inline tool sharing mode** is selected. Under Heartbeats, select **Enable Regular Heartbeat**.
   e. Click **OK**. The inline tool is configured.
4. **Create GigaSMART Group**
   a. Click **Next**. The next step in the Workflow is **GS Group**. Click **Create** to configure a GigaSMART group and associate it with a GigaSMART engine port. Enter an alias for the GigaSMART group and select a GigaSMART engine port from the Port List.
   b. Click **OK**. The GigaSMART group is added.
   c.
   > **NOTE:** You should always allow a maximum of 3 minutes time gap when you delete and recreate a gsgroup through GigaVUE-FM.
5. **Set Up the Virtual Port**
   a. Click **Next**. The next step in the Workflow is **Virtual Port**. Click **Create** to configure a virtual port.
   b. You cannot add multiple vports on the same gsgroup.
   c. Enter an alias for the virtual port, select the previously configured GigaSMART group, then select an Inline Failover Action.
   d. Click **OK**. The virtual port is added.
6. **Set Up the GigaSMART Operation**
   a. Click **Next**. The next step in the Workflow is **GS Operation**. Click **Create** to configure a GigaSMART operation.
   b. Enter an alias for the GigaSMART operation, select the previously configured GigaSMART group, select Inline SSL as the GigaSMART Operation (GSOP), then select the previously configured Inline SSL profile.
   c. Click **OK**. The GigaSMART operation is added.

7. **Set Up the Inline Rule-Based Map**
   a. Click **Next**. The next step in the Workflow is **Inline Rule Based Map**.
   b. Configure the inline rule-based map. This map directs traffic from the inline network to the inline tool, using a specified rule. It has the same source port as the inline first level map and the same destination port as the inline second level map. Enter an alias for the map, and select the map Type (Inline) and Subtype (By Rule), select the source inline network and the destination inline tool.
   c. Click **Add a Rule** to specify a map rule. Click **Bi-directional**, select IPv4 Protocol from the Rule drop-down menu, and select TCP from the Protocol drop-down menu. Select Port Destination from the Rule drop-down menu and enter 80 in the text box for **Min**.
   d. Click **OK**. The map is added.
8. **Set Up the Inline First-Level Map**
   a. Click **Next**. The next step in the Workflow is **Inline First Level Map**.
   b. Configure the inline first level map. This map directs TCP traffic from the inline network to a virtual port (and to GigaSMART). Enter an alias for the map, and select the map Type (Inline First Level) and Subtype (Ingress to Virtual Port). Under Map Source and Destination, select the inline network as the source and the virtual port as the destination. Under Map Rules, click **Add a Rule**. Select IPv4 Protocol from the Rule drop-down menu, and select TCP from the Protocol drop-down menu.
   c. Click **OK**. The map is added.
9. **Set Up the Inline Second-Level Map**
   a. Click **Next**. The next step in the Workflow is **Inline Second Level Map**.
   b. Configure the next map, which is the inline second level map. This map directs traffic from the virtual port, uses the inline TLS/SSL GigaSMART operation, and sends traffic to the inline tool. Enter an alias for the map and select the map Type (Inline Second Level) and Subtype (Egress from Virtual Port). Under Map Source and Destination, select the virtual port as the source and the inline tool as the destination, then select the inline SSL GigaSMART operation.
   c. Click **OK**. The map is added.
10. **Set Up the Collector Map**
    a. Click **Next**. The next step in the Workflow is **Collector Map (bypass)**.
    b. Configure a collector map for any unmatched traffic including non-TCP traffic, which is directed to bypass. Enter an alias for the map, and select the map Type (Inline) and Subtype (Collector), then select a Traffic Path of ByPass. Under Map Source and Destination, select the inline network as the source.
    c. Click **OK**. The map is added.

Inline TLS/SSL L3 Tool NAT/PAT Support

This feature focuses on a new approach in GigaSMART to offload TLS decryption from Layer 3 inline-tools performing NAT/PAT (Network Address Translation / Port Address Translation) on the traffic passing through them. The GigaSMART engine maintains two separate sessions towards the client and server-side to achieve this.

Supported Platforms

- Gen3 cards in GigaVUE-HC1 and GigaVUE-HC3
- GigaVUE-HC1-Plus

Topology

The following are the preferred topologies for connecting to L3 NAT/PAT tools. The tool must be connected to the Nb side of the inline-network pair and the network links must be connected to the Na side of the inline-network pair.

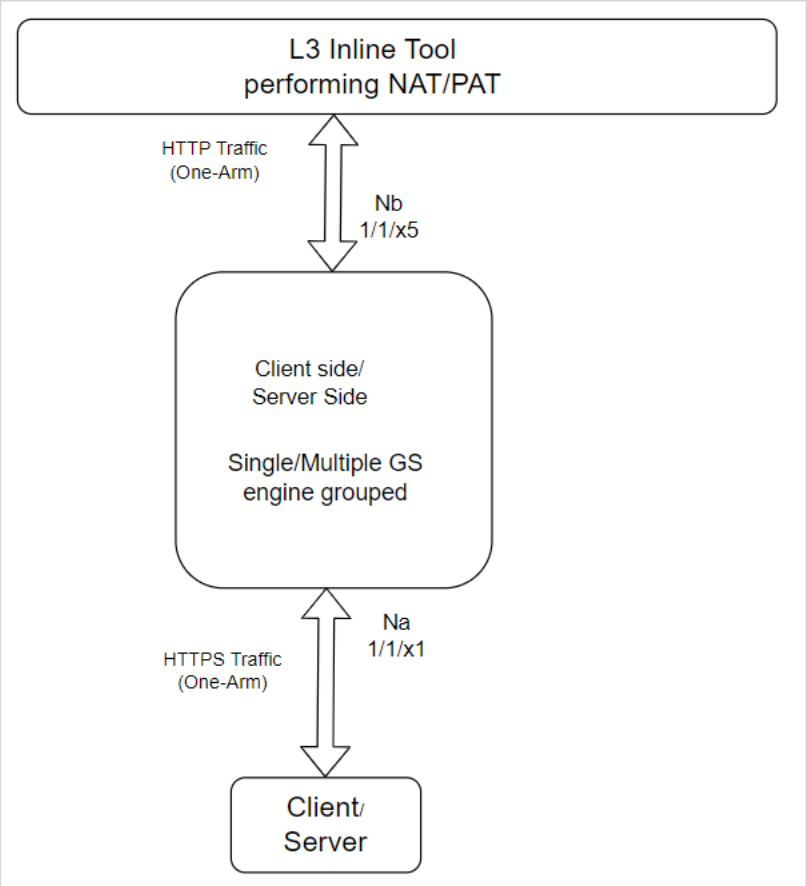

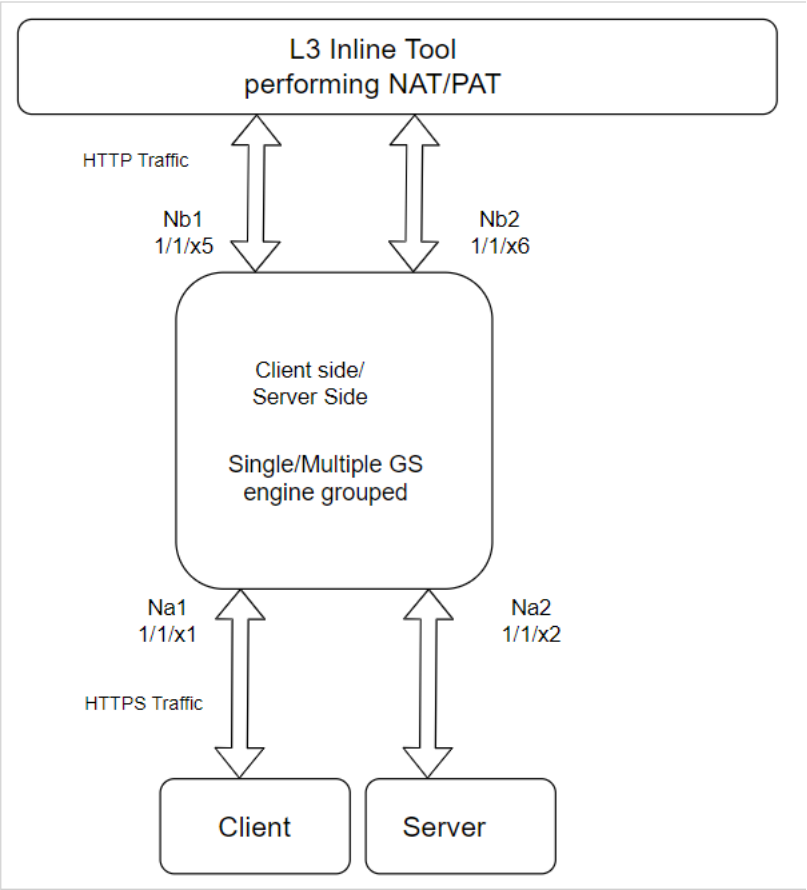

The traffic from the network is forwarded to the tool and the reverse traffic from the tool is forwarded to the server side.

Also due to separation of client and server side connections, it is possible to have the client and server connected across GS engines or devices as shown below.

However currently there is no communication mechanism between separate engines or devices. The fail-over action for these solutions must be vport-drop to avoid leaking of decrypted traffic to the network.

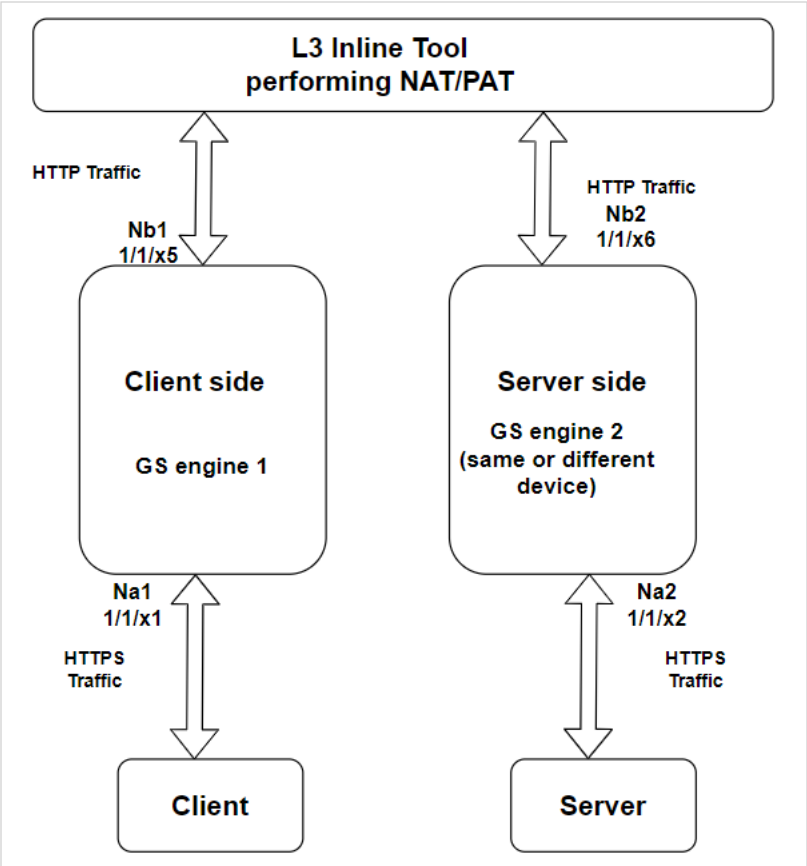
L3 Inline Tool
performing NAT/PAT
HTTP Traffic
Nb1
1/1/x5
HTTP Traffic
Nb2
1/1/x6
Client side
GS engine 1
Server side
GS engine 2
(same or different
device)
Na1
1/1/x1
Na2
1/1/x2
HTTPS
Traffic
HTTPS
Traffic
Client
Server

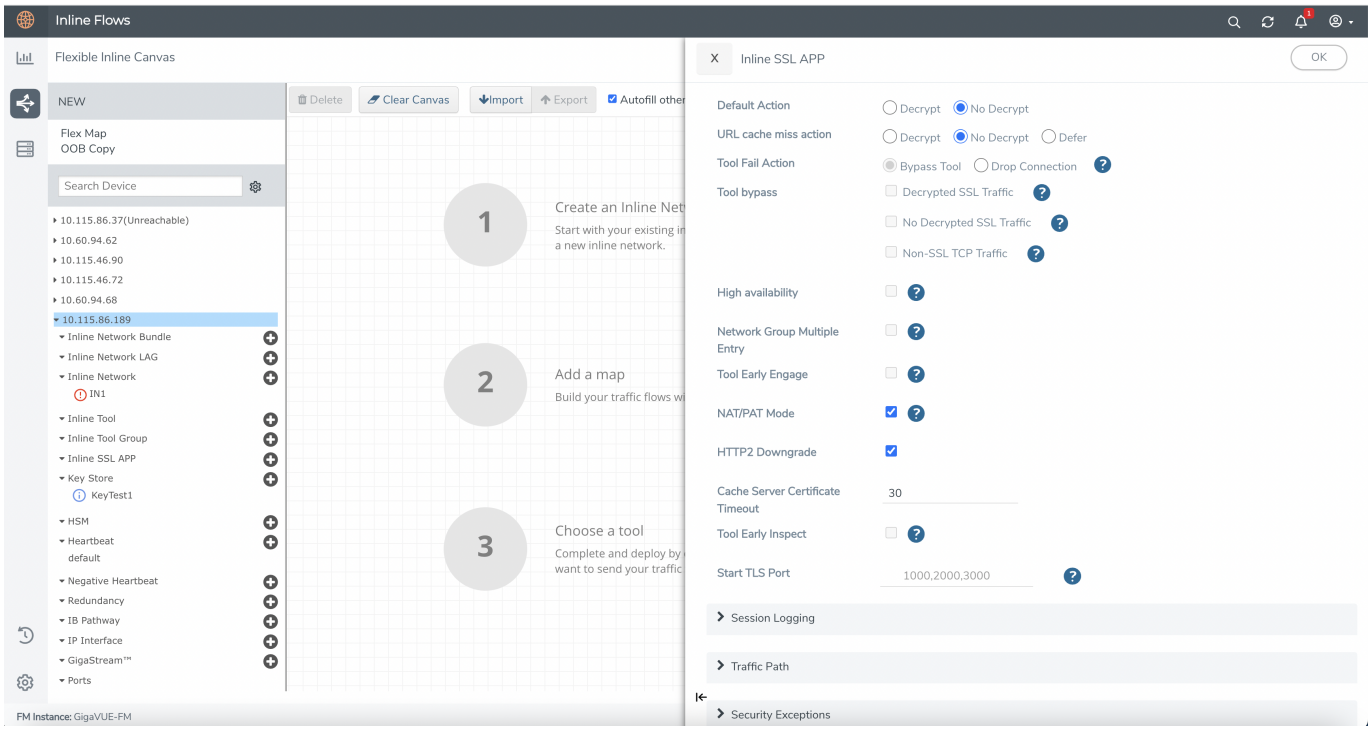
Inline Flows
Flexible Inline Canvas
NEW
Flex Map
OOB Copy
Search Device
Delete
Clear Canvas
Import
Export
Autofill other
1
Create an Inline Net
Start with your existing in
a new inline network.
2
Add a map
Build your traffic flows w
3
Choose a tool
Complete and deploy by
want to send your traffic
Inline SSL APP
OK
Default Action
Decrypt
No Decrypt
URL cache miss action
Defer
Tool Fail Action
Bypass Tool
Drop Connection
Tool bypass
Decrypted SSL Traffic
No Decrypted SSL Traffic
Non-SSL TCP Traffic
High availability
Network Group Multiple Entry
Tool Early Engage
NAT/PAT Mode
HTTP2 Downgrade
Cache Server Certificate Timeout
30
Tool Early Inspect
Start TLS Port
1000,2000,3000
Session Logging
Traffic Path
Security Exceptions
FM Instance: GigaVUE-FM

HTTP 2.0 Downgrade

In the Inline SSL APP, the **HTTP2 Downgrade** option is enabled by default, when the **Nat-Pat Mode** is enabled. The HTTP 2.0 traffic is downgraded to HTTP 1.1 and decrypted. When **HTTP2 Downgrade** option is disabled, the HTTP2 traffic is forwarded without decryption.

Cache Timeout

The server information is cached for performance optimization. The default time out is 30 minutes. The cache is flushed when the cache timeout value is set to zero. The cache is disabled when the timeout value is set to zero.

Refer to the following Gigamon Validated Design for more information:

- Offloading TLS Decryption for an One-Armed Inline Tool in L3 with NAT/PAT Mode
- Enabling GigaSECURE TLS Decryption to Offload SSL Inspection from Next-Generation Firewall

Limitations

- The decrypted data to inline tool is limited to HTTP/1.1 protocol over TLS, inline tool will only see encrypted data on other application protocols.
- The StartTLS traffic will not be decrypted as it is non HTTP traffic.
- It does not support bypass tool since all the packets from client or server need to pass to inline tool for NAT/PAT.
- This feature cannot co-exist with features such as Network group multiple entries, Inline network high availability, RIA, Tool early engage, Tool early inspect, one-arm, and Destination port translation.
- The IPv6 version is not supported in software release version 6.1.00. The IPv6 version is supported from software release version 6.2.00.

View Statistics

You can view the following inline TLS/SSL decryption statistics:

Topics:

- Inline TLS/SSL Session Statistics
- Monitor Statistics
- Certificate Statistics

Inline TLS/SSL Session Statistics

To display the inline TLS/SSL summary details, go to **GigaSMART > Inline SSL > Session Statistics**, view the Summary details under **Summary** tab which is displayed initially. Refer to Figure 60Inline TLS/SSL Session Statistics in GigaVUE-FM.

There are four sections: Session Statistics, Performance Statistics, Policy Statistics, and Certificate Statistics. Click **Show Summary** to view these sections. Click **Clear Session Summary** to clear all the displayed summary details.

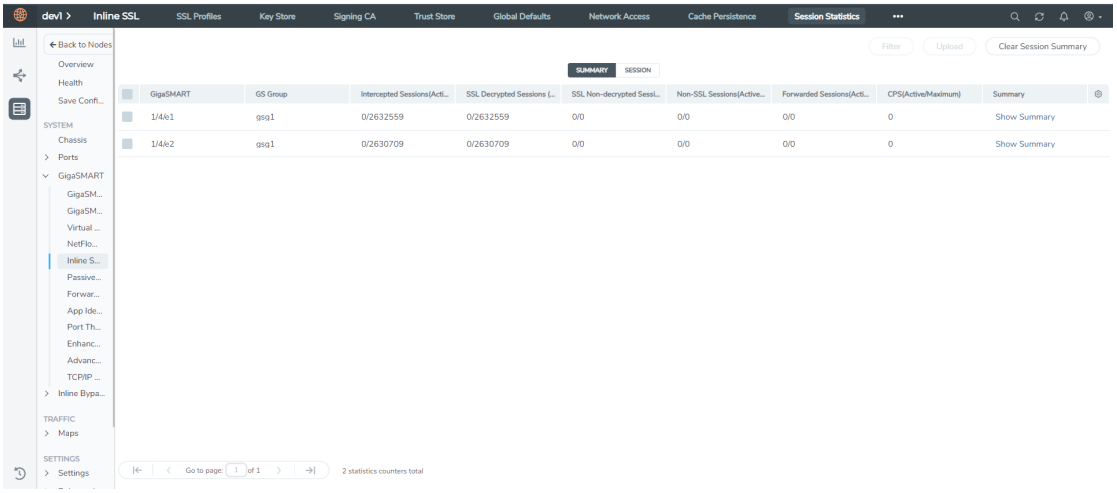

Figure 60 *Inline TLS/SSL Session Statistics in GigaVUE-FM*

To view the inline TLS/SSL session details, go to **GigaSMART > Inline SSL > Session Statistics**. Click on the **Sessions** tab. The list of available sessions will be displayed. To search and filter the session details, click **Filter** and enter an IPv4 source or destination, an L4 port source or destination, or a host name.

Monitor Statistics

To display monitor statistics, go to **GigaSMART > Inline SSL > Monitor Statistics**. Select the required GigaSMART engine from the drop-down menu.

There are three sections. The first section, which has a graph for INTERFACE TRAFFIC and Interface Packet statistics, is displayed initially. To return to this display, click the small graph, TOTAL INCOMING PACKETS. Refer to Figure 61 Inline TLS/SSL Session Monitor—Interface Packet Statistics.

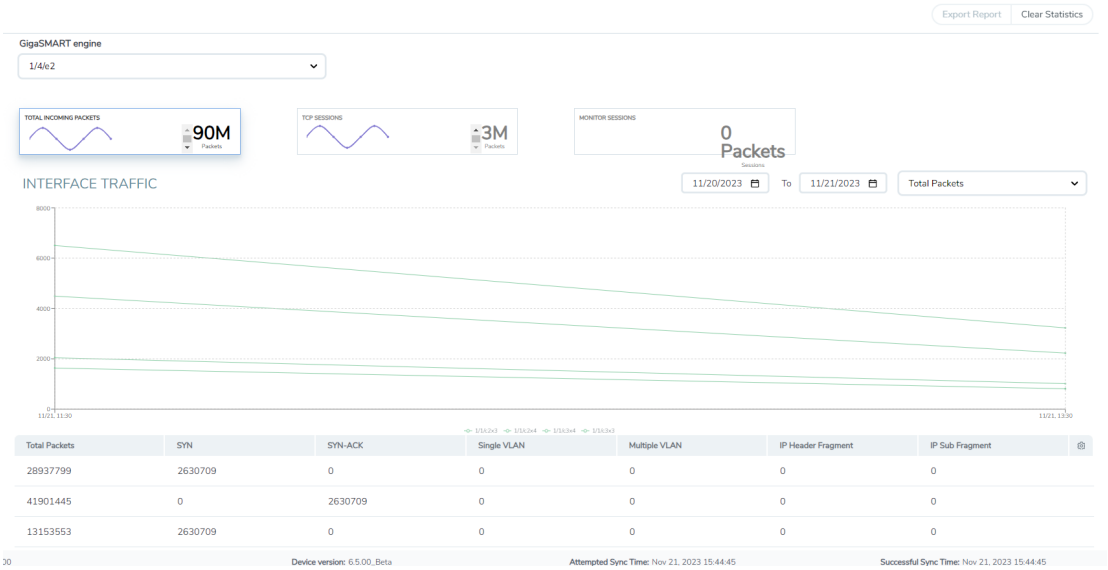

Figure 61 *Inline TLS/SSL Session Monitor—Interface Packet Statistics*

To display the graph and statistics for TCP Sessions, click the small graph, TCP SESSIONS. Refer to Figure 62 Inline TLS/SSL Session Monitor—TCP Sessions.

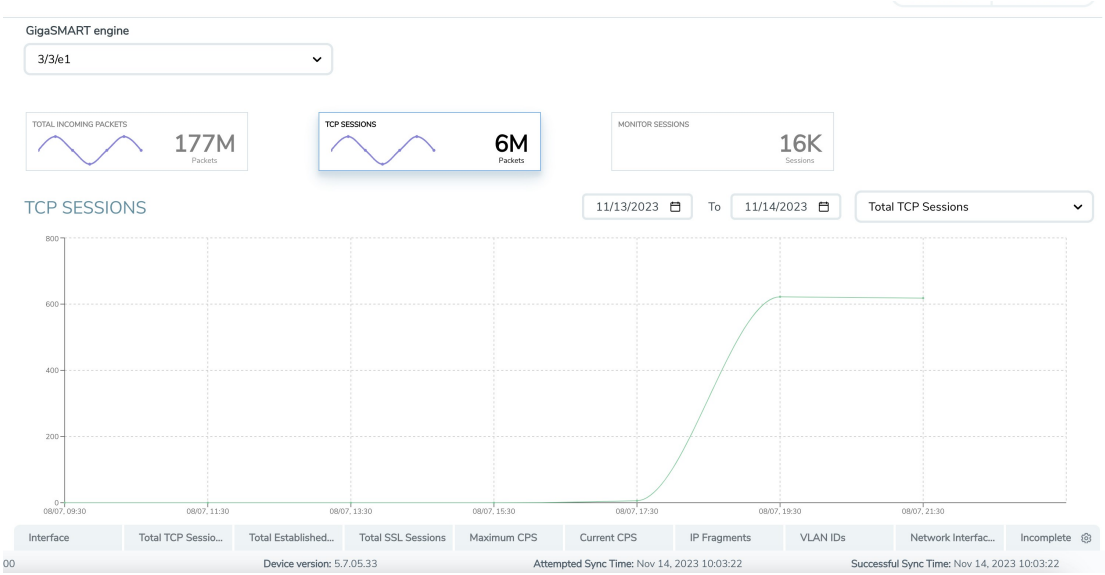

Figure 62 *Inline TLS/SSL Session Monitor—TCP Sessions*

To display Monitor Sessions, click the small graph, MONITOR SESSIONS. Refer to Figure 63 Inline TLS/SSL Session Monitor—Monitor Sessions.

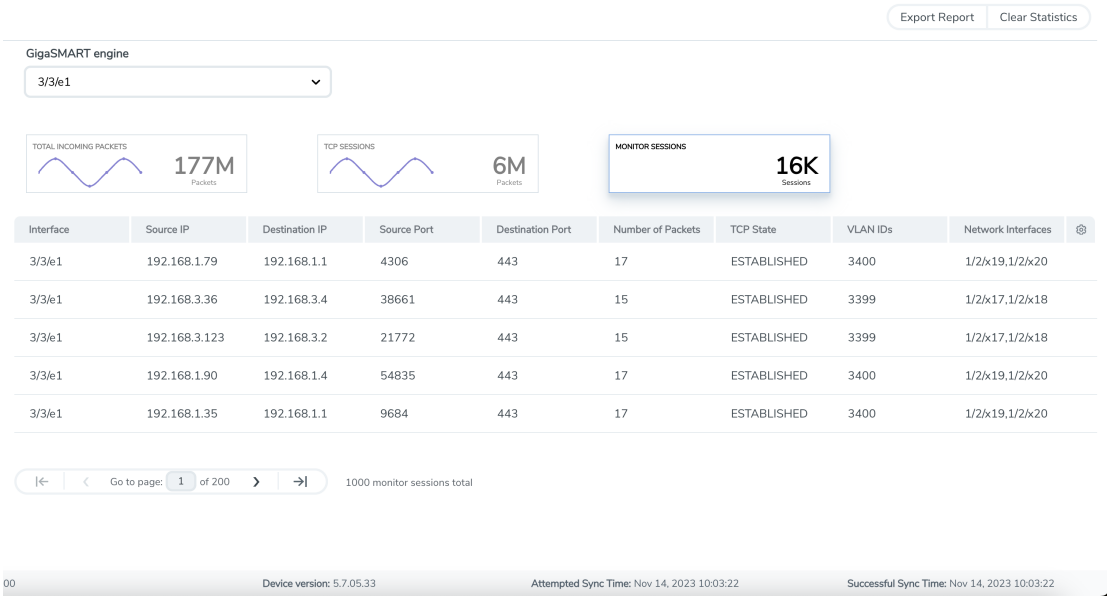

Figure 63 *Inline TLS/SSL Session Monitor—Monitor Sessions*

Click **Clear Statistics** to clear all the displayed statistics details.

Certificate Statistics

To view certificate statistics, go to **GigaSMART > Inline SSL > Certificate Statistics**. The page displays the hit count of each key store certificate and allows to track whether the certificate is actively used or not. The hit count is numbered only if the policy is set for decryption. If the policy is set to no-decryption or if the deployment is outbound, the hit count will not be considered.

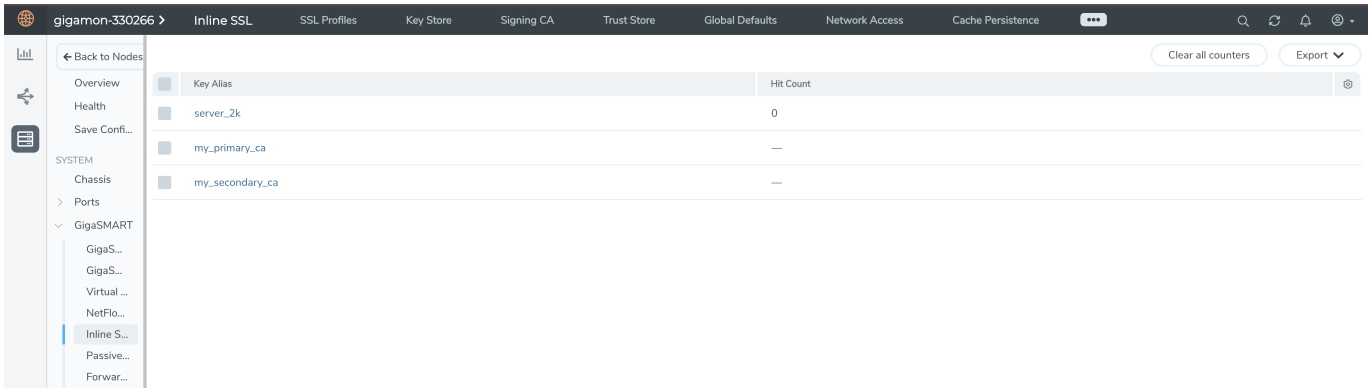

Figure 64 *Inline SSL - Certificate Statistics*

Click **Clear all Counters** to clear the hit counter of all the certificates. Click **Export** to export the available hit count details.

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 are the basic dashboards and its visualizations:

Table 12: Session Overall Dashboard

Dashboard	Description	Visualizations	Details
Session Overall	Displays visualizations on the overall details of encrypted traffic.	***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	Provides an insight into the TLS version

Dashboard	Description	Visualizations	Details
		Trend	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.

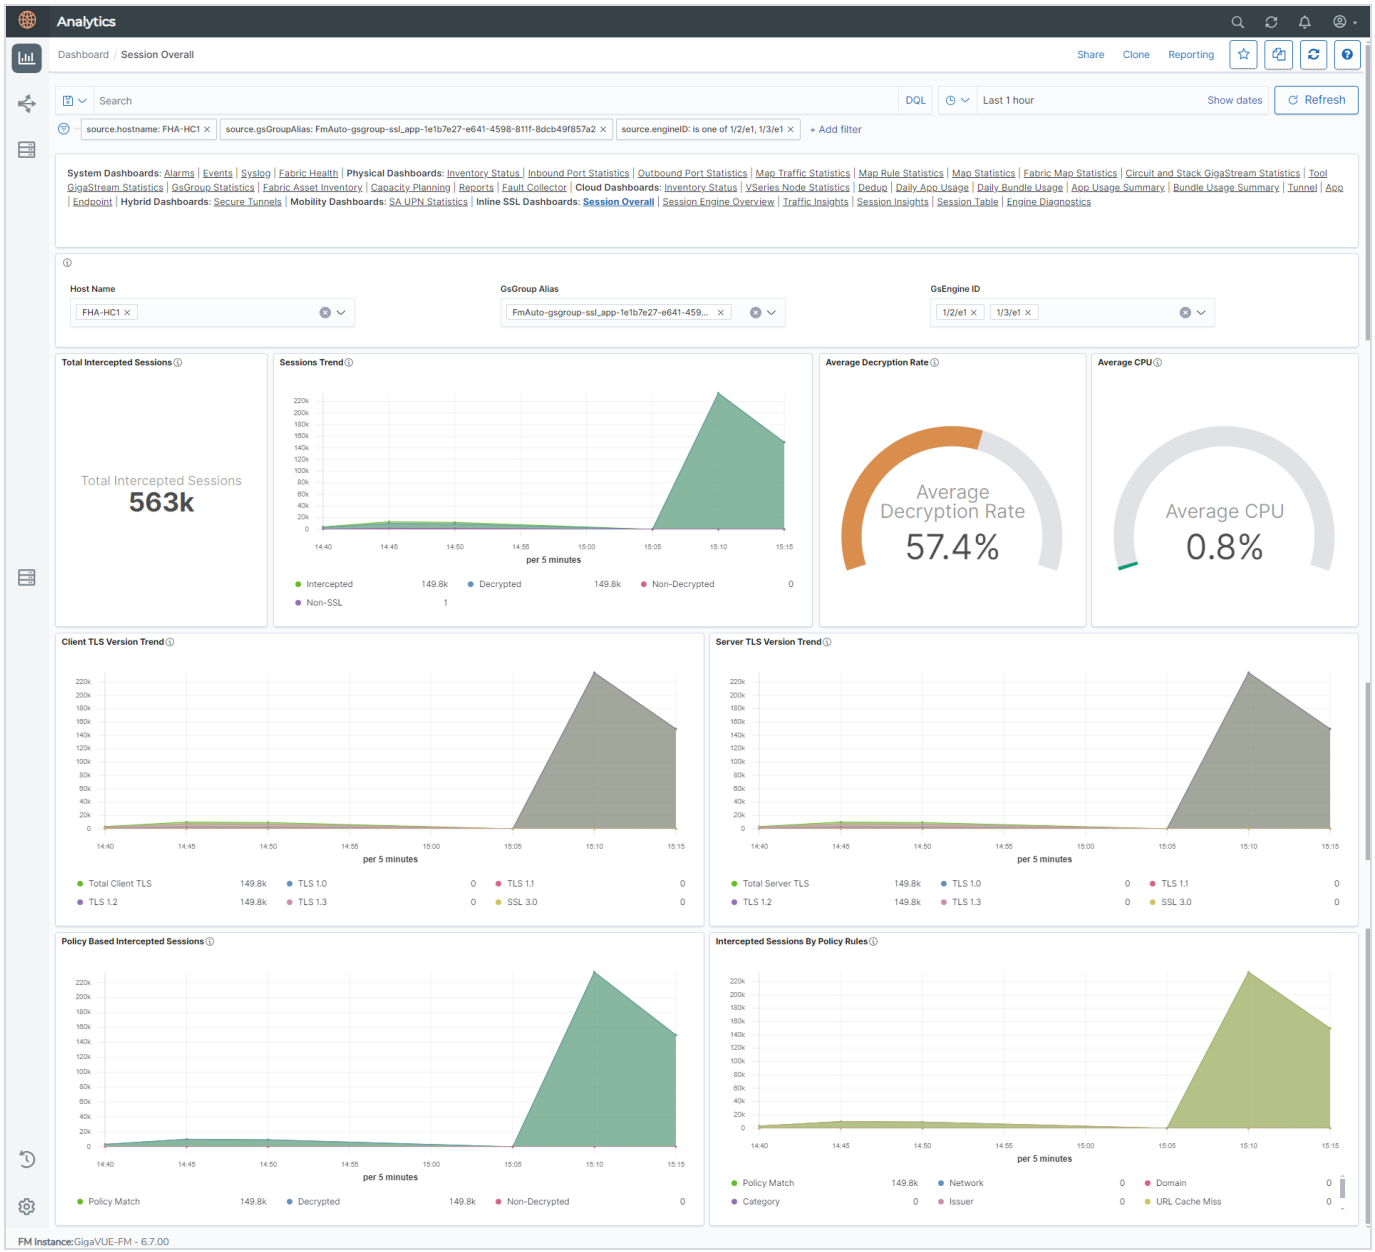

Table 13: Session Engine Overview Dashboard

Dashboard	Description	Visualizations	Details
Session Engine Overview	Displays visualizations related to Inline TLS/SSL Sessions per Engine	***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.
		Average CPS per Engine	Displays the average Connections per Second

Dashboard	Description	Visualizations	Details
			(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)

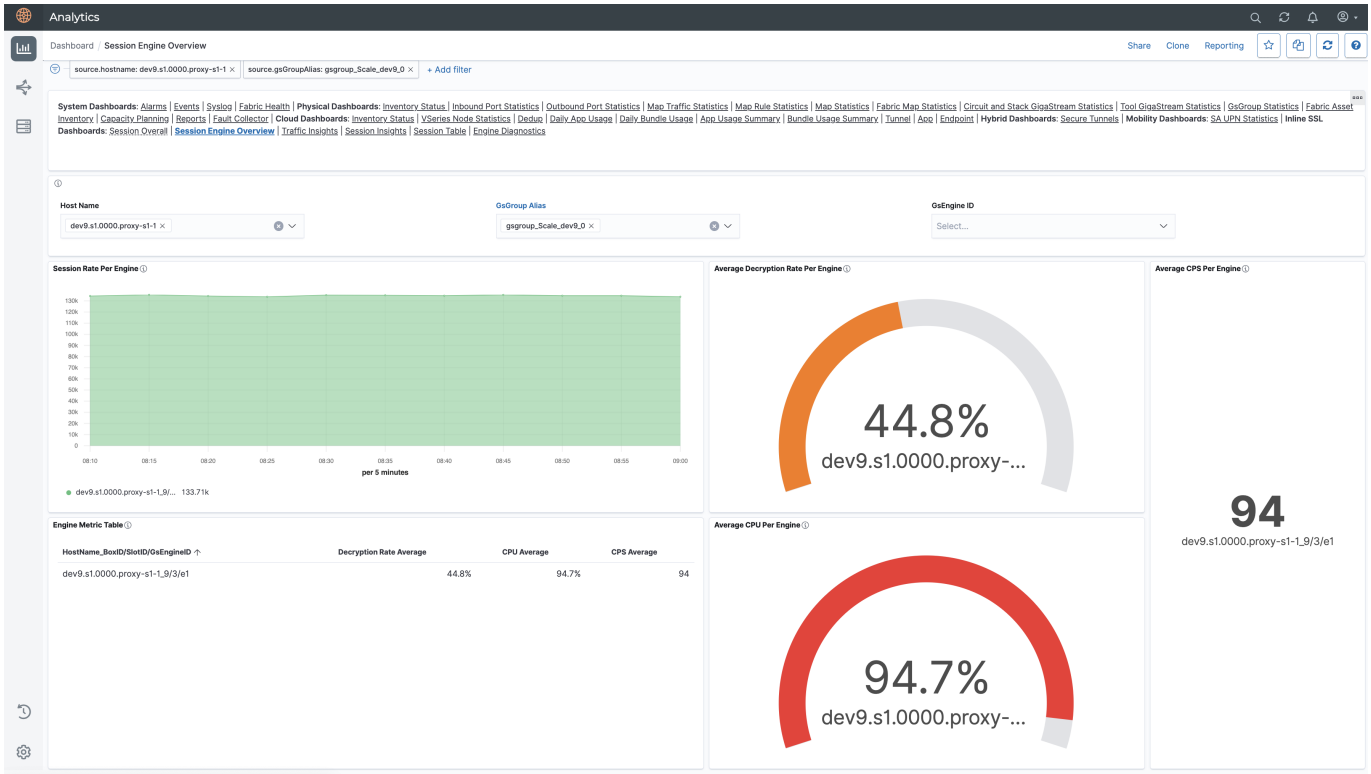

Table 14: Traffic Insights Dashboard

Dashboard	Description	Visualizations	Details
Traffic Insights	Displays visualizations related to the traffic that is handled with a Inline TLS/SSL sessions	***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 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.

Dashboard	Description	Visualizations	Details
		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.

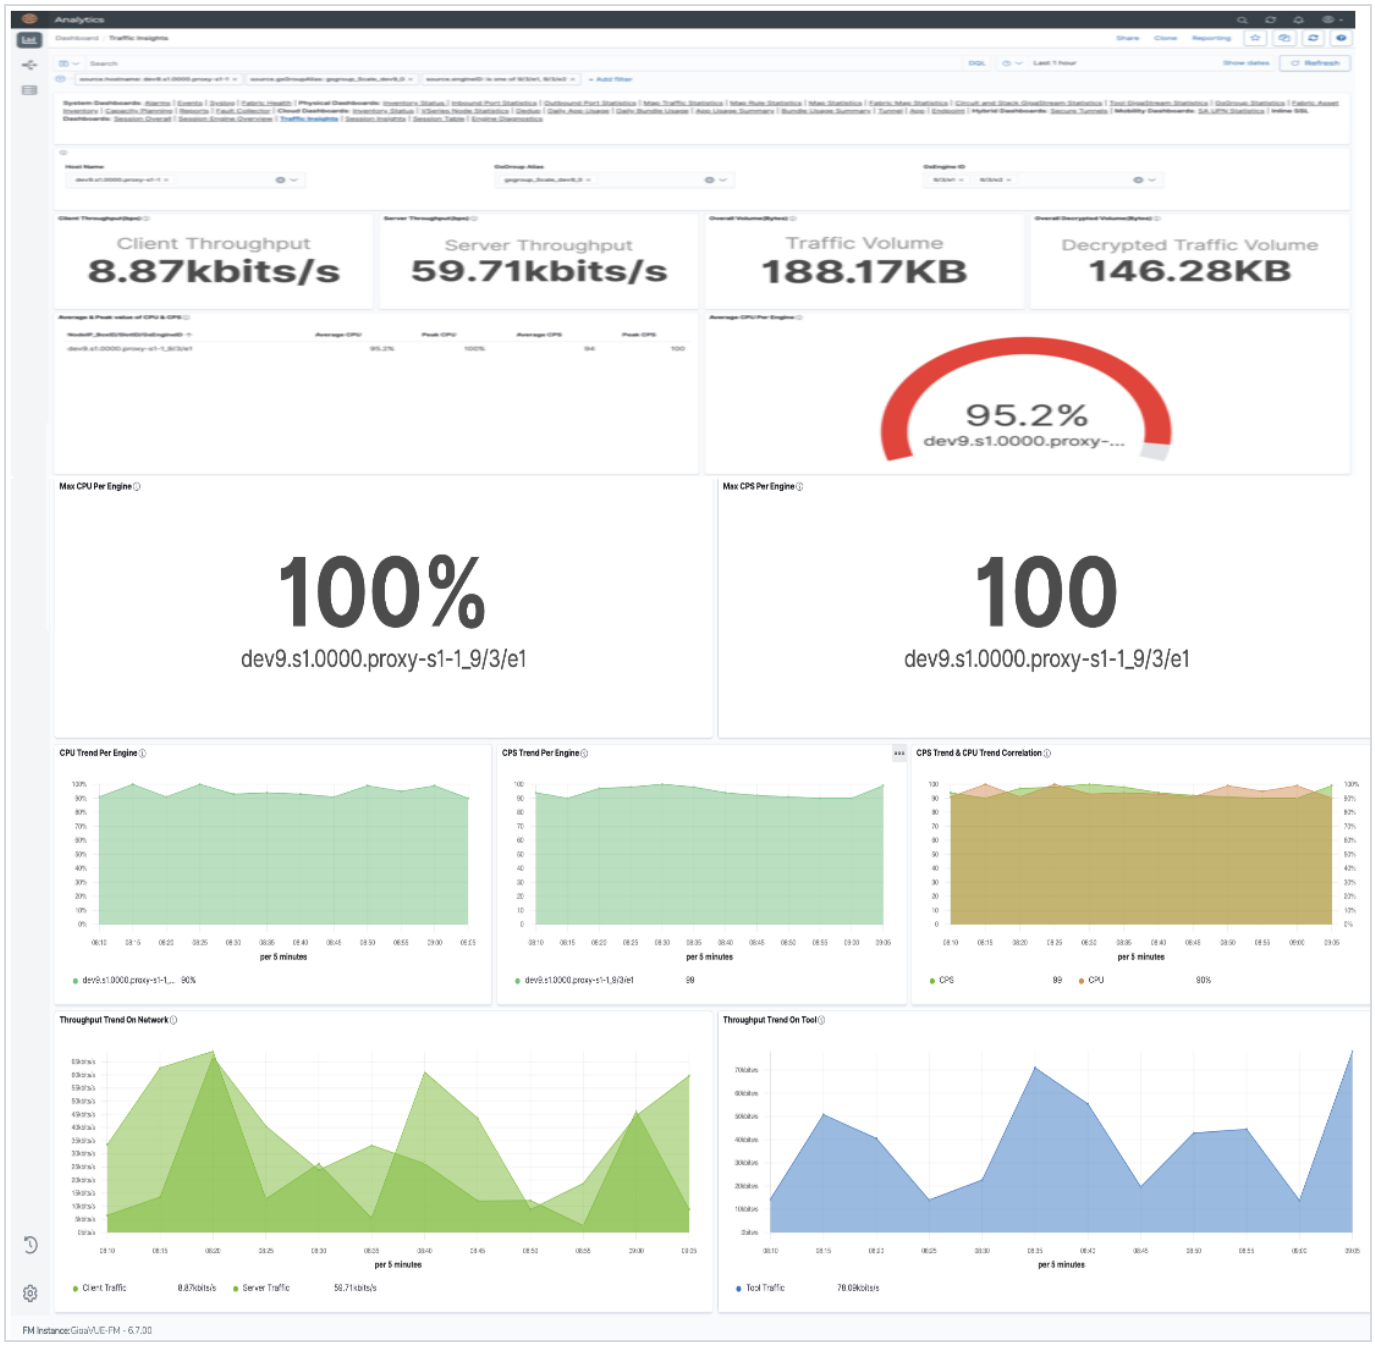
Analytics
Client Throughput
8.87kbits/s
Server Throughput
59.71kbits/s
Traffic Volume
188.17KB
Decrypted Traffic Volume
146.28KB
95.2%
dev9.s1.0000.proxy-...
Max CPU Per Engine
100%
dev9.s1.0000.proxy-s1-1_9/3/e1
Max CPS Per Engine
100
dev9.s1.0000.proxy-s1-1_9/3/e1
CPU Trend Per Engine
CPS Trend Per Engine
CPS Trend & CPU Trend Correlation
per 5 minutes
CPS
CPU
Throughput Trend On Network
Throughput Trend On Tool
Client Traffic
Server Traffic
Tool Traffic
FM Instance:GigaVUE-FM - 6.7.00

Table 15: Engine Diagnostics Dashboard

Dashboard	Description	Visualizations	Details
Engine Diagnostics	Displays the certificates and SSL alerts related to a GigaSMART engine	***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.

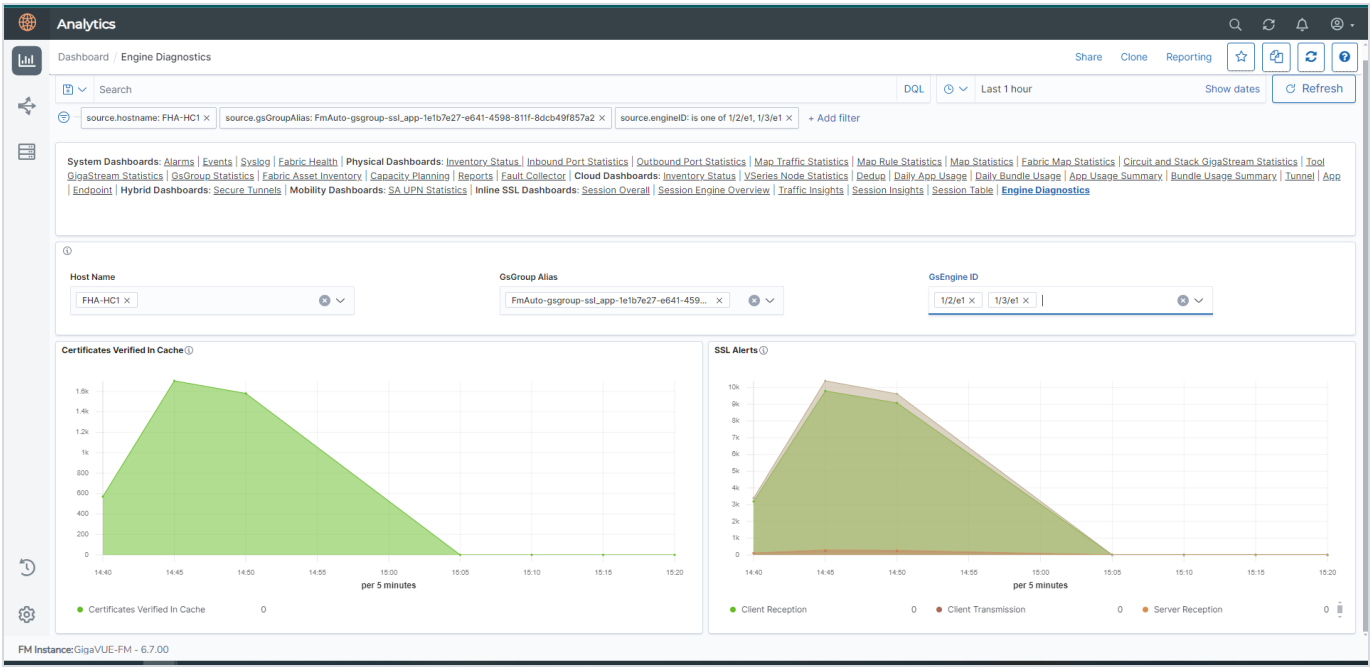

Advanced Dashboards

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 eight 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 are the advanced dashboards and its visualizations:

Table 16: Session Insight Dashboard

Dashboard	Description	Visualizations	Details
Session Insights	Displays visualizations on the details of an Inline TLS/SSL session.	***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

Dashboard	Description	Visualizations	Details
			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.

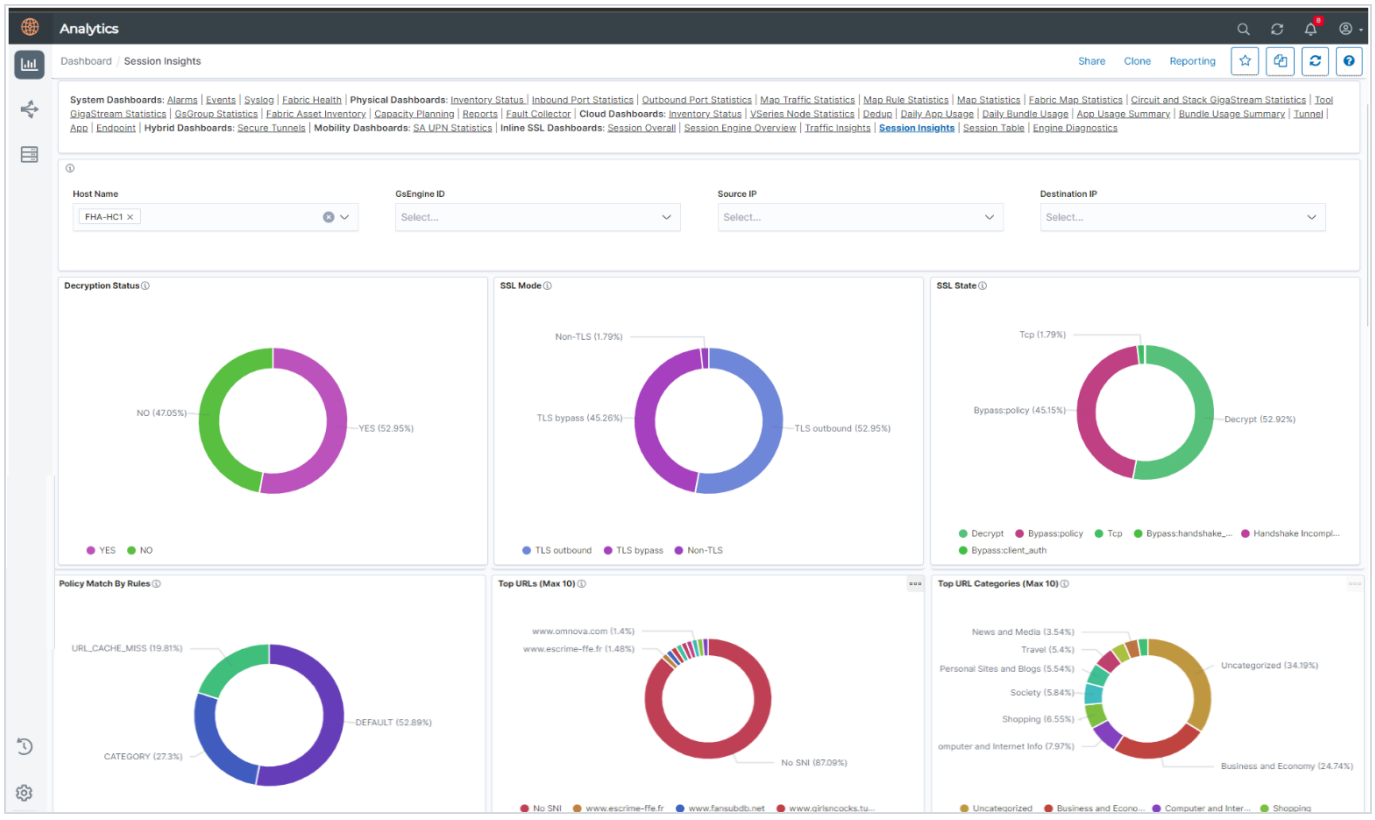

Table 17: Session Table Dashboard.

Dashboard	Description	Visualizations	Dashboard
Session Table	Displays visualizations related to Inline TLS/SSL Sessions per Engine in a tabular format.	***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

Dashboard	Description	Visualizations	Dashboard
			option by using (+) (−)l button.

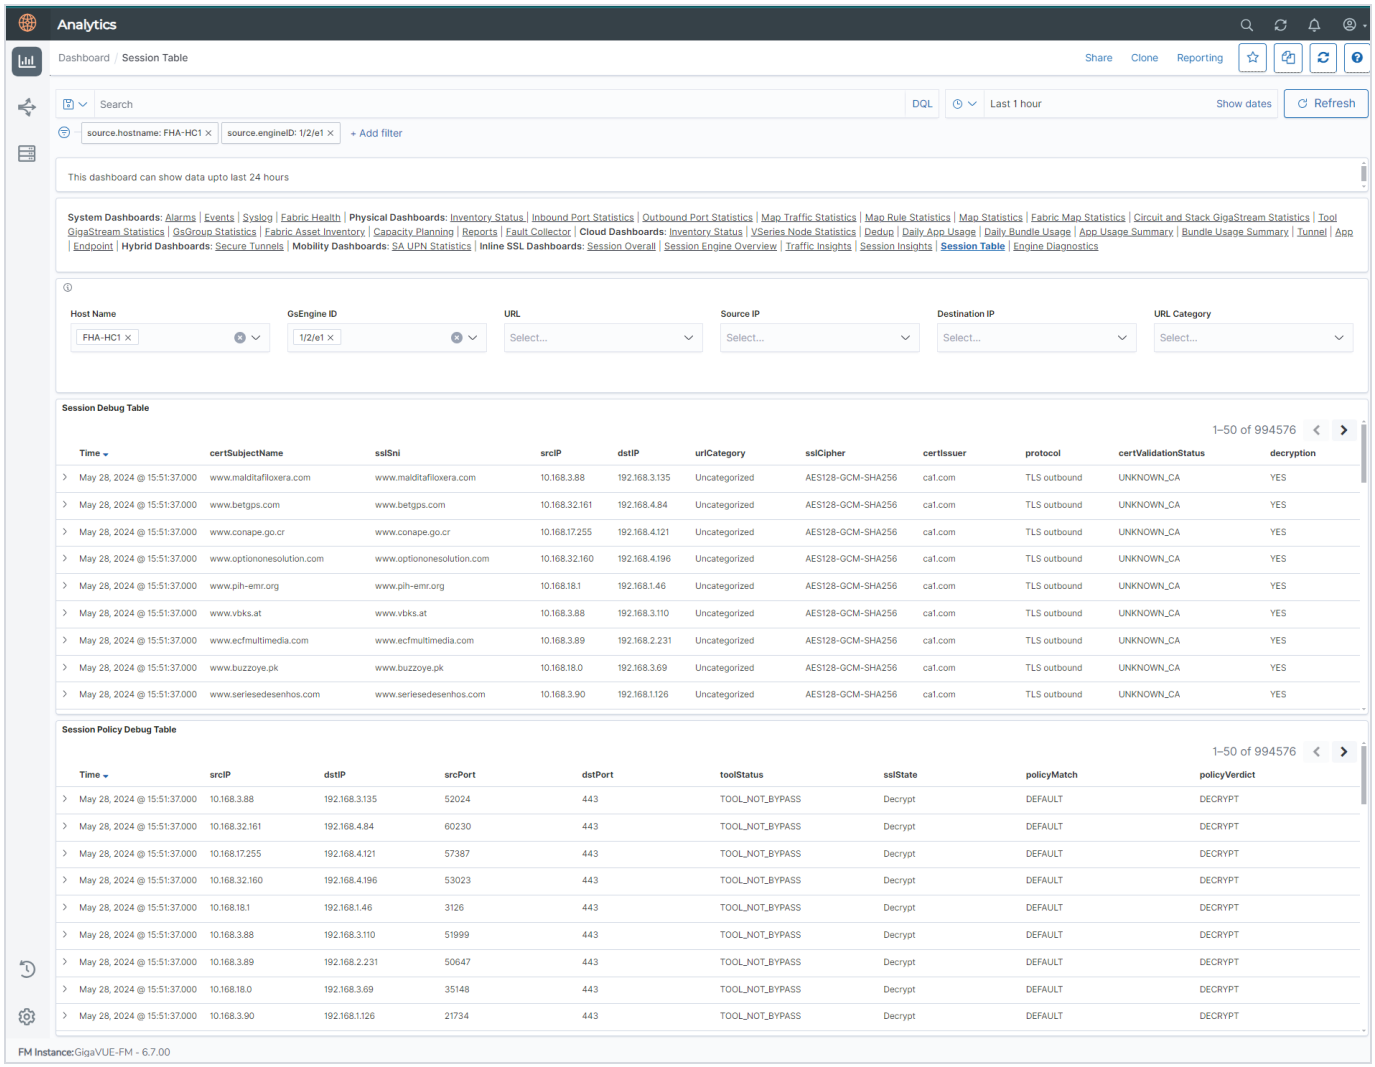

Resilient Inline Arrangement with GigaSMART Flex Inline Solution

GigaSMART Flex Inline Solution can now be configured in a Resilient Inline Arrangement (RIA). It is supported on all GigaVUE HC Series devices. To learn more about Resilient Inline Arrangements refer to Configure Resilient Inline Arrangement.

A resilient inline arrangement uses two nodes for traffic management of dual-path high availability environments. The nodes will process traffic at the same time using source and destination IP. For traffic received from the top network interfaces are decided based on the source IP, whereas the traffic received from bottom network interfaces use the destination IP of incoming traffic. If the IP address end with an even number, then the traffic will be forwarded to one node whereas if it is an odd number then it will be sent to another node.

For example:

In the below Resilient Inline Arrangement , interface (I/F) 1 and I/F 2 are connected to node 1, I/F 3 and I/F 4 are connected to node 2. For connection 1 client 1, traffic from I/F 1 with VLAN A source, IP A will be forwarded to GS1 in node 1 since last decimal digit of IP A is even. The Traffic of server 1 of the same connection from I/F 2 with VLAN A destination IP A will also be forwarded to GS1 in node 1. Since last digit of IP C is even, connection 2 traffic from I/F 3 and I/F 4 will also be forwarded to GS1 in node 1. Connection 2 traffic between GS1 and client 2 will be on I/F 3 with VLAN B, these packets need to pass through IBP connected between node 1 and node 2. Similarly, traffic between GS1 and server 2 will be on I/F 4 after passing through IBP.

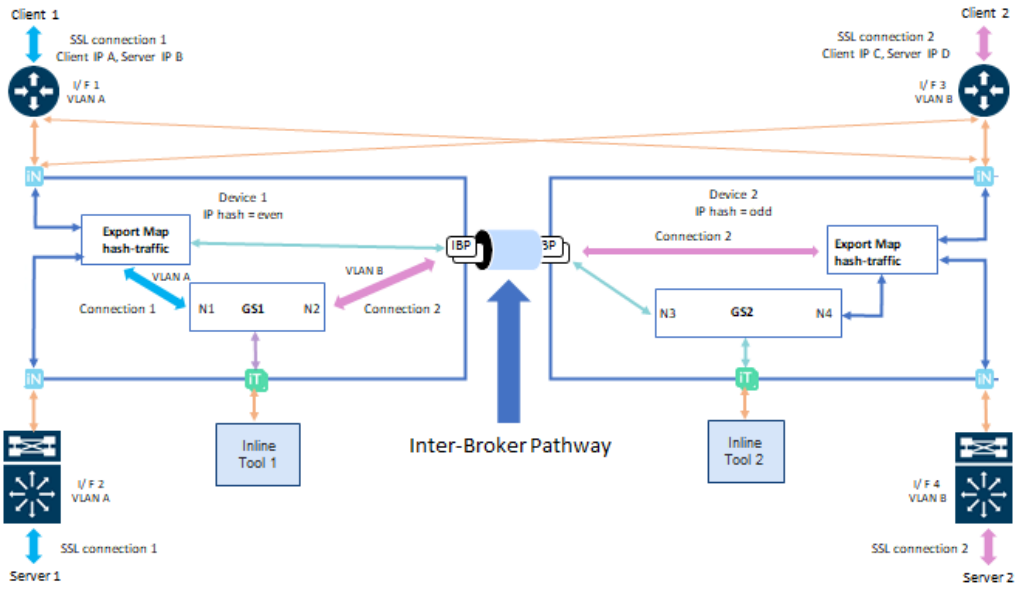

Symmetric Traffic in RIA

In a Symmetric connection the SYN packet from client 1 is received on I/F 1 with VLAN A. This incoming SYN packet is forwarded to GS1 in node 1 since source IP (IP A) is even. GS1 will initiate TCP connection to server 1 by sending SYN packet out from I/F 2 with VLAN A as

shown below.

Server 1 responds SYN ACK from I/F 2 with VLAN A and the packet is forwarded to GS1 since destination IP (IP A) is even. All inbound and outbound traffic between GS1 and client 1 will be on I/F 1 and all traffic between GS1 and server 1 will be on I/F 2 attached to node 1. Similarly, traffic from connection 2 on I/F 3 and I/F 4 with VLAN B are processed by GS1 on node 1 since client 2 IP (IP C) is even. All traffic between GS1 and client 2 will be on I/F 3 with VLAN B through IBP and all traffic between GS1 and server 2 will be on I/F 4 with VLAN B attached to node 2.

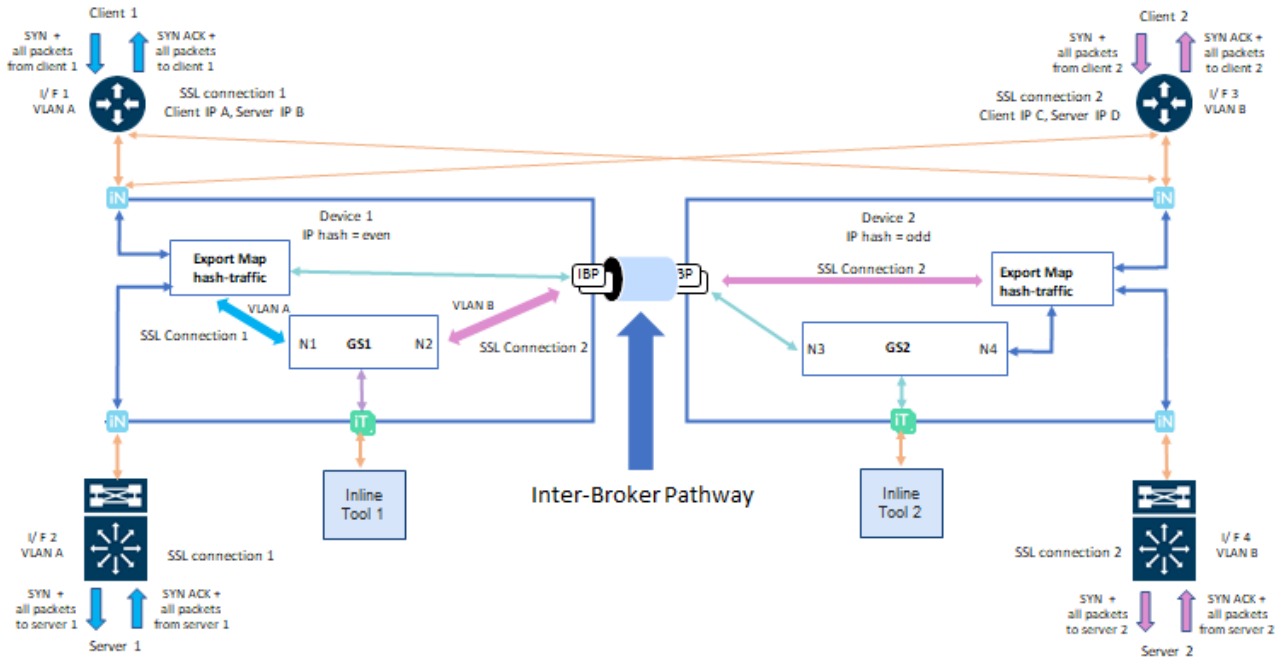

Asymmetric Traffic in RIA

An asymmetric connection can occur in the following scenarios:

1. Let's consider the below resilient inline arrangement ,the SYN packet of TLS/SSL connection 3 from client 3 is received on I/F 1 with VLAN A. This incoming SYN packet is forwarded to GS1 in device 1 since source IP (IP A) is even. Server 3 responds SYN ACK from I/F 2 with VLAN A on device 1 and the packet is forwarded to GS1 since destination IP (IP A) is even. After that, GS1 will send SYN ACK to client 3 on I/F 1 with VLAN A. Due to asymmetric routing or load balancing, some subsequent client or server traffic of connection 3 may arrive at I/F 3 or I/F 4 with VLAN B which is then forwarded to node 1

through IBP. This is to achieve symmetric inspection of traffic at the same tool. Regardless of incoming traffic interface, all outgoing traffic from GS1 to client 3 will be sent in I/F 1 with VLAN A and from GS1 to server 3 will be sent in I/F 2 with VLAN A.

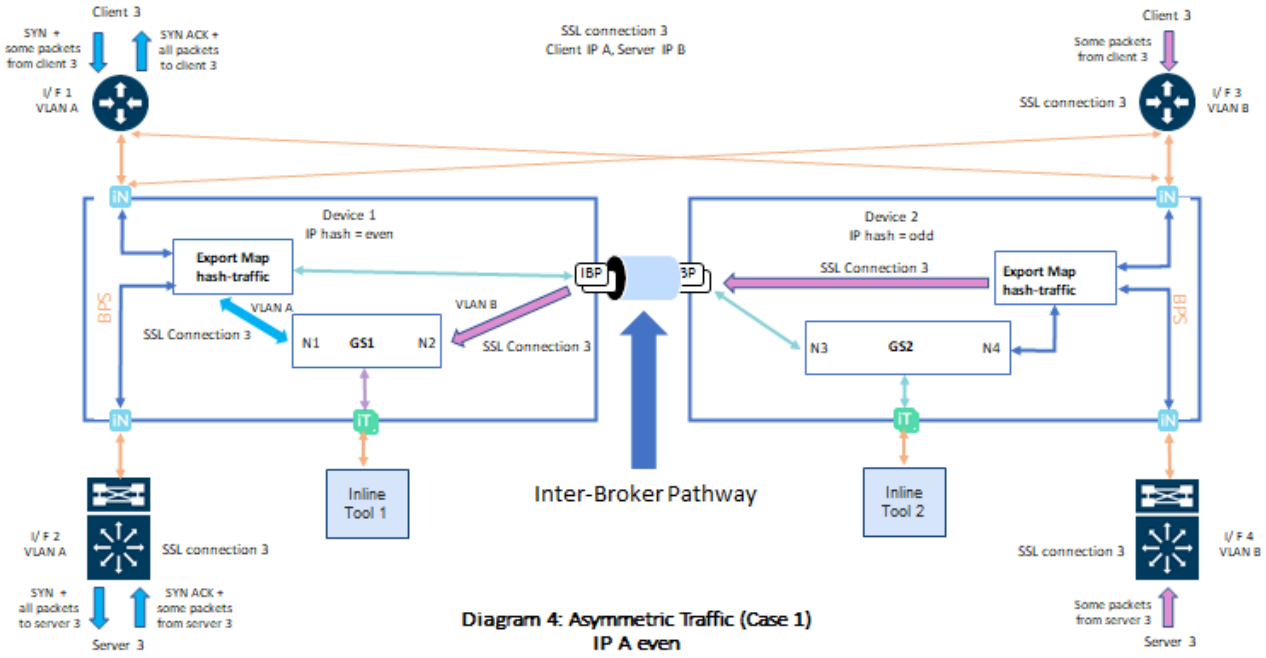

Diagram 4: Asymmetric Traffic (Case 1) IP A even

2. Consider the below arrangement wherein, the SYN packet of TLS/SSL connection 5 from client 5 is received on I/F 1 with VLAN A and it is forwarded to GS1 in node 1 since source IP (IP A) is even. GS1 will initiate TCP connection to Server 5 attached to node 1 by sending SYN packet out from I/F 2 with VLAN A . Server 5 responds SYN ACK from I/F 4 with VLAN B on node 2 and the packet is forwarded to GS1 through IBP since destination IP (IP A) is even. As SYN ACK is received from server 5 on I/F 4 with VLAN B in node 2, GS1 will send SYN ACK to client 5 on I/F 3 with VLAN B connected to node 2. Due to asymmetric routing or load balancing, some subsequent incoming traffic from client 5 may arrive I/F 3 with VLAN B on device 2, these incoming traffic will be forwarded to GS1 through IBP since source IP (IP A) is even. Similarly, some subsequent incoming server 5 traffic may arrive at I/F 2 with VLAN A on node 1. These traffic will also be forwarded to GS1 on node 1 since the destination IP (IP A) is even .The outgoing traffic from GS1 to client 5 will be on I/F 3 with VLAN B attached to node 2 and all outgoing traffic from GS1 to server 5 will be on I/F 2 with VLAN A attached to node 1 regard less of the incoming traffic interface.

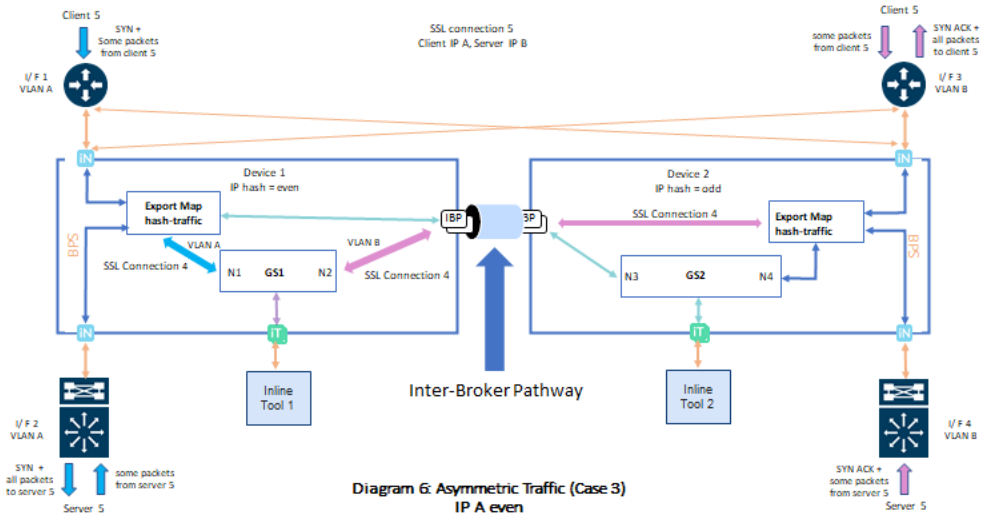

Diagram 6: Asymmetric Traffic (Case 3) IP A even

VLAN Tagging Behavior for Decrypted Traffic.

Inline Tool placed outside the SSL App can receive the original Map's VLAN tag for local node traffic or the import map's VLAN tag for remote node traffic. In asymmetric traffic, the Inline Tool will receive both the original Map's VLAN and the import map's VLAN.

Inline Tool placed inside the SSL App will always receive the tool tag configured in the SSL App for both local node traffic and remote node traffic.

VLAN Tagging Behavior for Non-Decrypted Traffic.

Inline Tool placed outside the SSL App can receive the original Map's/non-proxy Map's VLAN tag for local node traffic or the import map's/import non-proxy Map's VLAN tag for remote node traffic. In the case of asymmetric traffic, the Inline Tool will receive both the original Map's/non-proxy Map's VLAN and the import map's/import non-proxy Map's VLAN.

Inline Tool placed inside the SSL App will receive the non-proxy Map's Tool tag for local node traffic and import non-proxy Map's tool tag for remote node traffic.

Setup Resilient Inline TLS/SSL

To configure and deploy Resilient Inline TLS/SSL , ensure the following prerequisites are done for the nodes such as:

1. Configure the required inline networks. Refer to Configure Inline Network Ports and Inline Network.
2. Configure the required inline tools. Refer to Configure Inline Tool Ports and Inline Tools.

3. Configure the required inline tool group. Refer to Configure Inline Tool Group.
4. Create Inter-broker Pathway Refer to Configure Resilient Inline Arrangement
5. Click on  > **Inline Flows**> **Configuration Canvas** and select your node. Configure the Resilient Inline TLS/SSL profile:
   - Begin with configuring the Inline TLS/SSL App alias name.
   - Enable Resilient Inline Arrangements checkbox.
   - Select the nodes that would be configured and the respective GigaSMART modules.
   - Click on **Add Keys** under Deployment Type, to configure Key Store Certificates .The keys added will be pushed to both nodes. To delete the key you will have to do so from individual nodes.
   - Configure the Inline TLS/SSL profile fields for decryption and click on OK. For details about the Inline TLS/SSL Policy Profile fields and their descriptions, refer to *Inline TLS/SSL Policy Profile—Field References* in Configure Inline TLS/SSL Decryption Using GigaVUE-FM.
   - You can configure Inline SSL App for any one of the nodes. It will be available for the second node as well.

Once the necessary pre-requisites are configured , in the Flexible Inline Canvas do the following:

1. Drag and drop Inline Network/Inline Network Bundle into the canvas.
2. Drag and drop a Flex map, Inline Tools and Inline SSL APP that are available on both the nodes with same alias, to configure the Flex Inline TLS/SSL maps.
3. Under the Settings option, enable the *'Show Resilient Inline Menu'* checkbox and setup the Node, IB Pathway and Hashing configurations.
4. Click on **Deploy**.

Limitations

1. At a time, only one RIA enabled TLS/SSL app will be supported for the given set of two nodes needed for Resilient Inline TLS/SSL solution.
2. A combination of RIA enabled SSL app and normal Flex SSL app(RIA disabled) is not supported and will be blocked in GigaVUE- FM. If the node has a normal flex TLS/SSL app, you cannot add an RIA enabled TLS/SSL app and vice-versa.
3. Inline SSL apps in GigaSMART Flex Inline Solutions can be modified by removing the existing TLS/SSL app from all the solutions and deploying the same, followed by adding or replacing the TLS/SSL apps.
4. Inline-tool with shared mode false is not supported.
5. Inline NetLag as a source is not supported.
6. Editing tool side VLAN tag is not allowed.
7. Using protected ports (BPS ports) in would result half of the traffic to be un-inspected in case of node failure.

8. One-arm topology is not supported.
9. Tool early-engage is not supported.
10. Inline classic MAPs will not be supported.
11. Iboss and resilient hashing are not supported.
12. Cluster is not supported.
13. The deploying a flexible iSSL solution with single VLAN tag feature does not support double-tagged encrypted traffic.
14. When configuring a non Single VLAN Tagged iSSL Map along with a Single VLAN Tag that enabled iSSL Maps in a solution, the non SVT iSSL map must be configured as the lowest priority.

Refer the following Gigamon Validated Designs for more detailed information:

- Enabling SSL/TLS Inspection on Asymmetric Multi-Path Traffic
- Enabling SSL/TLS Inspection on a Network Having Asymmetric Routing

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, Entrust nShield Hardware Security Module (HSM) is integrated with Passive TLS/SSL decryption. Refer to Entrust nShield and 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 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 65Basic 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 65Basic TLS/SSL Handshake. The messages are encrypted after step 6, including the step 9 Finished message.

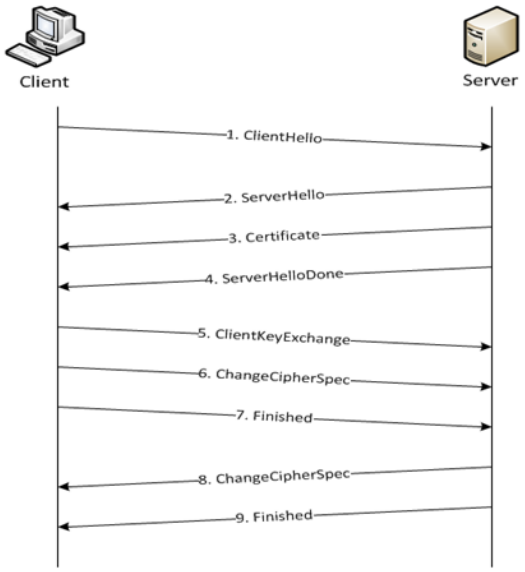

Figure 65 *Basic TLS/SSL Handshake*

Once a session has been established, the keys are saved so a session can be resumed efficiently later. Figure 66Resumed TLS/SSL Handshake shows the resumed TLS/SSL handshake, with fewer steps.

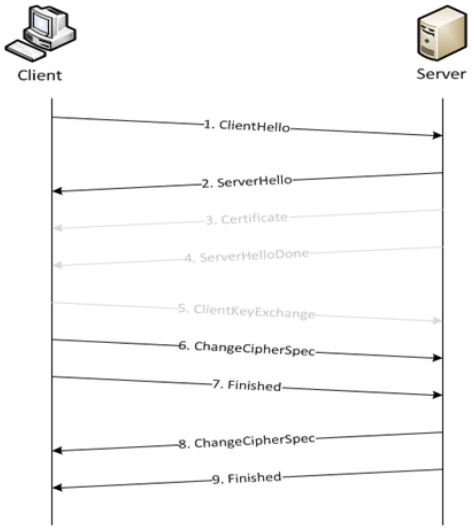

Figure 66 *Resumed TLS/SSL Handshake*

Passive TLS/SSL decryption can be deployed close to the server, as shown in Figure 67Inbound (Server Side).

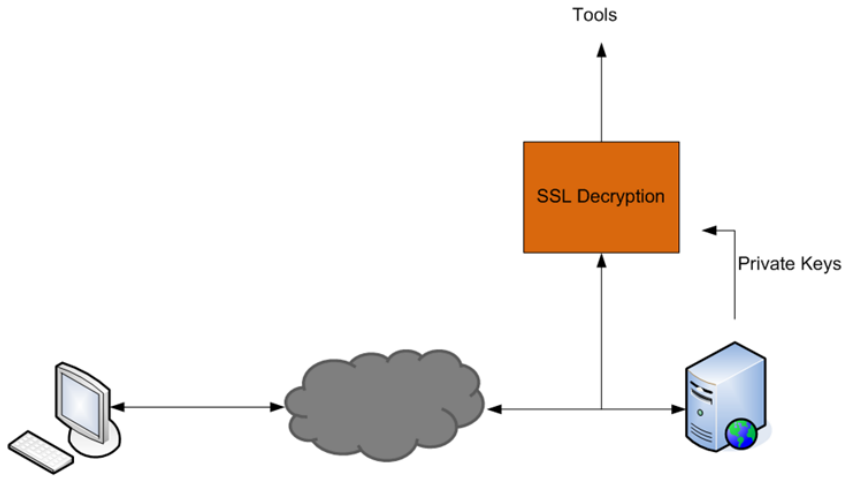

Figure 67 *Inbound (Server Side)*

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 68Outbound (Client/Enterprise Side).

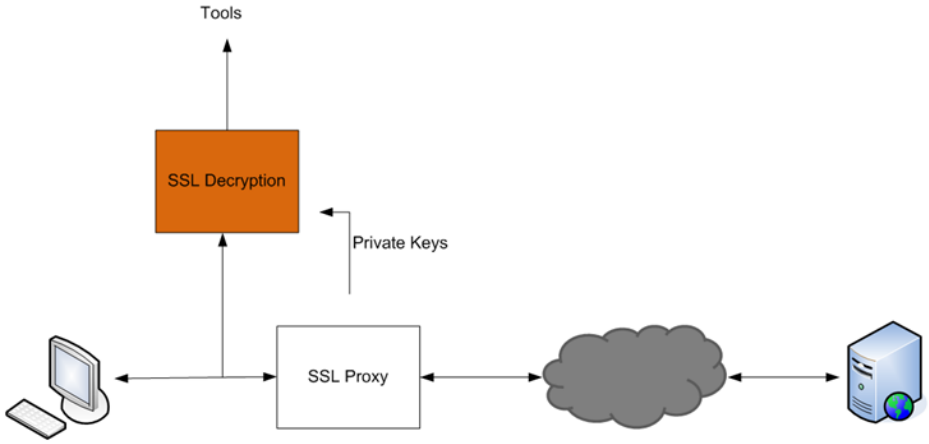

Figure 68 *Outbound (Client/Enterprise Side)*

In Figure 67Inbound (Server Side), tap traffic to the server and then send it for decryption. In Figure 68Outbound (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 18: Supported Ciphers for Passive SSL decryption.

Table 18: 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 140-2 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 GS 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. On the left navigation pane, click  , and then select **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 **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. On the left navigation pane, click  , and then select **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 Type**, select **RSA** or **ECDSA**.
   - For **File Type**, select **PEM**, **PKCS12** or **PKCS11**.
   - (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 Entrust nShield and 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. On the left navigation pane, click  , and then select **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 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.

ECODE	Description
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 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:

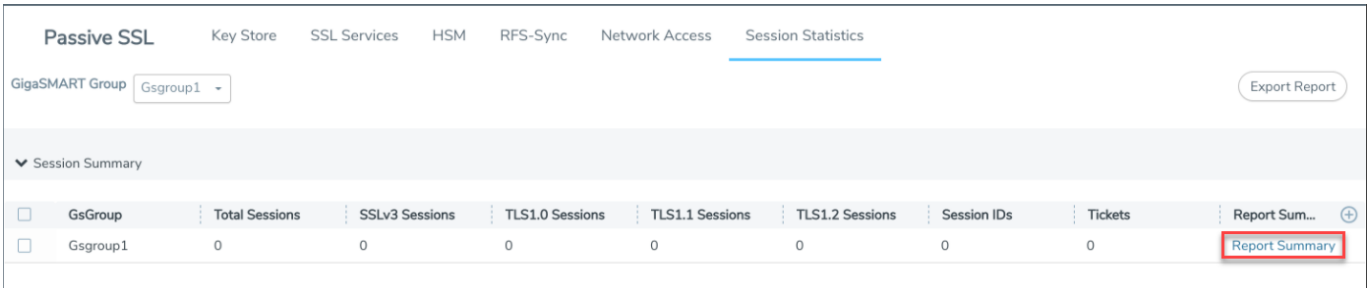



You can also choose to export the session summary to a remote server. Select the required row in the Session Summary table, and then click **Export Report**. The Upload Flow Ops Report page appears. Enter the remote server path and password to access the server. Click **Upload**. The details are exported to the remote server.

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 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**.

Entrust nShield and 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)

Starting in software version 5.3, Entrust nShield Hardware Security Module (HSM) is integrated with decryption for Out-of band Tools (Passive SSL/TLS Decryption). 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 Entrust nShield 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.

Entrust nShield HSM is supported on GigaVUE-HC1, GigaVUE-HC3,Generation 3 GigaSMART card (SMT-HC1-S) and GigaVUE-HCT.

Entrust nShield and 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 Entrust nShield HSM will use one Entrust nShield 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 non-HSM based decryption to Entrust nShield HSM based decryption.
    - From Entrust nShield 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
- Configure a GigaSMART Group
- 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**.

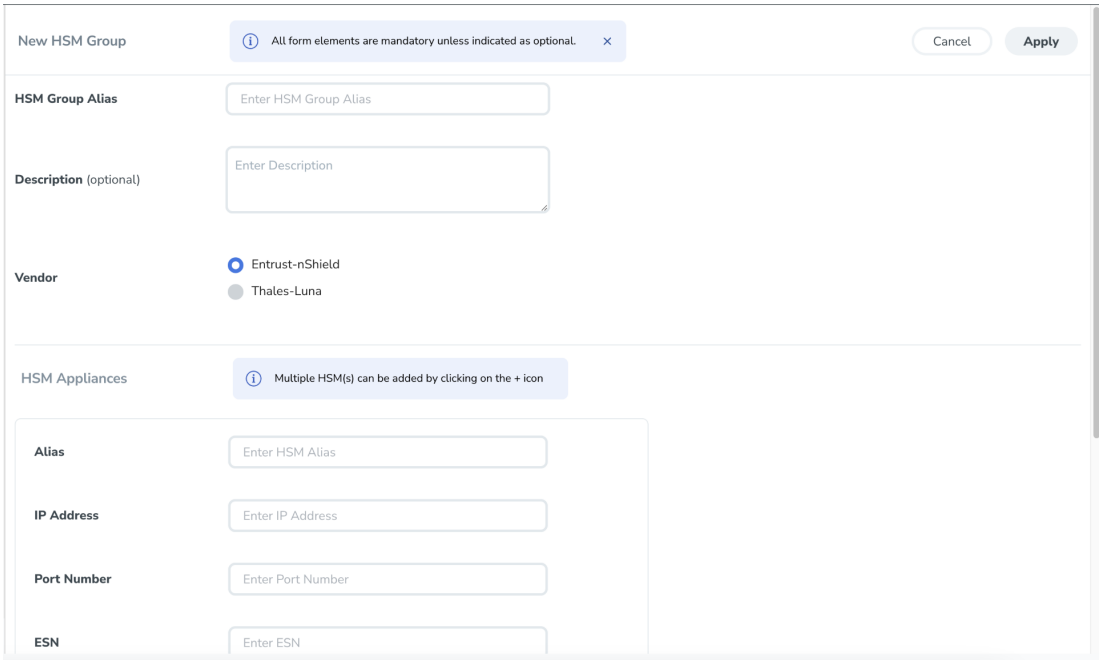

Figure 69 *Adding a New HSM Appliance*

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 (**Entrust-nShield** or **Thales-Luna**) to create the respective HSM Group. By default, **Entrust-nShield** will be selected as the vendor type.
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 **Entrust-nShield**:
   a. Enter the **ESN** (Electronic Serial Number) and **KNETI** that you obtained from the HSM administrator.
   b. Choose one of the following methods to select the required 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.

8. 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"

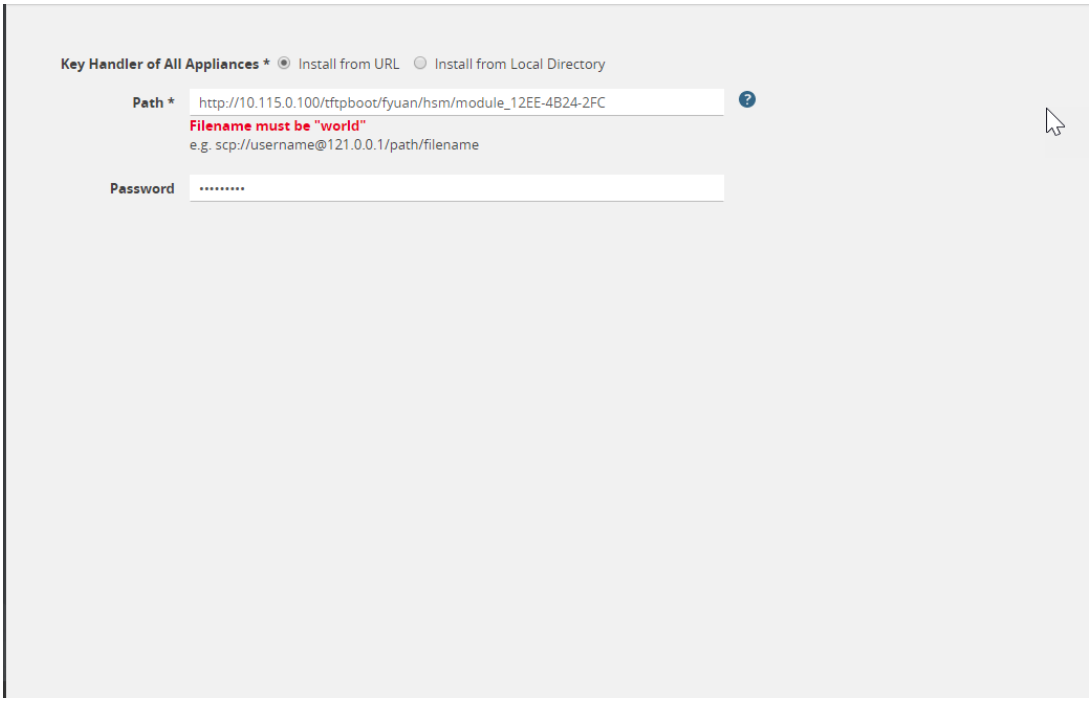

Figure 70 *HSM-Configure Key Handler*

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.

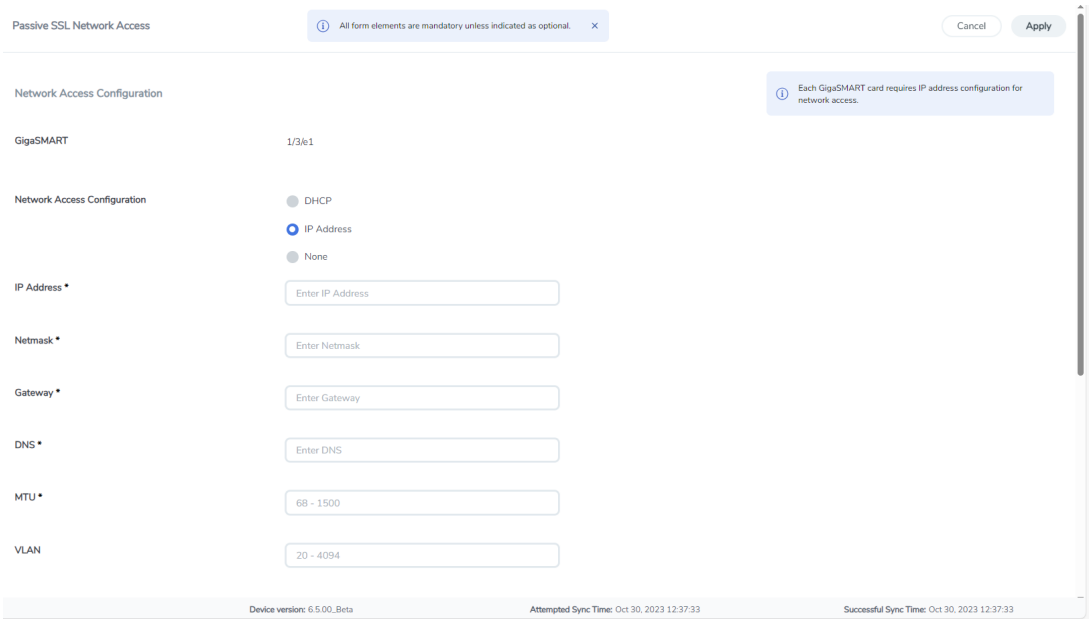

Figure 71 *Passive TLS/SSL Network Access - IP Configuration*

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 > Key Mapping**.
2. In the RFS section, click **New**. The Add RFS page appears.
3. Enter the IP address of the RFS where the encrypted keys are stored.
4. Select the **Enable** check box next to the **Automatic Sync** field to automatically synchronize the RFS with the GigaVUE-OS device.
5. In the **Sync Period** field, enter the time interval for synchronization in hours.
6. Click **OK**.

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 **Sync Period** and **Automatic Sync** fields for the RFS you have added. Click the server IP address to open the RFS quick view, and then click **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.

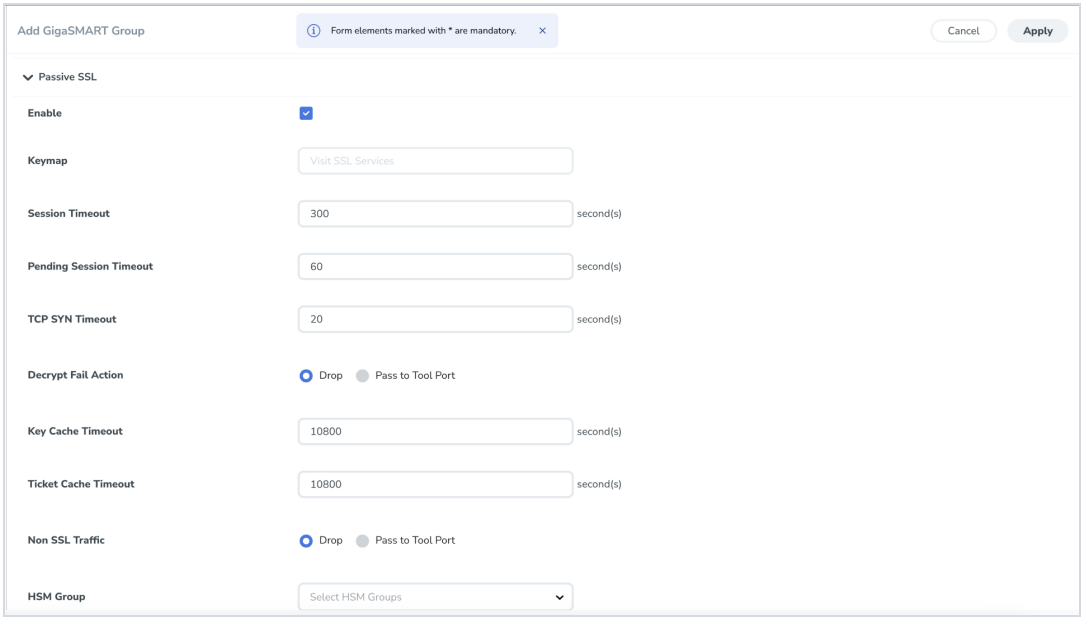

Figure 72 *GigaSMART Group Setup Page*

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**.

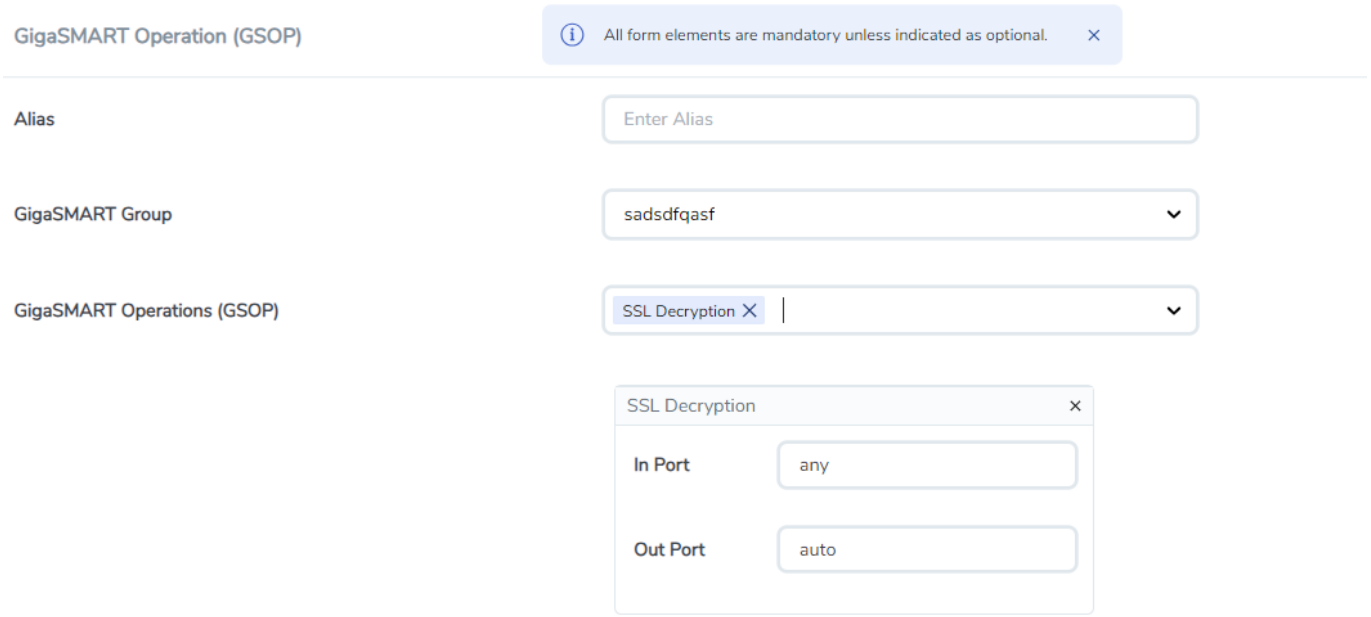

Figure 73 *GigaSMART Operations - Setup Page*

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.

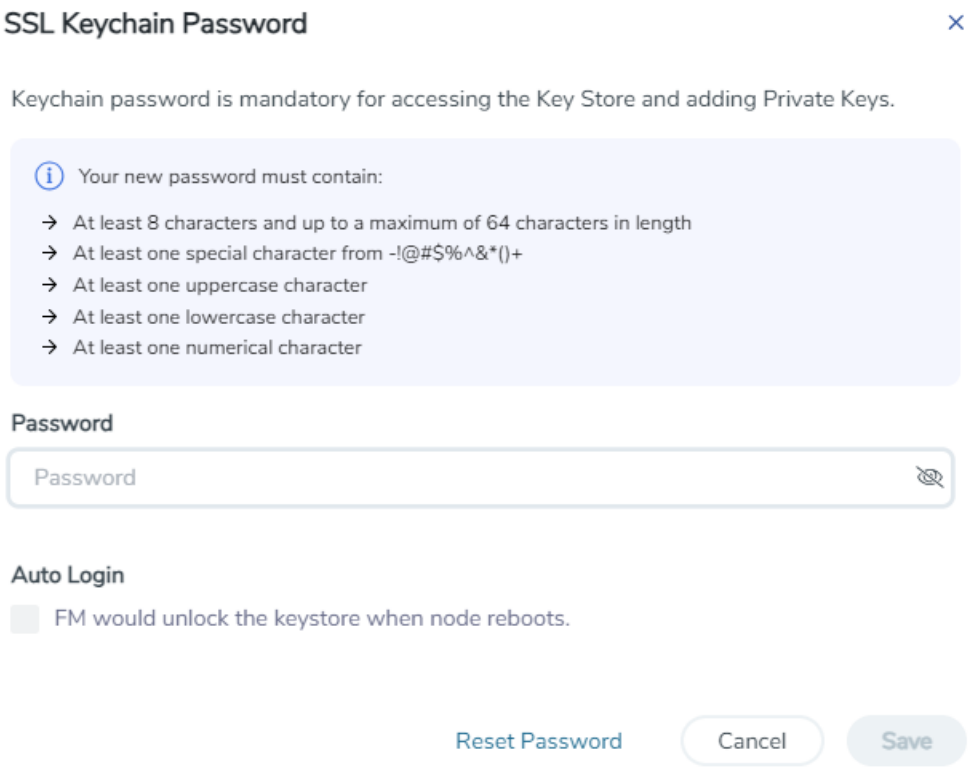

Figure 74 *SSL Keychain Password Setup Page*

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.

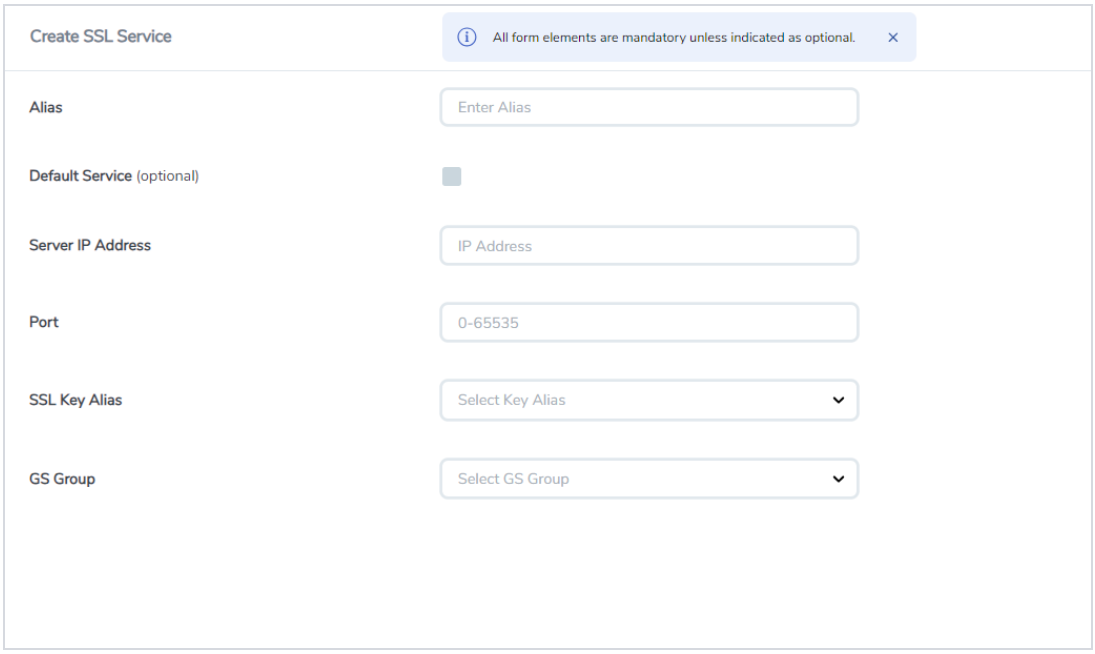

Figure 75 *SSL Service*

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.

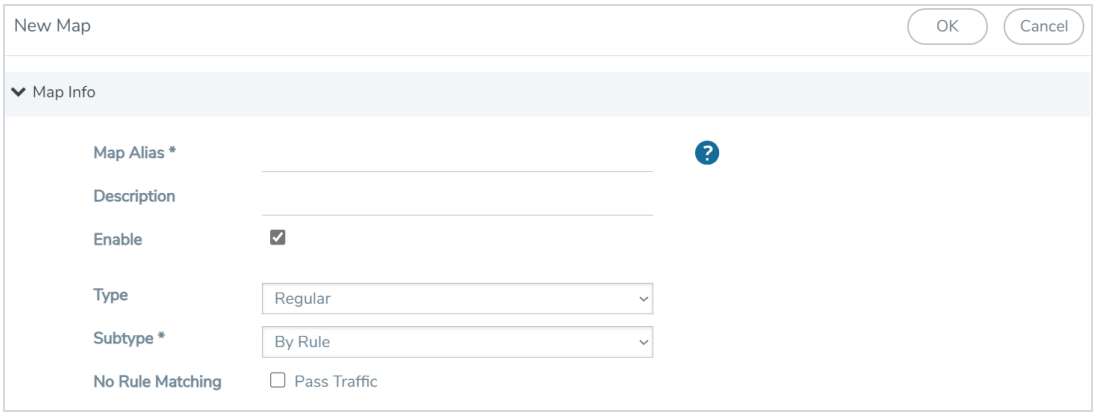

Figure 76 *Create New Map*

- 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.

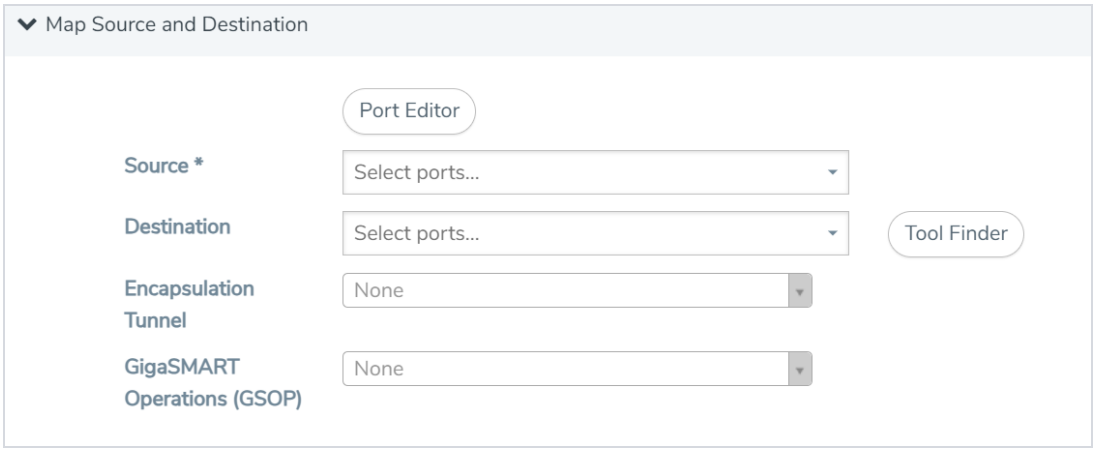

Figure 77 *Configure Map Details*

4. Add **Configuration & Rules**.

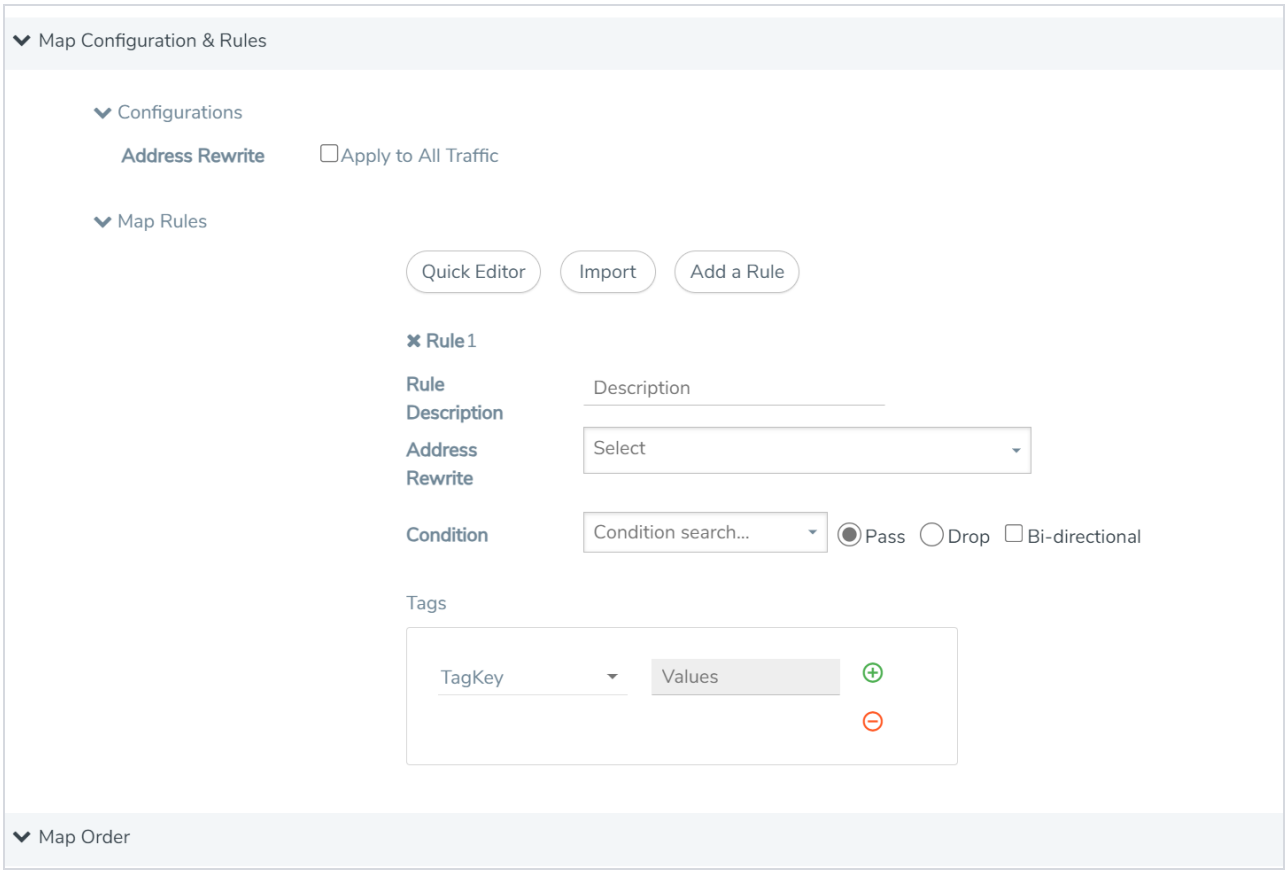

Figure 78 *Figure 20-123: Map Details - Create Rule*

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

Entrust nShield and Thales- Luna HSM for TLS/SSL Decryption for iSSL

The purpose of this feature is to provide the capability for inline-TLS/SSL feature to work with HSM hardware.

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.

Current Inline TLS/SSL decryption is enhanced to include Thales-Luna HSM support in addition to the current already supported Entrust nShield HSM solution.

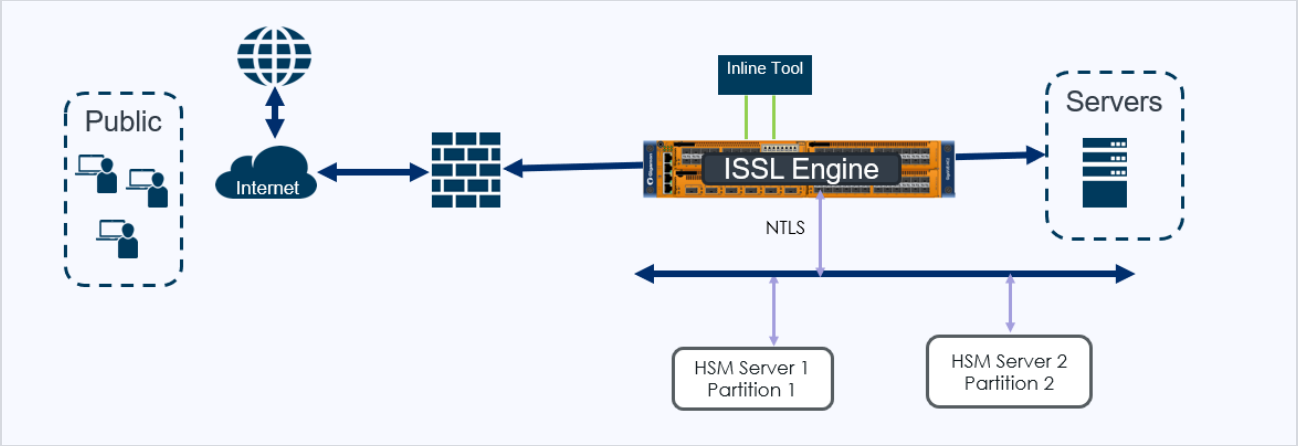

Figure 79 *ISSL with Thales Luna - Inbound support*

Supported Platforms

- GigaVUE-HC1 Gen3
- GigaVUE-HC3 Gen3
- GigaVUE-HC1-Plus

Limitations

Keep in mind the following limitations when configuring HSM:

- Gen3 Entrust nShield HSM iSSL and Gen3 Thales-Luna HSM iSSL cannot be configured together on the same GigaVUE HC Series device. (Also, you cannot configure each HSM on separate GigaSMART Gen3 cards of the same GigaVUE HC Series device.)
- Gen3 non-HSM iSSL and Gen3 HSM iSSL cannot be configured together on the same GigaVUE HC Series device. (Also, cannot be configured on separate GigaSMART Gen3 cards of the same GigaVUE HC Series device.)
- Gen2 non-HSM iSSL and Gen3 HSM iSSL cannot be configured together on the same GigaVUE HC Series device. (Also, cannot be configured on separate GigaSMART Gen2 and Gen3 cards of the same GigaVUE HC Series device.)
- 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

PKCS#11 Library

The PKCS#11 (Public Key Cryptography Standards) is a standard programming interface to communicate with HSMs. This standard specifies an application programming interface (API), called “Cryptoki” to devices which hold cryptographic information and perform cryptographic functions.

Proprietary interfaces using Secure Object Library are provided to interact with the HSM for:

- Generating key pair within the HSM.
- Installing existing key in the HSM.
- Manufacturing Protection key operations.

Configure HSM for TLS/SSL Decryption for iSSL

Refer to Configure Hardware Security Model (HSM) and Configure Flexible Inline TLS/SSL Decryption Solution sections for more details on how to configure HSM Group and use HSM Group for TLS/SSL decryption for iSSL.

If a HSM Decryption deployment is modified follow the below steps:

1. Move the Inline Network traffic path to bypass mode.
2. Make the desired deployment change such as:

- From non-HSM based decryption to Thales-Luna HSM based decryption.
- From non-HSM based decryption to Entrust nShield HSM based decryption.
- From Entrust nShield HSM based decryption to Thales-Luna HSM based decryption.
- From HSM based decryption to non-HSM based decryption.

3. . Reboot the GigaSMART card.
4. Move the Inline Network out of bypass mode to 'To inline Tool' mode.
5. There should be at least one active Luna HSM in the High Availability to ensure that the decryption is not interrupted.
6. Do not attempt to reload the device, if a partition fails in the Luna HSM configuration and at least one active partition is present.
7. 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.

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, GigaVUE-HC1 Plus Gen3, and GigaVUE-HCT Gen3.

Refer to Supported GigaSMART Operations for 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 five (5) maps can be attached to a virtual port (vport). Each map can contain up to 25 gsrules.

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. On the left navigation pane, click  , and then select **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.

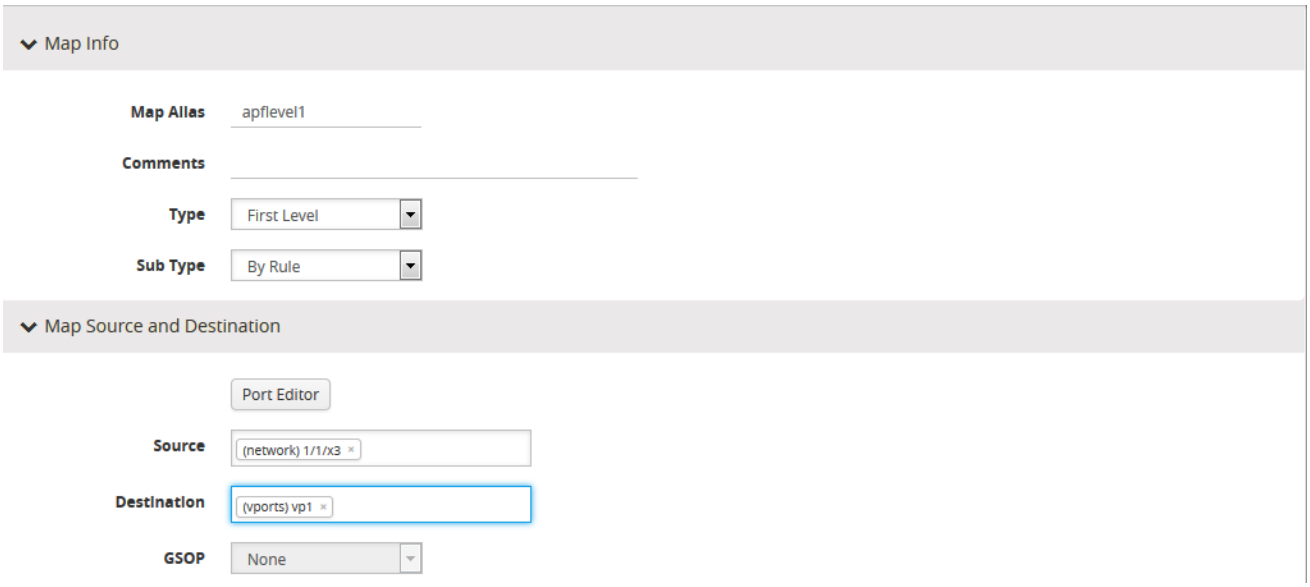

7. Create a second level map with the APF GigaSMART operation, the virtual port as the source, and a rule. The following figure shows an example.

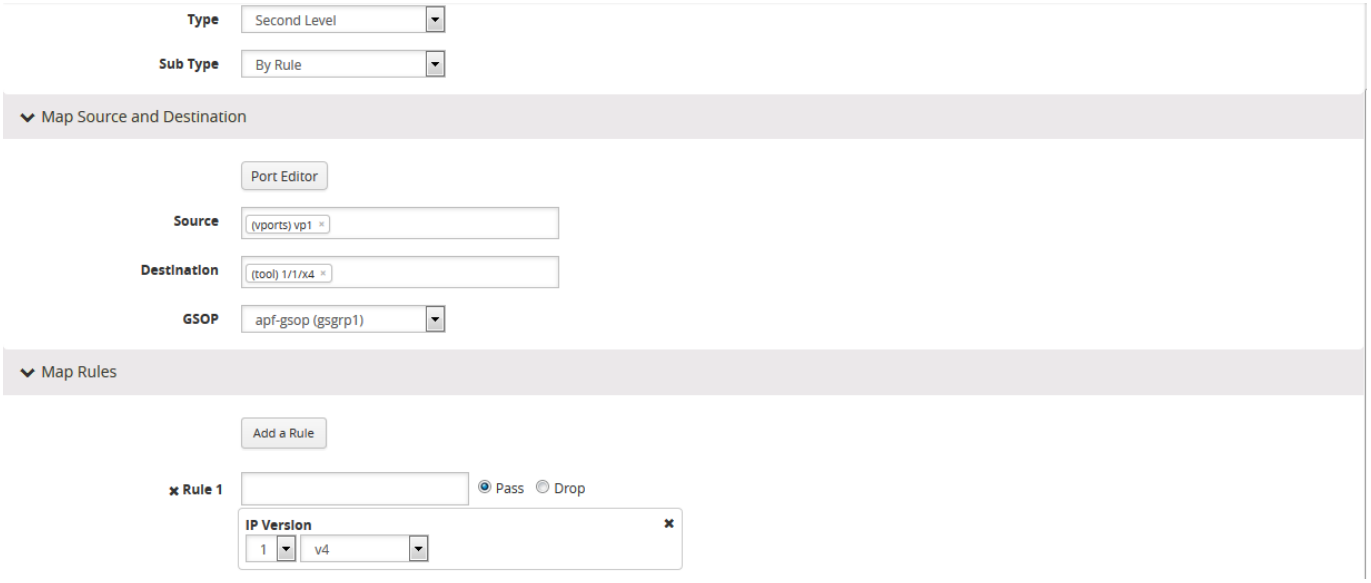

This completes the process to create an APF GigaSMART operation and corresponding rules. To learn more about the rules applicable for APF, see following sections.

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.

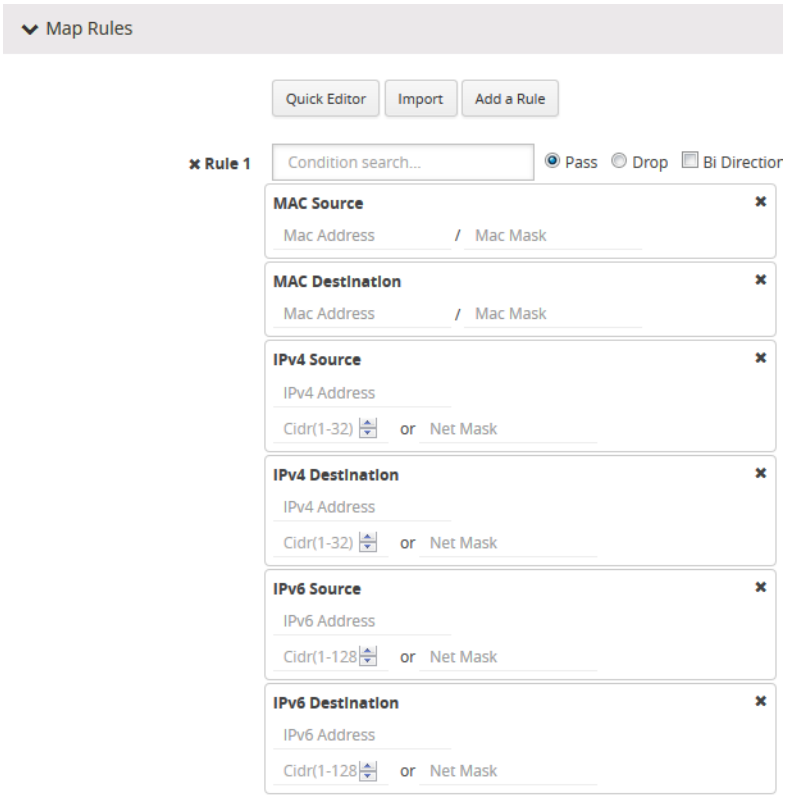

The second group of attributes has the following format:

```
<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.

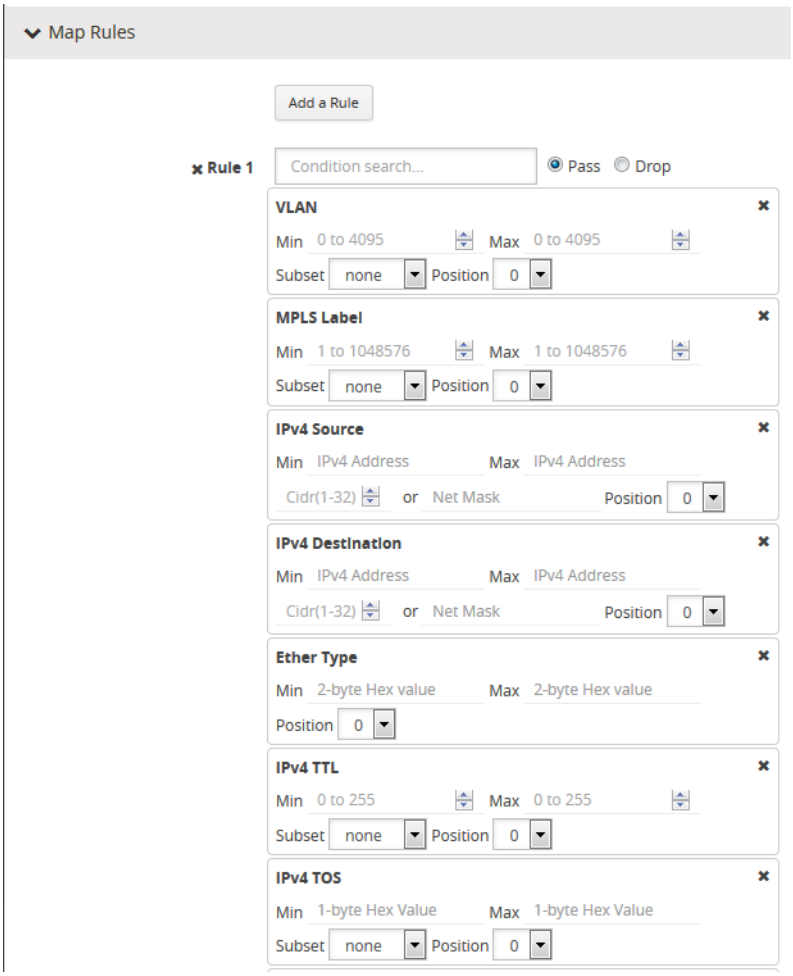

The third group of attributes has the following format: <value> <position>

```
<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.

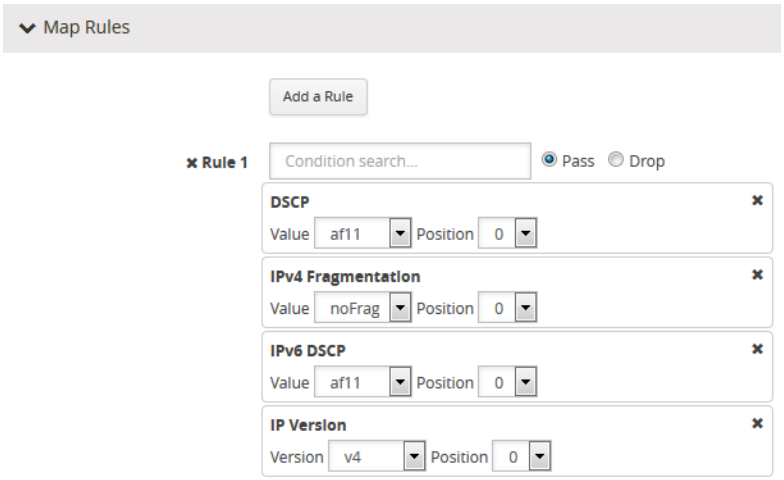

The fourth group of attributes has the following format:

```
<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.

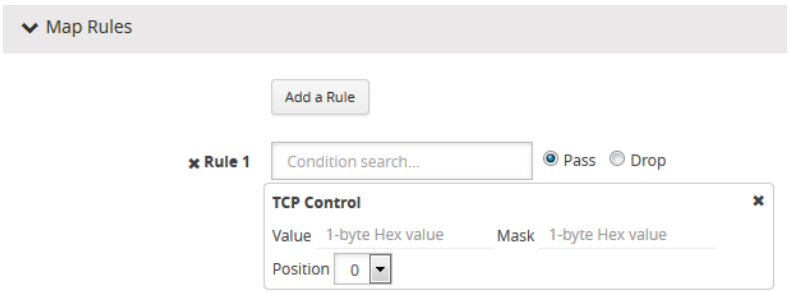

The maximum occurrences of each attribute supported are as follows:

Attribute	Maximum Occurrences
Attributes in IPv4 header	3
Attributes in IPv6 header	3
Attributes in MAC header	3
VLAN ID	4
MPLS label	4

Attribute	Maximum Occurrences
Attributes in L4port	3
Ethertype	6
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.

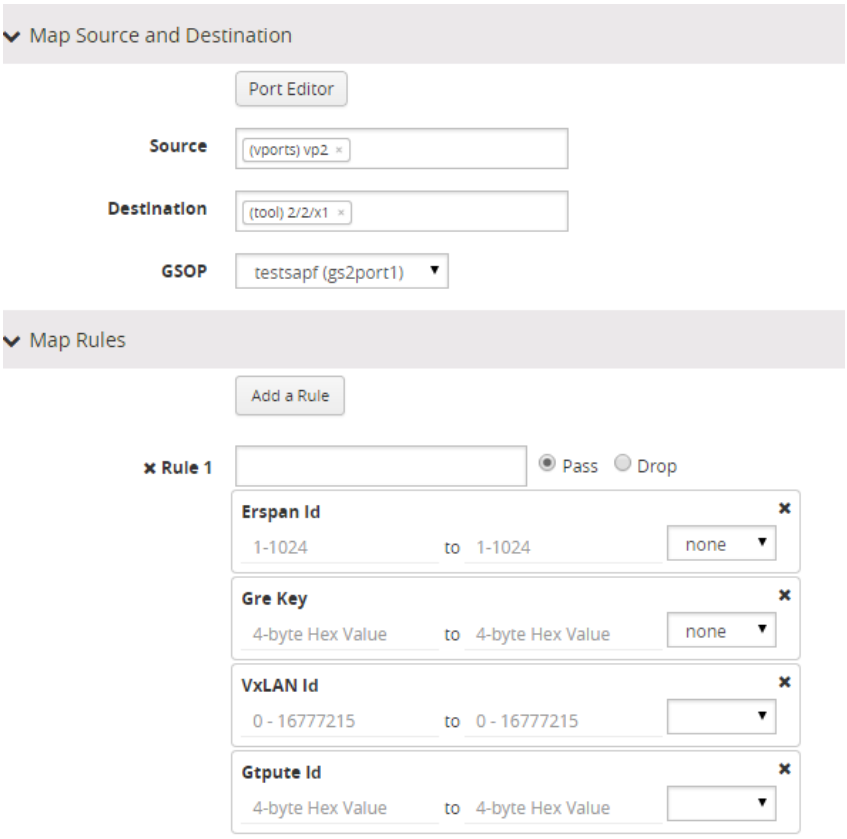

Pattern Matching on Gen 2 GigaSMART modules

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 80Use Pattern Match Under Maps for Pattern Matching.

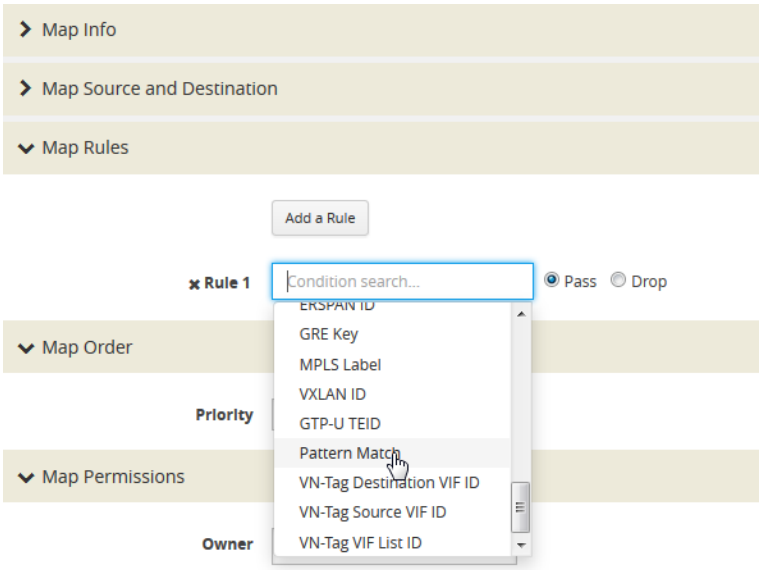

Figure 80 *Use Pattern Match Under Maps for Pattern Matching*

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 81Pattern Match with Type String.

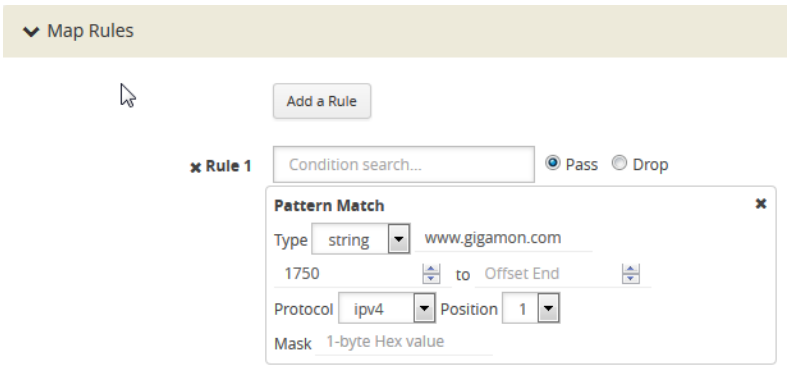

Figure 81 *Pattern Match with Type String*

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 82Pattern Match with Type RegEx.

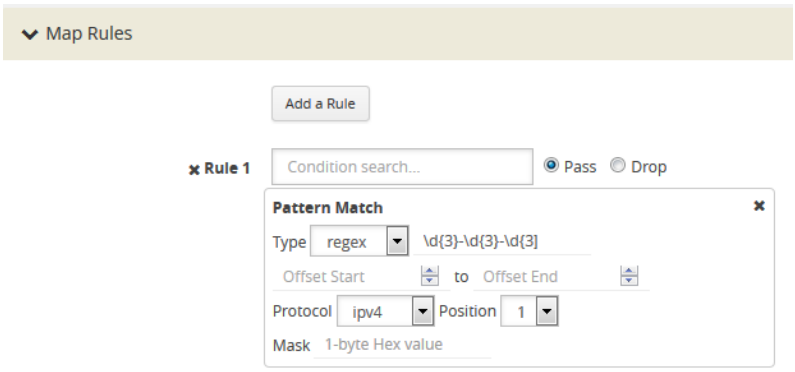

Figure 82 *Pattern Match with Type RegEx*

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 83Pattern Match for SSL Client Hello Packet.

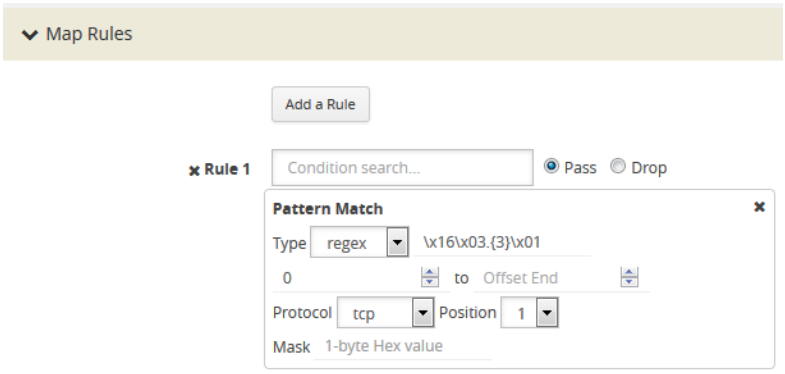

Figure 83 *Pattern Match for SSL Client Hello Packet*

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.

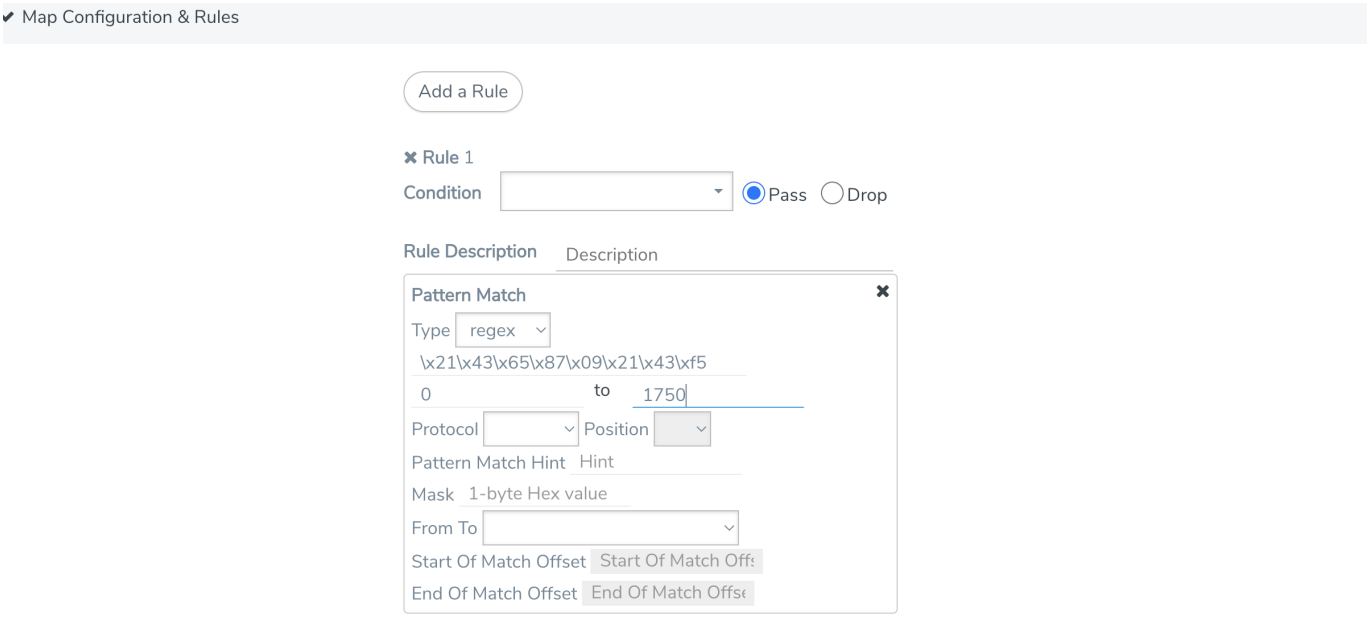

Masking with Pattern Matching

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 84Map Rule with RegEx for Masking SSNs.

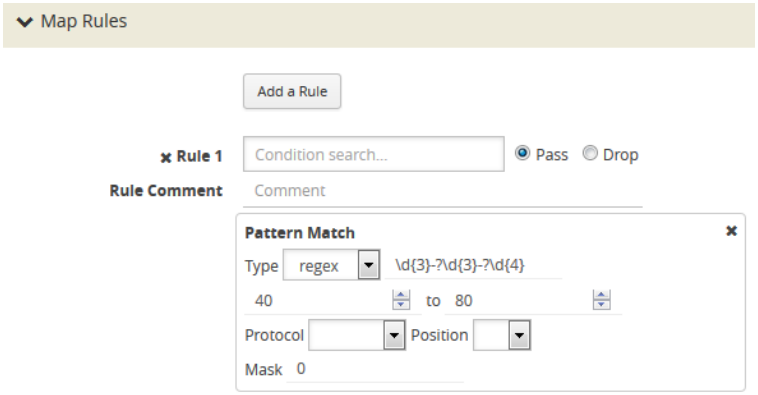

Figure 84 *Map Rule with RegEx for Masking SSNs*

Pattern Matching Hint

To optimize APF pattern matching performance in second level maps with gsrules, you can optionally use a pattern matching hint.

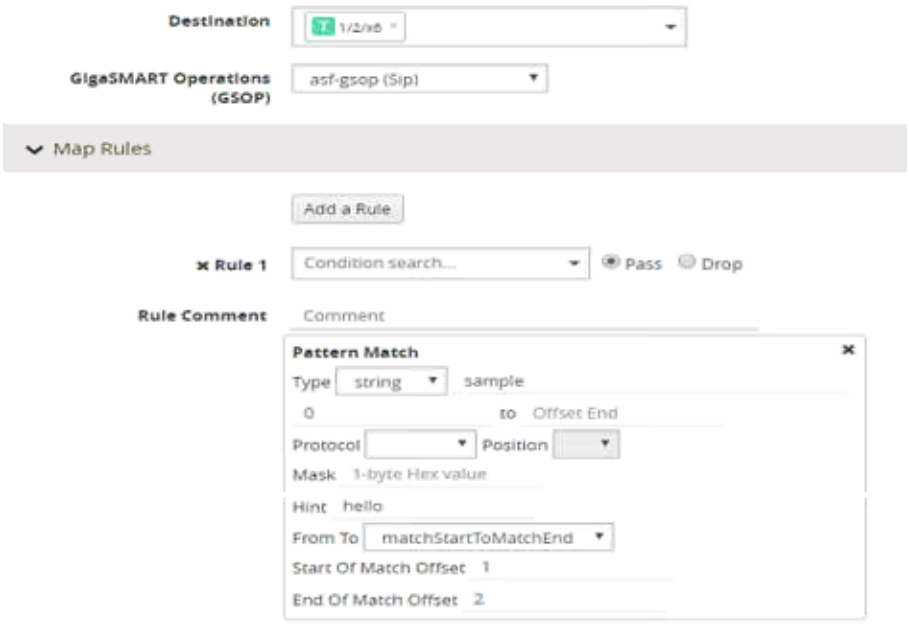

Figure 85 *Pattern Match with Hint*

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.

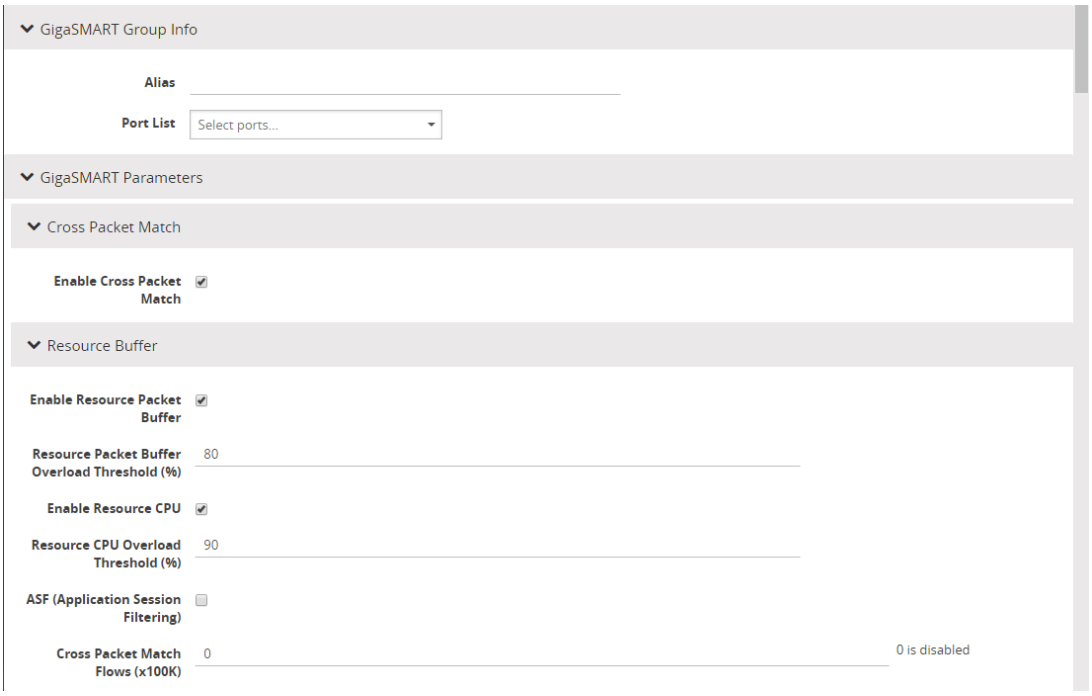

4. Click the **Enable Cross-packet Match** check box to enable.
5. Enter a range from 1 to 10 for the **Cross Packet Match Flows** parameter. Each unit is 100K bi-directional flows.
6. Click **OK**.

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 rules per second level MAP and 5 second level MAP. 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.

Unsupported PCRE RegEx patterns

The following PCRE RegEx patterns are not supported in Gen 3 GigaSMART rules:

- 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

The Pattern Match attribute in a map rule is shown in Figure 80Use Pattern Match Under Maps for Pattern Matching.

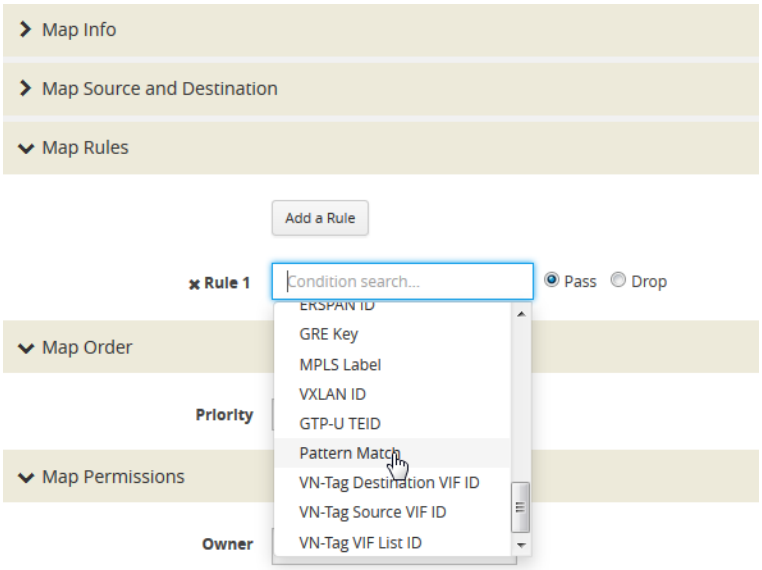

Figure 86 *Use Pattern Match Under Maps for Pattern Matching*

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 81Pattern Match with Type String.

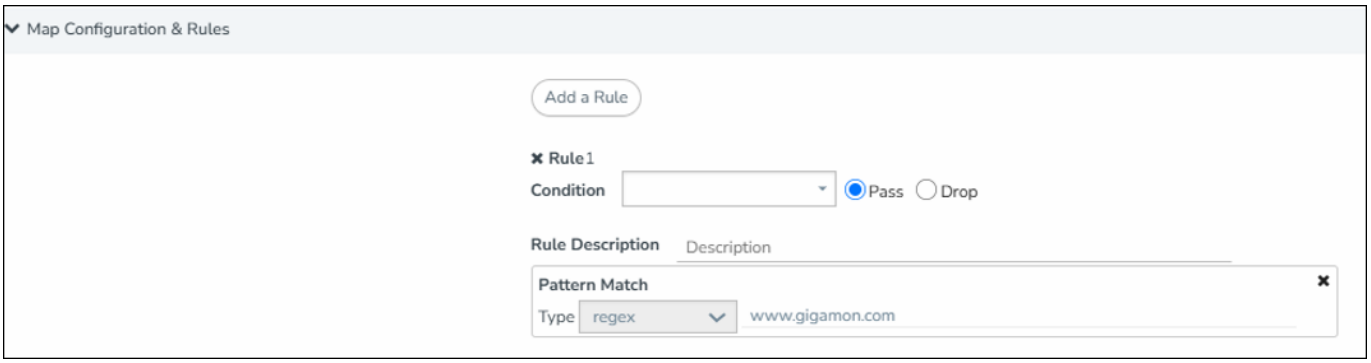

Figure 87 *Pattern Match with Type RegEx*

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 82Pattern Match with Type RegEx.

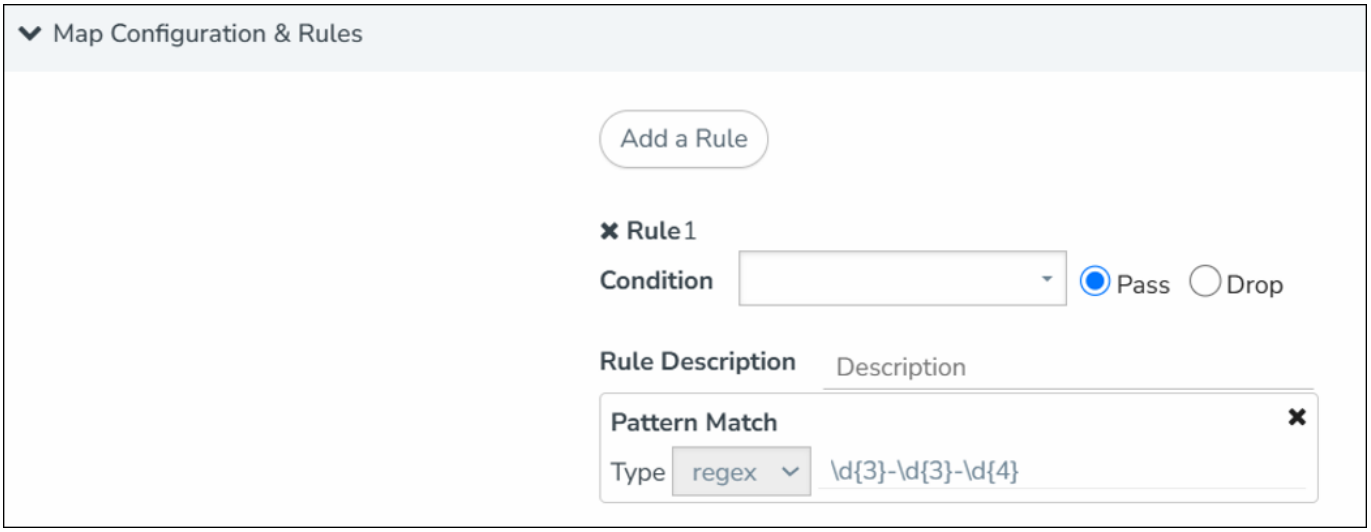

Figure 88 *Pattern Match with Type RegEx*

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	✓	×

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.

Task	Description	UI Steps
4	Create a virtual port.	**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 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.

Task	Description	UI Steps
		▪ Select Pass. ▪ Select Pattern Matching. ▪ Select regex for Type and enter the value \d{3}-?\d{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.

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.

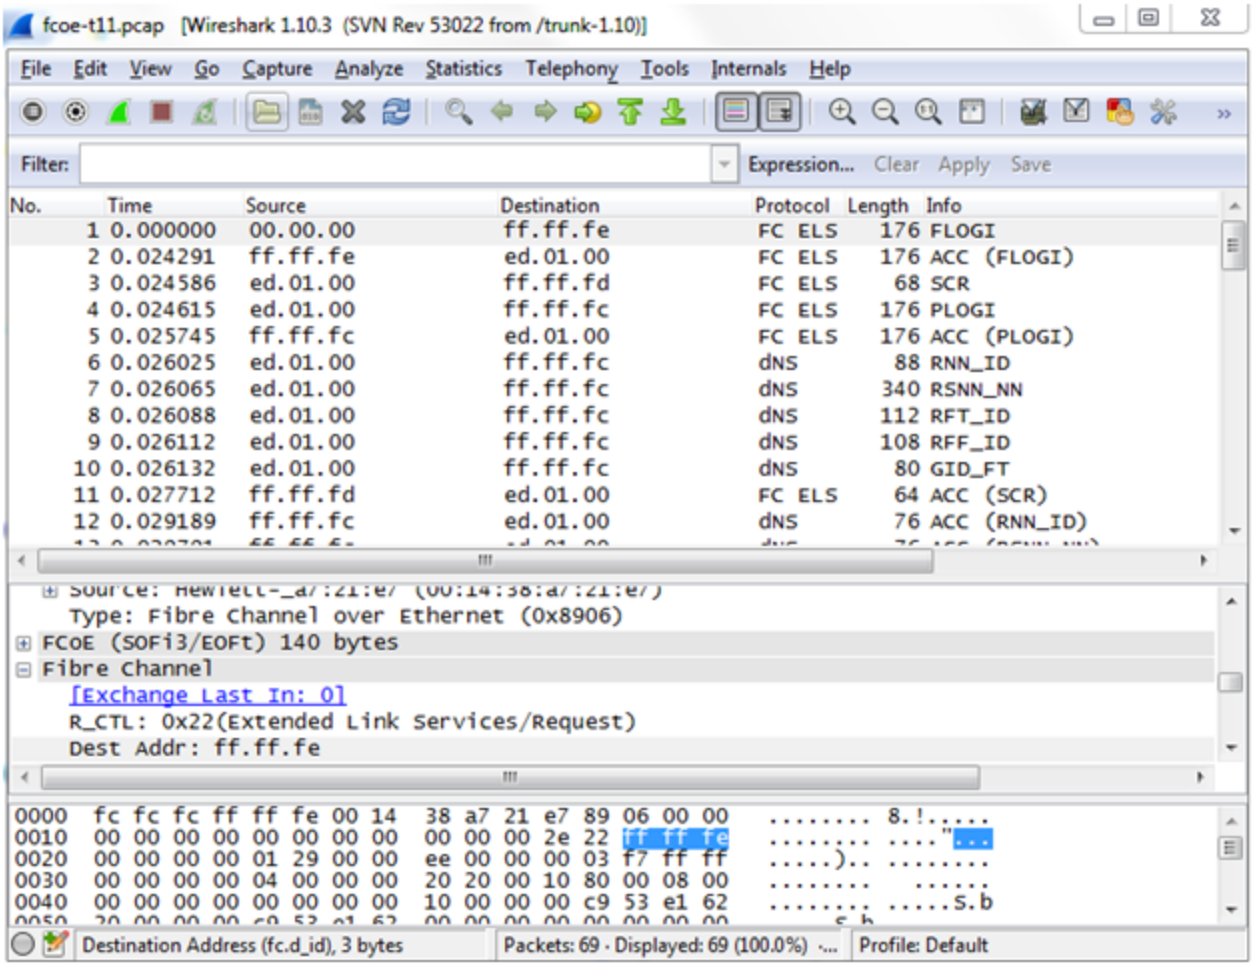

Task	Description	UI Steps
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3. Select Tool for port 1/1/x4 and port 1/1/x1. **4.** Select Enable for each port. **5.** 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. **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.

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 gsTraffic in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Click Save.
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.

Task	Description	UI Steps
		▪ 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 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 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

Task	Description	UI Steps
		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. ▪ 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.

Task	Description	UI Steps
		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/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

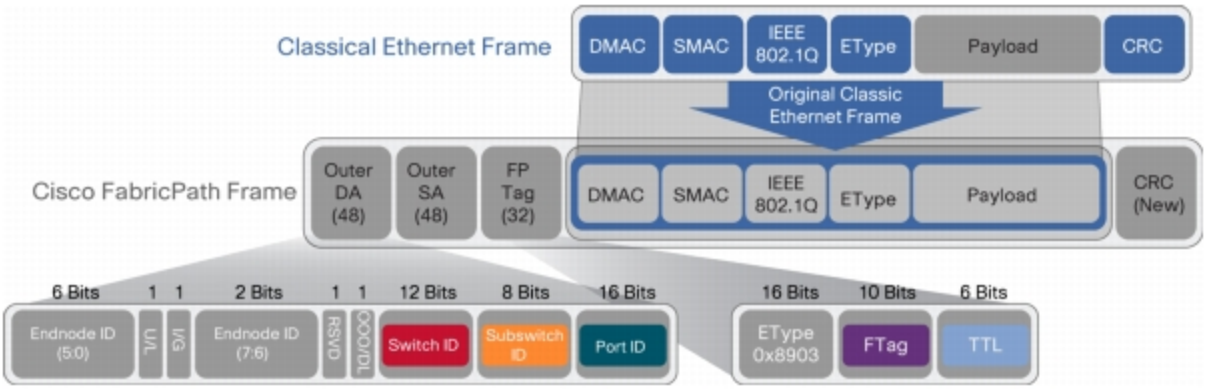

Task	Description	UI Steps
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3. Select Tool for port 1/1/x4 and port 1/1/x1. **4.** Select Enable for each port. **5.** 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.

Task	Description	UI Steps
		5. Click Save.
5	Create a first level map to forward fabric path packets 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 8903 in the Value field. **5.** Click Save.

Task	Description	UI Steps
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-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/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

Task	Description	UI Steps
		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

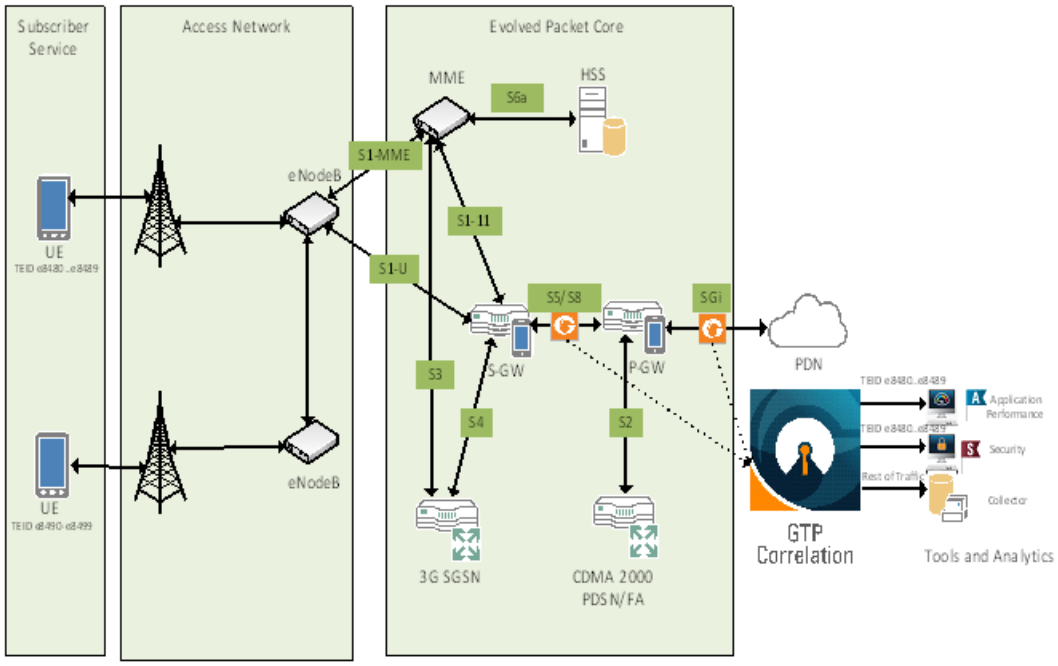

Task	Description	UI Steps
1	Configure one network and three tool type of ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x9. Select Tool for the port s 1/1/x13, 1/1/x14, and 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

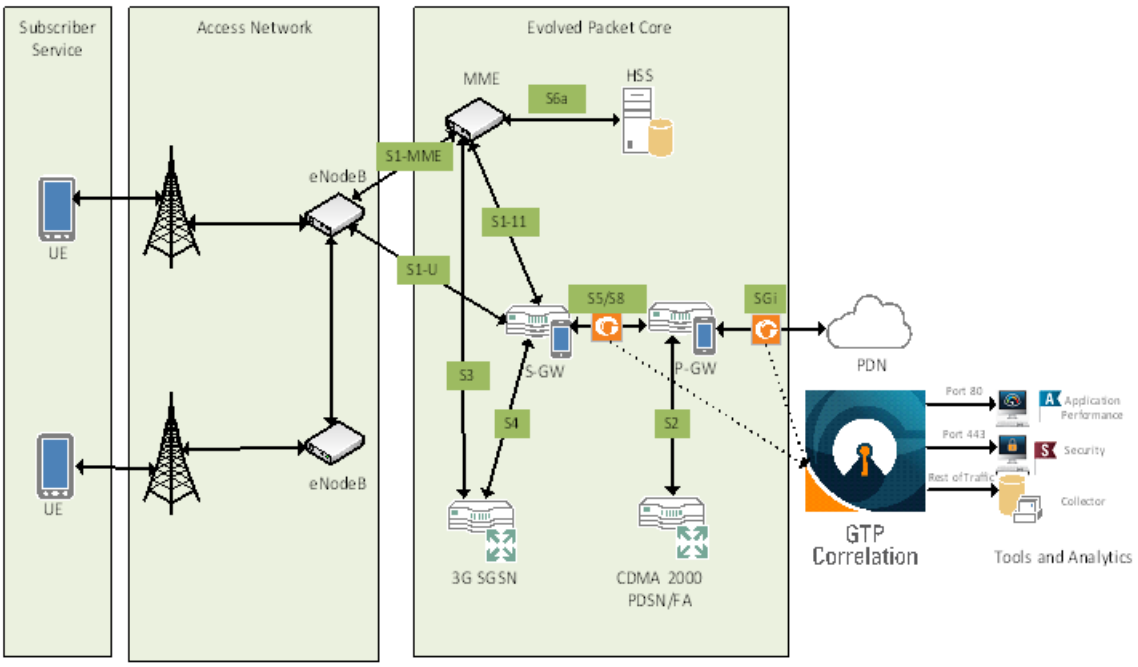

Task	Description	UI Steps
1	Configure one network and three tool type of ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and three tool ports. For example, select Network for port 1/1/x1. Select Tool for the port s 1/1/x10, 1/1/x11, and 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

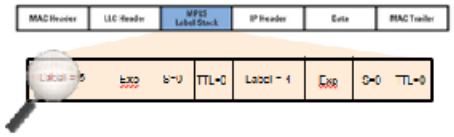

Step	Description	Command
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3 and select Tool for the port s 1/1/x4 and 1/1/x1 **4.** Select Enable for each port. **5.** 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.

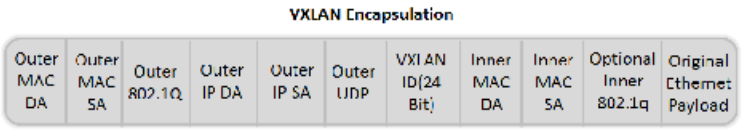

NOTE: This can be applied to any protocol that is supported through Header Stripping. GTP, VXLAN, ISL, MPLS, MPLS+VLAN, VLAN, VN-Tag, and fabric-path are all supported, as is Gigamon tunnel decapsulation.

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 for 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.

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.

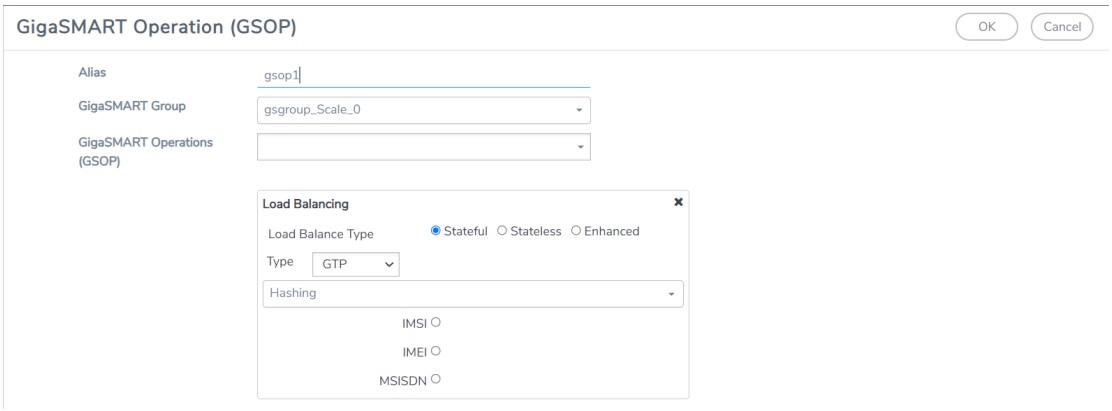

For details on stateful load balancing, refer to the following sections:

- 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.

Metric	Description
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. 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. Enable **Load Balancing.**

5. Select the port type.
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

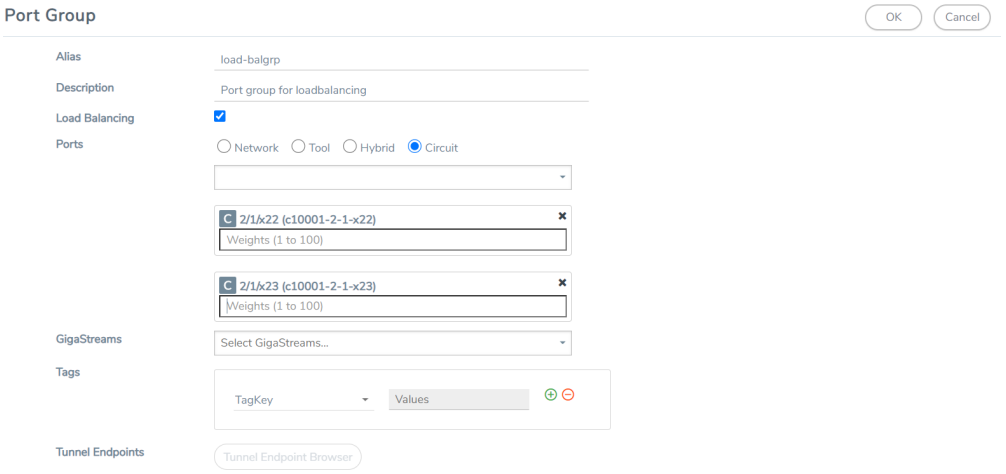

7. Click **OK** to save the configuration.

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.

Task	Description	UI Steps
		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. **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.

Task	Description	UI Steps
		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

Task	Description	Steps
		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 **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.

Task	Description	Steps
		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**.

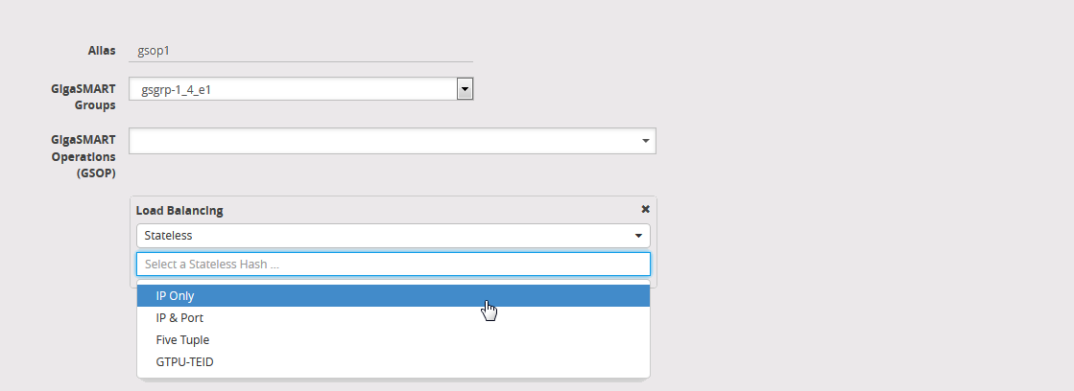

For details on stateless load balancing, refer to the following sections:

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 19: Stateless Load Balancing Metrics describes each of the metrics.
4. Click **Save**.

Table 19: 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

Metric	Description
	numbers.
Five Tuple	The source IP and destination IP addresses, source port and destination port numbers, and protocol field in the IP header.
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, and ERSPAN.

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. Enable **Load Balancing.**
5. Use the Ports field to select the ports for this port group.
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

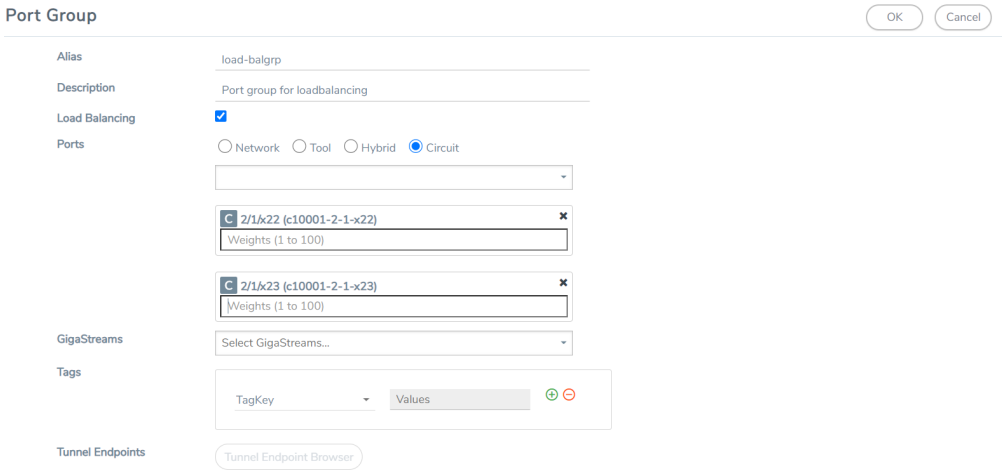

7. Click **OK** to save the configuration.

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.** 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, 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 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	**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 lbiponlyouter from the GSOP list. **h.** Click Add a Rule o Select Pass o Select IP Version for the condition.

Task	Description	Steps
	balancing.	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. **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:

Task	Description	Steps
		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.

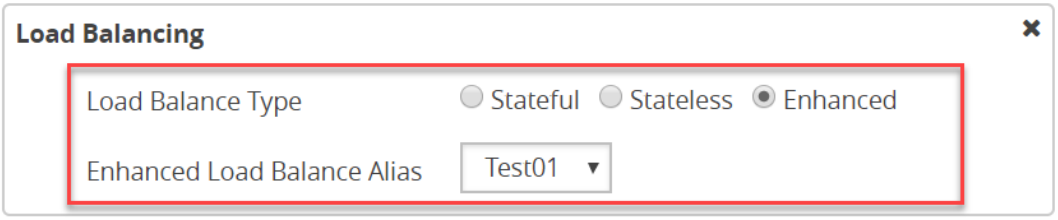

10. Click **OK**. The enhanced load balance GSOP is now available on the GSOP page.

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 for 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 89MPLS Header Between MAC and IP Address in Packet shows the MPLS example.

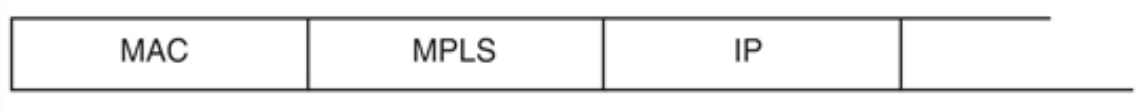

Figure 89 *MPLS Header Between MAC and IP Address in Packet*

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 for 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 90Strip 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 91Tunnel Stripping GigaSMART Operation Configured shows an example of the GS Operations page configured for tunnel stripping.

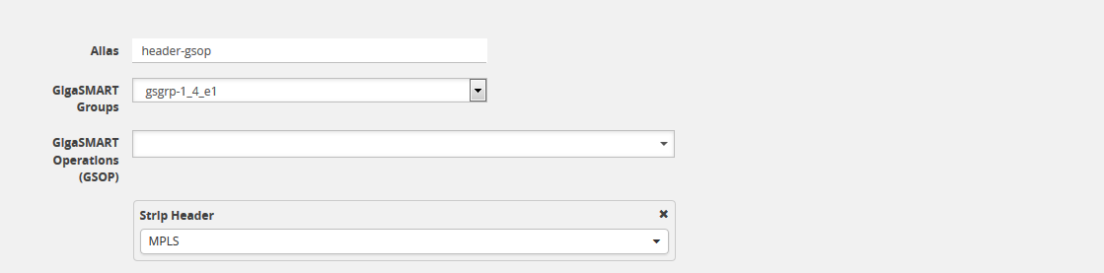

Figure 90 *Strip Header GigaSMART Operation Configured*

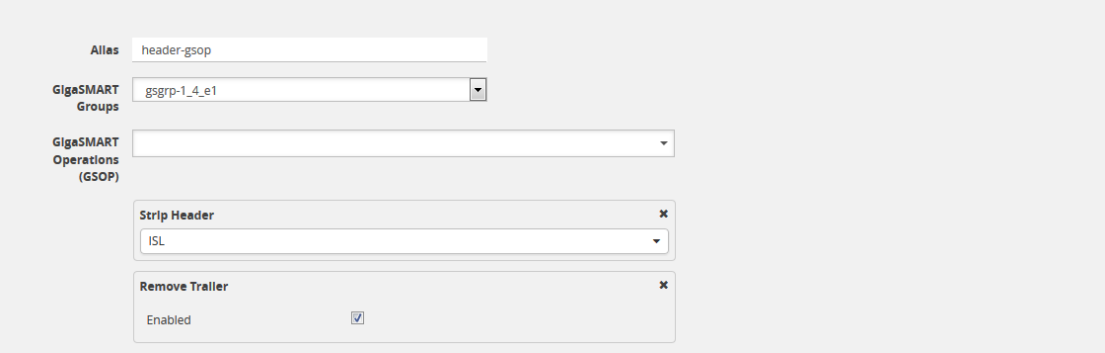

Figure 91 *Tunnel Stripping GigaSMART Operation Configured*

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.

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.

Summary	Description
	NOTE: When processing packets with multiple encapsulation layers – for example, an ERSPAN-tunneled packet with a VXLAN tag – a VXLAN header-stripping 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

Summary	Description
	View GigaSMART Statistics for more information.
Combine with Other Components	You can combine the **Strip 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 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.

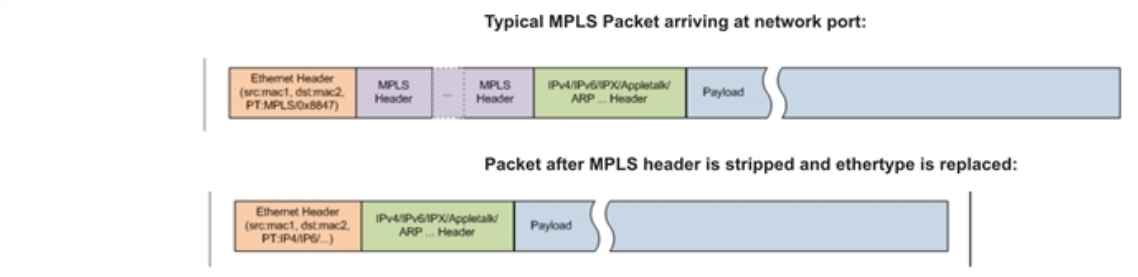

Example – Stripping MPLS Headers and Adding a VLAN ID

The example shown in Figure 92Strip 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.

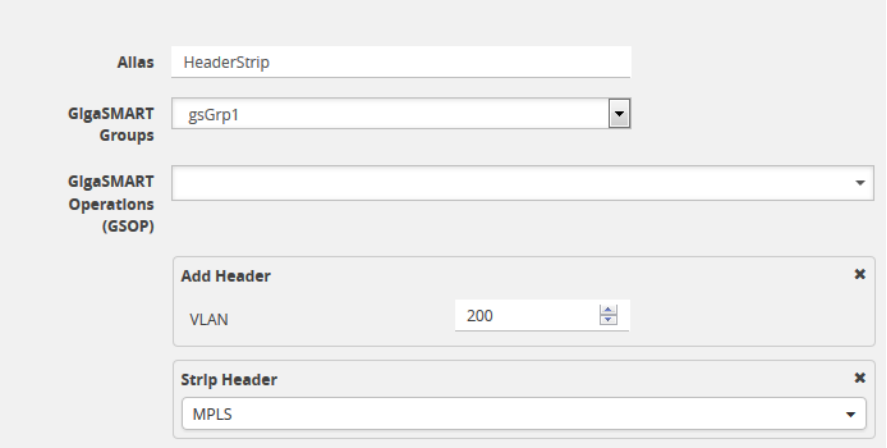

Figure 92 *Strip Header GigaSMART Operation for Stripping MPLS Headers*

The example shown in Figure 93Strip 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 93 *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.

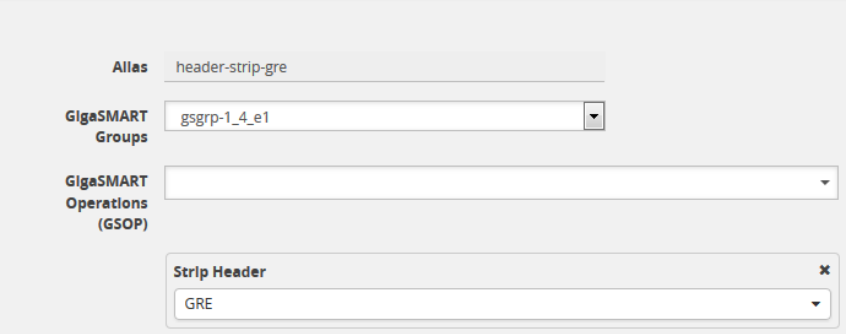

The following figure shows L3 GRE IP encapsulation before and after stripping.



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.

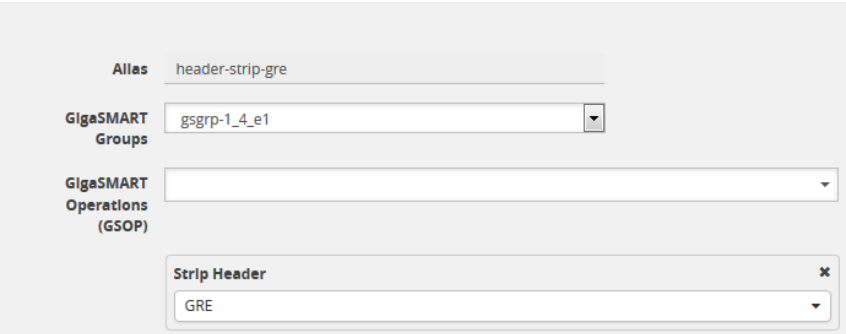

The following figure shows L2GRE Ethernet encapsulation before and after stripping.



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 94Strip Packets with FM6000 Timestamp shows an example GigaSMART Operation configured to strip packets containing the FM6000 timestamp:

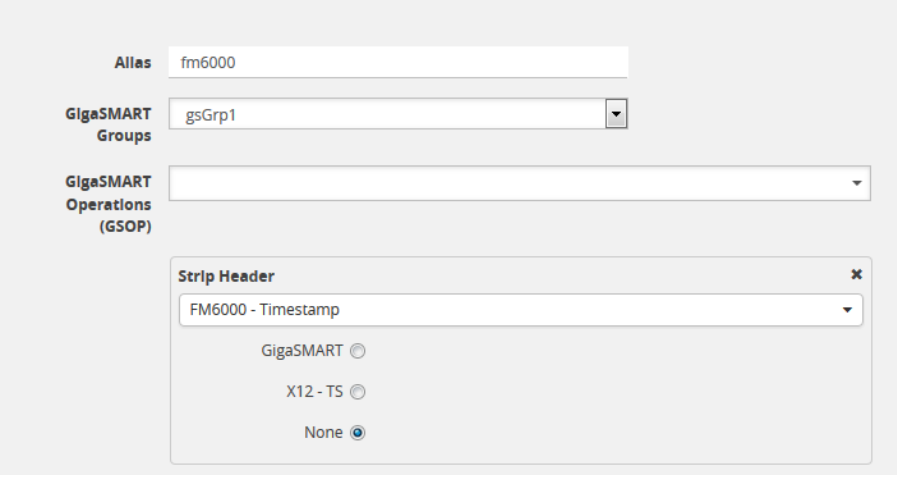

Figure 94 *Strip Packets with FM6000 Timestamp*

Figure 116GigaSMART 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:

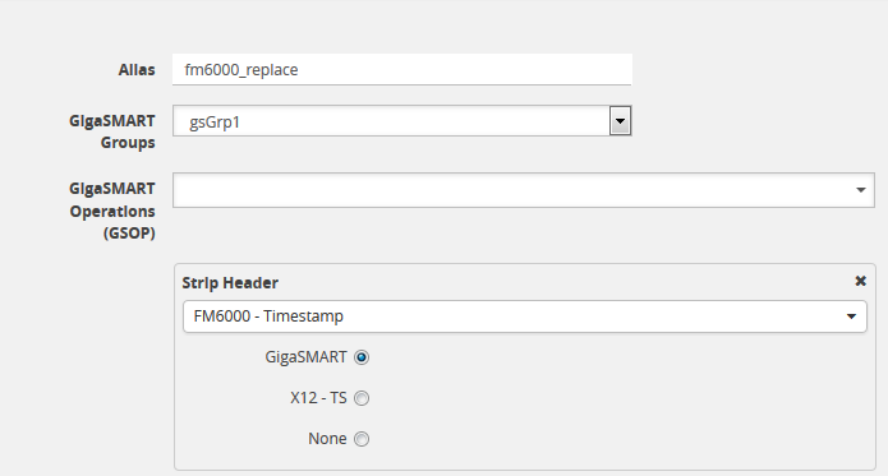

Figure 95 *Strip Packets and Append UTC Timestamp to the Gigamon Trailer*

Figure 96Map with Strip Header GigaSMART Operation is an example map using the strip header GigaSMART operation in Figure 95Strip Packets and Append UTC Timestamp to the Gigamon Trailer.

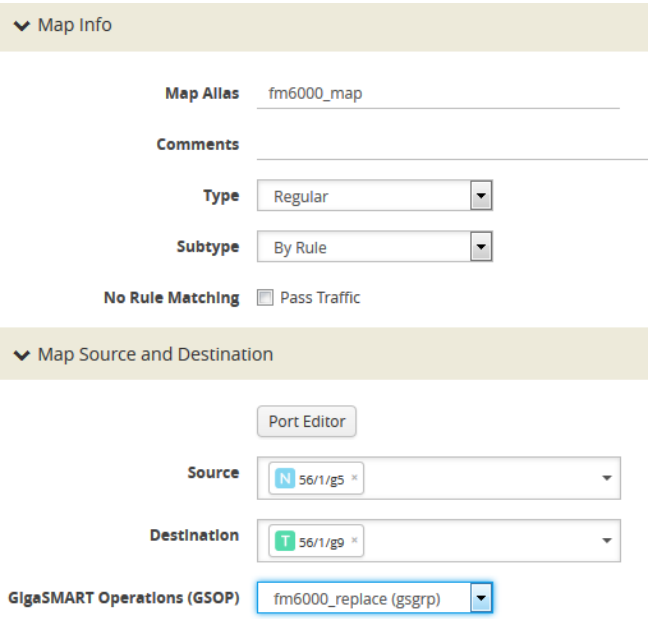

Figure 96 *Map with Strip Header GigaSMART Operation*

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.

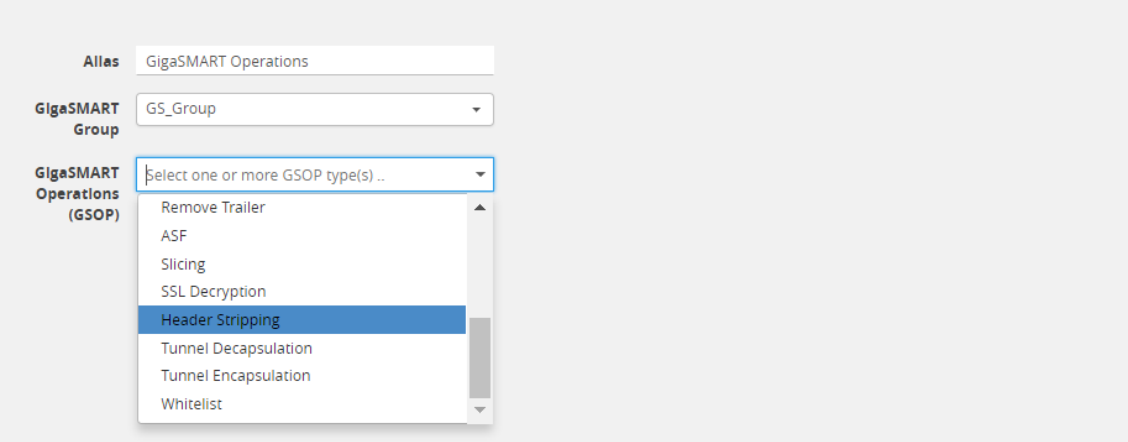

Figure 97 *GigaSMART Header Stripping*

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.

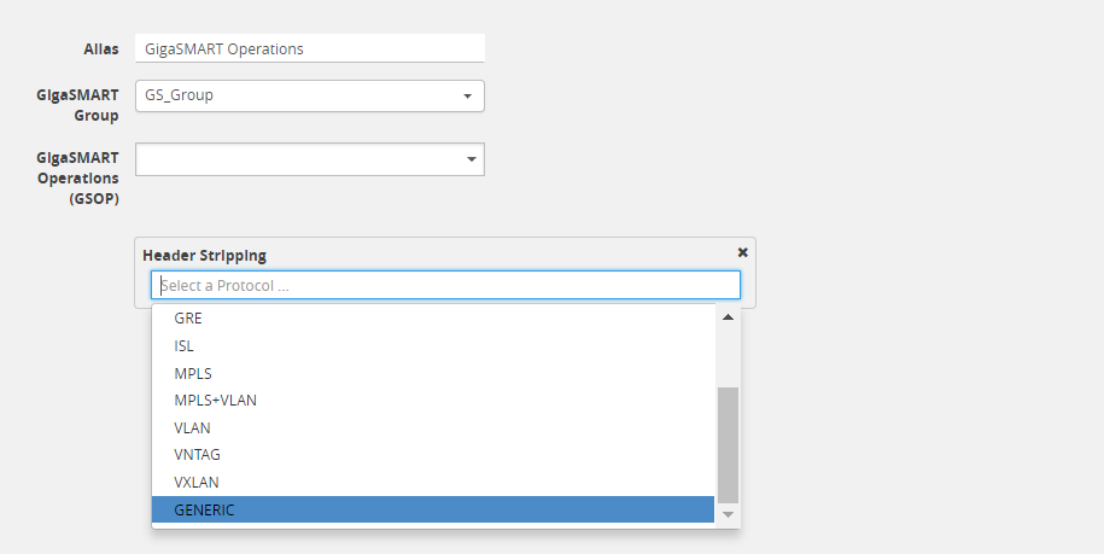

Figure 98 *GigaSMART Generic Header Stripping*

5. From the **Header Stripping** drop-down list, select **Generic**.

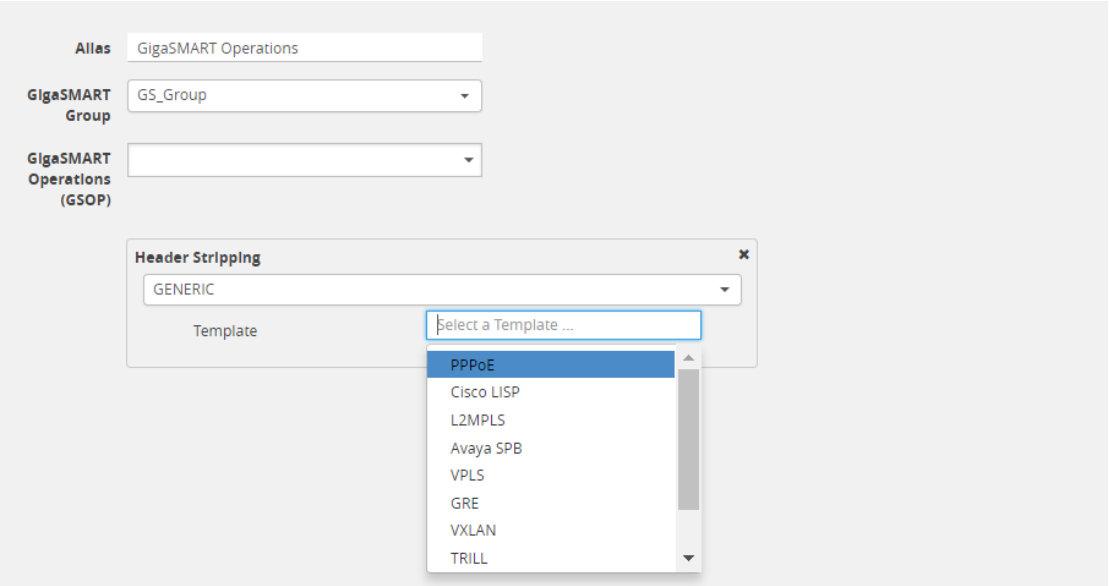

Figure 99 *Generic Header Stripping Template*

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.**

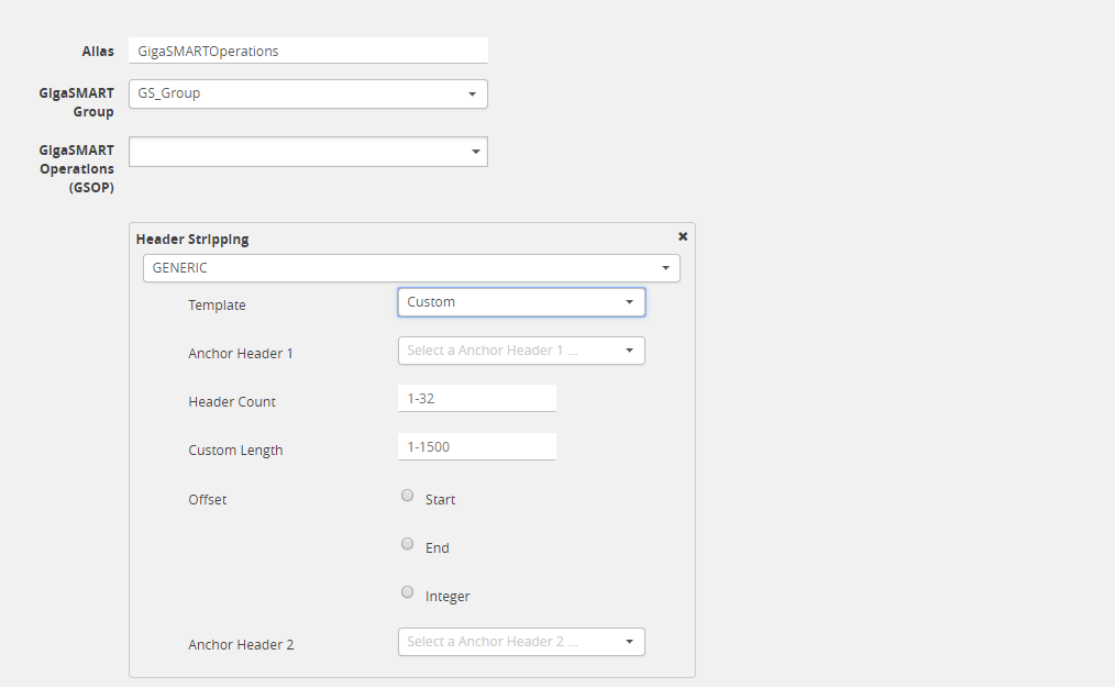

Figure 100*Generic Header Stripping - Custom*

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 101PPPoE Encapsulated Packets illustrates a red outline around the frame that needs to be striped.



Figure 101 *PPPoE Encapsulated Packets*

To strip the PPoE encapsulated packets:

1. From the **Template** drop-down list, select **PPPoE**. Refer to Figure 102Generic Header Stripping - PPPoE.

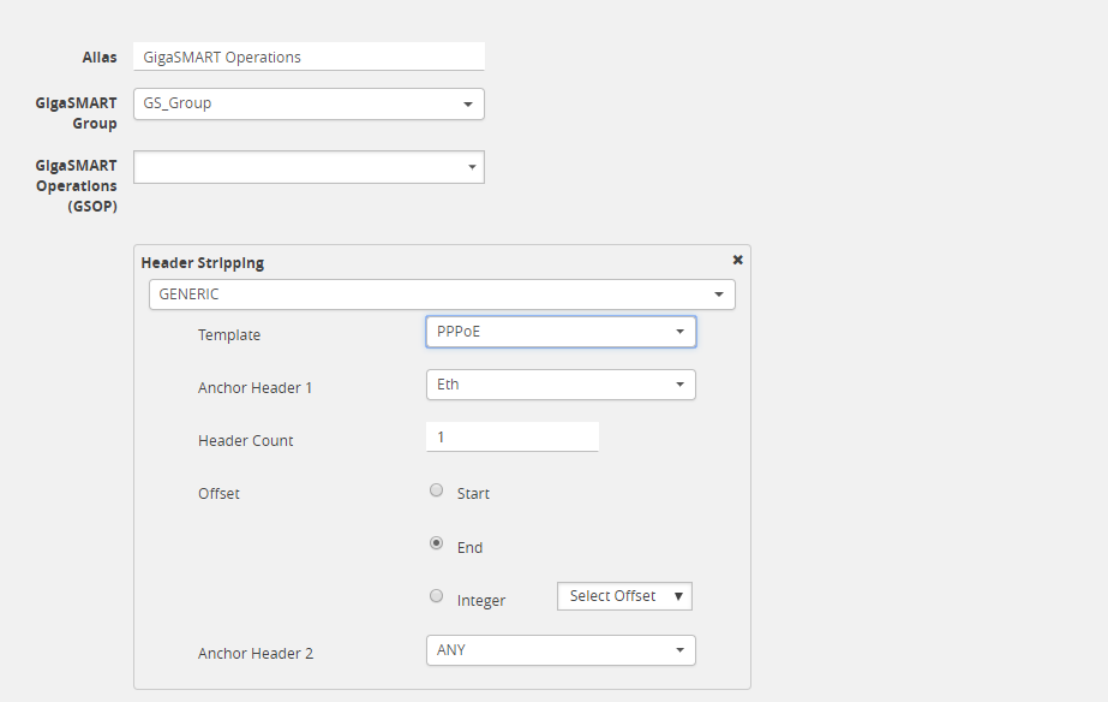

Figure 102*Generic Header Stripping - PPPoE*

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 103Cisco LISP Encapsulated Packets illustrates a red outline around the frame that needs to be striped.



Figure 103*Cisco LISP Encapsulated Packets*

To strip the Cisco LISP encapsulated packets:

1. From the **Template** drop-down list, select **Cisco LISP**.

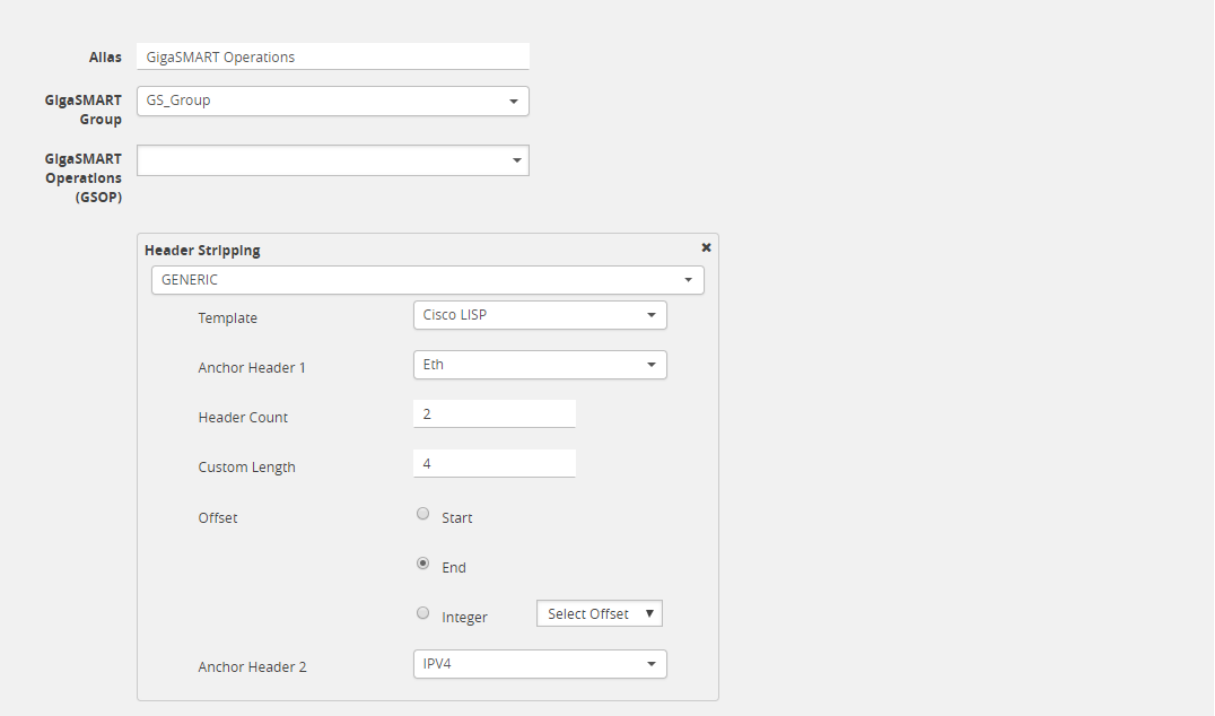

Figure 104*Generic Header Stripping - Cisco LISP*

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 105L2 MPLS Encapsulated Packets illustrates a red outline around the frame that needs to be striped.



Figure 105*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**.

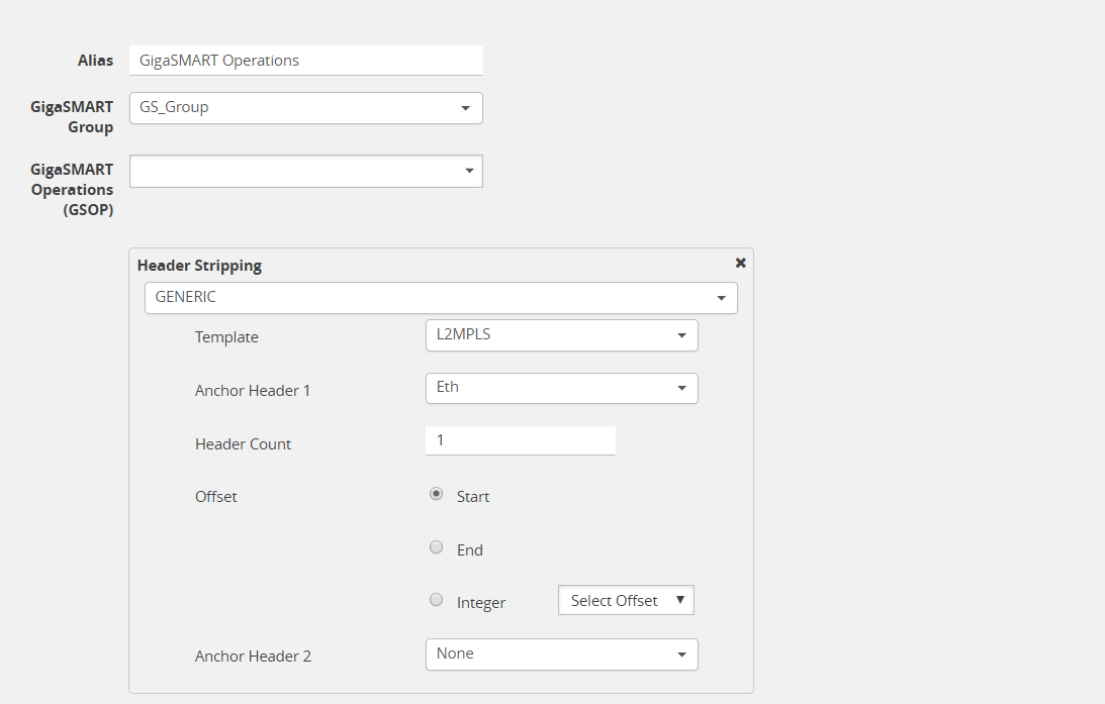

Figure 106*Generic Header Stripping - L2 MPLS*

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 107VXLAN Encapsulated Packets illustrates a red outline around the frame that needs to be striped.



Figure 107*VXLAN Encapsulated Packets*

To strip the outer Ethernet frame from the VXLAN encapsulated packets:

1. From the **Template** drop-down list, select **VXLAN**.

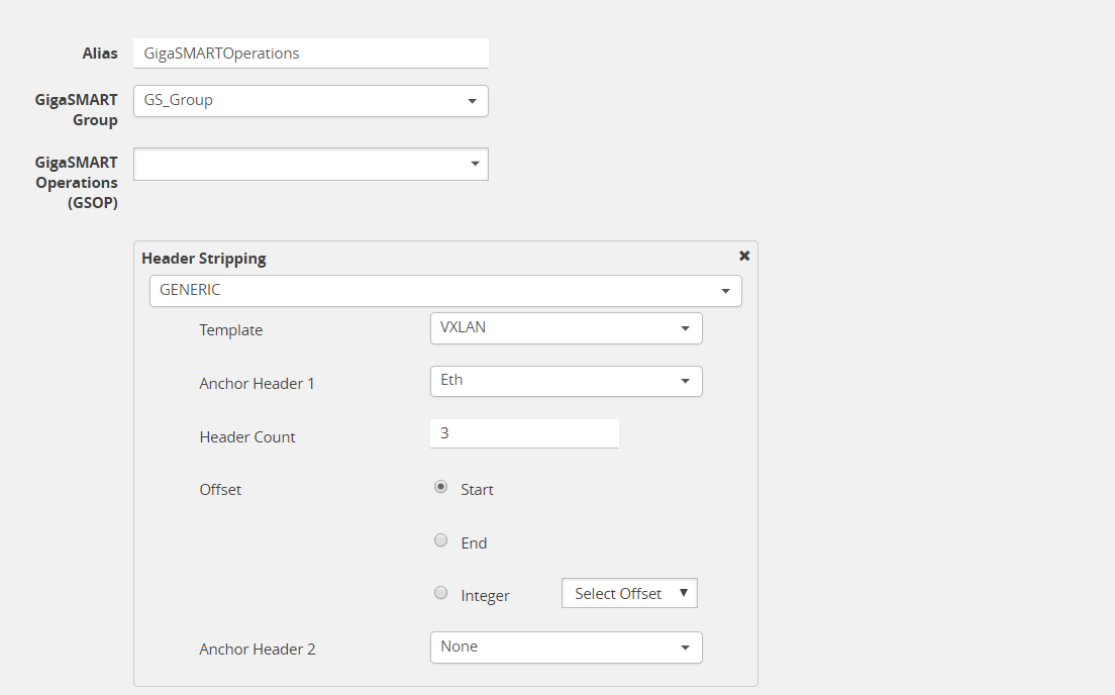

Figure 108*Generic Header Stripping - VXLAN*

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 109TRILL Encapsulated Packets illustrates a red outline around the frame that needs to be striped.



Figure 109*TRILL Encapsulated Packets*

To strip TRILL from the encapsulated packets:

1. From the **Template** drop-down list, select **TRILL**.

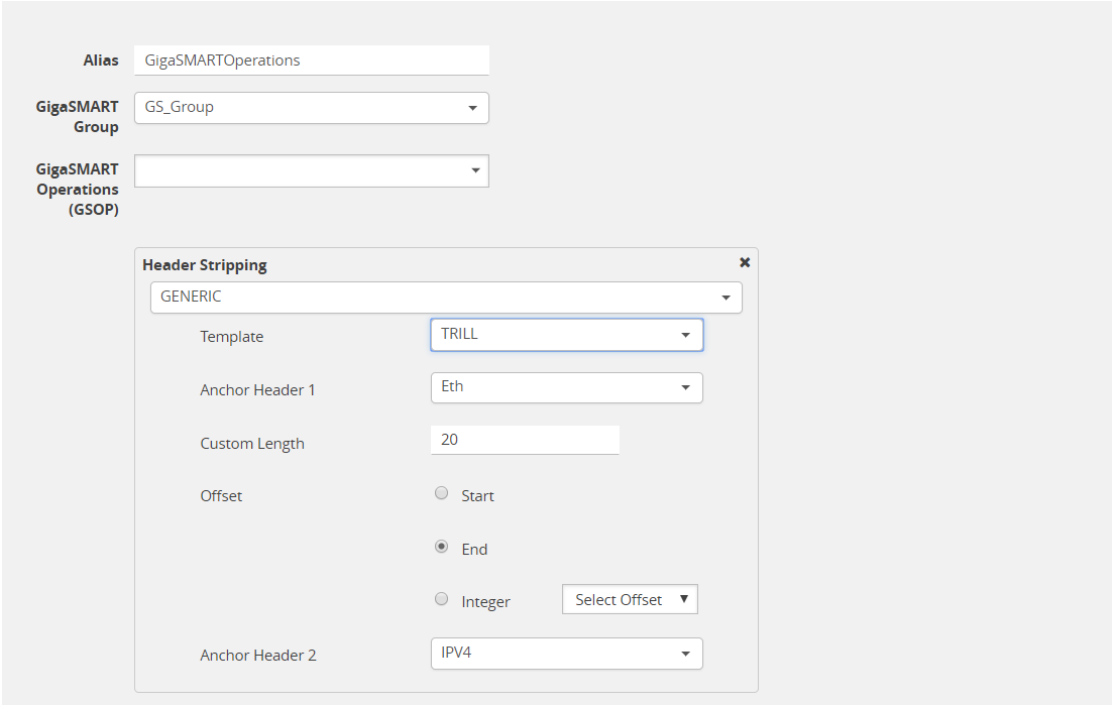

Figure 110 *Generic Header Stripping - TRILL*

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 111Avaya SPB Encapsulated Packets illustrates a red outline around the frame that needs to be striped.

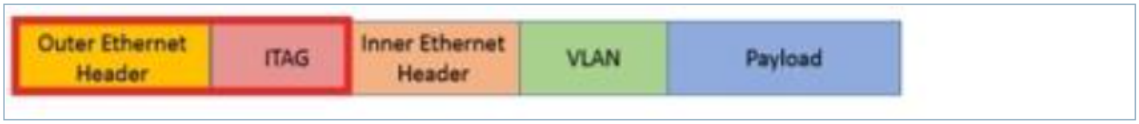

Figure 111 *Avaya SPB Encapsulated Packets*

To strip the outer Ethernet headers from the Avaya SPB encapsulated packets:

1. From the **Template** drop-down list, select **Avaya SPB**.

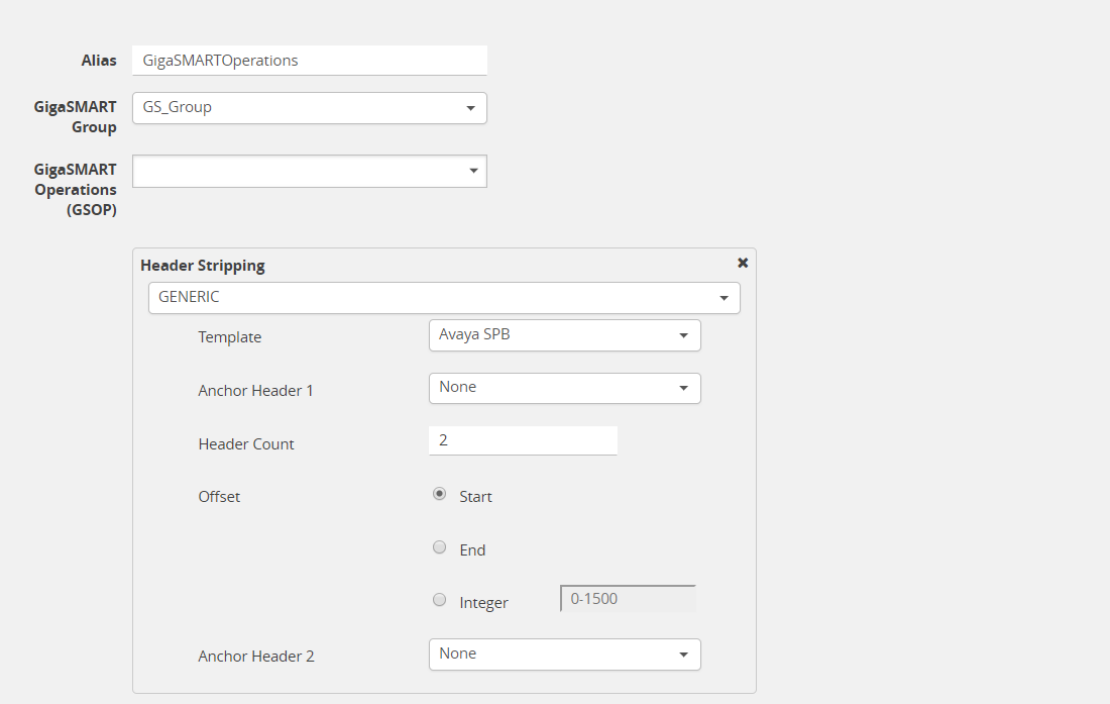

Figure 112 *Generic Header Stripping - Avaya SPB*

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 113Avaya SPB Encapsulated Packets illustrates a red outline around the frame that needs to be striped.



Figure 113 *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 for 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 114GigaSMART Operations Page with Masking Selected.

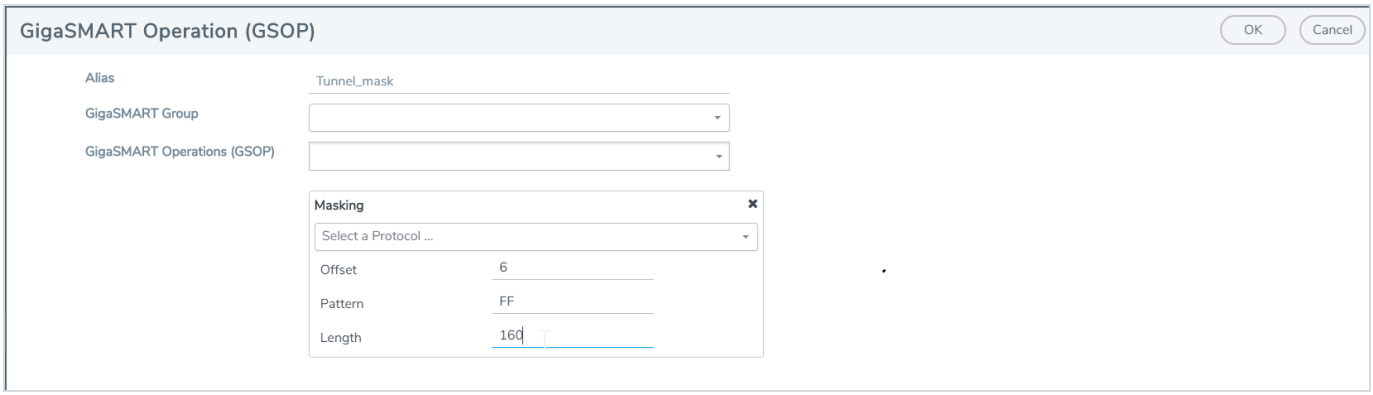

Figure 114 *GigaSMART Operations Page with Masking Selected*

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 for more details on the devices that support GigaSMART operations.

Examples – GigaSMART Masking

The example shown in Figure 115GigaSMART 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.

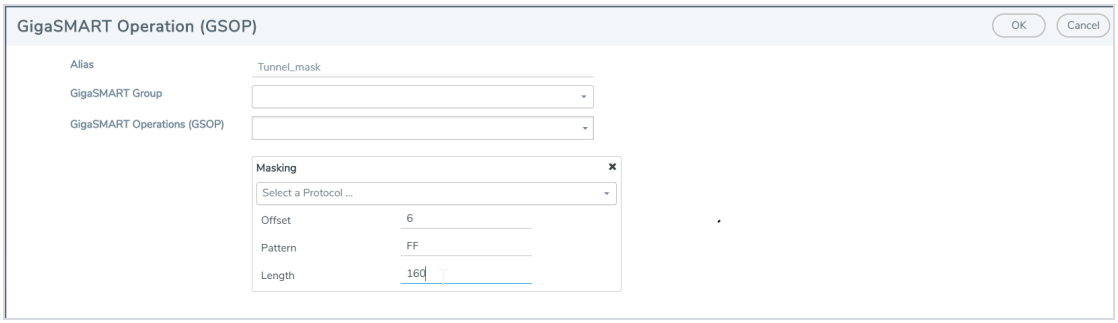

Figure 115 *GigaSMART Masking Operation*

This example shown in Figure 116GigaSMART 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.

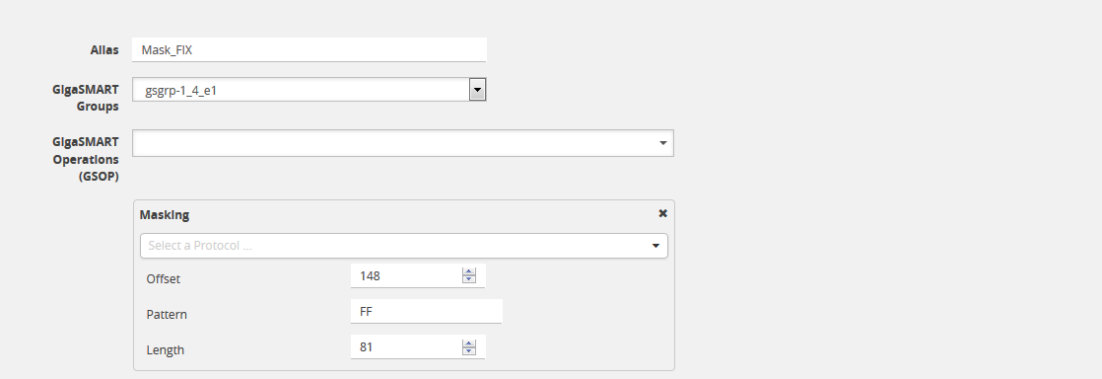

Figure 116 *GigaSMART Operation with a Static Offset*

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 for 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 117NetFlow Generation Gigamon Solution.

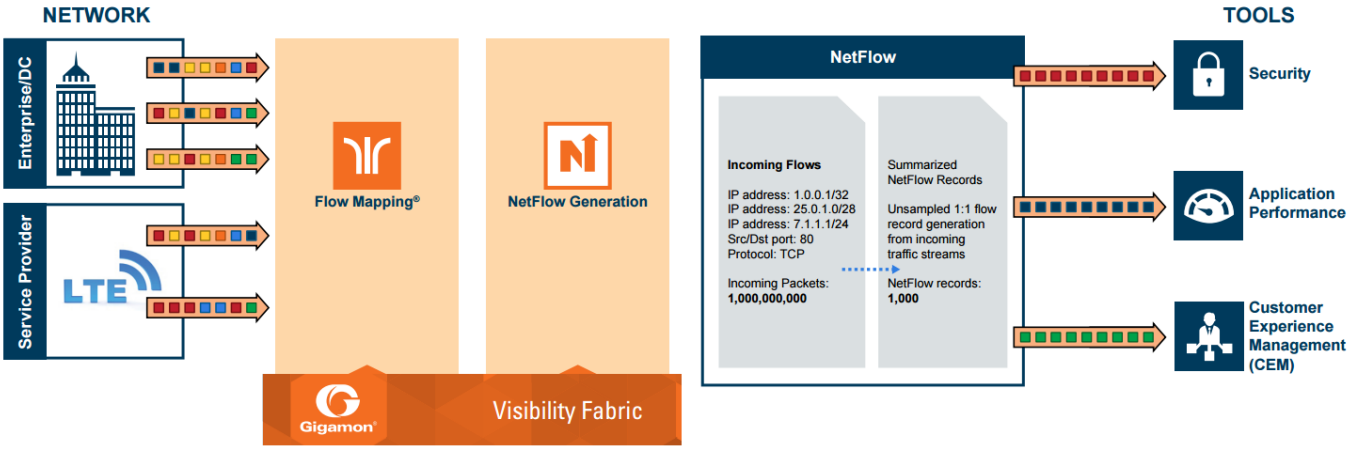

Figure 117 *NetFlow Generation Gigamon Solution*

In Figure 117NetFlow 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.

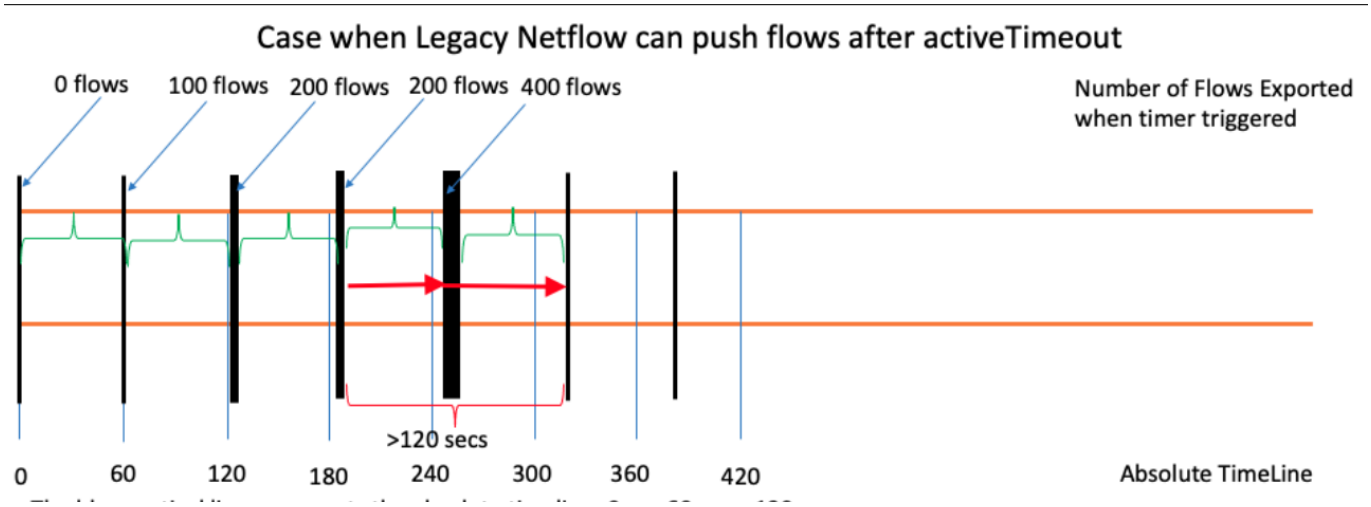

In the above diagram, the blue vertical line represents the absolute time. The black vertical line represents the time at which GigaSMART engine starts exporting data after inactive time out. The green curly brackets represent the active timeout. The width of the black line represents the time taken 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 118NetFlow Generation Components shows the NetFlow Generation components.

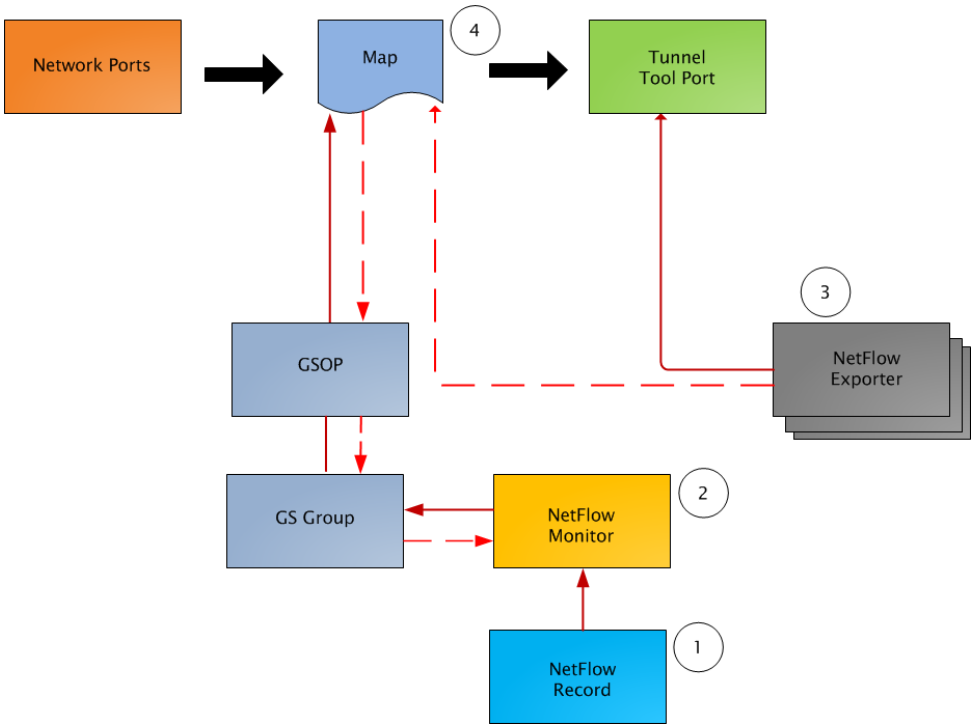

Figure 118 *NetFlow Generation Components*

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 119Collect Field Interface input Name.
5. Click **Save**.

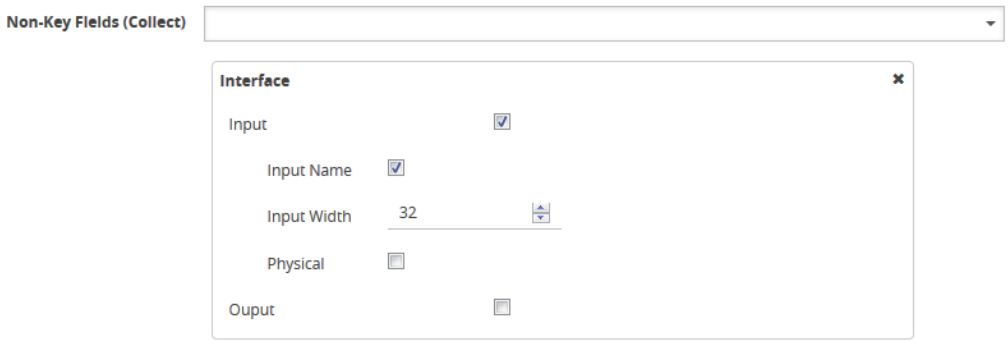

Figure 119 *Collect Field Interface input Name*

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 120NetFlow Record Configuration for HTTP Response Code.

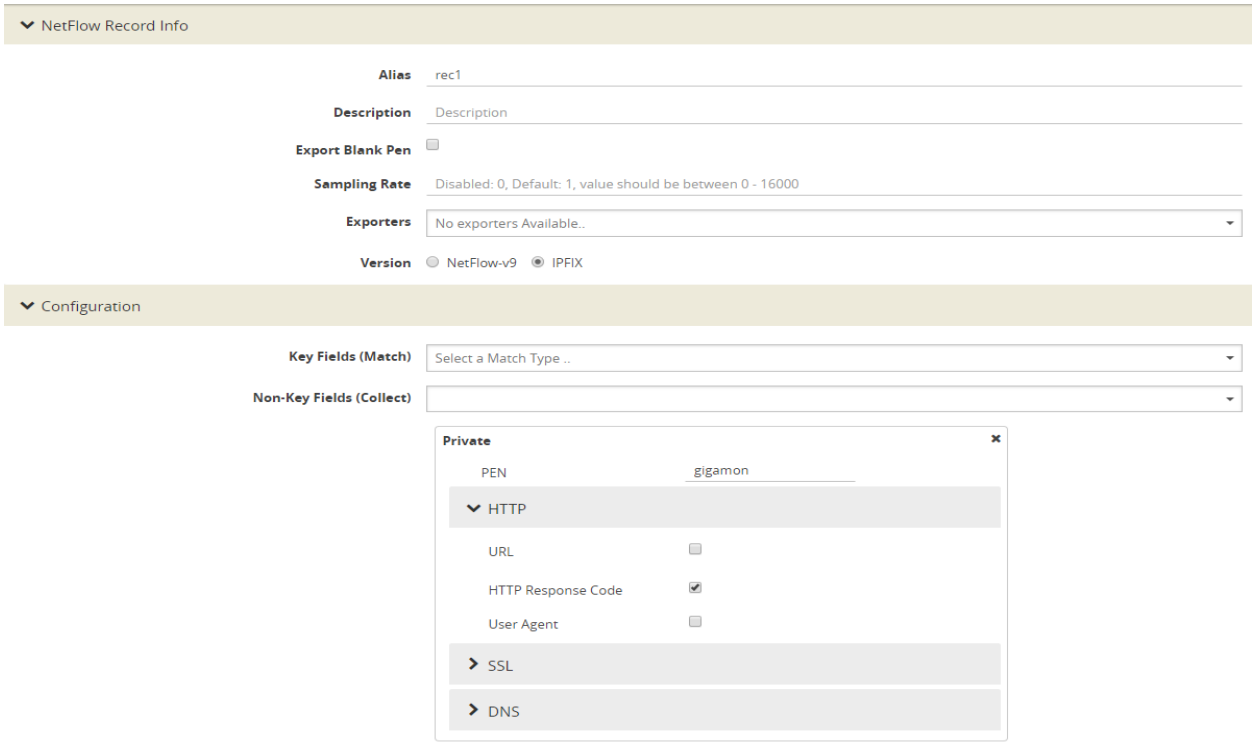

Figure 120 *NetFlow Record Configuration for HTTP Response Code*

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 120NetFlow 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 120NetFlow 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 121NetFlow Record Configuration for DNS.

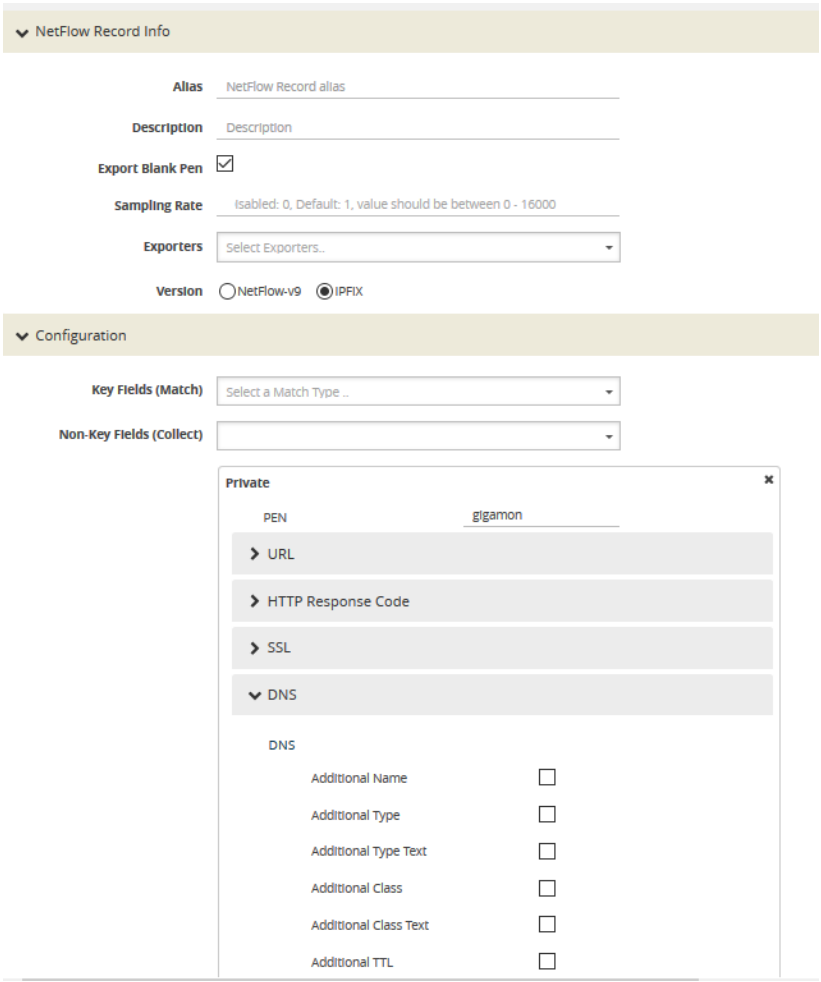

Figure 121 *NetFlow Record Configuration for DNS*

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 20: 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 20: 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 122NetFlow Record Configuration for SSL.

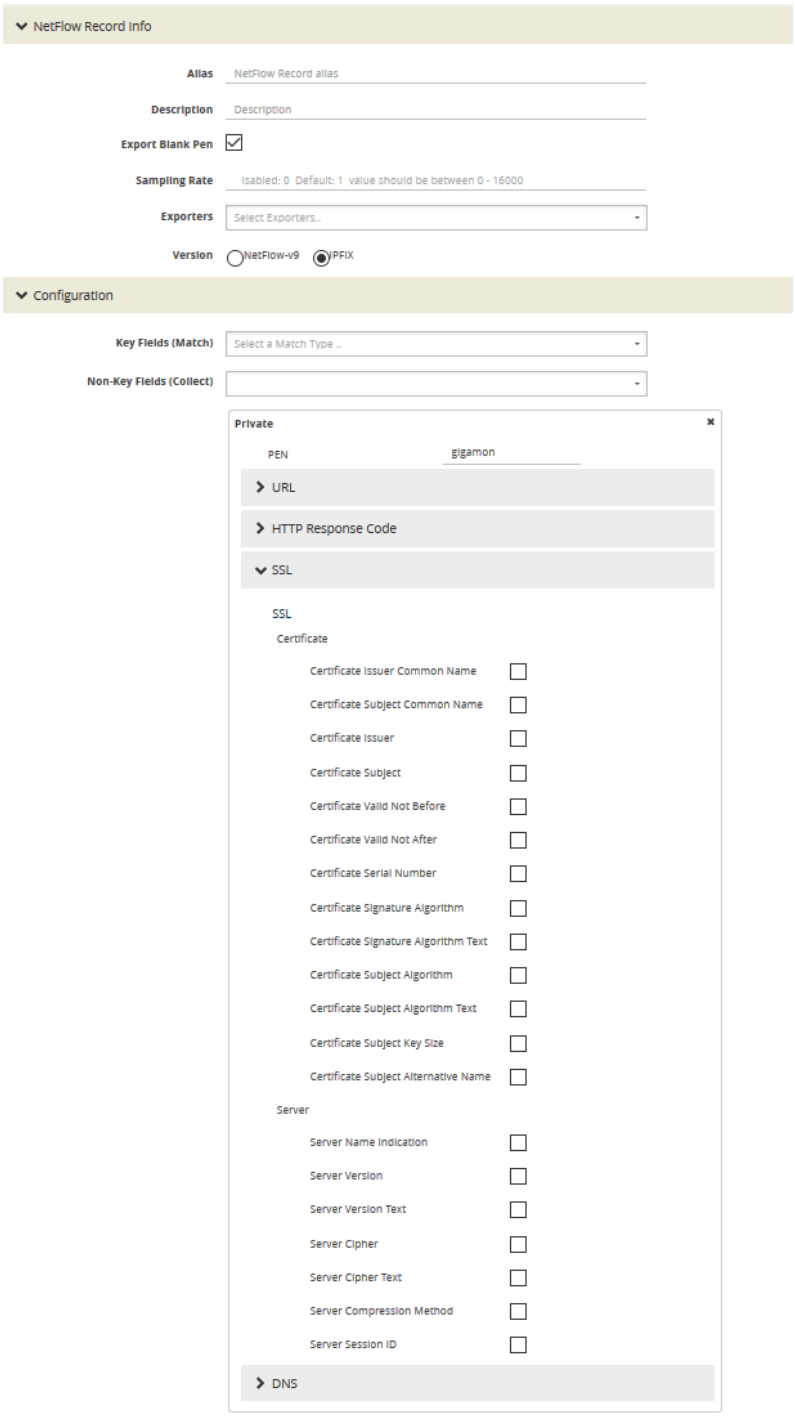

Figure 122 *NetFlow Record Configuration for SSL*

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=31373031303630303030303 05a
GigamonMdataSslSerialNo=0118ee3c2167b99e1b718c6eadb8fb4d00000000
GigamonMdataSslValidNotAfter=32303031313531323030303 05a
GigamonMdataSslCertSigAlgo=2a864886f70d01010b
GigamonMdataSslCertSubAlgo=2a864886f70d010101
GigamonMdataSslCertSubKeySize=270 GigamonMdataSslServerVersion=771
GigamonMdataSslCertSubAltName=*.stickyadstv.com
GigamonMdataSslServerCompressionMethod=192 GigamonMdataSslServerCipher=49199
GigamonMdataSslServerVersionText=TLSv1.2 GigamonMdataSslServerSessionId=63
GigamonMdataSslIssuer=2f433d55532f4f3d446967694365727420496e632f4f553d7777772e64
69676963
6572742e636f6d2f434e3d446967694365727420534841322048696768204173737572616e636
52053657276 6572204341 GigamonMdataSslCertSubCommonName=*.stickyadstv.com
GigamonMdataSslSub=2f433d55532f53543d4e657720596f726b2f4c3d4e657720596f726b2f4f3
d4672656
```

5776865656c204d6564696120496e632f4f553d46726565776865656c2f434e3d2a2e737469636b796164737 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 123NetFlow Record predefined_netflow_v5_record.

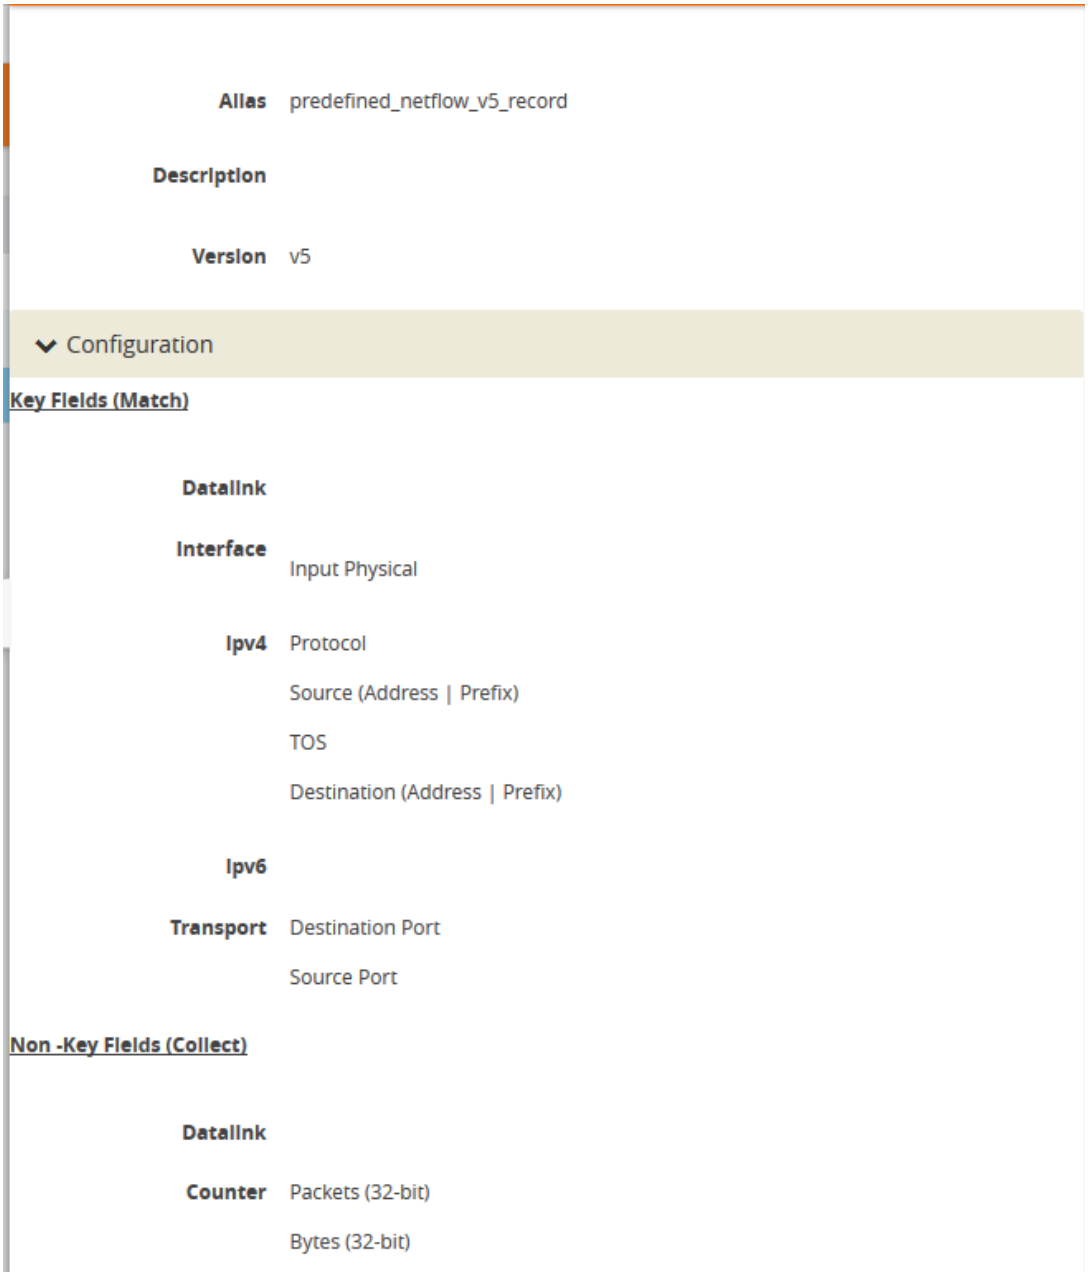

Figure 123 *NetFlow Record predefined_netflow_v5_record*

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.

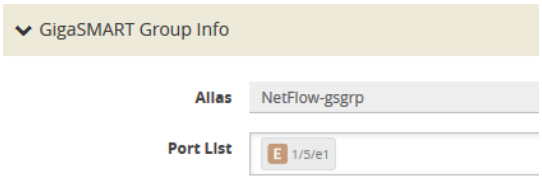

5. Click **Save**.

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 21: 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 21: 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 124NetFlow Record Page displays.

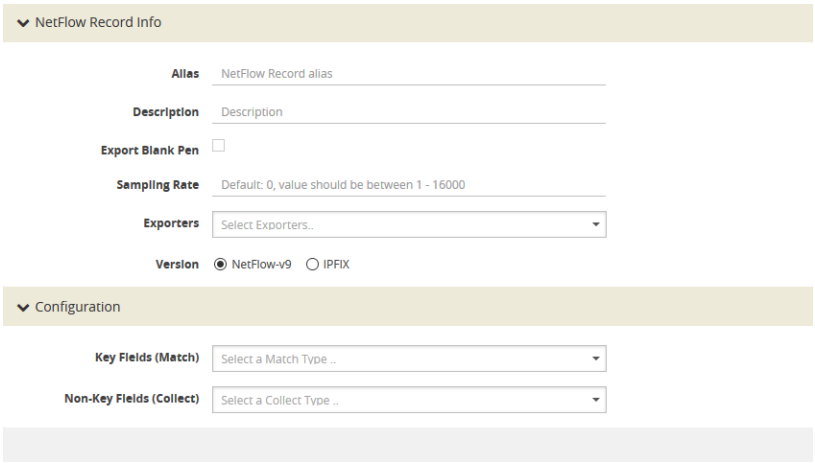

Figure 124*NetFlow Record Page*

3. On the NetFlow Record page, do the following:

Table 22: 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 23: NetFlow Monitor Parameters.
   d. Configure the **Sampling** parameters. Refer to Table 23: NetFlow Monitor Parameters.
   e. Select the **Record** that you configured in Step 4: Configure the NetFlow Record.
4. Click **Save**.

Table 23: 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.

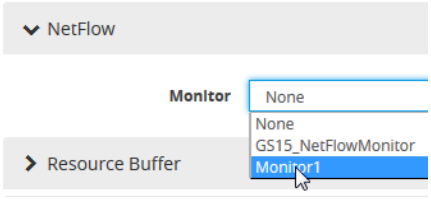

4. Click **Save**.

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 125GigaSMART 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 125GigaSMART Operation (GSOP) Page.

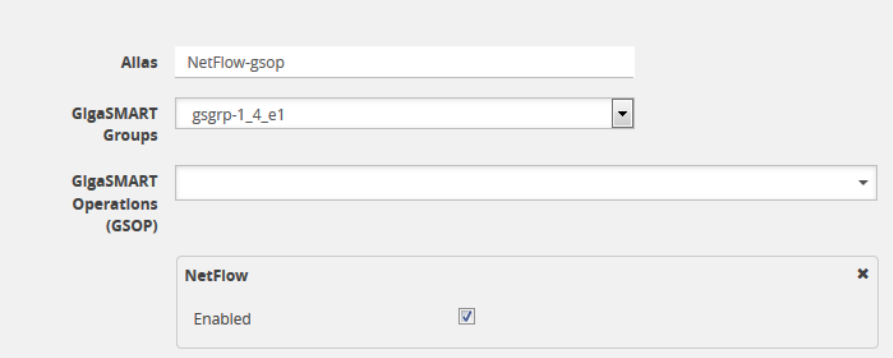

Figure 125 *GigaSMART Operation (GSOP) Page*

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 for 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.

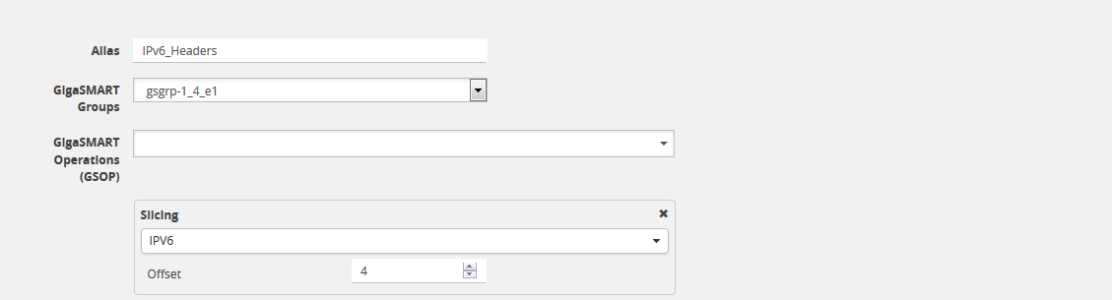

Figure 126 *GigaSMART Operations (GSOP) Page with Slicing Selected*

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 for 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.

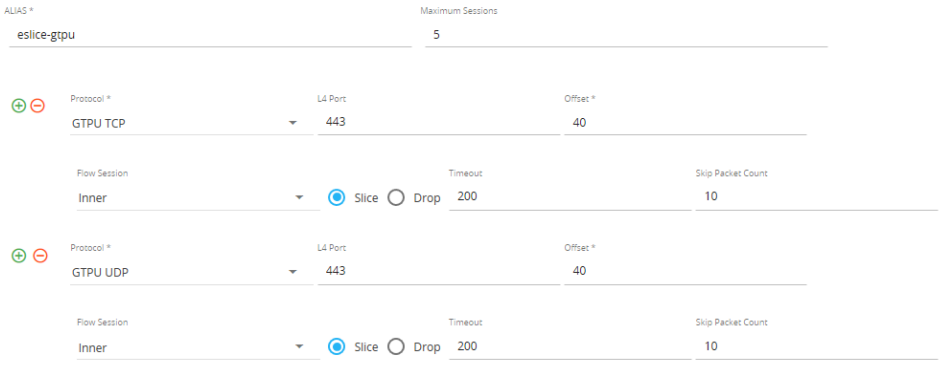

Figure 127 *Advanced Flow Slicing*

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 Inline and 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:

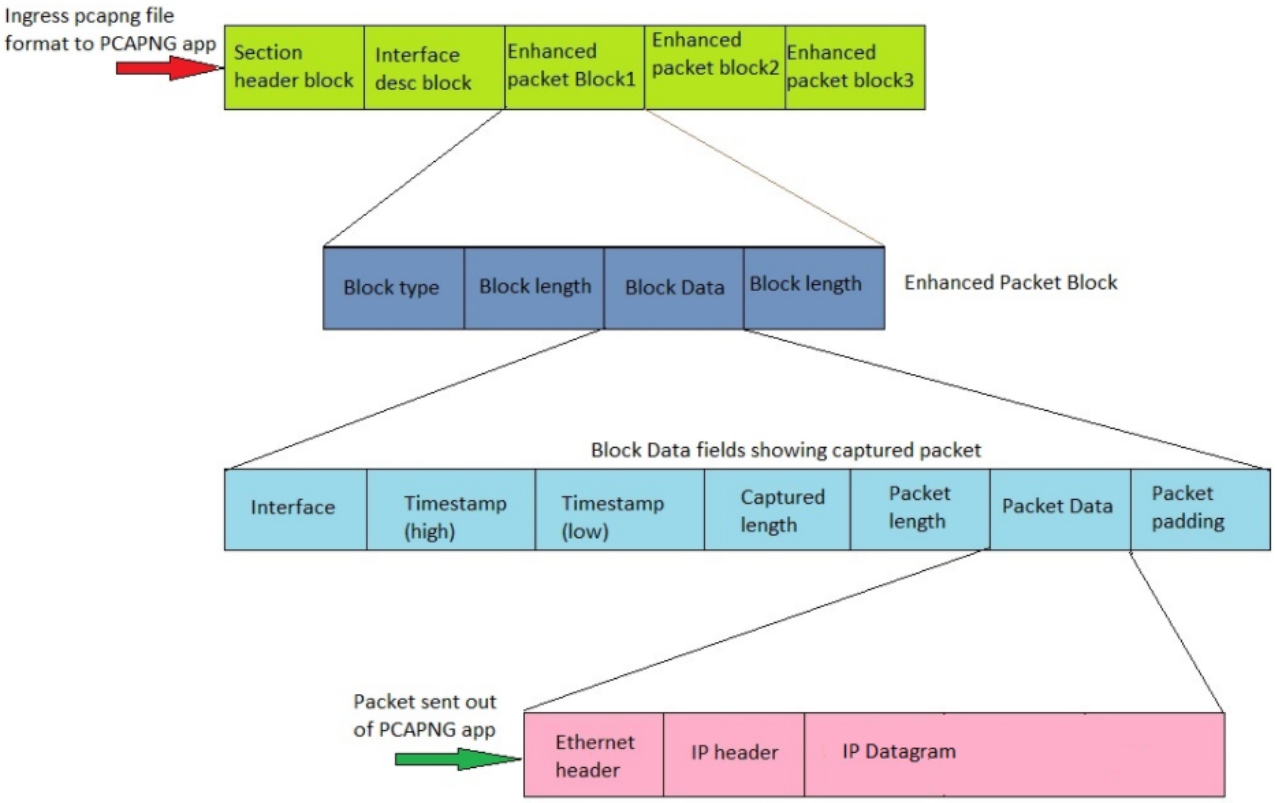

Tunneling Operations

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-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**

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-VM)

The following figures demonstrate how to create a sample IP tunnel between a sending GigaVUE-VM node in Sydney and a receiving GigaVUE H Series in Melbourne. First, the overall tunnel is summarized, followed by configuration descriptions for the sending and receiving ends.

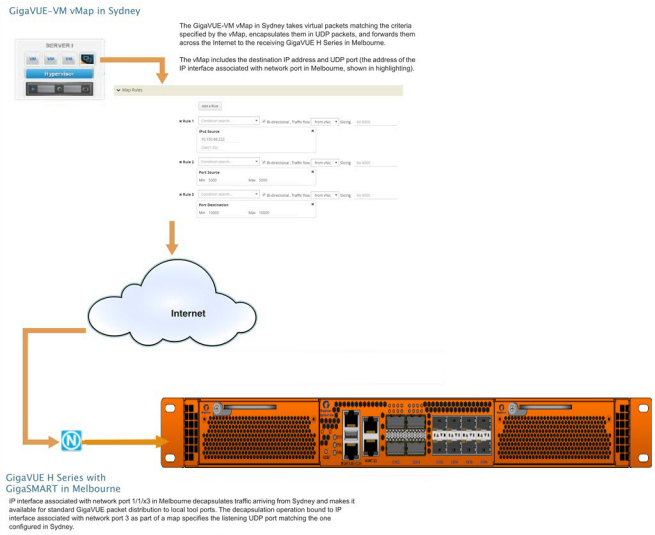

Configure Sending End of Tunnel: GigaVUE-VM vMap in Sydney

A GigaVUE-VM node in this location is configured with a vMap that will send data over the Internet to the IP interface associated with a network port on a GigaVUE H Series with a GigaSMART-HD0 line card installed.

VMaps 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 vMap "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 H 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**

Task	UI Steps
	▪ 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 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.	**1.** From the device view, select GigaSMART > GigaSMART Operations > GigaSMART Operation. **2.** Click New. **3.** Configure the GigaSMART Operation: ▪ 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.

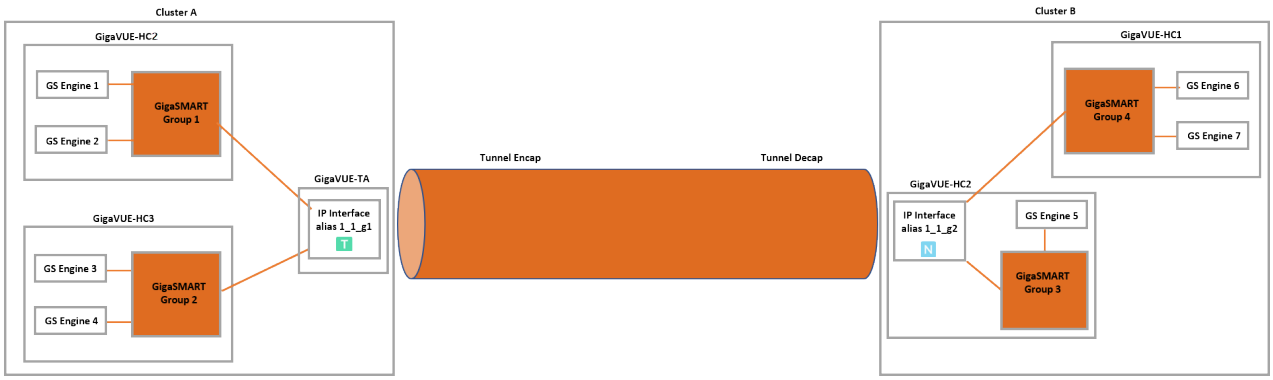

Figure 128 *Tunnel Centralization*

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 the IP Interfaces for the 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.

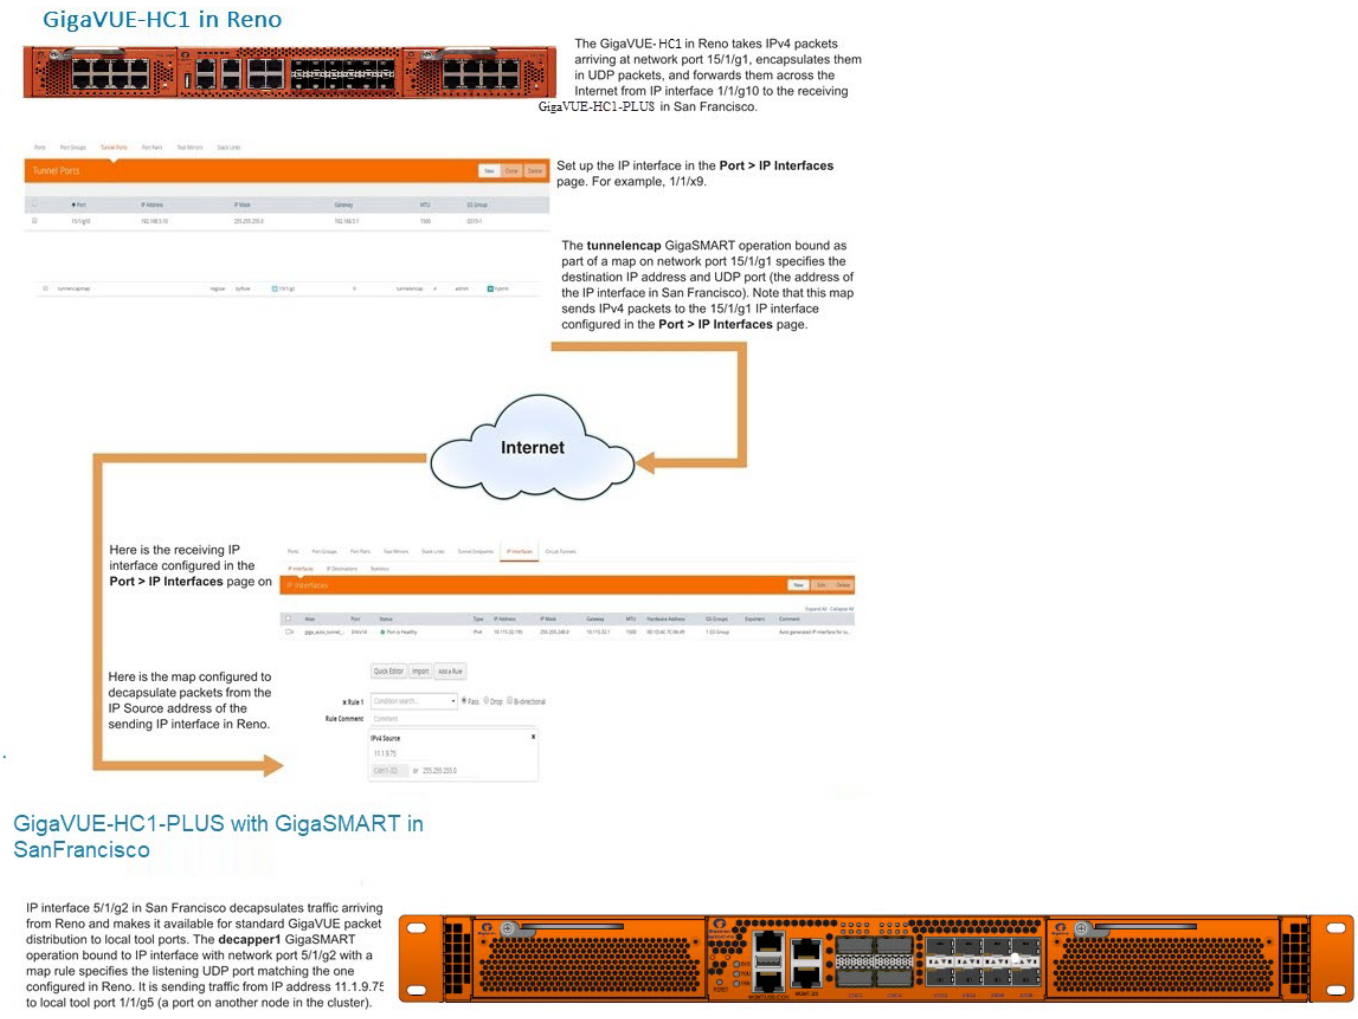

Configure Sending End of Tunnel: GigaVUE-HC1 in Reno

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 ▪ GigaSMART Group: gsport1 **4.** Click OK.

Task	UI Steps
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

Refer to Supported GigaSMART Operations for 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 129L2GRE 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.

NOTE: IPv6 addresses are not supported on SMT-HC1-S (Generation 3 GigaSMART module on GigaVUE-HC1.

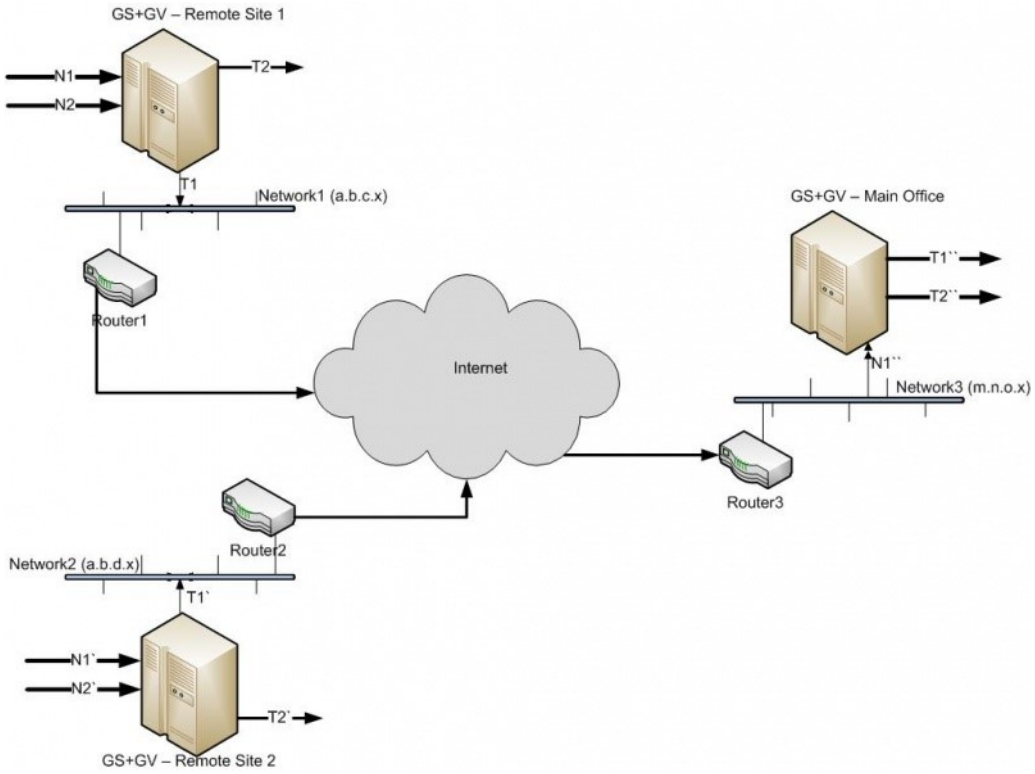

Figure 129 *L2GRE Tunnel Encapsulation/Decapsulation*

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.

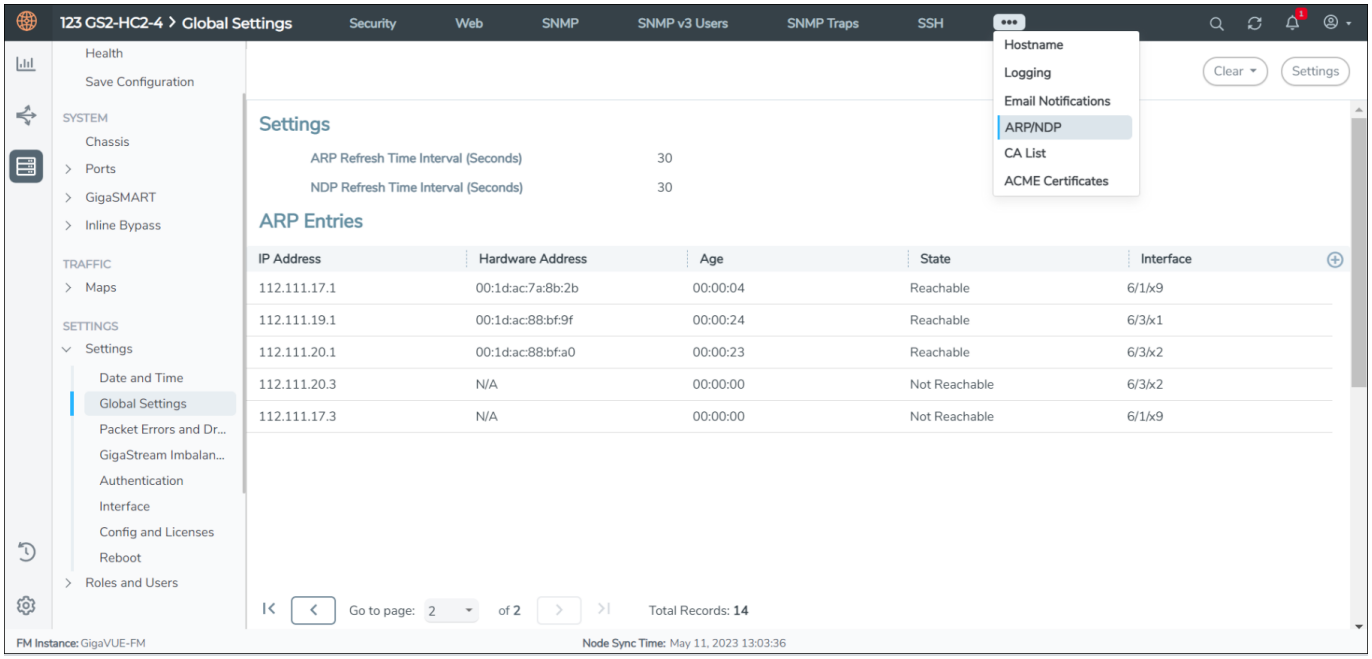

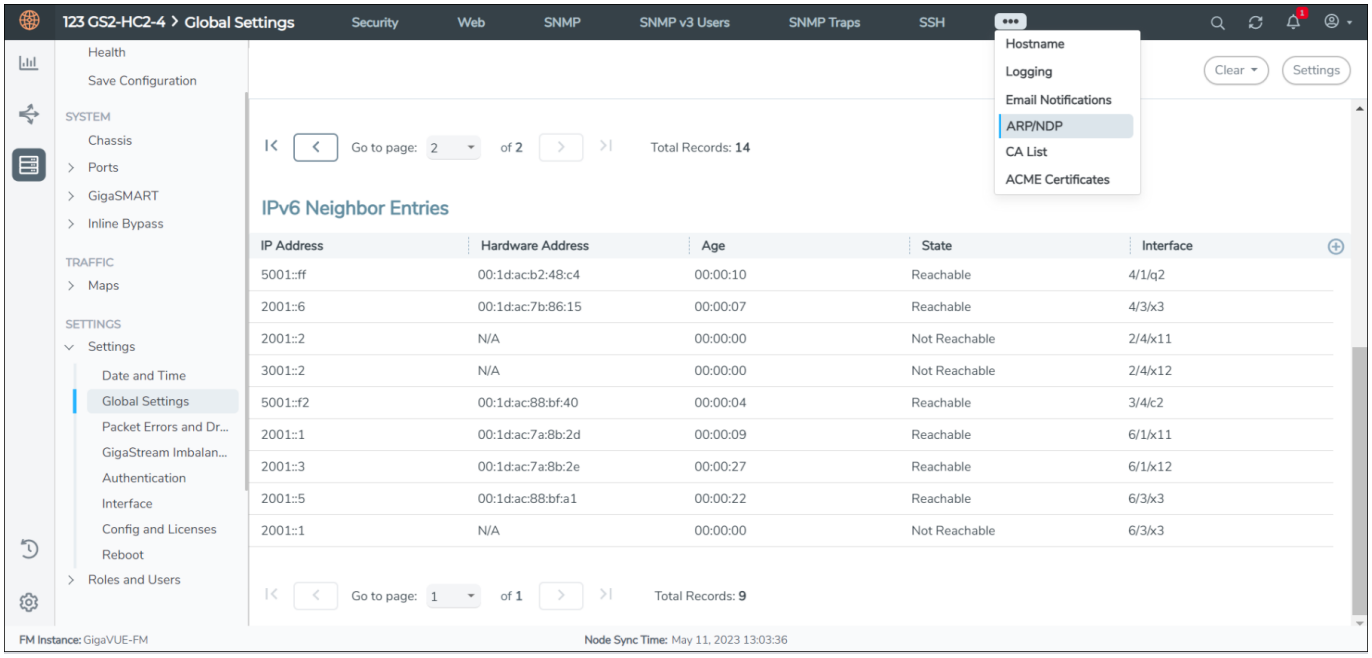

For a configuration example, refer to GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation

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.

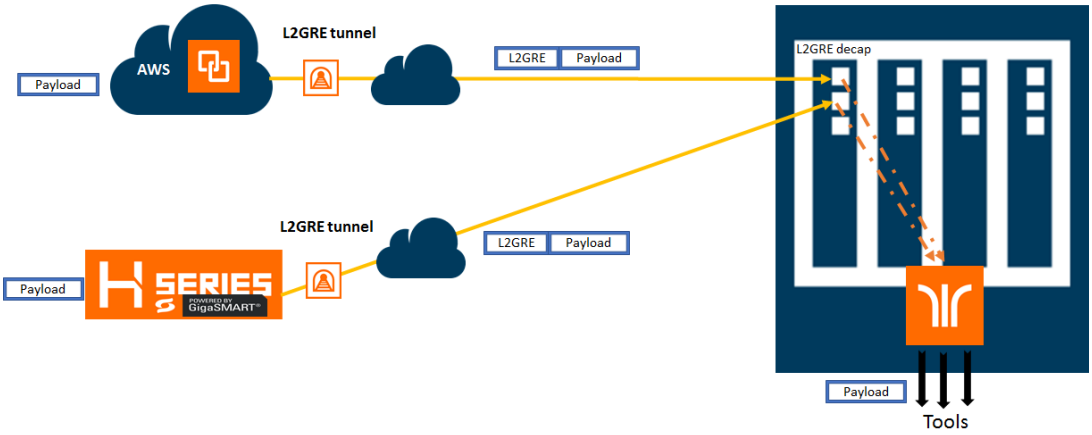

Configure L2GRE Tunnel Encapsulation and Decapsulation

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 130New GigaSMART Operation for Stateful Load Balancing.

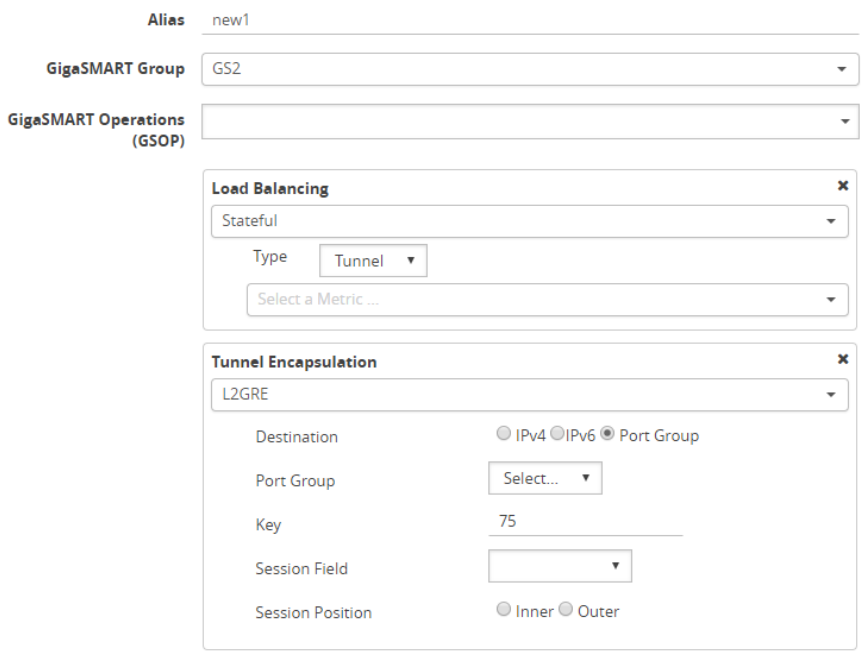

Figure 130*New GigaSMART Operation for Stateful Load Balancing*

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 131New Map Configuration.

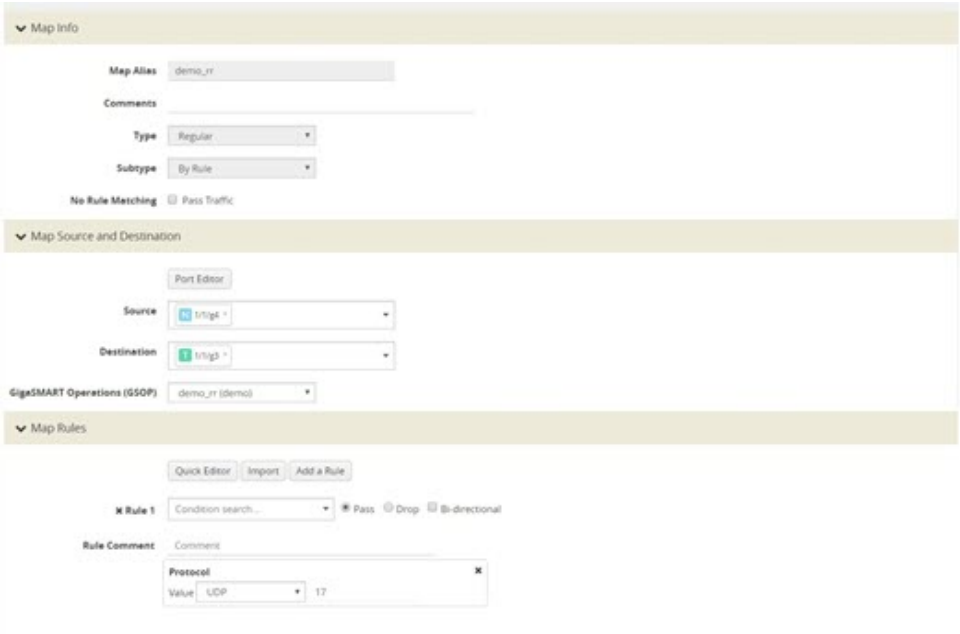

Figure 131 *New Map Configuration*

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 132New GigaSMART Operation for Stateless.

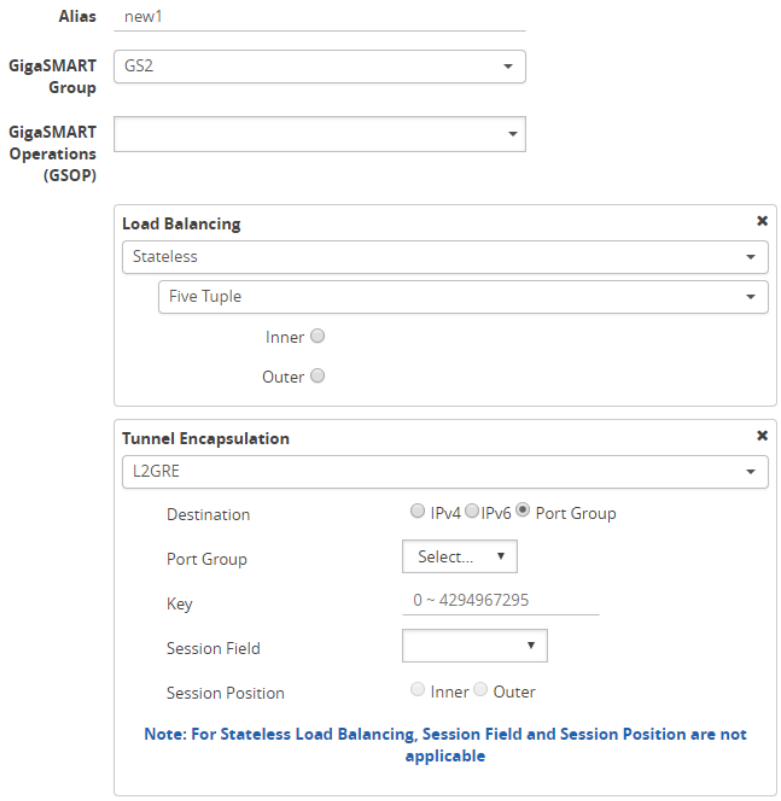

Figure 132 *New GigaSMART Operation for Stateless*

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**.

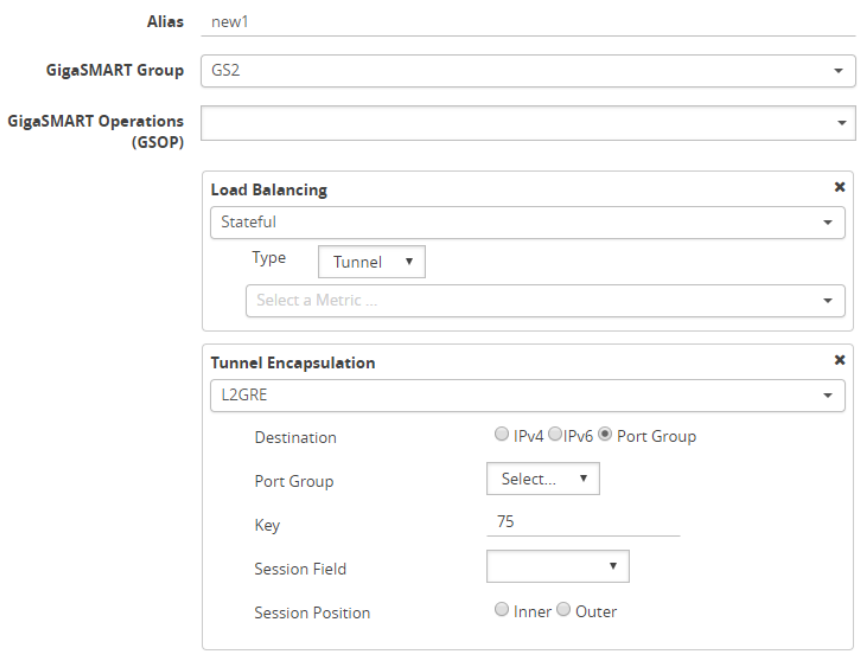

Figure 133 *GigaSMART Operation (GSOP)*

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**.

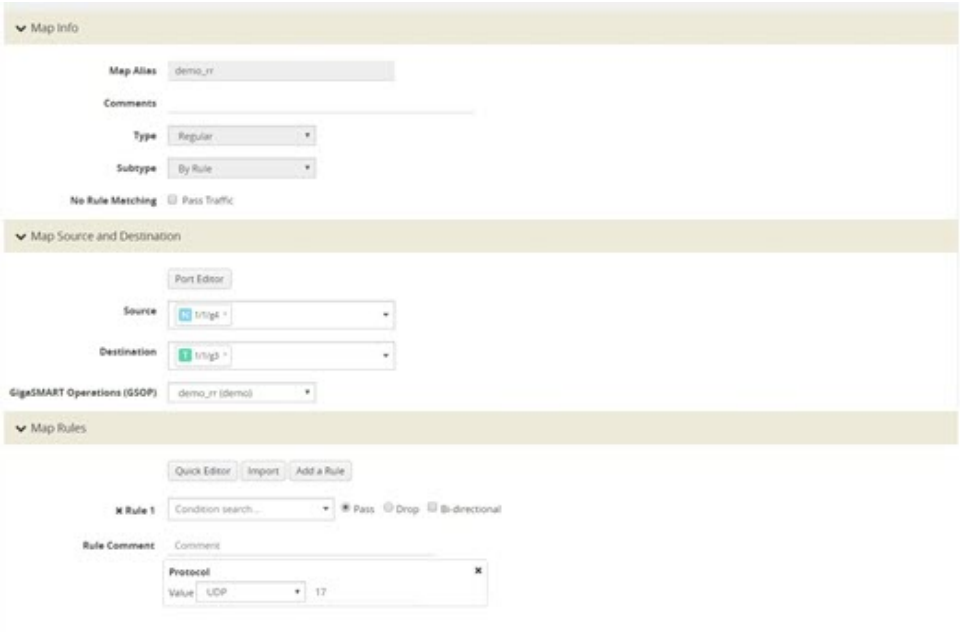

Figure 134*Map Configuration*

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 for 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.

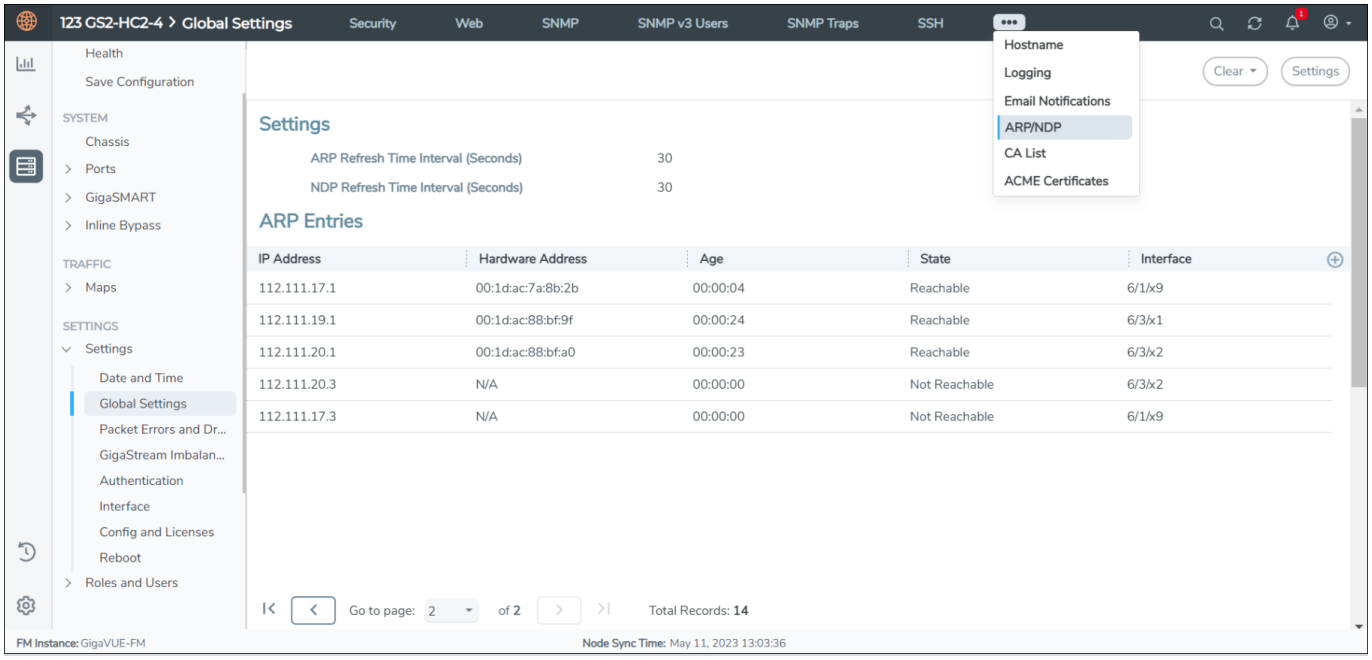

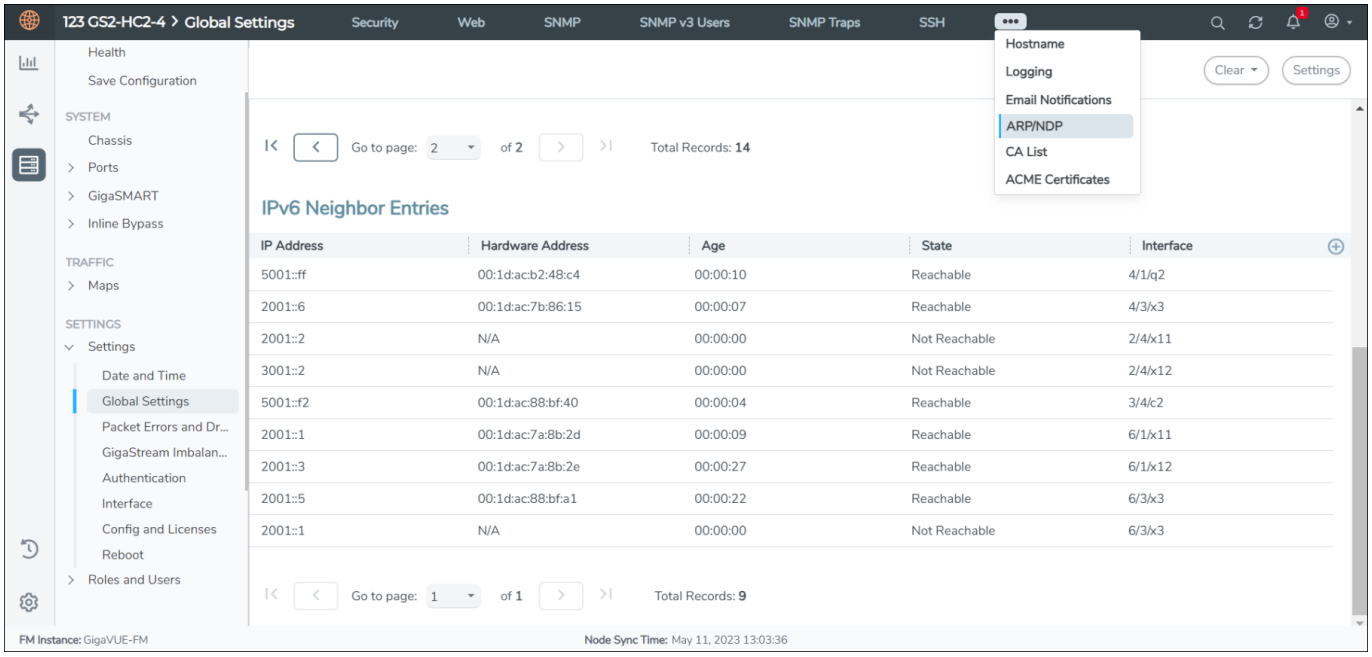

VXLAN IPv6 Encapsulation/Decapsulation

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 use the Port Editor window.

Field	Description
	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 for 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 135ERSPAN 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.

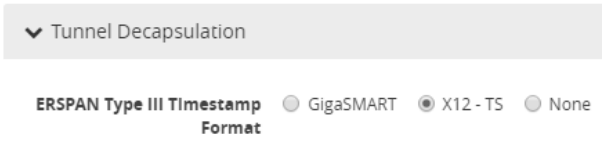

Figure 135 *ERSPAN Type III Timestamp Formats on GigaSMART Groups Page*

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 136Decapsulation Flow ID Set to Zero..

NOTE: A flow ID of zero is a wildcard value that matches all flow IDs.

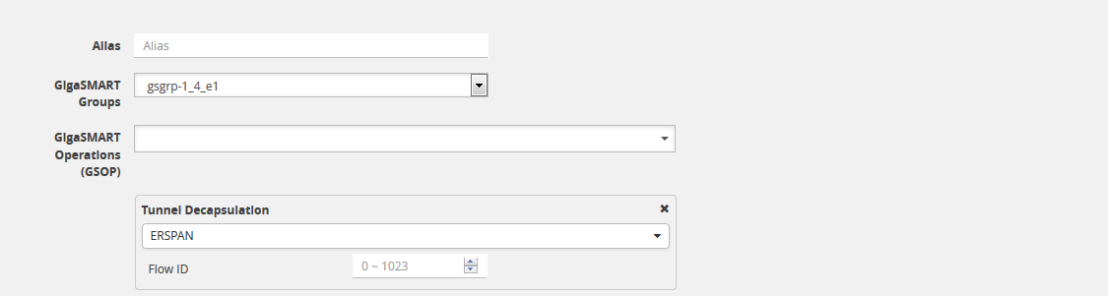

Figure 136 *Decapsulation Flow ID Set to Zero.*

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 136Decapsulation 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 136Decapsulation 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
		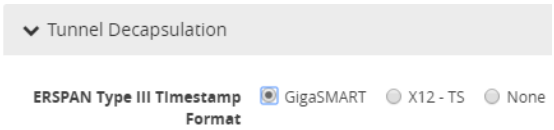 **e.** Click Save.
6.	Create a map. The map contains a rule to allow marker packets (UDP) to be processed.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type an alias in the Map Alias field that will help you identify this map. **d.** Select Regular and By Rule for the map type and subtype. **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 for 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.

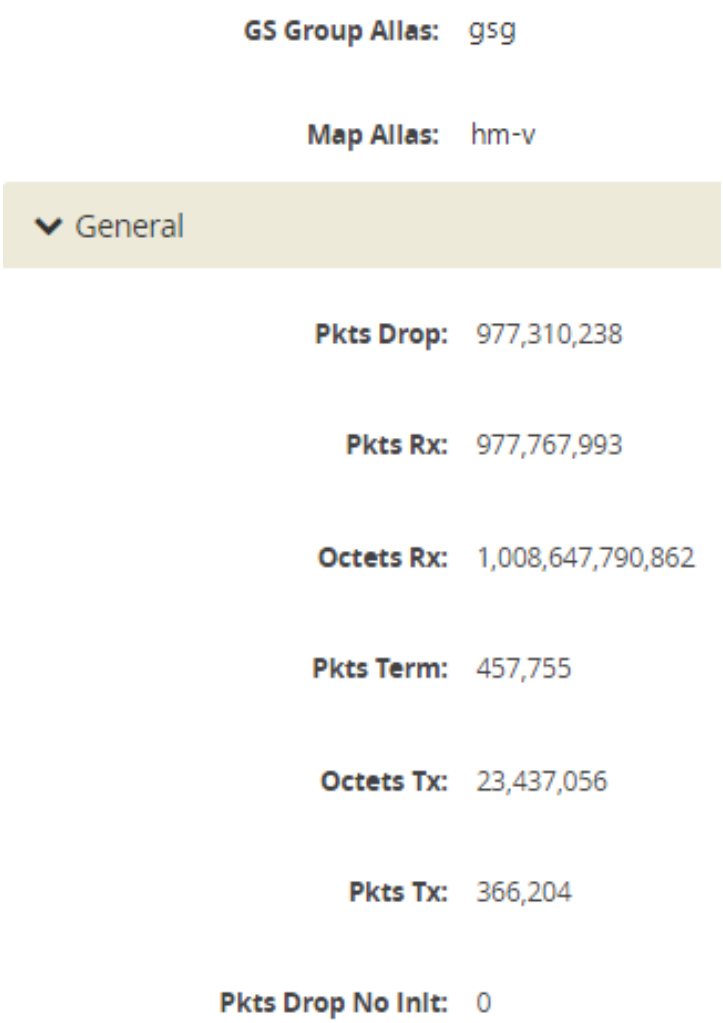

Figure 137 *GSOP VXLAN Statistics Quick View*

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.

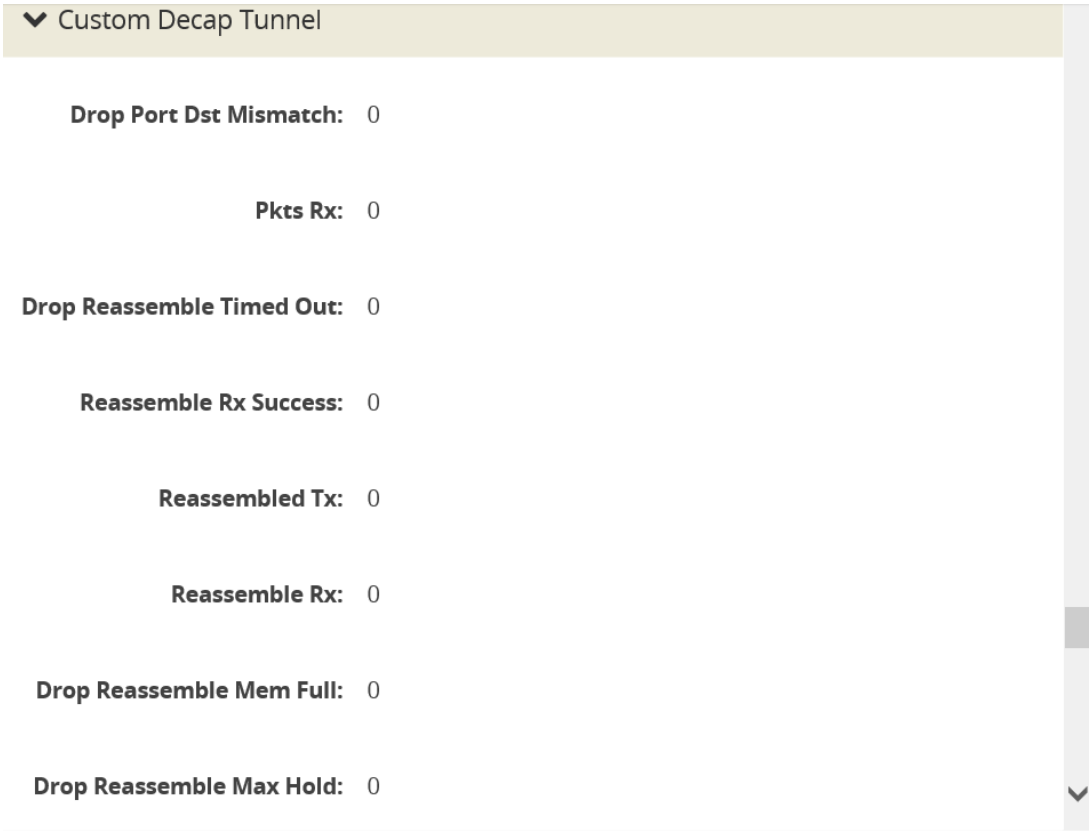

Figure 138

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.

GRE-In-UDP Tunnel Decapsulation

Required License- Advanced Tunneling license on GigaVUE-HC3. Referred to as “Tunneling license” on GigaVUE-HC1.

The GRE-In-UDP tunnel is a new tunnel type that supports decapsulating the GRE-in-UDP headers in the vTAP captured data packets. The decapsulation of the tunnel headers is

common for both control and data packets. However, the control packets (with JSON metadata format) are dropped based on the GRE Key value '0'. The data packets (either fragmented or unfragmented) are handled as follows:

S. No	Process	Description
1	**Reassembly of fragmented data packets (if fragmented packets are received)**	The fragmented packets are reassembled based on standard IP header-based reassembly using the following fields: • More Fragments (MF) field in IPv4 packets • Identification field in IPv6 packets
2	**Decap tunnel lookup**	The ingress packet fields are parsed to find a matching Tunnel End Point (TEP) based on the following header fields: • Destination IP • Source IP • UDP port numbers • GRE key
3	**Decapsulation of headers**	If a matching TEP is found, the encapsulation headers in the packet are removed and the packet is forwarded to the PCAPng application. Refer to the PCAPng Application section for more details.

NOTE: Refer to the "Add Applications to Monitoring Session" topic in the respective Cloud (GigaVUE V Series Solution) Guides for detailed information.

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.

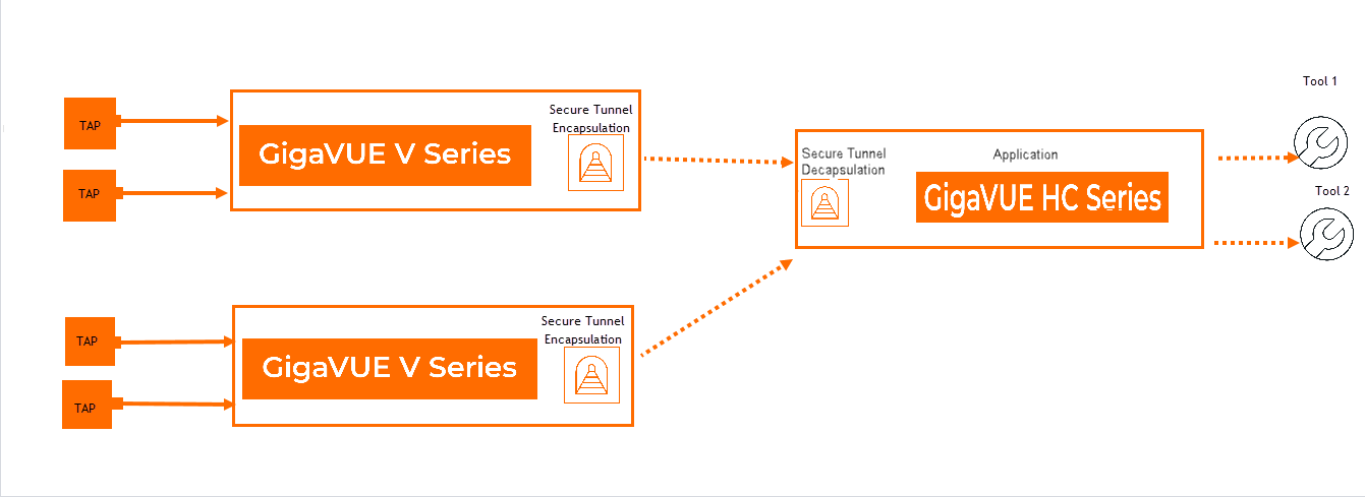

Figure 139 *GigaVUE HC Series with a single IP Interface and multiple Map rules*

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.

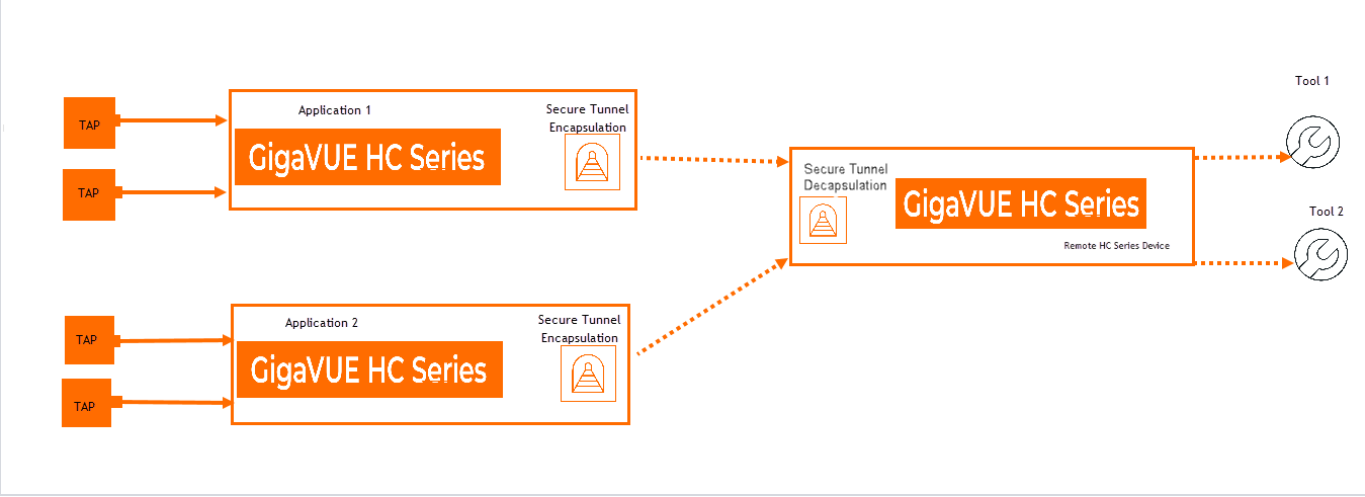

Secure Tunnels between GigaVUE HC Series to a GigaVUE V Series.

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.

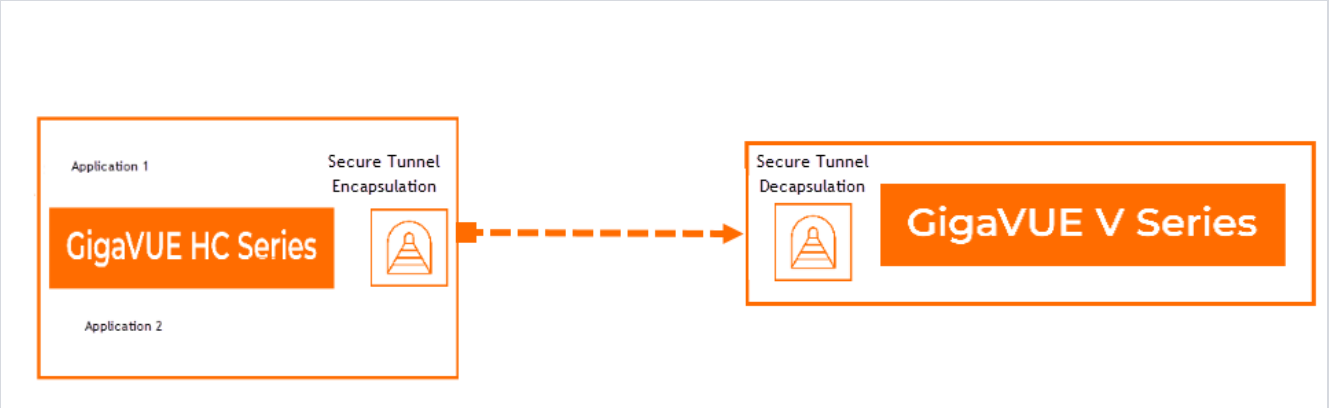

Supported Platforms

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.

Field	Description
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 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

Field	Description
	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.
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.

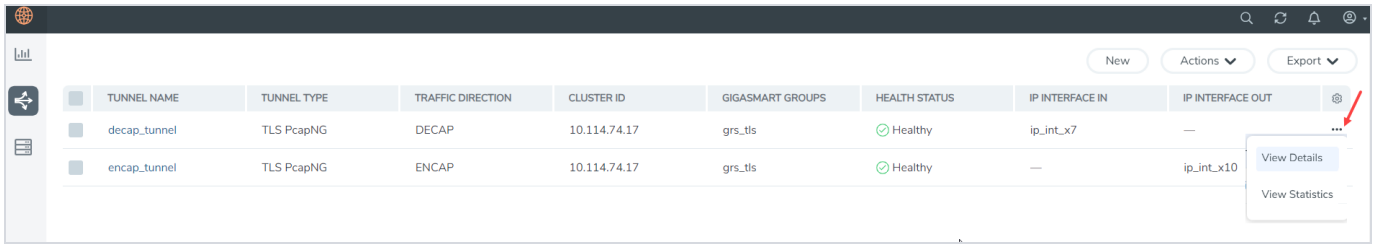

NOTE: The Tunnels page accessible here at **Traffic>Physical> Orchestrated Flows> Tunnels** is solely for configuring Secure tunnels in GigaVUE HC Series and is an Early Access page. To learn more about configuring other Tunneling operations, refer to the Tunneling Operations section.

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).

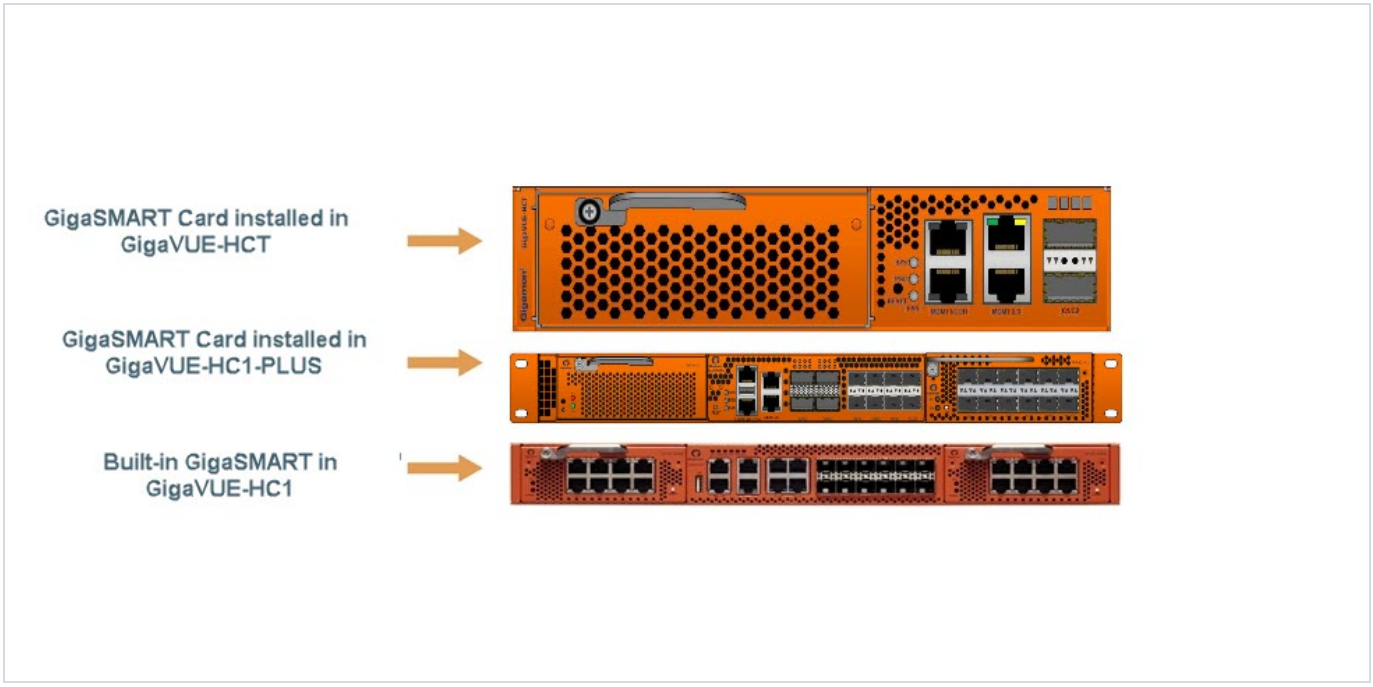

Refer to the following sections for details:

- 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
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	• Configure Inline TLS/SSL Decryption Using GigaVUE-FM • 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

GigaSMART Application	Refer to

De-duplication	• De-duplication Configuration Steps
	  ○ Configure GigaSMART Parameters for Packet De-duplication
	  ○ Example – GigaSMART De-duplication

Flow Masking	• GigaSMART Encapsulated Traffic Performance Enhancement
Header Stripping	• GigaSMART Header Stripping • Generic Header Stripping
Masking	• GigaSMART Masking
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

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**.)

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-C05 module on GigaVUE-HC3 can process packets at up to **100Gb**.
- 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-C05 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.

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:

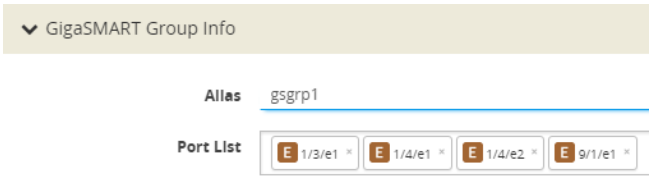

How to Use GigaSMART Operations – Example

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.

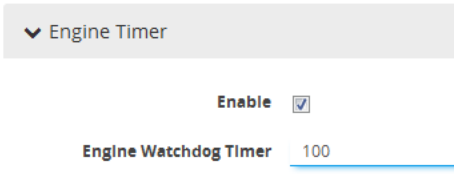

Figure 140*Engine Watchdog Timer*

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 141Configure ICMP Health Check Parameters.

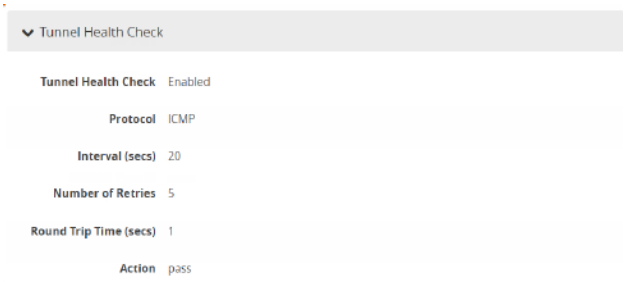

Figure 141 *Configure ICMP Health Check Parameters*

To configure UDP health check parameters for the GigaSMART group, refer to Figure 142Configure UDP Health Check Parameters.

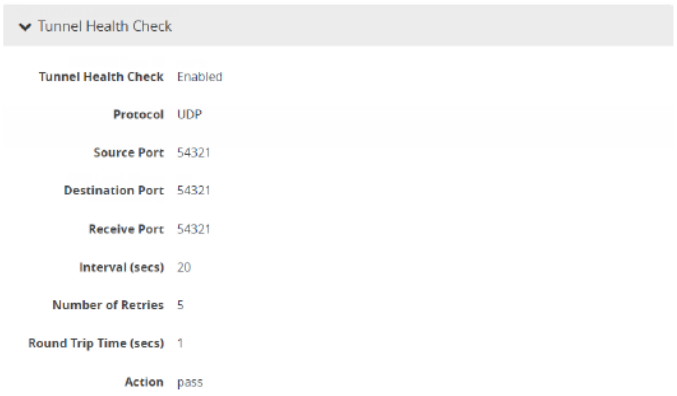

Figure 142*Configure UDP Health Check Parameters*

To view IP interface status, refer to Figure 143View IP Destination Status.

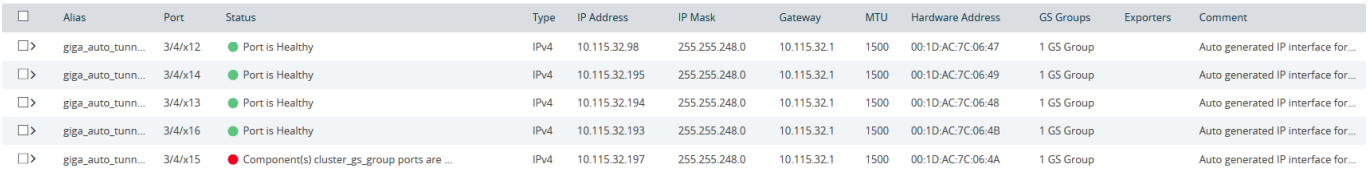

Alias	Port	Status	Type	IP Address	IP Mask	Gateway	MTU	Hardware Address	GS Groups	Exporters	Comment
giga_auto_tunn...	3/4/x12	Port is Healthy	IPv4	10.115.32.98	255.255.248.0	10.115.32.1	1500	00:1D:AC:7C:06:47	1 GS Group		Auto generated IP interface for...
giga_auto_tunn...	3/4/x14	Port is Healthy	IPv4	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 143*View IP Destination Status*

To view tunnel endpoint status, refer to Figure 144View 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 144*View Tunnel Endpoint Status*

To view IP interface statistics, refer to Figure 145View IP DestinationStatistics.

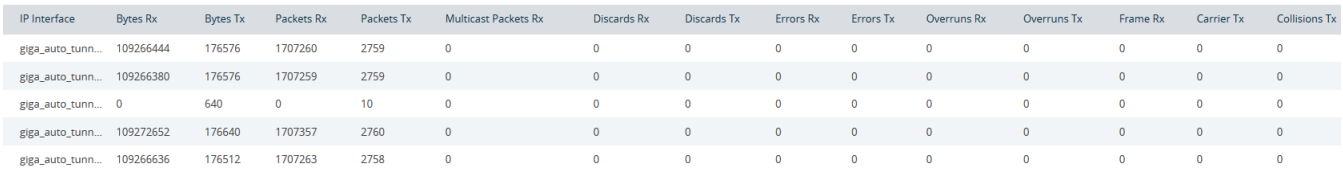

IP Interface	Bytes Rx	Bytes Tx	Packets Rx	Packets Tx	Multicast Packets Rx	Discards Rx	Discards Tx	Errors Rx	Errors Tx	Overruns Rx	Overruns Tx	Frame Rx	Carrier Tx	Collisions Tx
giga_auto_tunn...	109266444	176576	1707260	2759	0	0	0	0	0	0	0	0	0	0
giga_auto_tunn...	109266380	176576	1707259	2759	0	0	0	0	0	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 145*View IP DestinationStatistics*

To view GigaSMART operation statistics, refer to Figure 146View 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 146*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-C05 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-C05 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 GigaSMART Application Session Filtering (ASF) and Buffer ASF.
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 147GigaSMART Group on Different Chassis than Network/Tool Ports and Figure 148GigaSMART 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 147GigaSMART Group on Different Chassis than Network/Tool Ports and Figure 148GigaSMART Operations in Clusters illustrate examples of network, tool, and GigaSMART group ports on different nodes in a clustered environment. For example, Figure 147GigaSMART 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.

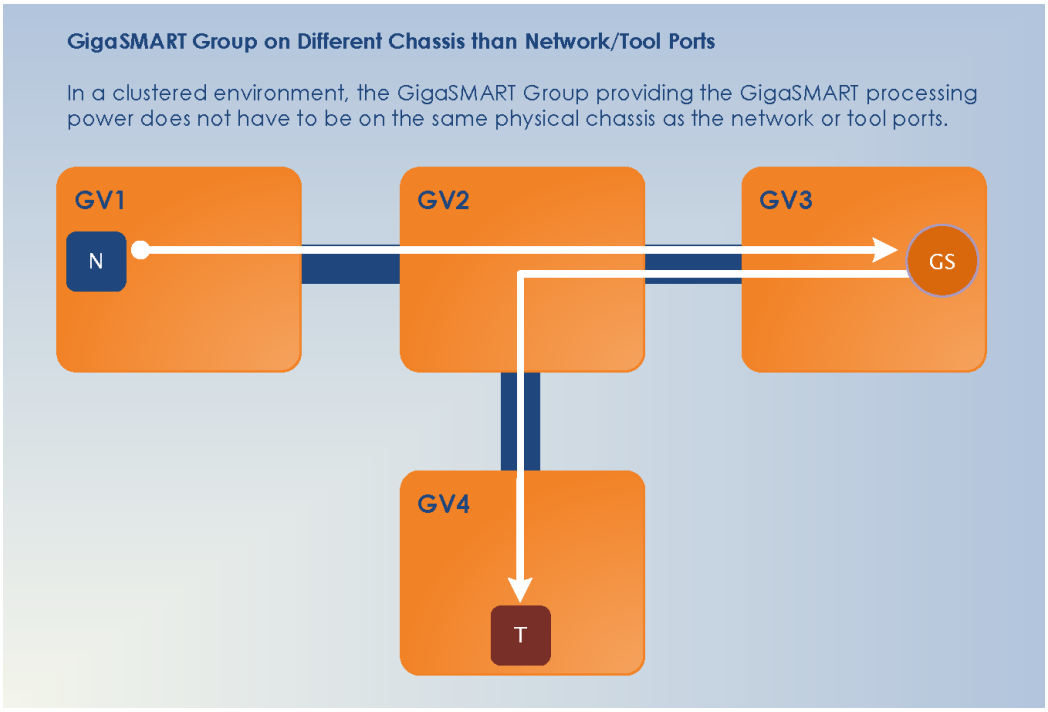

Figure 147*GigaSMART Group on Different Chassis than Network/Tool Ports*

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 148GigaSMART Operations in Clusters. The GigaSMART group on GV3 is on the same chassis as the egress port.

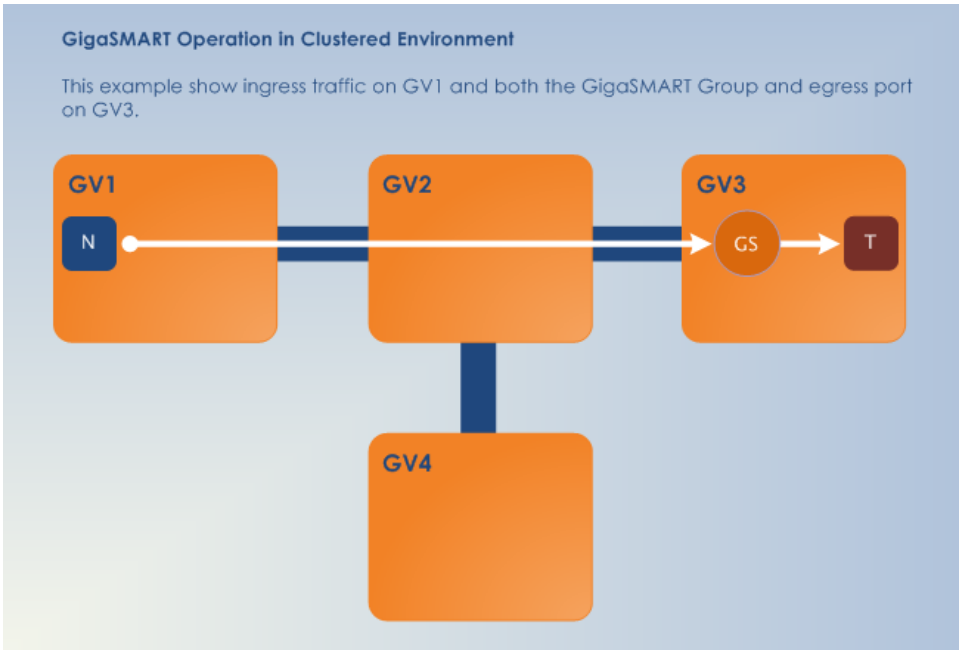

Figure 148*GigaSMART Operations in Clusters*

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.

Supported GSOP Combinations		Masking	Source Port Labelling	De- Dup	Load Balance	A PF	App Filtering			Flow-Ops				Headers			NetFlow		Slicing		SSL /TLS			Tunnels	
							A SF	A FI (ASF)	A MI	Flo wV UE	GTP Flow Filter	GTP Whit elist	GTP Flow Sampling	Strip Head ers	Add Head ers	Rem ove H/T	NetFlow 1st level	NetFlo w 2nd level	Sli cing	Advanced Flow Slicing	SSL Decr ypt	iS SL	ICAP	Tunn el Encap	Tunn el Decap
Masking			$	$	$	$	$	$	NS	$	NS	NS	NS	$	$	$	NS	NS	$	$	NS	NS	NS	$	$
Source Port Labelling		$		$	$	$	$	NS	NS	$	NS	NS	NS	$	$	$	NS	NS	$	$	NS	NS	NS	$	$
De-Dup		$	$		$	$	$	$	$	$	NS	NS	NS	$	$	$	$	$	$	$	$	NS	NS	$	$
Load Balance		$	$	$		$	$	$	NS	$	$	$	$	$	$	$	NS	NS	$	$	NS	NS	NS	$	$
APF		$	$	$	$		$	NA	NS	$	NS	NS	NS	$	$	NS	NS	NS	$	$	NS	NS	NS	$	NS
App Filtering	AFI (ASF)	$	NS	$	$	NA	NS		$	NS	NS	NS	NS	NS	NS	NS	NS	$	$	$	NS	NS	NS	NS	NS
	AMI	NS	NS	$	NS	NS	$	$		NS	NS	NS	NS	NS	NS	NS	NA	NA	NS	NS	NS	NS	NS	NS	NS
Flow- Ops	FlowVUE	$	$	$	$	$	$	NS	NS		NS	NS	NS	$	$	$	NS	NS	$	NS	NS	NS	NS	NS	NS
	GTP Flow Filter	NS	NS	NS	$	NS	NS	NS	NS	NS		$	NS	NS	NS	NS	NS	NS	$	NS	NS	NS	NS	NS	NS
	GTP Whitelist	NS	NS	NS	$	NS	NS	NS	NS	NS	$		NS	NS	NS	NS	NS	NS	$	NS	NS	NS	NS	NS	NS
	GTP Flow Sampling	NS	NS	NS	$	NS	NS	NS	NS	NS	NS	NS		NS	NS	NS	NS	NS	$	NS	NS	NS	NS	NS	NS
Headers	Strip Headers	$	$	$	$	$	$	NS	NS	$	NS	NS	NS		$	$	NS	NS	$	$	NS	NS	NS	$	$
	Add Headers	$	$	$	$	$	$	NS	NS	$	NS	NS	NS	$		$	NS	NS	$	$	NS	NS	NS	NS	$
	Remove H/T	$	$	$	$	NS	NS	NS	NS	$	NS	NS	NS	$	$		NS	NS	NS	NS	NS	NS	NS	$	NS
NetFlow	NetFlow 1st level	NS	NS	$	NS	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	NS	$	$	NA	NS	NS	NS	NS	NS	NS	NS	NS		NS	NS	NS	NS	NS	NS	NS
Slicing	Slicing	$	$	$	$	$	$	$	NS	$	$	$	$	$	$	NS	NS	NS		$	NS	NS	NS	$	$
	Advanced Flow Slicing	$	$	$	$	$	$	$	NS	$	$	$	$	$	$	NS	NS	NS	$		NS	NS	NS	$	$
SSL/ TLS	SSL Decrypt	NS	NS	$	NS	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	$	$	$	$	$	$	NS	NS	NS	NS	NS	NS	$	NS	$	NS	NS	$	$	NS	NS	NS		NS
	Tunnel Decap	$	$	$	$	NS	NS	NS	NS	NS	NS	NS	NS	$	$	NS	NS	NS	$	$	NS	NS	NS	NS	

$ = Supported NS NS = Not Supported NA = Not Applicable

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.

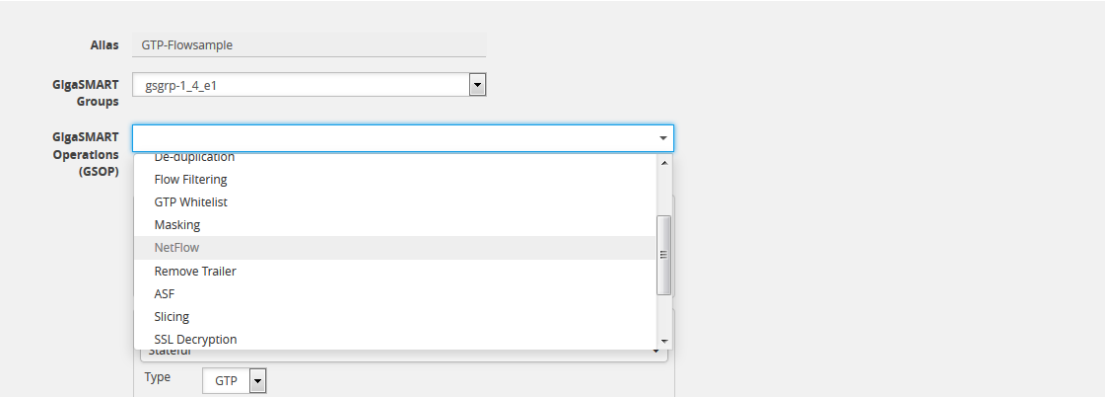

Figure 149*GigaSMART Operation Combinations Not Available*

However, in the following 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 2 (C05)	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	✓	×	✓	×	×	×	×

GigaSMART Operation	GigaVUE-HC1		GigaVUE-HC3		GigaVUE-HC1 Plus		GigaVUE-HCT
	Gen 2 On-board	Gen 3 SMT-HC1-S	Gen 2 (C05)	Gen 3 (C08)	Rear Gen3 SMT-HC1A-R	Front Gen 3 SMT-HC1-S	Gen 3 SMT-HC1-S
Remove							
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 Header Stripping	✓	✓	✓	✓	✓	✓	✓
GigaSMART Header Addition	✓	✓	✓	✓	✓	✓	✓
GigaSMART IP FlowVUE	✓	✓	✓	✓	✓	✓	✗
GTP Flow Filtering	✗	✗	✓	✗	✗	✗	✗
GigaSMART GTP Rotational Sampling	✗	✗	✓	✗	✗	✗	✗
GTP Whitelisting (GigaSMART GTP Rotational Sampling)	✗	✗	✓	✓ Non-CUPS GTP Correlation is supported.	✓	✓	✗

GigaSMART Operation	GigaVUE-HC1		GigaVUE-HC3		GigaVUE-HC1 Plus		GigaVUE-HCT
	Gen 2 On-board	Gen 3 SMT-HC1-S	Gen 2 (C05)	Gen 3 (C08)	Rear Gen3 SMT-HC1A-R	Front Gen 3 SMT-HC1-S	Gen 3 SMT-HC1-S
				LTE CUPS is released as BETA* in 6.1.			
GTP Flow Sampling (GigaSMART GTP 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.	✓	✓	✓	✓	✓	✓	✓
GigaSMART Application Session Filtering (ASF) and Buffer ASF	✓	✓	✓	✓	✓	✓	✓
Application Filtering Intelligence(AFI)	✓	✓	✓	✓	✓	✓	✓
Application Metadata Intelligence (AMI)	✓	✓	✓	✓	✓	✓	✓
GigaSMART NetFlow Generation (Application Intelligence)	✓	✓	✓	✓	✓	✓	✓
Application	✓	✓	✓	✓	✓	✓	✓

GigaSMART Operation	GigaVUE-HC1		GigaVUE-HC3		GigaVUE-HC1 Plus		GigaVUE-HCT
	Gen 2 On-board	Gen 3 SMT-HC1-S	Gen 2 (C05)	Gen 3 (C08)	Rear Gen3 SMT-HC1A-R	Front Gen 3 SMT-HC1-S	Gen 3 SMT-HC1-S
Visualization							
GigaSMART NetFlow Generation (Traffic Intelligence)	✓	×	✓	×	×	×	×
GigaSMART Load Balancing (Stateless)	✓	✓	✓	✓	✓	✓	✓
GigaSMART Load Balancing (Stateful)	✓	✓	✓	✓	✓	✓	✓
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	×	✓	×	✓	✓	✓	✓
GRE-In-UDP Tunnel Decapsulation	×	×	×	×	×	×	×

*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 Third Party Orchestration	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	✓	✗	✓	✓	✓	✓	✓

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 Third Party Orchestration	GigaVUE Cloud Suite for Nutanix
Deployment Guide.							
VXLAN Tunnel Encapsulation Refer to *Create Ingress and Egress Tunnels* section on the retrospective GigaVUE Cloud Suite Deployment Guide.	✓	✓	✓	✓	✓	✓	✓
L2GRE Tunnel Decapsulation Refer to *Create Ingress and Egress Tunnels* section on the retrospective GigaVUE Cloud Suite Deployment Guide.	✓	✗	✓	✓	✓	✓	✓
VXLAN Tunnel Decapsulati	✓	✓	✓	✓	✓	✓	✓

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 Third Party Orchestration	GigaVUE Cloud Suite for Nutanix
on 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*	✓	✗	✓	✓	✓	✓	✗

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 Third Party Orchestration	GigaVUE Cloud Suite for Nutanix
Tunnels section on the retrospective GigaVUE Cloud Suite Deployment 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	✓	✓	✓	✓	✓	✓	✓

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 Third Party Orchestration	GigaVUE Cloud Suite for Nutanix
Balancing							
SSL Decrypt	✓	✓	✓	✓	×	✓	✓
5G-Service Based Interface Application	×	×	✓	✓	✓	✓	×
5G-Cloud Applications	×	×	✓	✓	×	✓	×

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 150GigaSMART 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 150*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 150GigaSMART 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.

Statistic	Description
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

Statistic	Description
	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 151IP Interfaces Statistics shows an example.

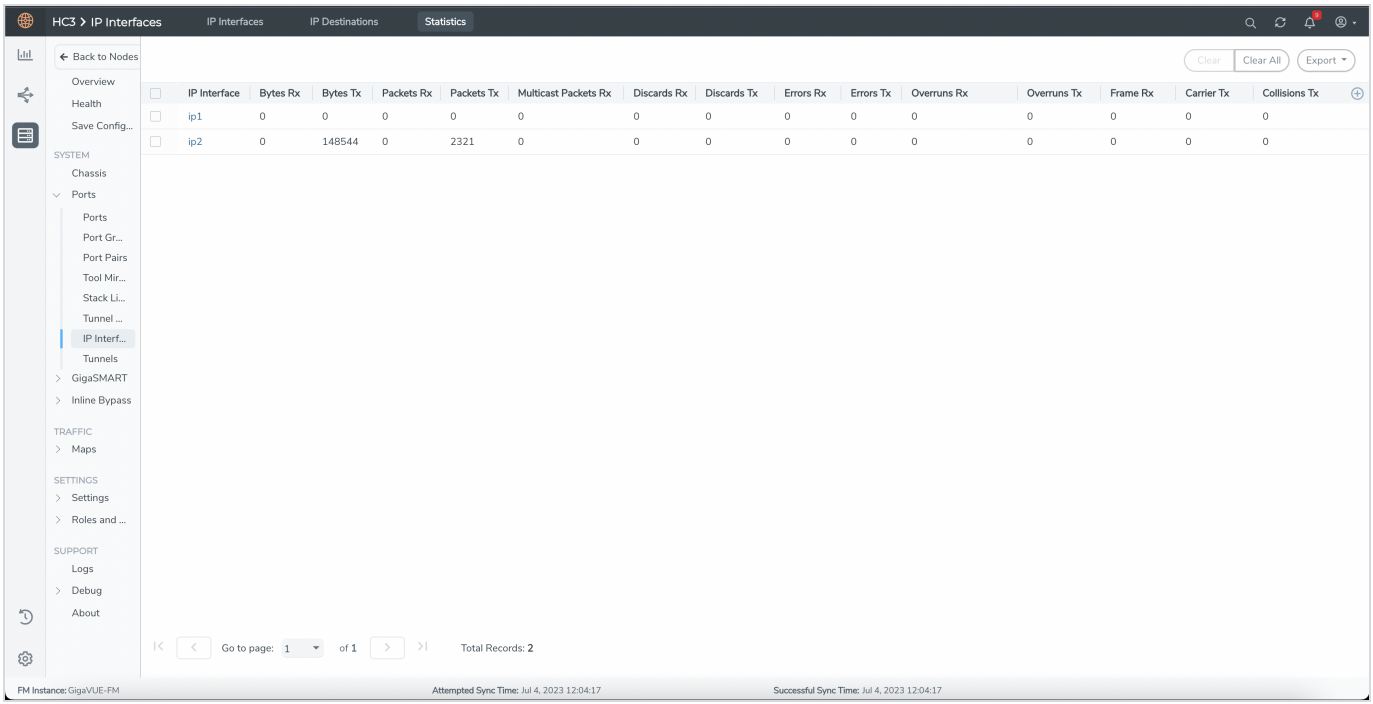

Figure 151 *IP Interfaces Statistics*

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.

Statistic	Description
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.
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.

Statistic	Description
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.)

Statistic	Description
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.)
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 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 152GigaSMART Group Statistics shows an example.

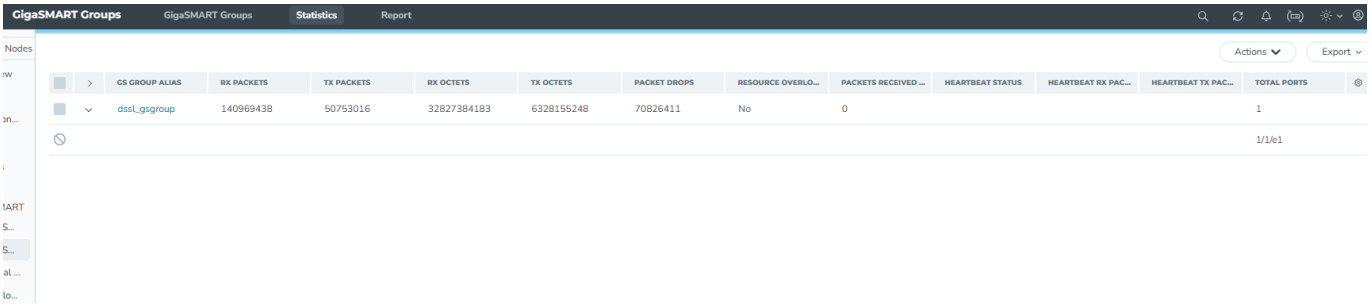

Figure 152 *GigaSMART Group Statistics*

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 Errors	The cumulative number of received packets with errors at a 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.

Statistic	Description
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.
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.

Statistic	Description
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
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

Statistic	Description
	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.
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	

Statistic	Description
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 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.

Statistic	Description
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 ▪ 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

Statistic	Description
	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.
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.

Statistic	Description
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.

Statistic	Description
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.
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.

Statistic	Description
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 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.

Statistic	Description
	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

Statistic	Description
	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 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.

Statistic	Description
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 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 and dropped due to overlapping reassembly fragments.

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.

Statistic	Description
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.
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

Statistic	Description
	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.
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.

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.

Statistic	Description
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.
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)

Statistic	Description
	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).
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 24: 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.

Name	Format
% 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-HC3 Gen 2

GigaSMART operations can add the GigaSMART Trailer to packets, providing metadata on where the packet was processed.

GigaSMART Trailers are optional for some 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 as another. The Source ID indicates where the packet entered the GigaVUE H Series node and how it was processed. Refer to About the Source ID Field for information included in the Source ID field.

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 for more details on the devices that support GigaSMART operations.

About the 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 TLVfor 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

Feature	Description
	H Series uses cluster IDs. o **Box ID** – The box ID configured for the GigaVUE node on which the packet was received. Box IDs are used for unique identification of nodes in a cluster. **NOTE:** The box ID field in the Gigamon Trailer supports box ID values from 1-64, 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 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 153GigaSMART Operation with Masking and Trailer Componentsshows the GigaSMART operation with masking an trailer components.

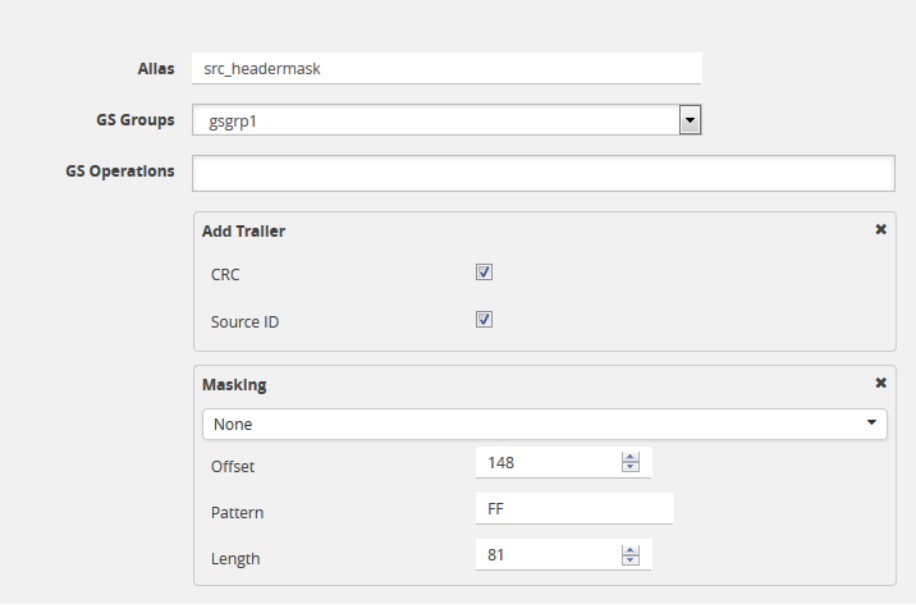

Figure 153 *GigaSMART Operation with Masking and Trailer Components*

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 154Remove Trailer Enabled shows Remove Trailer enabled for a GigaSMART operation.

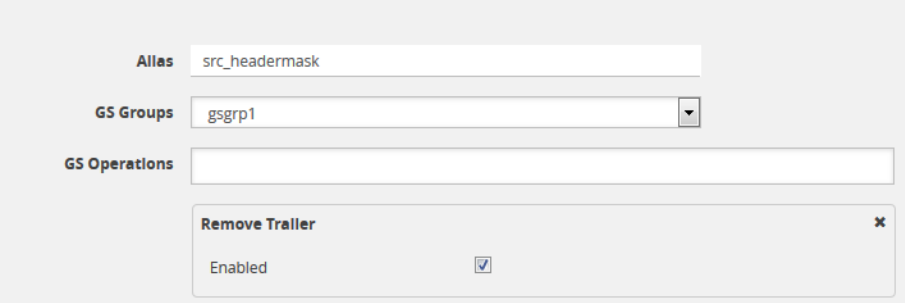

Figure 154*Remove Trailer Enabled*

Example: Removing GigaSMART Trailers in a Cascade

Figure 155Cascade Physically Cabled from Tool to Network Port and Figure 156Removing 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 155Cascade 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.

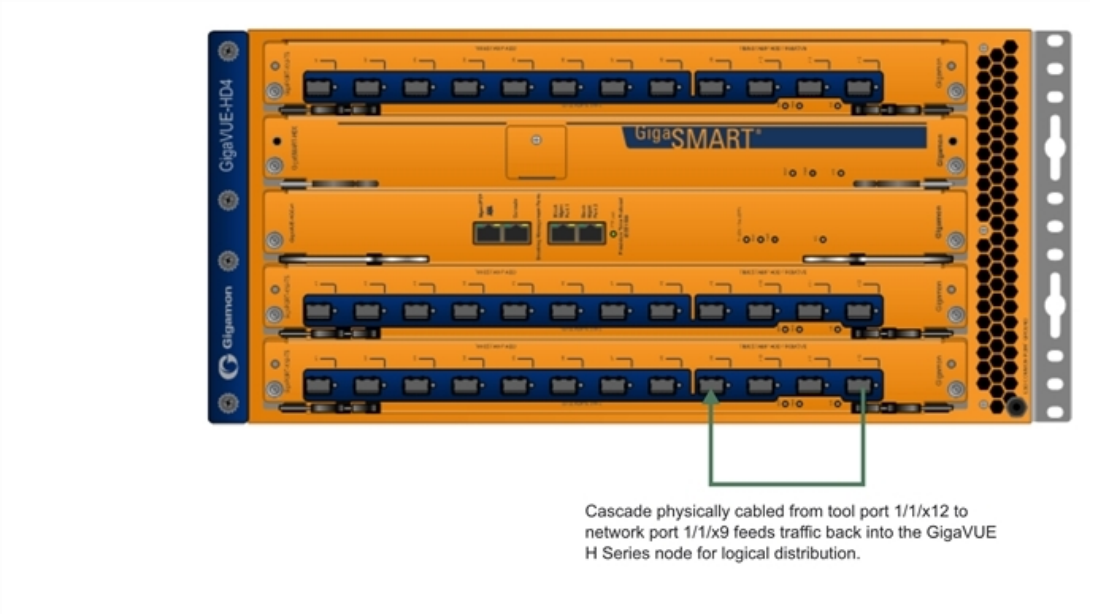

Figure 155 *Cascade Physically Cabled from Tool to Network Port*

For example, consider the packet distribution shown in Figure 156Removing 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.

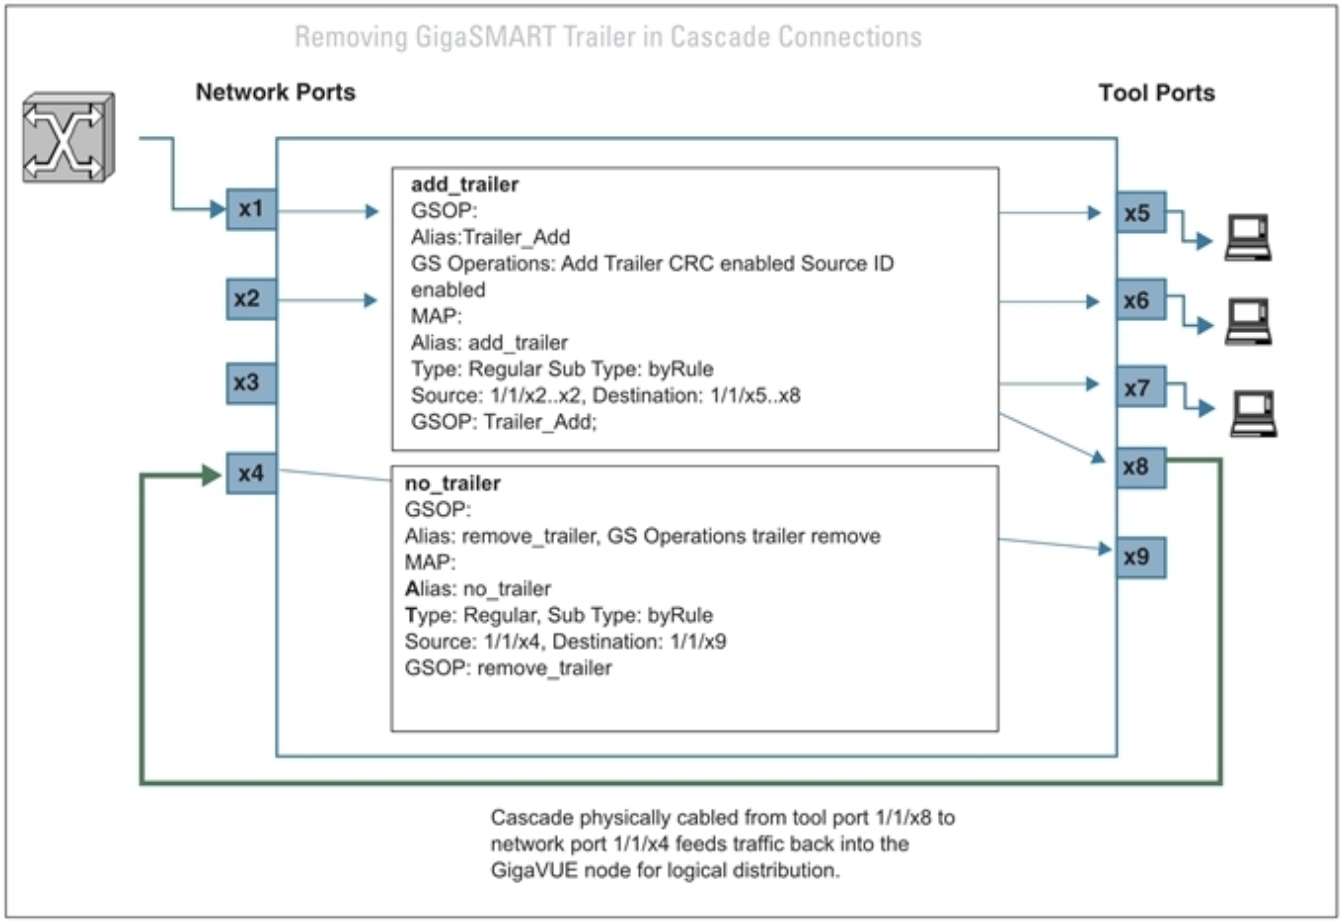

Figure 156 *Removing the GigaSMART Trailer from a Cascade Connection*

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 156Removing the GigaSMART Trailer from a Cascade Connectionand 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 that 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 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 157GigaSMART Trailer Format summarizes the position and contents of the GigaSMART Trailer and the Gigamon Ethertype field (0x22E5).

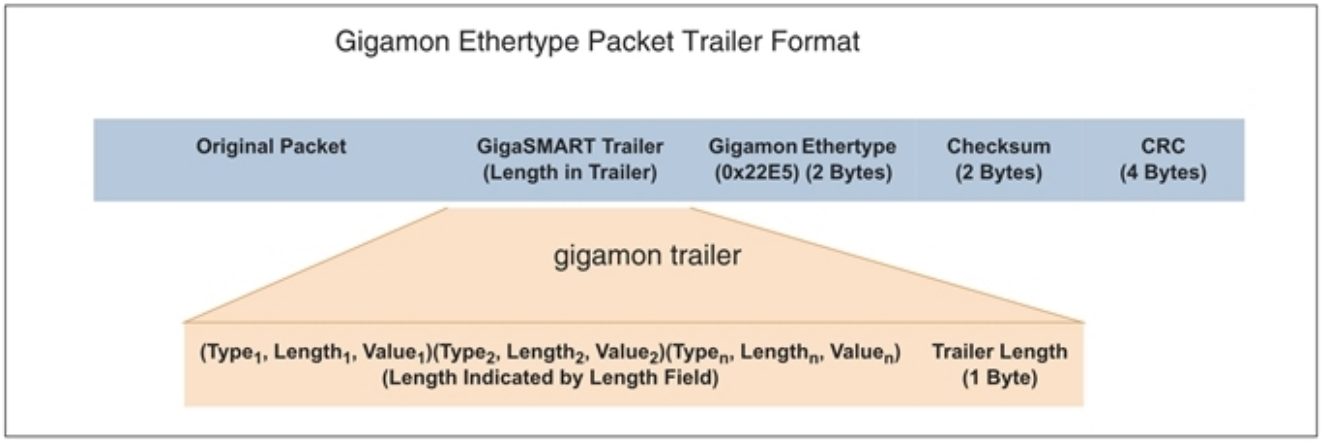

Figure 157 *GigaSMART Trailer Format*

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 25: Gigamon Trailer Format lists and describes each of the fields in the GigaSMART Trailer:

Table 25: 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 size of the Value field is indicated by the Length 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 includes the Length byte itself.
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 left to find 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 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 26: 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 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_ TIMESTAMP	3	8	**Timestamp – System Clock**	**Timestamp** A Timestamp TLV is included in any packet processed by a timestamp GigaSMART operation. The Timestamp always indicates the UTC time at which the last byte of the packet was seen by the GigaSMART line card or module. In addition, the exact type of Timestamp TLV indicates the time source used (NTP, GPS, and so on). Note that timestamp GigaSMART operations are not yet supported in this release of the GigaVUE H Series node.
GIGAMON_ TIMESTAMP_ NTP	4	8	**Timestamp – NTP Clock Source**	
GIGAMON_ TIMESTAMP_ GPS	5	8	**Timestamp – GPS Clock Source**	
GIGAMON_ TIMESTAMP_ 1588	6	8	Timestamp Based on PTP (IEEE 1588)	
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 (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 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_ TIMESTAMP_ 1588	6	8	**Timestamp – PTP (IEE 1588) Clock Source**	

Format of the GIGAMON_SRCID TLV

The GIGAMON_SRCID TLV is included in any packet processed by a **Source ID**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-TA100 o GigaVUE-HC1 o GigaVUE-HC3 o GigaVUE-TACS o GigaVUE-TA200 o GigaVUE-TA25	6
Group ID	Group IDs are used to identify a particular cluster. If a packet enters a cluster with a group ID configured, it is reflected here.	4
Box ID	Box IDs are used to uniquely identify nodes in a cluster. Standalone systems typically have the default box ID of **1** here. **NOTE:** The box ID field in the Gigamon Trailer supports box ID values from 1-64, inclusive.	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

Port ID Values by Line Card Type

The following table summarizes the values inserted in the GIGAMON_SRCID TLV for port ID by GigaVUE device and line card type:

GigaVUE Device / Line Card	Ports	Port IDs Inserted in Gigamon Trailer
GigaVUE-HC1		
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		
PRT-HC3-X24	x1..x24	1..24
PRT-HC3-C08Q16	c1..c8	1, 5, 9, 13, 17, 21, 25, 29
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
SMT-HC3-C05	c1..c5	1, 5, 9, 13, 17
SMT-HC3-C05–4x10Gb Mode	c1x1–x4.. c5x1–x4	1..20
SMT-HC3-C05–4x25Gb Mode on Control Card Version 2	c1x1–x4.. c5x1–x4	1..20
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

GigaVUE Device / Line Card	Ports	Port IDs Inserted in Gigamon Trailer
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 TA Series		
GigaVUE-TA25	x1..x48 c1..c8 c1x1..c1x4 c5x1..c5x4	1..48 49,53,57,61,65,69,73,77 49..52 65..68
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..c64	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, 42, 43, 54, 55, 68, 59, 60, 46, 47, 48, 50, 51, 52, 62, 63, 64, 66, 67, 78, 70, 79, 82, 83, 84, 71, 86, 72, 74, 75, 76, 87, 88

GigaVUE Device / Line Card	Ports	Port IDs Inserted in Gigamon Trailer
	c1x1..c32x4	1..128

GigaSMART Trailer Example

The following figures show a sample GigaSMART trailer with the packet length and source ID 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)



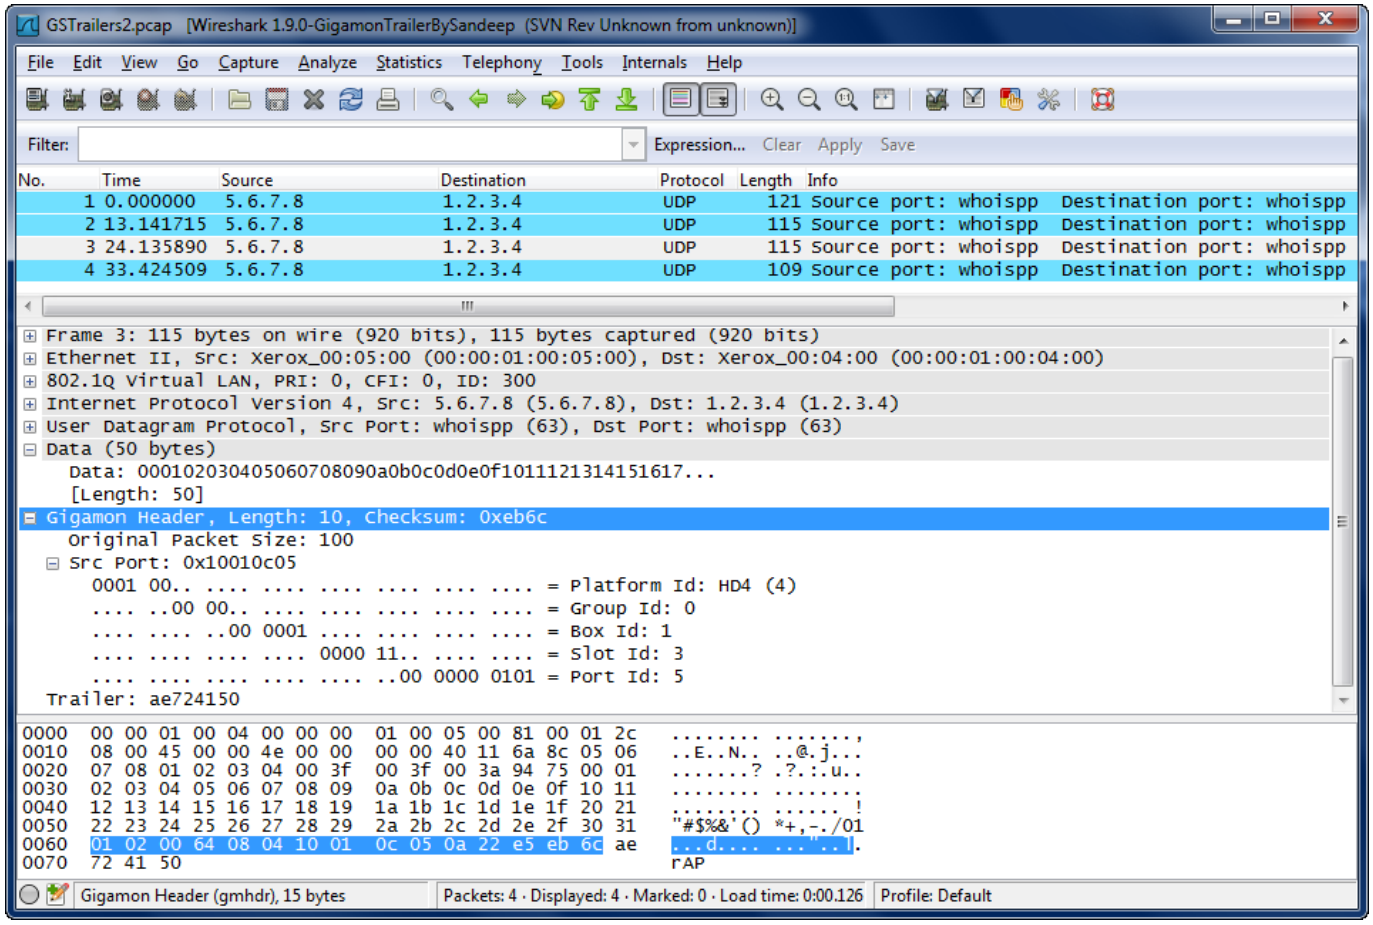

Figure 158 *Wireshark Decode of GigaSMART Trailer*

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.

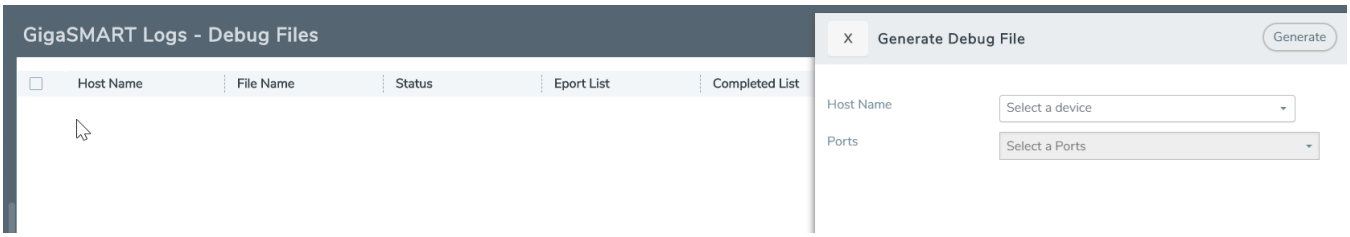

   The system generates a new GS Log file.

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. On the left navigation pane, click  .
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  > **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. On the left navigation pane, click  .
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  > **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 or the immediate next 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.

On the left navigation pane, click  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 6.9 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 6.9 Hardware and Software Guides
GigaVUE-TA200 Hardware Installation Guide
GigaVUE-TA200E Hardware Installation Guide
GigaVUE-TA400 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
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 Cloud Suite Deployment Guide - Third Party Orchestration
Universal Cloud TAP - Container Deployment Guide
Gigamon Containerized Broker Deployment Guide

GigaVUE 6.9 Hardware and Software Guides
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)