



GigaVUE Traffic Filtering Guide

GigaVUE

Product Version: 6.14

Document Version: 1.0

(See Change Notes for document updates.)

Copyright © 2026 Gigamon Inc. All rights reserved.

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

Trademark Attributions

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

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

Change Notes

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

Product Version	Document Version	Date Updated	Change Notes
6.14	1.0	02/25/2026	The original release of this document with 6.14.00 GA.

Contents

GigaVUE Traffic Filtering Guide	1
Change Notes	2
Contents	3
GV-Traffic-Filtering-Guide	20
Traffic Filtering	20
Ports and GigaStreams	21
About Ports	21
About Network and Tool Ports	21
Network (Ingress) Ports Defined	22
Tool (Egress) Ports Defined	22
Ports on GigaVUE® TA Series Traffic Aggregator Nodes	23
Concepts Illustrated in Figure 3GigaVUE-TA Packet Distribution	24
Notes and Consideration on GigaVUE® TA Series Nodes	25
Hybrid Ports	25
Stack Ports	26
Inline Network Ports	26
Inline Tool Ports	26
Circuit Ports	26
Required License: Advanced Feature License for GigaVUE® TA Series Nodes	26
GigaSMART Engine Ports	27
Port Lists	27
Port Aliases	27
Work with Hybrid Ports	28
Port Filters	30
Port Filter—Rules and Notes	30
Port-Filter Maximums	31
How to Apply Port Filters	32
View Port Filter Statistics	33
View Filter Resources for a Slot	34
Status of Line Cards/Nodes and Ports	35
How to Check Port Status with Ports Page	35
How to Check Port Status with Chassis Page	36

Managing Ports	36
Ports	37
All Ports	37
Port Quick View	40
Port List Filter	40
Quick Port Editor	43
Configure Ports	44
Ports Discovery	46
Statistics	47
Port Groups	48
All Port Groups	48
Create Port Groups	49
Edit Port Group	49
Clone Port Group	50
Port Group Statistics	51
Port Pairs	51
Create Port Pair	52
Tool Mirrors	53
Create Tool Mirror	54
Edit Tool Mirror Description	54
Clone Tool Mirror	54
Stack Links	55
IP Interfaces	56
Configure IP Interface	57
Prerequisites	57
IP Interface Statistics	58
Circuit Tunnels	58
Port Discovery	59
Port Discovery with LLDP and CDP	59
Enable Port Discovery	61
Limits of Discovery Information	61
Port Discovery Support	62
Port Discovery for a Cluster	62
Port Discovery Supported for SNMP	62
Ingress and Egress VLAN	62
About Ingress Port VLAN Tagging	63
Ingress Port VLAN Tagging	64
Adding VLAN Tags	64

Deleting VLAN Tags	65
Using VLAN Tags in Maps	65
Ingress Port VLAN Tag Limitations	66
Second Level Maps	66
Double-Tagged Packets	66
IP Interfaces	66
Local Tool Port Ingress VLAN Tag	67
Configure Egress Port VLAN Stripping	67
Enable Egress Port VLAN Stripping	67
Disable Egress Port VLAN Stripping	68
Display Egress Port VLAN Stripping	68
Egress Port VLAN Stripping Limitations	68
How to Use Both Ingress Tagging and Egress Stripping ..	69
How to Use GigaStream	69
GigaStreams	70
Accessing the GigaStream page	70
About GigaStream	71
Regular GigaStream	72
Controlled GigaStream	72
Regular GigaStream	72
Regular Tool GigaStream	73
Regular Stack GigaStream	74
Configure Regular GigaStream	75
Edit Regular GigaStream	76
Edit Regular Stack GigaStream	77
Traffic Distribution Across Regular GigaStream	77
Regular GigaStream Failover Protection—Resiliency	78
Regular Circuit GigaStream	78
Controlled GigaStream	79
Notes and Considerations for Controlled GigaStream	81
Controlled GigaStream Configuration	82
Edit Controlled GigaStream	84
Traffic Distribution Across Controlled GigaStream	85
Failover and Controlled GigaStream	87
Advanced Hashing	88
How to Change Advanced Hash Criteria	88
Advanced Hash Settings	88
Advanced Hash Examples	90
Hashing Behavior	92

Notes and Considerations for Advanced Hashing	94
Advanced Hashing with MPLS	94
Advanced Hashing with GTP TEID	96
Packet Distribution and the Advanced Hash Algorithm	98
Weighted GigaStream	98
GigaStream Rules and Maximums	99
Maximum Ports per GigaStream	100
Port Statistics and Counters	101
Display Port Statistics	101
Display Port Statistics for a Single Port	101
Display Statistics for All Ports	103
How to Clear Traffic Counters	105
Header Stripping	105
About VXLAN Header Stripping	105
VXLAN Tunnel Decapsulation Versus VXLAN Header Stripping	107
VXLAN Header Stripping – Rules, Notes, and Limitations	107
Configure VXLAN Header Stripping	109
Configure Ageing Interval for VXLAN Header Stripping ...	110
View Header Stripping Statistics	110
About MPLS Header Stripping	111
Advanced MPLS Header Stripping for GigaVUE-TA400	112
MPLS Header Stripping – Rules, Notes, and Limitations	113
Configure MPLS Header Stripping	115
Create a New Map	115
Configure MPLS Labels on a Chassis	118
Enable Header Stripping Protocol on Ports	118
Tunnels	121
Required License: Advanced Feature License on GigaVUE TA Series Nodes	121
About Circuit-ID Tunnels	122
Circuit-ID Tunnels—Rules and Notes	123
Circuit-ID Tunnel Encapsulation	124
Circuit-ID Tunnel Decapsulation	124
About Layer 2 Generic Routing Encapsulation (L2GRE) Tunnels	125
L2GRE Tunnel Configuration—Rules and Notes	126
Limitation	127

Configure L2GRE Tunnel to Encapsulate Traffic	127
Configure L2GRE Tunnel to Decapsulate Traffic	128
Orchestrated Workflow Configuration of L2GRE Tunnels	129
About Virtual Extensible LAN (VXLAN) Tunnels	131
VXLAN Tunnel Configuration—Rules and Notes	131
Configure VXLAN Tunnel to Encapsulate Traffic	133
Configure VXLAN Tunnel to Decapsulate Traffic	134
Orchestrated Workflow Configuration of VXLAN Tunnels	135
Create Tunnel	136
Create VXLAN / L2GRE Group	137
Configure L2GRE / VXLAN Identifier	138
Configure L2GRE / VXLAN Identifier for a Chassis	138
Configure L2GRE / VXLAN Identifier for a Network Port	139
View VXLAN / L2GRE ID Statistics	139
Tunnel Monitoring	141
Rules, Notes, and Limitations	142
View Tunnel Monitoring	143
Packet Capture (PCAP)	145
Rules, Notes, and Limitations	147
Configure PCAP Profile	148
View PCAP	151
Delete PCAP	151
View PCAP Files	151
Configure PCAP on IP Interface	152
View PCAP on IP Interface	153
Delete PCAP on IP Interface	154
Flow Mapping®	154
About Flow Mapping®	154
Flow Mapping® Overview	155
Get Started with Flow Mapping®	156
Check Status of Nodes and Ports	156
Designated Port Types	156
About Shared Collectors	157
No Map Statistics for Shared-Collector Only	157
Shared Collector Configuration	157
About Map-passall Maps	159

Map-Passall and Regular Byrule Map	159
Map-Passall and Shared-Collector Only	159
Map-passalls Configuration	159
Map-passalls	160
Define Map Source Port Lists	160
Share Network Ports Between Maps	160
Share Tool Ports Between Maps	161
Map Priority	161
Adjust Map Priority in GigaVUE-FM	161
Flow Map Syntax and Construction	162
Map Types	162
Define Map Source Port Lists for First Level Maps	163
Null Port in Maps	165
Map Subtypes	166
Map Type and Subtype Modification	167
Backwards Compatibility	167
Minimum Requirements for Map Creation	167
Map Rules	167
Other Types of Map Rules for GigaSMART Operations	167
IPv4/IPv6 and Map Rules	168
Protocol Map Rules and IPv6	169
Set Map Rules for TCP Control Bits	169
Examples	170
How to Use Bit Count Netmasks	171
How to Combine Rules and Rule Logic	173
Examples of Map Rule Logic	173
How to Mix Pass and Drop Rules	173
Drop Rules Have Precedence!	174
Configuring Port Criteria and Bi-Directional Rules in By Rule Maps	174
Using the Bi-Directional Option	175
Work with User-Defined Pattern Match Rules	175
User-Defined Pattern Match Syntax	176
User-Defined Pattern Match Rules	176
User-Defined Pattern Match Examples	177
Map #1 – GTP_Map218	178
Map #2 – GTP_Map228	178
User-Defined Pattern Match on GigaVUE-TA400, and GigaVUE-TA400E	179
Inner Header and MPLS Header Filtering	180
Inner Header Attributes	180
Example:	181
MPLS Header Attributes	181
Example:	182

How to Handle Q-in-Q Packets in Maps	183
Comparison of Q-in-Q Tagging	184
Priority Tagged Packets	185
Flow Mapping® on Inner VLAN Tags	185
Inner VLAN Limitation	186
Work with Map-Passalls and Port Mirroring in GigaVUE-FM	187
Syntax for Maps-passalls and Port Mirroring	187
Rules for Map-Passalls and Port Mirroring	187
View and Delete Map-passalls	188
Port Access and Map Sharing	188
Port-based Access Levels	188
How to share Maps	189
Map Examples	190
Example: How to Create a Simple Map	190
Example: How to Handle Overlaps when Sending VLANs and Subnets to Different Tools	192
Manage Maps	194
Map Views	194
List View	194
Map Topology View	195
Quickview	196
Free Layout	197
Reset to Default Layer	198
Manage Maps	198
Create a new map	199
Prerequisites:	199
Create a new map – Field References	199
Clone Map	202
Edit Maps	202
Delete Maps	203
Create Map Groups	203
Description to Map Rules	204
Error Messages	204
Edit Map Rule Description	204
How to Use the Quick Editor for Pass and Drop Rules	205
How to Use the Quick Editor to Add Port Numbers	205
How to Use the Quick Port Editor to Add IP Address	207
Map Templates	207
Create Map Templates	207

Edit Map Templates	208
Map Template Quick View	208
Manage Map Rule Resources	208
Template Groups	208
Add Tags to Map Rules	211
Notes	212
Map Rule Statistics Dashboard	212
Filter Templates	213
Filter Template Qualifiers and Defaults	214
Filter Template Configuration	215
Filter Template Limits	215
How to Understand Map Filter Resources	215
Filter Template Rules and Recommendations	216
Filter Template Best Practices	216
Filter Templates in a Cluster	217
Filter Templates Formulas	218
Troubleshoot Custom Filter Template Issues	220
Review Map Statistics with Map Rule Counters	222
Viewing Map Statistics with the Statistics Page	223
Viewing Map Statistics with Quick View	223
Flow Mapping® FAQ	223
How Many Map Rules are Supported?	223
How Many Rules Can Each Map Have?	228
How Many Maps Can Run at Once?	229
What Criteria can be Filtered in Q-in-Q Packets?	229
How Many Maps Can Share a Network Port?	229
How Many Network Ports and Tool Ports Can Be in a Map?	229
Are Port-Filters Supported?	229
Does Flow Mapping® Support Passalls?	230
Does Flow Mapping® Support port-pairs?	230
Does Flow Mapping® Support UDA Pattern Matches?	231
Are Maps Supported Across Nodes in a Cluster?	231
Can a GigaStream Act as a Shared Collector?	231
What Are the GigaStream Maximums?	231
DoesGigaVUE-FM Provide the Same Features as the GigaVUE-OS CLI?	231
What order are the map rules displayed in "show running config"?	232
Active Visibility	232
MAC Address Rewrite	235
Configuring MAC Address Re-write	236
License	237

Limitations	237
IP Address Rewrite	238
Configuring IP Address Re-write	239
License	240
Limitations	240
VLAN Manipulation	240
Required License:	240
Overview	241
Rule-Level Add VLAN Manipulation	241
Rule-Level Modify VLAN Manipulation	242
Rule-Level Strip VLAN Manipulation	243
Map-Level VLAN Manipulation	244
Limitations	247
Monitor Port Utilization	247
Port Utilization Availability by Port Type	248
Set Port Utilization Thresholds	248
Utilization Alarm/SNMP Trap Generation	249
Configure Alarm Buffer Thresholds	250
Set Alarm Buffer Thresholds	251
Microburst	252
Best Practices to Improve Burst Tolerance	253
Logical Grouping of Ports in GigaVUE-HC3, GigaVUE-TA100, and GigaVUE-TA200	253
Flexible Inline Arrangements	254
Flexible Inline Arrangements	255
Flexible Inline Arrangement vs Inline Bypass Solution	256
Supported Platforms	257
Flexible Inline Arrangement License	258
Licensing Rules and Notes	258
Software Version	258
Flexible Inline Solution Supported in Clustered Nodes	258
Limitations	260

Functionalities Not Supported by Flexible Inline Arrangement	260
Benefits of Flexible Inline Arrangements	260
About Flexible Inline Maps	261
Types of Flexible Inline Maps	262
Configure Flexible Inline Maps	263
Flexible Inline Arrangements—Rules and Notes	263
Visualize the Flexible Inline Arrangements Canvas	264
Configure Flexible Inline Flows	265
Troubleshoot Flexible Inline Flows	265
Status	265
Rules and Notes	266
Statistics	267
HSM Statistics	268
ICAP Show Stats	269
Troubleshoot	271
Example: Troubleshoot Traffic Issues Between Side A and Side B	272
View the Forwarding States of Inline Networks	273
Import and Export Flexible Inline Solution	280
Backup and Restore Flexible Inline Flows	281
Inline Bypass Solutions	282
About Inline Bypass Solutions	282
Introduction to Inline Bypass Solutions	283
Capabilities of Inline Bypass Solutions	283
Logical Bypass and Physical Bypass	284
Simple and Complex Inline Bypass Solutions	286
Typical Configuration	287
Distribution to Multiple Inline Tools	287
Inline Tools in a Series	288
Multiple Inline Networks	289
Inline Flow Mapping®	289
Send Traffic to Out-of-Band Tools	290
Configure Inline Bypass Solutions	291
Inline Networks	291
Unprotected Inline Network	292

Protected Inline Network	292
Mix of Protected and Unprotected	293
Network Port Link Status Propagation Parameter	294
Traffic Path Parameter	294
Physical Bypass Parameter	297
Redundancy Profile Parameter	298
Display Current State of Inline Bypass Solution	298
How to Use SNMP Polling to Obtain Inline Network State	303
SNMP Notification of Forwarding State Change	303
How to Use Syslog to Obtain Inline Network State	303
Inline Network Groups	304
Inline Tool Sharing	305
Configurable VLAN Tagging	306
Add VLAN Tag	306
Tools in Bridge Mode	306
Inline Tools	307
Inline Tool Failover Action	308
Inline Tool Failover Action with Inline Flow Mapping®	309
Inline Tool Recovery Mode	309
Inline Tool Sharing Mode	311
How to Use SNMP Polling to Obtain Inline Tool State	312
Heartbeats	312
Rules and Notes	313
Heartbeat Profiles	313
Standard Heartbeat	314
Standard or Custom Heartbeat Packet	315
Detect Inline Tool Failure	315
Negative Heartbeat Profiles	315
Heartbeat Support Between GigaVUE Nodes	316
Troubleshoot Heartbeat Failures	318
Configure a Heartbeat Profile	318
View Heartbeat Profile Statistics	319
Heartbeat Status after System Reload	320
Inline Tool Groups	320
Inline Tool Group Failover Action	323
Inline Tool Group Spare Inline Tool	324
Symmetrical and Asymmetrical Hashing	324
Asymmetrical Hashing Restrictions	326
Resilient Weighted Hashing	327
Inline Serial Tools	327
Inline Serial Tools Global Failover Action	330

Inline Tool Series Local Failover Action	332
Inline Tool Series Per-Direction Order	334
Associate Inline Networks with Inline Tools Using Inline	
Maps	336
Inline Map Passall	337
Inline Map	337
Inline Map Shared Collector	338
Inline Maps to Individual Members of an Inline Tool Group	339
Map, Inline Tool, and Inline Tool Group Configuration Restrictions	340
Inline Tool Failures and Failover Actions	341
Maps That May Lead to Selective Traffic Drops	342
Out-of-Band (OOB) Map	342
Symmetric and Asymmetric Maps	344
Configuration Steps	344
Configuration Step Details	347
Configure Inline Network Ports	347
Configure Inline Network (Unprotected)	347
Configure Inline Network Group	348
Create Heartbeat Profile	348
Create Negative Heartbeat Profile	349
Configure Inline Tool Ports	350
Configure Inline Tool	351
Create Inline Tool Group	352
Configure Inline Tool Series	353
Configure When GigaVUE® HC Series Modules are Operationally Up	354
Avoid Over subscription	354
Configure Gigamon Resiliency for Inline Protection	354
How to Cable GigaVUE Nodes	355
How to Handle Recovery	356
Both Nodes Go Down and Only Secondary Comes Up	356
How to Cable GigaVUE Nodes	356
Configure GRIP Solution Software	356
Limitation for Suspended Role	358
Configure Synchronization	358
Display Redundancy Control State	359
How to Use Suspended Role for Maintenance	359
Upgrade Procedure Recommendations	360
Rules and Notes	360
Limitations	361
Troubleshoot	361
Signaling Ports Down	361
Traffic outage in Inline Tool	361

Network Traffic Outage	362
FAQs	362
Migrate Inline Bypass(Classic) Solution from version 6.5 to 6.12	363
Migration Scenarios	363
Migration Considerations	363
Migrate Classic Inline Bypass Solution using GigaVUE-OS CLI	364
Migrate Classic Inline Bypass Solution using GigaVUE-FM	365
Inline Bypass Solution Examples	366
Example 1: Unprotected Inline Bypass with an Inline Tool Group	367
Example 2: Unprotected Inline Bypass with Default Heartbeat	369
Example 3: Protected Inline Bypass Using Combo Modules	370
Example 4: Gigamon Resiliency for Inline Protection	373
Configure Primary Role GigaVUE-HC2	373
Configure Secondary Role GigaVUE-HC2	375
Timestamps	377
Required License: Advanced Feature License	377
About Timestamps	378
Using Timestamps	379
Why PTP?	379
Limitations	381
Synchronization of the PTP and Local System Clock	381

Using PTP to Timestamp Packets	382
Configuring PTP Globally	382
Enabling PTP on a Network Port	383
Enabling Timestamp on a Port	384
Viewing PTP Details	385
Viewing the PTP Configuration on a Port	386
Viewing the PTP Clock Details	388
Viewing the PTP Foreign Source Details	392
Viewing PTP Statistics	393
PTP and Timestamp—Rules and Notes	394
Fabric Maps	395
Supported Topologies	395
Multihop Topology	396
Supported Multihop Topology	397
Unsupported Multihop Topology	397
Leaf and Spine Topology	398
Spine Source Map Mode	399
Dual Multipath Leaf and Spine Topology	399
Supported Failovers	400
Scenario 1	401
Scenario 2	402
UnSupported Failovers	402
Scenario 1	403
Scenario 2	404
Fabric Maps Prerequisites	404
Notes on Circuit Ports and Circuit GigaStream in Fabric	
Maps	405
Notes for Circuit ID	406
	407
Create Links between Clusters	407
Create Manual Links	407
Create Gigamon Discovery Protocol (GDP) Based Links	408
Create Fabric Maps	409
Edit and Delete Fabric Maps	412
Edit Fabric Maps	413
Edit Fabric Map Component	413

Delete Fabric Maps	414
Prioritize Fabric Maps	415
Configure Fabric Map Settings	416
Share Circuit ID Resource	416
Spine Source Map Mode	416
Limitations of Spine Source Map Mode	417
Fabric Maps Statistics	418
Display Fabric Map Statistics	418
Display Fabric Map Details	418
Filter Fabric Maps List View	419
Fabric Port Group	420
How Fabric Port Groups Work?	420
Rules and Notes	421
Create Fabric Port Group	422
Configure Fabric Map Using Fabric Port Group	423
Troubleshooting	424
How to Troubleshoot Partial Success Errors	427
Example 1:	427
Config Audit	429
Limitations of Fabric Maps	430
Backup and Restore Fabric Maps and Orchestrated Configurations	431
Orchestrated Configurations	431
About Intent Based Orchestrated Configurations	432
Benefits of Orchestrated Configurations	433
Orchestrated Configuration: Examples	434
Priority-Free Policies	434
Overlapping Rules in Policies	436
Overlapping Sources in Policies	437
Supported Topologies	438
Rules and Notes for Orchestrated Configurations	438
Create Orchestrated Policies	439
Prerequisites	439
How to Create a Policy	440
Tagging and RBAC Support	442
Drop Rules	443

Import and Export Orchestrated Policies	445
Rules and Notes	447
Import and Export Orchestrated Configuration	447
Egress Filters for Additional Filtering Capabilities	448
Glossary	449
Explore Traffic Policy Interface	449
Overview	450
Dashboard	450
Map Details Grid	452
Menu	455
View Migration Results	456
Filters	456
Settings	458
Actions	458
Config Audit	458
Resubmit All	458
Edit	458
Delete	459
Clone	459
Deploy	459
Undeploy	460
View Details	460
View Statistics	461
View Troubleshoot	463
Resources Tab	463
Fabric Port Groups	463
Tool Templates	463
Create and Manage Fabric Port Group	464
Create Fabric Port Group	464
Edit Fabric Port Groups	465
Delete Fabric Port Groups	465
Export Fabric Port Groups	465
Configure Traffic Policy using Fabric Port Groups	465
Configure a Virtual Port	466
Configure a First-Level Map:	466
Configure a Second-Level Map:	467
Configure Traffic Policy Without GigaSMART	
Operations	469
Configure VLAN Manipulation	471
Map-Level VLAN Manipulation	471

Migrate Maps to Traffic Policy	473
Dry Run before Migration	474
Migrating Traffic Intelligence and Subscriber Intelligence Solutions	475
Migrating Flow Maps	475
Migrating Fabric Maps	476
Migrating Application Intelligence Solutions	476
View Migration Results	477
Configure Map Settings	478
Share Circuit ID Resource	479
Spine Source Map Mode	479
Limitations of Spine Source Map Mode	480
Configuration Sync	480
Troubleshoot Traffic Policy Issues	480
View the Troubleshooting Page	481
Troubleshooting Statistics and Health Indicators	481
Policy Level Health	481
Component Level Health	482
FAQ Traffic Policy	482
Additional Sources of Information	485
Documentation	485
How to Download Software and Release Notes from My Gigamon	488
Documentation Feedback	488
Contact Technical Support	489
Contact Sales	490
Premium Support	490
The VUE Community	490
Glossary	491

GV-Traffic-Filtering-Guide

Traffic Filtering

This section describes the core intelligence provided by GigaVUE-OS and managed through GigaVUE-FM that can be used to optimize your network traffic flow.

Core intelligence is...

- Flow Mapping® that defines policies to extract flows of interest
- GigaStream load balancing that distributes flows across tools
- Terabit-scale configurations with Clustering and Fabric Maps
- Inline Bypass configurations that optimize threat prevention tools, enforcement point
- Visibility across physical, virtual and cloud infrastructure

You can use the GigaVUE-OS traffic management capabilities within GigaVUE-FM by accessing the device that has been added to the GigaVUE-FM fabric manager from the GigaVUE-FM interface. The Traffic option appears in the navigation pane of the device view on supported devices.

You can access traffic operations from the GigaVUE-FM interface.

Go to **Traffic**  and select **Physical**. The following options are available:

- Orchestrated Flows
- Flow Maps
- Active Visibility
- Tunnel Monitoring

Topics:

- [Ports and GigaStreams](#)
- [Flow Mapping®](#)
- [Flexible Inline Arrangements](#)
- [Inline Bypass Solutions](#)
- [Timestamps](#)
- [Fabric Maps](#)
- [Orchestrated Configurations](#)

See also:

- [GigaSMART®](#)
- [Clusters](#)

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 but down for port 1/1/x4.

Filtered By : **None**

☐	Port Id	Alias	Type	Speed	Admin Enabled	Link Status	Transceiver Type	Utilization (Tx/Rx)	Port Filter	Discovery Protocol
☐	1/1/x1	ColaSoft_Dedicated_Link_ESX12	IT	1G	✓	up	sfp cu	0 / 0	—	Off
☐	1/1/x2	WireShark_Dedicated_Link_ESX12	IT	1G	✓	up	sfp cu	0 / 0	—	Off
☐	1/1/x3	TunnelPort_From_ESX12	N	1G	✓	up	sfp cu	0 / 0	—	Off
☐	1/1/x4		T		—	down		0 / 0	—	Off

Figure 2 Port Link Status

Ports on GigaVUE® TA Series Traffic Aggregator Nodes

In earlier software versions, GigaVUE® TA Series Traffic Aggregator nodes did not support tool ports. Instead, they supported gateway ports as displayed in [Figure 3GigaVUE-TA Packet Distribution](#) and described in [Concepts Illustrated in Figure 3GigaVUE-TA Packet Distribution](#).

All gateway ports on GigaVUE® TA Series nodes are tool ports. For details, refer to [Notes and Consideration on GigaVUE® TA Series Nodes](#).




Figure 3 GigaVUE-TA Packet Distribution

Concepts Illustrated in Figure 3GigaVUE-TA Packet Distribution

Figure 3GigaVUE-TA Packet Distribution illustrates the concept of data flows. Data arrives from different sources at the network ports on the left and is forwarded to different Gigamon nodes connected to the tool ports (formerly gateway ports) on the right.

The following are important points about setting up packet distribution on GigaVUE® TA Series nodes:

- Traffic from multiple ingress ports can be sent to the same tool port for aggregated uplink to the Gigamon Platform fabric.

In this example, the traffic from Inputs A, B, and C is all sent to the same tool port. In turn, this tool port is connected to a GigaVUE® HC Series node so that the combined traffic from these inputs is available to the full suite of Flow Mapping® tools provided by the GigaMonDeep Observability Pipeline.

- Traffic arriving at a single network port can be sent to multiple destination tool ports. Note that in [Figure 3GigaVUE-TA Packet Distribution](#), the traffic arriving on **Input D** is sent to two different tool ports.
- Filters can be applied to tool ports:
Filters applied to tool ports, inline network hybrid and circuit ports are called **egress-filters**. Egress-filters are useful if you want to send the same traffic to multiple tool ports and have each one allow or deny different packets based on specified criteria. You can use up to **20 egress-filters** at a time on GigaVUE® TA Series nodes.

NOTE: In [Figure 3GigaVUE-TA Packet Distribution](#), egress-filters are set to focus on different parts of the data stream arriving at **Input D** – traffic on a VLAN range, a subnet range, and so on.

Notes and Consideration on GigaVUE® TA Series Nodes

GigaVUE® TA Series nodes support network, tool, stack, and hybrid port types. Refer to the following notes and considerations for GigaVUE® TA Series nodes (including GigaVUE-TA100, GigaVUE-TA200 and GigaVUE-TA400):

- Gateway ports on GigaVUE TA Series nodes are removed and converted to tool ports. In addition, gateway mirrors are removed and converted to tool mirrors.
- Tool ports on GigaVUE® TA Series nodes can continue to be used to aggregate traffic (as displayed in [Figure 3GigaVUE-TA Packet Distribution](#) and described in [Concepts Illustrated in Figure 3GigaVUE-TA Packet Distribution](#)).
- Tool ports on GigaVUE® TA Series nodes can also be used to directly connect to tools, such as firewalls, Intrusion Prevention Systems, or Application Performance Monitors.
- Hybrid ports are fully supported in both standalone and cluster mode on GigaVUE® TA Series nodes.
- GigaVUE® TA Series nodes can continue to be clustered with GigaVUE® HC Series nodes.
- When GigaVUE® TA Series nodes are in a cluster, bidirectional traffic flow is enabled on the stack links of GigaVUE® TA Series nodes.
- Map rules using GigaVUE® TA Series tool ports in the egress direction are supported.

Hybrid Ports

Hybrid ports are created by creating a dual function tool port. A physical tool port is set as a virtual network port which can then send traffic to other tool ports using secondary maps. A hybrid port is operated in loopback mode. This is only available if the GigaVUE H series node

is upgraded to minimum of 4.2 release. For more details on how to setup hybrid ports and the caveats, refer to the *GigaVUE-OS CLI Reference Guide*.

Stack Ports

Stack ports are used to carry traffic arriving at a network port on one GigaVUE node to a tool port on another GigaVUE node in a cluster.

Inline Network Ports

Inline networks, inline tools, and inline maps work together to form an Inline Bypass solution. The Inline Bypass solution has an overall state, which can change in response to hardware conditions and user configuration. Inline network ports are ports to which end-point devices are attached in an Inline Bypass solution.

NOTE: Inline network ports are supported only on GigaVUE-HC1, GigaVUE-HC1P, GigaVUE-HCT, GigaVUE-HC3, GigaVUE-TA25, GigaVUE-TA200, GigaVUE-TA400, and GigaVUE-TA400E.

Inline Tool Ports

Inline tool ports are ports to which inline tools are attached in an Inline Bypass solution.

NOTE: Inline tool ports are only supported on GigaVUE-HC1, GigaVUE-HC1P, GigaVUE-HCT, GigaVUE-HC3, GigaVUE-TA25, and GigaVUE-TA200, GigaVUE-TA400, and GigaVUE-TA400E.

Circuit Ports

Required License: Advanced Feature License for GigaVUE® TA Series Nodes

Circuit ports are used to send or receive traffic between two clusters. The circuit ports are configured at the sending and receiving ends of two clusters and the clusters are connected through a circuit tunnel. Circuit ports send or receive only the traffic that is tagged with a circuit-ID. In a map, if a circuit port is used as a source port, it acts as a network port, and decapsulates the traffic that contains a circuit-ID. If a circuit port is used as a destination port, it acts as a tool port, encapsulates the traffic, and strips the circuit-ID.

Circuit ports are supported on the following:

- All GigaVUE HC Series and TA Series nodes.
- As a source port in a regular map and as a destination port in a regular collector map.

- GigaStreams, port filter, and port groups.

GigaSMART Engine Ports

GigaSMART Engine ports are used when configuring GigaSMART groups. These ports cannot be edited. On the Ports page, the GigaSMART engine ports populates only the Port ID, Type, and Link Status fields.

Port Lists

Many map commands require a port-list (for example, rule and shared-collector arguments all require them). You can define the port lists using any combination of port IDs and port aliases. In GigaVUE-FM, port lists are created in the **Source** and **Destination** fields when editing or creating a new map. The following are considerations when creating a port list:

- When creating a **Pass All** map, you can specify a network port list, hybrid port list, or an inline network alias in the **Source** field. In the **Destination** field for a **Pass All** map, you can specify a tool port list, hybrid port list, an inline tool alias, an inline tool group alias, or an Inline Bypass.
- Circuit ports are supported as source ports on Regular maps and as destination ports on Regular Collector maps.
- The **Source** and **Destination** fields lets you select multiple non-contiguous ports. To enter port IDs in a list, simply select the port from the drop-down list after clicking in the field. If the port has a alias, it is shown in the list along with the ID.
- GigaSMART load balancing port groups can have ports with different rates.

Port Aliases

GigaVUE-FM lets you configure textual aliases for all port types. Aliases can be used in place of the numerical **bid/sid/pid** identifier required in many packet distribution.

To set up Port Aliases in GigaVUE-FM:

1. Select **Ports > Ports > All Ports**
2. Select the **Port ID** that needs an alias
3. Select **Edit**.
4. In the **Alias** field, enter an alias for the port, and the click **Save**.

Work with Hybrid Ports

A hybrid port is a physical port that has a dual function as an indirect traffic source port and a tool port. Hybrid means that a network port (ingress) can become a tool port (egress) to which map rules can be applied. Hybrid ports are introduced in software version 4.2. Hybrid ports are supported on GigaVUE HC Series nodes and GigaVUE TA Series devices.

A hybrid port is operated in loopback mode. The network data coming from the internal loopback is available to be used in maps.

Hybrid ports help alleviate the number of ports needed. For example, without hybrid ports, if you had traffic coming in with an MPLS header, but wanted to filter on a particular subnet, you would create a map to remove the MPLS header, physically loop the traffic back from the tool port to a new network port, and create another map to filter on the subnet. This same functionality can now be accomplished with hybrid ports.

If you have been using IP/UDP tunneling to encapsulate whole Ethernet frames and want to filter packets to destination tool ports after being decapsulated by GigaSMART, you can now use hybrid ports.

Hybrid ports can also be used to duplicate traffic from a network source. Using hybrid ports, you can create maps in parallel. For example, all HTTP traffic can be sent to one tool port unmodified and the same HTTP traffic can be sent to another tool port sliced at 100 bytes.

Using hybrid ports, you can create maps in a GigaSMART chaining.

As soon as a hybrid port is configured, it is internally changed to loopback mode. This means that the link is *Up* with or without SFPs inserted. (If SFPs are not inserted, the traffic runs at the maximum speed supported.) Traffic flows out of a hybrid port (Tx direction) and the duplicated flow loops back to it (Rx direction). This is similar to tool mirrors.

WARNING: Do not connect cables to hybrid ports coming from network ports. All cabling attached to hybrid ports must be attached to tools.

When a port is configured as a hybrid port type, it can be used as follows:

- as a map source and destination (for regular maps, as well as map-passall, and map-scollector)
- in a GigaStream
- in a port group
- with an egress port filter

Maps using hybrid ports, regardless of source or destination, can be applied to a GigaSMART operation.

When using hybrid ports, consider the following:

- Be aware not to configure traffic loops, such as $H1 \rightarrow H2$, $H2 \rightarrow H3$, $H3 \rightarrow H1$. Do not use the same hybrid port as ingress as well as egress on all maps, such as $H1 \rightarrow H1$.
- 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.

The screenshot shows the 'Quick Port Editor' interface. At the top, there is a search field containing '1/1/x10'. Below it is a table with columns: Port Id, Alias, Type, and Admin. The first row shows '1/1/x10' in the Port Id column, 'port alias' in the Alias column, and a dropdown menu in the Type column. The dropdown menu is open, showing options: Tool, Network, Tool, Hybrid (highlighted), Stack, Inline Network, and Inline Tool. In the Admin column, there is a checkbox labeled 'Enable' which is checked.

Port Id	Alias	Type	Admin
1/1/x10	port alias	▼ Tool Network Tool Hybrid Stack Inline Network Inline Tool	☒ Enable

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 with multiple tool ports or port groups.
- In cases of inline network LAG and inline network groups, the port filters must be applied on each of the inline network ports that are part of the inline network LAG or inline network group.
- Port filters for inline network ports are supported on GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA400, and GigaVUE-TA400E.
- The GigaVUE-TA25 and GigaVUE-TA25E ports cannot be part of destination ports of first level maps if the source port is on another node (i.e combination of VPort and GigaVUE-TA25 and GigaVUE-TA25E destination port in the “to” ports list) in legacy cluster.
- The outer VLAN tool port filter cannot be used to match the ingress VLAN tag configured on the source port in GigaVUE-TA400 in 5.14 release.
- IP fragmentation tool port filter is not supported on GigaVUE-TA400 in 5.14 release.

- The following limitation is applicable only for double tag mode (software version 5.14.00 onwards).

Egress port filters are supported on GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-HCT, and GigaVUE-HC1-Plus, except that

- VLAN rules are not supported with port filters, and
- Either IPv4 or IPv6 type port filter rules are supported only if L2 circuit encapsulation tunnels or GigaSMART maps are used else both IPv4 and IPv6 rules are supported.

Port-Filter Maximums

Table 1: Port-Filter Maximums per GigaVUE Node provides the maximum port-filters for the different GigaVUE nodes:

Table 1: Port-Filter Maximums per GigaVUE Node

GigaVUE Node	Maximum Number of Port-Filters
GigaVUE-HC1	• 448 for IPv4 rules • 255 for IPv6 rules • 448 for IPv4+IPv6 Pass rules.
GigaVUE-HC3 (CCv1 and CCv2)	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.

GigaVUE Node	Maximum Number of Port-Filters
GigaVUE-TA200	NOTE: For an IPv4 and IPv6 combination the maximum filters allowed is 448. In such combination the maximum limit is 254 for IPv4 filters and 255 for IPv6 filters. While configuring an IPv4 + IPv6 combination ensure that the individual filter limits are not crossed.
GigaVUE-TA200E	
GigaVUE-TA400	Without Advanced Feature License: • 20 per chassis With Advanced Feature License: • 448 per pseudo slot
GigaVUE-TA400E	Without Advanced Feature License: • 20 per pseudo slots With Advanced Feature License: • 475 for IPv4 rules per pseudo slot • 475 for IPv6 rules per pseudo slot.

NOTE: A single filter applied to multiple tool ports counts multiple times against the 100-filter limit.

How to Apply Port Filters

To apply a port filter, do the following:

1. From the device view, go to **Ports > Ports > All Ports**.
2. Select the egress port (tool, hybrid, circuit, and inline network) to which you want to apply a filter, and then click **Edit**.
3. Under the Filters section on the Ports page, click **Add Rule**.
4. Select and configure any of the following required rule:

Table 2: Port-Filter Rule

Rule	Action
circuit-id	Configure circuit id
Description	Add a description to the Map Rule
dscp	Configure DiffServ Code Point bits

Rule	Action
ethertype	Configure Layer 2 ethernet type
ip6dst	Configure destination IPv6 address
ip6src	Configure source IPv6 address
ipdst	Configure destination IPv4 address
ipfrag	Configure IP fragmentation bits
ipsrc	Configure source IPv4 address
ipver	Configure IP version number
l2gre-id	Configure l2gre id
macdst	Configure destination MAC address
macsrc	Configure source MAC address
portdst	Configure destination port number or port range
portsrc	Configure source port number or port range
protocol	Configure internet protocol 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:
 - o Use a **drop** rule to deny packets matching the specified criteria.
 - o 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:

The screenshot shows the GigaVUE-FM interface for device 'gigamon-a302dd (H Series)'. The left sidebar has a 'TRAFFIC' section with 'Ports' highlighted (1). The main area shows a table of ports. The 'All Ports' tab is selected (2). The port '1/1/x10' is selected (3). A quick view for 'Port: 1/1/x10 - @ gigamon-a302dd' is open (4), showing filter rules and statistics.

Port Id	Alias	Status	Type	Speed	Admin
1/1/x12		✓ Po...	N		Disable
1/1/x11		✓ Po...	N		Disable
1/1/x10		✓ Po...	T	10G	Enabled
1/1/x9		✓ Po...	T		Disable
1/1/x8		✓ Po...	N		Disable
1/1/x7		✓ Po...	N		Disable
1/1/x6		✓ Po...	T		Disable
1/1/x5		✓ Po...	T		Disable
1/1/x4		✓ Po...	N		Disable
1/1/x3		✓ Po...	N		Disable
1/1/x2		✓ Po...	N		Disable

Type	Number of Filters	Statistics
Power	6	
Vendor		
Product No		
Serial No		
Temperature		
Pass Rules		
Drop Rules		

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:

The screenshot displays the GigaVUE-FM interface for 'Box ID 1 - GigaVUE-HC2'. The left navigation pane has 'Chassis' selected (1). The main content area shows a table of chassis components (2) and a detailed view of Slot ID 1 (3).

Slot ID	Hardware Type	Configured	Health Status	Operational
1	PRT-HC0-X24	✓	Slot is health...	Up
2	PRT-HC0-X24	✓	Slot is health...	Up
cc1	HC2-Main-Board	✓	Slot is health...	Up

Slot ID	Hardware Type	Board(°C)	Exhaust(°C)	Intake(°C)
1	PRT-HC0-X24	24	-	-
2	PRT-HC0-X24	27	-	-
cc1	HC2-Main-Board	28	33	28

Power Module Id	Hardware Type	Status
1	HC2-Power-Supply-PDB	✓
2	HC2-Power-Supply-PDB	✓

Hardware Type	Configured	Status	Fabric Hash
PRT-HC0-X24	Yes	Up	N/A

Filter Template	Alias	Limit	Qualifiers
✓	defaults		

Filter Resource	Map Rule	Port Filter	App-Filter	In-use qualifiers
Limit	4000	100	504	
Used	0	2	0	

App	Reserved	Reserved Used	Free Pool Used
discovery	0	0	0
mapSrc	50	0	0
mapRule	0	0	0
gsd	0	0	0
ipinterface	0	0	0

Hardware Rev: A1-a2
PLD Current Revision: N/A
Product Code: 132-008D

Status of Line Cards/Nodes and Ports

You can review the current status of the node, line cards, modules, and ports through either the CLI or the GigaVUE-FM. This will ensure that all units have been properly configured and that the node is ready for further configuration.

To check the line cards/modules and ports with GigaVUE-FM, using the Ports page or the Chassis page (select **Chassis** in the navigation pane).

How to Check Port Status with Ports Page

To check the port status with the Ports page, do the following:

1. Select **Ports > Ports > All Ports** to open the Ports page.
2. Locate the port to check by entering the port ID or port alias in the search field.
3. If you need to change the port type or enable the port:
 - a. Click **Quick Port Editor**.
 - b. In the Quick Port Editor, enter the port ID or port alias in the Quick search field to find the port.

- c. Set the port type by selecting type from the drop-down list in the **Type** field. enable the by selecting **Enable** as needed.

How to Check Port Status with Chassis Page

To check the status of the ports and cards with the Chassis page, do the following:

1. Select **Chassis** from the navigation pane to open the Chassis view shown in [Figure 4Chassis Page](#).
2. Use the view buttons on the Chassis page to check the status of the cards as well as the ports. When viewing a node in cluster, there is a drop down option to select a specific node in a cluster configuration.

For details about the Chassis page, refer to the “Chassis” section in the *GigaVUE Administration Guide*.




Figure 4 Chassis Page

Managing Ports

The Ports pages allows you to manage and configure ports for various functions. In case of inline supported devices, the pages for managing Inline Bypass ports are available by selecting Inline Bypass from the main navigation pane.

Before beginning with managing and configuring ports, make sure that the user role, which is assigned to a User Group, has the permission for the specific ports on the system. For details on the port-based access levels, refer to the *GigaVUE-OS CLI Reference Guide* and “Managing Roles and Users” in the *GigaVUE Administration Guide*.

This section provides a description of the Ports pages in the GigaVUE-FM UI. It covers the following topics:

- [Ports](#)
- [Port Groups](#)
- [Port Pairs](#)
- [Tool Mirrors](#)

- [Stack Links](#)
- [IP Interfaces](#)
- [Circuit Tunnels](#)

NOTE: Starting in software version 5.5.01, any change in the port health status is indicated immediately in the following port pages:

- Ports > All Ports
- Port Groups > All Port Groups
- Port Groups > GigaStream
- Port Pairs
- Tool Mirrors
- Stack Links
- IP Interfaces

The link in the notification opens the port quick view.

Ports

The Ports tab lets you select the All Ports and Ports Discovery pages. You can also control which ports display.

All Ports

The All Ports page displays when you select All Ports. The Ports page shows a table with detailed information about each port ID on a specific device. Only the GigaVUE HC Series and GigaVUE TA Series devices are presented in the Port Page view as shown in [Table 3: Descriptions of Ports Page Columns](#). You can control which ports display on the page by selecting a set of filters or configure the ports through the Quick Port Editor or selecting Edit for a selected port. For details about filtering ports, refer to [Port List Filter](#). For details about the Quick Port Editor, refer to [Quick Port Editor](#).




Figure 5 Ports Page

The port type determines which columns are populated with data in the table. The columns are populated as follows:

- **Engine** ports populate the Port ID, Type, and Link columns.
- **Network** ports populate the Port ID, Alias, Type, Speed, Admin Enabled, Link Status, Transceiver Type, Utilization, Port Filter, and Discovery Protocol.
- **Tool** port populate the Port ID, Alias, Type, Speed, Admin Enabled, Link Status, Transceiver Type, Utilization, Port Filter, and Discovery Protocol.
- **Stack** port populates Port ID, Alias, Type, Admin Enable, Link Status, Utilization, Port Filter, and Discovery Protocol.
- **Hybrid** port populates Port ID, Alias, Type, Admin Enable, Link Status, Utilization, Port Filter, and Discovery Protocol.
- **Circuit** port populates Port ID, Alias, Type, Admin Enable, Link Status, Utilization, Port Filter, and Discovery Protocol.
- **Inline Network** port Port ID, Alias, Type, Admin Enable, Link Status, Utilization, Port Filter, and Discovery Protocol.
- **Inline Tool** port Port ID, Alias, Type, Admin Enable, Link Status, Utilization, Port Filter, and Discovery Protocol.

NOTE: Not all port types are supported on all platforms. Inline network and inline tool ports are supported on GigaVUE HC Series nodes, GigaVUE-TA25, GigaVUE-TA25E and GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA400, and GigaVUE-TA400E.

[Table 3: Descriptions of Ports Page Columns](#) provides descriptions of the columns on the ports page.

Table 3: Descriptions of Ports Page Columns

Column	Description
Port ID	The port number is in <box ID>/<slot ID>/<port_ D> format in the CLI and GigaVUE-FM. For the GigaVUE TA Series slot ID is always 1 as they are not modular. For the GigaVUE-HC1, and GigaVUE-HC3, the line cards or modules are identified by the slot number. In a standalone (default) configuration, box ID is always designated as 1 but can be changed through the CLI (it cannot be changed through GigaVUE-FM). In a cluster configuration, the box ID can vary.
Alias	Alias name of the port, if any.
Status	Health status of the port. Note: When a ports administrative status changes from 'enabled' to 'disabled', the health status of the administratively disabled ports remains in Green. However, the health status of administratively disabled port is indicated as NA if Exclusion Rules are enabled for the ports that are admin disabled in the Alarms Page. Refer to the 'Manage Alarms' section in the GigaVUE Administration Guide for more details.
Type	List the type of port such as network, tool, stack, inline network, inline tool, circuit, or hybrid. You can set the port type through the Quick Port Editor or selecting the port on the Port page and clicking Edit. The port type is set by selecting the type in the Type field.
Speed	Current setting for the port's speed.
Admin	Indicates whether the port is administratively enabled or disabled.
Force Link Up	Indicates the 'force link up' setting for the port. When enabled, this option forces connection on the optical port.
Ude	Indicates whether the port is enabled for unidirectional (Ude) or bidirectional traffic. Enabled means unidirectional; Disabled means bidirectional. Note: Do not disable UDE on Out Ports when using BiDi RX/TX optics connecting towards passive TAPs. For example, if BiDi RX/TX optics are connected to both the Tap-facing and Network-facing ports of the TAP-M506/506T, disabling UDE on the TAP-facing ports will cause a production link failure.
FEC	Configures forward error correction (FEC) on the port to ensure error-free traffic over long distance. The values are: • CL91—Enables FEC on the port and supports 25Gb, 100Gb, and 4x25Gb speeds. • CL74— Enables FEC on the port and supports 25Gb, and 4x25Gb speeds. • CL108 — Enables FEC on the port and supports 25Gb speed only on GigaVUE-TA25 and GigaVUE-TA25E. • OFF—Disables FEC on the port and supports 25Gb, 100Gb, and 4x25Gb speeds.
Link Status	The current status of the link connected to the port, either port link up or port link down.
Transceiver Type	The type of transceiver installed in this port.
SFP Power	SFP power for copper transceivers

Column	Description
	Note: When a new device is added to GigaVUE-FM, it takes one stats cycle for the SFP power value to be reflected in the GigaVUE-FM GUI.
Port Filter	Indicates if the egress port filter (tool, hybrid, circuit, and inline network), is associated with this port.
Discovery Protocol	Protocol used to discover neighboring nodes using CDP or LLDP. This feature is available for network, tool and circuit ports.
Box Hostname	Host name of the device
Gigamon Discovery	Indicates if Gigamon Discovery protocol is enabled
Tags	Tag associated with the port.

Port Quick View

The Quick View for ports displays when you click on a row in the Ports page to quickly get more information about a specific port. The quick view shows the port properties, statistics information on receiving (Rx) and/or Transmitting (Tx) ports and alarms information. The quick view also shows a graphical representation of port statistics. Refer to [Figure 6 Ports 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 7 Port List Filter](#) where you can specify how to filter ports displayed on the Ports page.

Ports All Ports Ports Discovery Fabric Statistics Header Stripping Statistics

Total Ports: 59 | Filtered By: none Mar 29, 2020 17:18:16

☐	Port Id	Alias	Status	Type	Speed	Admin
☐	1/1/x4		Port is down	T	10G	Inactive
☐	1/1/x7		Port is down	T	10G	Inactive
☐	1/1/x10		Port is healthy	N		Inactive
☐	1/1/x11		Port is healthy	T		Inactive
☐	1/1/x13		Port is healthy	N		Inactive
☐	1/1/x14		Port is healthy	N		Inactive
☐	1/4/e1		Port is healthy	E		Inactive
☐	1/1/x3		Port is down	T	10G	Inactive
☐	1/1/x5		Port is down	T	10G	Inactive
☐	1/1/x8		Port is down	T	10G	Inactive
☐	1/4/e2		Port is healthy	E		Inactive
☐	1/1/x1		Port is down	N	10G	Inactive

1 < Go to page: 1 of 5 > > Total Records: 59

Filter [Clear]

Box ID/Slot ID
Select a Box/Slot ID

Port Alias
Type Port Alias

Port ID
Type port #

Type
Select Port Type...

Admin Status
☒ All ☐ Enabled ☐ Disabled

Link Status
☒ All ☐ Up ☐ Down

Speed
Select Port Speed»

Transceiver Type
Select a Transceiver Type»

Tags

Figure 7 Port List Filter

The criteria that you can use to filter the port list is as follows:

Criteria	Description
Box/Slot ID	Display only those ports that match the specified box and slot IDs.
Port Alias	Display port with the specified alias.
Port ID	Display ports with specified number in the port ID. For example, if you specify 3 the result will also display ports that include the number 3, 13, 23, 30, and so on.
Type	Display ports with the specified port type. Select one of the following: • Network • Tool • Inline Network • Inline Tool • GigaSMART • Hybrid • Circuit • Stack
Admin Status	Display ports based on their current admin status. The possible selections are: All — display ports with a status of Enabled or Disabled. This is the default. Enabled — display ports with admin enabled Disabled — display ports with admin disabled
Link Status	Display ports based on their current link status. The possible selections are: All — display ports with a status of Up or Down. This is the default.

Criteria	Description
	Enabled — display ports with a link status of up. Disabled — display ports with a link status of down.
Speed	Display ports with the selected port speed. The port speeds available depend on the node.
Transceiver Type	Display ports with the selected transceiver type. The transceivers available selection depend on the type of transceivers connected to the ports.
Tags	Display ports associated with the selected tag key and tag value.

To filter the ports, enter the information to use for filtering the ports and select the radio buttons. For example, in [Figure 8 Filtering 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.

The screenshot shows the 'Ports' page in the GigaVUE interface. The 'Filter' overlay is active, showing the following settings: Box ID/Slot ID (dropdown), Port Alias (text input), Port ID (text input), Type (dropdown set to 'Network'), Admin Status (radio buttons: All, Enabled (selected), Disabled), Link Status (radio buttons: All (selected), Up, Down), Speed (dropdown set to '10G'), Transceiver Type (dropdown), and Tags (text input). The 'Ports' table shows 2 filtered ports, both with status 'Port is down' and 'Inactive'.

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 9 Filtered Ports List](#) shows the Port pages with two ports that correspond to the current filters: Network Type and Admin Status Enabled.

Total Filtered Ports : 2 Clear Filter					
☐	Port Id	Alias	Type	Speed	Admin Enabled
☐	10/1/x1		N	10G	✓
☐	10/1/g1	1G	N	1G	✓

Figure 9 Filtered Ports List

Quick Port Editor

From the Ports page, you can open the Port Type Editor to quickly change the port types in a chassis. To set the port type for ports in a chassis, do the following:

- 1. Click **Quick Port Editor**.
- 2. For each port on which you want to set the port type, select the type from the drop-down list. In [Figure 10 Port Type Selection](#), port 1/1/x9 is being changed from a network port to at tool port.

To find a specific port, you can use the Quick Search to find a specific port by entering the port ID or alias in the Quick search field.




Figure 10 Port Type Selection

- 3. To enable the port, select **Enable**.
- 4. Click **OK**.

Each port can also be assigned an alias. Any port types set in the CLI or through the GigaVUE-FM APIs are reflected on this page. For more information, refer to [Port Aliases](#) for port aliases and to the *GigaVUE-FM REST API Reference in GigaVUE-FM User's Guide* and the *GigaVUE-FM Reference for APIs*.

Configure Ports

From the Ports page, you can either configure or edit a specific port by selecting a port and clicking the **Edit** button on the Ports page or on the Quick Port Editor. [Table 4: Port Configuration Options](#) describes the options on the configuration page.

Table 4: Port Configuration Options

Field	Description
Alias	The alias configured for this port, if any. Aliases can be used in place of the numerical bid/sid/pid identifier required in many packet distribution commands in the CLI. For example, instead of configuring a connection between, say, 1/1/x1 and 1/2/x4, you could connect Gb_In to Stream-to-Disk. Note that aliases can only be applied to single ports. They cannot be applied to groups of ports. Port alias can be up to 128 characters long including special characters. Aliases are case sensitive.
Admin	Check to enable the port.
Type	Specifies whether the port is configured as an Inline Network port, Inline Tool port, Network port, Tool port, Stack port, Circuit port, or Hybrid port.
Speed	Specifies the speed for the selected port. For copper ports, you can click to change the speed as long as Auto Negotiation is disabled.
Duplex	Specifies the port's duplex configuration. Only full duplex is supported. Starting in software version 5.2, half duplex support is removed from all GigaVUE nodes. If half duplex was configured in a previous software version, it will remain intact following the upgrade to 5.2 or higher release. Update to full duplex, if required.
Auto Negotiation	Select to enable auto-negotiation for the selected port. When auto-negotiation is enabled, duplex and speed settings are ignored. They are set through auto-negotiation. For 1Gb fiber ports, auto-negotiation is not supported on Gigamon Platforms..
Force Link Up	When enabled, this option forces connection on an optical port. Use this option when an optical GigaPORT tool port is connected to a legacy optical tool that does not transmit light. This option is not available for 10Gb capable ports with a 1Gb SFP installed.
Ude	When selected, this option indicates the port is unidirectional (UDE). When deselected (disabled), the port is bidirectional. UDE is enabled by default. Note: This option is available for GigaVUE-HC3, GigaVUE-TA100, GigaVUE-TA200, GigaVUE-TA25, GigaVUE-TA400, and GigaVUE-TA400E platforms with 100Gb BiDi (QSB-512) transceiver. If used with passive taps, ports used for monitoring should be set to Network port with UDE enabled. Important: If you clear the UDE check box, the laser will start to transmit, which may affect the remote connectivity.
FEC	Configures forward error correction (FEC) on the port to ensure error-free traffic over long distance. The values are: CL91—Enables FEC on the port. CL74— Enables FEC on the port .

Field	Description
	CL108 — Enables FEC on the port for GigaVUE-TA25. OFF—Disables FEC on the port. This option is available only on 25Gb and 100Gb transceivers.
Timestamp	Use the timestamp options when a GigaPORT-X12-TS line card is installed. For details about the GigaPORT-X12-TS line card, refer to the <i>GigaVUE-OS CLI Reference Guide</i>. The timestamp options are as following: Append Ingress—Use this option to add a timestamp to ingress packets a GigaPORT-X12-TS ports. This applies to ports x1..x12 when configured as network ports. Strip Egress—Use this option to strip timestamps from egress packets. Source ID Egress—Use this option to specify a custom source ID to be included in the timestamp appended by the GigaPORT-X12-TS. The source ID identifies the ingress port on the GigaVUE HC Series node where the timestamped packet arrived. The timestamp always includes a source ID. If you do not specify a custom value, the GigaPORT-X12-TS generates one automatically using the following formula: (Box ID * 2048) + (Slot ID * 256) + Port Number Important: Only apply the Strip Egress option to packets with time stamps appended. The strip egress feature strips the last 14 bytes of each packet regardless of whether a timestamp has been added.
VLAN Tag	Use VLAN tags to identify, differentiate, or track incoming sources of traffic. When the traffic reaches the tools or the maps, you can filter on the VLAN tags for the corresponding ports you want to measure. The port must be a network or inline-network or hybrid type of port. Ingress port VLAN tagging is supported for IPv4 and IPv6 packet types, including non-tagged packets, tagged packets, and Q-in-Q packets. Ingress port VLAN tagging is not supported on inline network ports, hybrid ports, or on network ports that are connected via port-pairs. The same VLAN tag can be assigned to multiple network ports. However, each port can only have one VLAN tag. VLAN tagging is supported in a cluster. • To add, VLAN IDs for a Port, enter the VLAN ID in this field. • To modify, update the VLAN ID in this field and Save. It will take effect. • To delete, remove any values for the VLAN ID in this field and Save. It will remove the VLAN Tag from this Port. VLAN tags are only available on network ports. Note: On GigaVUE-TA25 and GigaVUE-TA25E, ingress VLAN tagging is not supported on network ports associated with GSOP maps, and GSOP maps are not supported on network ports configured with ingress VLAN tag. This limitation is not applicable to the E-tag mode.
TPID	Select the TPID for the VLAN Tag. The default value of TPID is 0x8100. You can select the other supported values 0x9100 and 0x88a8 from the drop-down list.
Port Discovery	Select to enable discovery of neighbors associated with the port. Neighbor discovery is only available on network ports
Discovery Protocols	When port discovery is enabled, use the Discovery Protocol options to set up CDP or LLDP or both (All) on the port. The results are shown on the Ports Discovery page.
Buffer Threshold	Specifies the alarm buffer threshold on a port. You can specify the alarm buffer

Field	Description
	threshold in the Rx and Tx directions on network and stack type ports and in the Tx direction on tool type ports. By default, the threshold is set to 0 , which disables the threshold
Utilization Threshold	Sets the utilization percentage for this port at which the GigaVUE HC Series node will generate high or low utilization alarms for the port. For more information about port utilization, refer to Monitor Port Utilization . By default, the threshold is set to 0 , which disables the threshold.
Lock Port	Restricts use of the port for only your user account as follows: • Users with the admin role can lock any port in the system. Users with the Default/Operator role assigned can only lock ports to which their account has been granted access. • Administrators can lock a port for another user by including the optional user. You can optionally share a locked port by specifying users in the Lock shared with Users field or selecting users to share the lock with through their assigned roles. For more information about who to set lock sharing, refer to Managing Ports .
Tags	Select the required tag key and tag value to which the port must be associated to. The tag key and the tag value will be displayed depending on the role and the corresponding access rights of the user.

Ports Discovery

The Ports Discovery page displays the port neighbor information for each port that has discovery enabled. For each network port, tool port or circuit port on which discovery is enabled, neighbor information is collected. Information for up to five of the most recent neighbors is retained for each port.

The following are limits on the amount of discovery information that is retained:

- For each port, discovery information for a minimum of two neighbors and a maximum of 20 neighbors is retained.
- For a chassis, discovery information for a maximum of 2K neighbors is retained.

Neighbor information is removed or replaced as follows:

- When the neighbor information expires due to the TTL.
- When the number of neighbors for the chassis reaches the 2K maximum and a new neighbor is discovered. In this case, the following can occur:
 - If there are currently two or more discovered neighbors for a port, the newly discovered neighbor replaces the neighbor information for the least recently updated neighbor.
 - If there are currently less than two discovered neighbors for a port, the newly discovered neighbor is added (actually exceeding the 2K limit to guarantee a minimum of two neighbors per port).

NOTE: Aging (the discovery protocol time-to-live) determines how long neighbor information is valid.

For information about the discovery protocols and enabling port discovery, refer to [Port Discovery](#)

Statistics

The Statistics page displays the statics for all the ports on the node, providing the following information about a packets transmitted or received on a port:

Column	Definition	Notes
Octets (Rx/Tx)	The count of packets/bytes received and transmitted by this port.	Error packets are not transmitted, therefore they are not counted. Excludes undersize frames.
Octets/sec (Rx/Tx)	The count of packets/bytes received and transmitted by this port per sec.	Error packets are not transmitted, therefore they are not counted.
Unicast Packets (Rx,Tx)	The count of packets/bytes received and transmitted by this port.	Excludes multicast packets, broadcast packets, packets with FCS/CRC errors, MTU exceeded errors, oversize packets, and pause packets.
Non-Unicast Packets (Rx/Tx)	Total Non-unicast packets received or transmitted.	
Packets/sec (Rx/Tx)	The rate which packets are received or transmitted.	
Packet Drops (Rx)	Total Dropped Packets	Packets are dropped when a network port's bandwidth is exceeded due to oversubscription. Packets are dropped when they reach the port but before they are sent out.

Column	Definition	Notes
Discards (Rx/Tx)	Total received and transmitted packets discarded. This counter increments when a packet is discarded at the tool port due to egress port filter.	Discards are counted in the following cases: Traffic arriving at a network port that is not logically connected using a map or map passall. Map rules applied on a network port. In packets on a tool port. Pause frames.
Errors (RX/TX)	Total Error Packets Received or Transmitted. Error packets include undersize, FCS/CRC, MTU exceeded, fragments, and oversize packets.	Excludes oversize packets without FCS/CRC. Packets larger than the MTU setting arriving on a network port are counted twice in the counter. So 1000 oversize packets would show up as 2000. This double-counting only happens with Oversize error packets.
Utilization (Rx/Tx)	Percentage of port utilization by packets received or transmitted	

Port Groups

The Port Groups selection in the top navigation bar provides access to the All Port Groups and GigaStream pages, for creating port groups and GigaStreams, respectively.

All Port Groups

Selecting **Port Groups** under **Port Groups** opens the **Ports Group** page by default. This page is used to create a port group, which can simplify administration of GigaVUE ports.

Administrators can create groups of ports that can then be quickly assigned to different user groups. With clustered ports potentially numbering into the hundreds, port groups provide a useful shorthand when assigning multiple ports to different user groups. (To create user groups, select **Roles and User** from the navigation pane, and refer to *Managing Roles and Users* in the *GigaVUE Administration Guide* for more details.) The following are the different types of port groups:

- **Network Port Group**—contains only network ports.
- **Tool Port Group**—contains only tool ports or tool GigaStream, which is a combination of multiple tool ports.
- **Hybrid Port Group**—contains only hybrid ports or hybrid GigaStream, which is a combination of multiple hybrid ports.
- **Circuit Port Group**—contains only circuit ports or circuit GigaStream, which is a combination of multiple circuit ports.

- **Load Balancing Port Group**—contains tool ports for load balancing. The maximum number ports allowed for port load balancing is 16.

However, port groups that include GigaStream can only be used with GTP Overlap Flow Sampling maps. For more information about GTP Flow Sampling, GTP Whitelisting and GTP Overlap Flow Sampling maps, refer to [GigaSMART GTP Whitelisting and GTP Flow Sampling](#).

Create Port Groups

To create a Port Group, do the following:

1. Select **Ports > Port Groups > Port Groups**.
2. Click **New**.
The Port Group configuration page appears.
3. Configure the Port Group:
 - a. Enter an alias in the **Alias** field.
 - b. (Optional) Enter a description in the **Description** field.
 - c. Select the Port type from Network , Tool , Hybrid or Circuit.
 - d. Click in the **Ports** field and select the ports for this port group .

NOTE:

- Use only ports from the same map when you create a Network Port Group. The port group won't work as expected if you select ports from different maps. Even when you use ports from the same map, ensure you use the exact set of ports defined in the map. Do not modify the port set when creating the port group.
- The port should not be part of the IP interface.
- Before upgrading to the latest version, you must remove either the port group or the IP interface if the ports are identical.

- e. (Optional) Click in the **Tags** field and select the tags for this port group.
4. Click **Apply**.

The Port Group created is added to the list view.

Edit Port Group

You can edit the optional fields in the Port Group such as Description, Ports, Tags, Tunnel endpoints etc. To edit, do the following:

1. Select **Ports > Port Group**. The Port Group page displays .

2. Select a Port Group and then click **Actions > Edit**.

NOTE: To edit a port group that's part of a map, first remove it from the map. You can't make changes to the port group while it's associated with a map.

3. After editing your configurations ,click **Apply**.

Clone Port Group

In some cases, you may want to create a Port Group that is similar to an existing one. To do this use the Clone feature.

1. Select **Ports > Port Group**.
2. Select the Port Group that you want to copy, and the click **Actions > Clone**.
3. Enter a new alias in the **Alias** field.
4. (Optional) Add or update Description in the **Description** field.
5. (Optional) Click in the **Ports** field and select the ports for this port group .
6. (Optional) Click in the **Tags** field and select the tags for this port group.
7. Click **Apply**.

Use the following buttons to manage the Port Groups.

Button	Description
Tags	Use to associate tags to Port Groups and to remove the tags from Port Groups when not required. To add tags: 1. Select the required Port Groups to which the tags must be applied. 2. Click the Tags drop-down menu button and select Add. 3. In the Add Tags to Resources page, select the required Tag Keys and Tag Values. 4. Click Ok to save the configuration. The Port Group is associated with the selected tags. To remove tags: 1. Select the required Port Groups that are already associated with tags. 2. Click the Tags drop-down menu button and select Delete. 3. In the Delete Tags from Resources page, select the required Tag Keys and Tag Values that must be removed from Port Groups. 4. Click Ok to save the configuration.
New	Use to create a new Port Group. Refer to Create Port Groups for

Button	Description
	detailed information.
Actions	Use the Actions drop-down button to perform the following options: • Clone: Use to clone the selected Port Group. Refer to Clone Port Group for details. • Edit: Use to edit the selected Port Group. Refer to Edit Port Group for details. • Delete: Use to delete the selected Port Group.
Export	Use to export the Port Groups in CSV or XLSX format. The following options are available: • Export All: All Port Groups in the list view are exported. • Export Selected: Only the selected Port Groups are exported.

Port Group Statistics

To view the statistical details of the port group members:

1. Select **Ports > Port Groups > Statistics**. The Port Group Statistics page is displayed.
The following fields are displayed:
 - Port Group Members
 - Total Bytes
 - Total Packets
 - Total Sessions
 - Active Sessions
2. Use the **Actions** drop-down button to perform the following:
 - **Clear:** Select a Port Group Member and click **Clear** to clear the port group members statistics counters.
 - **Clear All:** Click **Clear All** to clear the port group member statistics counters of all the port group members.
3. Use the **Export** button to export the statistical details.

Port Pairs

A port-pair is a bidirectional connection in which traffic arriving on one port in the pair is transmitted out the other (and vice-versa) as a passthrough TAP. Keep in mind the following rules and notes for port-pairs:

- You can configure whether a port-pair uses link status propagation. Link port propagation does the following:
 - Enabled—when one port in the pair goes down, the other port goes down.
 - Disabled—when one port in the pair goes down, the other port is unaffected.
- Port-pairs can be established between ports using different speeds. For example, from a 100Mb port to a 1Gb port. However, the system will warn you when creating such port-pairs. Depending on traffic volume, port-pairs between ports using different speeds can cause packet loss when going from a faster port to a slower port. For example, going from 1Gb to 100Mb, from 10Gb to 1Gb, and so on.

Create Port Pair

To configure a port pair, do the following:

1. Select **Ports > Port Pairs**.
2. Click **New**.
3. On the Port Pair page, do the following:
 - a. (Optional) Type an alias in the **Alias** field to help identify this port pair.
 - b. (Optional) Type a description in the **Description** field.
 - c. Click in the **First Port** field and select a network port.
 - d. Click in the **Second Port** field and select another network port.
 - e. (Optional) Enable **Link Failure Propagation**.

Port pairs can operate with or without line failure propagation (LFP) as follows:

- With LFP enabled, link failure on one of the ports in the port pair automatically brings down the opposite side of the port pair.
- With LFP disabled, the opposite port is not brought down automatically.

Note: A port pair created on a copper TAP has LFP enabled by default.

4. Click **Apply**.

Use the following buttons to manage the Port Pairs.

Button	Description
Tags	Use to associate tags to Port Pairs and to remove the tags from Port Pairs when not required. To add tags: 1. Select the required Port Pairs to which the tags must be applied. 2. Click the Tags drop-down menu button and select Add.

Button	Description
	3. In the Add Tags to Resources page, select the required Tag Keys and Tag Values. 4. Click Ok to save the configuration. The Port Pair is associated with the selected tags. To remove tags: 1. Select the required Port Pairs that are already associated with tags. 2. Click the Tags drop-down menu button and select Delete. 3. In the Delete Tags from Resources page, select the required Tag Keys and Tag Values that must be removed from Port Pairs. 4. Click Ok to save the configuration.
New	Use to create a new Port Pair. Refer to Create Port Pair for detailed information.
Actions	Use the Actions drop-down button to perform the following options: • Edit: Use to edit the selected Port Pair. • Delete: Use to delete the selected Port Pair.
Export	Use to export the Port Pairs in CSV or XLSX format. The following options are available: • Export All: All Port Pairs in the list view are exported. • Export Selected: Only the selected Port Pairs are exported.

Tool Mirrors

In addition to maps, the GigaVUE-OS also includes a special Tool Mirror packet distribution feature. A Tool Mirror can be used to send all packets on one tool port to another tool port (or multiple tool ports) or GigaStream on the same box. Tool Mirrors can still be applied to network ports even if they are already in use with an existing connection or map. Use tool-mirror connections between tool ports/GigaStreams on the same node, cross-box tool-mirror connections are not supported.

Tool-mirror can be created from:

- Tool port to tool port or ports on the same node.
- Tool port to GigaStream or GigaStreams on the same node.
The destination for a tool-mirror must always be either a tool port or a GigaStream.
- Tool Mirrors can cross line cards/modules – they can start on one line card and end on another in the same node. However, they cannot cross nodes in a cluster.
- Tool Mirrors on GigaVUE-HC1 can be created on tool ports or GigaStream ports.
- Tool Mirrors are not allowed from Tool GigaStream to tool port.

- Tool Mirrors are not supported on tool ports with copper SFPs installed.

Create Tool Mirror

To create a Tool Mirror, do the following:

1. Select **Ports > Tool Mirrors**.

The Tool Mirrors page displays a list of the currently configured Tool Mirrors.

2. Click **New**.

The Tool Mirror configuration page appears.

3. Configure the Tool Mirror:

- a. Enter an alias in the **Alias** field.

- b. (Optional) Enter a description in the **Description** field.

- c. Click in the **Source Tool Ports** field and select the source tool ports for this tool mirror.

- d. Click in the **Mirror Destination Ports** field and select the destination tool ports for this tool mirror.

4. Click **Apply**.

Edit Tool Mirror Description

Description in the Tool Mirror configuration is optional. However, you can add description at any time, or edit the existing description. To add or edit description, do the following:

1. Select **Ports > Tool Mirrors**. The Tool Mirrors page displays a list of the currently configured Tool Mirrors.
2. Select a Tool Mirror in the list of Tool Ports, and then click **Actions>Edit**.
3. Enter or change a description in the **Description** field. (You cannot make any other changes to the Tool Mirror.)
4. Click **Apply**.

Clone Tool Mirror

In some cases, you may want to create a Tool Mirror that is similar to an existing one. To do this use the Clone feature.

1. Select **Ports > Tool Mirrors**.
2. Select the Tool Mirror that you want to copy, and then click **Actions>Clone**.
3. Enter a new alias in the **Alias** field.

4. (Optional) Add or update Description in the **Description** field.
5. Make change to the **Source Tool Ports** and **Mirrored Destination Ports** as needed.
6. Click **Apply**.

Use the following buttons to manage the Tool Mirrors.

Button	Description
New	Use to create a new Tool Mirror. Refer to Create Tool Mirror for detailed information.
Actions	Use the Actions drop-down button to perform the following options: • Clone: Use to clone the selected Tool Mirror. Refer to Clone Tool Mirror. • Edit: Use to edit the selected Tool Mirror. Refer to Edit Tool Mirror Description. • Delete: Use to delete the selected Tool Mirror.
Export	Use to export the Tool Mirrors in CSV or XLSX format. The following options are available: • Export All: All Tool Mirrors in the list view are exported. • Export Selected: Only the selected Tool Mirrors are exported.

Stack Links

Use stack-links to connect multiple GigaVUE nodes in a unified cluster. The stack-links carry traffic entering one system and bound for another via a map. Stack management traffic uses its own dedicated network connections through the Stacking ports on the Control Cards.

You can construct stack-links either out of single stack ports or a stack GigaStream. However, because of the incredible 10Gb port density offered by the GigaVUE HC Series, using only one 10Gb port for a stack connection could cause a serious bottleneck.

NOTE: Packet loss may be seen on stack links when traffic exceeds 95% of the line rate. This is because each packet has a 16 bytes higit header added to it, which reduces the throughput.

NOTE: In a stack port configuration, with if a Gen 2 and Gen 3 card from the same chassis, use different VLAN IDs. You cannot use the default VLAN IDs.

A stack GigaStream dramatically increases the bandwidth available for stack connections, letting you connect GigaVUE nodes in a cluster and still take advantage of the 10Gb port density. Alternatively, nodes with 40Gb or 100Gb ports can take advantage of their high bandwidth for stack-links. (For more details about clustering, refer to the *GigaVUE-OS CLI Reference Guide*.)

Stack links are supported at speeds of 10Gb, 40Gb, and 100Gb. Refer to the *Hardware Installation Guide* for each GigaVUE node for information on stack link support.

Stacking is not supported on GigaVUE-TA400 in 5.14 release.

When using stack GigaStream for stack links, you must create a stack GigaStream on each side of the stack link and each must consist of the same number of ports running at the same speed.

To create a stack link, do the following:

1. Select **Ports > Stack Links**.
2. Click **New**.
3. Enter an alias for the stack link in the **Alias** field.
4. Select the **Type** for this stack link.
 - **Stack Ports** specifies that the stack link is between two ports.
 - **Stack GigaStream** specifies that the stack link is between GigaStream.
5. In the First Member and Second Member fields, select the ports or GigaStream for the stack link, depending on the type selected in [Step 4](#)
6. Click **Save**.

IP Interfaces

You can configure IP interface in the control card. All the control operations such as the gateway resolution, tunnel health check, and NetFlow exporter SNMP requests are handled in the control card. Similarly, the ARP/NDP timer configuration is also moved to the control card.

You can configure the IP Interfaces for the following tunnels:

- GigaSMART tunnel - Refer to the [Configure IP Interfaces for GigaSMART Tunnels](#).
- Layer 2 Generic Routing Encapsulation (L2GRE) tunnel - Refer to the [About 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 the GigaSMART group.

To configure an IP interface:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**. Select the node for which you want to configure the IP interface.
2. From the left navigation pane, go to **Ports > IP Interfaces**.
3. In the IP Interfaces page, click **New**. The New IP Interface page appears.
4. In the **Alias** and **Description** fields, enter a name and description for the IP interface.
5. Click **Port Editor**. A quick port editor appears. You can enable the port for which you want to configure the IP interface. Click **Save**.
6. From the **Ports** drop-down list, select the tool or the network port that you had enabled.

NOTE:

- The port should not be part of the port group.
- Before upgrading to the latest version, you must remove either the port group or the IP interface if the ports are identical.

7. Select the type as either **IPv4** or **IPv6**.
8. Enter the **IP Address**, **IP Mask**, and **Gateway** for the IP interface.

You can specify the subnet mask using either of the following formats:

- netmask – For example, 255.255.255.248
- mask length – For example, /29

NOTE: For IPv6 interface configured on Generation 3 GigaSMART card (SMT-HC1-S), prefix length of 128 is not supported.

9. Enter the **MTU** for the IP interface (100 - 9600 bytes). The MTU is fixed at 9600 for all network/tool ports on the following platforms:
 - Certified Traffic Aggregation White Box
 - GigaVUE-HC1
 - GigaVUE-HC3

- GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA100, GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA400, and GigaVUE-TA400E.

RECOMMENDATION: Set the MTU to 9400 on all platforms.

10. If you are configuring IP Interface for GigaSMART tunnels:

- a. From the **GS Groups** drop-down list, select the GigaSMART groups that you have configured.

NOTE: You can associate multiple GigaSMART groups to the IP interface.

- b. From the **Exporters** drop-down list, select the NetFlow exports that you have created.

11. Click **Apply** to configure the IP interface associated with a network or tool port and add it to the list of currently configured IP interfaces.

To edit the existing IP Interfaces, select the IP Interface and click **Actions > Edit**.

To delete the IP interfaces, select the IP Interface and click **Actions > Delete**.

IP Interface Statistics

To view the statistics such as the packets and bytes decapsulated on the IP interface:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**. Select the node for which you want to view the IP interface statistics.
2. From the left navigation pane, go to **Ports > IP Interfaces > Statistics**. The IP Interface Statistics page appears.

To clear the statistics for a specific IP Interface, select the IP Interface and click **Actions > Clear**.

To clear all the statistics, select all the IP interfaces and click **Actions > Clear All**.

To export the IP Interface statistics, click **Export**.

Circuit Tunnels

Circuit tunnels are used to route traffic between two clusters. The traffic is tapped and sent through network ports on the TAP landing nodes in a cluster. Based on the flow map configuration, traffic is filtered at the TAP landing nodes and sent to the circuit ports. The circuit ports encapsulate the traffic with a Circuit ID and routes the encapsulated traffic

through a circuit tunnel. At the receiving end, the traffic is decapsulated and sent to the tool ports. The circuit tunnels are bidirectional. For more information about circuit tunnels, refer to [About Circuit-ID Tunnels](#).

Port Discovery

This section describes port discovery for the GigaVUE HC Series, providing information about discovery protocols and how to enable discovery through GigaVUE-FM. For details refer to the following:

- [Port Discovery with LLDP and CDP](#)
- [Enable Port Discovery](#)
- [Port Discovery Support](#)

Port Discovery with LLDP and CDP

The GigaVUE HC Series and GigaVUE TA Series devices are capable of snooping Link Layer Discovery Protocol (LLDP) packets and Cisco Discovery Protocol (CDP) packets. If the devices in your network use either of these protocols, a GigaVUE HC Series or a GigaVUE TA Series node can identify its immediate neighbors and their capabilities. Snooped LLDP and CDP information includes the remote port and chassis IDs, as well as other selected information, if it is included by the sender. This information can be used to determine the origin of traffic flows.

All GigaVUE HC Series and GigaVUE TA Series nodes support LLDP and CDP port discovery,

LLDP and CDP are physical topology discovery protocols (Layer 2). The protocols are unidirectional. Devices send their identity and capabilities in a packet. The GigaVUE HC Series or GigaVUE TA Series node receives the packet and extracts information from it, such as the chassis ID and port ID of a neighbor. The information from the neighbors varies depending on what is sent in the packet.

An LLDP packet supports the following capabilities in a type-length-value (TLV) structure. The first four capabilities are mandatory.

- Chassis ID
- Port ID
- Time-to-Live (TTL)
- End of TLVs
- Port description

- System name
- System description
- System capabilities available
- System capabilities enabled
- VLAN name
- Management address
- Port VLAN ID
- Management VLAN ID
- Link Aggregation port ID
- Link Aggregation status
- Maximum Transmission Unit (MTU)

A CDP packet supports the following capabilities in a TLV structure:

- Device ID
- Port ID
- Platform
- Software version
- Native VLAN ID
- Capabilities
- Network prefix address
- Network prefix mask
- Interface address
- Management address

The LLDP/CDP discovery packets are copied and parsed by the node, and the neighbor information is cached. Discovery packets are not terminated on the GigaVUE HC Series or the GigaVUE TA Series node, nor are they removed from the ingress data stream.

Notes:

- Port discovery can be enabled on network, tool, and circuit type ports.
- Use port discovery on ports fed by SPAN ports or aggregators with caution. LLDP/CDP information received from a SPAN port may be misleading, depending on how it is configured. When a large range of ports are SPANed, different and conflicting LLDP/CDP information may be received. LLDP/CDP is best used on TAPed network interfaces.

Enable Port Discovery

Port discovery is disabled by default. It can be enabled on network, tool, and circuit type ports.

NOTE: The ports do not have to be included in a map.

1. Select **Ports > All Ports**.
2. On the Ports page, click on the Port ID of the port on which you want to enable port discovery.

The Quick View window displays for the port ID.

3. Select **Edit** from the top right corner of the Quick View Window.
4. To enable ports discovery do the following under **Device Discovery**:

- a. Select **Enable**
- b. For Discovery Protocols, select one of the following: **All**, **LLDP**, or **CDP**.

[Enable Port Discovery](#) Figure shows ports discovery enabled using the LLDP protocol for the port 1/2/x1.

5. Click **OK**.

Limits of Discovery Information

The following are limits on the amount of discovery information that is retained:

- For each port, discovery information for a minimum of two neighbors and a maximum of 20 neighbors is retained.
- For a chassis, discovery information for a maximum of 2K neighbors is retained.

Neighbor information will be removed or replaced as follows:

- when the neighbor information expires due to the TTL
- when the number of neighbors for the chassis reaches the 2K maximum and a new neighbor is discovered. In this case, the following can occur:
 - if there are currently two or more discovered neighbors for a port, the newly discovered neighbor will replace the neighbor information for the least recently updated neighbor
 - if there are currently less than two discovered neighbors for a port, the newly discovered neighbor will be added (actually exceeding the 2K limit in order to guarantee a minimum of two neighbors per port)

NOTE: Aging (the discovery protocol time-to-live) determines how long neighbor information is valid.

Port Discovery Support

This section describes port discovery for a cluster and port discovery for SNMP.

Port Discovery for a Cluster

LLDP and CDP discovery can be enabled on any network, tool, or circuit ports in a cluster. The discovery information will be aggregated and available on the cluster leader.

Port Discovery Supported for SNMP

The information from LLDP discovery is supported in the standard MIB and can be retrieved with SNMP **Get**.

The name of the MIB file that needs to be loaded in order to poll the LLDP information with SNMP is as follows:

- LLDP-MIB

The information from CDP discovery is supported in Cisco private MIBs and can be retrieved with SNMP **Get**.

The names of the Cisco MIB files that need to be loaded in order to poll the CDP information with SNMP are as follows:

- CISCO-CDP-MIB
- CISCO-SMI
- CISCO-SMI-MIB
- CISCO-TC
- CISCO-TC-MIB
- CISCO-VTP-MIB

Ingress and Egress VLAN

This section describes ingress port VLAN tagging and egress port VLAN stripping. Refer to the following sections for details:

- [About Ingress Port VLAN Tagging](#)

- [Ingress Port VLAN Tagging](#)
 - [Adding VLAN Tags](#)
 - [Deleting VLAN Tags](#)
- [Using VLAN Tags in Maps](#)
- [Ingress Port VLAN Tag Limitations](#)
 - [Second Level Maps](#)
 - [Double-Tagged Packets](#)
 - [IP Interfaces](#)
- [Configure Egress Port VLAN Stripping](#)
 - [Enable Egress Port VLAN Stripping](#)
 - [Disable Egress Port VLAN Stripping](#)
 - [Display Egress Port VLAN Stripping](#)
- [Egress Port VLAN Stripping Limitations](#)
- [How to Use Both Ingress Tagging and Egress Stripping](#)

About Ingress Port VLAN Tagging

You can add VLAN tags to ingress packets on a per-port basis. You manually associate VLAN IDs with specific ports of type network, inline-network, or hybrid or hybrid GigaStream.

Use VLAN tags to identify, differentiate, or track incoming sources of traffic. When the traffic reaches the tools or the maps, you can filter on the VLAN tags for the corresponding ports you want to measure.

Ingress port VLAN tagging is supported for IPv4 and IPv6 packet types, including non-tagged packets, tagged packets, and Q-in-Q packets. Ingress port VLAN tagging is not supported on network ports that are connected via port-pairs.

Each port can only have one VLAN tag. The same VLAN tag can be assigned to multiple network ports or hybrid ports or to both ports in an inline network port pair.

VLAN tagging is supported in a cluster.

Refer to [Figure 11Using Ingress Port VLAN Tagging](#) for an example. In the example, traffic from San Jose is tagged with VLAN 1001 and traffic from San Francisco is tagged with VLAN 1002.




Figure 11 Using Ingress Port VLAN Tagging

Ingress Port VLAN Tagging

The port type must be a network or hybrid or inline network type of port. Each network port can only have one VLAN tag. Once a VLAN tag is configured, it can be modified by overriding the existing one with a new VLAN ID.

The VLAN ID is specified in the **VLANTag** field of the port configuration page for an network or inline network port. The value of the VLAN ID is specified as a number between 2 and 4000.

Adding VLAN Tags

To add/modify VLAN tags, follow these steps:

1. Go to **System > Ports > All Ports**.
2. Click on the **Port ID**. Ensure that this Port ID is set as a network port. The Quick View window for the Port ID displays.
3. Select **Edit** from the top right corner of the Quick View Window.
4. Add the **VLAN ID** to the parameter field **VLAN Tag** and click **Save**.

Alias:

▼ Parameters

Admin: ☒ Enable

Type:

Speed:

Duplex: ☒ Full ☐ Half

Auto Negotiation: ☐ Enable

VLAN Tag:

Force Link Up: ☐ Enable

Timestamp: ☐ Append Ingress ☐ Strip Egress ☒ Source ID Egress

Figure 12 Port-ID Configuration

Deleting VLAN Tags

Once a VLAN tag is configured, it can be deleted by removing the value from the **VLAN Tag** field and saving the port configuration.

Using VLAN Tags in Maps

Ingress port VLAN tags are supported in first level maps, including the following:

- map
- map-passall
- map-scollector
- GigaSMART operation (gsop-enabled) maps

For example, if the traffic from network port 2/1/q3, (which has VLAN tag 1001 configured), is forwarded to tool port 2/1/q4. The traffic at tool port 2/1/q4 will have the added VLAN tag 1001. (Even though the VLAN tag is configured on the network port, it is added when the traffic exits the tool port.)

NOTE: Traffic from a network port will not match a map rule that filters on a VLAN tag configured on the network port.

Ingress Port VLAN Tag Limitations

The following sections describe limitations of ingress port VLAN tagging:

- [Second Level Maps](#)
- [Double-Tagged Packets](#)
- [IP Interfaces](#)
- [Local Tool Port Ingress VLAN Tag](#)

Second Level Maps

VLAN tagging is not supported for second level maps, which are maps from a virtual port (vport).

For tagged network ports, if the ingress traffic is going to a second level map, the packets will not be tagged at the egress ports of the second level map. This is a limitation of GigaSMART operations using maps with vports.

Double-Tagged Packets

If incoming packets already have two VLAN tags, such as with Q-in-Q, the addition of a third VLAN tag can cause problems with the following:

- Layer 3/Layer 4 filtering
- GigaStream hashing (all packets may be sent to only one tool port)

IP Interfaces

For IP interfaces, a VLAN tag added at the network port of the encapsulation path (n1 in [Figure 13 IP Interfaces](#)) will become part of the payload going to the decapsulation path. But a VLAN tag added at the network port of the decapsulation path (n2 in [Figure 13 IP Interfaces](#)) will be available at the end tool port for filtering (t2 in [Figure 13 IP Interfaces](#)).

Refer to [Figure 13 IP Interfaces](#). VLAN tag (vlan1) added at the encap network port (n1) is encapsulated in the tunnel payload and cannot be used for filtering at the decap side. VLAN tag (vlan2) added at the decap network port (n2) can be used in a filter rule to send packets to tool port (t2).




Figure 13 IP Interfaces

Local Tool Port Ingress VLAN Tag

Ingress VLAN tag is not working on the local tool port if both MPLS and Ingress VLAN tag are enabled on the same network port of GigaVUE-HC1-Plus, GigaVUE-HCT, GigaVUE-TA25, and GigaVUE-TA25E.



Note: Circuit ports to make Circuit Mapping functionality to work is not supported.

Configure Egress Port VLAN Stripping

You can enable or disable outer VLAN stripping on specified egress ports. The port type must be tool or hybrid.

Use egress port VLAN stripping to strip an outer VLAN tag without using a GigaSMART stripping operation.

Enable Egress Port VLAN Stripping

To enable egress port VLAN stripping:

1. Select a tool or hybrid port on the Ports page.
2. Click **Edit**.
3. Under Parameters, for Egress Vlan Tag, select Strip. Refer to [Figure 14 Select Strip Egress Vlan Tag](#).

Admin ☒ Enable

Type

Speed

Duplex ☒ Full ☐ Half

Auto Negotiation ☐ Enable

Egress Vlan Tag ☐ None ☒ Strip

Force Link Up ☐ Enable

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 14 Select 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-HCI-Plus, GigaVUE-HCT, and GigaVUE-TA25E engines, Port level egress VLAN stripping is only supported if the network and tool ports are not configured on the same node in double tag mode.

How to Use Both Ingress Tagging and Egress Stripping

When ingress port VLAN tagging is enabled on a network port and egress port VLAN stripping is enabled on a tool port on the same GigaVUE node, refer to the [Table 5: VLAN Stripping Table](#):

Table 5: VLAN Stripping Table

Tool →	Stripping Enabled			Stripping Disabled		
Network	Untagged	Single Tag	Double Tag	Untagged	Single Tag	Double Tag
Ingress VLAN tag enabled	None	None	Customer VLAN tag	Ingress VLAN tag	Ingress VLAN tag + Customer VLAN tag	Ingress VLAN tag + Customer VLAN tag + Service VLAN tag
Ingress VLAN tag disabled	None	None	Customer VLAN tag	None	Customer VLAN tag	Service VLAN tag

NOTES:

- The inner VLAN tag is classified as the Customer VLAN tag (ethertype 0x8100)
- The outer VLAN tag is classified as the Service VLAN tag (ethertype 0x8100, 0x88A8, or 0x9100)

How to Use GigaStream

This section describes how to create and manage GigaStream. A GigaStream groups multiple ports into a logical bundle. Refer to the following sections for details:

- [About GigaStream](#)
- [Regular GigaStream](#)
- [Controlled GigaStream](#)
- [Advanced Hashing](#)
- [Weighted GigaStream](#)
- [GigaStream Rules and Maximums](#)

GigaStreams

A GigaStream is a bundle of multiple ports on a GigaVUE-OS node. You can create the following types of GigaStreams:

- **Tool GigaStream**—It is a bundle of multiple tool ports used as a single logical group. This type of GigaStream as a single addressable destination, allowing you to overcome tool port oversubscription issues.
- **Hybrid GigaStream**—It is a bundle of multiple hybrid ports that are combined into a single logical group in all H Series nodes.
- **Stack GigaStream**—It is a bundle of multiple stack ports used as a single logical group. Stack-links can use GigaStream to distribute data between multiple GigaVUE® HC Series nodes operating in a cluster. With the number of 10Gb/40Gb/100Gb ports possible in a GigaVUE HC Series chassis, using only one 10Gb port for a stack-link could cause a serious bottleneck. A GigaStream dramatically increases the bandwidth available for stack-link connections, letting you connect H Series nodes in a cluster and still take advantage of the 10Gb port density.
- **Circuit GigaStream**—It is a bundle of multiple circuit ports that are combined into a single logical group. The circuit ports send or receive traffic that is tagged with the circuit ID.
- **Controlled GigaStream**—Controlled GigaStream provides more control of the traffic stream by specifying the size of a hash table and allowing the assignment of hash IDs to the ports in a GigaStream. This makes it possible to keep the hashing algorithm from reapplying the algorithm to the ports if one of the ports in the GigaStream goes down.

Accessing the GigaStream page

To access the GigaStream page:

1. From the left navigation pane, go to **Inventory > Physical > Nodes..**
2. In the Physical Nodes page, click the required standalone node or the cluster ID for which you want to configure a GigaStream.
3. On the left-navigation pane, go to **Ports > Port Groups > GigaStream™.**

The GigaStream page displays the list of configured GigaStreams.

The GigaStream page has the following buttons, which can be used to perform specific options.

Button	Description
Tags	Use to associate tags to GigaStream and to remove the tags from

Button	Description
	GigaStream when not required. To add tags: 1. Select the required GigaStream to which the tags must be applied. 2. Click the Tags drop-down menu button and select Add. 3. Click Ok. To remove tags: 1. Select the required GigaStream that are already associated with tags. 2. Click the Tags drop-down menu button and select Delete. 3. Click Ok.
New	Use to create a new GigaStream. Refer to Configure Regular GigaStream for detailed information.
Actions	Use the Actions drop-down button to perform the following options: • Clone: Use to clone the selected GigaStream. • Edit: Use to edit the selected GigaStream • Delete: Use to delete the selected GigaStream
Advanced Hash Settings	Use to configure the advanced hash settings. Refer to Advanced Hashing section for detailed information.
Export	Use to export the GigaStream in CSV or XLSX format. The following options are available: • Export All: All GigaStreams in the list view are exported. • Export Selected: Only the selected GigaStreams are exported.

About GigaStream

There are two types of GigaStream:

- Regular GigaStream
- Controlled GigaStream.

Both types of GigaStream bundle multiple ports to provide logical bandwidth. Packets arriving through network ports are processed with various map rules and then directed to ports. All traffic streams destined to a GigaStream are hashed among the bundled ports.

Regular GigaStream

Regular GigaStream groups multiple ports running at the same speed into a single logical bundle called a GigaStream. Regular GigaStream can be used as either a packet egress destination (tool GigaStream) or as a stack-link between two GigaVUE-OS nodes operating in a cluster (stack GigaStream).

NOTE: The existing tool and stack GigaStream are now referred to as regular GigaStream. The term GigaStream is used when something applies to both types

For details on regular GigaStream, refer to [Regular GigaStream](#).

Controlled GigaStream

Controlled GigaStream provides GigaStream controlled traffic distribution. Controlled GigaStream samples traffic based on hash settings and helps to ensure that traffic sent to each tool is within the capacity of the tool.

For details on controlled GigaStream, refer to [Controlled GigaStream](#).

Controlled GigaStream provides greater flexibility in allocating the bandwidth assigned to tools within the GigaStream. Regular GigaStream assumes that each tool in the GigaStream is sent an equal fraction of the traffic. Controlled GigaStream allows different tools to be sent different fractions of the traffic.



Note: Regular GigaStream and controlled GigaStream differ in the following ways:

- traffic distribution based on hashing
- traffic fail over when a port goes down
- editing a configuration, such as adding ports on the fly

Regular GigaStream

Regular GigaStream can be used as either a packet egress destination (tool GigaStream) or as a stack-link between two GigaVUE-OS nodes operating in a cluster (stack GigaStream).

All ports in a GigaStream must be running the same speed, such as 10Gb or 40Gb, and must use the same port type, either tool or stack. All ports in a GigaStream can be on different modules of the same GigaVUE-HC3 node.

With regular GigaStream, the hashing is computed based on traffic. Incoming packets arriving through network ports are processed with various map rules and then directed to ports. The result of the hash distributes traffic equally across the GigaStream members. Refer to [Figure 15 Regular 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 16 Regular 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 17 Regular Stack GigaStream Illustrated](#).




Figure 17 Regular Stack GigaStream Illustrated

Configure Regular GigaStream

All ports combined in a GigaStream must be running at the same speed, using the same port types. Port speeds less than 1000Mb are not supported.

Also, refer to [Advanced Hashing](#) for optional advanced hashing settings and [Weighted GigaStream](#) for optional weighting settings.

Before you configure a regular GigaStream, ensure that you have configured the required ports—tool, hybrid, stack, or Circuit. For information about configuring a port, refer to [Configure Ports](#).

To configure a regular GigaStream:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**.
2. In the Physical Nodes page, click the required cluster ID for which you want to configure a regular GigaStream.
3. On the left-navigation pane, go to **Ports > Port Groups > GigaStream™**.
4. Click **New**. The **GigaStream™** page appears.
5. In the **Alias** and **Description** fields, enter the name and description of the regular GigaStream that you want to configure.
6. Select the type of GigaStream that you want to configure. For example, if you want to configure a regular tool GigaStream, select **Tool GigaStream**.
7. From the **Ports** drop-down list, select the ports that you have configured. For example, select the required hybrid ports to configure a regular hybrid GigaStream.
8. From the **Weighting** drop-down list, select one of the following options:
 - **Equal**—Traffic is distributed equally to all the ports in the regular GigaStream.



You can enter a value for the Variance Threshold in % only if:

- Equal Weighting is selected
- GigaStream is a regular GigaStream
- Device Version is 6.0 and above

If the threshold is above or below the Variance threshold, and if the traffic distribution is not uniform across the ports in the GigaStream, traps will be sent from the devices to GigaVUE-FM

- **Relative**—Traffic is distributed to the ports in the regular GigaStream based on the relative weight or ratio assigned to the respective ports. The valid range is 1–256.
- **Percentage**—Traffic is distributed to the ports in the regular GigaStream based on the percentage assigned to the respective ports. The valid range is 1–100.

If you select **Relative** or **Percentage** as the weighting option, enter the hash weights for the ports that appear in the table below the **Weighting** drop-down list.

- 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.

- 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.
- 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:

- First level GigaSMART map
- Second level GigaSMART map (includes GTP overlap sampling maps)
- 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:

- Select **Ports > Port Groups > GigaStreams**.

2. Select a GigaStream and click **Actions > Edit** to open the GigaStream configuration page.
3. Click in the **Ports field** and **Change, Add** or **Delete** tool ports associated with the GigaStream as needed.
4. Click **Apply**.

Edit Regular Stack GigaStream

You can edit regular stack GigaStreams that are configured on either sides of a stack link. When a stack GigaStream is attached to a map, you can directly add or delete stack ports from the stack GigaStream.

To edit a regular stack GigaStream:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**. In the Physical Nodes page, click the required cluster ID.
2. On the left-navigation pane, go to **System>Ports > Stack Links**. The Stack Links page appears.
3. Select the alias of the stack link that is grouped in the stack GigaStream that you want to modify, and then click **Edit**. The Stack Link page appears.
4. From the **Ports** drop-down list, add or delete the required stack ports, and then click **OK**.

Traffic Distribution Across Regular GigaStream

All the traffic streams destined to the GigaStream are distributed among the bundled ports based on hashing, as defined in the advanced hash settings or Weighted GigaStream, as defined in the **Weighting** field.

The hash is performed across multiple fields, such as IP address, port number, protocol, MAC address, and other criteria. The best practice is to include both the source and destination fields, such as source IP address and destination IP address, within the advanced hash settings. Because the hash calculation is symmetrical with respect to source and destination addresses, all packets belonging to the same session will be sent to the same tool.

For more information on hashing, refer to [Advanced Hashing](#).

The GigaVUE-OS nodes distribute traffic between the ports in a regular GigaStream using one of the following criteria:

- The criteria configured using the Advanced Hash Setting page for the selected line card or chassis. (Click **Advanced Hash Setting** on the GigaStream page to open the Advance Hash Setting.) Because traffic is hashed across member ports rather than

divided evenly, the bandwidth available for a regular GigaStream is not a straight multiple of the number of ports in the bundle – some flows will use more bandwidth than others.

NOTE: The GigaVUE HC Series node tries to distribute incoming traffic evenly across all tool ports in the GigaStream. However, live network traffic is often unpredictable, including bursty periods for certain sessions. Because of this, the distribution patterns described are not ironclad – variations in traffic will result in variations in distribution.

The distribution described in this section applies to GigaVUE-HC3 nodes, GigaVUE-HC1 nodes, and GigaVUE TA Series nodes for regular tool GigaStream and regular stack GigaStream.

- Weighting mode and hash weights assigned to the different ports in the regular GigaStream. For more information about Weighted GigaStream, refer to [Weighted GigaStream](#).

Regular GigaStream Failover Protection—Resiliency

Regular GigaStream has built-in failover protection or resiliency. When there is a failover of a port that is part of a regular GigaStream, the traffic is redistributed to the other tool ports without disturbing the session continuity.

When a tool port goes down, the sessions allocated to the failed port are redistributed among the remaining available ports. With the resiliency functionality, the redistribution of traffic occurs without disturbing the session continuity of the active ports.

Recovery of a regular GigaStream is automatic. When a down link returns, the traffic will be reassigned to their original ports automatically.

The resiliency functionality is supported on all GigaVUE HC Series, and TA Series nodes and on Hybrid GigaStream, Tool GigaStream, Circuit GigaStream and Cluster GigaStream.
(missing or bad snippet)

Regular Circuit GigaStream

A regular circuit GigaStream uses circuit ports to pass traffic between two clusters. For more information refer to the following topics:

- [About Circuit-ID Tunnels](#)
- [Fabric Maps Prerequisites](#)

Controlled GigaStream

Controlled GigaStream provides controlled traffic distribution, which gives more granular control over hashing to the tool ports.

- All GigaVUE® HC Series and TA Series nodes support controlled GigaStream.
- GigaVUE nodes with controlled GigaStream are supported in a cluster environment.
- Controlled GigaStream can only be used as a packet egress destination (tool GigaStream). All port speeds are supported.

Controlled GigaStream delivers traffic flows based on hash settings and hash bucket assignments. Assigning varying numbers of hash buckets helps to ensure that traffic sent to each tool is within the capacity of the tool. Leaving some hash buckets unassigned creates a simple traffic sampling configuration, as traffic flows for unassigned hash buckets are discarded.

With controlled GigaStream, hashing is computed based on traffic. There is a configurable number of hash buckets, from 1 to 256. The hash size of a controlled GigaStream specifies the number of logical tools the traffic will be distributed across. Refer to [Figure 18 Controlled 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 18 Controlled 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 19 Controlled 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.

Type

☒ Tool GigaStream

☐ Hybrid GigaStream

☐ Stack GigaStream

☐ Circuit GigaStream

Controlled GigaStream

☒

Hash Size *

5

Hash ID	Port ID
1	T 6/1/x4 ▼
2	T 6/1/x4 ▼
3	Select... ▼
4	Select... ▼
5	Select... ▼

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.

Add GigaStream™

Form elements marked with * are mandatory. x

Alias *

Description

Type

- ☒ Tool GigaStream
- ☐ Hybrid GigaStream
- ☐ Stack GigaStream
- ☐ Circuit GigaStream

Controlled GigaStream ☒

Hash Size *

Hash ID	Port ID
1	T 6/1/x12
2	T 6/1/x13
3	T 6/1/x14
4	T 6/1/x15
5	Select...

Figure 21 Ports Assigned to Hash Bucket IDs

- 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, $500\text{Gb}/5\text{Gb} = 100$ tools. But if you only have 10 tools, you create a controlled GigaStream of 100 logical tools or 100 logical hash buckets, and then map only 10 of them.

Taking the number of buckets and dividing it by 100 tools ($256/100 = 2.56$) results in some buckets of 3, some buckets of 2.

The recommendation is to round to an even divisor of 256 (2, 4, 8, 16, 32, 64, 128, or 256). In this example, instead of using 100, use 128, so each bucket will be 2 ($256/128 = 2$).

Hash buckets IDs are mapped to ports as follows:

Buckets	Ports
1	x1
2	x2
3	x3
4	x4
5	x5
6	x6
7	x7
8	x8
9	x9
10	x10
11	unmapped
...	unmapped
128	unmapped

The hashing to the 10 connected tools captures the traffic associated with those sessions.

There are no tools associated with the remaining buckets, so that traffic is black holed. Unlike regular GigaStream, you do not have to allocate ports to the remaining buckets.

Multiple buckets can be mapped to one physical port as follows:

Buckets	Ports
1	x1

Buckets	Ports
2	x1
3	x2
4	x3
...	...

In this example, port x1 receives 5Gb of traffic, while ports x2 and x3 receive 2.5Gb each.

Note that the mapping does not have to be consecutive as follows:

Buckets	Ports
1	x1
2	x2
3	x3
4	x1
...	...

In this example, port x1 also receives 5Gb of traffic, while ports x2 and x3 receive 2.5Gb each.

Through the mapping of buckets to ports, you can control the overall distribution of traffic to a given port.

Failover and Controlled GigaStream

Unlike the regular GigaStream, failover will not be triggered during a port down event. With controlled GigaStream, there is no rehashing or redistribution of traffic. In other words, the sessions flowing to other tool ports will not be disturbed and do not risk becoming oversubscribed.

Controlled GigaStream maintains hashing. When a port goes down, traffic is not re-hashed, but is black holed. Unlike with regular GigaStream, you do not need to enable **Force Link Up** on ports in order to counteract the default failover protection.

Advanced Hashing

Both regular GigaStream and controlled GigaStream use advanced hashing, which lets you select the criteria on which the hash is based, such as source and destination IP address, source and destination MAC address, source and destination port, and protocol.

GigaVUE-OS nodes distribute traffic between the ports in a GigaStream based on the hashing criteria configured using the **Advanced Hash Settings** page for the selected line card or chassis. GigaStream hashing is applicable for the following port types:

- Tool port
- Hybrid port
- Circuit port

To open the Advanced Hash Settings page, select **Ports > Port Groups > GigaStream** and click **Advanced Hash Settings**. (For more details, refer to [Advanced Hash Settings](#).) On the GigaVUE nodes, GigaStream hashing is per chassis, not per line card.

The **Advanced Hash Settings** let you select the different packet criteria used to send matching flows to the same destination port within a GigaStream.

By default, the GigaVUE HC Series node hashes traffic based on source and destination IP addresses, IP protocol, and source and destination ports.

How to Change Advanced Hash Criteria

You can select the criteria for the advanced hash algorithm by using **Advanced Hash Settings**. The advanced hash method you specify is used for all GigaStream in place on the specified line card or chassis.

Advanced Hash Settings

The Advanced Hash Settings page is where criteria for the advanced hash algorithm is set. To open the page, select **Ports > Port Groups > GigaStream** and click **Advanced Hash Settings**.

The following table describes the fields in the Advanced Hashing Settings page.

Field	Description
Box	Identifies chassis to which the advanced algorithm will be applied.

Field	Description
Slot	Identifies the line card to which the advanced hash algorithm will apply. Each line card in certain GigaVUE HC Series nodes has its own individual advanced hash algorithm. On GigaVUE-HC1, and GigaVUE-HC3, GigaStream hashing is per chassis, not per line card. For example, the slot field will only show cc1 when configuring an GigaVUE-HC3.
Type	Type can be one of the following: • Default: Sets the advanced hash algorithm to its default settings. By default, the advanced hash algorithm includes source/destination IPv4/IPv6 addresses and ports. • Custom: Clears the settings from the advanced hash and allows you to select your own criteria. • All: Selects all criteria. • None: Clears all fields from the advanced hash.
IPv4	This area of the page lets you select the following criteria. Use the toggle button to select your options: • IPv4 Source Address—Adds IPv4 source IP • IPv4 Destination Address—Adds IPv4 destination IP • IPv4 Protocol—Adds IPv4 protocol • IPv4 Source Port—Adds IPv4 source port • IPv4 Destination Port—Adds IPv4 destination port
IPv6	This area of the page lets you select the following criteria. Use the toggle button to select your options: • IPv6 Source Address—Adds IPv6 source IP • IPv6 Destination Address—Adds IPv6 destination IP • IPv6 Next Header—Adds IPv6 next header field. • IPv6 Source Port—Adds IPv6 source port • IPv6 Destination Port—Adds IPv6 destination port
Layer2	This area of the page lets you select the following criteria: • Source MAC Address—Adds L2 source MAC • Destination MAC Address—Adds L2 destination MAC • EtherType—Adds L2 ethertype field.
MPLS	This area of the page lets you select the following criteria: • MPLS Hash—Adds MPLS labels (up to three)
GTP TEID	This area of the page lets you select the following criteria: • GTP TEID—Adds GTP tunnel endpoint identifier
Ingress Port	This area of the page lets you select the following criteria:

Field	Description
	Ingress Ports—Adds ingress port. For advanced hash fields, the "ingressport" refers to the ingress port specified in the map's from port-list. In a multi-path leaf and spine cluster, when a spine node uses the "ingressport" field for advanced hashing, it references from the port-list defined in the map(s) responsible for forwarding traffic across the spine.

Advanced Hash Examples

The following are some different advanced hash examples. Note that the advanced hash method usually combines multiple criteria.

The example in [Figure 22Advanced Hash with IPv4 Source and Destination Addresses](#) sets a **Custom** advanced hash method for slot cc1 in box ID 6 that distributes traffic based on matching IPv4 source and destination addresses.

The screenshot displays the 'Advanced Hash Settings' configuration page for GigaStream™. At the top, there are tabs for 'Port Groups', 'GigaStream™', and 'Statistics'. The 'Advanced Hash Settings' section includes the following fields:

- Box:** A dropdown menu set to '6'.
- Slot:** A dropdown menu set to 'cc1'.
- Type:** Radio buttons for 'Default', 'Custom' (selected), 'All', and 'None'.
- IPv4:** A section header for the following criteria.
- IPv4 Source Address:** A toggle switch set to 'On'.
- IPv4 Destination Address:** A toggle switch set to 'On'.
- IPv4 Protocol:** A toggle switch set to 'Off'.
- IPv4 Source Port:** A toggle switch set to 'Off'.

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.

5 >

Port Groups

Port Groups

GigaStream™

Statistics

Advanced Hash Settings

Box

6

Slot

cc1

Type

☒ Default

☐ Custom

☐ All

☐ None

IPv4

IPv4 Source Address

☒ On

IPv4 Destination Address

☒ On

IPv4 Protocol

☐ Off

IPv4 Source Port

☒ On

IPv4 Destination Port

☒ On

IPv6

IPv6 Source Address

☒ On

IPv6 Destination Address

☒ On

IPv6 Next Header

☐ Off

Attempted Sync Time: Aug 14, 2023 04:58:37

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.

The screenshot displays the 'Advanced Hash Settings' window. At the top, there are tabs for 'Port Groups', 'GigaStream™', and 'Statistics'. Below the title bar, the settings are organized into sections:

- Box:** A dropdown menu set to '3'.
- Slot:** A dropdown menu set to 'Select...'.
- Type:** Radio buttons for 'Default', 'Custom' (selected), 'All', and 'None'.
- IPv4:** A section containing seven toggle switches, all currently set to 'Off':
 - IPv4 Source Address
 - IPv4 Destination Address
 - IPv4 Protocol
 - IPv4 Source Port
 - IPv4 Destination Port
- IPv6:** A section containing five toggle switches, all currently set to 'Off':
 - IPv6 Source Address
 - IPv6 Destination Address
 - IPv6 Next Header
 - IPv6 Source Port
 - IPv6 Destination Port
- Layer 2:** A section containing three toggle switches:
 - Source MAC Address: Set to 'On'.
 - Destination MAC Address: Set to 'On'.
 - Ether Type: Set to 'Off'.
- GTP TEID:** A section at the bottom, currently empty.

Figure 24 Advanced Hash with Source and Destination MAC Address

Hashing Behavior

Table 6: Hashing Behavior Based on Hash Criteria Field Combinations shows the possible hash criteria field combinations and the corresponding hashing behavior based on packet type for advanced hashing for non-MPLS packets. (Refer to Table 7: Hashing Behavior Based on Hash Criteria Field Combinations for MPLS packets.)

Table 6: Hashing Behavior Based on Hash Criteria Field Combinations

Hash Criteria Fields	Packet Type	Hashing Behavior
Source MAC Address, Destination MAC Address	MAC	Hash on Source MAC Address, Destination MAC Address
Source MAC Address, Destination MAC Address	MAC + IPv4	Hash on Source MAC Address, Destination MAC Address

Hash Criteria Fields	Packet Type	Hashing Behavior
Source MAC Address, Destination MAC Address	MAC + IPv6	Hash on Source MAC Address, Destination MAC Address
IPv4 Source Address, IPv4 Destination Address	MAC	No hash
IPv4 Source Address, IPv4 Destination Address	MAC + IPv4	Hash on IPv4 Source Address, IPv4 Destination Address
IPv4 Source Address, IPv4 Destination Address	MAC + IPv6	No hash
IPv6 Source Address, IPv6 Destination Address	MAC	No hash
IPv6 Source Address, IPv6 Destination Address	MAC + IPv4	No hash
IPv6 Source Address, IPv6 Destination Address	MAC + IPv6	Hash on IPv6 Source Address, IPv6 Destination Address
Source MAC Address, Destination MAC Address, IPv4 Source Address, IPv4 Destination Address	MAC	Hash on Source MAC Address, Destination MAC Address
Source MAC Address, Destination MAC Address, IPv4 Source Address, IPv4 Destination Address	MAC + IPv4	Hash on IPv4 Source Address, IPv4 Destination Address
Source MAC Address, Destination MAC Address, IPv4 Source Address, IPv4 Destination Address	MAC + IPv6	Hash on Source MAC Address, Destination MAC Address
Source MAC Address, Destination MAC Address, IPv6 Source Address, IPv6 Destination Address	MAC	Hash on Source MAC Address, Destination MAC Address
Source MAC Address, Destination MAC Address, IPv6 Source Address, IPv6 Destination Address	MAC + IPv4	Hash on Source MAC Address, Destination MAC Address
Source MAC Address, Destination MAC Address, IPv6 Source Address, IPv6 Destination Address	MAC + IPv6	Hash on IPv6 Source Address, IPv6 Destination Address

Hash Criteria Fields	Packet Type	Hashing Behavior
Source MAC Address, Destination MAC Address, IPv4 Source Address, IPv4 Destination Address, IPv6 Source Address, IPv6 Destination Address	MAC	Hash on Source MAC Address, Destination MAC Address
Source MAC Address, Destination MAC Address, IPv4 Source Address, IPv4 Destination Address, IPv6 Source Address, IPv6 Destination Address	MAC + IPv4	Hash on IPv4 Source Address, IPv4 Destination Address
Source MAC Address, Destination MAC Address, IPv4 Source Address, IPv4 Destination Address, IPv6 Source Address, IPv6 Destination Address	MAC + IPv6	Hash on IPv6 Source Address, IPv6 Destination Address

NOTE: No hash means that the packets will be sent to the first port in the GigaStream.

Notes and Considerations for Advanced Hashing

Refer to the following notes and considerations for advanced hashing:

- With symmetric hashing, packets with their source and destination IP addresses and Layer 4 (L4) ports interchanged will go to the same GigaStream port. It is recommended to enable source and destination IPv4 or IPv6 pairs and L4 source and destination ports.
- Symmetric hashing is enabled for all GigaStream on all GigaVUE® HC Series and TA Series nodes.
 - For non-MPLS IPv4 and IPv6 packets, the hashing is fixed to the following 3-tuple: ipsrc, ipdst, and protocol or ip6src, ip6dst, and protocol. All other traffic follows the advanced hash settings.
- Gigamon devices do not support hashing of the IP header fields when there is a PPPOE header. Only Layer 2 (L2) fields can be used for such packets.

Advanced Hashing with MPLS

Starting in software version 5.1, GigaStream MPLS hashing adds the ability to hash on MPLS labels as well as the following IP address fields inside an MPLS tunnel: **ipsrc**, **ipdst**, **ip6src**, and **ip6dst**.

Advanced hashing with MPLS is supported on all GigaVUE® HC Series and TA Series nodes.

Use the **Advanced Hash Settings** to specify MPLS, which can detect up to three MPLS labels. Packets with one to three MPLS labels can be hashed, along with IP address fields, if present. If a packet has more than three MPLS labels, IP address fields after the third MPLS label cannot be hashed. Refer to [Hashing Behavior Based on Hash Criteria Field Combinations](#) for details of the hashing behavior.

MPLS labels will be used as part of the GigaStream hash criteria if the MPLS field is configured and the packet has EtherType 0x8847.

MPLS hashing applies to the following:

- regular GigaStream
- controlled GigaStream
- stack GigaStream
- inline tool groups

[Table 7: Hashing Behavior Based on Hash Criteria Field Combinations](#) shows the possible hash criteria field combinations and the corresponding hashing behavior based on packet type for advanced hashing with MPLS packets. (Refer to [Table 6: Hashing Behavior Based on Hash Criteria Field Combinations](#) for non-MPLS packets.)

Table 7: Hashing Behavior Based on Hash Criteria Field Combinations

Hash Criteria Fields	Packet Type	Hashing Behavior
IPv4 Source Address, IPv4 Destination Address	MAC + Label1 + Label2 + Label3 + IP + L4 + Payload	Hash on IP
Source MAC Address, Destination MAC Address	outer MAC + Label1 + Label2 + Label3 + inner MAC + IP + L4 + Payload	Hash on outer MAC
Source MAC Address, Destination MAC Address	MAC + Label1 + Label2 + Label3 + IP + L4 + Payload	Hash on MAC
MPLS Hash	MAC + Label1 + Label2 + Label3 + IP + L4 + Payload	Hash on MPLS Label1, Label2, Label3
MPLS Hash	MAC + Label1 + Label2 + IP + L4 + Payload	Hash on MPLS Label1, Label 2
MPLS Hash	MAC + Label1 + IP + L4 + Payload	Hash on MPLS Label1
MPLS Hash, EtherType, IPv4 Protocol	MAC + Label1 + Label2 + Label3 + Label4 + IP + L4 + Payload	Hash on MPLS Label1, Label2, Label3
MPLS Hash, IPv4 Source Address, IPv4 Destination Address	MAC + MPLS Label1 + IP + L4 + Payload	Hash on MPLS Label1, IPv4 Source Address, IPv4 Destination Address
MPLS Hash, IPv4 Source Address, IPv4 Destination Address	MAC + Label1 + Label2 + IP + L4 + Payload	Hash on MPLS Label1, Label2, IPv4 Source Address, IPv4 Destination Address

Hash Criteria Fields	Packet Type	Hashing Behavior
MPLS Hash,IPv4 Source Address,IPv4 Destination Address	MAC + Label1 + Label2 + Label3 + IP + L4 + Payload	Hash on MPLS Label1, Label2, Label3, IPv4 Source Address, IPv4 Destination Address
MPLS Hash,IPv4 Source Address,IPv4 Destination Address	MAC + Label1 + Label2 + Label3 + Label4 + IP + L4 + Payload	Hash on MPLS Label1, Label2, Label3
MPLS Hash,IPv4 Source Address,IPv4 Destination Address	MAC + MPLS Label1 + MAC + IP + L4 + Payload	Hash on MPLS Label1
MPLS Hash,IPv4 Source Address,IPv4 Destination Address	MAC + Label1 + Label2 + MAC + IP + L4 + Payload	Hash on MPLS Label1, Label2
MPLS Hash,IPv4 Source Address,IPv4 Destination Address	MAC + Label1 + Label2 + Label3 + MAC + IP + L4 + Payload	Hash on MPLS Label1, Label2, Label3
MPLS Hash,Source MAC Address,Destination MAC Address	MAC + Label1 + Label2 + Label3 + IP + L4 + Payload	Hash on MPLS Label1, Label2, Label3
MPLS Hash,IPv4 Source Port,IPv4 Destination Port	MAC + Label1 + Label2 + Label3 + Label4 + IP + L4 + Payload	Hash on MPLS Label1, Label2, Label3
IPv4 Source Port,IPv4 Destination Port	MAC + Label1 + Label2 + Label3 + IP + L4 + Payload	No Hash

NOTES:

- No hash means that the packets will be sent to the first port in the GigaStream.
- If an MPLS packet has a router alert label as one of its labels, the router alert label is skipped and the other available labels are used for hashing. For example, if a packet has four labels and the second label is the router alert, the first, third, and fourth labels are used for hashing.
- Starting with 6.1 release, all platforms (excluding TA25, TA25E, HC1P) support symmetric hashing of MPLS and non-MPLS mixed IP traffic (both IPv4, IPv6) using Source IP address, Destination IP address (2-tuple), Source Port, Destination Port (4-tuple), Protocol/Next-header (5-tuple) hash field combinations.

TA400 supports this with or without MPLS Header Stripping, other platforms require MPLS Header Stripping. TA400 supports this with MPLS layer2 as well as MPLS layer3 IP traffic, other platforms support this with MPLS layer 3 traffic only.

Advanced Hashing with GTP TEID

Starting in software version 5.2, GigaStream GTP TEID hashing adds the ability to hash on GTP tunnel endpoint identifiers (TEIDs). GPRS Tunneling Protocol (GTP) is an IP/UDP-based protocol for mobile data.

The TEID field in a GTP header is a unique identifier for mobile subscribers and is used to multiplex different connections on the same GTP tunnel. Use GTP TEID advanced hashing to load balance GTP packets across all GigaStream ports.

Advanced hashing with GTP TEID is supported on GigaVUE® HC Series and TA Series nodes, with the following distinctions:

- The following nodes are supported: GigaVUE-HC1, GigaVUE-HC1-Plus, GigaVUE-HCT, GigaVUE-HC3 (with control card version 1 or 2), GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA100, GigaVUE-TA200, GigaVUE-TA200E, and GigaVUE-TA400.

Use the **Advanced Hash Settings** to specify the **GTP TEID** field. It must be specified with one of the port source and destination pairs: either **IPv4 Source Port** and **IPv4 Destination Port** for IPv4 or **IPv6 Source Port** and **IPv6 Destination Port** for IPv6.

When the **GTP TEID** field is configured, GTP packets will use it for hashing along with configured IP fields, instead of the Layer 4 (L4) source and destination for GTP packets. The hashing for all non-GTP packets will be based on L4 source and destination port.

GTP TEID hashing is supported only for GTP-User packets, version 1 (V1), with L4 source and destination port 2152. The hashing functionality works only for non-fragmented packets.

Refer to [Hashing Behavior Based on Hash Criteria Fields: GTP TEID](#) for details of the hashing behavior.

GTP TEID hashing applies to the following:

- regular GigaStream
- controlled GigaStream
- stack GigaStream (except GigaVUE-HC3)
- inline tool groups

NOTE: GTP TEID hashing is not supported for stack GigaStream on GigaVUE-HC3 in this software version due to the 3-tuple limitation listed in [Notes and Considerations for Advanced Hashing](#).

[Table 8: Hashing Behavior Based on Hash Criteria Fields: GTP TEID](#) shows the possible hash criteria field combinations and the corresponding hashing behavior based on packet type for advanced hashing with GTP packets.

Table 8: Hashing Behavior Based on Hash Criteria Fields: GTP TEID.

Hash Criteria Fields	Packet Type	Hashing Behavior
IPv4 Source Address,IPv4 Destination Address,IPv4 Source Port,IPv4 Destination Port	MAC + IP + L4 + Payload	Hash on IPv4 Source Address,IPv4 Destination Address,IPv4 Source Port,IPv4 Destination Port
IPv4 Source Address,IPv4 Destination Address,IPv4 Source Port,IPv4 Destination Port,GTP TEID	MAC + IP + L4 + GTP + Payload	Hash on IPv4 Source Address,IPv4 Destination Address, GTP TEID
IPv4 Source Address,IPv4 Destination Address,IPv4 Source Port,IPv4 Destination Port,GTP TEID	MAC + IP + L4 + Payload	Hash on IPv4 Source Address,IPv4 Destination Address,IPv4 Source Port,IPv4 Destination Port
IPv6 Source Address,IPv6 Destination Address,IPv6 Source Port,IPv6 Destination Port	MAC + IP + L4 + Payload	Hash on IPv6 Source Address,IPv6 Destination Address,IPv6 Source Port,IPv6 Destination Port
IPv6 Source Address,IPv6 Destination Address,IPv6 Source Port,IPv6 Destination Port,GTP TEID	MAC + IP + L4 + GTP + Payload	Hash on IPv6 Source Address,IPv6 Destination Address, GTP TEID
IPv6 Source Address,IPv6 Destination Address,IPv6 Source Port,IPv6 Destination Port,GTP TEID	MAC + IP + L4 + Payload	Hash on IPv6 Source Address,IPv6 Destination Address,IPv6 Source Port,IPv6 Destination Port

Packet Distribution and the Advanced Hash Algorithm

- When an **IPv4 Fragmentation** map rule is used to send traffic to an advanced hash tool GigaStream, all fragments are consolidated to a single port within the GigaStream.
- Packets with multiple VLAN tags (such as Q-in-Q) will experience uneven traffic distribution. For this traffic, GigaSMART load balancing is recommended.

Weighted GigaStream

Weighted GigaStream provides you the ability to distribute traffic to the ports by assigning either an equal weight or a custom weight to the ports. You can assign custom weight in percentage or ratio. If a port in a weighted GigaStream goes down, the traffic from the port will be redistributed to other healthy ports in the weighted GigaStream. The port assigned with maximum weight receives more traffic than the ports assigned with lesser weight.

Weighted GigaStream is supported on the following:

- All GigaVUE HC Series and GigaVUE TA Series nodes.
- Regular tool GigaStream, regular hybrid GigaStream, and regular circuit GigaStream.

Use the **Weighting** and **Drop Weight** fields in the GigaStream configuration page to configure a weighted GigaStream. For instructions, refer to [Configure Regular GigaStream](#).

You can also choose to rehash the traffic when you find that the traffic distribution is not ideal. When you rehash the traffic, GigaVUE-FM reassigns the hash buckets to the ports. For example, the following table shows that the ports are assigned with sequential hash buckets:

Port	Hash Buckets
x1	1, 2, 3, 4
x3	5, 6, 7, 8
x4	9, 10, 11, 12

In such cases, the traffic distribution may not be ideal. You can choose to rehash the traffic. The following table shows how the hash buckets are reassigned when you rehash the traffic:

Port	Hash Buckets
x1	1, 4, 7, 10
x3	2, 5, 8, 11
x4	3, 6, 9, 12

GigaStream Rules and Maximums

The following rules apply to regular GigaStream and controlled GigaStream:

GigaStream Rule	Description
Port Location	All participating ports must be on the same GigaVUE node. On the GigaVUE-HC3, GigaStream can be across modules.
Speed Requirements	- All participating ports must be running the same speed (1Gb, 10Gb, 40Gb, or 100Gb) and must use the same port types (for example, all g, x, q, or c). - A stack GigaStream must consist of ports with 10Gb speed or higher. - The system will not let you change the speed of any port participating in a

GigaStream Rule	Description
	tool GigaStream. Keep in mind that the only ports that allow speed changes through are gx ports. You can use gx ports in a regular tool GigaStream, but only with ports running at the same speed (and no slower than 1000Mb).
Addressing	Once a port belongs to a GigaStream, it must be addressed by its GigaStream alias. It can no longer be addressed as an individual port. For example, if tool port 1/1/x4 is part of a tool GigaStream, the GigaVUE® HC Series node prevents you from using it as the destination for a map rule.
Stack Ports	Stacking ports must be 10Gb or higher. Therefore, the Maximum Stack Ports per GigaStream for any 1Gb port is N/A in : Maximum Ports per GigaStream to GigaStream Rules and Maximums
SFP+	For SFP+ ports that can operate at 10Gb or 1Gb, refer to the values in the Maximum Tool Ports per GigaStream column for the 10Gb Ports rows in : Maximum Ports per GigaStream to GigaStream Rules and Maximums .

For GigaStream maximums, refer to [Maximum Ports per GigaStream](#), which have been updated in software version 5.1.

Maximum Ports per GigaStream

[Table 9: : Maximum Ports per GigaStream](#) lists the maximum ports per GigaStream for GigaVUE nodes.

Table 9: : Maximum Ports per GigaStream

Platform	Maximum Stack Ports per GigaStream	Maximum Tool Ports per GigaStream
GigaVUE-TA 10	32	256
GigaVUE-TA 40	64	256
GigaVUE-TA 100	64	256
GigaVUE-TA200	64	256
GigaVUE-TA200E	128	128
GigaVUE-TA25	56	56
GigaVUE-TA25E	64	64
GigaVUE-TA400	128	128

Platform	Maximum Stack Ports per GigaStream	Maximum Tool Ports per GigaStream
GigaVUE-TA400E	128	128
GigaVUE-HC1-PLUS	64	256
GigaVUE-HCT	64	256
GigaVUE-HC1	64	256
GigaVUE-HC 3	64	256

Port Statistics and Counters

This section describes the counters displayed for the Port Statistics information. This page provides information similar to the output from the **show port stats** command from the CLI.

The major sections in This section include:

- [Display Port Statistics](#)
- [How to Clear Traffic Counters](#)

Display Port Statistics

From the **Ports** page, you can view the port statistics for either of the following:

- A single port. Refer to [Display Port Statistics for a Single Port](#)
- All the ports. Refer to [Display Statistics for All Ports](#)

Display Port Statistics for a Single Port

From the **Ports > Ports > All Ports**, select any port by clicking on the row. The quick view window that appears provides port statistics for that particular port. Each field is color-coded in the graphical representation.

By hovering over the graph, numerical value for each of the data points is visible as shown in [Figure 25 Numerical Values for Data Points on the Statistics Graph](#).




Figure 25 Numerical Values for Data Points on the Statistics Graph

NOTE: GigaVUE-FM will display 9-11 data points in the graph.

You can modify the time lapse for measuring various data points by selecting the Minutes, Hour, Day, or Week button.

[Table 10: Port Statistics Definitions](#) describes the port statistics available.

Table 10: Port Statistics Definitions

Counter	Definition	Notes
Packet Errors	Total Error Packets Received or Transmitted This indicates hardware detected errors. Error packets include undersize, FCS/CRC, MTU exceeded, fragments, and oversize packets. All packets that list under this counter are discarded and not processed further.	Excludes oversize packets without FCS/CRC. Packets larger than the MTU setting arriving on a network port are counted twice in the counter. For example, 1000 oversize packets would show up as 2000. This double-counting only happens with Oversize error packets.
Discards	Discards Received or Transmitted	Discards are counted in the following cases: - Traffic arriving at a network port that is not logically connected using a map or map passall. - Map rules. - Packets on a tool port enabled by force link up. - Drop rules in an egress filter. - Pause frames.
Packet Drops	Total Dropped Packets Received or Transmitted	Packets are dropped when a port's bandwidth is exceeded due to oversubscription. Packets are dropped when they reach the port but before they are sent out.
Octets	Total Bytes Received or Transmitted	Excludes undersize frames.

Counter	Definition	Notes
	Includes all valid and error frames with the exceptions noted in the adjacent columns.	
Buffer Usage	Percentage of buffer space used by packets transmitted or received	The buffer is used when the port reaches 100 percent utilization such as during a microburst. If the buffer reaches 100 percent utilization, packets will be dropped.
Utilization	Percentage of port utilization by packets received or transmitted	
Packets	Total Packets Received or Transmitted Excludes multicast packets, broadcast packets, packets with FCS/CRC errors, MTU exceeded errors, oversized packets, and pause packets.	Excludes packets with FCS/CRC errors.

Display Statistics for All Ports

To view the statistics for all ports select, **Ports > Ports > Statistics**. The **Ports Statistics** page displays, which shows a table with the statistics for each port as shown in [Figure 26Port Statistics for All Ports](#). For the definitions of the statistics shown in the table, refer to [Table 10: Port Statistics Definitions](#).

	Port ID	Octets		Octets /sec		Unicast Packets		Non-Unicast Packets		Packets /sec		Packet Drops	Discards		Error		Utilization	
		Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Rx	Tx	Rx	Tx	Rx	Tx
☐	 1/1/x20	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
☐	 vTunnelEn...	2.38 G	83.03 K	36.42 K	0	2.26 M	1.18 K	310.9 K	0	49	0	0	0	0	0	0	0.03	0
☐	 toRSA5ecu...	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	-	-
☐	 1/2/q5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	-	-

Figure 26 Port Statistics for All Ports

The information on the statistics page can be filtered as well as downloaded to an Excel spreadsheet. To export the statistics, click **Export**. The statics table is downloaded with a filename in the format Port_Stats_<yyyymmddhhmmss>; for example, Port_Stats_20161003172336.

To filter the port statistics, click **Filter**. A Filter Quick View opens, where you can specify how to filter ports displayed on the Statistics page.

The criteria that you can use to filter the port statistics is as follows:

Criteria	Description
Box/Slot ID	Display only those ports that match the specified box and slot IDs.
Port Alias	Display port with the specified alias.
Port ID	Display ports with specified number in the port ID. For example, if you specify 3 the result will also display ports that include the number 3, 13, 23, 30, and so on.
Type	Display ports with the specified port type. Select one of the following: • Network - ports that receive incoming traffic from various sources. • Tool - egress ports where data is sent out to monitoring and analysis tools. • Inline Network - ports to which end-point devices are attached in an Inline Bypass configuration. • Circuit port- ports to send or receive traffic between clusters. • Inline Tool - ports that are used to connect inline security or monitoring tools in an Inline Bypass solution. • GigaSMART - ports are internal ports on GigaSMART line cards or modules used to power GigaSMART features. • Hybrid - physical port that serves a dual function as both an indirect traffic source port and a tool port. • Stack - ports to connect multiple GigaVUE nodes in a unified cluster. • Backplane - transfer traffic from front ports to the GigaSMART engines through XAU11 interfaces. • XAU11 - facilitate communication between GigaSMART engines, particularly when a second-level map is implemented.
Admin Status	Display ports based on their current admin status. The possible selections are: • All — display ports with a status of Enabled or Disabled. This is the default. • Enabled — display ports with admin enabled

Criteria	Description
	Disabled — display ports with admin disabled
Link Status	Display ports based on their current link status: The possible selections are: - All — display ports with a status of Up or Down. This is the default. - Up— display ports with a link status of up. - Down— display ports with a link status of down.

How to Clear Traffic Counters

To clear the traffic counters, do the following:

1. On the left navigation pane, go to  **Ports> Ports**.
2. Go to **Fabric Statistics** page and select a port.
3. Click **Clear** to clear the traffic counters in the Hardware.

The traffic counters will now be cleared.

Header Stripping

With the Header Stripping functionality enabled on the network and hybrid ports, the GigaVUE devices can identify and remove headers from tagged or tunneled (encapsulated) packets. Refer to the following sections for details:

- [About VXLAN Header Stripping](#)
- [About MPLS Header Stripping](#)


About VXLAN Header Stripping

The VXLAN Header Stripping functionality can identify and remove headers from the VXLAN tagged packets that are tapped from the VXLAN-based enterprise networks and are routed to the respective tools for analysis. This functionality is useful when working with tools that either cannot recognize the VXLAN headers or must engage in additional processing to analyze the VXLAN traffic. The GigaVUE-OS device is configured with the required traffic intelligence capability to strip the VXLAN header from the incoming packets. It then sends the inner payload to the tools based on the map rules configured.

The incoming VXLAN packets are discarded until the GigaVUE-OS device acquires the capability to remove the VXLAN header.

Once the required intelligence is acquired, Header Stripping can happen at line rate.

The following figure illustrates the VXLAN Header Stripping functionality:



In this diagram, the VXLAN encapsulated traffic from the network is tapped on the network port, 1/1/c5 in the GigaVUE-TA200 device. The VXLAN Header Stripping functionality is enabled in this network port. Based on the traffic intelligence capability, the GigaVUE-TA200 device strips the VXLAN header from the incoming packets that has the L4 destination port as 4789. You can choose to configure the L4 destination port at the node level. You can also define non-standard port (47889) instead of standard port (4789). The non-standard port is also known as iVXLAN. The GigaVUE-TA200 device routes the inner payload to the respective tools based on the map rules configured.

Any other traffic that enters the VXLAN Header Stripping-enabled network port will also be processed similar to a normal by-rule map.

The following table provides the capabilities available for VXLAN tunnel decapsulation as against the capabilities available for VXLAN Header Stripping:

Capabilities	With GigaSMART	Without GigaSMART
IPv4 support	Yes	Yes
IPv6 support	Yes	No
Header stripping on GigaVUE-TA Series	No	Yes
Header stripping at line rate	No	Yes

GigaVUE-OS has a scan interval between 300 to 1000000 seconds to optimize the VXLAN ID processing capability. You must configure 0 to disable the scan interval.

NOTE: GigaVUE-OS restarts the VXLAN traffic intelligence capability on every reload.

VXLAN Tunnel Decapsulation Versus VXLAN Header Stripping

The difference between the VXLAN tunnel decapsulation and VXLAN Header Stripping is that in the case of VXLAN tunnel decapsulation, the traffic originates from and terminates at the GigaVUE-OS devices. So, the GigaVUE-OS devices are aware of the VXLAN IDs based on which the traffic is decapsulated. In the case of VXLAN Header Stripping, the GigaVUE-OS devices are configured with traffic intelligence capability to strip the VXLAN header from the incoming packets with any VXLAN IDs.

VXLAN Header Stripping – Rules, Notes, and Limitations

Keep in mind the following rules and notes when working with VXLAN Header Stripping:

- VXLAN Header Stripping is supported on GigaVUE-HC1-Plus, GigaVUE-HCT, GigaVUE-HC1, GigaVUE-HC3, GigaVUE-TA100, GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA400, and GigaVUE-TA400E devices.
- The destination IP based statistics is not supported on GigaVUE-TA25E, GigaVUE-HC1-Plus, GigaVUE-TA25, and GigaVUE-HCT.
- The VXLAN header stripped packet does not match the VLAN qualifier in maps due to chip set limitation in GigaVUE-TA25E, GigaVUE-HC1-Plus, GigaVUE-TA25, and GigaVUE-HCT.
- Gigamon®'s traffic intelligence processing may lead to initial packet drops. This is not applicable for GigaVUE-TA400.
- On a network or hybrid port that taps the network traffic, you can enable either VXLAN Header Stripping or MPLS Header Stripping, but you cannot enable both the functionalities.

- Network ports configured with VXLAN Header Stripping and MPLS Header Stripping functionalities cannot be part of the same map. You must create separate maps for these ports.
- After a header is removed from the packet, FCS is recomputed by the hardware and not by the GigaVUE-OS.
- VXLAN Header Stripping is not supported on network ports that are part of pass-all maps except on GigaVUE-TA400.
- VXLAN Header Stripping is not supported with IPv6 addresses except on GigaVUE-TA400.
- You cannot associate an IP interface with a network or hybrid port that is enabled with VXLAN Header Stripping.
- Filter rule is not supported on hybrid port that is enabled with VXLAN Header Stripping.
- Ingress VLAN tagging is not supported.
- VXLAN Header Stripping is not supported for Q-in-Q traffic except on GigaVUE-TA400.
- You cannot enable VXLAN Header Stripping on a port that is part of a port-pair.
- Reassembly of fragmented packets after VXLAN Header Stripping is not supported.
- A maximum of up to 4096 dynamic VXLAN IDs are supported for VXLAN Header Stripping. On GigaVUE-TA400, and GigaVUE-TA400E, all VXLAN IDs are supported.
- If a map has both Header Stripping-enabled ports and other network ports, the VXLAN traffic that enters the other network ports will not be sent to the shared collector except on GigaVUE-TA400.
- VXLAN Header Stripping does not work if you configure a map with any rule that includes the qualifying attributes of the VXLAN header because such rules override the traffic intelligence capability. For example, if you configure a map with pass rule as IPv4 Destination, IPv4 Source, MAC Destination, or MAC Source and the source port of the map is overlapped/matched with VXLAN headers, the Header Stripping functionality does not work. This is an exception for GigaVUE-TA400.
- VXLAN Header Stripping does not work if you configure a map with only drop rules and choose the **Pass Traffic** option so that the traffic is passed through the port when there are no matching rules. For more information, refer to [Map Types](#). This is an exception for GigaVUE-TA400.
- Following table provides the maximum number of static IP addresses that can be configured for each platform for the VXLAN Header Stripping functionality:

Platform	Maximum number of static IP addresses supported
Platform	Maximum number of static IP addresses supported
GigaVUE-HC1	3966
GigaVUE-HC3, GigaVUE-TA100, and GigaVUE-TA200	1918

Configure VXLAN Header Stripping

Before you configure VXLAN Header Stripping, refer to [VXLAN Header Stripping – Rules, Notes, and Limitations](#).

The following table summarizes the required tasks to configure VXLAN Header stripping to strip the incoming VXLAN packets:

S.No	Task	Refer to..
1.	Configure the required network or hybrid port.	Configure Ports
2.	Enable VXLAN Header Stripping on the network or hybrid port that you configured in task 1.	Enable Header Stripping Protocol on Ports
3.	It is recommended that you configure the aging interval for the VXLAN Header Stripping to refresh the traffic intelligence capability for the device.	Configure Ageing Interval for VXLAN Header Stripping
4.	Configure a new map. Keep in mind the following details when you configure a map: • Ensure that you select the network or hybrid port on which you enabled VXLAN Header Stripping as the source port of the map. • Ensure that you do not select any rule that includes the qualifying attributes of the VXLAN header because such rules override the traffic intelligence capability. For more details, refer to VXLAN Header Stripping – Rules, Notes, and Limitations .	Create a new map
5.	View the VXLAN and destination IP statistics to know the number of packets transmitted for a VXLAN ID and a destination IP address, and the transmitted bytes per packet.	View Header Stripping Statistics

Configure Ageing Interval for VXLAN Header Stripping

It is recommended that you configure the ageing interval to refresh the traffic intelligence capability based on which GigaVUE-OS devices strips the VXLAN headers from the incoming VXLAN traffic.

To configure the ageing interval:

1. From the left navigation pane, go to **System > Chassis**. Toggle to view the list view.
2. Select the Box ID on which you want to configure the Header Stripping protocol.
3. From the **Actions** drop-down list, select **Configure Header Stripping**. The Configure Header Stripping page appears.
4. From the **Header Stripping Protocol** drop-down list, select **VXLAN**.
5. In the **Aging Interval** field, enter the interval in seconds to refresh the traffic intelligence capability for the chassis.
6. Click **Save**.

The screenshot shows the 'Configure Header Stripping' configuration page. At the top, there is a title bar with the text 'Configure Header Stripping' and two buttons: 'save' and 'Cancel'. Below the title bar, there are two main configuration fields. The first field is 'Header Stripping Protocol', which has a dropdown menu currently showing 'VXLAN'. The second field is 'Aging Interval', which has a text input field containing the value '0 to 16777215' and a unit label 'seconds (0 to disable)'.

View Header Stripping Statistics

To view the Header Stripping statistics for a VXLAN ID and destination IP address:


1. From the device view, go to **Ports > Ports > Header Stripping Statistics**.
2. From the **Statistics** drop-down list, select one of the following options:
 - **IP Statistics** – View the number of packets transmitted and the bytes per packet for a destination IP address.
 - **VXLAN Statistics** – View the number of packets transmitted and the bytes per packet for a VXLAN ID.

NOTE: The destination IP based statistics is not supported on GigaVUE-TA25E, GigaVUE-HC1-Plus, GigaVUE-TA25, and GigaVUE-HCT.

About MPLS Header Stripping

The MPLS Header Stripping functionality can identify and remove headers from the MPLS traffic tapped at the network and routed to the respective tools for analysis. The GigaVUE-OS device strips the MPLS header based on the MPLS label configured for the device. The outer and the inner MPLS labels must be configured for the device. The MPLS traffic is then routed to the respective tools based on the map rules configured to match the MPLS payload.

The following figure illustrates the MPLS Header Stripping functionality.




In this diagram, the L3-MPLS encapsulated traffic from the network is tapped on the network port, 1/1/c1 in the GigaVUE-TA200 device. The L3-MPLS Header Stripping functionality is enabled on this network port. Based on the MPLS label configured for the GigaVUE-TA200 device, the device strips the MPLS header from the incoming packets and routes the MPLS payload to the respective tools based on the map rules configured.

Any other traffic that enters the L3-MPLS Header Stripping-enabled network port will also be processed similar to a normal by-rule map.

Advanced MPLS Header Stripping for GigaVUE-TA400

The GigaVUE-TA400 introduces advanced MPLS Header Stripping capabilities, enabling support for all Layer 2 and Layer 3 MPLS traffic types. Header stripping applies only to the outer MPLS header in nested MPLS scenarios, and up to 7 MPLS labels can be stripped. You can configure the Advanced MPLS Header Stripping options as required. If multiple options are configured, stripping logic will be applied per the diagram below.

NOTE: When L3 MPLS options are configured, then any traffic that is not L3 MPLS will be as is without stripping.



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

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

L3 MPLS IPv4 without Options - If the first byte after the BOS label is 0x45, it strips MPLS headers.

L3 MPLS IPv4 with Options - If the first byte after the BOS label is in the range 0x46 to 0x4F, it strips MPLS headers.

L2 MPLS with CW/ACH

- If the first nibble after the BOS label is 0x1, it strips the outer ethernet and MPLS headers along with Associated Channel Header(ACH).
- If the first nibble after the BOS label is 0x0, it strips the outer ethernet and MPLS headers along with Control Word (CW).

L2 MPLS without CW/ACH - If the first nibble after the BOS label does not have a value matching in the list [0x0, 0x1, 0x45, 0x46 to 0x4F, and 0x6], it strips the outer ethernet and MPLS headers.

The table below describes the Header strip configuration and its potential misclassification.

Header strip configuration	MPLS packet type classification
L2 MPLS without CW/ACH, L3 MPLS IPv4 without options	L2 MPLS packets without CW/ACH having Packet DMAC starting with 0x45 will be misidentified as MPLS IPv4 without options.
L2 MPLS without CW/ACH, L3 MPLS IPv4 with options	L2 MPLS packets without CW/ACH having Packet DMAC starting with ranges 0x46-0x4F will be misidentified as MPLS IPv4 with options.
L2 MPLS without CW/ACH, L3 MPLS IPv6	L2 MPLS packets without CW/ACH having Packet DMAC starting with 0x6 will be misidentified as MPLS IPv6.
L2 MPLS with CW/ACH, L2 MPLS without CW/ACH	L2 MPLS packets without CW/ACH having Packet DMAC starting with 0x0 or 0x1 will be misidentified as L2 MPLS with CW/ACH.
L2 MPLS without CW/ACH	L2 MPLS packets with CW/ACH will be misidentified as L2 MPLS without CW/ACH and only the Outer Ethernet, MPLS headers will be stripped without the PW CW from the packet.

To configure Advanced MPLS Header Stripping options, refer to [Enable Header Stripping Protocol on Ports](#).

MPLS Header Stripping – Rules, Notes, and Limitations

Keep in mind the following rules and notes when working with MPLS Header Stripping:

- MPLS Header Stripping is supported on GigaVUE-HC1-Plus, GigaVUE-HCT, GigaVUE-HC1, GigaVUE-HC3, GigaVUE-TA100, GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA400, and GigaVUE-TA400E devices.
- On GigaVUE-TA25, the outer tag is also stripped along with MPLS labels of incoming tagged MPLS traffic.
- On GigaVUE-HC1-Plus, GigaVUE-TA25, GigaVUE-TA25E, and GigaVUE-HCT devices, ingress VLAN tag functionality will not be supported on MPLS header strip enabled ports.
- In legacy stacking mode, the MPLS header stripped packets cannot traverse a cluster stack link and reach the destination ports in remote node on GigaVUE-TA25E, GigaVUE-HC1-Plus, GigaVUE-TA25, and GigaVUE-HCT. It is expected that the source and destination ports reside in the same device. This is not a limitation in case of ETAG

mode. The customer VLAN gets stripped off in ETAG mode when the source port is enabled for MPLS Header Stripping in GigaVUE-TA25E, GigaVUE-HC1-Plus, GigaVUE-TA25, and GigaVUE-HCT.

- In GigaVUE-TA25E, the MPLS traffic received on the MPLS header strip enabled port for which the corresponding MPLS labels are not programmed in the device already, hits the shared collector and gets forwarded to tools with an additional VLAN tag. This is not applicable in Etag
- MPLS Header Stripping is supported only on traffic with a maximum of up to two MPLS labels (inner and outer MPLS labels) configured for a device. On GigaVUE-TA400, MPLS Header Stripping is supported with a maximum of seven MPLS labels per packet.
- MPLS Header Stripping is not supported with IPv6 addresses except on GigaVUE-TA400.
- A maximum of up to 4096 MPLS labels can be configured for a device. If the label configuration exceeds the maximum limit specified and if the new labels overlap with the already configured labels, you must delete the existing labels and reconfigure the complete set of labels without exceeding the maximum limit. This is not applicable for GigaVUE-TA400.
- All the MPLS labels that are used to strip the MPLS header must be configured at the chassis-level so that all the network or hybrid ports enabled with the MPLS Header Stripping functionality is aware of the labels. 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, and GigaVUE-TA400E, both L3 MPLS outer Header Stripping and L2 MPLS outer Header Stripping of traffic carrying with and without Pseudowire MPLS control word (PWMCW) or Pseudowire associated channel header (PWACH) are supported.

Configure MPLS Header Stripping

Before you configure MPLS Header Stripping, refer to [MPLS Header Stripping – Rules, Notes, and Limitations](#).

The following table summarizes the required tasks to configure MPLS Header stripping to strip the incoming MPLS packets:

S.No	Task	Refer to..
1.	Configure the required MPLS labels at the chassis level.	Configure MPLS Labels on a Chassis
2.	Configure the required network or hybrid port.	Configure Ports
3.	Enable MPLS Header Stripping on the network or hybrid port that you configured in task 2.	Enable Header Stripping Protocol on Ports
4.	Configure a map. Keep in mind the following details when you configure a map: - Ensure that you select the network or hybrid port on which you enabled MPLS Header Stripping as the source port of the map. - Configure the map rules with the following parameters: - L2 parameters such as EtherType, MAC Destination, MAC Source, VLAN, or Inner VLAN. - MPLS payload parameters such as IP Version, IPv4 Destination, IPv4 Source, IPv6 Destination, IPv6 Source, Protocol, UDA1, or UDA2. For more details, refer to MPLS Header Stripping – Rules, Notes, and Limitations .	Create a New Map

Create a New Map

The following are the steps for creating a map:

- 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](#).
- From the device view, go to **Maps > Maps** to open the Maps page.
- Click **New**.

4. Enter the Map Information:

a. Enter an alias for the map.

Use an alias that helps identify the task and destination. For example, netflix_traffic_to_wireshark.

b. (Optional) Enter a description of the map. When adding a description, consider the following:

- Use up to 128 characters, including spaces.
- Enclose the description in quotation marks if the description is longer than one word.
- To include double quotation marks (") inside the quotation marks, precede it with a backslash (\).

See also [Configure MPLS Header Stripping](#).

c. Select the **Type**.

The map type can be Regular, First Level, Second Level, or Inline.

For detailed information about the types of maps, refer to [Map Types](#)

d. Select the map **Subtype**.

The map subtype can be **By Rule**, **Pass All**, or **Collector**.

For detailed information about Pass All, refer to [About Map-passall Maps](#). For detailed information about Collector, refer to [About Shared Collectors](#).

e. Enable the Pass Traffic checkbox if no rules match.

f. Enable the Control Traffic checkbox to pass the GTP control traffic (GTP-c) to all GigaSMART engines in a GTP engine group. A GTP engine group has multiple GigaSMART engine port members.

NOTE: The Control Traffic checkbox is applicable only for GTP and is displayed only if the map type is configured as First level, and the map sub type is configured as By Rule.

5. Specify the **Map Source and Destination**.a. From the **Source** and **Destination** drop-down list, select the required source and destination ports for the map. To create a port list, click **Port Editor**.

NOTE: You can add a maximum of 324 ports in the **Source** drop-down list, if the ports are not attached to a GigaStream.

NOTE: For details about port types that are supported for the different types of maps, refer to [Port Lists](#).

- b. If you have selected a circuit port in the **Destination** drop-down list, select the required circuit tunnel from the **Encapsulation Tunnel** drop-down list to encapsulate the traffic.

NOTE: For details about circuit tunnels, refer to [About Circuit-ID Tunnels](#).

- c. If the map is used to redirect the decapsulated traffic to the required tool ports, ensure that you select the IP interface in the **Source** drop-down list. You must have attached the IP interface to the VXLAN or L2GRE tunnel. For details about VXLAN or L2GRE tunnels, refer to [About Virtual Extensible LAN \(VXLAN\) Tunnels](#) and [About Layer 2 Generic Routing Encapsulation \(L2GRE\) Tunnels](#). From the **Destination** drop-down list, select the required tool ports.
 - d. If the map will use a GigaSMART operation, select the Header Stripping (MPLS) operation that you have created earlier from the **GigaSMART Operations (GSOP)** drop-down list.
 - e. If you select the tool GigaStream for the destination, you can view the utilization value for the GigaStream.
6. Add rules to the map.

To add rules to the map, do any of the following:

- Use the **Quick Editor**. For details, refer to the [Configure MPLS Header Stripping](#).
- Import a map template by clicking **Import**.
- Create a rule by clicking **Add a Rule**.

For detailed information about map rules, refer to [Map Rules](#).

7. Set the **Map Order** by selecting the priority from the Priority list.

For details about map priority, refer to [Map Priority](#).

8. Set the **Map Permissions**.

For details about map permissions, refer to [Port Access and Map Sharing](#).

9. Set the **Map Tag**.

Select the required tag key and tag value to which the map must be associated. The tag key and the associated tag values must be created in advance in GigaVUE- FM. Refer to the "Tags" and "Role Based Access Control" sections in the GigaVUE Administration Guide for more details.

NOTE: When you associate a map to a tag value, then the ports, port groups, port pairs, GigaStreams that belong to the map are also associated to the tags.

Configure MPLS Labels on a Chassis

You must configure MPLS labels based on which the device strips the MPLS header from the incoming traffic. This configuration is not applicable for GigaVUE-TA400.

To configure MPLS labels on a chassis:

1. From the left navigation pane, go to **System > Chassis**. Toggle to view the list view.
2. Select the Box ID on which you want to configure the Header Stripping protocol.
3. From the **Actions** drop-down list, select **Configure Header Stripping**. The Configure Header Stripping page appears.
4. From the **Header Stripping Protocol** drop-down list, select **MPLS-L3**.
5. Select the **Apply to All Box** check box to configure the MPLS labels on all the chassis.
6. In the **Labels** field, enter the MPLS labels that you want to configure on the required box. You can also choose to add a range of labels.
7. Click **Save**.

Configure Header Stripping save Cancel

Header Stripping Protocol MPLS-L3

Apply to All Box ☒

Labels

Enable Header Stripping Protocol on Ports

You can enable Header Stripping protocol on any network or hybrid ports. If you enable the Header Stripping protocol on a port, you cannot configure the VXLAN ID or L2GRE ID for that port.

To enable a Header Stripping protocol on a network or hybrid port:

1. From the device view, go to **Ports > Ports > All Ports**.
2. Select the required network or hybrid port on which you want to enable the Header Stripping protocol, and then click **Edit**.

3. Under the Parameters section, from the **Header Stripping Protocol** drop-down list, select one of the following options:

- **None** – Header stripping protocol is not enabled on the port.
- **MPLS-L3** – Port is enabled with MPLS Header Stripping protocol to strip the MPLS header from the MPLS layer 3 packets tapped on this port. Select this option if L3 MPLS traffic type only is present in network..
- **VXLAN** – Port is enabled with VXLAN Header Stripping protocol to strip the VXLAN header from the VXLAN packets tapped on this port.
- **MPLS-Advanced** – Enable advanced MPLS Header Stripping on the port. This functionality applies only to the GigaVUE-TA400 devices. Refer to [Advanced MPLS Header Stripping for GigaVUE-TA400](#). This option supports the stripping of various MPLS packet formats and provides additional sub-options:
 - **IPv4 Without Options** – Strips the MPLS Header from the L3 MPLS IPv4 (without options) packets. Select this option only if L3 MPLS IPv4 (without options) traffic type is present in network. IPv4 without options has a standard 20-byte header for efficient routing.
 - **IPv4 With Options** – Strips the MPLS Header from the L3 MPLS IPv4 (with options) packets. Select this option only if L3 MPLS IPv4 (with options) traffic type is present in network. IPv4 with options includes additional fields (up to 60 bytes) for advanced functions like timestamping.
 - **IPv6** – Strips the MPLS header from the L3 MPLS IPv6 packets. Select this option only if L3 MPLS IPv6 traffic type is present in network. The standard IPv6 header (40 bytes) follows the MPLS label stack and contains source and destination IPv6 addresses, flow labels, and other fields.
 - **L2 Without CW/ACH** – Strips the outer Ethernet and MPLS headers from the incoming L2 MPLS without Control Word(CW) or Associated Channel Header(ACH) packets. Select this option only if L2 MPLS (without CW/ACH) traffic type is present in network. L2 MPLS without CW or ACH provides a basic MPLS label-switched path for transporting Layer 2 packets without the extra control, sequencing, or signaling functionalities provided by the CW and ACH headers.
 - **L2 With CW/ACH** – Strips the outer Ethernet and MPLS headers along with Control Word (CW) or Associated Channel Header(ACH) from the incoming L2 MPLS with CW or ACH packets. Select this option only if L2 MPLS (with CW/ACH) traffic type is present in network. L2 MPLS with CW ensures packet sequencing and integrity for Layer 2 frames, while ACH enables Operations, Administration, and Maintenance (OAM) transport for network monitoring and management, especially in MPLS traffic.

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

Edit Port(s): 3/1/d2

Parameters

Admin ☐ Enable

Type

Duplex ☒ Full ☐ Half

Auto Negotiation ☐ Enable

VLAN Tag

Egress Vlan Tag ☐ None ☐ Strip

Force Link Up ☐ Enable

Receive Only ☒ Enable

VXLAN ID

L2GRE ID

Header Stripping Protocol

MPLS-Advanced

☒ ip4-w/o-options ⓘ ☒ ip4-options ⓘ ☐ ipv6 ⓘ

☐ l2-w/o-cw-ach ⓘ ☐ l2-cw-ach ⓘ

ⓘ Choose only those MPLS traffic types ingressing on these port(s) to prevent any misidentifications.

Tunnels

Tunneling is a communication protocol that is used to transmit data from one network to another by encapsulating the data. A tunnel is a virtual interface. You can create tunnels for both encapsulation and decapsulation.

Required License: Advanced Feature License on GigaVUE TA Series Nodes

This chapter describes about the different types of native tunnels, which are independent of GigaSMART operations. It also describes how to configure these tunnels for encapsulating and decapsulating traffic. Refer to the following sections for details:

- [About Circuit-ID Tunnels](#)
- [About Layer 2 Generic Routing Encapsulation \(L2GRE\) Tunnels](#)
- [About Virtual Extensible LAN \(VXLAN\) Tunnels](#)
- [Create Tunnel](#)
- [Create VXLAN / L2GRE Group](#)
- [Configure L2GRE / VXLAN Identifier](#)
- [View VXLAN / L2GRE ID Statistics](#)

About Circuit-ID Tunnels

Circuit-ID tunnels are used to route traffic between two clusters. The traffic is tapped and sent through network ports that are configured on the cluster in the encapsulation side. Based on the flow map configuration, traffic is filtered and then sent through the circuit ports that are configured as the destination port in the map. These circuit ports encapsulate the traffic with a circuit-ID and transmit the encapsulated traffic through the circuit tunnel that connects two clusters. At the receiving end of another cluster in the decapsulation side, the circuit port that is configured as the source port decapsulates the traffic and sends the traffic to the appropriate tools through the tool ports.

The following figure illustrates the circuit Flow Mapping® between two clusters using circuit-ID.




Figure 27 Circuit Flow Mapping®

In this example, the GigaVUE TA Series and GigaVUE HC Series nodes reside in cluster A. The tapped traffic is sent through network ports, 1/1/n1 and 1/1/n2. Based on the rules configured in the map, the traffic is filtered at the nodes in cluster A. The filtered traffic is then sent through the circuit port, 1/1/g1 and circuit GigaStream that are configured as the destination port in the map. These circuit ports encapsulate the traffic with circuit-ID and transmit the encapsulated traffic through the circuit tunnel that connects cluster A and cluster B. At the receiving end of cluster B, the circuit ports, 1/1/g2, 1/1/g3, 1/1/g4, and 1/1/g5 that are configured as the source ports decapsulate the traffic, strip the circuit-ID, and send the traffic to the appropriate tools through the tool ports, 1/1/x1, 1/1/x2, and 1/1/x3.

Refer to the following sections for details about the Circuit-ID tunnel encapsulation and decapsulation:

- [Circuit-ID Tunnels—Rules and Notes](#)
- [Circuit-ID Tunnel Encapsulation](#)

- [Circuit-ID Tunnel Decapsulation](#)

Circuit-ID Tunnels—Rules and Notes

Keep in mind the following rules and notes when working with Circuit-ID tunnel encapsulation and decapsulation:

- A maximum of 1024 fabric maps can be created between any two given clusters or nodes in non-shared mode.

NOTE: Prior to GigaVUE-OS release 6.12, a maximum of 512 unique circuit IDs are supported within each cluster or standalone device for encapsulation and decapsulation.

- If a network port receives a double-tagged packet that is encapsulated with a circuit-ID, the five tuple hashing will not work only in the second cluster, that is the cluster in the decapsulation side over stack GigaStream or tool gigastream. Hence, traffic cannot be filtered using the IP/L4 parameters. After decapsulation, Flow Mapping® filters the traffic based on circuit-ID.
- It is not supported for inline scenarios.

Keep in mind the following rules and notes when working with Circuit-ID tunnel encapsulation:

- Circuit-ID tunnel encapsulation is not supported on Pass All maps.
- Port filter configured on circuit port for VLAN pass/drop will try to match the encapsulation circuit-id instead of packet outer VLAN.

Keep in mind the following rules and notes when working with Circuit-ID tunnel decapsulation:

- Circuit-ID tunnel decapsulation is not supported on Pass All and Shared Collector maps.
- A circuit-ID must be paired with a circuit port, only in one circuit-ID tunnel.

Keep in mind the following rules and notes when working with Fabric map circuit ID allocation:

- Circuit ID is used internally to pass traffic from one cluster to another. You can configure Circuit ID ranges from 2 to 4000, or set your own custom range, with each topology able to reuse ID ranges (e.g., 2-513) across multiple topologies.
- Circuit ID allocation is managed globally, beginning from the lower limit of the defined range. This ensures that IDs are allocated efficiently across different topologies. It is allocated automatically.

Circuit-ID Tunnel Encapsulation

Before creating a Circuit-ID tunnel for encapsulation, refer to the [Circuit-ID Tunnels—Rules and Notes](#).

The following table summarizes the required tasks to configure a circuit-ID tunnel for encapsulating the traffic:

S.No	Task	Refer to...
1.	Configure the required circuit ports and circuit GigaStream.	Configure Ports Configure Regular GigaStream
2.	Configure a circuit-ID tunnel for encapsulation. Ensure that you select the mode as Encap .	Create Tunnel
3.	Configure a map to encapsulate the traffic and attach the circuit-ID tunnel to the map.	Create a new map

Circuit-ID Tunnel Decapsulation

Before creating a Circuit-ID tunnel for decapsulation, refer to the [Circuit-ID Tunnels—Rules and Notes](#).

The following table summarizes the required tasks to configure a circuit-ID tunnel for decapsulating the traffic:

S.No	Task	Refer to...
1.	Configure the required circuit ports and circuit GigaStream.	Configure Ports
2.	Configure a circuit-ID tunnel for decapsulation. Ensure that you select the mode as Decap and attach the circuit ports or circuit GigaStream that you configured in step 1 to the circuit-ID tunnel.	Create Tunnel
3.	Configure a map to decapsulate the traffic and ensure that you specify the circuit-ID as a pass/drop rule in the map.	Create a new map

About Layer 2 Generic Routing Encapsulation (L2GRE) Tunnels

L2GRE tunnels are used to route traffic from any remote device to a GigaVUE HC Series or GigaVUE-TA Series device over the internet. The device at the remote site encapsulates the filtered packets, adds a L2GRE encapsulation header, and forwards it to the corresponding circuit port that is used for GRE encapsulation. The encapsulation header consists of Ethernet + IP + GRE headers. The parameters of the encapsulated header are user-configurable, such as the IPv4 address of the IP interface on the destination GigaVUE device and the GRE key that identifies the source of the tunnel.

The encapsulated packet is sent out of the circuit port, which is connected to the public network (the Internet). This packet is routed in the public network to reach the main office site. The packet is ingressed at the circuit port of the GigaVUE device at the main office. The destination IP address of the received packet is checked against the IP configured for the circuit port. If they match, decapsulation is applied. The Ethernet + IP + GRE header is stripped and the remaining packet is sent to the tool port.

If the destination IP address of the received packet does not match with the IP address configured for the IP interface, the packet is dropped.

The following figure illustrates the L2GRE tunnel encapsulation and decapsulation.




Figure 28 L2GRE Tunnel Encapsulation and decapsulation

In this diagram, traffic is tapped on a GigaVUE-TA200 device at a remote site, and then it is tunneled through L2GRE encapsulation across the network before it reaches the GigaVUE-HC3 device at the main office site, which is connected to the actual tools. The tunnel decapsulation is executed on an ingress circuit port (IP interface). After tunnel decapsulation the packet is presented to the Flow Mapping® module to filter based on map rule parameters.

Refer to the following sections for details about the L2GRE tunnel configuration:

- [L2GRE Tunnel Configuration—Rules and Notes](#)
- [Limitation](#)
- [Configure L2GRE Tunnel to Encapsulate Traffic](#)
- [Configure L2GRE Tunnel to Decapsulate Traffic](#)

L2GRE Tunnel Configuration—Rules and Notes

Keep in mind the following rules and notes when working with L2GRE tunnels:

- L2GRE tunnels are supported only on GigaVUE-HC1-Plus, GigaVUE-HCT, GigaVUE-HC1, GigaVUE-HC3CCv1 and CCv2, GigaVUE-TA100, GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA400, GigaVUE-TA400E, and DELL S4112F-ON devices.
- A maximum of 1500 L2GRE IDs are supported.
- IPv6 L2GRE tunnels are supported only when using the GigaSMART engine and require an Advanced Tunneling license.
- Ingress VLAN tagging and Tool Mirror features are not supported with L2GRE tunnels.
- Filtering of Q-in-Q packets is not supported with L2GRE tunnels except on GigaVUE-TA400.
- Map-passall is not supported for the circuit port that encapsulates or decapsulates the L2GRE packet.
- When configuring a map for L2GRE encapsulation, you cannot configure a combination of a regular tool port and L2GRE encapsulation tunnel as part of the "To" ports.
 - Any encapsulated packet that exceeds the MTU value configured for the IP interface will be discarded because IP fragmentation and reassembly of packets are not supported.
 - L2GRE tunnel encapsulation is not supported on circuit GigaStreams.
 - Flow mapping that is configured on the circuit port used for L2GRE decapsulation will filter only the inner packet attributes along with L2GRE-ID. Any other non-tunneled packets that ingress on this circuit port will not be filtered or redirected to tool ports, even if it matches the rules configured on the map.
 - GigaSMART operations cannot be combined with L2GRE decapsulation in the same map.
 - L2GRE tunnel decapsulation is supported only on encapsulated packets that are not tagged. On GigaVUE-TA400, GigaVUE-TA400E L2GRE tunnel decapsulation is supported on encapsulated packets that are both tagged and untagged.
 - Inner VLAN qualifier is not supported on the port in which the L2GRE tunnel decapsulation is enabled except on GigaVUE-TA400.
 - L2GRE ID qualifier is available as part of existing static templates. Following table provides details about the platforms for which the static templates are available:

Table 11: L2GRE ID Qualifier Availability

Template	Platform	
	GigaVUE-HC1	GigaVUE-HC3/ GigaVUE-TA100/ GigaVUE-TA25/ GigaVUE-TA25E/ GigaVUE-TA200/ GigaVUE-TA200E/ GigaVUE-TA400/ GigaVUE-TA400E
IPv4	Not available	Available
IPv6	Available	Available
IPv4+UDA	Not available	Available
IPv4+MAC	Available	Available
UDA	Available	Available

Limitation

When the encapsulation device fragments your traffic, the L2GRE Tunnels used to decapsulate the traffic does not support re-assembly. To avoid this, you can use GigaSMART L2GRE decapsulation, which reassembles the fragmented packets. Refer to [GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation](#) for more detailed information on how to configure GigaSMART L2GRE Tunnel Decapsulation. You can also configure the highest possible MTU value before tapping the traffic to the virtual machine so that packets are not fragmented.

Configure L2GRE Tunnel to Encapsulate Traffic

Before creating a L2GRE tunnel to encapsulate traffic, refer to the [L2GRE Tunnel Configuration—Rules and Notes](#).

The following table summarizes the required tasks to configure a L2GRE tunnel for encapsulating the traffic:

S.No	Task	Refer to...
1.	Configure the required circuit port.	Configure Ports
2.	Configure an IP interface and attach the circuit port that you created in task 1.	Configure IP Interface
3.	Configure a L2GRE tunnel for encapsulation and attach the IP interface that you created in task 2. Ensure that you select the Type as L2GRE and the Mode as Encapsulation .	Create Tunnel
4.	Create a L2GRE group for a device and add all the L2GRE IDs that are specific to the device.	Create VXLAN / L2GRE Group
5.	Configure the L2GRE ID either at the chassis-level or at the port-level to help identify the encapsulation tunnel. You must use one of the L2GRE IDs that you have already added in the L2GRE group for the device in task 4.	Configure L2GRE / VXLAN Identifier
6.	Configure a map to encapsulate the L2GRE traffic. Ensure that you add the following details when configuring a map: • If you have configured the L2GRE ID for the chassis, you can configure any network port on the chassis as the source port of the map. • If you have configured the L2GRE ID for a network port, you must configure that network port as the source port of the map. • Configure the encapsulation tunnel that you created in task 3 as the destination port of the map. • Specify the required pass/drop rule in the map to filter the traffic based on inner packet attributes or L2GRE ID for the template configured.	Create a new map

Configure L2GRE Tunnel to Decapsulate Traffic

Before creating a L2GRE tunnel to decapsulate traffic, refer to the [L2GRE Tunnel Configuration—Rules and Notes](#).

The following table summarizes the required tasks to configure a L2GRE tunnel for decapsulating the traffic:

S.No	Task	Refer to...
1.	Configure the required circuit port.	Configure Ports
2.	Configure an IP interface and attach the circuit port that you created in task 1.	Configure IP Interface
3.	Configure a L2GRE tunnel for decapsulation and attach the IP interface that you created in task 2. Ensure that you select the Type as L2GRE and Mode as Decapsulation .	Create Tunnel
4	Create a L2GRE group for a device and add all the L2GRE IDs that are specific to the device.	Create VXLAN / L2GRE Group
5.	Configure a map to decapsulate the traffic. Ensure that you add the IP interface that you created in task 2 as the source of the map and specify the required pass/drop rule in the map to filter the traffic based on inner packet attributes or L2GRE ID for the template configured.	Create a new map

Orchestrated Workflow Configuration of L2GRE Tunnels

The L2GRE tunnels can now be configured through the Orchestrated Configuration page. To configure L2GRE Tunnels, follow the below steps:

- In GigaVUE-FM go to **Traffic>Physical> Orchestrated Flows> Tunnels**.
- Click on **New**.
- Select **Embedded L2GRE**. The New Tunnel configuration page appears. Enter the required information as described below:

Field	Description
Tunnel Name	The name of the tunnel. Note: Alias must not have spaces or any of these characters: *"?;:/%@.
Tunnel Description	The description of the tunnel endpoint.
Traffic Direction	Select the traffic direction of the tunnel. Choose DECAP for decapsulation or choose ENCAP for encapsulation.

Field	Description
	- If you choose Decap, select the IP Interface from the list. You can create one using the Create an IP interface option if you do not have an IP interface. Refer to IP Interfaces to learn more about creating an IP Interface. - If you choose Encap enter the Source Port. You can select from the drop-down list. To edit the Port type or to enable the port for Admin privileges use the Port Editor window. - You can also add a new L2GRE ID by clicking on the ADD L2GRE ID. Note: The L2GRE ID configured at source port value will take precedence over the L2GRE ID that is configured at chassis level.
Nodes	Select the Nodes on which you intend to create the tunnel. For L2GRE Encapsulation tunnel you will have to configure the L2GRE ID. - Click on Config L2GRE ID> ADD L2GRE ID. - Fill in the require parameters. - Click on ADD. - The L2GRE IDs will list in the Configuration page.
ID	For L2GRE Decapsulation tunnel you will have to select the L2GRE ID. To configure a new L2GRE ID do the following: - Click on ADD L2GRE ID. - Fill in the require parameters. - Click on ADD. Note: Configuring a L2GRE ID is not mandatory if the map rules are configured. Either configure the rules or a L2GRE ID.
Rules	Configure the map rules that needs to be adhered to while decapsulating or encapsulating the traffic. Click on Rules Editor to configure the rules. Note: Configuring rules are mandatory for Encapsulation tunnel, For a Decapsulation tunnel rules are not mandatory if a L2GRE ID is already configured.
Encapsulation IP interface	Enter the interface IP address of the node (Destination IP) for an encapsulation tunnel.
Remote Tunnel IP	Enter the Remote Tunnel IP values for and encapsulation tunnel.
Destination Port	Select the destination port for the decapsulation tunnel. You can edit the port details from the Port editor screen.

The configured tunnels provide you a **Details View** and **Troubleshoot View**. Click on the tunnel profile to view the below:

- **Details View:** Displays the configured parameters of the configured tunnel.
- **Troubleshoot View:** Displays the tunnel's statistics. Use the **Clear Stats** button to reset the statistics.

About Virtual Extensible LAN (VXLAN) Tunnels

VXLAN is a simple tunneling mechanism that allows overlaying a Layer 2 (L2) network over a Layer 3 (L3) underlay with the use of any IP routing protocol. It uses MAC Address-in-User Datagram Protocol (MAC-in-UDP) encapsulation. A remote device, such as the Gigamon cloud, GigaVUE TA Series, GigaVUE HC Series, or a customer-specific device, encapsulates the filtered traffic, adds an encapsulation header that consists of Layer 2 + IP + UDP + VXLAN headers. The encapsulated packet is sent out of the circuit port, which is connected to the public network (the Internet). This packet is routed in the public network to reach the main office site. The packet is ingressed at the circuit port configured in the GigaVUE-H Series or GigaVUE-TA Series device at the main office. After validation of the source port, destination port, and VXLAN Network Identifier (VNI) of the packet, the VXLAN tunnel header will be removed and the inner payload will be sent to the tools based on the map rules configured.

The following figure illustrates the VXLAN tunnel encapsulation and decapsulation.




Figure 29 *VXLAN Tunnel Encapsulation and Decapsulation*

In this diagram, traffic is tapped on a GigaVUE-TA200 device at a remote site, and then it is tunneled through VXLAN encapsulation across the network before it reaches the GigaVUE-HC3 device at the main office site, which is connected to the actual tools. The tunnel decapsulation is executed on an ingress circuit port (IP interface). After tunnel decapsulation, the packet is presented to the flow mapping module to filter based on map rule parameters.

Refer to the following sections for details about VXLAN tunnels:

- [VXLAN Tunnel Configuration—Rules and Notes](#)
- [Configure VXLAN Tunnel to Encapsulate Traffic](#)
- [Configure VXLAN Tunnel to Decapsulate Traffic](#)

VXLAN Tunnel Configuration—Rules and Notes

Keep in mind the following rules and notes when working with VXLAN tunnels:

- VXLAN tunnels are supported only on GigaVUE-HC1, GigaVUE-HC3, GigaVUE-HCT, GigaVUE-HC1-Plus, GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA400, and GigaVUE-TA400E devices.
- A maximum of 1500 VXLAN IDs are supported.
- Flow mapping that is configured on the circuit port used for VXLAN decapsulation will filter only the inner packet attributes along with VXLAN-ID. Any other non-tunneled packets that ingress on this circuit port will not be filtered or redirected to tool ports, even if it matches the rules configured on the map.
- IPv6 protocol is not supported with VXLAN tunnels. However, IPv6 protocol is supported with GigaSMART VXLAN tunnels.
- Ingress VLAN tagging and Tool Mirror features are not supported with VXLAN tunnels.
- Any encapsulated packet that exceeds the MTU value configured for the IP interface will be discarded because IP fragmentation and reassembly of packets are not supported.
- VXLAN tunnel encapsulation is not supported on circuit GigaStreams.
- VXLAN tunnel decapsulation is supported only on encapsulated packets that are not tagged. On GigaVUE-TA400, and GigaVUE-TA400E VXLAN tunnel decapsulation is supported on encapsulated packets that are both tagged and untagged.
- GigaSMART operations cannot be combined with VXLAN decapsulation in the same map.
- Map-passall is not supported for the circuit port that encapsulates or decapsulates the VXLAN packet.
- When a circuit port is configured for VXLAN tunnel decapsulation, you cannot use the port in any other regular map in which a network port is configured as the source port.
- Inner VLAN qualifier is not supported on the port in which the VXLAN tunnel decapsulation is enabled except on GigaVUE-TA400, and GigaVUE-TA400E.
- VXLAN ID qualifier is available as part of existing static templates. Following table provides details about the platforms for which the static templates are available:

Table 12: VXLAN ID Qualifier Availability

Template	Platform	
	GigaVUE-HC1	GigaVUE-HC3/GigaVUE-TA100/GigaVUE-TA200/ GigaVUE-TA200E/ GigaVUE-TA25/GigaVUE-TA25E/ GigaVUE-TA400/GigaVUE-TA400E
IPv4	Not available	Available
IPv6	Available	Available

Template	Platform	
	GigaVUE-HC1	GigaVUE-HC3/GigaVUE-TA100/GigaVUE-TA200/ GigaVUE-TA200E/ GigaVUE-TA25/GigaVUE-TA25E/ GigaVUE-TA400/GigaVUE-TA400E
IPv4+UDA	Not available	Available
IPv4+MAC	Available	Available
UDA	Available	Available

Configure VXLAN Tunnel to Encapsulate Traffic

Before creating a VXLAN tunnel to encapsulate traffic, refer to the [VXLAN Tunnel Configuration—Rules and Notes](#).

The following table summarizes the required tasks to configure a VXLAN tunnel for encapsulating the traffic:

S.No	Task	Refer to...
1.	Configure the required circuit port.	Configure Ports
2.	Configure an IP interface and attach the circuit port that you created in task 1.	Configure IP Interface
3.	Configure a VXLAN tunnel for encapsulation and attach the IP interface that you created in task 2. Ensure that you select the Type as VXLAN and the Mode as Encapsulation .	Create Tunnel

S.No	Task	Refer to...
4	Create a VXLAN group for a device and add all the VXLAN IDs that are specific to the device.	Create VXLAN / L2GRE Group
5.	Configure the VXLAN ID either at the chassis-level or at the port-level to help identify the encapsulation tunnel. You must use one of the VXLAN IDs that you have already added in the VXLAN group for the device in task 4.	Configure L2GRE / VXLAN Identifier
6.	Configure a map to encapsulate the VXLAN traffic. Ensure that you add the following details when configuring a map: • If you have configured the VXLAN ID for the chassis, you can configure any network port on the chassis as the source port of the map. • If you have configured the VXLAN ID for a network port, you must configure that network port as the source port of the map. • Configure the encapsulation tunnel that you created in task 3 as the destination port of the map. • Specify the required pass/drop rule in the map to filter the traffic based on inner packet attributes or VXLAN ID for the template configured.	Create a new map

Configure VXLAN Tunnel to Decapsulate Traffic

Before creating a VXLAN tunnel termination, refer to the .

The following table summarizes the required tasks to configure a VXLAN tunnel for decapsulating the traffic:

S.No	Task	Refer to...
1.	Configure the required circuit port.	Configure Ports
2.	Configure an IP interface and attach the circuit port that you created in task 1.	Configure IP Interface
3.	Configure a VXLAN tunnel for decapsulation and attach the IP interface that you created in task 2. Ensure that you select the Type as VXLAN and Mode as Decapsulation .	Create Tunnel

S.No	Task	Refer to...
4	Create a VXLAN group for a device and add all the VXLAN IDs that are specific to the device.	Create VXLAN / L2GRE Group
5.	Configure a map to decapsulate the traffic. Ensure that you add the IP interface that you created in task 2 as the source of the map and specify the required pass/drop rule in the map to filter the traffic based on inner packet attributes or VXLAN ID for the template configured.	Create a new map

Orchestrated Workflow Configuration of VXLAN Tunnels

The VXLAN tunnels can now be configured through the Orchestrated Configuration page. To configure VXLAN Tunnels follow the below steps:

- In GigaVUE-FM go to **Traffic>Physical> Orchestrated Flows> Tunnels**.
- Click on **New**.
- Select **Embedded VXLAN**. The New Tunnel configuration page appears. Enter the required information as described below:

Field	Description
Tunnel Name	The name of the tunnel. NOTE: Alias must not have spaces or any of these characters: * '?',,/,/%@.
Tunnel Description	The description of the tunnel endpoint.
Traffic Direction	Select the traffic direction of the tunnel. Choose DECAP for decapsulation or choose ENCAP for encapsulation. • If you choose Decap, select the IP Interface from the list. You can create one using the Create an IP interface option if you do not have an IP interface. Refer to IP Interfaces to learn more about creating an IP Interface. • If you choose Encap enter the Source Port. You can select from the drop-down list. To edit the Port type or to enable the port for Admin privileges use the Port Editor window. ◦ You can also add a new VXLAN ID by clicking on the ADD VXLAN ID. NOTE: The VXLAN ID configured at source port will take precedence over the VXLAN ID configured at chassis level.
Nodes	Select the Nodes on which you intend to create the tunnel. For VXLAN Encapsulation tunnel you will have to configure the VXLAN ID.

Field	Description
	- Click on Config VXLAN ID > ADD VXLAN ID. - Fill in the require parameters. - Click on ADD. - The VXLAN IDs will list in the Configuration page.
ID	For VXLAN Decapsulation tunnel you will have to select the VXLAN ID. To configure a new VXLAN ID do the following: - Click on ADD VXLAN ID. - Fill in the require parameters. - Click on ADD. NOTE: Configuring a VXLAN ID is not mandatory if the map rules are configured. Either configure the rules or a VXLAN ID.
Rules	Configure the map rules that needs to be adhered to while decapsulating or encapsulating the traffic. Click on Rules Editor to configure the rules. NOTE: Configuring rules are mandatory for Encapsulation tunnel, For a Decapsulation tunnel rules are not mandatory if a VXLAN ID is already configured.
Encapsulation IP interface	Enter the interface IP address of the node (Destination IP) for an encapsulation tunnel.
Remote Tunnel IP, L4 Port	Enter the IP address , L4 Port values for an encapsulation tunnel.
Destination Port	Select the destination port for the decapsulation tunnel. You can edit the port details from the Port editor screen.

The configured tunnels provide you a **Details View** and **Troubleshoot View**. Click on the tunnel profile to view the below:

- **Details View:** Displays the configured parameters of the configured tunnel.
- **Troubleshoot View:** Displays the tunnel's statistics. Use the **Clear Stats** button to reset the statistics.

Create Tunnel

You can create tunnels on the encapsulation side and/or decapsulation side. To create a tunnel:

1. From the device view, go to **Ports > Tunnels**.
2. Click **New**. The Tunnel configuration page appears.
3. In the **Alias** and **Description** fields, enter the name and description of the tunnel.
4. From the **Type** drop-down list, select one of the following options:
 - **Circuit-ID**—to create Circuit-ID tunnels.

- **VXLAN**—to create VXLAN tunnels.
 - **L2GRE**—to create L2GRE tunnels.
5. Select one of the following modes:
 - **Encapsulation**—Select this mode to send the encapsulated traffic to the destination node that resides in another cluster.
 - **Decapsulation**—Select this mode to decapsulate the traffic received from the source node that resides in another cluster.
 6. In the **Circuit-ID** field, enter the circuit ID used to encapsulate or decapsulate the traffic. You can enter multiple circuit-IDs when you create a circuit-ID tunnel for decapsulation.

NOTE: This field is available only when you select **Circuit-ID** in the **Type** drop-down list.

7. In the **Destination IP Address** field, enter the IPv4 address of the destination node.

NOTE: This field is available only when you select **VXLAN** or **L2GRE** in the **Type** field and the **Mode** as **Encapsulation**.

8. In the **L4 Source Port** field, enter the layer 4 source port number.

NOTE: This field is available only when you select **VXLAN** in the **Type** field and the **Mode** as **Encapsulation**.

9. From the **Attached entity** drop-down list, select the required entity based on the tunnel type you are creating:
 - If you are creating a Circuit-ID tunnel, this field is available only when you select the **Mode** as **Decapsulation**. Ensure that you attach the required circuit ports or circuit GigaStreams.
 - If you are creating a VXLAN tunnel or a L2GRE tunnel, ensure that you select the IP interface to which you have attached the circuit port.
10. Click **OK**.

The newly added tunnel appears in the Tunnels listing page.

Create VXLAN / L2GRE Group

You can create a VXLAN or L2GRE group for a GigaVUE HC Series or a GigaVUE TA Series device and add all the VXLAN IDs or the L2GRE IDs that are specific to the device. You can add a maximum of 1500 IDs per system. To create a VXLAN or L2GRE group:

1. From the device view, go to **Ports > Tunnels > VXLAN / L2GRE Groups**.
2. Click **New**. The VXLAN / L2GRE Group page appears.
3. In the **Alias** and **Description** fields, enter the name and description of the group.
4. Select either **VXLAN** or **L2GRE** based on the group that you want to create.
5. From the **Box ID** drop-down list, select the required box ID of the device for which you want to create the group.
6. Click **Add ID**, and then enter the VXLAN ID or L2GRE ID for the device.
7. Repeat step 6 to add multiple VXLAN IDs or L2GRE IDs for the device, and then click **OK**.

NOTE: Ensure that you create a VXLAN or L2GRE group and add the required IDs to the group before you assign the VXLAN or L2GRE IDs to the map you create for tunnel termination.

The ingress packets to the VXLAN tunnels must have a matching L4 destination port configured without which the packets will be discarded. To configure the **L4 Destination Port** :

1. Click **Actions** on the VXLAN / L2GRE Group page and select **Configure VXLAN Global**.
2. Select the required **Box ID**.
3. Enter the **L4 Destination Port** value. Default value is 4789.
4. Click **OK**.

NOTE: The L4 Destination port is configured globally on the chassis and thus affects all the VXLAN tunnels on the box.

Configure L2GRE / VXLAN Identifier

The L2GRE or VXLAN IDs that you have added in the L2GRE or VXLAN group created for a specific device must be configured either at the chassis-level or at the network port-level. The ID that you configure for the network port overrides the ID configured for the chassis. The L2GRE or VXLAN IDs help to identify the respective tunnels.

Configure L2GRE / VXLAN Identifier for a Chassis

To configure L2GRE or VXLAN ID for a chassis:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**. In the Physical Nodes page, select the node for which you want to configure the L2GRE or VXLAN ID.

- From the left navigation pane, go to **System > Chassis**.
- Go to List View, and then select the Box ID for which you want to configure the L2GRE or VXLAN ID.
- From the **Actions** drop-down list, select **Configure L2GRE/VXLAN ID**.
- 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.

- Click **OK**.

Configure L2GRE / VXLAN Identifier for a Network Port

To configure L2GRE or VXLAN ID for a network port:

- From the left navigation pane, go to **Inventory > Physical > Nodes**. In the Physical Nodes page, select the node for which you want to configure the L2GRE or VXLAN ID.
- Go to **System > Ports > Ports > All Ports**.
- 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.
- 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.

- Click **OK**.

View VXLAN / L2GRE ID Statistics

You can view the statistics such as the number of packets and bytes that were decapsulated using a VXLAN ID or a L2GRE ID for a device. To view the VXLAN /L2GRE ID Statistics page, go to **Ports > Tunnels > Statistics**. [Figure 30VXLAN / L2GRE ID Statistics](#) illustrates the statistics.

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

Figure 30 *VXLAN / L2GRE ID Statistics*

The following table describes the VXLAN or L2GRE ID statistics:

Statistic	Description
VXLAN/L2GRE ID	The VXLAN or L2GRE identifier that is specific to the device.
Box ID	The box identifier of the device.
VXLAN/L2GRE Group	The name of the group created for the device.
Type	The type of group—VXLAN or L2GRE.
Packets Rx	The number of packets that were decapsulated using the VXLAN or L2GRE ID.
Bytes Rx	The total bytes that were decapsulated using the VXLAN or L2GRE ID.

To export the statistics to a CSV file, click the '+' icon and then select **Download all data as CSV**.

Note: The bulk CSV download window does not generate the correct CSV file.

To clear the VXLAN statistics, click **Clear Stats > Clear All VXLAN Stats**. To clear the L2GRE statistics, click **Clear Stats > Clear All L2GRE Stats**.

Tunnel Monitoring

A tunnel is a logical entity configured to establish secure connections across networks. Traffic is encapsulated at the sending end of the tunnel and decapsulated at the receiving end of the tunnel. GigaVUE-FM supports configuring various tunnel types. Refer to the [Tunnels](#) section for details.

With the Tunnel Monitoring feature, GigaVUE-FM uses a tunnel discovery mechanism to identify the tunnels created across the managed nodes and clusters. The following tunnel types are supported in this release:

- Embedded tunnels.
 - Circuit-id tunnels
 - L2GRE tunnels
 - VXLAN tunnels
- GigaSMART tunnels.
 - GigaSMART L2GRE tunnel
 - GigaSMART VXLAN tunnel
 - GigaSMART TLS_PCAPNG Tunnel
- Mixed Tunnels: Tunnel type on encapsulation end and decapsulation end can be either Embedded tunnel or GigaSMART tunnel.
- Hybrid Tunnels: Egress Tunnels (L2GRE/VxLAN/TLS_PCAPNG) originating from GigaVUE V Series Nodes and terminating at GigaVUE HC Series or GigaVUE TA Series devices.
- Secure Tunnels.

The identified tunnels are logically grouped based on the following grouping criteria and listed in the Tunnel Monitoring page. The Tunnel Monitoring page provides a unified view of both ends of the tunnels created across the devices managed by GigaVUE-FM.

Tunnel Type	Grouping Criteria
Embedded Tunnels	
- Circuit-ID Tunnels	• Circuit-ID • Physical Connectivity discovered by GDP. Note: Physical connectivity is an optional parameter for grouping the circuit-id tunnels
- L2GRE/VXLAN Tunnels	• Destination IP • L2GRE-ID/VXLAN-ID

Tunnel Type	Grouping Criteria
GigaSMART Tunnels	
- L2GRE/VXLAN Tunnels	- Destination IP - L2GRE-ID/VXLAN-ID - Application ports (VXLAN Tunnel)
-TLS_PCAPNG Tunnels	- Destination IP - Application Ports
Hybrid Tunnels	
- L2GRE/VXLAN/ from GigaVUE V Series Nodes and terminating at GigaVUE HC Series and GigaVUE TA Series devices	- Destination IP - L2GRE-ID/VXLAN-ID
-TLS_PCAPNG Tunnel form GigaVUE V series Nodes	- Destination IP - Application Ports

During addition of nodes to GigaVUE-FM, the tunnels configured in the devices are automatically detected, logically grouped and added to the Tunnel Monitoring page. The tunnel logical group gets deleted during node deletion. GigaVUE-FM identifies the tunnel related configuration changes during the config sync cycle and accordingly updates the tunnel logical groups in the Tunnel Monitoring page.

NOTE: For the tunnels listed in the tunnel monitoring page, you can identify the traffic drop at the decapsulation end of the tunnel by creating Alert Policies. Refer to the [Traffic Drop Identification](#) section for details.

Rules, Notes, and Limitations

Refer to the following Rules, notes, and limitations:


- The tunnel discovery mechanism identifies tunnels and the tunnel applications that were configured at the devices that are being managed by GigaVUE-FM. The tunnels are identified irrespective of the device software version.
- GigaVUE-FM identifies tunnels configured between the physical devices. GigaVUE-FM also identifies tunnels configured between GigaVUE V Series Nodes and physical devices. Tunnels deployed in VMware (ESXi and NSX-T) are only supported.
- The following tunnel types are not supported in this release:
 - Tunnels configured between the GigaVUE V Series Nodes.
 - Tunnels configured between GigaVUE V Series Nodes and UCT-Vs
- GigaVUE-FM detects tunnels that use 'Regular' maps for encapsulation and decapsulation.

- Circuit ID tunnels that are created manually will be discovered and monitored. Circuit ID tunnels created as part of Fabric Map configurations are not within the scope of the Tunnel Monitoring page.
- For embedded tunnels (L2GRE/VXLAN), GigaVUE-FM detects the decapsulation configuration only when pass rule is configured with L2GRE/VXLAN Id on the maps at the decapsulation end.

View Tunnel Monitoring

To access the tunnel monitoring page:

1. Go to **Traffic** 
2. Go to **Physical > Tunnel Monitoring**.



The screenshot displays the Tunnel Monitoring interface. At the top, there's a header with 'Tunnel Monitoring' and search, refresh, and user icons. Below the header is a table listing tunnels. The table has columns for Tunnel Logical Group, Tunnel ID, Source ID, Destination ID, Tunnel Type, Health Status, Tunnel Discovery State, Monitoring Status, and Alert Policy Created. The table shows several tunnels, all with 'Complete' discovery state and 'Active' monitoring status. Below the table, there's a pagination bar indicating '77 tunnel groups total'. A detailed view for a specific tunnel is shown below the table, with tabs for 'Source/Destination' and 'Statistics'. The 'Source/Destination' tab is active, showing details for Source ID: 7777333 and Destination ID: 7777333. The Source ID details include Encapsulation Details (Circuit ID: 501, Tunnel Alias: TCE, Map(s): TCE-L2CIRCUIT-MAP). The Destination ID details include Decapsulation Details (Circuit ID: 501, Tunnel Alias: TCE-d, Map(s): TCE-L2CIRCUIT-decap-map). The page footer shows 'FM Instance: GigaVUE-FM' and 'Last Updated At: Nov 15, 2022 12:54:46'.

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)

Field	Description
	- 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

Button	Description
	- Source ID - Destination ID
Export	Use to export the tunnels (either all or the selected tunnels) in CSV or XLSX format.

Packet Capture (PCAP)

Starting from software version 5.16.00, you can configure Packet Capture (PCAP) from GigaVUE-FM. Both GigaVUE-FM and the devices must be running software version 5.16.00 and greater. Use the PCAP feature to analyze the network traffic and to troubleshoot any performance issues.

GigaVUE-FM allows you to configure packet capture at the ingress port or egress port or both. The port must be a physical port. The port type used for packet capture can be network, tool, hybrid, inline tool, or inline network port. Packet capture is not supported on GigaSMART ports or back plane ports.

NOTE: PCAP feature is enabled by default. To disable or re-enable PCAP, contact Gigamon customer support. Once disabled, the corresponding PCAP configurations will not work.

Supported Devices

Packet capture functionality is supported on the following devices:

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

You can configure PCAP on both standalone nodes as well as on nodes that belong to a cluster. For non-leader ports, PCAP can be configured only from the leader node.

To configure packet capture, you must define filters to capture specific traffic based on rules. You can specify the following criteria in the rules:

Criteria	Description
Source MAC address	The source and destination MAC address.
Destination MAC address	
VLAN ID	VLAN ID value.
Inner-VLAN	Inner VLAN ID value.
Layer 2 ethernet type	Layer 2 Ethernet type value.
Source IPv4 address	The source and destination IPv4 address. You can also specify a wild card with an IP mask.
Destination IPv4 address	
Internet protocol	Valid Internet protocol.
IP version number	IP version for traffic, either IPv4 or IPv6.
IP fragmentation bits	Match IP fragments.
Time to Live (TTL) value	Time to Live (TTL—IPv4) or Hop Limit (IPv6) value in an IP packet.
DiffServ Code Point (DSCP) bits	Decimal DSCP value.
Layer 4 destination port number	Layer 4 destination port number.
Layer 4 source port number	Layer 4 source port number.
TCP flags	TCP flags to indicate the state of connection.

You can specify the criteria in any combination. Packets matching the defined criteria are captured and saved as pcap files.

Refer to the following sections for details:

- [Rules, Notes, and Limitations](#)
- [Packet Capture \(PCAP\)](#)
- [Configure PCAP Profile](#)
- [View PCAP Files](#)
- [Configure PCAP over IP Interface](#)
- [View PCAP over IP Interface](#)
- [Delete PCAP over IP Interface](#)

Rules, Notes, and Limitations

Refer to the following rules and notes:

- The PCAP feature supports up to 16 capture files per device. Each capture file can have up to 40000 packets. A capture file is maintained per PCAP session. Each session can have up to 64 filter rules per direction. Each capture file can be viewed, deleted, or uploaded out of the device for offline use.
- The PCAP feature supports up to 16 active capture sessions at a time per port on all GigaVUE platforms except GigaVUE-TA400, and GigaVUE-TA400E. The GigaVUE-TA400, and GigaVUE-TA400E platform currently supports one active PCAP session per port at a time.
- The PCAP configuration doesn't persist across node reboots or upgrades.
- The PCAP feature is not supported on stack ports in legacy stacking mode.
- The PCAP feature does not support 'vlan' and 'inner-vlan' filter rules on a tool, hybrid or circuit port in the 'tx' direction.
- The PCAP feature on tool ports does not capture the vlan tag specified with the ingress-vlan-tag feature. To overcome this, redirect the traffic to another hybrid port along with other tool ports and capture the packets on the hybrid port ingress.
- The PCAP feature on GigaVUE-HC1-Plus, GigaVUE-HC1, GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA400, and GigaVUE-TA400E platforms contain extra VLAN header added in capture files. Untagged packet captures contain vlan-tag 1 header added and tagged packet captures contain an outer tag duplicated.
- The PCAP feature on port discovery protocols (LLDP/CDP/GDP) enabled ports will not capture the discovery protocol control packet in the PCAP file.
- The PCAP feature may miss some packets in the capture file depending on the rate of traffic being captured.
- When the PCAP feature is enabled for both directions to capture TX and RX, bi-directional traffic must be sent.
- The PCAP feature supports a maximum of four active capture sessions when the direction is configured as TX or Both.
- GigaVUE-OS cannot capture ARP/ICMP control packets on an IP Interface.

- When PCAP is configured to capture TX direction traffic on an IP interface for Embedded Tunnels, only the actual packets are visible, not the L2GRE/VxLAN header. This is because the packet is mirrored to the PCAP interface before encapsulation occurs. This does not affect packet capture for GigaSMART tunnels, where the actual packets with headers can be seen.
- Q-in-Q packets cannot be captured on the egress port when any filter rule is applied.

Configure PCAP Profile

To configure PCAP through GigaVUE-FM:

1. From the device view, go to **Ports > Ports > All Ports**.
2. Select the required port/ports for which you need to configure PCAP.

NOTE: You can configure PCAP only for a maximum of four ports at a time.

3. Click **Action** and select **Configure PCAP**.

The **Action** button is disabled:

- If you select more than four ports.
- If you do not select any port.
- If you select GigaSMART Engine ports or other unsupported port types.


If the PCAP feature is disabled by the customer support team, a banner notification is displayed.

The **Action** button is hidden:

- For G-TAP devices.
- For devices running software version less than 5.16.00 and managed by GigaVUE-

FM.

4. Select or enter the following details:



Field	Description
Alias	Name of the packet capture filter
Direction	The direction of traffic. Can be: • Rx • Tx • Both
Channel Port	The channel port identifier for the packet capture filter. The channel port is any unused port that does not have any map configuration. The channel port must be on the same node as the capture port. The channel port must be administratively enabled and must remain enabled while a packet capture filter is configured. You must specify one channel port for each transmitted or both direction. channel port is not needed for received direction. Note: If a PCAP configuration is deleted, the channel ports configured in the PCAP will go down.
Packet Limit	The number of packets to capture. The valid range is 1 to 40000 for all

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. From the left navigation pane, go to **Inventory > Physical > Nodes**.
2. Select a cluster ID, and then from the left navigation pane, go to **Support > Debug > PCAP**.



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

Configure PCAP on IP Interface


The PCAP over IP Interface feature extends GigaVUE-OS packet capture capabilities to ports associated with IP interfaces, enabling capture of traffic received from and transmitted to those interfaces for deep analysis and troubleshooting. It preserves existing PCAP behavior and performance characteristics while ensuring that critical control traffic such as ARP and ICMP remains unaffected, and is supported across all GigaVUE HC Series and GigaVUE TA Series devices.

To configure PCAP IP Interface, do the following:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link where you want to create the packet capture. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > Ports > IP Interfaces**. The IP Interfaces page appears.
4. Select the alias checkbox, that you want to configure the PCAP IP Interface.

NOTE: A maximum of four alias checkbox can be selected to configure the PCAP IP Interface.

5. Click **Troubleshoot**.



6. Select **Configure PCAP** from the list.

- Alias - Enter the name for the alias.
- Direction - (Optional) Select either Rx, Tx, or Both.

NOTE: If you select Tx or Both directions, then you need to select the channel port from the list.

- Channel Port - Select the channel port from the list.
- Packet Limit - (Optional) Enter the value from 1 to 40000. If no value is provided, by default it will take as 40000.
- Rules - (Optional) You can add a rule. Click Add a Rule, and select Condition from the list, and provide the values.

7. Click **Save**.

View PCAP on IP Interface

To view PCAP IP Interface, do the following:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link where you want to view the packet capture. The **Overview** page of the selected node appears.

3. In the left menu, go to **System > Ports > IP Interfaces**. The IP Interfaces page appears.
4. Select the alias checkbox, that you want to view the PCAP IP Interface. The view PCAP configuration page appears.
5. Click **Close**.

Delete PCAP on IP Interface

To delete PCAP IP Interface, do the following:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link where you want to delete the packet capture. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > Ports > IP Interfaces**. The IP Interfaces page appears.
4. Select the alias checkbox, that you want to delete the PCAP IP Interface. The delete PCAP configuration page appears.
5. Select the checkbox on the right-side to delete the PCAP over IP Interface.
6. Click **Delete**.

Flow Mapping®

This chapter provides the following information about flow mapping:

- [About Flow Mapping®](#)
- [Manage Maps](#)
- [Flow Mapping® FAQ](#)
- [Configure Active Visibility](#)

About Flow Mapping®

This section describes what is Flow Mapping® and how to apply it to GigaVUE® HC Series and TA Series nodes. Refer to the following sections for details:

- [Flow Mapping® Overview](#)
- [Get Started with Flow Mapping®](#)
- [Flow Map Syntax and Construction](#)
- [Work with Map-Passalls and Port Mirroring in GigaVUE-FM](#)
- [Port Access and Map Sharing](#)
- [Map Examples](#)

Flow Mapping® Overview

Flow Mapping® is the technology found in GigaVUE nodes that takes line-rate traffic at 1Gb, 10Gb, 40Gb, or 100Gb from a network TAP or a SPAN/mirror port (physical or virtual) and sends it through a set of user-defined map rules to the tools and applications that secure, monitor, and analyze IT infrastructure. Flow Mapping® provides **superior granularity and scalability** above and beyond the capabilities of connection and ACL filter based technologies by addressing the problems inherent when going beyond small numbers of connections or when more than one traffic distribution rule is required.

Flow Mapping® can granularly filter and forward traffic to specific monitoring tools through thousands of map rules with criteria based on over 30 predefined Layer 2, Layer 3, and Layer 4 parameters including IPv4/IPv6 addresses, application port numbers, VLAN IDs, MAC addresses and more. Users can also define custom rules that match specific bit sequences in the traffic streams, applying Flow Mapping® to tunneled traffic, specialized applications and even higher-layer protocols.

In addition, IT operations can deploy Visibility as a Service (**VaaS**), taking advantage of the Flow Mapping® **role based access control (RBAC)** features on GigaVUE nodes. Users can be given access to traffic based on their needs without interfering with the monitoring operations of the other teams. This helps protect compliance and privacy protocols and allows teams to dynamically react when needed for increased efficiency.

Flow Mapping® can be combined with **GigaSMART** technology to provide packet modification and intelligent capabilities like de-duplication and packet slicing, making tools more efficient by reducing the number and size of packets they have to store and process. Header stripping and de-tunneling functions provide tools access to protocols and data they would otherwise be blind to.

When multiple GigaVUE nodes are in a **stacked or clustered configuration**, Flow Mapping® enables traffic to be sent from any network port to any tool port, expanding visibility beyond a single rack, row, or data center. GigaSMART can be leveraged on all traffic flow, accepting traffic from any network port and regardless of where the GigaSMART hardware is located within the stack or cluster.

To access Flow Maps

1. Go to **Traffic**  and select **Physical > Flow Maps**. This displays the list of Devices/Cluster Nodes managed by this instance of GigaVUE-FM.
2. Click the Cluster ID of any node to open the node. The following options are displayed on left navigation pane:

- Maps
- Workflows
- Map Templates
- Filter Templates

Get Started with Flow Mapping®

You can manage packet distribution in both the GigaVUE-OS command-line interface (CLI) and GUI (GigaVUE-FM). Both interfaces allow you to perform all packet distribution tasks:

- designating ports as network or tool ports
- setting up map rules
- mapping network ports to tool ports
- many other such functions

NOTE: These tasks can also be performed with the GigaVUE-FM APIs. For more information about the APIs, refer to *GigaVUE-FM REST API Reference in GigaVUE-FM User's Guide*.

For the setup of a few simple maps, refer to the following sections for examples:

- [Example: How to Create a Simple Map](#)
- [Example: How to Handle Overlaps when Sending VLANs and Subnets to Different Tools](#)

Check Status of Nodes and Ports

Before configuring maps, check the status of line cards, modules, and ports. To view the status of line cards and modules, select **Chassis** from navigation pane and select Table view. For more information about the Chassis page, refer to the “*Chassis*” section in the *GigaVUE-OS and GigaVUE-FM Administration Guide*. Go to **System > Ports > Ports > All Ports** to see a table with details about each port. For more information about ports, refer to [About Ports](#) and [Managing Ports](#).

Designated Port Types

Ports on GigaVUE-OS nodes can be one of the following types:

- network
- tool
- stack
- inline-network

- inline-tool
- hybrid
- circuit

NOTE: Not all port types are supported on all platforms. Inline-network and inline-tool port types are only supported on GigaVUE® HC Series nodes, GigaVUE-TA25, GigaVUE-TA200, GigaVUE-TA400, and GigaVUE-TA400E.

About Shared Collectors

GigaVUE nodes let you create map rules that direct traffic on any network port or ports to any tool ports. Traffic can be dropped intentionally using the drop rule or any packets that do not match any other rule in the map can be sent to the collector. Shared collectors are set up to capture any packets that do not fulfill the map criteria but may be required by other tools.

NOTE: If a shared-collector destination for a set of network ports is not defined, non-matching traffic is silently discarded.

When assigning the priorities to map rules on GigaVUE® HC Series nodes, GigaVUE-TA25 and GigaVUE-TA200, the first rule setup will also have the highest priority unless specified by the user. The shared collector rule is the only exception because it will always have the lowest priority even if configured first. This means that an incoming packet will be matched against all the rules in the same map and when not matched with any rules, it be forwarded to the designated tool port for the collector.

A GigaStream or multiple sets of GigaStream can also be set as destination for a collector port by using the GigaStream alias.

In cases, where multiple network ports are sharing the multiple maps, packets that do not 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.

Map Info

Map Alias *

Description

Enable ☒

Type Regular

Subtype * By Rule

No Rule Matching ☐ Pass Traffic

Map Source and Destination

Port Editor

Source * N 4/1/x1 N 4/1/x2

Destination T 5/1/x6 T 5/1/x23 Tool Finder

Encapsulation Tunnel None

GigaSMART Operations (GSOP) None

Map Configuration & Rules

Configurations

Address Rewrite ☐ Apply to All Traffic

Map Rules

Quick Editor Import Add a Rule

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 1 Shared 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 where you want to send any packets not matching the normal flow maps.

Alias	Map Status	Source	Destination	Encapsulat...	Description	Enabled	Type	Subtype	Number of ...	GSOP	Priority	Access Level	Tags	Tx Cluster ...	Rx ...
newma...	Map is ...	2 Ports	2 Ports			true	Regular	By Rule	1		1	admin		1 Stack Po...	1 Stack Po...
		N 4/1/1	T 5/1/6											S 4/1/9	S 5/1/8
		N 4/1/2	T 5/1/23												
newma...	Map is ...	1 Port	2 Ports			true	Regular	By Rule	1		1	admin		1 Stack Po...	1 Stack Po...
		N 4/1/4	T 5/1/24											S 4/1/9	S 5/1/8
			T 5/1/23												

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 in GigaVUE-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 in GigaVUE-FM and the CLI, the two features are identical.

Define Map Source Port Lists

You can configure more than one map with the same source ports in the **Source** field in a map.

The source port list of one map must be exactly the same as the source port list of another map (have the same ports as well as the same number of ports) and it must not overlap with the source port list of any other map.

For example if map1 has **Source** ports 1/1/x4, 1/1/x5, and 1/1/x6 already configured:

- map2 **Source** ports 1/1/x5, and 1/1/x6 cannot be configured because ports x5 and x6 are configured already in map1 as source ports.
- map3 **Source** ports 1/1/x3, 1/1/x4, 1/1/x5 cannot be configured because ports x4 and x5 are configured already in map1 as source ports.

Share Network Ports Between Maps

Network ports can be shared between a regular map and a map passall, as follows:

- the regular map has network ports 1/1/x3..x4
- the passall map has network port 1/1/x4

When there are overlapping network ports and shared tool ports between a regular map and passall map, the map passall tool ports or GigaStream will receive traffic from the network ports configured on the regular map, in addition to the traffic from its own network port or ports.

In the configuration above, the map passall will also receive traffic from 1/1/x3.

NOTE: A shared collector map and a regular map should have exactly the same set of network ports. (This is the correct use case.) Overlapping network ports should not be configured for collector maps. (This is an incorrect use case whether the overlapping ports are a subset or a superset of the network ports.) Network ports cannot be shared between a regular map and a shared collector map.

Share Tool Ports Between Maps

When a map passall and shared collector share the same tool ports, removing the shared tool ports from the passall map may affect the shared collector traffic. The workaround is to not share the same tool ports between a map passall and a shared collector.

Map Priority

Packets matching multiple maps in a configuration are sent to the map with the highest priority when the network ports are shared among multiple maps with pass-by map rules. By default, the first map configured has the highest priority; however, you can adjust this.

In GigaVUE-FM, the UI displays the maps from highest priority to lowest as top to bottom.

Maps sharing the same source port list are grouped together for the purpose of prioritizing their rules. Traffic is subjected to the rules of the highest priority map first and then the rules of the next highest priority map and so on. Within a map, drop rules are applied first and then pass rules, in other words, drop rules always have higher priority than pass rules. Currently when a map's source port list is defined the map is grouped/prioritized with other maps sharing the same source port list. Newly configured maps are added as the lowest priority map within the group when initially configured unless changed by the user.

NOTE: Shared collector will always go to the lowest priority when setting up maps.

Adjust Map Priority in GigaVUE-FM

Before you get started adjusting map priority, start by reviewing the current map priorities in place by opening the Maps page and viewing the priority of the maps in the Priority field. For example, [Figure 2Map Priorities](#) shows three maps MyMap1, MyMap2, and MyMap3 with the same source port 1/1/x8. The Priority column in the table shows the current priority of each map.

☐	Alias	Comments	Type	Sub Type	Source	No of Rules	GSOP	Priority	Access Level	Destination
☐	GTP-Sampling-2		secondLevel	flowSample	vport_elias_test	1	gtp_flow_sampling_elias_test	1	admin	1/1/x4
☐	MyMap1		regular	byRule	N 1/1/x8	0		1	admin	1/1/x7
☐	MyMap2		regular	byRule	N 1/1/x8	0		2	admin	1/1/x10
☐	MyMap3		regular	byRule	N 1/1/x8	0		3	admin	1/1/x15
☐	OpenStack_vTraffic_toWireshark		regular	byRule	N vTunnelEndpointForOpenStack	1	GigavueVM_Tunnel	1	admin	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 Filtering (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

Type	Subtype	Map Rule	GSOP	Source	Destination
					inline-serial tools
First Level Note: The Gen3 GigaSMART cards support a maximum of 230 first-level maps per gsgroup and vport, when gsop is associated with flow sample or flow whitelist applications.	By Rule	Rule	no	network ports, or hybrid ports	virtual ports, tool ports, or hybrid ports, or GigaStream, (collector is not allowed)
Transit Level	By Rule	gsrule	yes	virtual ports	virtual ports
Second Level	By Rule, or Collector	gsrule	yes	virtual ports	tool ports, or hybrid ports, or GigaStream, or port group
	Flow Filter	flowrule	yes		
	Flow Sample	flowsample	yes		
	Flow Sample Overlap	flowsample	yes		
	flowSample-sip	flowsample	yes		
	Flow Whitelist	whitelist	yes		
	Flow Whitelist Overlap		yes		
flexinline	byRule collector	ordered list of inline-tool or inline- tool-group in a-to-b or b-to-a direction or both	rule	inline-network alias	inline-tool pairs, or inline-tool-groups, or bypass
			-		

Null Port in Maps

A map can be configured with only the source port without specifying any destination port (null-port in command-line reference). A map is still created, but the traffic is dropped.

If a null-port is specified in map, the GigaSMART operation is performed on the traffic, and the traffic is finally dropped. After dropping the traffic, the GigaSMART packet drop counter is incremented. If GigaSMART operations are chained, then after applying all the GigaSMART operations on traffic, the traffic is dropped.

NOTE: This feature is supported on GigaVUE HC series devices on both stand-alone nodes and clusters. Null port is applicable for regular maps and second-level maps involving GigaSMART operations.

Map Subtypes

The map subtype describe in [Matrix of Map Types](#) is optional. It specifies the following:

- **By Rule**—Specifies a rule-based map subtype, which is supported on the following map types:
 - **First Level, inline, and Regular** map types.
 - **Transit Level** map type.
 - **Second Level** map type.
- **Pass All**—Specifies a passall map subtype, which applies to **regular** and **inline** map types. The **Pass All** subtype is not supported on **First Level** and **Second Level** map types. With this subtype, map priority cannot be configured or modified.
- **Collector**—Specifies a collector map subtype, which applies to **Regular, inline, and Second Level** map types. The **Collector** subtype is not supported on the **First Level** map type. With this subtype, map priority cannot be configured or modified.
- **Flow Filter**—Specifies a flow filtering map subtype, which applies to **Second Level** map types. Specify the **Flow Filter** map subtype when using a **Flow Filter** parameter.
- **Flow Sample**—Specifies a flow sampling map subtype, which applies to **Second Level** map types. Specify the **Flow Sample** map subtype when using a **Flow Sample** parameter.
- **Flow Sample Overlap**—Specifies a flow sampling overlap map subtype, which applies to **secondLevel** map types. Specify the **flowSample-ol** map subtype when using a **flowsample** rule.
- **Flow Sample sip**—Specifies a SIP flow sampling map subtype, which applies to **secondLevel** map types.
- **Flow Sample Overlap**—Specifies a flow sampling overlap map subtype, which applies to **Second Level** map types. Specify the **Flow Sample Overlap** map subtype when using a **flowsample** rule.
- **Flow Whitelist**—Specifies a whitelist map subtype, which applies to **Second Level** map types. Specify the **Flow Whitelist** map subtype when using a whitelist rule.
- **Flow Whitelist Overlap**—Specifies a whitelist overlap map subtype, which applies to **Second Level** map types, Specify **Flow Whitelist Overlap** when using a whitelist rule.
- **Flow Whitelist-sip**—Specifies a SIP flow whitelist map subtype, which applies to **secondLevel** map types.

The default map subtype is **By Rule**.

NOTE: Maps with subtype **Flow Sample Overlap** or **Flow Whitelist Overlap** do not support map editing.

If a By Rule map which has overlapping network and tool ports with pass-all map is disabled, traffic will still be forwarded to tool ports. Re-configure the maps to avoid traffic loss.

Map Type and Subtype Modification

Once a map is created, the map **Type** and **Subtype** cannot be modified. However, you can delete the map and recreate it with a different **Type** and **Subtype**.

Backwards Compatibility

For backwards compatibility, the map **type** parameter does not have to be configured. The **type** and **subtype** will be determined by the system based on the remainder of the map configuration parameters. If not enough information is available, the default values of **regular** and **byRule** will be assumed for the type and subtype.

Minimum Requirements for Map Creation

A map must be configured with at least a **from** parameter. Even if other parameters such as **to**, **rule**, or **use** are configured, without **from**, the map will not be created.

Map Rules

This section provides information about the different types of map rules that you can specify when creating maps in GigaVUE-FM. The following topics are covered:

- [Other Types of Map Rules for GigaSMART Operations](#)
- [IPv4/IPv6 and Map Rules](#)
- [Set Map Rules for TCP Control Bits](#)
- [How to Use Bit Count Netmasks](#)
- [How to Combine Rules and Rule Logic](#)
- [How to Mix Pass and Drop Rules](#)
- [Configuring Port Criteria and Bi-Directional Rules in By Rule Maps](#)
- [Work with User-Defined Pattern Match Rules](#)
- [Inner Header and MPLS Header Filtering](#)

Other Types of Map Rules for GigaSMART Operations

There are other types of rules for GigaSMART operations as follows:

- Adaptive Packet Filtering (APF) and Adaptive Session Filtering (ASF). For details, refer to [GigaSMART Adaptive Packet Filtering \(APF\)](#) and [GigaSMART Application Session Filtering \(ASF\)](#) and Buffer ASF.

- GTP Correlation. For details, refer to [GigaSMART GTP 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. 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.	When used with the <1-byte-hex> argument, matches against the Next Header field in the standard IPv6 header.
IPv4 TTL	Matches against the standard TTL (time-to-live) field in the IPv4 header. Note: These fields perform essentially the same service in both versions, specifying how long a datagram can exist.	Matches against the standard Hop Limit field in the IPv6 header.

NOTE: The **IP Version** argument is implicitly set to **4** – if you configure a map rule without **IP Version** specified, the GigaVUE® HC Series node assumes that the IP version is 4.

Protocol Map Rules and IPv6

The predefined protocol map-rules available for IPv4 (GRE, RSVP, and so on) are not allowed when **IP Version** is set to **6**. This is because with the next header approach used by IPv6, the next layer of protocol data is not always at a fixed offset as it is in IPv4.

To address this, the GigaVUE HC Series node provides the <1-byte-hex> option to match against the standard hex values for these protocols in the Next Header field. The standard 1-byte-hex values for both IPv4 and IPv6 are set by selecting the option from the **Value** list or specifying a decimal value when **IPv4 Protocol** is selected. The following are options or values that you can select from the **Value** list.

- IPv6Hop 0
- ICMP_IPv4 1
- IGMP 2
- IPv4 4
- TCP 6
- UDP 17
- IPv6 41
- RSVP 46
- GRE 47
- ICMP_IPv6 58

You can also specify a custom value (0-255) by selecting **Custom**. The following are some additional values and their meanings:

- 43: Routing Option (v6 only)
- 44: Fragment (v6 only)
- 50: Encapsulation Security Payload (ESP) Header
- 51: Authentication (v6 only)
- 59: No Next Header (v6 only)
- 60: Destination Option (v6 only)

Set Map Rules for TCP Control Bits

Select the **TCP Control** to set map rules matching one-byte patterns for the standard TCP control bits. The following table summarizes the bit positions of each of the flags, along with their corresponding hexadecimal patterns.

NOTE: Rules using the **TCP Control** must also include **IPv4 Protocol** and the Value set to **TCP 6**.

Flag	Bit Position	Pattern	TCP Control Mask
Congestion Window Reduced	X... ..	0x80	0x3f
ECN Echo	.X.. ..	0x40	0x3f
Urgent Pointer	..X.	0x20	0x3f
Acknowledgment	...X	0x10	0x3f
Push	 X...	0x08	0x3f
Reset	 X..	0x04	0x3f
SYN	X.	0x02	0x3f
FIN	X	0x01	0x3f

Examples

The map rule shown in [Figure 3Map Rule with SYN Bit Set](#) matches packets with only the SYN bit set:

Quick Editor
Import
Add a Rule

✕ Rule1

Condition
☒ Pass
☐ Drop
☐ Bi-directional

Rule Description
Description

Tags

TagKey
Values

TCP Control
✕

Value 82
Mask 3f

✕ Rule2

Condition
☒ Pass
☐ Drop
☐ Bi-directional

Rule Description
Description

Tags

TagKey
Values

Protocol
✕

Value
 IPv4
4

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 4](#) **Map Rule Matching All Packets with Both ACK and SYN Bits set** matches all packets with both the ACK and SYN bits set.

✕ Rule1

Rule Description

Description

Address Rewrite

Select

Condition

☒ Pass ☐ Drop ☐ Bi-directional

Tags

TagKey

Values

+

−

Protocol

✕

Value

TCP

6

✕ Rule2

Rule Description

Description

Address Rewrite

Select

Condition

☒ Pass ☐ Drop ☐ Bi-directional

Select a Rule

Tags

TagKey

Value

+

−

TCP Control

✕

Value

12

Mask

3f

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.

Quick Editor Import Add a Rule

✕ Rule 1
Condition ☒ Pass ☐ Drop ☐ Bi-directional

Rule Description Description

Tags

TagKey Values + -

Port Source ✕
Min 82 Max 0 to 65535
Subset

✕ Rule 2
Condition ☐ Pass ☒ Drop ☐ Bi-directional

Rule Description Description

Tags

TagKey Values + -

Protocol ✕
Value ICMP_IPv4 1

Figure 5 Pass and Drop Rules in a Map

Drop Rules Have Precedence!

Keep in mind that within a map, drop rules have precedence over pass rules. So, if a packet matches both a pass and a drop rule in the same map, the packet is dropped rather than passed.

Configuring Port Criteria and Bi-Directional Rules in By Rule Maps

When configuring a By Rule map in the device, it's essential to understand how to accurately set port criteria and leverage the Bi-Directional option to ensure comprehensive traffic matching.

Defining Port Ranges and Single Ports

Some rule conditions, such as Port Destination or Port Source, require you to input Minimum and Maximum values to define port ranges. Here's how to use these fields effectively:

Single Port Matching:

To match a specific port (e.g., port 80), enter the same value in both the Min and Max fields.

Example:

Min: 80

Max: 80

Port Range Matching:

To match a range of ports (e.g., 1000 through 2000), enter the lower value in Min field and the higher value in Max field.

Example:

Min: 1000

Max: 2000

This method ensures precision in traffic filtering and is applicable to all rule types using Min/Max configurations.

Using the Bi-Directional Option

The Bi-Directional checkbox significantly simplifies rule setup by enabling mirrored traffic matching. Typically, you must choose between Source or Destination for conditions like ports or IPs. Enabling Bi-Directional creates a reverse match automatically.

Example:

If a rule matches traffic with destination port 80, enabling Bi-Directional will also match traffic where port 80 is the source. This eliminates the need to manually define both directions.

Work with User-Defined Pattern Match Rules

GigaVUE-OS lets you create pass and drop map rules with *pattern matches* to search for a particular sequence of bits at a specific offset in a packet. You can configure up to two user-defined, 16-byte **pattern matches** in a map rule. A **pattern** is a particular sequence of bits at a specific location in a frame.

NOTE: Refer to [User-Defined Pattern Match Examples](#) for step-by-step instructions on creating a real-world pattern-match map rule.

User-defined pattern matches consist of the following components:

Step	Description
Pattern	Use the UDA1 Value and UDA2 Value fields for map rule to set up the actual bit patterns you want to search for. Refer to User-Defined Pattern Match Examples for details.
Mask	Use the UDA1 Mask and UDA2 Mask fields for map rules to specify which bits in the pattern must match to satisfy the map rule.
Offset	Use the UDA1 Offset and UDA2 Offset fields for map rules to specify where in the packet bits

Step	Description
	must match. Note: The GigaVUE® HC Series node accepts a maximum of two offsets per device. When both of the available offsets for the device are in use with existing map rules, you will not be able to add a new rule with a different value for UDAx Offset until at least one of the UDAx Offset is freed up from all existing map rules.

User-Defined Pattern Match Syntax

The user-defined pattern match syntax is as follows:

- Both the **UDAx Value** and **UDAx Mask** arguments are specified as 16-byte hexadecimal sequences. Specify the pattern in four 4-byte segments separated by hyphens. For example:
0x01234567-89abcdef-01234567-89abcdef
- Masks specify which bits in the pattern must match. The mask lets you set certain bits in the pattern as wild cards – any values in the masked bit positions will be accepted.
 - Bits masked with binary 1s must match the specified pattern.
 - Bits masked with binary 0s are ignored.
- You can set up the two global offsets allowed per device at 4-byte boundaries beginning at frame offset 2 and ending at offset 110. The resulting data range for pattern matches is from byte 3 through byte 126.
 - Multiple offsets must be set either equal to one another, or set beyond the boundaries of each other. For example, if **UDA1 Offset** starts at byte 2, the **UDA2 Offset** can only start either at byte 2 or at any point beginning with byte 18 (which would be the next 4-byte boundary after the 16-byte pattern used at **UDA1 Offset**).
 - Offsets are always frame-relative, not data-relative.
 - In many cases, you will be looking for patterns that do not start exactly on a four-byte boundary. To search in these position, you would set an offset at the nearest four-byte boundary and adjust the pattern and mask accordingly.

User-Defined Pattern Match Rules

Keep in mind the following rules when creating user-defined pattern matches:

- Offsets are specified in decimal; patterns and masks are specified in hexadecimal.
- All hexadecimal values must be fully defined, including leading zeroes. For example, to specify 0xff as a 16-byte value, you must enter 00000000-00000000-00000000-000000ff.
- User-defined pattern-match criteria are not allowed in egress port-filters (tool, hybrid, circuit, and inline network).
- You can use user-defined pattern matches as either standalone map rules or in tandem with the other available predefined criteria for map rules (for example, port numbers, IP addresses, VLAN IDs, and so on).

- You can use up to two separate user-defined pattern matches in a single map rule. When two user-defined pattern matches appear in the same map rule, they are joined with a logical AND. However, the two patterns cannot use the same offset.
- User-defined pattern matches are combined in map rules using the same logic described in [How to Combine Rules and Rule Logic](#).
- Avoid using user-defined pattern matches to set map rules for elements that are available as predefined criteria (for example, IP addresses, MAC addresses, and so on).
- GigaVUE HC Series nodes accept a maximum of two offsets per line card. When both of the available offsets for a line card are in use with existing map rules, you will not be able to add a new rule with a different value for **UDAx Offset** until at least one of the **UDAx Offsets** is freed up from all existing map rules.
- On GigaVUE-HC1-Plus, GigaVUE-HCT, GigaVUE-TA25, and GigaVUE-TA25E, UDA1 supports a 12-byte data match of tagged traffic with an offset starting from 16 to 116. UDA1 does not support untagged traffic. UDA2 supports a 4-byte data match with an offset starting from 0 to 60.
- Due to limitations in the platform, map rules to match the data pattern on Inner-L3 and Inner-L4 qualifiers with UDA base does not match the Inner headers for the following types of encapsulation:
 - ERSPAN
 - LISP
 - L2GRE

However, the inner headers of ERSPAN, LISP, L2GRE can be made to match with outer L3 and Outer L4 UDA base by appropriately adjusting the offsets.

- When configuring a map rule with common UDA offset, you cannot combine source ports from GigaVUE-HC1-Plus or GigaVUE-HCT or GigaVUE-TA25 devices together with source ports from other devices in a single map. You must create separate maps for GigaVUE-HC1-Plus or GigaVUE-TA25 devices and other devices.

User-Defined Pattern Match Examples

In this example, a 3G carrier is monitoring the Gn interface between the SGSN and the GGSN in the mobile core network and wants to split traffic from different subscriber IP address ranges to different tool ports. However, because the subscriber IP addresses are tunneled using the GPRS Tunneling Protocol (GTP), standard IP address map rules will not work. The addresses are always at the same offsets, though, so we can construct UDA pattern match rules to match and distribute the traffic correctly.

For example, suppose we want to apply the following rules to all traffic seen on network port 1/5/x1:

- Send all traffic to and from the 10.218.0.0 IP address range inside the GTP tunnel to tool port 1/5/x4.

- Send all traffic to and from the 10.228.0.0 IP address range inside the GTP tunnel to tool port 1/5/x9.

Keep in mind that we also know the following about tunneled GTP traffic:

- The offset for source IP addresses inside the GTP tunnel is 62.
- The offset for destination IP addresses inside the GTP tunnel is 66.

The following example explains how to construct two maps that will distribute traffic using UDA pattern match rules.

Description	UI Step
Map #1 – GTP_Map218 Our first map will send traffic to and from the 10.218.0.0 IP address range inside the GTP tunnel to tool port 1/5/x4.	
Create a map with the alias GTP_Map218 .	1. Select Maps > Maps > Maps. 2. Click New. 3. Enter GTP-MAP218 in the Alias field.
Specifies the map type and subtype.	4. Select Regular for Type. 5. Select By Rule for Subtype.
Specify that this map will match packets arriving on network port 1/5/x1.	6. Select 1/5/x1 for Source.
Specify that packets matching this map will be sent to tool port 1/5/x4.	7. Select 1/5/x4 for Destination.
Next, add the map rules for our first address range – 10.218.0.0. This IP address translates to 0ada in hex. The first rule matches the 10.218.0.0 address at the source address offset of 62 in the GTP tunnel.	8. Click Add a Rule to add the first rule. 9. Select Pass. 10. Select UDA1 and set the values: Value: 0ada0000-00000000-00000000-00000000 Mask: ffff0000-00000000-00000000-00000000 Offset: 62
The second rule matches the same address range (10.218.0.0) but at the destination address offset of 66 in the GTP tunnel. Notice that we have still specified the offset as 62 and have simply masked out to the correct location of the destination address. This way, we have still only used one of the two possible offsets in place for the GigaVUE HC Series node at any one time.	11. Click Add a Rule to add the second rule. 12. Select UDA1 and set the values: - o Value: 00000000-0ada0000-00000000-00000000 - o Mask: 00000000-ffff0000-00000000-00000000 - o Offset: 62
Save the map.	13. Click Save.
Map #2 – GTP_Map228	

Description	UI Step
Our second map will send traffic to and from the 10.228.0.0 IP address range inside the GTP tunnel to tool port 1/5/x9.	
Create a map with the alias GTP_Map228 .	1. Select Maps > Maps > Maps. 2. Click New. 3. Enter GTP-MAP228 in the Alias field.
Specifies the map type and subtype.	4. Select Regular for Type. 5. Select By Rule for Subtype.
Specify that this map will match packets arriving on network port 1/5/x1.	6. Select 1/5/x1 for Source.
Specify that packets matching this map will be sent to tool port 1/5/x9.	7. Select 1/5/x9 for Destination.
Now, create rules for the second address range – 10.228.0.0 (0ae4 in hex). As with the first range, create separate rules for the source and destination offsets inside the GTP tunnel. This address range is being sent to 1/1/x4.	8. Click Add a Rule to add the first rule. 9. Select Pass. 10. Select UDA1 and set the values: Value: 0ae40000-00000000-00000000-00000000 Mask: ffff0000-00000000-00000000-00000000 Offset: 62
Here is the companion rule for the destination address offset of 66.	11. Click Add a Rule to add the second rule. 12. Select UDA1 and set the values: - o Value: 00000000-0ae40000-00000000-00000000 - o Mask: 00000000-ffff0000-00000000-00000000 - o Offset: 62
Save the map.	13. Click Save.

User-Defined Pattern Match on GigaVUE-TA400, and GigaVUE-TA400E

The following table outlines the attributes, base offsets, offset ranges, and data widths for user-defined pattern matching (UDA) on the GigaVUE-TA400, and GigaVUE-TA400E chassis:

Attribute	Base Offset	Offset Range	Data Width
uda1	Starting of Outer Ethertype header	0-160	8 bytes
uda2	Starting of Outer L3 header	0-160	8 bytes
inner-uda1	Starting of inner L3 header	0-120	8 bytes
inner-uda2	Starting of inner L3 header	0-120	8 bytes

The data to be matched must be specified with trailing 8-byte zeros. For example to specify 0x54206874 as a match value, enter 00000000-54206874-00000000-00000000.

Flow Map Rule Combinations

The following table shows the unsupported UDA combinations:

Map1	Map2	Support Status
uda2	inner-uda2	Not Supported
inner-uda1	uda1	Not Supported
uda1	inner-uda1	Not Supported
inner-uda2	uda2	Not Supported

NOTE: When working with UDA configurations, the combination of (uda1 + inner-uda1) and (uda2 + inner-uda2) is not supported at the chassis level. You may upgrade devices with these unsupported combinations, but further editing of the rules will be restricted due to potential traffic impacts, as inner-uda takes precedence. You are allowed to delete the mapping and rules for these unsupported configurations.

Inner Header and MPLS Header Filtering

GigaVUE-OS now supports the following capabilities for GigaVUE-TA400, and GigaVUE-TA400E devices:

Inner Header Attributes

GigaVUE-OS can filter based on inner payload packet fields for tunneled traffic. This enables the extraction of flows based on inner header packet attributes without stripping the encapsulated header. GigaVUE-OS can filter traffic based on the following inner header packet attributes:

- **Inner Ethertype:** Filter based on the Ethertype of the inner packet.
- **Inner IP Version:** Filter based on the IP version (IPv4 or IPv6) of the inner packet.
- **Inner IP Fragmentation Bits:** Filter based on the fragmentation bits in the inner IP header.
- **Inner Source IP:** Filter based on the source IPv4 address of the inner packet.
- **Inner Destination IP:** Filter based on the destination IPv4 address of the inner packet.
- **Inner IP Protocol:** Filter based on the protocol (e.g., TCP, UDP, ICMP) in the inner IP header.
- **Inner Source IPv6:** Filter based on the source IPv6 address of the inner packet.
- **Inner Destination IPv6:** Filter based on the destination IPv6 address of the inner packet.

- **Inner L4 Source Port:** Filter based on the Layer 4 (TCP/UDP) source port of the inner packet.
- **Inner L4 Destination Port:** Filter based on the Layer 4 (TCP/UDP) destination port of the inner packet.
- **Inner UDA1:** Filter based on the first user-defined attribute (UDA) in the inner packet.
- **Inner UDA2:** Filter based on the second user-defined attribute (UDA) in the inner packet.

For filtering based on the DSCP value in the payload IPv4 or IPv6 headers, the inner UDA attributes can be used. These attributes take the offset from the start of the inner L3 header, ensuring accurate filtering.

Example:

The map rule shown below matches the Inner Header packets:

The screenshot shows the 'Map Rules' configuration window. At the top, there are buttons for 'Quick Editor', 'Import', and 'Add a Rule'. Below these, 'Rule 1' is selected. The configuration fields include:

- Rule Description:** A text input field.
- Address Rewrite:** A dropdown menu currently set to 'Select'.
- VLAN Action:** A dropdown menu currently set to 'None'.
- Condition:** A dropdown menu with radio buttons for 'Pass' (selected), 'Drop', and 'Bi-directional'.
- Tags:** A section with a 'TagKey' dropdown and a 'Values' input field, accompanied by '+' and '-' icons.
- Inner IPv4 Destination:** A filter box with 'IPv4 Address' input, a radio button for 'Cidr(1-32)' or 'Net Mask', and a close button 'x'.
- Inner Port Destination:** A filter box with 'Min' and 'Max' input fields (both set to '0 to 65535'), a 'Subset' dropdown, and a close button 'x'.
- Inner IPv4 Protocol:** A filter box with a 'Value' dropdown (set to 'IPv6Hop') and a close button 'x'.

Figure 6 Map Rule with Inner IPv4 Destination, Inner Port Destination, and Inner Ipv4 Protocol

MPLS Header Attributes

GigaVUE-OS also supports filtering based on MPLS header attributes, allowing traffic to be forwarded without stripping the MPLS header. The following attributes can be filtered:

- **MPLS Label ID:** Filter based on the MPLS Label ID.
 - Position any: Specifies all levels (1-7) of the MPLS header in the packet.

- Position 1-7: Specifies the depth of the MPLS header in the packet (1 = outermost label, 7 = innermost label).
- **MPLS Label EXP (Experimental):** Filter based on the EXP bits in the MPLS header.
 - Position any: Specifies all levels (1-7) of the MPLS header in the packet.
 - Position 1-7: Specifies the depth of the MPLS header in the packet (1 = outermost label, 7 = innermost label).
- **MPLS Label BOS (Bottom of Stack):** Filter based on the BOS bit in the MPLS header.
 - Position any: Specifies all levels (1-7) of the MPLS header in the packet.
 - Position 1-7: Specifies the depth of the MPLS header in the packet (1 = outermost label, 7 = innermost label).
- **MPLS Label TTL (Time to Live):** Filter based on the TTL value in the MPLS header.
 - Position any: Specifies all levels (1-7) of the MPLS header in the packet.
 - Position 1-7: Specifies the depth of the MPLS header in the packet (1 = outermost label, 7 = innermost label).

Traffic can be filtered based on these attributes and sent to the desired tools without modifying or stripping the MPLS header.

Example:

The map rule shown below matches the MPLS Header packets:

The screenshot displays the 'Map Rules' configuration window. At the top, there are buttons for 'Quick Editor', 'Import', and 'Add a Rule'. Below these, a rule named 'Rule1' is shown with the following configuration:

- Rule Description:** A text field labeled 'Description'.
- Address Rewrite:** A dropdown menu currently set to 'Select'.
- VLAN Action:** A dropdown menu currently set to 'None'.
- Condition:** A dropdown menu with radio buttons for 'Pass' (selected), 'Drop', and 'Bi-directional'.
- Tags:** A section with a 'TagKey' dropdown and a 'Values' input field, accompanied by add (+) and remove (-) icons.
- Filters:** Three filter blocks are listed:
 - MPLS Label BOS:** Includes a 'Position' dropdown set to 'Select Position...'.
 - MPLS Label ID:** Includes a 'Position' dropdown set to 'Select Position...'.
 - MPLS Label TTL:** Includes a 'Position' dropdown set to 'Select Position...'.

Figure 7 Map Rule with MPLS Label BOS, MPLS Label ID and MPLS Label TTL

To configure an Inner Header qualifier and MPLS Header qualifier using GigaVUE-OS CLI, refer to the [Configure Inner Header Qualifier and MPLS Header Qualifier](#) chapter in the GigaVUE-OS CLI Reference Guide.

How to Handle Q-in-Q Packets in Maps

The Q-in-Q packets in maps are handled as follows:

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

The screenshot shows the configuration for a rule named 'Rule1'. At the top, there are buttons for 'Quick Editor', 'Import', and 'Add a Rule'. Below these, the 'Condition' is set to a dropdown menu, and the action is set to 'Pass' (selected with a radio button), with options for 'Drop' and 'Bi-directional'. The 'Rule Description' is set to 'Description'. Under the 'Tags' section, there is a 'TagKey' dropdown and a 'Values' input field with add (+) and remove (-) icons. Below this, a 'VLAN' section is expanded, showing 'Min 100' and 'Max 1 to 4094', with a 'Subset' dropdown menu.

- 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.

The screenshot shows the configuration for a rule named 'Rule1'. At the top, there are buttons for 'Quick Editor', 'Import', and 'Add a Rule'. Below these, the 'Condition' is set to a dropdown menu, and the action is set to 'Drop' (selected with a radio button), with options for 'Pass' and 'Bi-directional'. The 'Rule Description' is set to 'Description'. Under the 'Tags' section, there is a 'TagKey' dropdown and a 'Values' input field with add (+) and remove (-) icons. Below this, a 'VLAN' section is expanded, showing 'Min 100' and 'Max 1 to 4094', with a 'Subset' dropdown menu.

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

Invalid ethertype : '0x8100'. Please use attribute 'vlan' instead.

Invalid ethertype : '0x88A8'. Please use attribute 'vlan' instead.

Invalid ethertype : '0x9100'. Please use attribute 'vlan' instead.

NOTE: The **Value** field accepts values with out the leading 0x only.

In summary, for single-tagged (0x8100) or double-tagged (0x88A8 and 0x9100) VLAN packets, you only configure the VLAN as the matching criteria, not the ethertype.

For handling of priority tagged packets, refer to [Priority Tagged Packets](#).

For filtering of Q-in-Q packets on inner VLAN tag, refer to [Flow Mapping® on Inner VLAN Tags](#).

Comparison of Q-in-Q Tagging

The following table details the various combinations and corresponding behaviors depending on the packet content:

Packet Content	Rule: pass vlan 100	Rule: pass ethertype 0x0800	Rule: pass vlan 100 ethertype 0x0800
No tags, ethertype 0800	drop	pass	drop
One tag: TPID 8100, VID 100, ethertype 0800	pass	pass	pass
One tag: TPID 9100, VID 100, ethertype 0800	pass	pass	pass
One tag: TPID 88a8, VID 100, ethertype 0800	pass	pass	pass
Two tags: outer TPID 8100 VID 100, inner TPID 8100 VID 200, ethertype 0800	pass	pass	pass
Two tags: outer TPID 9100 VID 100, inner TPID 8100 VID 200, ethertype 0800	pass	pass	pass
Two tags: outer TPID 88a8 VID 100, inner TPID 8100 VID 200, ethertype 0800	pass	pass	pass
Two tags: outer TPID 8100 VID 200, inner TPID 8100 VID 100, ethertype 0800	drop	pass	drop
Two tags: outer TPID 8100 VID 200, inner TPID 88a8 VID 100, ethertype 0800	drop	drop	drop
Two tags: outer TPID 88a8 VID 200, inner TPID 8100 VID 100, ethertype 0800	drop	pass	drop

Packet Content	Rule: pass vlan 100	Rule: pass ethertype 0x0800	Rule: pass vlan 100 ethertype 0x0800
Two tags: outer TPID 8100 VID 100, 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	✗	✓	✓	✓	✓
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/Owner** – Can perform all the Read/Write functions and assign map sharing permission levels.

***Requires Level 2 or Level 3 access, based on User Group membership.**

[Table 3: Permission Levels for Map Sharing](#) summarizes the permission levels for map sharing.

Table 3: Permission Levels for Map Sharing.

Permissions	Read Only	Listen	Read/Write	Read/Write/Owner
View map	✓	✓	✓	✓
Add tool port*	✗	✓	✓	✓
Remove tool port	✗	✓*	✓	✓
Remove network port	✗	✗	✓	✓
Add network port*	✗	✗	✓	✓
Delete/edit map	✗	✗	✓	✓
Share map	✗	✗	✗	✓

***Only applies to ports to which the user has Level 2 or Level 3 access.**

NOTE: In [Table 3: Permission Levels for Map Sharing](#), tool port includes ports of type tool and inline-tool. Network port includes ports of type network and inline-network.

The admin user can also assign map sharing permissions.

Users with Level 1 (or greater) access to a given network port can also view, but not edit, maps associated with that network port. This is independent of the map sharing permissions.

Map sharing permissions override and supersede role based access controls. Thus, a user group can be assigned Read/Write access to map even if they do not have any access rights to any of the associated network or tool ports. However, adding tool ports to a map or removing network or tool ports from a map requires Level 2 or Level 3 permissions, as defined by the user group, for the ports to be added or removed.

Map Examples

This section provides the following map examples:

- [Example: How to Create a Simple Map](#)
- [Example: How to Handle Overlaps when Sending VLANs and Subnets to Different Tools](#)

Example: How to Create a Simple Map

In this example, a few simple maps are illustrated to show how to create and display the packet distribution in place on the node.

When you set up flow maps from the perspective of your tools, start by asking yourself which traffic you would like a particular tool to see. Then, select the necessary traffic from network ports.

For example, the scenario in this example is as follows:

- An application performance management tool is connected to **tool port 2/4/x6** that focuses on traffic from **VLANs 100..199** on **network ports 2/2/x10** and **2/2/x12**.
- An application performance Management tool connected to **tool port 2/4/x18** that focuses on traffic from **VLANs 200..299** on **network ports 2/2/x10** and **2/2/x12**.

Consider a GigaVUE-HC3 device with the following configurations:

- BoxID of the GigaVUE-HC3 device is set as 2.
- The network ports are set on the second blade in the node.
- The tool ports are set on the fourth blade in the node.

Therefore, the ports in this scenario are represented as follows:

- The network port IDs are set as 2/2/x2, 2/2/x10, and 2/2/x12.
- The tool port IDs are set as 2/4/x6, 2/4/x8, 2/2/x20, and 2/4/x24

The maps for this scenario the map types are set as follows:

- The maps with the VLAN rules specified have the subtype set as **By Rule**.
- The shared collector maps are set as subtype **Collector** and no rules added.
- The passall maps are defined as subtype **Pass All** and no rules added.

For details about the different map types, refer to [Matrix of Map Types](#) and [Map Subtypes](#).

Except for the following map types, the map type defaults to **Regular**

- Only maps with Inline Bypass solutions can be set as type **Inline**.
- Only maps that have GigaSMART operations defined for second level maps can be set as **First Level** or **Second Level** maps.

The following are the steps to create the simple maps as shown in [Map Examples](#):

1. Check the port types for the each of the ports that the maps will use and set them if necessary.

To set the port types, select Port > Ports > All Ports and use the Port Type Editor to set the port types. For more information about port types, refer to

2. Create a **Regular** map with a **By Rule** subtype to pass the VLAN100 traffic as shown in [Figure 8Map for VLAN 100..199](#).

The screenshot displays the GigaVUE configuration interface for creating a map. It is divided into three main sections:

- Map Info:**
 - Map Alias:** vlan100s
 - Comments:** (empty text field)
 - Type:** Regular (dropdown menu)
 - Sub Type:** By Rule (dropdown menu)
 - Rule Matching:** ☐ Blacklist
- Map Source and Destination:**
 - Port Editor:** (button)
 - Source:** 1/1/x1, 1/1/x4 (with 'x' icons for removal)
 - Destination:** 1/1/x2 (with 'x' icon for removal)
 - GSOP:** None (dropdown menu)
- Map Rules:**
 - Quick Editor, Import, Add a Rule:** (buttons)
 - Rule 1:**
 - Condition search...** (text field)
 - Pass, Drop, Bi Directional:** (radio buttons, with 'Pass' selected)
 - VLAN:**
 - Min:** 100 (with up/down arrows)
 - Max:** 199 (with up/down arrows)
 - Subset:** none (dropdown menu)

Figure 8 Map for VLAN 100..199

3. Create a **Regular** map with a **By Rule** subtype to pass the VLAN200 traffic.
4. Create a map with type **Regular** and subtype **Collector** to create a shared collector map. For more information about shared collector maps, refer to [About Shared Collectors](#).
5. Create a map with type **Regular** and subtype **Pass All** to create the passall map.

Example: How to Handle Overlaps when Sending VLANs and Subnets to Different Tools

[Figure 9 Sending 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 9 Sending 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 9 Sending 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](#).

Splitting Subnets and VLANs

In this example, we want to send all packets on the 172.16.0.0 subnet to 1/2/x1, all packets on the 172.17.0.0 subnet to 1/2/x2, and all packets on VLAN 100 to 1/2/x3. Our concern is how to handle packets that are on **both** VLAN 100 **and** one of those two subnets.

To handle this, we give our highest priority to packets matching both VLAN 100 and either one of the two subnets. Notice how the first two maps entered -- the maps with the highest priority -- combine the subnet and VLAN criteria in a single line. Packets matching **both** of these criteria will be sent to the ports both for their subnet and for their VLAN criteria. Because we entered these maps first, they have higher priority than the maps that simply match the subnet or VLAN criteria.




Figure 9 *Sending Subnets and VLANs to Different Ports*

Manage Maps

This section provides a description of the Maps pages in the GigaVUE-FM UI. It covers the following topics:

- [Map Views](#)
- [Manage Maps](#)
- [Map Templates](#)
- [Manage Map Rule Resources](#)
- [Filter Templates](#)
- [Review Map Statistics with Map Rule Counters](#)

Map Views

The Maps page displays the maps created using the CLI, or GigaVUE-FM.

To view the map, do the following:

1. Go to **Inventory > Physical > Nodes**.
2. Click the Cluster ID link, go to **Traffic > Maps > Maps**.

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

The maps can be displayed in List or Map Topology views.


List View

The List View is the default view of the maps when the Maps page is opened. This view shows the basic information about each map:

- Alias
- Type
- Subtype
- Source and destination ports
- Control Traffic: Displays if the control traffic is enabled or disabled for a particular map. This is applicable only for First Level, By Rule maps.

Map Topology View

Click the  icon. The Topology View of the maps appears. Each map in the Maps page gets displayed as Topologies. The Map Topology presents a graphical representation of all maps available from the device, including both system-defined and user-created maps. The system-defined maps are displayed when you click the **filter**. A pop-up appears, and then you enable **Show Auto-Generated Maps**.



Alias	Map ...	Sour...	Dest...	Enca...	Des...	Ena...	Type	Subt...	Num.
☒ > FmAuto-TP1-YAO-d90be24f-46dc-4eb5-b440-efe0...		1 Port	1 Port		Crea...	true	Regu...	By R...	2
☐ > Test		1 Port	1 Port			true	Regu...	By R...	0
☐ > Test2		1 Port	1 Port			true	Regu...	By R...	0
☒ > FmAuto-AMX-seu-1072ed73-7ee3-4d86-b149-32...		1 Port	1 Vir...		Crea...	true	First ...	By R...	1

Go to page: 1 of 1 Total Records: 4

Filter

Map Alias

Type

Sub Type

Priority

Source Ports

Destination Ports

Tags


Control Traffic

Alt Enabled Disabled

Show Auto Generated Maps ☒

You can interact with the topology to efficiently search for ports, navigate through the layout, and inspect individual elements in detail.

If a particular map is selected, only that map is displayed with its details, instead of showing all maps. You can use the back and forward arrows to navigate to the previous and next maps, respectively. The **Home** option allows you to return to the main maps page.



GigaVUE-FM provides multiple options for viewing map topology, including QuickView and Free Layout.


Quickview

If QuickView is enabled and you select a port, its details are displayed in a pane. In QuickView, you can move ports up and down only.


- When you click on any node in the Topology View, a Quick View panel opens, displaying details for that element.
- The Quick View panel opens, showing details such as port details, map configuration, and associated connections.
- Use the zoom and scroll buttons at the bottom of the page to zoom in and out.
- Click on a map to view the connections.
- Use the Search field to enter the name or type of map you want to locate. The topology view filters and displays only matching maps.

Without enabling QuickView, selecting a map display only that map in the topology view.

The topology section provides access to all maps created on the device by the user. You can filter for a specific type of map to view its details. Only the searched map will be displayed. You can search using port alias, port ID, etc. For example, if you search for the x23 port, only the map containing that port will be shown.




NOTE: Searching narrows the display to matching items only, making it easier to locate specific maps in large deployments.



Free Layout

NOTE: Starting from 6.14, a new **Free Layout** option in the topology view lets you freely rearrange these canvas elements to match your preferred visual layout. Your layout is stored as a user preference and is automatically restored when you return to the Maps topology view.

1. Click on the  icon. The Topology View of the maps appear. Each map in the Maps page gets displayed as Topologies.
2. In the topology view, turn on the Free Layout option (toggle). When Free Layout is enabled, the elements on the topology canvas become freely movable.
3. Click and drag map elements (such as network, map layers, tool ports, virtual ports, and hybrid ports) to the desired positions.
4. Arrange the elements in any logical pattern you prefer. For example, vertically from network (source) to tool port (destination) instead of the default horizontal layout.
5. Use the zoom and scroll buttons at the bottom of the page to zoom in and out.
6. Click on a map to view the connections.



Reset to Default Layer

You can return to the default topology layout, do the following:

1. In the Map Topology view, disable the **Free Layout**.
2. The topology resets to the default layout, clearing any applied filters or navigation state.

Manage Maps

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

- [Create a new map](#)
- [Clone Map](#)
- [Edit Maps](#)
- [Delete Maps](#)
- [Create Map Groups](#)

It also includes information about the following:

- [Description to Map Rules](#)
- [How to Use the Quick Editor for Pass and Drop Rules](#)

Create a new map

Prerequisites:

1. Check the status of the nodes and ports that you plan to use with the map. For information about how to check the status of the nodes and ports, refer to [Status of Line Cards/Nodes and Ports](#).
2. From the device view, go to **Maps > Maps** to open the Maps page.
3. Click **New**.

Create a new map – Field References

The following table lists and describes the fields you must complete to configure new map.

Field	Description
Map Info	
Map Alias	A map alias specifies the name of the map. The alias must be unique and can contain up to 128 alphanumeric characters. Aliases are case-sensitive. Most special characters are supported in map aliases. However, map aliases that are only one period (.) or two periods (..) should not be created. These aliases cannot be accessed for editing. Use an alias that helps identify the task and destination. For example netflix_traffic_to_wireshark.
Description	(Optional) The description for the map. When adding description, consider the following: • Use up to 128 characters, including spaces. • Enclose the description in quotation marks, if the description is longer than one word.

Field	Description
	To include double quotation marks (") inside the quotation marks, precede it with a backslash (\). See also Description to Map Rules.
Enable	Checkbox to enable or disable the map.
Type	Select the map type. The map type can be Regular, First Level, Second Level, Inline, Inline First Level, Inline Second Level, or Transit Level. For detailed information about the types of maps, refer to Map Types
Subtype	Select the subtype. The map subtype can be By Rule, Pass All, or Collector. For detailed information about Pass All, refer to About Map-passall Maps. For detailed information about Collector, refer to About Shared Collectors.
No Rule Matching	Checkbox for excluding rule-based matching. Enable the Pass Traffic checkbox if no rules are matching.
Control Traffic	Enable the Control Traffic checkbox to pass the GTP control traffic (GTP-c) to all GigaSMART engines in a GTP engine group. A GTP engine group has multiple GigaSMART engine port members. Note: The Control Traffic checkbox is applicable only for GTP and is displayed only if the map type is configured as First level, and the map sub type is configured as By Rule.
Map Source and Destination Map	
Source	From the Source drop-down list, select the required source ports for the map. To create a port list, click Port Editor. Note: You can add a maximum of 324 ports in the Source drop-down list, if the ports are not attached to a GigaStream. Note: For details about port types that are supported for the different types of maps, refer to Port Lists.
Destination	From the Destination drop-down list, select the required destination ports for the map. To create a port list, click Port Editor .
Encapsulation Tunnel	If you have selected a circuit port in the Destination drop-down list, select the required circuit tunnel from the Encapsulation Tunnel drop-down list to encapsulate the traffic. If the map is used to redirect the decapsulated traffic to the required tool ports, ensure that you select the IP interface in the Source drop-down list.

Field	Description
	You must have attached the IP interface to the VXLAN or L2GRE tunnel. For details about VXLAN or L2GRE tunnels, refer to About Virtual Extensible LAN (VXLAN) Tunnels and About Layer 2 Generic Routing Encapsulation (L2GRE) Tunnels . From the Destination drop-down list, select the required tool ports.
GigaSMART Operations (GSOP)	If the map will use a GigaSMART operation, select the operation from the GigaSMART Operations (GSOP) drop-down list.
Map Configuration and Rules	
Address Rewrite	(Optional) Enable the address rewrite checkbox and select either MAC Address or IPv4 Address. For detailed information about MAC Address Rewrite, Refer to MAC Address Rewrite. For detailed information about IP Address Rewrite, Refer to IP Address Rewrite.
VLAN Action	Select Add option from the VLAN Action field to add a new VLAN tag to the outgoing traffic. For detailed information about VLAN Manipulation, Refer to VLAN Manipulation.
Map Rules	To add rules to the map, do any of the following: • Use the Quick Editor. For details, refer to the How to Use the Quick Editor for Pass and Drop Rules. • Import a map template by clicking Import. • Create a rule by clicking Add a Rule. For detailed information about map rules, refer to Map Rules.
Map Order	
Priority	Set the Map Order by selecting the priority from the Priority list. For details about map priority, refer to Map Priority.
Map Permissions	
Owner	Select user group(s) who own the map.
Edit	Select user group(s) who can edit the map.
Listen	Select user group(s) who can listen to the traffic.

Field	Description
View	Select user group(s) who can view the map.
Map Tag	
Tags	Select the required tag key and tag value to which the map must be associated. The tag key and the associated tag values must be created in advance in GigaVUE- FM. Refer to the "Tags" and "Role Based Access Control" sections in the GigaVUE Administration Guide for more details. Note: When you associate a map to a tag value, then the ports, port groups, port pairs, GigaStreams that belong to the map are also associated to the tags.

Clone Map

You can create a copy of an existing map by doing the following:

1. Select **Maps > Maps** to open the Maps page.
2. Select the check box of the map that you to clone.
3. Click **Clone**.
4. Make changes to the map as necessary, such as specifying an alias.
5. Click **Save**.

Edit Maps

To edit an existing map:

1. Click the List view button.
2. Select the map on the Maps page by either selecting the check box and then clicking **Edit**, or click on the row in the table and clicking **Edit** in the Map Quick View.
3. Make the changes to the map.

When editing a map, you can only modify the following:

- Map alias
 - Description
 - Change the source and destination.
 - Select a different GS Operation.
 - Modify rules.
 - Add rules.
 - Delete rules.
 - Permissions as allowed.
4. Click **Save**.

Delete Maps

Keep in mind the following rules and notes before you delete a map:

- If you have configured an Active Visibility policy action for a map, ensure that you delete the Active Visibility policy before you delete the map. If you delete the map before deleting the Active Visibility policy, it would result in errors in the policy action.
- If you delete a map, which has a network port that is shared with another map, the map statistics will be reset for the other map.

To delete a map or maps, do the following:

1. Go to **Maps**, select **Maps > Maps** to open the Maps page.
2. Select the check box for the map or maps to delete, and then Click **Delete**.
3. When the message appears, asking if you want to delete the selected maps, click **OK**.

NOTE: In the GRIP configuration, when you delete a map on the primary node, irrespective of the inline-network traffic-path, the traffic is switched to the secondary node. The port utilization must be 0% on the primary node and active on the secondary node.

Create Map Groups

Map Groups are a collection of maps that are used with GTP Flow Sampling Overlap and GTP Whitelisting Overlap GigaSMART solutions.

Use the Map Groups page to create a group of maps for GTP forward listing and GTP flow sampling. All the maps in a map group receive traffic according to map rules, rather than map priority. Thus, multiple copies of a GTP packet can be sent to more than one tool. This functionality is referred to as overlapping maps.

The virtual port for specified as the source for GTP forward listing and GTP flow sampling must have **GTP Overlap** enabled.

When creating Map Groups keep the following in mind:

- A map group can be associated with only one GigaSMART group (gsgroup).
- All maps within a map group must be connected to the same vport.
- A map group can consist of only one GTP forward listing map or only one GTP flow sampling map but it cannot contains two maps of the same type.
- Once you have created a map group, you cannot edit it to change the type or subtype of the map. However, you can add or edit the map rules for a map, which is configured in a map group. Similarly, you cannot edit the map group alias but you can edit the aliases of the maps that are configured in the map group.

- If multiple map groups are configured, the maps within each map group must point to the same port groups as the other map groups.

To create a Map Group, do the following:

1. Go to **Maps**, select **Maps > Map Groups** to open the Maps page.
2. Click **New**.
3. Enter an alias for the map group.
Use an alias that helps identify the map group.
4. (Optional) Enter a description about the map group. Refer to [Description to Map Rules](#) for the considerations regarding description.
5. Click in the **Maps** field and select the maps to add to the map group.
6. Click **Apply**.

Description to Map Rules

Use description to label the purpose of a rule or the type of traffic covered by a rule. To add a map rule description to a map select **Maps > Maps > Edit**.

Consider the following when adding map rule description:

- Use up to 128 characters, including spaces.
- Enclose the description in quotation marks, if the description is longer than one word.
- To include double quotation marks (") inside the quotation marks, precede the quote mark with a backslash (\).

Error Messages

Error messages are displayed when a description is invalid, for example:

- if the description is longer than one word and does not include double quotation marks
- if the description is longer than 128 characters
- if the rule with which the description is included is not valid.

Edit Map Rule Description

To edit a map rule description, do the following:

1. Select **Maps > Maps**
2. Select the map to edit.
3. Click **Edit**.
4. Change the description in the **Description** field

5. Click **Save** to recreate with a different description.

How to Use the Quick Editor for Pass and Drop Rules

When creating a map, you can use the Quick Rule Editor to quickly select custom port numbers for a map rule or add a range of IP address.

When manually configuring an Inline Bypass rule-based map in the device, it's important to understand how the Quick Rule Editor behaves based on the map's configuration mode:


- a. If the map type is set to '**Inline**' mode before using the Quick Rule Editor. Any rules added through the Quick Rule Editor will be applied with default, non-editable values, resulting in "**read-only**" rules.
- b. If the map remains in its default mode (not set to 'Inline') when using the Quick Rule Editor. Rules can be added and subsequently edited as needed.

How to Use the Quick Editor to Add Port Numbers

To use the Quick Rule Editor, do the following:


1. While on an Edit Map or New Map page, click **Quick Editor** under Map Rules. This opens the **Quick Rule Editor**.
2. On the Quick Rule Editor view, select the port number or numbers to add for a pass or drop rule or both.

The Quick Rule Editor has two columns of custom port numbers, one for pass rules and one for drop rules. In each column, the ports are categorized by type. For example, **Web** provides a list of HTTP ports as shown in the following figure, where HTTPS port 443 is selected for pass and HTTP port 80 is selected for drop.



3. Click **Add**.

A rule with a port source is added for each custom port number selected in the Quick Rule Editor. If the port was selected from the custom port numbers under Pass, the rule is a pass rule. If the port was selected from the port numbers under Drop, the rule is a drop rule. The following figure shows two rules added by the example shown in the previous step.



How to Use the Quick Port Editor to Add IP Address

Rather than enter IP address for map rules one at a time, the Quick Rule Editor makes it possible to enter a range to quickly add the IP addresses, saving time.

To enter a range of IP address with the Quick Rule Editor, do the following:

1. While on an Edit Map or New Map page, click **Quick Editor** under Map Rules. This opens the **Quick Rule Editor**.
2. Enter an IP address range in the **List** field under IPv4 Source or IPv4 Destination or both. For example, 10.10.10.9..11/32 in the **List** field under IPv4 Source .
3. Click **Add**.

The Quick Rule Editor adds the IPv4 Source rules with the IP addresses. For example, if you entered 10.10.9..11 for the IPv4 Source, the editor adds three Ipv4 Source rules with the IP addresses 10.10.10.9/32, 10.10.10.10/32, 10.10.10.11/32.

Map Templates

Map templates can be created by admin users. Once created, any user creating a map can use any template.

Admin users can define standardized traffic flows, applications, and rules that will be convenient for users when creating their maps. To do this, the admin creates map templates; later, users can use one or more templates as the basis for their maps.

Templates are created using the same rules and parameters as regular maps, but they do not have any network or tool ports. GigaSMART operations are also not included in templates.

The rules defined in the template become the starting point for the map. The rules can be edited or removed and new rules can be added to the map.

No changes made to the map will be reflected back in the original template. Once the map is created, it has no association with the original template from which it was created. Any changes to a template will not be reflected in any maps created with the previous version of the template.

Create Map Templates

To create a map template, do the following:

1. Select **Traffic > Maps > Map Templates**.
2. Click **New**. The New Map Template opens.
3. Enter an alias in the **Map Template Alias** field.
4. (Optional) Enter description about the template.
5. Add map rules by clicking **Add a Rule** for each rule that you want to add.
6. After you are done creating rules, click **Save**.

The new map is included on the Map Templates page.

Edit Map Templates

To edit a map template, do the following:

1. Select **Traffic > Maps > Map Templates** to open the Map Templates page.
2. On the Map Template page, select the template to edit.
3. Click **Edit**.
4. Modify the map template by adding or deleting rules or description. (You cannot change the alias.)
5. Click **Save**.

Map Template Quick View

When you click on the alias of a map template on the Map Templates page, a Quick View displays that shows the comments (if any) and rules.

Manage Map Rule Resources

The resources available on a GigaVUE HC Series line card or node changes depending on the combination of criteria in place on the line card or node as a whole. In general, adding or removing MAC address, UDA pattern match, or IPv6 criteria in the map rules bound to a line card or node changes the type of filter template used on the line card or node. This can result in a brief interruption of traffic as the new template is applied.

Template Groups

The template groups are listed sequentially from least resource-intensive to most resource-intensive:

- **IPv4 Only** – This is the default filter template, including all IPv4 arguments without any MAC addresses, UDA data patterns, or IPv6 arguments. This template can support the IPv4 and related filter criteria, including VLAN tags, source/destination ports, protocol criteria, and so on.
- **IPv4 and MAC Addresses** – This template combines MAC address criteria with the standard IPv4-related criteria. When MAC address criteria are in use, map rule resources are decreased.
- **IPv4 and Single UDA Data Pattern** – This template combines one of the two available UDA data patterns with the standard IPv4-related criteria. Using a single UDA criteria does not affect the total number of drop map rules available, but it does decrease the number of pass map rules available.
- **Both UDA Data Patterns** – This template uses both UDA data patterns but removes the ipv4 argument. Drop map rule availability is not affected by adding a second UDA data pattern, but pass map rules are decreased again from what was available when only a single UDA was used.
- **IPv6 Arguments** – This template adds the use of the IPv6 argument. IPv6 criteria are resource-intensive, significantly decreasing both drop and pass map rule capacity, as shown in the following table. Note also the changes in available criteria and available resources.

The map rule criteria available in each filter template (or “group”) is shown in the following table.

Table 4: Map Rule Criteria for Default Templates

Rules - Pass and Drop	IPv4	IPv4 + MAC	IPv4 + UDA	UDA	IPv6
IPv4	✓	✓	✓		✓
IPv6					✓
MAC		✓			
UDA1			✓	✓	
UDA2				✓	
VLAN	✓	✓	✓	✓	✓
Inner VLAN	✓	✓	✓	✓	✓
L4 Port dst and src	✓	✓	✓	✓	✓
Ethertype		✓	✓		
IP Ver	✓	✓	✓	✓	✓
Protocol	✓	✓	✓	✓	✓
DSCP	✓	✓	✓	✓	✓
ToS	✓	✓	✓	✓	✓

Rules - Pass and Drop	IPv4	IPv4 + MAC	IPv4 + UDA	UDA	IPv6
TCP Ctl	✓	✓	✓	✓	✓
IP Frag	✓	✓	✓	✓	✓
TTL	✓	✓	✓	✓	✓
IPv6 Flow Label		✓	✓		

The number of rule entries in a cluster is shown in the following table.

Table 5: Rule Entries

# Rule Entries	IPv4	IPv4 + MAC	IPv4 + UDA	UDA	IPv6
GigaVUE-HCT	3071	3071	3071	3071	3071
GigaVUE-HC1-Plus	3071	3071	3071	3071	3071
GigaVUE-HC3	4096(1024 per slot)	4096(1024 per slot)	4096(1024 per slot)	4096(1024 per slot)	4096(1024 per slot)
GigaVUE-HC1 Node	16384	8192	8192	8192	8192
GigaVUE-TA100 Node	1024	1024	1024	1024	1024
GigaVUE-TA200 Node	1024	1024	1024	1024	1024
GigaVUE-TA200E	1024	1024	1024	1024	1024
GigaVUE-TA25	3071	3071	3071	3071	3071
GigaVUE-TA25E	3071	3071	3071	3071	3071
GigaVUE-TA400	10229	6133	10229	10229	6133
GigaVUE-TA400E	10229	6133	10229	10229	6133

NOTE: The above table lists the values for the Default template with license.

Refer to the following notes:

- You can use flexible filter templates on GigaVUE-HC3 to support up to 6K rules per slot.
- GigaVUE-TA100, GigaVUE-TA200 and GigaVUE-TA200E support up to 6K rules per pseudo-slot or 24K rules per node.
- GigaVUE-TA25 supports up to 18K rules per node and GigaVUE-TA25E supports up to 36k rules per node.

Refer to [Filter Templates](#) for details.

Add Tags to Map Rules

You can associate tags to the map rules. Each rule within a map can be associated with different tag keys and tag values. Tagging at rule level allows you to view the statistics of the flow of traffic based on the rules associated to the traffic.

Consider the following maps and map rules associated to tags:

Maps	Tag Key	Tag Value
Map 1 Src: 1/1/x1 Dst: 1/1/x2	Maptag	M1
- Rule 1 Vlan 100	Flow	Flow1
- Rule 2 Vlan200		Flow2
Map 2 Src: 1/1/x3 Dst: 1/1/x4	Maptag	M2
- Rule 1 Vlan 100	Flow	Flow1
- Rule 2 Vlan200		Flow2

With rule-level tagging, you can view the traffic flow associated with Rule 1 of both the Maps based on the tag value Flow1, even though the rules are defined as part of different maps. Refer to the [Map Rule Statistics Dashboard](#) section for viewing statistics related to the flow of traffic based on the rules defined.

To associate tags to map rules:

1. When creating a map, click on **Add a Rule**.
2. Add the required tags to the rule.

Map Rules

Quick Editor

Import

Add a Rule

✖ Rule 1

Condition search...

☒ Pass
☐ Drop
☐ Bi-directional

Rule Description

Description

Tags

TagKey

Values

+

-

Notes

Refer to the following notes:

- Rule level tagging is applicable for:
 - Flow Maps
 - Fabric Maps
- Rule level tagging is applicable for the following map types:
 - Regular Maps
 - First Level Maps
- You can associate only aggregation tags with map rules. However, traffic statistics displayed in the Map Rule Statistics dashboard is based on the following
 - RBAC tags associated at the map level.
 - Aggregation tags associated at the rule level.

Map Rule Statistics Dashboard

The Map Rule Statistics dashboard displays statistical data related to the map rules and the associated traffic. To view the Map Rule Statistics dashboard:

1. On the left navigation pane, select  -> **Analytics** -> **Dashboards**.
2. Click **Dashboards** from the top menu. Select **Map Rule Statistics** from the list of dashboards.

The **Map Rule Statistics** dashboard is displayed. The visualizations are displayed in the following two tabs:

- Metric
- Trend

The Metric tab displays the following visualizations

- **Total Traffic:** Total traffic in packets per second
- **Average Traffic:** Average traffic in packets per second
- **Associated Map Rules:** Map rules associated with the traffic

The Trend tab displays the following visualizations:

Visualizations	Displays...
Maximum vs. Average Traffic in pps	Maximum traffic versus average traffic in packets per second on an hourly interval.
Maximum vs. Average Traffic in bps	Maximum traffic versus average traffic in bytes per second on an hourly interval.
Traffic Distribution at Rule Level in pps	Traffic distribution at rule level in packets per second.
Top-N Rules by Maximum Traffic in pps	Top 5 rules contributing to the maximum traffic in packets per second
Top-N Rules by Average Traffic in pps	Top 5 rules contributing to the average traffic in packets per second

Use the following 'Control Visualizations' to filter the traffic:

- Flow Sample
- Cluster ID
- Host Name
- Map Rule Sample

Customize the **Flow Sample** and **Map Rule Sample** options to filter the data based on the tags configured in GigaVUE-FM. Refer to the [Filter Data Using Tags in Control Visualizations](#) section for more details.

NOTE: Refer to the [Fabric Health Analytics](#) section for information about filtering and searching in the dashboards page.

Filter Templates

Filter templates maximize the number of map rules, optimize filter resources, and enhance the scalability and flexibility of flow mapping. Flexible filter template is supported in GigaVUE-HC1-Plus, GigaVUE-HCT, GigaVUE-HC1, GigaVUE-HC3, GigaVUE-TA100, GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA400, and GigaVUE-TA400E.

Refer to [Manage Map Rule Resources](#) for template groups on other GigaVUE nodes.

Flexible filter templates increase the number of map rules and also eliminate current restrictions on map rule combinations, such as ipv6+MAC or ipv6+UDA.

Refer to the section [Flow Mapping® FAQ](#) for the number of map rules supported.

Flow Mapping® uses filter templates to determine the traffic to filter based on qualifiers specified in the template. A filter template has a specific set of qualifiers used to apply to map rules. You can control the template that you apply to a specific slot on GigaVUE-HC3 or a specific pseudo-slot on GigaVUE-TA100, GigaVUE-TA200, GigaVUE-TA400, and GigaVUE-TA400E. For GigaVUE-HC1, you can apply the filter template only at the control card level which will be applied across all the line cards.

Flexible filter templates offer five default templates. Custom templates can also be created that have a qualifier set selected from the list of available qualifiers.

For the GigaVUE-TA400, and GigaVUE-TA400E platform, custom templates can be created with the inner qualifiers and MPLS headers. Inner qualifiers provide extended support for packet processing in the Flexible filter template. MPLS header attributes now support up to 7 levels, allowing for detailed and specific MPLS header field mapping and processing.

Refer to the following sections for details:

- [Filter Template Qualifiers and Defaults](#)
- [Filter Template Configuration](#)
- [Filter Template Limits](#)
- [Filter Template Rules and Recommendations](#)
- [Filter Template Best Practices](#)
- [Filter Templates in a Cluster](#)
- [Filter Templates Formulas](#)
- [Troubleshoot Custom Filter Template Issues](#)

Filter Template Qualifiers and Defaults

Refer to the rows in [Map Rule Criteria for Default Templates](#) for the list of qualifiers for filter templates. Refer to the columns in [Table 4: Map Rule Criteria for Default Templates](#) for the default templates and the qualifiers that are predefined for the defaults.

NOTES:

- The default templates cannot be deleted.
- The **ipver** qualifier is implicitly included in all default and custom templates.

Filter Template Configuration

To configure filter templates:

1. On the left navigation pane, go to **Inventory**  and select **Physical**.
2. Select the required GigaVUE node.
3. Select **Maps > Filter Templates**.
4. To add a custom template, click **New**.
5. Specify an alias, an optional description, then select qualifiers. Click **OK**.
6. To apply a custom template to a slot or pseudo-slot, select it and click **Apply**.

NOTE: For GigaVUE-HC1, you can apply a filter template only at the control card level which will be applied across all the line cards

7. Select the slot or pseudo-slot and click **OK**.

The Filter Templates page displays the applied slot or pseudo-slot. You can edit an existing custom filter or delete it. A template can be deleted if it is not currently in use, meaning that it has not been applied.

8. To display filter templates, click on a row in the Filter Templates page.

Filter Template Limits

The number of qualifiers in a template limits the total number of rules that can be defined. The maximum rule limit on the GigaVUE-HC3, GigaVUE-TA100, or GigaVUE-TA200 is 1K (1024) per slot or pseudo-slot when using the default templates.

Custom templates allow the creation of templates with only those qualifiers needed for the rules that you plan to use in flow maps. The qualifiers specified in a flexible template can increase or decrease the maximum rule limit, depending on the qualifiers selected. With flexible filter templates, it is possible to reach a maximum limit of 6K rules per slot on the GigaVUE-HC3 node and 6K rules per pseudo-slot, or 24K total rules on the GigaVUE-TA100, GigaVUE-TA200 and GigaVUE-TA200E node, 18k rules per node on GigaVUE-TA25 node, 36k rules per node on GigaVUE-TA25E node, on GigaVUE-TA400, and GigaVUE-TA400E node, 22K rules per pseudo-slots.

How to Understand Map Filter Resources

Starting in software version 5.0, when a filter template is applied, filter resources display the total number of map rules used in a map as well as the limit. If the limit is 1024, 1023 is displayed, even though the actual limit is 1022, or two less than the limit. This discrepancy is

due to extra resources needed for internal usage.

Filter Template Rules and Recommendations

When creating flexible filter templates, keep the following rules and recommendations in mind:

- Filters are applied to a specific slot or pseudo-slot, not to the node.
- By default, all slots will be in the pre-defined **ipv4** template.
- There is a limit of 512 custom templates.
- Custom templates can have duplicate sets of qualifiers.
- The filter limit is calculated when the template is created. In most cases, a higher-cost qualifier set (for example, IPv6, UDA, or MAC are higher cost) consumes more resources and leads to a lower filter limit.
- Flexible filter templates have no effect on existing flow mapping behavior, including pass versus drop map rules, map priority, network port sharing, GigaSMART operations, or first level and second level maps.
- When deploying a Resilient Inline Arrangements (RIA) map in IPv6, specifying MAC source address, MAC destination address and EtherType as qualifiers will not be accepted.
- When configuring filter templates, certain combinations of qualifiers are not supported on some of the platforms even though the total bits consumed by the qualifiers is less than (480-54) bits. This is due to the limitations in the hardware. For example, a flexible filter template configured with qualifiers 'ipsrc ipdst portsrc portdst uda1 uda2' is supported on GigaVUE-HC3.
- To use Flexible Filter Template along with inline-netlag, both VLAN and inner VLAN must be added in addition to other needed qualifiers.

NOTE: To verify if a flexible filter template is supported on a specific platform and the number of rules supported, you must create the template with the desired qualifiers and execute the `show filter-template limit` command. If the number of rules is N/A, then it indicates that this combination of qualifiers is not supported on the corresponding platform.

Filter Template Best Practices

The following are best practices for optimizing filter resources using filter templates.

First determine all the needed qualifiers, then create a template, apply the template, and configure the map rules.

Example 1 - Applicable for the all the platforms

- Connect network ports of a slot to flows of the same application. For example, if you have two flows:
 - one is filtered on **macsrc** and **macdst**

- the second one is filtered on **ipdst** and **ipsrc**
- In case both flows connect to ports on the same slot, that slot will have to have a template of **macsrc**, **macdst**, **ipsrc**, and **ipdst**, with a limit of 1024 rules.
- However, filter resources can be optimized by connecting these two flows to ports on different slots with one template for **macsrc** and **macdst** and the other template for **ipsrc** and **ipdst**. Both templates will have a limit of 3072 rules.

Example 2 - Applicable for the GigaVUE-TA400 platform

- Connect network ports of a slot to flows of the same application. For example, if you have two flows:
 - one is filtered on **mpls-label-id** and **mpls-label-exp**
 - the second one is filtered on **inner ipsrc** and **inner ipdst**
- In case both flows connect to ports on the same slot, that slot will have to have a template of **mpls-label-id**, **mpls-label-exp**, **inner ipsrc**, and **inner ipdst**.
- However, filter resources can be optimized by connecting these two flows to ports on different slots with one template for **mpls-label-id** and **mpls-label-exp** and the other template for **inner ipsrc** and **inner ipdst**.

Example 3 - Applicable for the GigaVUE-TA400 platform

- Connect network ports of a slot to flows of the same application. For example, if you have two flows:
 - one is filtered on **macsrc** and **macdst**
 - the second one is filtered on **inner ipdst** and **inner ipsrc**
- In case both flows connect to ports on the same slot, that slot will have to have a template of **macsrc**, **macdst**, **inner ipsrc**, and **inner ipdst**.
- However, filter resources can be optimized by connecting these two flows to ports on different slots with one template for **macsrc** and **macdst** and the other template for **inner ipsrc** and **inner ipdst**.

The following are best practices for adding more rules if a limit has been reached:

- Create a new template with all the qualifiers that are in use on a specified slot.
- Issue the **show filter-resource** command to obtain the list of qualifiers in use.
- Issue the **filter-template alias <alias> qualifiers add** command with that list of qualifiers.
- Issue the **show filter-template limit** command to check if the new template allows a higher limit. If it does, apply the filter using the **card slot <slot ID> filter-template** command.

Filter Templates in a Cluster

Filter template configuration is synchronized across the cluster. However, a cluster can have different GigaVUE nodes, so one set of qualifiers may or may not be valid on all nodes.

Filter Templates Formulas

The formulas in this section can help you determine the number of map rules that are supported, based on the qualifiers specified in the filter template. Use the formulas as guidelines.

The number of map rules depends on the number of qualifiers a template can support.

Table 6: Maximum Cost Supported per Map Rule for Different Qualifier Combinations

Qualifier Combination	Maximum cost supported per map rule
Outer Header attribute with IPv6 qualification	10
Outer Header attribute without IPv6 qualification	12
Inner Header and MPLS Header attribute	18

NOTE: Combining outer header attributes with inner or MPLS header attributes in a map rule is supported, and the maximum cost per map rule will increase accordingly.

The cost of each qualifier depends on the number of bits it consumes. The following table lists the number of bits each qualifier consumes and the cost for each qualifier.

Table 7: Bits Consumed and Cost per Qualifier

Qualifier	Bits	Cost
dscp	6	1
ethertype	16	1
ip6src	128	4
ip6dst	128	4
ip4src	32	1
ip4dst	32	1
macdst	48	2
macsrc	48	2
macsrc and macdst	96	3

Qualifier	Bits	Cost
portsrc	16	1
portdst	16	1
portsrc and portdst	32	1
protocol	8	1
tos	8	1
ttl	8	1
vlan	16	1
vxlan	48	1
l2gre	24	1
uda1	128	4
uda2	128	4
inner-vlan	16	1
qset1*	58*	*

* qset1 is made up of the following: tos: 8, ipfrag: 2, tcpctl: 8, ttl: 8, ip6fl: 32. The cost depends on the combination of the qualifiers used.

Table 8: Bits Consumed and Cost per Inner Qualifier and MPLS header for the GigaVUE-TA400

Qualifier	Bits	Cost
inner ethertype	16	1
inner ip6src	128	4
inner ip6dst	128	4
inner ipfrag	32	2
inner ip4dst	32	1
inner ip4src	32	1
inner portdst	16	1
inner portsrc	16	1
inner portdst and inner portsrc	32	1

Qualifier	Bits	Cost
inner protocol	8	1
inner uda1	64	4
inner uda2	64	4
mpls-label-id	20	1 per level
mpls-label-exp	3	1 per level
mpls-label-bos	1	1 per level
mpls-label-ttl	8	1 per level

Refer to Filter Template Limits section to know the maximum rules supported for the selected filter template.

Troubleshoot Custom Filter Template Issues

Starting with release 6.10, certain combinations of custom filter-template configurations assigned to slots on GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-HCT, and GigaVUE-HCI-Plus devices are no longer supported. This change is due to the introduction of new feature sets, which increase the overall size of the filter-template. Specifically, unsupported configurations include custom filter-templates that use the maximum combination of qualifier sets exceeding the supported TCAM width in BCM.

Examples for combinations which have become unsupported: [Any one of existing qualifier can be removed to support this combination]

{SIP6, MACSRC, MACDST, PROTOCOL, L4 PORTDST, UDA1, UDA2}

{SIP6, DIP6, PROTOCOL, L4 PORTSRC, UDA1, UDA2}

{SIP6, DIP6, MACSRC, ETHERTYPE, UDA1, UDA2}

Whenever you configure an unsupported custom filter-template and assign it to a slot, the system will restrict the configuration and display an error message.

Sample CLI Output:

```
HC1P (config) # filter-template alias ft1 qualifiers add ip6src macsrc macdst ipdst
portdst protocol uda1 uda2
HC1P (config) # card slot 1/1PS1 filter-template ft1
% Filter-template 'ft1' is not valid for the specified slot.
```


A CLI command is available to help determine whether a custom filter-template is supported on a GigaVUE chassis without applying it to a slot. A result of “N/A” indicates that the template is unsupported.

```
HC1P (config) # show filter-template limit
Total Filter-Template-Limit Count: 6
```

Template Alias	HC1 CC1	HCT CC1	HC2 CC1	HC2P CC1	----- C16	C08	C05	HC3 C08FGS	X24	X16	X16Q4LR	X16C4LR	12X04	X04G44	HD Q02X32	C02X08	C06X24	Q08	TS
ipv4	16384	3071	4096	16384	1024	1024	1024	1024	1024	1024	1024	1024	2048	2048	2048	2048	16384	2048	2048
ipv6	8192	3071	2048	8192	1024	1024	1024	1024	1024	1024	1024	1024	512	512	512	512	4096	512	512
ft1	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
uda	8192	3071	2048	8192	1024	1024	1024	1024	1024	1024	1024	1024	512	512	512	512	4096	512	512
ipv4+mac	8192	3071	2048	8192	1024	1024	1024	1024	1024	1024	1024	1024	1024	1024	1024	1024	8192	1024	1024
ipv4+uda	8192	3071	2048	8192	1024	1024	1024	1024	1024	1024	1024	1024	1024	1024	1024	1024	8192	1024	1024

Template Alias	-----HC1P----- PS1	PS2	LY2 CC1	TA1 CC1	TA10 CC1	TA40 CC1	TAC CC1	TACK CC1	TA200 CC1	TA25 CC1	DELL S4112F CC1	TA400 CC1	TA200E CC1	TA25E CC1	TA400E CC1
ipv4	3071	3071	2048	2048	2048	4096	1024	1024	1024	3071	8192	10229	1024	3071	10229
ipv6	3071	3071	512	512	512	2048	1024	512	1024	3071	4096	6133	1024	3071	6133
ft1	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	6133	N/A	N/A	6133
uda	3071	3071	512	512	512	2048	1024	512	1024	3071	4096	10229	1024	3071	10229
ipv4+mac	3071	3071	1024	1024	1024	2048	1024	1024	1024	3071	4096	6133	1024	3071	6133
ipv4+uda	3071	3071	1024	1024	1024	2048	1024	1024	1024	3071	4096	10229	1024	3071	10229

Upgrade process (as applicable)

If you were using any unsupported combination of custom filter-templates in your environment prior to release 6.13 and then upgraded to 6.13, the restoration process will complete successfully, but internal errors may be observed. After the upgrade, you must manually verify whether any unsupported custom filter-templates were applied to a slot by running the CLI command below. Check the “Map Rules Limit” value in the output—if it is displayed as “0”, the template is unsupported.

```
HC1P (config) # show filter-resource
```

Box/ Slot	Map Rules Used	Filter Limit	Filter Template	App-filter Used	Filter Limit	Port Filter Used	Filter Limit	In-Use Qualifiers
1/1PS1	9		ft1	2	504	0	448	ipdst, ip6src, udal and uda2
1/1PS2	6	3071	ft2	2	504	0	448	ip6src and macsrc

```
HC1P (config) # show filter-template alias ft1 brief
Total Filter-Template Count: 1
```

Template Alias	Qualifier List
ft1	circuit-id, ip6src, ipdst, macdst, macsrc, portdst, protocol, udal, uda2

Recovery Procedure

You must remove all flow map configurations created on top of the unsupported filter-template, configure the filter-template as “defaults” for the respective slot, and then follow one of the options below:

- Replace the existing custom filter-template qualifier set with a reduced combination, or split it into two different custom filter-templates that are supported on the chassis. Then associate the supported filter-template to the slot and manually reconfigure all flow map configurations.
- Migrate the unsupported custom filter-template to a different chassis such as GigaVUE-HC3, GigaVUE-TA200, and GigaVUE-TA400, which may support the combination (verify using the `show filter-template limit` CLI command). After migration, reconfigure the flow map configurations manually. For example,


```
HC1P (config) # filter-template alias ft1 qualifiers add ip6src macsrc macdst
ipdst portdst protocol uda1 uda2
HC1P (config) # card slot 1/1PS1 filter-template ft1
% Filter-template 'ft1' is not valid for the specified slot.
```

As shown in the example above, filter-template ft1 is unsupported on the GigaVUE-HC1-Plus chassis but is supported on GigaVUE-TA400/GigaVUE-TA400E.

Review Map Statistics with Map Rule Counters

Map Statistics can be viewed in the following ways:

- The Statistics page.
For details, refer to [Viewing Map Statistics with the Statistics Page](#)
- The Map Quick View.
For details, refer to [Viewing Map Statistics with Quick View](#)

A single packet may match multiple rules in the map and will not cause multiple rule counters to increment. Only the last rule which is the highest priority in the order will increment. The flow map rule priority is based on the order it was created. Thus, the sum of the map rule counters across all the map rules may be higher than the total number of packets received and transmitted by the map.

NOTE: Drop rules have a higher priority than pass rules.

For example, consider the following three map rules:

- **Rule 1** – ipsrc 10.10.0.0 /24 bidir
- **Rule 2** – ipsrc 10.10.0.100 /32 bidir
- **Rule 3** – portsrc 80

A packet with ipsrc 10.10.0.100 and portsrc 20 will match Rule 1 and Rule 2, which will be forwarded to the tool port or ports. The counters for Rule 2 will only be incremented.

There are several reasons a map rule counter may not increment:

- Traffic matching the rule is not currently present in the production network.
- The network port is not monitoring the network at the proper location to see the traffic specified by the map.
- A higher-priority map is forwarding the packet before it can be inspected by this particular map.
- The map rule itself may be specified incorrectly.

Viewing Map Statistics with the Statistics Page

To review map statistics indicating the total packets and total octets handled by maps and the number of rules in the map, select **Maps > Maps > Statistics**. The Statistics page displays the map statistics in a table format listing the maps by their alias. Clicking on a map alias opens a Quick View for that map. To clear map counters, click the **Clear** button.

Viewing Map Statistics with Quick View

To review map rule counters indicating the number of rule matches for a map as packets are inspected and forwarded, select the map and view the information in the Quick View window.

Flow Mapping® FAQ

This section answers frequently asked questions by users migrating to the Flow Mapping® model.

How Many Map Rules are Supported?

The maximum number of map rules supported per line card or standalone node are shown in [Maximum Number of Nodes and Map Rules Supported in a Cluster](#).

Table 9: Maximum Map Rules

Node Type	Maximum Combined Rules for Default Template Without a License	Maximum with Flexible Filter Templates Without a License	Maximum with Advanced Features License for Flexible Filter Templates
GigaVUE-HCT Node	6k (6142)	6k (6142)	N/A
GigaVUE-HCT Module	3k (3071)	3k (3071)	N/A
GigaVUE-HC1-Plus Node	6k (6142)	36k (36862)	N/A
GigaVUE-HC3 Node	4K (4096)	24K (24576)	N/A
GigaVUE-HC3 Module	1K (1024)	6K (6144)	N/A
GigaVUE-HC1	16K (16383)	16K (16383)	N/A
GigaVUE TA Series	256	The Flexible Filter Templates are available for the following TA platforms: GigaVUE-TA100 GigaVUE-TA200 GigaVUE-TA25 GigaVUE-TA400 GigaVUE-TA200E GigaVUE-TA25E	The Advanced Features License is available only for the GigaVUE-TA Series. The extended max rule limit varies based on the GigaVUE-TA Series node.
GigaVUE-TA25	256	256	18k (18431) rules per node
GigaVUE-TA25E	256	256	36k (36862) rules per node
GigaVUE-TA100	256	256	6K (6144) per pseudo-slot with Flexible Filter Template and Advanced Feature License. 1K (1024) per pseudo-slot with only Advanced Feature License.
GigaVUE-TA200	256	256	6K (6144) per pseudo-slot with Flexible Filter Template and Advanced Feature License. 1K (1024) per pseudo-slot with only Advanced Feature License.
GigaVUE-TA200E	256	256	6K (6144) per pseudo-slot with Flexible Filter Template and Advanced Feature License. 1K (1024) per pseudo-slot

Node Type	Maximum Combined Rules for Default Template Without a License	Maximum with Flexible Filter Templates Without a License	Maximum with Advanced Features License for Flexible Filter Templates
			with only Advanced Feature License.
GigaVUE-TA400	256	256	22k map rules (22517) per pseudo-slot (or) 88k map rules in total for 4 available pseudo slots of TA400 with Flexible Filter Template (depends on the qualifier list. Reduces when more qualifiers included) and Advanced Feature License. 10K map rules (10299) per pseudo-slot (or) 40k map rules in total for 4 available pseudo slots of TA400 with only Advanced Feature License.
GigaVUE-TA400E	256	256	22k map rules (22517) per pseudo-slot (or) 88k map rules in total for 4 available pseudo slots of GigaVUE-TA400E with Flexible Filter Template (depends on the qualifier list. Reduces when more qualifiers included) and Advanced Feature License. 10K map rules (10299) per pseudo slot (or) 40k map rules in total for 4 available pseudo slots of GigaVUE-TA400E with only Advanced Feature License.

GigaVUE-TA400: The reference logs as below:

With Pseudo slot:

ST-32N-HC3v2-5-22 [32N-LSCLUSTER999999999: leader] (config) # show filter-resource box-id 38

Slot : 38/1PS1

Filter-Template : portdst (portdst)

In-Use Qualifiers : NONE

Map Rules Usage (used/limit) : 0/22517

Port Filter Usage (used/limit) : 0/950

App-filter usage (used/limit) : 0/1016

Slot : 38/1PS2

Filter-Template : portdst (portdst)

In-Use Qualifiers : NONE

Map Rules Usage (used/limit) : 0/22517

Port Filter Usage (used/limit) : 0/950

App-filter usage (used/limit) : 0/1016

Slot : 38/1PS3

Filter-Template : portdst (portdst)

In-Use Qualifiers : NONE

Map Rules Usage (used/limit) : 0/22517

Port Filter Usage (used/limit) : 0/950

App-filter usage (used/limit) : 0/1016

Slot : 38/1PS4

Filter-Template : portdst (portdst)

In-Use Qualifiers : NONE

Map Rules Usage (used/limit) : 0/22517

Port Filter Usage (used/limit) : 0/950

App-filter usage (used/limit) : 0/1016

ST-32N-HC3v2-5-22 [32N-LSCLUSTER999999999: leader] (config) #

Without Pseudo Slots:

ST-32N-HC3v2-5-22 [32N-LSCLUSTER999999999: leader] (config) # show filter-resource box-id 38

Slot : 38/1PS1

Filter-Template : NONE (NONE)

In-Use Qualifiers : NONE

Map Rules Usage (used/limit) : 0/10229

Port Filter Usage (used/limit) : 0/950

App-filter usage (used/limit) : 0/1016

Slot : 38/1PS2

Filter-Template : NONE (NONE)

In-Use Qualifiers : NONE

Map Rules Usage (used/limit) : 0/10229

Port Filter Usage (used/limit) : 0/950

App-filter usage (used/limit) : 0/1016

Slot : 38/1PS3

Filter-Template : NONE (NONE)

In-Use Qualifiers : NONE

Map Rules Usage (used/limit) : 0/10229

Port Filter Usage (used/limit) : 0/950

App-filter usage (used/limit) : 0/1016

Slot : 38/1PS4

Filter-Template : NONE (NONE)

In-Use Qualifiers : NONE

Map Rules Usage (used/limit) : 0/10229

Port Filter Usage (used/limit) : 0/950

App-filter usage (used/limit) : 0/1016

ST-32N-HC3v2-5-22 [32N-LSCLUSTER9999999999: leader] (config) #

The limit for GigaVUE TA Series standalone nodes is 256 combined pass/drop rules but is up to 2048 with the Advanced Features License installed.

Refer to [Manage Map Rule Resources](#) for managing map rules.

The maximum number of nodes and map rules supported when in a cluster is as follows:

Table 10: Maximum Number of Nodes and Map Rules Supported in a Cluster

When a Cluster is Configured with:	Number of Nodes	Maximum Map Rules
Out-of-Band Cluster Management	32	38000
Inband Cluster Management	16	38000

The maximum number of map rules supported in a cluster apply to all nodes in the cluster including GigaVUE HC Series nodes: GigaVUE-HC3, GigaVUE-HC1, GigaVUE-HC1-Plus, GigaVUE-HCT and GigaVUE TA Series nodes: GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-TA100, GigaVUE-TA200, GigaVUE-TA200E, and GigaVUE-TA400 including Certified Traffic Aggregation White Box (white box).

How Many Rules Can Each Map Have?

The maximum number of rules per map is 4K (4096), except on products that only support a total of 2K map rules. GigaVUE-HC3 supports 6K rules per map. Refer to [Rules Per Map](#).

Table 11: Rules Per Map

Node Type	Rules per map
GigaVUE-HCT Node	4K
GigaVUE-HCT Module	4K
GigaVUE-HC1-Plus Node	4K
GigaVUE-HC1-Plus Module	4K
GigaVUE-HC3 Node	6K
GigaVUE-HC3 Module	6K
GigaVUE-HC1	4K
GigaVUE TA Series	4K
GigaVUE-TA25	4K
GigaVUE-TA25E	4K
GigaVUE-TA100	4K

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

How Many Maps Can Run at Once?

A maximum of 16000 maps can be configured on a cluster or standalone device.

What Criteria can be Filtered in Q-in-Q Packets?

Maps on GigaVUE nodes can match Layer 3/Layer 4 criteria in packets using Q-in-Q with up to two tags. For more information refer to [How to Handle Q-in-Q Packets in Maps](#).

How Many Maps Can Share a Network Port?

A maximum of 512 maps can share the same source port.

How Many Network Ports and Tool Ports Can Be in a Map?

If the ports are not in a GigaStream, the number of individual map ports in the **Source** or **Destination** field of a map is limited to 64 on all GigaVUE HC Series nodes. The individual ports can be any of the following port types: network or tool.

On GigaVUE-HC3, if the ports are in a GigaStream, the limit is 95.

Are Port-Filters Supported?

Yes.

Egress port-filters (tool, hybrid, circuit, and inline network), are less efficient and scalable than flow maps, but they do provide a convenient way to narrow down the traffic seen by the tools/GigaVUE® HC Series nodes without having to change an entire map. Refer to [Port Filters](#) for details.

Each GigaVUE-HC3 module, or GigaVUE-HB1 node supports 100 combined egress port-filters. A single filter applied to multiple tool ports counts multiple times against the 100-filter limit.

In the the GigaVUE-HC1 or GigaVUE-HC3 node, the limit is 400 filters.

The GigaVUE TA Series can only support 20 egress port-filters. When the GigaVUE-TA100 or GigaVUE-TA200 , GigaVUE-TA200E are in a cluster, they can support 400 filters. GigaVUE-TA25, GigaVUE-TA25E supports 20 port filters by default and up to 100 port filters with advanced feature license.

Does Flow Mapping® Support Passalls?

Yes.

Flow Mapping® supports passalls with the following:

- The map **Subtype Pass All** option for network to tool port passalls.
- A tool-mirror connection between to tool ports. The Tool Mirrors page replaces the tool-to-tool port passalls. To create a Tool Mirror, select **Ports > Tool Mirrors** to go to the Tool Mirror page shown in [Figure 10Tool Mirror Page](#).



The screenshot shows a form for creating a Tool Mirror. It includes four fields: 'Alias' (a text input), 'Comment' (a text input), 'Source Tool Ports' (a dropdown menu with 'Click to select' text), and 'Mirrored Destination Ports' (a dropdown menu with 'Click to select' text).

Figure 10 Tool Mirror Page

Does Flow Mapping® Support port-pairs?

Yes.

Select **Ports > Port Pairs** to go to the Port Pairs page shown in [Figure 11Creating Port Pairs](#). For more information about port pairs and details on configuring two ports as an inline TAP, refer to [Port Pairs](#).

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

Figure 11 Creating Port Pairs

Does Flow Mapping® Support UDA Pattern Matches?

Yes.

Pattern match rules are still supported in map rules. They can be used in both pass rules and drop rules. Refer to [Work with User-Defined Pattern Match Rules](#) for details.

Are Maps Supported Across Nodes in a Cluster?

Yes.

Clusters of GigaVUE nodes operate as a unified fabric. Use the standard box ID/slot ID/port ID syntax to create packet distribution, just as on a standalone node. Maps can have network ports and tool ports on different physical nodes within the cluster.

Similarly, a map does not need to keep its network and tool ports on the same physical node in order to take advantage of GigaSMART operations – a GigaVUE TA Series node in a cluster can take advantage of the GigaSMART processing available on a GigaVUE-HC3 module, or GigaVUE-HC1 node elsewhere in the same cluster.

Can a GigaStream Act as a Shared Collector?

Yes.

Multiple individual ports for a Shared Collector can be setup or a GigaStream tool group. Refer to [About Shared Collectors](#) for details.

What Are the GigaStream Maximums?

The number of GigaStream per line card, module, or node varies by product. Refer to [GigaStream Rules and Maximums](#) for the details.

Does GigaVUE-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 events and take actions in response to events in your visibility network.

An Active Visibility policy defines conditions and actions. When conditions are met, actions are executed. The policy specifies both the conditions and the actions and ties them together.

The conditions and actions are pre-defined. Conditions are events that are used to trigger changes to configuration. Actions can notify users of certain events and/or modify the configuration in response to certain events. Conditions can be port-based or time-based.

For example, if a tool port is overloaded, you might want to reduce the traffic sent to the tool. You can configure two maps, one targeted for your high priority traffic and the other targeted for your low priority traffic. Then you can specify conditions to monitor tool port utilization. When traffic is below a threshold, both maps can be enabled, thus all traffic will be sent to the tools. But when traffic is above the threshold, you can specify an action so that only the map targeted for the high priority traffic is enabled and thus only that traffic will be sent to the tool.

Another example is that you might want to use a different set of maps and map rules to provide visibility during different times, such as during working hours, or on weekends. If there is going to be a backup on the weekend, you can specify a policy to disable a map at a specific time or day.

NOTE: GigaVUE-FM GUI for Active Visibility is available only until software version 5.4.00. Refer to the *"Configuring Active Visibility"* section in the *GigaVUE-OS CLI Reference Guide* for details on further software enhancements of this feature.

Active Visibility—Rules and Notes

Keep in mind the following rules and notes when working with Active Visibility:

- A policy must have at least one condition. An action is not required.
- Up to five (5) conditions can be specified in a policy. The parameters and values that are specified in the condition depends on the condition. The policy is executed only when all conditions are met.
- Up to five (5) actions can be specified in a policy. The parameters and values that are specified in the action depends on the action.
- Within a policy, there is one unique condition. For example, there can only be one **PortUp** condition, not multiple **PortUp** conditions within a policy.
- Within a policy, there can be multiple actions, including the same action. For example, there can be multiple **PortEnable** actions within a policy.
- A policy is triggered if all the conditions are met. Then all the actions are executed. If there are five conditions in the policy, they all have to be met before the action or actions are executed. If there are multiple actions, they are executed in sequence, as specified in the policy. The policy specifies the priority of the actions.
- When there are multiple actions, they will continue to be executed even if there is an error. For example, if there are three actions and there is an error with the second action, the first and third actions will be executed. When the policy executes, each action will have their status reported (success or failure).
- Multiple policies can be in effect at the same time. The policies can monitor different conditions and take different actions. Up to 100 policies can be defined for a cluster.
- A policy must be enabled for it to become active. However, when you first create a policy, you might want to disable it so that it does not become active right away.
- When a policy is triggered, an SNMP event and email notification (policytrigger) can optionally be generated.
- When you want to delete a map that is attached to a policy, you must first delete the policy and then delete the map. If you delete the map without deleting the policy, there will be configuration synchronization issues in GigaVUE-FM. In such a case, delete the policy using GigaVUE-OS CLI. Refer to the *"policy"* command in the *GigaVUE-OS CLI Reference Guide*.

Configure Policies

To configure a policy:

1. Go to **Traffic** 
2. On the left navigation pane, go to **Physical > Active Visibility**, and then select the required cluster or node ID.
3. In the Policies page, click **New**.
4. Enter an alias and description for the policy.

5. When you first define a policy, you will probably want the policy to be disabled, so clear the **Enable** check box.
6. Select the required condition(s) from the drop-down list. The parameters and values are populated based on the condition(s) you select. For example, if you select the condition, PortTxUtilLow, there is an **Any** checkbox and fields for **Port ID**, **Threshold (%)**, and **Period (seconds)**. To view the list of pre-defined conditions along with their descriptions and comments, refer to [View List of Conditions and Actions](#).
7. Select the required action(s) from the drop-down list. The parameters and values are populated based on the action(s) you select. For example, if you select the action, Port Disable, there is a **Porting** checkbox and a field for **Port ID**. To view the list of pre-defined actions along with their descriptions and comments, refer to [View List of Conditions and Actions](#).
8. When you have added the required conditions and actions, click **OK**.

OK
Cancel

Policy

▼ Policy Info

Alias * Policy1

Enable ☐

Description

▼ Conditions

+
-

Port Down

Any ☒ i

Port ID *

N 1/1/x1

N 1/1/x3

N 1/1/x5

i

Period (seconds) 3 i

▼ Actions

+
-

Port Disable

Porting ☐ i

Port ID *

N 1/1/x1

N 1/1/x3

N 1/1/x5

i

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

Following are the additional actions that you can perform in the Policies page:

- Enable a policy—When you are ready to enable a policy, select the policy, and then go to **Actions > Enable**.
- Edit a policy—Select the policy and then click **Edit**. Make the required changes, and click **OK**.
- Clone a policy—Select the policy and click **Clone**. Change the policy alias, make any changes to conditions or actions, and click **OK**.
- Delete a policy—Select the policy and click **Delete**. A confirmation displays. To confirm, click **OK**.
- View policy details—Click the policy alias to view the details.
- View policy report—To view information on a policy that has been executed, click **Report History**. The policy report history is displayed.

View List of Conditions and Actions

To view the list of pre-defined conditions and actions along with their descriptions and comments:

1. Go to **Traffic** 
2. On the left navigation pane, go to **Physical > Active Visibility**, and then select the required cluster or node ID. You can view the policies details.
3. On the top, select **Help > Conditions** to view the list of pre-defined conditions. Select **Help > Actions** to view the actions that need to be performed.

MAC Address Rewrite

Media Access Control (MAC) address rewrite converts the incoming traffic's MAC address (source, destination, or both) with a user configured MAC address. The modified packets are then delivered as per flow mapping configurations. This allows the user to maintain confidentiality of the outgoing MAC address.

MAC address rewrite can be enabled in two ways:

- **Rule based**- The MAC address rewrite functionality is enabled for traffic that qualifies a specific rule in a map. This can be enabled only for pass rules. Rule based MAC address re-write allows modifying the rule, source, and destination MAC address.

- **Map Based-** The MAC address rewrite functionality is enabled for traffic that qualifies any of the rules configured in regular by-Rule maps and shared collectors. The configuration applies to all the rules that are part of the map except for drop rules. Map based MAC address re-write allows modifying the source and destination MAC address and can also be applied to a deployed map. Refer to [Map MAC Address Source and Destination Compatibility Matrix](#) for more information.

Table 12: Map MAC Address Source and Destination Compatibility Matrix

Source	Destination	Supported
Network	Tool/Hybrid, Tool GigaStream/Hybrid GigaStream, Tool with Egress VLAN strip/Tool with Egress Port filters.	Yes
Network	L2 Circuit Encapsulation Tunneling	No
Hybrid	Tool/Hybrid, Tool GigaStream/Hybrid GigaStream, Tool with Egress VLAN strip/Tool with Egress Port filters.	Yes
Network /Hybrid	Port-group(without smart-lb enabled).	Yes
IP interface (decapsulation Tunnel)	Tool/Hybrid	Yes
L2GRE/VXLAN	L2GRE/VXLAN encapsulation tunnel.	No
VXLAN Header/MPLS Header stripping	Tool/Hybrid	No
Network Port with Ingress VLAN tag	Tool/Hybrid	Yes
L2-Circuit Tunnel	Tool/Hybrid/GigaStream	Yes
VXLAN/L2GRE Tunnel decapsulation with IP interface	Tool/Hybrid/GigaStream	No
Port-Group	Tool/Hybrid/GigaStream	Yes

NOTE: If you have configured both map level and rule level MAC address rewrite functionality in the same map, then rule-based configuration takes priority.

Configuring MAC Address Re-write

Media Access Control address is a six byte hardware identification field with 12 hexadecimal digits that uniquely identifies a device in the network. You can rewrite the MAC source and destination fields to configurable MAC address as follows:

1. To enable MAC address rewrite functionality through GigaVUE-FM:
 - a. **Map based Configuration-** To configure MAC address rewrite based on maps follow the below steps:
 - Navigate to  > **Physical > Nodes.**
 - Select the required cluster or device. Navigate to **Maps** and click create **New Map**. Scroll down to **Map Configuration & Rules.**

- Under **Configuration**, enable the '**Address Rewrite**' checkbox.
 - Select either MAC Source, Mac Destination, or both.
 - Specify the MAC Source and Destination.
 - Click on **OK** to complete the configuration.
- b. **Rule based Configuration**- To configure MAC address rewrite based on map rules follow the below steps

- Navigate to  > **Physical > Nodes**.
 - Select the required cluster or device. Navigate to **Maps** and click create **New Map**. Scroll down to **Map Configuration & Rules**.
 - Under **Map Rules**, click **Add a Rule**.
 - Enable the '**Address Rewrite**' checkbox.
 - From Map Rules section select either MAC Source , Mac Destination, or both.
 - Specify the MAC Source and Destination.
 - Click on **OK** to complete the configuration.
2. To enable MAC address rewrite through GigaVUE-OS -CLI enter the map prefix mode with the command **config map alias<map>** and then enter any one of the following commands such as:

```
rewrite-dstmac xx:xx:xx:xx:xx:xx
rewrite-srcmac xx:xx:xx:xx:xx:xx
no rewrite-dstmac
no rewrite-srcmac
```

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

License

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

Limitations

The following are the limitations of MAC Address rewrite.

- Pass-all maps are not supported.
- GSOP enabled maps are not supported.
- VXLAN/L2GRE decapsulation and encapsulation tunnels are not supported
- Inline, Flex Inline maps and OOB copy maps are not supported.
- First level, second level and transit maps are not supported.

- This feature is not supported with Fabric Maps, L2 Circuit Tunnel Encapsulation, MPLS and VXLAN Header Stripping enabled-port configurations.
- A paired port receives rewritten mac address when creating a port-pair with a network port in map/rule-based mac-rewrite byrule map.

IP Address Rewrite

Internet Protocol (IP) address rewrite converts the incoming traffic's IP address (source, destination, or both) with a user configured IP address. The modified packets are then delivered as per flow mapping configurations. This allows the user to maintain the confidentiality of the outgoing IP address.

IP address rewrite can be enabled in two ways:

- **Rule based-** The IP address rewrite functionality is enabled for traffic that qualifies a specific rule in a map. This can be enabled only for pass rules. Rule based IP address re-write allows modifying the rule, source, and destination IP address.
- **Map Based-** The IP address rewrite functionality is enabled for traffic that qualifies any of the rules configured in regular by-Rule maps and shared collectors. The configuration applies to all the rules that are part of the map except for drop rules. Map based IP address re-write allows modifying the source and destination IP address and can also be applied to a deployed map. Refer to [Map IP Address Source and Destination Compatibility Matrix](#) for more information.

Table 13: Map IP Address Source and Destination Compatibility Matrix

Source	Destination	Supported
Network	Tool/Hybrid, Tool GigaStream/Hybrid GigaStream, Tool with Egress VLAN strip/Tool with Egress Port filters.	Yes
Hybrid	Tool/Hybrid, Tool GigaStream/Hybrid GigaStream, Tool with Egress VLAN strip/Tool with Egress Port filters.	Yes
Network with L2GRE/VXLAN enabled	L2GRE/VXLAN Encapsulation Tunnel	No
Network with VXLAN Header Strip enabled Port/MPLS Header Strip	Tool/Hybrid	No
Network Port with Ingress VLAN tag	Tool/Hybrid	Yes
Network	L2 Circuit Encapsulation Tunneling	No
L2-Circuit Tunnel	Tool/Hybrid/GigaStream	Yes
VXLAN/L2GRE Tunnel decapsulation with IP interface	Tool/Hybrid/GigaStream	No
Port-Group	Tool/Hybrid/GigaStream	Yes
Network /Hybrid	Port-group (without smart-lb enabled).	Yes

NOTE: If you have configured both map level and rule level IP address rewrite functionality in the same map, then rule-based configuration takes priority.

Configuring IP Address Re-write

Internet Protocol address is a four-byte hardware identification field with 8 hexadecimal digits that uniquely identifies a device in the network. You can rewrite the IP source and destination fields to configurable IP address as follows:

1. To enable IP address rewrite functionality through GigaVUE-FM:
 - a. **Map based Configuration-** To configure IP address rewrite based on maps follow the below steps:
 - Navigate to  > **Physical > Nodes.**
 - Select the required cluster or device. Navigate to **Maps** and click create **New Map**. Scroll down to **Map Configuration & Rules.**
 - Under **Configuration**, enable the '**Address Rewrite**' checkbox.
 - Select either IPv4 Source, IPv4 Destination, or both.
 - Specify the IPv4 Source and Destination.
 - Click on **OK** to complete the configuration.
 - b. **Rule based Configuration-** To configure IP address rewrite based on map rules follow the below steps
 - Navigate to  > **Physical > Nodes.**
 - Select the required cluster or device. Navigate to **Maps** and click create **New Map**. Scroll down to **Map Configuration & Rules.**
 - Under **Map Rules**, click **Add a Rule.**
 - Select the IPV4 address from the Address Rewrite drop down list.
 - From Map Rules section select either IPv4 Source , IPv4 Destination, or both.
 - Specify the IPv4 Source and Destination.
 - Click on **OK** to complete the configuration.
2. To enable IP address rewrite through GigaVUE-OS -CLI enter the map prefix mode with the command **config map alias<map>** and then enter any one of the following commands such as:

```
rewrite-dstip x.x.x.x
rewrite-srcip x.x.x.x
no rewrite-dstip
no rewrite-srcip
```

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

License

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

Limitations

The following are the limitations of IP Address rewrite.

- Pass-all maps are not supported.
- GSOP enabled maps are not supported.
- VXLAN/L2GRE Encapsulation and decapsulation tunnels are not supported
- Inline, Flex Inline maps and OOB copy maps are not supported.
- First-level, second-level, and transit maps are not supported.
- This feature is not supported with Fabric Maps, L2 Circuit Tunnel Encapsulation, MPLS, and VXLAN Header Stripping enabled-port configurations.
- A paired port receives rewritten IP address when creating a port-pair with a network port in map/rule-based ip-rewrite byrule map.
- IP rewrite is not supported in GigaVUE-TA400 and GigaVUE-HCT devices.
- IP rewrite with single tagged traffic and ingress VLAN tag is not supported in GigaVUE-HC1-Plus, GigaVUE-TA25, and GigaVUE-TA25E.
- IP rewrite with IPv6 and L3-mpls traffic is not supported.
- Cluster and double-tagged traffic are not supported.
- When a pass-all map shares the same source ports with a by-rule map with IP rewrite enabled, the IP rewrite effect will be applied to the destination port of both the pass-all map and the by-rule map.

VLAN Manipulation

Required License:

- GigaVUE HC Series - Base License
- GigaVUE TA Series - Advanced Feature License

Overview


The VLAN manipulation adds a new VLAN tag, modifies/strips the existing outer VLAN tag based on the user configured VLAN manipulation operation. The modified packets are then delivered according to the Flow Mapping® configurations. This process ensures that the confidentiality of outgoing traffic is maintained while allowing efficient traffic segmentation.

NOTE: From 6.8 Release, VLAN manipulation add is supported in standalone mode. From 6.10 Release, VLAN manipulation add is supported in Cluster mode.

NOTE: Starting from 6.13 Release, VLAN manipulation Modify, VLAN manipulation Strip is supported only in standalone mode. It is not supported in cluster mode.

VLAN manipulation can be configured based on map rules and maps:

- Rule-Based-** In this method, the VLAN manipulation is configured for traffic that matches a specific rule on a map. This can be configured only for pass rules. Rule-based VLAN manipulation adds a new VLAN tag, Modifies/Strips the existing outer VLAN tag based on the user configured VLAN manipulation operation. The figure below illustrates how the Outer VLAN (O-VLAN) can be added to incoming traffic using the VLAN Add operation.



Rule-Level Add VLAN Manipulation

To add VLAN manipulation on rule-level,

1. On the left navigation pane, go to  > **Physical** > **Nodes**.
2. Select the required cluster or device.
3. Go to **Maps** and click **New Map** to create the new map.

NOTE: By default, a Regular ByRule Map will be created.

4. In the **Map Info** dialog, do the following:
 - **Map Alias** - Enter the alias for the new map.
 - (Optional) **Description** - Enter the description for the new map.
 - **Subtype** - Select the **By Rule** from the list.
5. In the **Map Source and Destination** section, do the following:
 - Select the **Source** port from the list.
 - (Optional) Select the **destination** port from the list.
6. In the **Map Configuration & Rules** section, go to the Map Rules section do the following:
 - Click **Add a Rule**.
 - (Optional) **Rule Description** - Enter a rule description.
 - From the **VLAN Action** field, select **Add**.
 - In the **VLAN ID** field, enter a **VLAN ID** value between 1 and 4095.
 - From the **Tag Protocol Id** drop-down list, select the **TPID** value for the VLAN Tag. The default value is 0x8100, but you can also select the other supported values 0x9100 and 0x88a8 from the drop-down list.
 - **Condition** - Select the **condition** from list, and choose **Pass**. Without a condition, the rule will throw an error.
7. Click **OK** to complete the configuration.

To configure VLAN manipulation using GigaVUE-OS CLI, refer to the “[Configure VLAN Manipulation](#)” section in the GigaVUE-OS CLI Reference Guide.

Rule-Level Modify VLAN Manipulation

To modify VLAN manipulation on rule-level,

1. On the left navigation pane, go to  > **Physical** > **Nodes**.
2. Select the required cluster or device.
3. Go to **Maps** and click **New Map** to create the new map.
4. In the **Map Info** dialog, do the following:

- **Map Alias** - Enter the alias for the new map.
 - (Optional) **Description** - Enter the description for the new map.
 - **Subtype** - Select the **By Rule** from the list.
5. In the **Map Source and Destination** section, do the following:
 - Select the **Source** port from the list.
 - (Optional) Select the **destination** port from the list.
 6. In the **Map Configuration & Rules** section, go to the Map Rules section do the following:
 - Click **Add a Rule**.
 - (Optional) **Rule Description** - Enter a rule description.
 - From the **VLAN Action** field, select **Modify**.
 - In the **VLAN ID** field, enter a **VLAN ID** value between 1 and 4095.
 - From the **Tag Protocol Id** drop-down list, select the **TPID** value for the VLAN Tag. The default value is 0x8100, but you can also select the other supported values 0x9100 and 0x88a8 from the drop-down list.
 - **Condition** - Select the **condition** from list, and choose **Pass**. Without a condition, the rule will throw an error.
 7. Click **OK**.

Rule-Level Strip VLAN Manipulation

To delete VLAN manipulation on rule-level,

1. On the left navigation pane, go to  **> Physical > Nodes**.
2. Select the required cluster or device.
3. Go to **Maps** and click **New Map** to create the new map.
4. In the **Map Info** dialog, do the following:
 - **Map Alias** - Enter the alias for the new map.
 - (Optional) **Description** - Enter the description for the new map.
 - **Subtype** - Select the **By Rule** from the list.
5. In the **Map Source and Destination** section, do the following:
 - Select the **Source** port from the list.
 - (Optional) Select the **destination** port from the list.
6. In the **Map Configuration & Rules** section, go to the **Map Rules** section do the following:

- Click **Add a Rule**.
- (Optional) **Rule Description** - Enter a rule description.
- From the **VLAN Action** field, select **Strip**.
- **Condition** - Select the **condition** from list, and choose **Pass**. Without a condition, the rule will throw an error.

7. Click **OK** to complete the configuration.


- **Map-Based**- In this method, VLAN manipulation is configured for traffic that qualifies under rules defined in regular by-rule maps and shared collectors. This configuration applies to all rules within the map except for drop rules. Map-based VLAN manipulation adds a new VLAN tag, Modifies/Strips the existing outer VLAN tag based on the user configured VLAN manipulation operation. For more information, refer to the [Map VLAN manipulation Source and Destination Compatibility Matrix](#) .

NOTE: If you have configured both map level and rule level VLAN manipulation functionality in the same map, then rule-based configuration takes priority.


Map-Level VLAN Manipulation

▪ Map-Level Add VLAN Manipulation


The figure below illustrates how the Outer VLAN (O-VLAN) can be added to the rules of the incoming traffic using the VLAN Add operation within the map-based configuration.



▪ Map-Level Modify VLAN Manipulation



Map-Level Strip VLAN Manipulation



To configure VLAN manipulation on map-level,

1. On the left navigation pane, go to  > **Physical > Nodes**.
2. Select the required cluster or device.
3. Go to **Maps** and click **New Map** to create a new map.

NOTE: By default, a Regular ByRule Map will be created.

4. In the **Map Info** dialog, do the following:
 - **Map Alias** - Enter the alias for the new map.

- (Optional) **Description** - Enter the description for the new map.
 - **Subtype** - Select the **By Rule** from the list. By default, the subtype will be By Rule.
5. In the **Map Source and Destination** section, do the following:
 - Select the **Source** port from the list.
 - (Optional) Select the **Destination** port from the list.
 6. In the **Map Configuration & Rules** section, go to the **Configurations** section do the following:

Add:

 - From the **VLAN Action** field, select **Add**.
 - In the **VLAN ID** field, enter a **VLAN ID** value between 1 and 4095.
 - From the **Tag Protocol Id** drop-down list, select the **TPID** value for the VLAN Tag. The default value is 0x8100, but you can also select the other supported values 0x9100 and 0x88a8 from the drop-down list.

Modify:

 - From the **VLAN Action** field, select **Modify**.
 - In the **VLAN ID** field, enter a **VLAN ID** value between 1 and 4095.
 - From the **Tag Protocol Id** drop-down list, select the **TPID** value for the VLAN Tag. The default value is 0x8100, but you can also select the other supported values 0x9100 and 0x88a8 from the drop-down list.

Strip:

 - Select **Strip**, from the **VLAN Action** list.

NOTE: For Strip, there is no need for a VLAN ID, and TPID is not required.
 7. Once you add the map rule configuration, do the following:
 - Click **Add a Rule**.
 - In the **Condition**, select **IPv4 Source** from the list.
 - Enter the **IPv4 Address**, and **Cidr**, or **Net Mask** values.
 8. Click **OK** to complete the map configuration.

Table 14: Map VLAN manipulation Source and Destination Compatibility Matrix

Source	Destination	Supported
Network	Tool/Hybrid, Tool GigaStream/Hybrid GigaStream, Tool with egress VLAN strip/Tool with egress Port filters.	Yes
Hybrid	Tool/Hybrid, Tool GigaStream/Hybrid GigaStream, Tool with egress VLAN strip/Tool with egress Port filters.	Yes

Source	Destination	Supported
Network/Hybrid Port with ingress VLAN tag	Tool/Hybrid	Yes
Port-Group	Tool/Hybrid/GigaStream	Yes
Network /Hybrid	Port-group (without smart-lb enabled).	Yes

Limitations

The following are the limitations of VLAN manipulation.

- Pass-all maps are not supported.
- GSOP-enabled maps are not supported.
- VXLAN/L2GRE Encapsulation and Decapsulation Tunnels are not supported.
- Inline, Flex Inline maps, and OOB copy maps are not supported.
- First-level, second-level, and transit maps are not supported.
- Fabric Maps, L2 Circuit Encapsulation and Decapsulation Tunnels are not supported.
- MPLS, and VXLAN Header Stripping enabled-port configurations do not support this feature.
- VLAN manipulation is not supported in GigaVUE-HC3 ccv1 device.
- VLAN manipulation with IP rewrite is not supported.
- Port filter with VLAN Qualifier is not supported.
- When VLAN manipulation with ingress VLAN Tag is configured, VLAN Manipulation will take higher precedence.
- When VLAN manipulation with egress VLAN Strip is configured, VLAN Manipulation will take higher precedence.
- If advanced VLAN manipulation is configured on either regular by-rule maps or shared collector maps, then both Passall maps or Port Pair should not use the same network ports as those deployed in regular and collector maps. Similarly, if a collector map or regular map is configured on the same network ports as a Passall map or Port Pair, VLAN manipulation should not be configured. Configuring VLAN manipulation in the above two scenarios may result in traffic discards on the destination ports of the Passall map.

Monitor Port Utilization

This section describes how to monitor port utilization and buffer thresholds on the GigaVUE HC Series and GigaVUE TA Series nodes. Refer to the following sections for details:

- [Port Utilization Availability by Port Type](#)
- [Set Port Utilization Thresholds](#)
- [Configure Alarm Buffer Thresholds](#)

- Set Alarm Buffer Thresholds
- Microburst

Port Utilization Availability by Port Type

You can view port utilization for all network, tool, hybrid, and stack link ports on the GigaVUE HC Series or GigaVUE TA Series nodes. From the device view, go to **Ports > All Ports** and view the Utilization (Tx/Rx) column in the table.

☐	Port Id	Alias	Status	Type	Speed	Admin	Link Status	Transceiver T...	SFP Power	Avg Util Tx/Rx	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...	IN	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.

Alias			
Comment			
Port Role			

Parameters

Admin	☐ Enable	
Type	Network	
Duplex	☐ Full ☐ Half	
Auto Negotiation	☐ Enable	
VLAN Tag		
Egress Vlan Tag	☒ None ☐ Strip	
Force Link Up	☐ Enable	
Receive Only	☒ Enable	
VXLAN ID ⓘ	0 ~ 16777215	0 is disabled
L2GRE ID ⓘ	0 ~ 4294967295	0 is disabled
Header Stripping Protocol	None	

Ports Discovery

Network Discovery ⓘ	☐ Enable	
Discovery Protocols	☐ All ☐ LLDP ☐ CDP	
Gigamon Discovery ⓘ	☐ Enable	

Alarms

Buffer Threshold (%)	Rx	0	Tx	0
Utilization Threshold (%)	High	0	Low	0

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**.


The screenshot shows the configuration page for a port. The 'Parameters' section includes fields for Alias, Comment, and Port Role. The 'Ports Discovery' section includes checkboxes for Network Discovery, Discovery Protocols (All, LLDP, CDP), and Gigamon Discovery. The 'Alarms' section is expanded, showing the Buffer Threshold (%) and Utilization Threshold (%) settings. The Buffer Threshold (%) has Rx and Tx fields, both set to 0. The Utilization Threshold (%) has High and Low fields, both set to 0.

Section	Parameter	Value	Notes
Parameters	Alias		
	Comment		
	Port Role		
	Admin	☐ Enable	
	Type	Network	
	Duplex	☐ Full ☐ Half	
	Auto Negotiation	☐ Enable	
	VLAN Tag		
	Egress Vlan Tag	☒ None ☐ Strip	
	Force Link Up	☐ Enable	
Ports Discovery	Network Discovery	☐ Enable	
	Discovery Protocols	☐ All ☐ LLDP ☐ CDP	
	Gigamon Discovery	☐ Enable	
	Buffer Threshold (%)	Rx: 0, Tx: 0	
	Utilization Threshold (%)	High: 0, Low: 0	

Microburst

A microburst is a very intense traffic transmitted in a short period of time. It is a situation in which a large amount of burst data is received in milliseconds, where the burst data rate is much higher than the egress port's bandwidth.

For example, two 100Gb ingress ports transmit traffic to a 10Gb egress port at an average rate of 1Gbps each. The aggregate bandwidth utilized will be 2Gbps (1Gbps x 2 network ports), which is well within the bandwidth of the 10Gb egress port. However, if there is a sudden spike in traffic from one of the ingress ports, that is, traffic is transmitted at a rate of 3.5Gbps in 4.5 milliseconds (as shown in the figure below), it is referred to as a microburst.



When the traffic is aggregated, packets are received at the egress port, usually from multiple ingress ports, in parallel rather than one after another. This results in queuing of packets. GigaVUE-OS devices have some amount of buffer to queue these packets. Each port accesses this buffer dynamically, sharing this resource. This shared buffer is allocated to individual ports based on threshold values. If a port's shared buffer usage reaches its threshold value, the subsequent packets that the port receives will be dropped by the device due to unavailability of buffer resources.

During microbursts, the port utilization will appear to be low, but packets will be dropped and the `IfOutPktDrops` counter will continue to increment.

Following factors impact the packets drop rate:

- Intensity of the microburst traffic.
- Number of egress ports having microbursts at a given point in time.
- Buffer absorption capability of the GigaVUE-OS device.

Best Practices to Improve Burst Tolerance

Keep in mind the following best practices when you design your topology to improve burst tolerance:


- If you have traffic flowing from a port with higher Gigabit to ports with lower Gigabit, that is, from a 40Gb port to four 10Gb ports, and bursty traffic is expected, one of the port can be internally modified into 4 x 25Gb and made as a hybrid port so the traffic flows from a 40Gb port to a 25Gb port, and then to a 10Gb port. However, you must create more than one map to implement this type of configuration. For more information about Hybrid ports, refer to [Work with Hybrid Ports](#).
- Increase the tool ports bandwidth by adding multiple tool ports to tool GigaStream so that traffic is distributed across multiple ports, and thereby minimizes the risk of microburst. Use the Advanced Hashing feature to select the hashing criteria. For more information, refer to [Advanced Hashing](#).
- Ensure that the egress ports are spread across different logical ports so that the load balancing and buffer utilization is effective. For example, you have a 40Gb egress port that is broken out in to four 10Gb ports, ensure that two of the 10Gb ports are in one logical group and the remaining two 10Gb ports are in another logical group. For details about the logical grouping of ports, refer to [Logical Grouping of Ports in GigaVUE-HC3, GigaVUE-TA100, and GigaVUE-TA200](#).
- Ensure that both ingress and egress ports through which traffic is flowing into and out of the network must be in the same logical group. Following table provides an example of the ingress and egress ports mapping in GigaVUE-TA100. For details about the logical grouping of ports, refer to [Logical Grouping of Ports in GigaVUE-HC3, GigaVUE-TA100, and GigaVUE-TA200](#)

Egress Port	Ingress Port	Logical Group
C1	C5	Group 0
C9	C14	Group 1
C18	C23	Group 2
C25	C32	Group 4


Logical Grouping of Ports in GigaVUE-HC3, GigaVUE-TA100, and GigaVUE-TA200

The ports in GigaVUE-HC3, GigaVUE-TA100, and GigaVUE-TA200 are grouped in to four logical groups—Group 0, Group 1, Group 2, and Group 3. You must ensure that the ingress and egress ports are spread across these logical groups to improve burst tolerance.

- GigaVUE-HC3—The ports in slot 1 are logically grouped into Group 0, ports in slot 2 into Group 1, ports in slot 3 into Group 2, and ports in slot 4 into Group 3.
- GigaVUE-TA100—Following diagram illustrates the logical grouping of ports in GigaVUE-TA100:



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



Legend

- Logical Group 0
- Logical Group 1
- Logical Group 2
- Logical Group 3

Flexible Inline Arrangements

This chapter provides an overview about the flexible inline arrangements, the supported platforms and software versions, the supported and unsupported functionalities, and limitations. It also provides details about how to configure the flexible inline maps and how to visualize the flexible inline arrangements canvas.

IMPORTANT: It is recommended to define inline configurations through GigaVUE-FM.

- If you configure a flexible inline arrangement solution using the GigaVUE-OS CLI, you cannot view or manage it using GigaVUE-FM.
- If you modify a flexible inline arrangement solution using the GigaVUE-OS CLI, you cannot view the changes in GigaVUE-FM.

NOTE: If your nodes and GigaVUE-FM instance are running software version 5.6.xx or earlier, ensure that you add all the nodes to GigaVUE-FM, upgrade the GigaVUE-FM to 5.8.xx, and then upgrade the nodes to 5.8.xx. This is to ensure that the flexible inline maps added to the nodes before the upgrade are visible in the canvas in GigaVUE-FM. Refer to the GigaVUE-OS Upgrade guide for details.

Refer to the following sections for details:

- [Flexible Inline Arrangements](#)
- [Benefits of Flexible Inline Arrangements](#)
- [About Flexible Inline Maps](#)
- [Flexible Inline Arrangements—Rules and Notes](#)
- [Visualize the Flexible Inline Arrangements Canvas](#)
- [Configure Flexible Inline Flows](#)
- [Troubleshoot Flexible Inline Flows](#)
- [Backup and Restore Flexible Inline Flows](#)

NOTE: Flexible inline arrangement is an advanced approach than the (Classic) Inline Bypass Solutions . Classic Inline Bypass functionality remains intact for backwards compatibility. For information on configuring classic Inline Bypass solutions , refer to [Inline Bypass Solutions](#) .

Flexible Inline Arrangements

Flexible inline arrangements are an approach to guide multiple inline traffic flows through a user-defined sequence of inline tools and inline tool groups. It uses the same software constructs as the existing Inline Bypass solution, such as inline network, inline tool, and inline tool group. Flexible inline arrangements support physical protection based on the specialized hardware on BPS modules. It also supports both protected and unprotected inline network links.

Flexible inline arrangements offer an alternative to classic Inline Bypass. Classic Inline Bypass functionality remains intact for backwards compatibility. For information on configuring Inline Bypass solutions (classic), refer to [Inline Bypass Solutions](#).

Using flexible inline maps, traffic from the same inline network can traverse different sequences of inline tools and share tools across traffic flows or with other inline networks.

You can identify specific flows of traffic using Layer 2 to Layer 4 rules, then designate the tools that will inspect the traffic, and specify the order of the tools. For example, you can send Web traffic (defined by L4 port, 80 and/or 443) through a Web Application Firewall (WAF) and an Intrusion Prevention System (IPS), have backup traffic that might bypass inspection, and send all other traffic through the same IPS.

[Figure 12Flexible Inline Arrangements Scenario](#) illustrates a complex inspection scenario that is enabled by flexible inline arrangements.




Figure 12 Flexible Inline Arrangements Scenario

In this example, on the left, there are three network links, all of which share the IPS. The details are as follows:

- At the top of the figure, the yellow line represents 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 12 Flexible Inline Arrangements Scenario](#), the traffic in the reverse direction can have a different order of tools than the west-east traffic.

Refer to the Gigamon Validated Design for more information.

- [Deploying GigaSECURE Inline SSL Solution using Flexible Inline](#)
- [Guiding Relevant Inline Traffic to Tools](#)

Flexible Inline Arrangement vs Inline Bypass Solution

A brief difference between Flexible and Classic Inline Arrangements.

Flexible Inline Arrangement	Inline Bypass Solution (Classic)
Flexible inline arrangements offer flexibility in how traffic is guided through inline tools and the order in which they can be defined to move.	Inline Bypass solutions involve bidirectional traffic between two networks, intercepted by a Gigamon node, and guided through one or more inline tools. The classic Inline Bypass solution allows users to define what traffic is inspected by which inline tools and which should be bypassed and not inspected.
Guides traffic through any arbitrary sequence of inspection tools.	The Gigamon node sends uninspected traffic from one side of the network to an inline tool (such as an IPS), and then sends the inspected traffic to the other side of the network.
Shares inline tools across multiple inline network links and across multiple inline maps.	A sequence of inline tools / inline tool groups used for inspecting traffic on a given inline network / inline network group cannot be reused for inspecting traffic on another inline network / inline network group.
Distributes traffic across multiple tools to meet bandwidth and throughput demands.	The decrypted traffic needs to be received at all the attached inline tools and cannot provide flexibility to guide the traffic to individual tools.

Supported Platforms

Flexible inline arrangements are supported on the following nodes:

- GigaVUE-HC1-Plus
- GigaVUE-HCT
- GigaVUE-HC1
- GigaVUE-HC3(CCV1 and CCv2)
- GigaVUE-TA200
- GigaVUE-TA25
- GigaVUE-TA200E
- GigaVUE-TA25E
- GigaVUE-TA400
- GigaVUE-TA400E

The inline networks, inline tools, and inline tool groups involved in the flexible inline maps must be on the same node.

Flexible Inline Arrangement License

Flexible Inline Arrangements on the GigaVUE TA Series requires an Inline Bypass License. Refer the table below:

GigaVUE TA Series Node	Monthly Subscription (Term License)	Perpetual Subscription
GigaVUE-TA25	IBP-TAX20-SW-TM	IBP-TAX20
GigaVUE-TA25E	IBP-TAx20-SW-TM	IBP-TAX20
GigaVUE-TA200	IBP-TAC20-SW-TM	IBP-TAC20
GigaVUE-TA200E	IBP-TAC20E-SW-TM	IBP-TAC20E
GigaVUE-TA400	IBP-TAC40-SW-TM	IBP-TAC40
GigaVUE-TA400E	IBP-TAC40E-SW-TM	IBP-TAC40E

Licensing Rules and Notes

- Without the license you will not be allowed to create Inline-net and inline-tool ports.
- If your license has expired or uninstalled then the Inline-net and inline-tool ports will move to inactive state port params edit such as speed change and admin enable option will be blocked except for port type modification.

Software Version

GigaVUE-FM and GigaVUE-OS running on software version 5.3.xx or higher support the flexible inline arrangement functionality.

Flexible Inline Solution Supported in Clustered Nodes

The GigaVUE® HC Series, GigaVUE-TA25, GigaVUE-TA200, GigaVUE-TA400, and GigaVUE-TA400E nodes can now be clustered with Flexible Inline solution for the configuration path. The traffic path is limited to a single node such that all inline ports should reside in the same node. This feature is applicable in Out-of-Band and Leaf-Spine cluster. The Flexible Inline features that can be configured are as follows:

Supported Feature	Reference
Flexible Inline Solution	Flexible Inline Arrangements
Single VLAN Tag	Configure Inline Single Tag.
Non Shared Tool	Configure Inline Tool Ports and Inline Tools.
Resilient Weighted Hashing	Refer to 'Resilient Weighted Hashing' in Configure Inline Tool Group.
Out-of-Band Copy The oob-copy target ports must reside on the same node as the respective flexible inline type maps. The traffic originated from oob-copy can be exposed to out-of-band tool residing on other nodes by using hybrid ports. This feature also allows the flexibility to have different oob-copies on each direction (a-to-b and b-to-a) from GigaVUE-FM for bidirectional flex maps. Starting in software version 6.4, the tool/hybrid port which is part of a regular byRule or passall map with ingress VLAN tag enabled on the network port can also be used as an OOB-copy port of flexible inline maps.	Refer to Example 7—Protected Flexible Inline, Out-of-Band Copy in GigaVUE-OS CLI Reference Guide.
Resilient Inline Arrangement. In a Resilient Inline Arrangement, both GigaVUE® HC Series nodes should be in different clusters or either one of the node can be a standalone node.	Refer to 'Resilient Inline Arrangement' in Configure Resilient Inline Arrangement.
Network Link Aggregation Group (LAG)	Refer to Configure Inline Network Link Aggregation Group (LAG).
Hashing	Refer to Asymmetrical Hashing options in Configure Inline Tool Group.

For 5.13.00 version, only GigaVUE® HC Series devices can be configured for Flexible Inline solution in a cluster. Flexible Inline solution in a cluster is not applicable for GigaVUE TA Series device but, it can be part of a GigaVUE® HC Series cluster that is configured for this functionality.

For 5.14.00 version, only GigaVUE® HC Series and GigaVUE-TA200 devices can be configured for Flexible Inline solution in a cluster. Flexible Inline solution in a cluster is not applicable for other GigaVUE TA Series device but, it can be part of a GigaVUE® HC Series cluster that is configured for this functionality.

For 5.16.00 version, only GigaVUE® HC Series, GigaVUE-TA200 and GigaVUE-TA25 devices can be configured for Flexible Inline solution in a cluster. Flexible Inline solution in a cluster is not applicable for other GigaVUE TA Series device but, it can be part of a GigaVUE® HC Series cluster that is configured for this functionality

Limitations

The inline-network, inline-tool and destination tool ports in oob-copy must reside on the same node and cannot be chosen from across devices in a cluster.

Functionalities Not Supported by Flexible Inline Arrangement

Certain functionality is not supported by flexible inline arrangements. In the following cases, use classic inline bypass instead:

- for inline tools that cannot tolerate the addition of VLAN tags to the traffic

Also not supported:

- **Combining Classic and Flex:** Using both Classic Inline and Flexible Inline arrangements on the same device is not supported even in a standalone environment.

Benefits of Flexible Inline Arrangements

Flexible inline arrangements offer flexibility in how traffic is guided through inline tools. It has the following benefits compared to classic Inline Bypass:

- Guides traffic through any arbitrary sequence of inspection tools.
- Shares inline tools across multiple inline network links and across multiple inline maps.
- Distributes traffic across multiple tools to meet bandwidth and throughput demands.

Flexible inline arrangements use the same software constructs as the classic Inline Bypass solution, such as inline network, inline tool, and inline tool group. However, inline network group and inline serial constructs are not needed.

Inline network groups have changed with flexible inline arrangements. Now every inline network is independent and can share any combination of tools in any order. The concept of inline network group is supported by creating multiple flexible inline maps. Also, multiple inline networks can be grouped into an inline network bundle. You can configure one inline map for the network bundle with the inline network bundle as the source.

Inline serial is not needed because the flexibility offered with flexible inline arrangements allows for the same configuration without the inline serial construct.

[Figure 12Flexible Inline Arrangements Scenario](#) illustrates the benefits of flexible inline arrangements by showing the kinds of deployment scenarios that can be enabled with this approach.

About Flexible Inline Maps

Traffic flows are the building blocks of flexible inline arrangements. Flows can be based on any flow mapping criteria, such as TCP port, IP subnet, or VLAN. There is a one-to-one correspondence between a traffic flow and a flexible inline map.

A flexible inline map is a new map type. Flexible inline arrangements allow inline maps from inline networks to arbitrary sequences of shared (overlapping) sequences of inline tools and inline tool groups.

Using flexible inline maps, you can identify specific flows of traffic using Layer 2 (L2) to Layer 4 (L4) rules, then designate the tools that will inspect the traffic, and specify the order of traffic to the tools.

You can configure a flexible inline map with a specific inline tool that is part of an inline tool group, which is associated with another flexible inline map. For example, you have created an inline tool group, ITG1 in which inline tools, IT1, IT2, and IT3 are grouped together. You can configure a flexible inline map, Map1 with inline network, IN1 as the source and ITG1 as the destination. You can configure a second flexible inline map, Map2 with IN1 as source and IT1 as destination. Such configuration is useful to guide specific traffic to a particular inline tool and the rest of the traffic to the inline tool group in which the inline tool is associated.

To properly guide traffic through the inline tools, each flow of traffic is assigned a VLAN tag. VLAN tags can be automatically assigned or can be user-defined. You can use flexible inline single tags to map incoming VLANs on the network side to the outgoing VLANs on the tool side.

With flexible inline arrangements, VLAN tags are associated with each inline map, not with each inline network port as in the case of classic Inline Bypass. A single inline network port can have multiple inline maps, each with a separate VLAN tag.

A VLAN tag automatically assigned to an inline map can be manually added to another inline map. For example, if an inline map is given the VLAN tag 2000, the same VLAN tag can be manually added to another inline map. Here, the system would prioritize the manually added VLAN tag and modify the auto-assigned VLAN Tag. This is not applicable for VLAN's assigned to internal inline maps such as import/export maps and inner non-proxy maps configured by GigaVUE-FM for Resilient Inline Arrangement (RIA) and RIA SSL solutions.

For example, traffic flows can be defined with the following VLAN tags:

- Unspecified traffic—VLAN 101
- Web traffic—VLAN 102
- Email traffic—VLAN 103
- Database traffic—VLAN 104

NOTE: The VLAN tags are added to the traffic before it is sent to the tools and are removed before it is sent back to the network.

Types of Flexible Inline Maps

To define a traffic flow, you must configure a flexible inline map. Following are the two types of flexible inline maps:

- **byRule**—Use the **byRule** map type to define a flow using map rules. Any standard L2-L4 mapping rules can be specified in the map, such as, IPv4, IPv6, L4 port, or UDA.
- **collector**—Use the **collector** map type for all other traffic. A collector is the lowest priority of map and does not have a map rule definition. Use a collector to catch any traffic that does not go to any other map. You can define a flexible inline collector map without any other maps in place. This provides a map passall, provided there are no rule-based maps. If you want all the traffic to go to the same tools, you only need to configure a collector.

Flexible inline arrangements guide rule-based or collector-based inline traffic flows through unidirectional or bidirectional sequences of inline tools or inline tool groups. The traffic path can be set up independently for side A to side B and side B to side A directions, meaning that the traffic flow can be either symmetrical or asymmetrical.

You specify the ordered list of inline tools or inline tool groups that will inspect a particular flow of traffic. Additionally, you can specify if the A-to-B and B-to-A directions have the same order or the reverse order. Reverse order is the order of inline tools as they are wired in a physical network if a Gigamon network packet broker was not present.

For example, in the A-to-B direction, if the tools are specified in the following order: T1, T2, T3, the same order in the B-to-A direction will be: T1, T2, T3, while the reverse order in the B-to-A direction will be: T3, T2, T1. Or, you can specify the order of the tools explicitly, for example, the B-to-A direction can be: T2, T1, T3.

You can create separate flexible inline maps for each flow of traffic to be inspected by a sequence of inline tools. Create maps until you have accounted for all the flows of traffic. Any unspecified traffic will go to the collector. You can also specify map priorities for the flexible inline maps.

Configure Flexible Inline Maps

Before you configure a flexible inline map, ensure that you create the required inline network. For information, refer to [Configure Inline Network Ports and Inline Network](#).

To configure a flexible inline map:

1. Go to **Physical > Orchestrate > Inline Flows**, and then click **Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the Flexible Inline Canvas that is displayed, select the required device for which you want to configure the flexible inline map.
3. Drag and drop the required inline network into the Flexible Inline Canvas.
4. Drag and drop the Flex Map into the Flexible Inline Canvas.
5. In the **Properties** pane that appears, enter a name and description for the map in the **Alias** and **Description** fields.
6. Enter the required **Tool Side VLAN Tag**.
7. Select the **TPID** for the Tool Side VLAN Tag. The default value of TPID is 0x8100. You can select the other supported values 0x9100 and 0x88a8 from the drop-down list.
8. Add the required rules for the inline map. You can also choose to import an existing map template. For instructions on how to create a map template, refer to [Create Map Templates](#).
9. Click **OK** to save the configuration.

Flexible Inline Arrangements—Rules and Notes

The table below lists the main rules, limits, and unsupported features for flexible inline arrangements, organized for quick reference.

Rule / Feature	Details / Limitations
Exclusive Use	Inline tools in a flexible inline map cannot be used in classic inline or inline decryption maps. All inline networks and tools must belong to only one type of map.
Collector Maps	Only one unidirectional collector map is allowed for the same inline network. To use different VLANs in each direction, create separate unidirectional maps with unique VLAN tags. Tags can be set manually or assigned automatically by GigaVUE-FM.
Unsupported Features (	- Physical Bypass (no BPS card)

Rule / Feature	Details / Limitations
GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA25E, GigaVUE-TA25, GigaVUE-TA400, GigaVUE-TA400E)	- Flexible and Resilient Inline SSL (no GigaSMART card) - GRIP (no BPS card) - ICAP (no GigaSMART card) - Classic Inline Bypass
VLAN Tagging and OOB Copy Limits (GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-HC1-Plus)	Flexible Inline Single VLAN Tag with monitoring mode may send incorrect VLAN tags. OOB copy packets may also have wrong tags. You cannot use BYPASS WITH MONITORING with MONITORING mode on the tool. OOB copy from inline network is not allowed in this mode.
Inline Map Limits — Bidirectional	GigaVUE-TA25,GigaVUE-TA25E, GigaVUE-HC1-Plus → 126 maps GigaVUE-HC1,GigaVUE-HC3 (CCv1 & CCv2), GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA400, GigaVUE-TA400E → 256 or 512 maps depending on setup.
Inline Map Limits — Unidirectional	GigaVUE-TA25,GigaVUE-TA25E, GigaVUE-HC1-Plus→ 252 maps GigaVUE-HC1,GigaVUE-HC3 (CCv1 & CCv2),), GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA400, GigaVUE-TA400E → 512 or 1024 maps depending on setup.
Flexible Inline SSL Limits	Not supported with Inline Network LAG. Setting inline tools to “Drop” in the chain does not block Inline SSL traffic.
Filtering Limits(GigaVUE-TA400, GigaVUE-TA400E)	VLAN-based filtering in the Egress Port Filter for OOB copies is not supported. If one tool in the map is in monitoring mode, all tools must use the same mode. Asymmetric hashing (a-srcip-b-dstip and b-srcip-a-dstip) is not supported.
Protocol Pass-Through(GigaVUE-TA400, GigaVUE-TA400E)	CDP pass-through is not supported when the source is an Inline Network LAG. Bypass for LACP, CDP, and LLDP is supported.
Scaling Limits — GigaVUE-TA400, GigaVUE-TA400E	Max Inline Networks and Tools: 48 Max Inline Network LAG list: 24 Max Inline tools or tool groups per direction: 16 Max OOB copy entries per direction: 17 Max OOB copy ports per entry: 128

Visualize the Flexible Inline Arrangements Canvas

The GigaVUE-FM user interface provides clear visualization of inline maps. The user interface makes it easy to visualize and configure inline maps and tools. The drag-and-drop capability lets you define and add tools to maps, in any order.

Configure Flexible Inline Flows

The different flexible inline flows are configured and managed using GigaVUE-FM. For end-to-end workflows, examples, and best practices, refer to the [GigaVUE Inline Solutions Guide](#).

Troubleshoot Flexible Inline Flows

The flexible inline canvas provides you with the ability to view the status of the flexible inline flow deployments, port statistics, and the details of the cluster-level maps used in the deployments. You can use these details to troubleshoot any issues or failures in your flexible inline flows.

The flexible inline canvas has the following tabs:

Status


The Status tab provides details of the forwarding states of inline network. Click the required component that is part of the selected inline network to view the component's properties. Refer to [View the Forwarding States of Inline Networks](#).

It also provides the status of the components that are part of the selected inline network. Hover over the status to view the description.

Click the **Show Stats/Hide Stats** toggle button to view the Rx/Tx rate for the components that are part of the flexible inline flow deployment. Refer to the following figure for details.

Inline Flows

Configuration Canvas



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.

STATUS		STATISTICS		TROUBLESHOOT	
Inline Network		IN1		Hide Stats	
Port A		1/1x12			
Port B		1/1x14			
Packet Counts		Packet Rate per Second		Drops	
PORT A	PORT B	PORT A	PORT B	PORT A	PORT B
Rx: 218992	219796 :Tx	Rx: 1	1 :Tx	Rx: 0	0 :Tx
Tx: 296447	334089 :Rx	Tx: 3	5 :Rx	Tx: 0	0 :Rx
Inline Tool		IT1			
Port A		1/1x6			
Port B		1/1x7			
Packet Counts		Packet Rate per Second		Drops	
PORT A	PORT B	PORT A	PORT B	PORT A	PORT B
Rx: 137804	136651 :Tx	Rx: 0	0 :Tx	Rx: 0	0 :Tx
Tx: 6152	5879 :Rx	Tx: 0	0 :Rx	Tx: 0	0 :Rx
Inline Tool		IT2			
Port A		1/1x5			
Port B		1/1x1			
Packet Counts		Packet Rate per Second		Drops	
PORT A	PORT B	PORT A	PORT B	PORT A	PORT B
Rx: 307082	307082 :Tx	Rx: 0	0 :Tx	Rx: 0	0 :Tx
Tx: 347138	347138 :Rx	Tx: 2	2 :Rx	Tx: 0	0 :Rx

HSM Statistics

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

STATUS
STATISTICS
TROUBLESHOOT

Search Device

hc1-left

- ▼ Inline Network Bundle
- ▼ Inline Network LAG
- ▼ Inline Network

INET1
- ▼ ICAP Client

SSL APP INSSL

Engine Port 1/2/e1

View HSM Stats

Monitor Statistics


Sessions (Active)		Traffic		Performance Metrics	
Metric	Value	Metric	Value	Metric	Value
Established Sessions	0	Packets/sec Rx	2	Active Connections Per Second (CPS)	1
SSL Sessions	0	Drop Packets/sec	0	CPU Utilization%	0
Incomplete Sessions	0	SYN Received	0	CPU Idle Time%	100
Multi VLAN Sessions	0	SYN ACK Received	0		

FM Instance: GigaVUE-FM

Last Updated At: Aug 2, 2023 11:45:21

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.

Inline Flows

Back to icap1 Statistics Session Troubleshoot

ICAP Client Alias: icap1 GigaSMART Group Inline Network Server Statistics

GigaSMART Group Clear GS Stats

Packets/sec Rx CPU Utilization%

Drop Packets/sec CPU Idle Time%

GS Engine

Inline Network

IP Interface Port

Server Statistics

Server Alias	Interface Id	GSGroup Alias	Total Connections	Active Connecto...	Average RTT	Requests Mapped	Request Sent	Request Inflight	Request Timeout	Responses Recei...	Session I
icap1	1_3_e1	FmAuto-gsgro...	689	0	0	689	689	0	0	684	563

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

Inline Flows

Back to icap1 Statistics Session Troubleshoot

Filter Action Export

Interface Id	GigaSMART Gro...	ICAP Server	Source IP	Destination IP	Source Port	Destination Port	Protocol	TCP Status	ICAP Session	Error	ICAP Status	Start Time	Total Duration	C25 Total Pkt Co...	S2C Total Pkt Co...
1/3/e1	FmAuto-gsgro...	icap1	5.5.5.5	34.107.221.82	38840	80	http1.1	BYPASS	BYPASS	ERR	REQMOD	1970-01-01 0...	28	7	4
1/3/e1	FmAuto-gsgro...	icap1	5.5.5.5	104.254.148...	39950	443	http1.1	BYPASS	BYPASS	ERR	REQMOD	1970-01-01 0...	14	18	17
1/3/e1	FmAuto-gsgro...	icap1	5.5.5.5	204.79.197.2...	43952	443	tcp	NO	NO	NO ERR		1970-01-01 0...	23	1	0
1/3/e1	FmAuto-gsgro...	icap1	5.5.5.5	40.126.62.130	47740	443	http1.1	YES	YES	NO ERR	RESPMOD	1970-01-01 0...	16	14	13
1/3/e1	FmAuto-gsgro...	icap1	5.5.5.5	52.37.154.49	54204	443	http1.1	YES	YES	NO ERR	RESPMOD	1970-01-01 0...	28	4	3
1/3/e1	FmAuto-gsgro...	icap1	5.5.5.5	34.117.65.55	43502	443	http1.1	BYPASS	BYPASS	ERR	REQMOD	1970-01-01 0...	26	6	5
1/3/e1	FmAuto-gsgro...	icap1	5.5.5.5	20.189.173.12	37036	443	http1.1	YES	YES	NO ERR	REQMOD	1970-01-01 0...	16	16	15
1/3/e1	FmAuto-gsgro...	icap1	5.5.5.5	20.189.173.12	37038	443	http1.1	YES	YES	NO ERR	RESPMOD	1970-01-01 0...	12	11	8
1/3/e1	FmAuto-gsgro...	icap1	5.5.5.5	20.189.173.12	36046	443	http1.1	YES	YES	NO ERR	REQMOD	1970-01-01 0...	20	16	15
1/3/e1	FmAuto-gsgro...	icap1	5.5.5.5	40.126.62.130	47736	443	http1.1	BYPASS	BYPASS	ERR	REQMOD	1970-01-01 0...	16	10	9
1/3/e1	FmAuto-gsgro...	icap1	5.5.5.5	68.67.129.19	33634	443	http1.1	BYPASS	BYPASS	ERR	REQMOD	1970-01-01 0...	14	136	136
1/3/e1	FmAuto-gsgro...	icap1	5.5.5.5	23.37.116.6	47334	443	http1.1	YES	YES	NO ERR	RESPMOD	1970-01-01 0...	15	48	48
1/3/e1	FmAuto-gsgro...	icap1	5.5.5.5	34.107.221.82	38830	80	http1.1	BYPASS	BYPASS	ERR	REQMOD	1970-01-01 0...	28	7	4
1/3/e1	FmAuto-gsgro...	icap1	5.5.5.5	69.192.139.2...	50548	443	http1.1	YES	YES	NO ERR	RESPMOD	1970-01-01 0...	20	4	3

Go to page: 1 of 1 14 sessions total

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.

STATUS

STATISTICS

TROUBLESHOOT

Cluster level Maps

▼

ING-Bundle (Inline Network Bundle)

▼

default_inline_net_1_1_3

▼

FmAuto-Web-Traffic_default_inline_net-20a73719-7412-41e4-8011-96df2b5b1409

Comment

CREATED BY GIGAVUE-FM. DO NOT MODIFY OR DELETE, processed map: Web-Traffic_default_inline_net_1_1_3

▼

Priority

1

Source

●

default_inline_net_1_1_3

AtoB Tools

●

SSL-VA1,

●

Fire-eye-APT,

●

SSL-VA2

BtoA Tools

Reverse

Tool Side VLAN Tag

auto 3997

▼

Rules

Rule	Type	Bi-directional	Packets	Octets	Conditions	Comments
Rule 1	Pass	No	73583	9839834	portDst: 443	

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

1. From the flexible inline canvas, go to the **Status** tab to check the forwarding states of the required inline network.
2. Go to the **Statistics** tab to check the port statistics to ensure that there are no drops, errors, or discards.

NOTE: Click the **Statistics** tab again to refresh the data.

3. Go to the **Troubleshoot** tab to check the required map configurations and isolate the issue.

For instructions on how to troubleshoot a specific issue, refer to [Example: Troubleshoot Traffic Issues Between Side A and Side B](#).

Example: Troubleshoot Traffic Issues Between Side A and Side B

This section provides you an example of how to troubleshoot a specific issue using the flexible inline canvas.

Consider that you have an inline network IN1 with one by-rule map and one collector map. The traffic is not flowing from Side A to Side B. To troubleshoot this issue:

1. Go to **Physical > Inline Flows**.
2. Select the required device, and then drill down to the inline network that has the traffic issue.
3. Click **Status** to view the forwarding states of the inline network. Ensure that the forwarding state of the inline network is Normal. For details of the forwarding states, refer to [View the Forwarding States of Inline Networks](#).
4. Click **Statistics** to view the total packet count of the inline network. The Rx count of Port A of the inline network must match the Tx count of Port B. If the count does not match, the traffic is blocked between Port A and Port B.

NOTE: To refresh the statistical data, click **Statistics** again.

5. Click **Troubleshoot** to view the configurations of the by-rule and collector maps configured for the inline network.
 - a. Check that the by-rule map has packet count.
 - b. If the by-rule map has packet count, check the inline tools from A to B and B to A directions of the by-rule map to ensure that the inline tools are in up state and the **Flex Traffic Path** is set to **To inline tool**.

6. Click **Statistics** to view the port statistics of the required inline tool. If the Rx count of Port A of the inline tool does not match the Tx count of Port B of the inline tool, the traffic is dropped at the inline tool. Check the inline tool and take the required action.
7. If the Rx count of Port A of the inline tool matches the Tx count of Port B of the inline tool, repeat steps 5 and 6 for the collector map configured for the inline network.
8. If the Rx count of Port A of the inline tool configured for the collector map matches the Tx count of Port B of the inline tool, contact Gigamon Technical Support.

View the Forwarding States of Inline Networks

To view the forwarding states of inline networks in the flexible inline canvas, choose the required inline network, and then click **Status**. Refer to the following figure.




Figure 13 Inline Network Forwarding States

Following inline network states are not explicitly shown in the flexible inline canvas:

- Normal—If the state of all inline tools are up, the inline network is in Normal state.
- Abnormal—If any inline tool involved in flexible inline maps (directly or indirectly as a member of an inline tool group) is operationally down and there is no network-level failover action in effect, the inline network is in an Abnormal state.

Following table provides the list of forwarding states of inline network and their description.

Table 15: Forwarding States of Inline Networks

Inline Network Physical Bypass	Inline Network Traffic Path	Far-End Status of Links Connected to Inline Network Ports	Operational State of Not Forced Inline Tools and Inline Tool Groups Involved in Maps from the Inline Network	Forwarding State	Description
enable	any inline network traffic path configuration	any combination of far-end ports status	any combination of operational state of the inline tools or inline tool groups involved in the maps originating from the inline network	PHYSICAL BYPASS	all traffic exchanged directly between the end nodes without being noticed by the switching fabric (GigaVUE node acting as a wire or fiber)
disable	traffic path set to drop	any combination of far-end ports status	any combination of operational state of the inline tools or inline tool groups involved in the maps originating from the inline network	DISABLED	all traffic arriving at the inline network ports is dropped
disable	traffic path set to bypass, monitoring, or to-inline-tool	at least one far-end port is down	any combination of operational state of the inline tools or inline tool groups involved in the maps originating from	DISCONNECTED	No traffic is exchanged between the nodes

Inline Network Physical Bypass	Inline Network Traffic Path	Far-End Status of Links Connected to Inline Network Ports	Operational State of Not Forced Inline Tools and Inline Tool Groups Involved in Maps from the Inline Network	Forwarding State	Description
			the inline network		
disable	traffic path set to bypass	both far-end ports are up	any combination of operational state of the inline tools or inline tool groups involved in the maps originating from the inline network	FORCED BYPASS	All traffic that matches any of the maps originating from the inline network is redirected through a logical bypass
disable	traffic path set to monitoring	both far-end ports are up	any combination of operational state of the inline tools or inline tool groups involved in the maps originating from the inline network	FORCED BYPASS WITH MONITORING	A copy of the traffic originating from the inline network bypasses the sequence of inline tools and inline tool groups and is re-directed to the opposite-side inline network port. Another copy of the traffic is directed to the sequence of inline tools and inline tool groups, except that no traffic of the second copy is sent to the exit port.

Inline Network Physical Bypass	Inline Network Traffic Path	Far-End Status of Links Connected to Inline Network Ports	Operational State of <i>Not</i> Forced Inline Tools and Inline Tool Groups Involved in Maps from the Inline Network	Forwarding State	Description
disable	traffic path set to to-inline-tool	both far-end ports are up	all inline tools involved (directly or indirectly as members of inline tool groups) in the maps originating from the inline network are in the <i>up</i> operational state	NORMAL	The traffic is guided between the source inline network port and the destination inline network port according to the status of the inline tools and inline tool groups. NOTE: The state of all inline tools must be <i>up</i>, including inline tools configured as spare in an inline tool group, inline tools or inline tool group members in the a-to-b and b-to-a lists configured with any traffic path other than to-inline-tool.
disable	traffic path set to to-inline-tool	both far-end ports are up	at least one of the inline tools or	NETWORK PORTS FORCED DOWN	No traffic is exchanged between the

Inline Network Physical Bypass	Inline Network Traffic Path	Far-End Status of Links Connected to Inline Network Ports	Operational State of <i>Not</i> Forced Inline Tools and Inline Tool Groups Involved in Maps from the Inline Network	Forwarding State	Description
			inline tool groups involved in the maps originating from the inline network configured with the traffic path parameter to-inline-tool and failover action of network-port-forced-down is in the <i>down</i> operational state		inline network ports, and the inline network ports are brought down
disable	traffic path set to to-inline-tool	both far-end ports are up	a. none of the inline tools or inline tool groups involved in the maps originating from the inline network configured with to-inline-tool and failover action network-port-forced-down is in the <i>down</i> operational state b. at least one of the inline tools or inline tool groups involved in the maps originating	FAILURE INTRODUCED DROP	All traffic arriving at the inline network ports is dropped

Inline Network Physical Bypass	Inline Network Traffic Path	Far-End Status of Links Connected to Inline Network Ports	Operational State of <i>Not</i> Forced Inline Tools and Inline Tool Groups Involved in Maps from the Inline Network	Forwarding State	Description
			from the inline network configured with to-inline-tool and failover-action of network-drop is in the <i>down</i> 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 <i>down</i> operational state b. at least one of the inline tools or inline tool groups involved in the maps originating from the inline network configured with	FAILURE INTRODUCED BYPASS	All traffic that matches any of the maps originating from the inline network is redirected through a logical bypass

Inline Network Physical Bypass	Inline Network Traffic Path	Far-End Status of Links Connected to Inline Network Ports	Operational State of <i>Not</i> Forced Inline Tools and Inline Tool Groups Involved in Maps from the Inline Network	Forwarding State	Description
			to-inline-tool and failover action of network-bypass is in the <i>down</i> 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 <i>down</i> operational state and there is no network-level failover action in effect, the inline network is in the ABNORMAL state.

NOTE: When the Inline Network traffic path is set to monitoring and the Inline Tool Failover action is configured as 'Network Port Force Down,' the Inline Tool will continue to receive traffic, even if it is disabled or in an operational down state. This behavior is expected.


Import and Export Flexible Inline Solution

The flexible inline canvas allows you to import and export a flexible inline solution. The exported solution gets downloaded in to your local folder as an YAML file, which can be imported and deployed again in the following scenarios:

- Retrieve a solution that was deleted unintentionally
- Deploy the solution in another device
 - Re-deploy a solution in the device after GigaVUE-FM is upgraded to a new version (in case of issues in the existing solution)

The downloaded YAML file contains the following information:

- Software version of the GigaVUE-FM instance and the device on which the solution was created
- Member, port and configuration information of the various inline configurations such as the Inline Network Alias, Member Ports, Inline Tool Group Alias, and other such details.



Rules and Notes

- When deploying a solution after GigaVUE-FM is upgraded to a new version, the source port in the target device must not have been used in any solution.
- When importing an YAML file, the following configurations specified in the YAML file must match the configurations in the target device:
 - Inline network alias
 - Inline network bundle alias
 - Inline tool alias
 - Inline network LAG alias
 - Inline SSL App alias

Import and Export a Flexible Inline Solution

To import a file:

1. Click the **Import** button on the canvas.
2. Select the .YAML file, containing the required flexible inline solution, saved in your local folder. The solution appears in the canvas.
3. Click **Deploy** to deploy the solution, again.

You can also edit an imported solution on the canvas, and re-export the same as a new solution.

To export a solution:

1. Create a flexible inline solution and **Deploy** it.
2. Click the **Export** button. The solution is downloaded as an YAML file.
3. Save the file to the required location.

Backup and Restore Flexible Inline Flows

You must backup the devices and GigaVUE-FM at the same time. Ensure that the devices are not undergoing configuration edits during backup or restore. For more information, refer to the “*Restore Devices and GigaVUE-FM for Traffic Management Solutions*” section in the “*GigaVUE Administration Guide*”.

Inline Bypass Solutions

This chapter provides the following information about inline bypass solutions:

- [About Inline Bypass Solutions](#)
- [Configure Inline Bypass Solutions](#)

About Inline Bypass Solutions

Security tools such as firewalls and intrusion protection systems (IPSs) are often connected inline on production networks, with traffic flowing from the network segment through the tool and then back onto the production network.

Inline bypass solutions involve bidirectional traffic between two networks, intercepted by a Gigamon node, and guided through one or more inline tools.

Inline bypass is a pillar of the GigaSECURE Security Delivery Platform. For an overview of GigaSECURE, refer to [About Inline Bypass Solutions](#).

Inline bypass is supported on GigaVUE® HC Series nodes: GigaVUE-HC3, GigaVUE-HC1 and GigaVUE-HC1-Plus.

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

This section describes inline bypass solutions. Refer to the following sections for details:

- [Introduction to Inline Bypass Solutions](#)
- [Logical Bypass and Physical Bypass](#)
- [Inline Networks](#)
- [Simple and Complex Inline Bypass Solutions](#)

NOTE: The configuration of inline bypass solutions can be complex. Follow the order of configuration outlined in [Configuration Steps](#) and demonstrated in [Inline Bypass Solution Examples](#).

Introduction to Inline Bypass Solutions

Inline bypass solutions place the Gigamon node inline between two sides of a network. The Gigamon node sends uninspected traffic from one side of the network to an inline tool (such as an IPS), and then sends the inspected traffic to the other side of the network.

The reasons for deploying Gigamon's inline bypass solutions are as follows:

- to protect against inline tool failures (including loss of link to inline tools and any inline tool problems detected through heartbeat monitoring)
- to bypass the inline tools for traffic that does not need to be examined by the tools or that cannot be processed by the tools
- to distribute the traffic load among multiple tools
- to send specific traffic to specialized inline tools
- to guide traffic serially through the inline tools, with the traffic from one tool flowing to the next, so that all tools see the same traffic
- to share inline tools among multiple inline network links
- to implement tiered network security by combining inline and out-of-band, allowing the traffic to be examined by both types of tools

For improved reliability and ease of maintenance, the inline bypass solutions also provide the following:

- protection against power loss through physical bypass
- protection against power loss through the high availability solution, Gigamon Resiliency for Inline Protection (GRIP™)

Capabilities of Inline Bypass Solutions

The inline bypass solution offers the following capabilities:

- Guides traffic through one or more inline tools according to user-defined bidirectional connectivity arrangements among respective inline network ports and inline tool ports.
- Reacts to failure conditions, such as the failure of inline tools, or the failure of links leading to the end-point devices between which the inline tool is inserted.
- Configures maps to associate inline networks with inline tools and inline tool groups.
- Supports bidirectional heartbeat and negative heartbeat to monitor inline tool health.
- Supports protected and unprotected inline networks.

- Supports physical bypass with specialized hardware on GigaVUE-HC3, , GigaVUE-HC1 and GigaVUE-HC1-Plus(bypass combo modules equipped with optical protection switches)
- Supports physical bypass on and GigaVUE-HC1 (TAP modules equipped with electrical relays).
- Distributes traffic across multiple inline tools, providing load sharing and traffic redistribution should a tool fail.
- Supports out-of-band maps for selective forwarding of inline traffic to monitoring tools.
- Supports sharing of inline tools by multiple inline networks.
- Supports N+1 and 1+1 inline tool redundancy for inline tool groups.
- Supports guiding inline traffic through inline tools in a serial fashion.
- Supports inline flow mapping through rule-based and shared collector inline maps.
- Provides a high availability solution, Gigamon Resiliency for Inline Protections (GRIP).

Logical Bypass and Physical Bypass

Logical bypass lets traffic bypass the inline tool should it experience a failure. A failure is declared if the GigaVUE node either loses connectivity to the inline tool or fails to receive a heartbeat from the tool.

Logical inline bypass does not require any specialized hardware and can be facilitated on regular ports on GigaVUE-HC3, GigaVUE-HC1 and GigaVUE-HC1-Plus modules. On GigaVUE-HC1, logical inline bypass can be configured on the base module.

Physical bypass provides protection against failure of the GigaVUE node, such as if power is lost. On GigaVUE® HC Series nodes, it is implemented with specialized hardware called bypass combo modules.

The specialized hardware triggers a bypass when power is lost to the node. When the physical bypass is activated, traffic flows from one side of the network to the other, but without monitoring.

The GigaVUE-HC3 bypass combo module is a bypass switch (BPS) module, as follows:

- Bypass Combo Module with two 100Gb/40Gb SR4 MPO inline network port pairs and sixteen regular SFP+ (10Gb) port cages (BPS-HC3-C25F2G)
- Bypass Combo Module with two 40Gb or 100Gb LR4 inline network port pairs and sixteen 10Gb/25Gb cages (BPS-HC3-Q35C2G, and BPS-HC3-C35C2G).

The GigaVUE-HC1 bypass combo module is a bypass switch (BPS) module, as follows:

- Bypass Combo Module with two SX/SR (50/125µm multimode) inline network port pairs and four regular SFP+ (1Gb/10Gb) port cages (BPS-HC1-D25A24)
- Bypass Combo Module with six single-mode LR 1Gb/10Gb inline network port pairs BPS-HC1-D35C60 (LX/LR) and BPS-HC1-D25A60 (SX/SR) .

The following bypass modules are available for the GigaVUE-HC1-Plus:

- Bypass Combo Module with two SX/SR (50/125µm multimode) inline network port pairs and four regular SFP+ port cages (BPS-HC1-D25A24). It requires GigaVUE-HC1-Plus to be running software version 6.3 or higher. It provides both physical and logical inline bypass, with automatic failover protection in case of power failure. The ports can operate at 1Gb/10Gb speeds and can be configured as any port type
- Bypass modules with six inline network port pairs each, supporting multi-mode SR and single-mode LR, respectively, at 1Gb/10Gb speeds (BPS-HC1-D25A60 (SX/SR) and BPS-HC1-D35C60 (LX/LR)).

All these modules are hot-swappable and provide physical bypass protection, ensuring network continuity in the event of a power failure or module removal. The BPS modules provide optical bypass, while the TAP-HC1-G10040 provides copper bypass functionality.

Physical bypass is also supported on the following copper TAP modules on GigaVUE-HC1 and GigaVUE-HC1-Plus:

- On the GigaVUE-HC1 copper TAP module, TAP-HC1-G10040
- On the GigaVUE-HC1-Plus copper TAP module with 8 x 1Gb ports for either four TAPs or bypass copper (BPC) port pairs (TAP-HC1-G10040). It can be used as a copper bypass module, providing protected inline networks for copper ports. It requires GigaVUE-OS version 6.4 or higher.

The GigaVUE-HC3, GigaVUE-HC1 and GigaVUE-HC1-Plus offer physical and logical inline bypass. Physical bypass provides automatic failover protection in the case of a power failure. On the GigaVUE HC Series nodes, the bypass combo modules provide the physical bypass function. Or on GigaVUE-HC1, the copper TAP modules provide the physical bypass function.

The physical bypass function, as it applies to a single pair of inline network ports, is as follows:

- When the module is not powered, (either the entire node is powered down or the module is removed from the node), the inline network port pair is in the physical bypass mode. That means that traffic is exchanged directly between network Port A and network Port B of the inline network pair.
- When the module is powered, the mode (inline or bypass) of the inline network port pair is controlled through software. In the physical bypass mode, the inline network port pair behaves exactly as if the module was not powered. In the inline mode, the inline network port pair behaves as any other inline network port pair configured for working with an inline tool.

For information on the bypass combo module on GigaVUE-HC3, refer to the *GigaVUE-HC3 Hardware Installation Guide*. This document also contains procedures for installing, removing, and replacing modules in the GigaVUE-HC3.

For information on the bypass combo module or the TAP-HC1-G10040 module on GigaVUE-HC1, refer to the *GigaVUE-HC1 Hardware Installation Guide*. This document also contains procedures for installing, removing, and replacing modules in the GigaVUE-HC1.

For information on the bypass combo module on GigaVUE-HC1-Plus, refer to the *GigaVUE-HC1-Plus Hardware Installation Guide*. This document also contains procedures for installing, removing, and replacing modules in the GigaVUE-HC1-Plus Hardware Installation Guide.

Refer also to [Figure 14 Physical Bypass Protection](#) and [Protected Inline Network](#).




Figure 14 Physical Bypass Protection

Simple and Complex Inline Bypass Solutions

Simple and complex inline bypass solutions are described in the following sections:

- [Typical Configuration](#)
- [Distribution to Multiple Inline Tools](#)
- [Inline Tools in a Series](#)
- [Multiple Inline Networks](#)
- [Inline Flow Mapping®](#)
- [Send Traffic to Out-of-Band Tools](#)

Typical Configuration

In the typical or most common configuration, a single inline tool is inserted in a Network A to Network B link. All traffic is sent to the inline tool and inspected in both directions.

A typical configuration is shown in [Figure 15 Simple Configuration](#). The physical protection switch is optional and present only when using bypass combo modules or copper TAP modules.




Figure 15 Simple Configuration

Distribution to Multiple Inline Tools

One of the challenges to inline monitoring is scaling the throughput of the inline tools with the speed of the network. One way to do this is to share the load across multiple inline tools.

Because inline inspection of network traffic is processor-intensive, it is common to combine the power of multiple inline tools to monitor traffic. This is especially true for 10Gb, 40Gb, and 100Gb networks and tools that only support 1Gb or 10Gb interfaces.

An inline tool group is an arrangement of multiple inline tools which share the traffic load. Traffic is sent to the tools based on standard hashing parameters. This is referred to as hash-based distribution. If a single tool in the group of tools fails, the packets will be redistributed to other tools. This ensures that all packets are inspected.

A multiple inline tool arrangement is shown in [Figure 16 Multiple Tool Arrangement](#).




Figure 16 Multiple Tool Arrangement

Multiple tool distributions can be non-redundant or have 1+1 or N+1 redundancy. [Inline Tool Groups](#)

Inline Tools in a Series

Tools can form an inline series, in which the traffic from one tool flows to the next, so all tools see the same traffic.

Refer to [Figure 17 Inline Tool Series](#) for an inline tool series. In [Figure 17 Inline Tool Series](#), traffic is only shown from A-to-B.




Figure 17 Inline Tool Series

For more information on inline tool series, refer to [Inline Tools in a Series](#).

Multiple Inline Networks

An inline network group is an arrangement of multiple inline networks that share the same inline tool, inline tool group, or inline tool series. The numbers of networks to tools can be many-to-one as shown in [Figure 18 Inline Network Group](#) or many-to-many. Traffic is guided to a particular inline network through internal VLAN ID tagging.




Figure 18 *Inline Network Group*

For more information on inline network groups, refer to [Configure Inline Bypass Solutions](#).

Inline Flow Mapping®

Some tools are optimized for particular inline traffic. With inline flow mapping, the GigaVUE node forwards packets to different inline tools based on criteria, such as TCP/UDP port number, or any other rule that can be defined in a map. Using these map rules, selected traffic can be sent to specific tools.

When inline flow mapping is combined with distribution to multiple tools, one application, such as Web traffic, can be sent to one tool or group of tools, while another application, such as email traffic, can be sent to another tool or group of tools. If there is traffic, such as encrypted traffic, that does not need to be or cannot be inspected, it can be bypassed.

An inline flow mapping based traffic distribution is shown in [Figure 19 Inline Flow Mapping® Based Traffic Distribution](#). In [Figure 19 Inline Flow Mapping® Based Traffic Distribution](#), traffic is only shown from A-to-B.




Figure 19 *Inline Flow Mapping® Based Traffic Distribution*

Send Traffic to Out-of-Band Tools

Traffic can be sent to out-of-band (OOB) tools. Any port used for inline functionality can be used as a source for a map to an out-of-band tool. This includes any inline network port or inline tool port. For example, you can use inline tool ports to inspect packets that have passed through the IPS.

An out-of-band arrangement is shown in [Figure 20 Out-of-Band Arrangement](#).




Figure 20 Out-of-Band Arrangement

Configure Inline Bypass Solutions

Inline bypass solutions are arrangements of inline networks that act to guide traffic through one or more inline tools attached to a GigaVUE node.

Configure inline bypass solutions using the following software constructs:

- **Inline Network.**
For configuring inline networks, which are pairs of network ports attached to the sides of the network, refer to [Inline Networks](#).
- **Inline network groups**
For configuring inline network groups, which are multiple inline networks sharing the same inline tool or tools, refer to [Inline Network Groups](#)
- **Inline tool**
For configuring inline tools, which are pairs of inline tool ports plus the inline tools attached to the ports, refer to [Inline Tools](#)
- **Heartbeat profile**
For configuring monitoring of inline tools using heartbeat and/or negative heart beat packets or both, refer to [Heartbeats](#).
- **Inline tool group**
For configuring multiple inline tools, which are arrangements of multiple inline tools to which traffic is distributed based on hashing, refer to [Inline Tool Groups](#)
- **Inline serial tools**
For configuring inline tools in a series, in which traffic is guided through the inline tools in a serial fashion. Refer to [Inline Serial Tools](#)

Inline Networks

Currently configured in-line networks are displayed on the Inline Networks page, which is opened by selecting **Inline Bypass > Inline Networks**.

Click **New** to open the **Inline Network** configuration page to configure an inline network, which is a pair of network ports arranged for inline monitoring. The arrangement facilitates access to a bidirectional link between the two networks (two far-end network devices) that need to be linked through an inline tool.

An inline network consists of inline network ports, always in pairs, running at the same speed, on the same medium (either fiber or copper). The inline network ports must be on the same GigaVUE-HC3, GigaVUE-HC1 or GigaVUE-HC1-Plus node.


The types of inline networks are as follows:

- [Unprotected Inline Network](#)
- [Protected Inline Network](#)
- [Mix of Protected and Unprotected](#)

Unprotected Inline Network

An unprotected inline network consists of two ports of the inline-network type, which facilitate access to a bidirectional link between two networks (or more precisely, two far-end network devices).

Any available network type ports on a GigaVUE HC Series node can be configured to be inline-network type ports and combined to form an unprotected inline network. An unprotected inline network is shown in the following figure.




Protected Inline Network

A protected inline network uses the port pairs associated with optical protection switches located on bypass combo modules on GigaVUE-HC3, GigaVUE-HC1 and GigaVUE-HC1-Plus. The specialized hardware triggers a bypass when power is lost to the node. On GigaVUE-HC1 and GigaVUE-HC1-Plus, protected inline networks for copper are supported on the TAP-HC1-G10040 module.

Physical bypass provides protection against failure of the Gigamon node due to a loss of power. When physical bypass is activated, traffic flows from one side of the network to the other, but without monitoring.

Because of the specialized hardware, when a bypass combo module initializes, the system creates inline-network type ports for each of the protected port pairs on the module. This does not happen automatically on the TAP-HC1-G10040 module on the GigaVUE-HC1.

A protected inline network is shown in the following figure.



Mix of Protected and Unprotected

Any combination of protected and unprotected inline networks is supported in inline bypass solutions.

The GigaVUE-HC3 supports a 100Gb bypass combo module and starting in software version 5.2, the GigaVUE-HC3 100Gb bypass combo module supports dual speeds: 100Gb or 40Gb.

NOTE: You do not need to create the inline network ports for the protected port pairs on the bypass combo modules because they are created automatically when the bypass combo module initializes.

Other ports on the bypass combo module or on other line cards or modules can be designated as inline-network type ports. These ports will automatically be unprotected.

You also do not need to create the protected inline networks on the bypass combo modules because they are created automatically when the bypass combo module initializes. In addition, you cannot delete the protected inline networks on the bypass combo modules.

On GigaVUE-HC1, inline networks can also be configured on the copper TAP-HC1-G10040 module.

Once an inline network is created, it is fully configurable regardless of whether or not it is associated with an inline tool or an inline tool group through a map or map passall.

Use the Inline Networks page to display the configuration of inline networks as well as the current state of the inline bypass solution. The overall state of the inline bypass solution consists of the inline networks involved in the solution as well as the associated inline maps and inline tools. For details, refer to the following:

- [Display Current State of Inline Bypass Solution](#)
- [How to Use SNMP Polling to Obtain Inline Network State](#)
- [SNMP Notification of Forwarding State Change](#)
- [How to Use Syslog to Obtain Inline Network State](#)

For details on the parameters of inline networks, refer to the following:

- [Network Port Link Status Propagation Parameter](#)
- [Traffic Path Parameter](#)
- [Physical Bypass Parameter.](#)

Network Port Link Status Propagation Parameter

One of the parameters of inline networks is link status propagation, which controls the behavior of the link status for the inline network ports involved in a given inline network. The default is enabled.

When enabled, an inline network link failure on one side of the inline network will be propagated to the other side. For example, when the link is lost on one side of the network such that traffic cannot be sent to the inline tools, the link on the opposite side of the network is also brought down.

When the link is restored to the side that originally went down, the link will automatically be restored to the other side of the network. The GigaVUE node will not forward packets to the inline tools until the link is restored on both sides.

Link status propagation is enabled by selecting **Link Failure Propagation** when configuring an inline network port.

Traffic Path Parameter

One of the parameters of inline networks is **Traffic Path**, which specifies the path of the traffic received at an inline network port. The values are as follows:

- **Drop**—no traffic is exchanged through the inline network ports. All traffic to these ports is dropped. No traffic is forwarded to or from the inline tool or tools. No traffic is passed from inline network port A to inline network port B or from inline network port B to inline network port A. Refer to [Figure 21Traffic Path of Drop](#).
- **ByPass**—there are two cases for bypass as follows, depending on the inline map configuration:
 - If there are no inline maps associated with the inline network or if the set of inline maps associated with the inline network guarantees that no traffic is dropped when the traffic path is set to **To Inline Tool**, then setting the traffic path to **Bypass** leads to the following: all traffic arriving at the side A inline network port is forwarded to the side B inline network port and all traffic arriving at the side B inline network port is forwarded to the side A inline network port through a logical bypass. Refer to [Figure 22Traffic Path of Bypass](#) for a simple scenario involving a map-passall in which there is no possibility of traffic drops.
 - If the set of inline maps associated with the inline network involves some traffic drop when the traffic path is set to **To Inline Tool**, then setting the traffic path to **Bypass** leads to the following: all traffic arriving at the side A inline network port that would not have been dropped with traffic path set to **To Inline Tool** is forwarded to the side B inline network port and all traffic arriving at the side B inline network port that would not have been dropped with traffic path set to **To Inline Tool** is forwarded to the side A inline network port through a logical bypass. In other words, packets that were dropped when the traffic path was set to **To Inline Tool** will also be dropped when set to **Bypass**.

In either of these two bypass cases, no traffic is forwarded to the inline tool or tools.

- **ByPass with Monitoring**—there are two cases for monitoring as follows, depending on the inline map configuration:
 - If there are no inline maps associated with the inline network or if the set of inline maps associated with the inline network guarantees that no traffic is dropped when the traffic path is set to **To Inline Tool**, then setting the traffic path to **ByPass with Monitoring** leads to the following: all traffic is forwarded as for bypass, but a copy of the traffic is forwarded to the inline tool or tools according to the configured maps between the inline network and the inline tool or tools. Refer to [Figure 23Traffic Path of Monitoring](#) for a simple scenario involving a map-passall in which there is no possibility of traffic drops.
 - If the set of inline maps associated with the inline network involves some traffic drop when the traffic path is set to **To Inline Tool**, then setting the traffic path to **ByPass with Monitoring** leads to the following: all traffic that would not have been dropped with traffic path set to **To Inline Tool** is forwarded as for bypass, but a copy of the traffic is forwarded to the inline tool or tools according to the configured maps between the inline network and the inline tool or tools.

In either of these two monitoring cases, no traffic is taken from the inline tools.

- **To Inline Tool**—traffic received at the inline network ports is forwarded according to the following factors:
 - the configured maps between the inline network and the inline tools
 - the failover actions of the inline tool or tools

- the health state of the inline tool or tools

Refer to [Figure 24Traffic Path of To-Inline-Tool](#).

The default is **Bypass**. This avoids any traffic loss when first configuring an unprotected inline network or when disabling physical bypass on a protected inline network.

[Figure 21Traffic Path of Drop](#) to [Figure 24Traffic Path of To-Inline-Tool](#) show the traffic path for a simple inline bypass solution with inline network ports (nA and nB), inline tool ports (tA and tB), and a map passall.




Figure 21 Traffic Path of Drop




Figure 22 Traffic Path of Bypass




Figure 23 Traffic Path of Monitoring




Figure 24 Traffic Path of To-Inline-Tool

Physical Bypass Parameter

One of the parameters of inline networks is physical bypass, which controls the state of the optical protection switch on the bypass combo module or copper TAP module when the module is powered on. The optical protection switch can have one of the following states:

- close—the fiber connected to the side A network port is passively coupled with the fiber connected to the side B port without any transceivers or switching fabric. Therefore, any traffic coming in on these fibers is exchanged between the two inline network ports without being noticed by the system.
- open—the fiber connected to the inline network ports is coupled through transceivers with the switching fabric that is under software control. Therefore, any traffic coming in on these fibers is subject to the traffic forwarding rules imposed by the current configuration as well as the current state of the inline tools.

When the bypass combo module or copper TAP modules is powered off, the optical protection switch is always in the close state.

When the bypass combo module or copper TAP modules is powered on, the state of the optical protection switch is as follows:

- the close state if **Physical Bypass** is set to enable (that is, selected on the configuration page)
- the open state if physical bypass is set to disable (that is, not selected on the configuration page)

Physical Bypass is enabled by default.

NOTE: Physical bypass only applies to protected inline networks.

Redundancy Profile Parameter

One of the parameters of inline networks is redundancy profile. A redundancy profile is used to configure Gigamon Resiliency for Inline Protection (GRIP). For more detailed information about GRIP, refer to the [Configure Gigamon Resiliency for Inline Protection](#).

To create a redundancy profile, do the following:

1. Select **Inline Bypass > Redundancies**.
2. Click New to open the Add Redundancy Profile page.
3. Type a name for the profile in the Alias field to help identify the profile.
4. Select a stack port by clicking the **Signaling Port** field and selecting an available port.

The Signaling Port field specifies the used to signal the state of the two GigaVUE-HC3 nodes to each other. The ports provide the mechanism to detect loss of power in one of the GigaVUE-HC3 nodes.

5. Click in the **Protection Role** field and select a role.

The protection role specifies the role of the GigaVUE-HC3, as primary, secondary, or suspended. The default is suspended. When suspended, the protection role is on hold. Changing a GigaVUE-HC3 from the primary role to the suspended role can be used to manually force one of the GigaVUE-HC3 nodes to become active. The suspended role is also used when performing maintenance.

6. Click **Save**.

To add the redundancy profile to the inline network, select the profile from the Redundancy Profile drop-down list on the Inline Network configuration page.

Display Current State of Inline Bypass Solution

Inline networks, inline tools, and inline maps work together to form an inline bypass solution. The inline bypass solution has an overall state, which can change in response to hardware conditions and user configuration.

Several factors make up the overall state of an inline bypass solution, as follows:

- The physical bypass configuration of the inline network is protected.
- The inline network configuration, in particular, if physical bypass is enabled or disabled, what traffic path is configured, and if link failure propagation is enabled or disabled.
- The inline tool configuration, in particular, if the state of the inline tool is enabled or disabled, if there is a heartbeat profile configured and if the heartbeat is enabled or disabled, and what failover actions are configured.

- The inline tool group configuration, in particular, if the state of the inline tool group is enabled or disabled, what failover mode is configured, what failover action is configured, and what number of minimum healthy tools is configured.
- The status of the links attached to the inline network ports and inline tool ports.

[Forwarding States and Determining Factors](#) describes each forwarding state and the factors determining that state.

Whenever link failure propagation is configured as false (disabled), the inline network port status reflects the status of the respective far-end ports. Only when link failure propagation is configured as true (enabled) does this behavior change. Refer to the note under the Forwarding states for DISABLED and DISCONNECTED in [Forwarding States and Determining Factors](#).

Table 16: Forwarding States and Determining Factors

physical-bypass	traffic-path	Status of far-end ports connected to inline network ports	Status of inline tool side	Description
enable	drop, bypass, monitoring, or to-inline-tool	any status	any status	Forwarding state—PHYSICAL BYPASS If physical bypass is enabled, all traffic is exchanged directly between side A network and side B network without any monitoring by the GigaVUE node. Only applies to protected inline networks.
disable	drop	any status	any status	Forwarding state—DISABLED If the inline network is configured with a traffic path of drop, no traffic is exchanged between side A network and side B network because all packets coming to the inline network ports are dropped. Note: If the inline network is configured with link failure propagation set to true (enabled), the status of the inline network ports will be determined by the status of the far-end ports connected to the inline network ports. If both far-end ports are up, then both inline network ports

physical-bypass	traffic-path	Status of far-end ports connected to inline network ports	Status of inline tool side	Description
				will be up. If any far-end ports are down, then both inline network ports will be down.
disable	bypass, monitoring, or to-inline-tool	at least one far-end port is down	any status	Forwarding state—DISCONNECTED When one of the inline network ports is down due to a link down caused by a far-end device, no traffic is exchanged between side A network and side B network. Note: If the inline network is configured with link failure propagation enable as true, the status of both inline network ports will be down. That is, if only one inline network port is down, the other will be brought down.
disable	bypass	both far-end ports are up	any status	Forwarding state—FORCED BYPASS If the inline network is configured with a traffic path of bypass, all traffic that would not have been dropped when the traffic path is set to to-inline-tool is exchanged directly between side A network and side B network through the switching fabric.
disable	monitoring	both far-end ports are up	any status	Forwarding state—FORCED BYPASS WITH MONITORING If the inline network is configured with a traffic path of monitoring, all traffic that would not have been dropped when the traffic path is set to to-inline-tool is exchanged directly between side A network and side B network through the switching fabric. If there is any inline map configured for the inline network, a copy of the respective traffic is directed to the respective inline tools according to the configured inline maps.

physical-bypass	traffic-path	Status of far-end ports connected to inline network ports	Status of inline tool side	Description
disable	to-inline-tool	both far-end ports are up	all inline tools involved in the inline bypass solution are operating as expected	Forwarding state—NORMAL If all inline tools involved in the inline bypass solution are enabled, all inline tool ports are up, and the inline tools operating with heartbeat protocol enabled have a heartbeat status of up, traffic flows as desired according to the configuration of the inline maps. Note: If there are no inline maps, setting the traffic path of a protected fiber inline network to to-inline-tool results in a NORMAL forwarding state, but the traffic sent to the source inline network ports will be dropped because there are no inline maps specifying destination tool ports to which to forward the traffic.
disable	to-inline-tool	both far-end ports are up	at least one inline tool involved in the inline bypass solution is disabled or associated with any inline tool port that is down or associated with the heartbeat protocol in a down state	Forwarding state—FAILURE-INTRODUCED BYPASS All traffic is exchanged directly between side A network and side B network through the switching fabric as a result of an inline tool port failure or heartbeat failure or inline tool being disabled that led to the network-level bypass due to the configured failover actions.
disable	to-inline-tool	both far-end ports are up	at least one inline tool involved in the inline bypass solution is disabled or	Forwarding state—FAILURE-INTRODUCED DROP No traffic is exchanged between side A network and side B network. All packets coming to the inline network ports are dropped as a result of an

physical-bypass	traffic-path	Status of far-end ports connected to inline network ports	Status of inline tool side	Description
			associated with any inline tool port that is down or associated with the heartbeat protocol in a down state	inline tool port failure or heartbeat failure or inline tool being disabled that led to the network-level drop due to the configured failover actions.
disable	to-inline-tool	both far-end ports are up	at least one inline tool involved in the inline bypass solution is disabled or associated with any inline tool port that is down or associated with the heartbeat protocol in a down state	Forwarding state—NETWORK PORTS FORCED DOWN No traffic is exchanged between side A network and side B network. The inline network ports are forced down as a result of an inline tool port failure or heartbeat failure or inline tool being disabled that led to the network ports being forced down due to the configured failover actions.
disable	to-inline-tool	both far-end ports are up	at least one inline tool involved in the inline bypass solution is disabled or associated with any inline tool port that is down or associated with the heartbeat protocol in a down state	Forwarding state—ABNORMAL Any condition of the inline bypass solution that does not meet the criteria of the forwarding states listed in this table. At least one inline tool passes traffic as desired. There are many scenarios that can lead to the abnormal forwarding state, for example, when one inline tool member of an inline tool group has failed, but the number of remaining healthy tools is still above the minimum required number of healthy tools.

How to Use SNMP Polling to Obtain Inline Network State

The overall inline network state can also be obtained through SNMP polling using an SNMP-compliant network management application or a MIB browser. The names of the MIB files that need to be loaded in order to poll the inline network state are: GIGAMON-COMMON-SMI and GIGAMONINLINEBYPASS.

The inline network forwarding states are described in [Forwarding States and Determining Factors](#). Some of the states have slightly different names in SNMP.

SNMP Notification of Forwarding State Change

Use the following steps to configure a notification that will be sent when the inline bypass forwarding state changes:

1. Select **Settings > SNMP Traps**.
2. Click **Trap Settings**. The Edit SNMP Trap Settings page displays.
3. On the Edit SNMP Trap Settings page, select **Enable** for **Inline Bypass Forwarding State Change**.
4. Click **Save**.

How to Use Syslog to Obtain Inline Network State

The overall inline network state can also be obtained through syslog.

When the inline network state change is triggered by a configuration change, a sample syslog message is as follows:

```
IBFE STATE CHANGE inline network default_inline_net_1_1_1 status NORMAL changed to PHYSICAL BYPASS triggered by config
```

When the inline network state change is triggered by a link status change or heartbeat status change, a sample syslog message is as follows:

```
IBFE STATE CHANGE inline network inNet1 status NORMAL changed to FAILURE INTRODUCED BYPASS triggered by 1/x10
```

In the syslog messages, IBFE indicates the Inline Bypass Failover Engine.

The inline network forwarding states are described in [Forwarding States and Determining Factors](#).

Inline Network Groups

Use the **Inline Network Group** configuration page to configure an inline network group. An inline network group is an arrangement of multiple inline networks that share the same inline tool or tools. Use this page to specify the list of inline networks in the inline network group.

The inline network ports that make up the inline networks participating in the inline network group are always in pairs, running at the same speed, on the same medium (fiber or copper). All inline network ports of the inline network group must be on the same GigaVUE-HC3, or GigaVUE-HC1 or GigaVUE-HC1-Plus node. The inline networks participating in the inline network group can be different speeds and different mediums.

An inline network group can share an inline tool, or tools in an inline tool group or inline tool series. Many-to-one means from an inline network group to an inline tool and is shown in [Figure 25 Inline Network Group Many-to-One](#). Many-to-many means from an inline network group to an inline tool group or inline tool series.




Figure 25 *Inline Network Group Many-to-One*

When an inline tool or inline tool group is configured between several pairs of networks, each pair operates independently of the others. The traffic coming from the inline tool or tools must be segregated into individual substreams according to the traffic source.

For example, the traffic that came into the GigaVUE node on inline network port net_A_3 must be sent through inline network port net_B_3 (with net_A_3 and net_B_3 belonging to the same inline tool that is the third member of the inline tool group). Refer to [Figure 26 Inline Tool Inserted in Multiple NetA-NetB Links](#).



Figure 26 *Inline Tool Inserted in Multiple NetA-NetB Links*

To accomplish the segregation, packets are tagged coming from an individual side A inline network port before they are sent to the inline tool or tools. When tagged packets come back from the inline tool or tools, the tag identifies the respective side B inline network port through which the packets should be sent. The tags are removed before the packets are sent through the side B inline network ports.

Traffic is guided to a particular inline network through internal VLAN ID tagging. This VLAN tagging affects packets only on their way from inline tool ports to the attached inline tool and back from the attached inline tool to the inline tool ports. The packets sent out from inline network ports remain untagged. Refer to [Figure 25](#) *Inline Network Group Many-to-One*.

NOTE: Internal VLAN ID tagging creates hidden VLAN ID tags. Explicit VLAN ID tagging is not needed, however starting in software version 4.6, explicit VLAN tagging is also available. Refer to [Configurable VLAN Tagging](#).

Inline Tool Sharing

Inline network groups require inline tool sharing to be enabled on the inline tool or the members of the inline tool group or inline tool series specified in the inline map.

When shared is enabled (true), the inline tool can receive traffic from multiple sources (the inline networks in the inline network group) and can be used in a map in which the source is an inline network group.

An inline tool group or inline series does not have its own shared setting. The shared setting is derived from the inline tools. Therefore all the members in an inline tool group or inline series must have the same setting. For example, if an inline tool group has three inline tool members, the shared setting of all three inline tools must be the same.

Configurable VLAN Tagging

Explicit VLAN tagging for inline network groups can be configured. For example, you can use VLAN tags for managing policies. A mixture of internal and explicit VLAN tags is supported.

The VLAN tags are configured on the ports of inline networks. They can be configured at any time, but are only applied when the inline networks are part of an inline network group. Across the inline network group, the VLAN tags must be unique; however, both ports of an inline network can have the same VLAN tag. Refer to [Tools in Bridge Mode](#).

An error message is displayed if the same VLAN tag is used for more than one inline network in an inline network group.

Refer also to [Ingress and Egress VLAN](#).

NOTE: For inline SSL decryption, the inline network group does not support ingress VLAN tagging on the member links.

For out-of-band maps from inline network group ports or inline tool ports mapped from an inline network group, the out-of-band tool ports will receive the following:

- tagged packets, if they originally come from an inline network port with an ingress VLAN tag configured
- untagged packets, if they originally come from an inline network port without an ingress VLAN tag configured

Add VLAN Tag

The following are the steps for adding a VLAN Tag.

1. Go to **System > Ports > Ports > All Ports**.
2. Select the port to configure as an inline network port and click Edit.
3. Set the following parameters to configure an inline network port with VLAN tagging:
 - Select **Inline Network** or **Network** for Type.
 - Enter a VLAN ID in the **VLAN Tag** field.
4. Click **Save**.

Tools in Bridge Mode

The same VLAN tag can be assigned to both ports in an inline network port pair.

The following example configures the same ingress port VLAN tag on the net-a and net-b ports of an inline network. When the net-a and net-b ports have the same VLAN tag, an inline tool will send packets back to the network from which it came.

1. Go to **System > Ports > Ports > All Ports**.
2. Configure the net-a port.
 - a. Select the port (for example, 1/1/x17) and click **Edit**.
 - b. Enter inline-network-port-a in the **Alias** field.
 - c. Select **Inline Network** for **Type**.
 - d. Enter 123 in the VLAN Tag field.
 - e. Click **Save**.
3. Configure the net-b port.
 - a. Select the port (for example, 1/1/x18) and click **Edit**.
 - b. Enter inline-network-port-a in the **Alias** field.
 - c. Select **Inline Network** for **Type**.
 - d. Enter 123 in the VLAN Tag field.
 - e. Click **Save**.
4. Configure the Inline Network.
 - a. Select Inline **Bypass > Inline Networks**.
 - b. Click **New**.
 - c. Enter an alias in the **Alias** field.
 - d. Select port 1/1/x17 for Port A.
 - e. Select port 1/1/x18 for Port B.
 - f. Click **Save**.

Inline Tools

There are two meanings to the term inline tool. The inline tool software construct consists of a pair of inline tool ports plus the inline tool attached to the ports. The software construct has attributes that are configured on the GigaVUE® HC Series node.

The term inline tool also refers to the pass-through device itself that performs packet inspection and selective forwarding, such as an Intrusion Protection System (IPS). This is a physical device, external to the GigaVUE node.

Use the Inline Tool page to configure the inline tool software construct. An inline tool consists of inline tool ports, always in pairs, running at the same speed, on the same medium (fiber or copper). The inline tool ports must be on the same GigaVUE-HC3, or GigaVUE-HC1 or GigaVUE-HC1-Plus node. The inline tool ports must also be on the same GigaVUE-HC3, or GigaVUE-HC1 or GigaVUE-HC1-Plus node as the inline network ports.

Inline tool ports are ports to which inline tools are attached. Inline tool ports are always in pairs, and must have the same line rate. Inline tool ports must be on the same GigaVUE-HC1 node as the inline network ports.

An inline tool consists of a pair of inline tool ports plus the inline tool attached to the ports. Refer to [Figure 27 Inline Tool and Inline Tool Ports](#).




Fig. 3 Inline Tool

Figure 27 *Inline Tool and Inline Tool Ports*

Use the Inline Tools page to display of configuration of inline tools.

Inline Tool Failover Action

One of the parameters of inline tools is failover action, which controls the action taken when a tool is unhealthy or in response to a failure of an inline tool. You can configure one of the following failover actions:

- **ToolBypass**—when the inline tool fails, the traffic that normally was directed to the inline tool is redirected to the bypass path. Use this failover action for configurations involving multiple inline tools associated with an inline network or inline network group using rule-based maps. For configurations using map passalls, tool-bypass is the same as network-bypass.
- **NetworkBypass**—when the inline tool fails, all traffic that would not have been dropped when the inline network or networks had a NORMAL forwarding state is directed to the bypass path. That is, all such traffic arriving at the side A inline network port or ports is forwarded to the side B inline network port or ports and all traffic arriving at the side B inline network port or ports is forwarded to the side A inline network port or ports.

- **ToolDrop**—when the inline tool fails, the traffic that normally was directed to the inline tool is dropped. Use this failover action for configurations involving multiple inline tools associated with an inline network or inline network group using rule-based maps. For configurations using map passalls, tool-drop is the same as network-drop.
- **NetworkDrop**—when the inline tool fails, all traffic coming to the respective inline network (or inline network group) is dropped.
- **NetworkPortForcedDown**—when the inline tool fails, the inline network ports of the respective inline network (or inline network group) are forced down.

NOTE: However, after the device reload, the inline network ports become operationally up, whereas the inline tool is still down. You need to manually force down the inline network ports.

The default is **ToolBypass**.

The bypass path is between side A and side B of the inline network ports.

Inline Tool Failover Action with Inline Flow Mapping®

When the inline bypass solution uses inline flow mapping, the failover actions of inline tools are as follows:

- **ToolBypass**—when the inline tool fails, the traffic that normally was directed to the inline tool is redirected to the bypass path. The traffic going to the healthy inline tools (through rule-based maps) remains unchanged.
- **ToolDrop**—when the inline tool fails, the traffic that normally was directed to the inline tool is dropped. The traffic going to the healthy inline tools (through rule-based maps) remains unchanged.

Inline Tool Recovery Mode

An inline tool detects failures in the traffic path between port pairs and automatically diverts traffic away to avoid disruption. After an inline tool goes down, the following modes specify how to bring it back up after it has recovered:

- **automatic**—Specifies automatic recovery, which redirects traffic back to the inline tool as soon as it has recovered from all faulty conditions.
- **manual**—Specifies manual recovery, which lets you control when to put an inline tool back into service after the tool has recovered. For example, you may wait for a maintenance window to return the inline tool to service.

The default is automatic.

By selecting the tool and selecting the Recover button, users can set the recovery of the inline tools to manual or automatic.

Refer to [Figure 28Automatic and Manual Inline Tool Recovery from Failover](#) for automatic and manual inline tool recovery from failover.




Figure 28 Automatic and Manual Inline Tool Recovery from Failover

The left side of [Figure 28Automatic and Manual Inline Tool Recovery from Failover](#) shows an inline tool that has failed and the bypass failover action has been executed.

Automatic recovery is shown on the right side of [Figure 28Automatic and Manual Inline Tool Recovery from Failover](#). When the inline tool recovers, traffic is automatically directed back to it.

Manual recovery is shown in the center of [Figure 28Automatic and Manual Inline Tool Recovery from Failover](#). When the recovery mode is configured as manual, an SNMP notification, if enabled, will send a notification when the inline tool is ready to be put back into service. The failover action, in this case, bypass, will be enforced until you manually put the inline tool back into service.

When the recovery mode is configured as manual, an SNMP notification, when enabled, will notify you when the inline tool is ready to be put back into service.

Use the following steps to configure notification that will be sent when the inline tool is ready:

1. Select **Settings > SNMP Traps**.
2. Click **Trap Settings**. The Edit SNMP Trap Settings page displays.
3. On the Edit SNMP Trap Settings page, select **Enable** for **Inlinetool Recovery**.
4. Click **Save**.

The default for **Inlinetool Recovery** is disabled.

Use the following steps to put an inline back in service when the recovery mode is manual and the inline tool has an operational state of ready.

1. Select Inline **Bypass > Inline Tools**.
2. On the Inline Tools page, select the inline tool.
3. Click **Recover**.

NOTE: Also use **Recover** after the GigaVUE® HC Series node is reloaded or rebooted (even though the inline tool has not failed). Issue the **Recover** on all the inline tools that are configured with manual recovery after a reload or **Actions > Shut Down** from the Chassis page on a selected card in Chassis Table View followed by **Actions > Start Up**. Refer to the “Chassis” section in the *GigaVUE Administration Guide* for more details.

NOTE: In certain circumstances while using classic inline, if the recovery mode is set to "manual" and the inline-tool is in “ready” state, traffic may still be forwarded to the inline-tool ports . This occurs because the recovery state shifts from “ready” to “up” when "inline-tool disable/enable" is used. When the recovery mode for the inline-tool is "manual" and is in “ready” state, do not disable/enable the inline-tool . Instead, disable/enable the port.

The **Inline Tool** page displays the operational state of each inline tool as up, down, disable, or ready. Refer to [Aggregate Inline Tool States](#) for detailed descriptions of the states.

Inline Tool Sharing Mode

Inline tool sharing mode specifies how an inline tool is going to be shared as follows:

- **Enable**—Specifies that the inline tool is going to be shared by different sources.
- **Not enabled**—Specifies that the inline tool will not be shared by different sources.

The default is **not enabled**.

When sharing is enabled, the inline tool can receive traffic from multiple sources (the inline networks in the inline network group) and can be used in a map in which the source is an inline network group.

An inline tool group or inline series does not have its own shared setting. The shared setting is derived from the inline tools. Therefore all the members in an inline tool group or inline series must have the same setting. For example, if an inline tool group has three inline tool members, the shared setting of all three inline tools must be the same.

When an inline tool has sharing mode enabled, the traffic will be VLAN tagged. The connected inline device is expected to receive VLAN tagged packets. When an inline tool does not have sharing mode enabled, the extra VLAN tag is not added.

Go to **Inline Bypass > Inline Tools** and select an inline tool or click **New**. Under Configuration, Inline tool sharing mode, select Enable.

How to Use SNMP Polling to Obtain Inline Tool State

The inline tool state can also be obtained through SNMP polling using an SNMP-compliant network management application or a MIB browser. The names of the MIB files that need to be loaded in order to poll the inline tool state are: GIGAMON-COMMON-SMI and GIGAMONINLINEBYPASS.

The inline tool states are described in [Table 17: Aggregate Inline Tool States](#). They are an aggregate of the inline tool port statuses and the heartbeat status.

Table 17: Aggregate Inline Tool States

Status	Description
up	The heartbeat is up and all the inline tool ports are up. The inline tool is operational and is forwarding traffic to the tool.
down	Either the heartbeat is down or one or more of the inline tool ports are down or disabled. The inline tool is in a failed state and is not ready to recover. The tool is not receiving any traffic.
disable	The inline tool is disabled.
ready	The inline tool is in a failed state but is ready to recover. The tool is not receiving any traffic.

Heartbeats

When heartbeat packets are sent to an inline tool, they are expected to be received back when the inline tool is healthy. Negative heartbeat packets complement heartbeat packets to verify the health of inline tools. When negative heartbeat packets are sent to an inline tool, they are not expected to be received back when the inline tool is healthy.

When some inline tools begin to fail, they allow packets though that should have been dropped. A negative heartbeat detects such a failure by sending a packet that should not pass through an inline tool. If the negative heartbeat packet passes through an inline tool, the tool is deemed to have failed. Therefore, a negative heartbeat packet received back from an inline tool indicates a tool failure.

Heartbeat packets and negative heartbeat packets can be used in any combination: heartbeat only, negative heartbeat only, neither, or both.

Rules and Notes

Keep in mind the following rules and notes when you configure the heartbeat mechanism between GigaVUE nodes:

- The heartbeat mechanism is supported in GigaVUE-HC1, GigaVUE-HC3 and GigaVUE-HC1-Plus devices.
- Negative heartbeat profiles are supported.
- Both standard and custom heartbeat packets are supported.
- Both protected and unprotected inline network pairs are supported.
- The heartbeat mechanism is not supported on inline netlag and inline network group components. However, it is supported on the underlying inline networks.
- The heartbeat mechanism is not supported for classic inline bypass solution.
- The minimum timeout for heartbeat sessions is 200 milliseconds, if the tool type is configured as GigaVUE Node.
- The size of a custom heartbeat packet must be less than 128 bytes.

Heartbeat Profiles

A heartbeat profile supports health monitoring of inline tools and inline networks. A heartbeat profile is a group of attributes applied to an inline tool or inline network to configure its heartbeat operation. Multiple inline tools can share a heartbeat profile. However, for inline network, heartbeat is supported only within an inline network pair; it is not supported on inline netlag or inline network groups.

To display the configured heartbeat profiles, select **Inline Bypass > Heartbeats** to open the Heartbeats page. An example is shown in [Figure 29Heartbeats Page with Heartbeat Profiles](#).

Heartbeats

Heartbeat Profiles

Statistics

New

Clone

Edit

Delete

☐	Alias	Type	Packet Fo...	Custom P...	Direction	Period (ms)	Timeout (...)	Recovery...	Re...	⊕
☐	DBL_hbpr...	regular	arp		aToB	900	200	5	3	
☐	DBL_hbpr...	regular	arp		both	900	200	5	3	
☐	DBL_hbpr...	regular	arp		bToA	900	200	5	3	
☐	default	regular	arp		both	1000	500	30	3	

Figure 29 *Heartbeats Page with Heartbeat Profiles*

The Heartbeats page includes a default heartbeat profile that has the following settings:

- Alias—default
- Type—Regular
- Packet Format—arp or custom
- Custom Packet—URL from which a PCAP file can be imported
- Direction—bi-directional
- Period—1000 milliseconds
- Timeout—500 milliseconds
- Recovery Period—30 seconds
- Retries—3

The highest frequency heartbeat that can be configured is as follows:

- period—30 milliseconds
- timeout—20 milliseconds
- retry-count—0

The heartbeat mechanism supports the maximum number of inline tools, at the highest frequency, which is 48 on the GigaVUE-HC3, 16 on the GigaVUE-HC1, and 48 in GigaVUE-HC1-Plus.

To display the heartbeat profile associated with an inline tool and the heartbeat status, open the Inline Tools page by selecting **Inline Bypass > Inline Tools**. There is also a combined heartbeat status, which combines heartbeat and negative heartbeat statuses and indicates the tool health used for inline tool failover actions or SNMP traps. The combined heartbeat status is the combination of both heartbeat statuses. If both are configured and one is down, the combined status will be down.

Standard Heartbeat

The standard heartbeat is a packet sent by the GigaVUE node that passes through the inline tool to verify that it is passing traffic, even if the link is *up*. If the packet is not passed through the tool, the tool is considered to have failed and a bypass action is triggered.

Even when the tool is considered down, heartbeat packets continue to be sent so that the bypass action can be reversed when the tool is healthy again.

Heartbeats are sent bidirectionally to the inline tool.

Standard or Custom Heartbeat Packet

The format of the heartbeat packet can be the standard ARP packet or a custom packet. For a custom packet, you must provide a URL from which a PCAP file can be imported. If the PCAP file contains several packets, the first packet present in the file is taken as the heartbeat packet. The size of a custom heartbeat packet must be less than 128 bytes.

NOTE: The system will overwrite the MAC address portion of the custom heartbeat packet.

If the inline tool through which the heartbeat packets are passed is expecting IPv6 traffic exclusively, you must select a custom heartbeat packet.

Custom heartbeat packets are needed in situations in which inline tools do not reliably pass standard ARP packets. For example, if an inline tool is configured to pass only IPv6 traffic, an ICMPv6 ARP packet might be appropriate.

If a custom heartbeat packet is specified, the **Heartbeats** page displays the name of the PCAP file from which it was imported.

Detect Inline Tool Failure

The health of the inline tool is critical to the proper handling of traffic. An inline tool is determined to have failed if:

- link is lost to the tool
- inline heartbeat fails

When the tool fails in one direction, it is considered to have failed in both directions. For example, if the heartbeat stops flowing in the northbound direction, neither northbound or southbound packets are sent to the tool.

Negative Heartbeat Profiles

A negative heartbeat profile is a group of attributes applied to an inline tool to configure its negative heartbeat operation. Multiple inline tools can share a negative heartbeat profile. The content of a negative heartbeat is configurable using the same PCAP file mechanism as for a custom heartbeat packet.

Use the **Add Heartbeats** page shown in [Heartbeats](#) to configure a negative heartbeat profile. A negative heartbeat profile can be created by selecting **Negative** in the **Type** field. The profile will have the following settings:

- **PacketFormat** set to **Custom**
- **Direction** set to **Bi-directional**
- **Period** set to **1000** (period is specified in milliseconds)
- **Recover Time** set to **30** (recovery is set in seconds)

You must provide a valid PCAP file when Packet Format is set to Custom before the negative heartbeat profile can be applied to an inline tool.

When a negative heartbeat is configured, the system will send packets specified by the **Custom Packet Format**, in the time specified by **Period**, in the direction specified by **Direction**. The inline tool absorbs the negative heartbeat packets until the number of seconds specified by **Recovery Time** has passed. **Recovery Time** specifies the number of seconds of not receiving negative heartbeat packets in order for the inline tool to be declared healthy.

The negative heartbeat mechanism supports the maximum number of inline tools, which is 48 on the GigaVUE-HC3, 16 on the GigaVUE-HC1, and 48 on the GigaVUE-HC1-Plus.


Use the Heartbeat page to display the configured negative heartbeat profiles. Use the Inline Tool page to display the negative heartbeat profile associated with an inline tool, the negative heartbeat status, and the counters of received packets in each direction. There is also a combined heartbeat status, which combines heartbeat and negative heartbeat statuses and indicates the tool health used for inline tool failover actions or SNMP polling. The combined heartbeat status is the combination of both heartbeat statuses. For example, if both are enabled and both are up, the combined status is up. If both are enabled and one is down, the combined status is down.

Heartbeat Support Between GigaVUE Nodes

The heartbeat mechanism focuses on providing extended heartbeat capability to monitor the following types of devices when the devices are connected to the inline-tool pair of ports as a tool:

- GigaVUE nodes
- GigaVUE nodes with GigaSMART operations configured

Following figure illustrates an example of a topology with GigaVUE nodes placed at three different layers.



-----> Heartbeat (Access Layer -> Distribution Layer -> Access Layer)

-----> Heartbeat (Distribution Layer -> Tool -> Distribution Layer)

——— Traffic

GigaVUE HC Series means it can be either GigaVUE-HC3, GigaVUE-HC1 or GigaVUE-HC1-PLUS

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

The GigaVUE node at the distribution layer distributes the traffic from the access layer to the tool layer.

The GigaVUE node at the tool layer acts as the SSL decryption tool.

In this topology, heartbeats are essential to monitor the traffic integrity at the distribution layer and to ensure automatic failover in case of a tool failure. In the access layer device, the ports that are connected to the distribution layer device are configured as inline tool ports because they face the tool side. In the distribution layer device, the ports that are connected to the access layer device are configured as inline network ports because they face the network side of the topology. The heartbeat packets will be sent from the inline tool port pair of the access layer device to the inline network port pair of the distribution layer device. If the forwarding state of the inline network pair is normal, the heartbeat packet is sent back to the inline tool port pair of the access layer device. Else, the packet is dropped.

The heartbeat mechanism is extended to support the GigaVUE node at the distribution layer to monitor the GigaVUE node that acts as a tool at the tool layer. In the distribution layer device, the ports that are connected to the tool layer device are configured as inline tool ports. In the tool layer device, the ports that are connected to the distribution layer device are configured as inline network ports. The heartbeat packets that are sent from the distribution layer device to the tool layer device will monitor the availability of both, the tool layer device and its GigaSMART engines.

Troubleshoot Heartbeat Failures

In your topology, you may sometimes notice heartbeat packets getting dropped at the inline network or the heartbeat status of the inline tool is not up. For details about how to isolate and troubleshoot such failures, refer to the following sections:

- [Heartbeat Status of Inline Tool is Not Operationally Up](#)
- [Heartbeat Packets are Dropped](#)

Configure a Heartbeat Profile

Use the Add Heartbeats page shown in the below figure to configure a regular heartbeat profile or a negative heartbeat profile by selecting **Inline Bypass > Heartbeats > Heartbeats**, and then clicking **New**.

Add HeartBeat Profile

OK

Cancel

Alias	Alias
Type	Type
Packet Format	Packet Format
Direction	Direction
Timeout	20-1000
Period	30-5000
Recovery Time	5-60
Retries	0-5

View Heartbeat Profile Statistics

Statistics about heartbeat profiles are displayed on the heartbeats statistics page. To open the Statistics page, select **Inline Bypass > Heartbeats > Statistics**. The page shows the following information:

Column	Description
Alias	The alias of the inline tool or inline network to which the heartbeat profile is associated with.
Inline Type	The inline type to which the heartbeat profile is associated with. The valid values are Inline Tool or Inline Network.
Heartbeat Profile	The alias of the heartbeat profile
Heartbeat Type	Indicates the type of heartbeat profile: Regular or Negative
A to B Packets	The number of packets sent/received from port A to port B of the inline tool.
B to A Packets	The number of packets sent/received from port B to port A of the inline tool.
A to B Packet Drops	The number of packets dropped from port A to port B of the inline tool or inline network.
B to A Packet Drops	The number of packets dropped from port B to port A of the inline tool or inline network.

You can clear the statistics for a specific heartbeat profile or all heartbeat profile. To clear the statistics for a heartbeat profile, select the required row, and then click **Clear**.

To clear the statistics for all heartbeat profiles, select **Clear All**, and then choose one of the following options:

- Select **Clear All Inline Network Heartbeat Stats** to clear the heartbeat statistics for all inline networks.
- Select **Clear All Inline Tools Heartbeat Stats** to clear the heartbeat statistics for all inline tools.
- Select **Clear All Negative Heartbeat Stats** to clear only the statistics for negative heartbeat profiles. The statistics for regular heartbeat profiles will remain.

Heartbeat Status after System Reload

Following a reload, there is a 5-minute delay for the system to stabilize before heartbeat packets are sent or received. During this delay, the heartbeat status is down.

Inline Tool Groups

Use the **Inline Tool Groups** configuration page to configure an inline tool group, which is an arrangement of multiple inline tools to which traffic is distributed based on hashing. In an inline tool group, traffic is shared. Each inline tool in the group receives a portion of the traffic. The distribution mechanism includes a way of dealing with failures of individual tools through traffic redistribution to the remaining healthy tools.

The inline tool ports that make up the inline tools participating in the inline tool group are always in pairs, running at the same speed, on the same medium (fiber or copper). All inline tool ports of the inline tool group must be on the same GigaVUE-HC3, or GigaVUE-HC1 and GigaVUE-HC1-Plus node, but can be on different modules on the node. The inline tool ports must also be on the same GigaVUE-HC3, or GigaVUE-HC1 and GigaVUE-HC1-Plus node as the inline network ports.

Inline tool groups can be configured as follows:

- non-redundant—multiple inline tools with no spare inline tool. Refer to [Figure 30 Inline Tool Group With No Spare, Non-Redundant](#).
- 1+1 redundancy—single inline tool with a spare inline tool. Refer to [Figure 31 Inline Tool Group With Spare, Redundant 1+1 Scenario](#).
- N+1 redundancy—multiple inline tools that are considered active, with a standby inline tool that is only used if one of the active inline tools fails. Refer to [Figure 32 Inline Tool Group With Spare, Redundant N+1 Scenario](#).

With 1+1 redundancy, an inline tool is paired with a standby tool. When there is a loss of link or a heartbeat failure on an active tool, the traffic will be sent to the standby tool with no loss. In addition, if the standby tool fails, you can configure what happens to the traffic in that case, such as drop it or forward it.

With N+1 redundancy, one tool is added to a group of N distributed inline tools. When any one of the N tools fails, the traffic from that tool is sent to the standby (or spare) tool with no loss. In addition, if the standby tool fails, you can configure what happens to the traffic in that case, such as redistribute it, send it to another tool, or declare a failure on the tool group.

Refer to [Figure 30 Inline Tool Group With No Spare, Non-Redundant](#) for a non-redundant inline tool group.




Fig.4 Inline-Tool-Group no-spare-tool scenario
(Only A-to-B traffic shown for diagram clarity)

Figure 30 *Inline Tool Group With No Spare, Non-Redundant*

Refer to [Figure 31 Inline Tool Group With Spare, Redundant 1+1 Scenario](#) for an inline tool group with a single inline tool and a spare inline tool configured. This is also referred to as 1+1 redundancy or N+1 redundancy where N equals 1. In [Figure 31 Inline Tool Group With Spare, Redundant 1+1 Scenario](#), traffic is only shown from A-to-B.




Figure 31 Inline Tool Group With Spare, Redundant 1+1 Scenario

Refer to [Figure 32](#) Inline Tool Group With Spare, Redundant N+1 Scenario for an inline tool group in an N+1 redundant scenario, in which N is greater than 1. In [Figure 32](#) Inline Tool Group With Spare, Redundant N+1 Scenario, traffic is only shown from A-to-B.




Figure 32 Inline Tool Group With Spare, Redundant N+1 Scenario

For details on the parameters of inline tool groups, refer to the following:

- [Inline Tool Group Failover Action](#)
- [Inline Tool Group Spare Inline Tool](#)
- [Symmetrical and Asymmetrical Hashing.](#)
- [Resilient Weighted Hashing](#)

Inline Tool Group Failover Action

One of the parameters of inline tool groups is the failover action, taken in response to a failure of an inline tool group, when the number of healthy inline tools in the inline tool group (including the spare inline tool, if configured) falls below the configured minimum. You can configure one of the following failover actions:

- **ToolBypass**—when the inline tool group fails, the traffic that normally was directed to the inline tool group is redirected to the bypass path. Use this failover action for configurations involving multiple inline tools or inline tool groups associated with an inline network or inline network group using rule-based maps. For configurations using map passalls, tool-bypass is the same as network-bypass.
- **NetworkBypass**—when the inline tool group fails, all traffic that would not have been dropped when the inline network or networks had a NORMAL forwarding state is directed to the bypass path. That is, all such traffic arriving at the side A inline network port or ports is forwarded to the side B inline network port or ports and all traffic arriving at the side B inline network port or ports is forwarded to the side A inline network port or ports.
- **NetworkPortForcedDown**—when the inline tool group fails, the inline network ports of the respective inline network (or inline network group) are forced down.
- **ToolDrop**—when the inline tool group fails, the traffic that normally was directed to the inline tool group is dropped. Use this failover action for configurations involving multiple inline tools or inline tool groups associated with an inline network or inline network group using rule-based maps. For configurations using map passalls, tool-drop is the same as network-drop.
- **NetworkDrop**—when the inline tool group fails, all traffic coming to the respective inline network (or inline network group) is dropped.

The default is tool-bypass.

The bypass path is between side A and side B of the inline network ports.

The failover action of all the inline tools specified by the inline tool list is overwritten by the failover mechanism of the inline tool group. This means that when a given inline tool specified by the inline tool list fails, the traffic originally directed to this inline tool is redirected to the spare inline tool (if one is configured and available) or handled according to the failover mode of the active tools, so long as the total number of healthy inline tools in the inline tool group is not smaller than the minimum required number of healthy inline tools.

When the total number of healthy inline tools in the inline tool group drops below the minimum required number of healthy inline tools, the failover action of the inline tool group determines the action to be taken.

Inline Tool Group Spare Inline Tool

One of the parameters of inline tool groups is a spare inline tool. If a spare is configured, the inline tool group becomes a redundant arrangement of inline tools. When the first failure occurs in a set of active inline tools, traffic will be forwarded to the spare with no loss, thus the spare will replace the failed tool in the active set.

The inline tools in the inline tool list are considered to be active inline tools. The traffic is hash-distributed over the active inline tools as long as all the inline tools are healthy. When one of the active inline tools fails, the spare inline tool takes the place of the failed inline tool and the new set operates as a new active set. If another inline tool fails, the traffic is redistributed according to the failover mode, as if there was no spare.

When the number of failed inline tools is such that the number of healthy inline tools is less than the minimum-group-healthy-size, the group heals itself by re-spreading the traffic over the healthy tools. When the number of healthy tools falls below the minimum-group-healthy-size, the failover action of the inline tool group takes place, while the failover action of the member inline tools is ignored.

The spare inline tool works with another parameter called release-spare-if-possible. When the inline tool that had been replaced with the spare inline tool recovers, the release-spare-if-possible parameter determines if the recovering inline tool is included in the active set of inline tools or if it becomes the new spare inline tool.

The default of the release-spare-if-possible parameter is disabled. Disabled means that even if the original inline tool recovers, the spare that replaced it will remain in the active set of inline tools. Enabled means that after the original inline tool recovers, the spare that replaced it will be released, if possible, from the active set of tools to become the spare again.

Configure the minimum-group-healthy-size and release-spare-if-possible parameters at the same time you configure the spare inline tool.

Symmetrical and Asymmetrical Hashing

One of the parameters of inline tool groups is hashing, which is used for distributing packets across a group of inline tools belonging to the inline tool group. The values for the hash parameter are as follows:

- **advanced**—Specifies symmetrical hashing, which is derived from the combination of packet fields based on the criteria selected for the advanced-hash algorithm configured under the **Ports > Port Groups > GigaStreams > Advanced Hash Settings** page.

For inline bypass applications, the most common choice of criteria for the advanced-hash algorithm is the combination of source IP and destination IP addresses. This produces a hash value that sends all traffic associated with the same session to the same inline tool in the inline tool group.

- **a-srcip-b-dstip**—Specifies asymmetrical hashing, which is derived from the source IP address for side A of the network and the destination IP address for side B of the network. This produces a hash value that sends all traffic associated with the same source address residing on side A to the same inline tool in the inline tool group, regardless of destination or session.
- **b-srcip-a-dstip**—Specifies asymmetrical hashing, which is derived from the destination IP address for side A of the network and the source IP address for side B of the network. This produces a hash value that sends all traffic associated with the same source address residing on side B to the same inline tool in the inline tool group, regardless of destination or session.

The default is **advanced**.

Use asymmetrical hashing if all traffic exchanged between a particular node on one side of the network and any nodes on the other side of the network that communicate with that node need to go to the same inline tool. The asymmetrical hashing options involve only source IP address (srcip) in one direction and only destination IP address (dstip) in the opposite direction. Bi-directional traffic, such as between a given user and all the Internet sites visited by the user, will be sent to the same inline tool in the group.

NOTE: When asymmetric hashing is configured, the portsrc and portdst packet fields are not included in the advanced-hash calculation for any GigaStream and inline tool groups across the GigaVUE node.

With symmetrical hashing, the inline network traffic path parameter can be configured to different values on the inline networks. With asymmetrical hashing, there is a restriction. Refer to [Asymmetrical Hashing Restrictions](#).

Refer to [Figure 33Asymmetrical Hashing](#) for asymmetrical hashing.




Figure 33 Asymmetrical Hashing

Use the hashing option **a-srcip-b-dstip** if the node is on side A of the network and the Internet is on side B. For example, Node A has IP address A. Traffic from Node A (from side A) will have IP address A. Traffic from side B (the Internet) destined for Node A, will have a destination of IP address A. This traffic will go to the same inline tool in the group.

If the network has the Internet on side A and the node on side B, use the hashing option **b-srcip-a-dstip**.

Asymmetrical Hashing Restrictions

The following are restrictions for asymmetrical hashing:

- If asymmetrical hashing is configured for the inline tool group, only rule-based maps or shared collector maps can be used to send traffic to the inline tool group. Inline map passalls cannot be used to send traffic to the inline tool group.
- For inline networks belonging to an inline network group, mapped to an inline tool group using asymmetrical hashing, the **Traffic Path** must be configured to the same value on all the inline networks, one of **Drop**, **Bypass**, **To Inline Tool**, or **ByPass with Monitoring**.

NOTE: For the inline networks belonging to an inline network group, mapped to an inline tool group using symmetrical hashing, the traffic path parameter can be configured to different values on the inline networks.

- If an inline network is involved in an inline map to an inline tool group configured with asymmetrical hashing, the inline network ports of the inline network cannot be used as the **Source** in any out-of-band maps. Also, if the traffic path parameter for the inline network is configured to **ByPass with Monitoring**, there will not be any bypass traffic. All traffic will be forwarded to the inline tool group.

- When an inline tool group is included as a member of an inline series, asymmetrical hashing is not supported.
- If an inline network is involved in a flex-inline map and the inline tool used in the map is a part of an inline tool group; the traffic path parameter for the inline network or the flex traffic path of the inline tool from the inline tool group is set to **ByPass with Monitoring**, there will not be any bypassed traffic irrespective of the hashing. All traffic will be forwarded to the inline tool used from the inline tool group.

Resilient Weighted Hashing

One of the parameters of inline tool groups is weighting that provides you the ability to distribute traffic to the inline tools by assigning either an equal weighting or a custom weighting to the inline tools. You can assign custom weight in percentage or ratio. If an inline tool in a group goes down and the group maintains the **Minimum Healthy Group Size** that is defined for the group, the traffic is redistributed to the remaining tools based on the equal weighting or the custom weighting assigned to the tools. If the inline tool group does not meet the **Minimum Healthy Group Size** defined for the group, the traffic is redistributed based on the **Failover Action** defined for the group.

NOTE: Resilient hashing is not supported for classic inline maps.

The values for the weighting parameter are as follows:

- **Equal**—Traffic is distributed equally to all the inline tools in the inline tool group.
- **Relative**—Traffic is distributed to the inline tools in the inline tool group based on the relative weight or ratio assigned to the respective inline tools. The valid range is 1–256.
- **Percentage**—Traffic is distributed to the inline tools in the inline tool group based on the percentage assigned to the respective inline tools. The valid range is 1–100.

If you select **Relative** or **Percentage** as the weighting option, enter the hash weights for the inline tools that appear in the table below the **Weighting** drop-down list. Ensure that you assign a hash weight for each inline tool in the inline tool group.

Inline Serial Tools

Use the **Inline Serial Tools Groups** configuration page to configure inline tools in a series, in which the traffic from one side of the inline network is guided through the members of the inline tool series before it is sent out the other side of the inline network. With inline tools and inline tools groups arranged in a series, the traffic from one inline tool or inline tool group flows to the next, so all tools see the same traffic.

The inline tool ports that make up the inline tools and inline tool groups participating in the inline tool series are always in pairs, running at the same speed, on the same medium (fiber or copper). All inline tool ports of the inline tool series must be on the same GigaVUE-HC3, or GigaVUE-HC1 and GigaVUE-HC1-Plus node. The inline tool ports and inline tool groups must also be on the same GigaVUE-HC3, or GigaVUE-HC1 and GigaVUE-HC1-Plus node as the inline network ports.

To view the currently configured inline tool series, select **Inline Bypass > Inline Serial Tools** to open the Inline Serial Tool Groups page shown in the following figure.

Refer to [Figure 34 Inline Serial Tools](#) for an illustration of an inline tool series. In [Figure 34 Inline Serial Tools](#), traffic is only shown from A-to-B.




Figure 34 *Inline Serial Tools*

Refer to [Figure 35 Inline Tool Series, Including Inline Tool Group](#) for an inline tool series that includes an inline tool group in addition to individual inline tools. The per-direction-order is set to forwarding. In [Figure 35 Inline Tool Series, Including Inline Tool Group](#), the inline tool group is placed as the middle member of the series, but it could be placed as the first member of the series or the last member of the series as well. In the inline tool group, traffic is shared.




Figure 35 *Inline Tool Series, Including Inline Tool Group*

NOTE: For inline SSL decryption, the inline tool series does not support an inline tool group in the series.

The number of inline tools and inline tool groups in the inline series is limited only by the number of ports available for creating the inline networks and inline tools participating in the inline bypass solution.

When an inline tool group is included as a member of a inline series:

- a spare inline tool can be configured on the inline tool group
- inline maps to individual members of an inline tool group are not supported
- asymmetrical hashing is not supported, which means that the hash options, ascrip-b-dstip and b-scrip-a-dstip, are not allowed on the inline tool group.

To configure an inline serial tool group, do the following:

1. Open the **Inline Serial Tool Group** configuration page by selecting **Inline Bypass > Inline Serial Tools** from the main navigation pane, and then clicking **New**.
2. Enter a name for the inline serial tool group in the **Alias** field to identify the group and an optional description in the **Description** field.
3. Select and order the inline tools for the inline tool group.
 - a. Click in the **Inline Tools** fields. The drop-down list shows the available inline tools
 - b. Select the inline tools or in or inline tool group to add to the inline serial tool group.

The inline tools are displayed in the order that they are selected. To change the order, click the up and down arrows.
4. Click **Save**.

NOTE: An inline serial tool group cannot be edited after it is saved.

Inline Serial Tools Global Failover Action

One of the parameters of inline tool series is the failover action taken in response to a failure of the inline tool series as a whole. This is referred to as the global failover action.

Each inline tool or inline tool group in the series can also have its own failover action. This is referred to as the local failover action. Refer to [Inline Tool Series Local Failover Action](#) for details.

For global failover actions, an inline tool series is declared to be in a failure condition as soon as any of its member inline tools goes into a failure condition. An inline tool series recovers from a failure condition after all the member inline tools recover from their failure conditions. The failover action attributes of the individual inline tools participating in an inline tool series are ignored. Instead, the **Failover action** configured on the Inline Serial Tool Group page for the inline tool series is executed. The values for global failover actions are as follows:

- **ToolBypass**—when the inline serial tools fails, the traffic that normally was directed to the inline tool is redirected to the bypass path. Use this failover action for configurations involving an inline serial tools that is associated with an inline network using rule-based maps. For configurations using map passalls, tool-bypass is the same as network-bypass. Refer to [Figure 36Inline Tool Series Global Bypass Failover Action](#).
- **NetworkBypass**—when the inline serial tools fails, all traffic that would not have been dropped when the inline network or networks had a NORMAL forwarding state is directed to the bypass path. That is, all such traffic arriving at the side A inline network port or ports is forwarded to the side B inline network port or ports and all traffic arriving at the side B inline network port or ports is forwarded to the side A inline network port or ports. Refer to [Figure 36Inline Tool Series Global Bypass Failover Action](#).
- **NetworkPortForcedDown**—when the inline serial tools fails, the inline network ports of the respective inline network are forced down. Refer to [Figure 37Inline Tool Series Global Network Ports Forced Down Failover Action](#).
- **ToolDrop**—when the inline serial tools fails, the traffic that normally was directed to the inline tool is dropped. Use this failover action for configurations involving an inline serial tools that is associated with an inline network using rule-based maps. For configurations using map passalls, tool-drop is the same as network-drop. Refer to [Figure 38Inline Tool Series Global Drop Failover Action](#).
- **NetworkDrop**—when the inline serial tools fails, all traffic coming to the respective inline network (or inline network group) is dropped. Refer to [Figure 38Inline Tool Series Global Drop Failover Action](#).

The bypass path is between side A and side B of the inline network ports.

The default is **ToolBypass**.

Any failure of any member leads to the failure of the inline serial tools, hence the failover action for an inline serial tools will overwrite the failover action of the inline tool members of the series.

Figure 36 Inline Tool Series Global Bypass Failover Action to Figure 38 Inline Tool Series Global Drop Failover Action show the global failover actions for the inline series when any individual inline tool in the series fails.




Figure 36 Inline Tool Series Global Bypass Failover Action




Figure 37 Inline Tool Series Global Network Ports Forced Down Failover Action




Figure 38 Inline Tool Series Global Drop Failover Action

Inline Tool Series Local Failover Action

Each inline tool in the series can have its own local failover action. When an individual inline tool or inline tool group in the series fails, the action taken depends on the failover action of the individual inline tool.

To configure local failover actions, configure a failover action of **Per Tool** for the series as a whole. Then the individual failover action for each inline tool in the series, as configured with the **inline Tool failover action**, takes effect. For details on the values, refer to [Inline Tool Failover Action](#).

For example, if the failover action of the inline series is configured as **Per Tool** and the failover action of an individual inline tool in the series is configured as **ToolBypass**, when that tool in the series fails, the traffic will skip over the failed tool.

NOTE: For inline SSL decryption, the **Per Tool** failover action is not supported.

The values for local failover actions are as follows:

- **ToolBypass**—when the inline tool fails, the traffic bypasses the failed tool. That is, the traffic originally coming to the inline tool is diverted to the next inline tool in the series or to the appropriate inline network port if the inline tool is the last in the series. Refer to [Figure 39Inline Tool Series Local Tool Bypass Failover Action](#).

NOTE: When all the inline tools in a series are configured as **ToolBypass** and they all fail, this is the same as the failover action of **ToolBypass** for the series.

- **ToolDrop**—when the inline tool fails, traffic to this inline tool stops being forwarded. Effectively, this has the same result as the failover action of **ToolDrop** for the series as a whole, although the healthy members of the series will still receive traffic in one of the directions. Refer to [Figure 40Inline Tool Series Local Tool Drop Failover Action](#).
- **NetworkBypass**—when the inline tool fails, a bypass is established between the inline network ports. Refer to [Figure 41Inline Tool Series Local Network Bypass Failover Action](#).
- **NetworkDrop**—when the inline tool fails, traffic is dropped at the inline network ports. Refer to [Figure 42Inline Tool Series Local Network Drop Failover Action](#).
- **NetworkPortForcedDown**—when the inline tool fails, the links for the inline network ports are brought down. Refer to [Figure 43Inline Tool Series Local Network Ports Forced Down Failover Action](#).

[Figure 39Inline Tool Series Local Tool Bypass Failover Action](#) to [Figure 43Inline Tool Series Local Network Ports Forced Down Failover Action](#) show the local failover actions when an individual inline tool in a series fails




Figure 39 *Inline Tool Series Local Tool Bypass Failover Action*




Figure 40 *Inline Tool Series Local Tool Drop Failover Action*




Figure 41 *Inline Tool Series Local Network Bypass Failover Action*




Figure 42 *Inline Tool Series Local Network Drop Failover Action*




Figure 43 *Inline Tool Series Local Network Ports Forced Down Failover Action*

Figure 44 Inline Tool Series Local Failover of Two Tools shows the failure of two individual inline tools in a series with different configured failover actions.




Figure 44 *Inline Tool Series Local Failover of Two Tools*

Inline Tool Series Per-Direction Order

One of the parameters of inline tool series is the per-direction order of the inline tool series. This parameter configures the traffic direction order of side B traffic with respect to the inline tool list, that is, the direction of the return traffic.

The **Return Direction** options on the Inline Serial Tool Group page specify the per-direction order of the side B traffic of the inline tool series as follows:

- **Reverse** specifies that the traffic from network B will flow through the inline tool list in reverse order, for example, from the third tool, to the second tool, to the first tool. This specifies the reverse order of inline tools for both directions.
- **Forward** specifies that the traffic from network B will flow through the inline tool list in the order it which it is defined, for example, from the first tool, to the second tool, to the third tool. This specifies the same order of inline tools for both directions of traffic.

The default is **Reverse**.

Figure 45 Inline Tool Series Per-Direction Order: Simplified Reverse and Forward shows a simplified view of the flow through the tools:




Figure 45 Inline Tool Series Per-Direction Order: Simplified Reverse and Forward

In Figure 45 Inline Tool Series Per-Direction Order: Simplified Reverse and Forward, **Reverse** is on the left and **Forward** is on the right. Traffic from network side A to network side B for both reverse and forward flows from the first tool, to the second tool, to the third tool. But traffic from network side B to network side A with reverse, flows from the third tool, to the second tool, to the first tool, whereas traffic from network side B to network side A with forward, flows from the first tool, to the second tool, to the third tool.

Figure 46 Inline Tool Series Per-Direction: Reversed shows the reverse direction with the tools connected to the GigaVUE node.




Figure 46 Inline Tool Series Per-Direction: Reversed

Figure 46 Inline Tool Series Per-Direction: Reversed shows the forward direction with the tools connected to the GigaVUE node.




Figure 47 Inline Tool Series Per-Direction: Forward

Associate Inline Networks with Inline Tools Using Inline Maps

Inline networks and inline network groups are associated with inline tools, inline tool groups, and inline serial tools through inline maps. An inline map is a regular map, but the **Source** and **Destination** fields specify inline software constructs instead of port lists.

On the **Edit Map** or **New Map** page, the **Source** field specifies the inline network alias or inline network group alias. The **Destination** field specifies an inline tool alias, an inline tool group alias, an inline tool series alias, or an inline bypass. An inline bypass is a special construct that is used as a pseudo-inline tool to allow a portion of traffic to bypass any inline tools.

Maps can associate a network with multiple inline tools or they can associate multiple inline networks with the same inline tool or with multiple inline tools.

With inline maps, only the traffic that meets the map rules is sent to the tools, to bypass, or to shared collectors. For example, you can send traffic to tools for which they are specialized and send the rest to bypass. Or, if there is a type of traffic in which the tools are not interested or do not understand, that traffic can be sent to a shared collector.

NOTES:

- When an inline network is mapped to an inline tool or inline tool group, a second inline network cannot be mapped to the same inline tool. (In other words, an inline tool can be used in only one map.) However, when there are multiple inline networks, use an inline network group to map to the same inline tool.
- If an inline tool is already specified in a map, that tool cannot be included in an inline tool group (unless the map is first deleted).
- Inline network ports, inline tool ports, and out-of-band tool ports that are used in map configuration must all be configured on the same GigaVUE-HC3, or GigaVUE-HC1 or GigaVUE-HC1-Plus node. Even if nodes are in a cluster, the inline ports cannot be on different nodes.

Refer to the following sections:

- [Inline Map Passall](#)
- [Inline Map](#)
- [Inline Map Shared Collector](#)
- [Inline Maps to Individual Members of an Inline Tool Group](#)
- [Out-of-Band \(OOB\) Map](#)
- [Symmetric and Asymmetric Maps](#)

Inline Map Passall

Use the **Pass All** subtype to configure a type of inline map that passes all traffic. Map-passalls facilitate the sending of traffic between inline network ports (through inline tools or bypass). All of the inline network ports and inline tool ports involved in an inline map must be located on the same Gigamon node.

Use the **Source** of the **Pass All** to specify a single inline network.

Use the **Destination** of the **Pass All** to specify an inline tool alias, an inline tool group alias, an inline tool series alias, or an inline bypass. An inline bypass is a special construct that is used as a pseudo inline tool to allow a portion of traffic to bypass any inline tools.

An inline bypass, using the physical bypass option, is only valid if the **Source** is an inline network port. This applies only to the **Pass All** subtype for asymmetrical scenarios.

Inline Map

Use the **New Map** page to configure a type of inline map that uses rules to direct traffic. These inline maps are referred to as rule-based maps. All of the inline network ports and inline tool ports involved in an inline map must be located on the same GigaVUE node.

Use the **Source** field to specify an inline network alias or an inline network group alias.

Support for rule-based maps is limited to symmetric scenarios, which means that the **Destination** of rule-based inline maps can only be aliases of inline networks or inline network groups (not individual ports).

Use the **Destination** to specify an inline tool alias, an inline tool group alias, an inline tool series alias, or an inline bypass. An inline bypass is a special construct that is used as a pseudo inline tool to allow a portion of traffic to bypass any inline tools.

For rule-based maps, the **Destination** can be configured to inline bypass with no restrictions, so long as the **Source** specifies either an inline network or an inline tool.

Use the **Priority** to order inline maps by priority. For example, you can specify the highest priority map to be for encrypted traffic and lowest priority map to be for a shared collector. You can also place inline maps before or after one another using the **Priority**.

Use the **Type** to specify any map rule.

Inline Map Shared Collector

Use the **Collector** subtype to configure a shared collector to which to send any packets that do not match the map rules in the inline maps. Use a shared collector with one or more rule-based inline maps. All of the inline network ports and inline tool ports involved in an inline shared collector map must be located on the same GigaVUE node.

Use the Source argument to specify an inline network alias or an inline network group alias.

Use the Destination argument to specify an inline tool alias, an inline tool group alias, an inline tool series alias, or an inline bypass. An inline bypass is a special construct that is used as a pseudo inline tool to allow a portion of traffic to bypass any inline tools.

For shared collector maps, the **Source** field can be configured to inline bypass with no restrictions, so long as the **Destination** field specifies either an inline network or an inline tool.

Support for shared collector inline maps is limited to symmetric scenarios, which means that the **Destination** field of rule-based inline maps can only be aliases of inline networks or inline network groups (not individual ports).

Inline Maps to Individual Members of an Inline Tool Group

Prior to software version 4.4, the **Map** page was used to configure an inline map that directed traffic to the inline tool group as a whole. The map could be rule-based or passall. There was only one type of hashing available, which distributed the traffic across the tools in the inline tool group.

Starting in software version 4.4, the **Map** page can also be used to configure inline maps to the individual members of an inline tool group. The maps must be rule-based. There are also more hashing options that can be specified for traffic that does not match any of the map rules. The hashing options are described in [Symmetrical and Asymmetrical Hashing](#).

The rule-based maps are defined with the inline tool group sharing the same source, either an inline network or an inline network group in the **Source** field of the map. The map destinations (the **Destination** field of the map) are the individual inline tools in the group. Traffic not matching any of the map rules is sent to a shared collector to be distributed according to the specified hashing value.

The shared collector must also have the same source as the maps to the individual members of the inline tool group. The destination for the shared collector is the inline tool group. The shared collector map is a mandatory part of the configuration.

Both configurations are available: either a single map to the inline tool group as a whole, or multiple rule-based maps to the individual members of the inline tool group plus a shared collector; however, they cannot both be configured at the same time.

Refer to [Figure 48Rule-Based Maps to Individual Tools in an Inline Tool Group](#) for the rule-based maps to the individual members of the inline tool group. In [Figure 48Rule-Based Maps to Individual Tools in an Inline Tool Group](#), traffic is only shown from A-to-B.




Figure 48 Rule-Based Maps to Individual Tools in an Inline Tool Group

Map, Inline Tool, and Inline Tool Group Configuration Restrictions

The following are map, inline tool, and inline tool group configuration restrictions for the rule-based maps to the individual members of the inline tool group:

- If there is a map to the inline tool group as a whole, there cannot also be rule-based maps to the individual inline tools in the group.
- Maps to the individual inline tools in a group must be rule-based. Map passalls to the individual tools cannot be configured.
- The source of rule-based maps to the individual inline tools in a group must be the same (either the same inline network or the same inline network group). The shared collector must have the same source as well.
- If there is a map configured to the individual inline tools in a group, the inline networks must have their traffic path set to **to-inline-tool**. This applies to individual inline networks as well as to inline networks involved in an inline network group.
- For the individual inline tools in a group, the recovery mode of the individual inline tools must be configured as **automatic**. A recovery mode of **manual** cannot be configured.
- For the inline tool group, the failover action must be either **NetworkBypass** or **NetworkDrop**. A failover action of either **ToolBypass** or **ToolDrop** for the inline tool group cannot be configured.
- Only one inline shared collector map can be configured (among the set of inline maps).
- Maps must be created in a specific order. The shared collector map must be configured last. For example, if there are three inline tools in the group, configure the three maps to the individual members of the group first, then configure the shared collector map.

- Maps must be deleted in a specific order. The shared collector map must be deleted first. Then the maps to the individual members of the group can be deleted.
- Once the shared collector map is configured, any changes to the maps to the individual members of the group are restricted. Only the map rules can be edited.

NOTE: All the rules in a map cannot be deleted. All maps must have attributes for **Source**, **Destination**, and at least one **rule** configured.

- When an inline tool group is included as a member of an inline series, inline maps to individual members of an inline tool group are not supported.
- If one of the inline tools is disabled in an inline-network-group to inline-tool-group map that has the same ingress VLAN tag on the ports of an inline-network pair then the traffic will be looped back to the same network. This behavior is seen only with passall maps. In above scenarios use rule-based maps instead of passall maps.

Inline Tool Failures and Failover Actions

An inline tool group has a failover action for the group as a whole. The failover action is taken in response to a failure when the number of healthy inline tools in the inline tool group (including the spare inline tool, if configured) falls below the configured minimum. In addition, the individual inline tools in the group have failover actions.

When there are maps to individual inline tool members of the group, an inline tool has both a group failover action and an individual failover action.

Refer to [Failover Actions When an Individual Tool in an Inline Tool Group Fails](#) for the failover actions when an individual inline tool in an inline tool group fails.

Table 18: Failover Actions When an Individual Tool in an Inline Tool Group Fails

Inline Tool Group Spare	Number of Healthy Tools	When Individual Inline Tool Fails:	
		Traffic to Inline Tool Group	Traffic to Failed Tool
no spare	equal to or greater than the configured minimum healthy size	is redistributed among the remaining healthy tools in the group	no action—the failover action of the individual inline tool is ignored. If there is an inline map to the failed inline tool, the connectivity arrangements supporting the individual map are maintained as if nothing happened.
	less than the configured minimum healthy size	fails over according to the failover action for the inline tool group	no action—the failover action of the individual inline tool is ignored. If there is an inline map to the failed inline tool, the connectivity arrangements supporting the individual map are maintained as if nothing happened.

Inline Tool Group Spare	Number of Healthy Tools	When Individual Inline Tool Fails:	
		Traffic to Inline Tool Group	Traffic to Failed Tool
spare	equal to or greater than the configured minimum healthy size	is redistributed among the remaining healthy tools	no action—the failover action of the individual inline tool is ignored. If there is an inline map to the failed inline tool, the connectivity arrangements supporting the individual map are maintained as if nothing happened.
	less than the configured minimum healthy size	fails over according to the failover action for the inline tool group	no action—the failover action of the individual inline tool is ignored. If there is an inline map to the failed inline tool, the connectivity arrangements supporting the individual map are maintained as if nothing happened.
failed spare	N/A	no action, if there is no inline map configured to the spare	no action—the failover action of the individual inline tool (in this case, the spare) is ignored. If there is an inline map to the failed inline tool, the connectivity arrangements supporting the individual map are maintained as if nothing happened.

Maps That May Lead to Selective Traffic Drops

With inline bypass solutions based on inline flow mapping, the use of rule-based maps can lead to selective traffic drops. Traffic drops can occur as follows:

- if a shared collector from the inline network or inline network group has not been configured. Packets not matching the criteria specified by the rules in the configured rule-based maps will be dropped.
- if drop rules have been included in the rule-based maps configured from the inline network or inline network group

In most inline flow mapping solutions, all traffic exchanged between the two end nodes of a given inline network are expected to be processed by the inline tool or tools associated with this inline network through the configured maps. Therefore, it is recommended to always configure a shared collector and to not include drop rules in the rule-based maps.

Out-of-Band (OOB) Map

All inline network ports and inline tool ports can be subject to monitoring by listen-only tools. This means that an inline network port or inline tool port can be listed in the **Source** field in which the **Destination** field is an arbitrary tool type of port located anywhere in the system and not limited to the same node.

Inline network ports and inline tool ports involved in rule-based inline maps can be used as network ports for monitoring (or out-of-band) maps.

Out-of-band (OOB) maps are supported as follows:

- If the inline bypass solution use passalls, the OOB arrangements can use any rule-based maps or map shared collectors, or passalls. Refer to [Figure 49 Out-of-Band Rule-Based Maps](#).
- If the inline bypass solution use rule-based maps or map shared collectors, the OOB arrangements can use only map passalls. Refer to [Figure 50 Out-of-Band Map Passalls](#).
- When the source port of an OOB map is associated with an inline network, a list of inline ports is supported in the port list (the **Source** field of the Map page).

The following restrictions apply to OOB maps:

- When the source port of an OOB map is associated with an inline network group, only a single inline port (network or tool) is supported in the port list. In this case, multiple OOB maps are needed because each OOB map only accepts one inline port (network or tool) in the **Source** field on the Map page.
- OOB maps from inline network ports of inline networks involved in maps to inline tool groups configured with asymmetrical hashing are not allowed. If an inline network is involved in an inline map to an inline tool group configured with asymmetrical hashing, the inline network ports of the inline network cannot be used as the **Source** attribute in any out-of-band maps.

Prior to software version 4.4, if an inline network was part of an inline network group, sending traffic to an out-of-band tool was not allowed.

Starting in software version 4.4, out-of-band maps from inline ports involved in inline network groups are supported. You can configure OOB maps originating from inline network ports or inline tool ports when these ports are involved in an inline network group, except for the following:

- GigaSMART operations
- tool ports located on a different node





Figure 49 Out-of-Band Rule-Based Maps**Figure 50** Out-of-Band Map Passalls

Symmetric and Asymmetric Maps

In a symmetric map configuration, the southbound and northbound forwarding to the tools is the same. For example, traffic from A to B goes through an inline tool, as does traffic from B to A. Rule-based maps are limited to symmetric configurations.

In an asymmetric map configuration, the southbound traffic is distributed to the inline tools, but northbound traffic can be sent through uninspected. For example, traffic from A to B goes through an inline tool, but traffic from B to A bypasses it.

Asymmetric configurations are only supported with map passalls. Traffic can come from individual inline networks, be sent to individual inline tools or inline tool groups, or to bypass.

Symmetrical combinations of two asymmetrical arrangements (that is, with both side A and side B pointing to the inline tool or to bypass) are not allowed.

Some IPSs do not need to inspect northbound traffic. For example, those focused on preventing Denial of Service (DoS) attacks, may not need to keep track of session flows and may only be concerned about southbound traffic.

Conversely, data loss prevention systems, those that are more concerned about what sensitive data is leaving the protected network than what is coming in, may focus solely on northbound traffic.


Configuration Steps

The configuration steps in summary for an inline bypass solution are as follows:

1. Configure inline network ports. (Optional for protected inline network.)

2. Configure inline network. (Optional for protected inline network.)
3. (Optional) Configure inline network group.
4. (Optional) Configure heartbeat or negative heartbeat profile.
5. Configure inline tool ports.
6. Configure inline tool.
7. (Optional) Configure inline tool group.
8. (Optional) Configure inline tool series.
9. Configure inline maps, either passall, map (rule-based), map shared collector, or bypass.
10. Configure non-default values for parameters of the inline networks or inline tools.

The summary steps are shown in [Figure 51 Configuration Steps for Inline Bypass Solutions](#).

**Configuration Steps**

Inline Maps to Individual Members of an Inline Tool Group

Figure 51 Configuration Steps for Inline Bypass Solutions

Configuration Step Details

This section provides detailed steps for configuring inline bypass.

Configure Inline Network Ports

Use the following procedure to create inline network ports:

1. Go to **System > Ports > Ports > All Ports**.
2. Open Quick Port Editor by clicking the **Port Editor** button.
3. Use the Quick search field to Find the ports to configure.
4. In the Alias field enter a name to help identify the inline port.
5. For Type, select **Inline Network**.

An Inline Network (unprotected) is a software arrangement of two network-type ports allocated to facilitate access to a bidirectional link between two networks (far end network devices) that are linked to inline tool ports.

6. Click **OK**.

NOTE: Any available network-type ports on a Gigamon node can be used to form an unprotected inline network.

Configure Inline Network (Unprotected)

An Inline Network Group is an arrangement of multiple inline network ports to which traffic is distributed based on calculated hash values used by a Gigamon node.

Perform the steps to configure an inline network:

1. After configuring the Inline Network ports, select **Inline Bypass > Inline Networks**.
2. Click **New**.
3. In the **Alias** field, enter an alias for the inline network to help identify the inline network.
4. From the **Port A** drop-down list select an inline network port.
5. An inline network port is automatically selected for **Port B**. To select a different port, select one from the **Port B** drop-down list if there is more than one inline network port.
6. Select a traffic path from the **Traffic Path** drop-down list. The types of traffic paths are:

- **Bypass** — All traffic arriving at the Port A inline network port is directly forwarded to the Port B inline network port and all traffic arriving at the Port B inline network port is directly forwarded to the Port A inline network port.
- **Drop** — No traffic is exchanged through the inline network ports (all traffic coming to these ports is dropped).
- **ByPass with Monitoring** — All traffic is forwarded as a forced bypass value and a copy of the traffic is also forwarded to the inline tools. A traffic map must first be configured between the inline network and the inline tool to have the traffic forwarded with no traffic taken from the inline tools.
- **To Inline Tool** — The traffic received at the inline network ports is forwarded based on:

1. The traffic map between the inline network and the respective inline tools.
 2. The failover action attributes of the inline tools
 3. The health state of the inline tools.
7. Select the **Link Propagation** check box to enable whether the inline network link on one side of the inline network gets propagated to the other side.
8. Click **Save**.

Configure Inline Network Group

An Inline Network Group is an arrangement of multiple inline network ports to which traffic is distributed based on calculated hash values used by Gigamon node.

Perform the following steps to configure an Inline Network Group:

1. After configuring the Inline Network Ports, select **Inline Bypass > Inline Network Groups**.
2. Click **New**.
3. In the **Alias** field, enter a name for the network group.
4. From the **Inline Network** drop-down list, select the Inline Network ports.
5. Click **Save**.

Create Heartbeat Profile

The Heartbeat Profile is a data structure that contains the heartbeat attributes that are applied to an Inline Tool for configuring its heartbeat. The Create Heartbeat Profile wizard allows you to apply attribute values for the heartbeat profile. Use the following procedure:

1. Select **Inline Bypass > Heartbeats**.
2. Click **New**.
3. In the **Alias** field, enter a name for your heartbeat profile.

4. In the **Type** field, select **Regular**.

For details about regular heartbeats, refer to [Standard Heartbeat](#).

5. Use **Packet Format** drop-down list to select a packet type. The formats are:
 - **ARP**—This protocol (Address Resolution Protocol) is used for resolution of network layer address into link layer addresses, which is critical for multiple-access network operation. ARP is the default.
 - **Custom**—This format is a binary packet content associated with a packet capture (pcap) file. For details about custom packet format, refer to [Standard or Custom Heartbeat Packet](#).

When you select **Custom**, a **Custom Format** field displays with a **Browse** button. Use the **Browse** button to upload the pcap file.

6. Use the **Direction** drop-down menu to select the direction for sending heartbeat. The directions are:
 - **A to B**—From Port A to Port B of the inline tool.
 - **B to A**—From Port B to Port A of the inline tool.
 - **Bi-directional**—Both directions.
7. In the **Timeout** field, enter a number in milliseconds to indicate a timeout period for heartbeat packets between sending and receiving. The acceptable range is 20 to 1000 milliseconds. The default is 500 milliseconds.
8. In the **Period** field, enter a number in milliseconds for sending subsequent heartbeat packets. The acceptable range is 30 to 5000 milliseconds. The default is 1000 milliseconds.
9. In the **Recovery Time** field, enter a number in seconds to indicate that the inline tool is declared up with successfully receiving packets. The acceptable range is 5 to 60 seconds. The default is 30 seconds.
10. In the **Retries** field, enter the number for consecutive timed-out heartbeat packets at which the system will trigger (retry) a fail over condition.
11. Click **Save**.

The heartbeat profile appears in the Heartbeat Profile table.

NOTE: Highlight the heartbeat profile and click **Edit** to modify the parameters, if needed.

Create Negative Heartbeat Profile

The Negative Heartbeat Profile is a data structure that contains the negative heartbeat attributes that are applied to an Inline Tool for configuring its negative heartbeat. The Create Heartbeat Profile wizard allows you to apply attribute values for the negative heartbeat profile. Use the following procedure:

1. Select **Inline Bypass > Heartbeats**.
2. Click **New**.
3. In the **Alias** field, enter a name for your heartbeat profile.
4. In the **Type** field, select **Negative**
 This is a negative heartbeat profile. For details about negative heartbeats, refer to [Negative Heartbeat Profiles](#).
5. Use **Browse** button in the **Custom Format** field to upload binary packet content associated with a packet capture (pcap) file. For details about the custom format, refer to [Standard or Custom Heartbeat Packet](#).
6. Use the **Direction** drop-down menu to select the direction for sending heartbeat. The directions are:
 - **A to B**—From Port A to Port B of the inline tool.
 - **B to A**—From Port B to Port A of the inline tool.
 - **Bi-directional**—Both directions.
7. In the **Period** field, enter a number in milliseconds for sending subsequent negative heartbeat packets. The acceptable range is 30 to 5000 milliseconds. The default is 1000 milliseconds.
8. In the **Recovery Time** field, enter the minimum number of seconds since the last negative heartbeat packet is received to declare that the inline tool is up.
 The inline tool is up from the standpoint of the negative heartbeat if the negative heartbeats sent are *not* received. When a tool is declared down, sent heartbeats should not be received for a number of seconds in order to declare the tool as being up.
 The acceptable range for the Recovery Time field is 5 to 60 seconds. The default is 30 seconds.
9. Click **Save**.

The heartbeat profile appears in the Heartbeat Profile table.

NOTE: Highlight the heartbeat profile and click **Edit** to modify the parameters, if needed.

Configure Inline Tool Ports

Use the following procedure to create inline tool ports:

1. Go to **System > Ports > Ports > All Ports**.
2. Open Quick Port Editor by clicking the **Port Editor** button.
3. Use the Quick search field to Find the ports to configure.
4. In th an Alias field enter a name to help identify the inline port.

5. For **Type**, select **Inline Tool**.

An Inline Tool represents a pair of inline tool ports.

6. Click **OK**.

Configure Inline Tool

An Inline Tool represents a pair of inline tool ports. To configure an Inline Tool, do the following:

1. Select **Inline Bypass > Inline Tools**.
2. Click **New**.
3. If needed, click **Port Editor** to open the Quick Port Editor to configure the inline tool ports.
4. Select the inline tool ports for the inline tool.
 - o For **Port A**, select an inline tool port configured in the previous procedure, [Configure Inline Tool Ports](#).
 - o For **Port B**, select another inline tool port configured in the previous procedure, [Configure Inline Tool Ports](#).
5. Select **Enabled** to set the inline tool ports as enabled for inline bypass traffic.
6. For **Failover action**, select one of the following:
 - o **Tool Bypass** — When the inline tool fails all traffic coming to the respective inline tool is directed via the bypass path.
 - o **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.
 - o **Tool Drop** — When the inline tool fails all traffic coming to the respective inline tool is dropped.
 - o **Network Drop**—When the inline tool fails all traffic coming to the respective inline tool is dropped.
 - o **Network Port Forced Down**—When the inline tool fails the inline network ports of the respective inline network are forced as "down".
7. Select the **Recovery Mode**. The recovery mode can be one of the following:
 - o **automatic**—Specifies automatic recovery, which redirects traffic back to the inline tool as soon as it has recovered from all faulty conditions.
 - o **manual**—Specifies manual recovery, which lets you control when to put an inline tool back into service after the tool has recovered. For example, you may want to wait for a maintenance window to return the inline tool to service.
8. Select **Enable Heartbeat** to set the heartbeat.

If the heartbeat is enabled, do the following:

- a. From **Profile**, select the desired heartbeat profile. The available heartbeat profiles are existing profiles from other inline tools. Once a heartbeat profile is selected, its attributes are displayed.
 - b. In the **HB IP Address A** field, enter the server's IP address to send the heartbeat packets.
 - c. In the **HB IP Address B** field, enter the server's IP address to send a second heartbeat packet.
9. Click **Save**.

Create Inline Tool Group

An Inline Tool Group is an arrangement of multiple Inline Tools to which traffic is distributed based on calculated hash values used by Gigamon node.

To configure an Inline Tool Group, do the following:

1. Use the **Quick Port Editor** to configure the inline tool ports.
2. Select **Inline Bypass > Inline Tool Groups**.
3. Click **New**.
4. In the **Alias** field, enter a name to help identify the inline tool group.
5. For **Inline Tools**, select the inline tool ports for the inline tool group.
6. (Optional) For the **Inline Spare Tool**, select another inline tool port.
If a spare is selected, the inline tool group becomes a redundant arrangement of inline tools. When the first failure occurs in a set of active inline tools, traffic will be forwarded to the spare with no loss, thus the spare will replace the failed tool in the active set.
7. Select **Enable** to enable the inline tool group.
8. Select the **Failover Action**. The failover actions are:
 - o **Tool Bypass**—When the inline tool group fails all traffic coming to the respective inline network is directed via the bypass path.
 - o **Tool Drop**—When the inline tool group fails all traffic coming to the respective inline network is dropped.
 - o **Network Bypass**—When the inline tool group fails all traffic coming to the respective inline network is directed to the inline tool group via the bypass path.
 - o **Network Drop**—When the inline tool group fails all traffic coming to the respective inline network group is dropped.
 - o **Network Port Forced Down**—When the inline tool group fails the inline network ports of the respective inline network are forced as "down".
9. For **Minimum Healthy Group Size**, select a number that represents the minimal amount of inline tools that are required to have a state of Normal.

10. Select the **Hash** for the inline tool group.

Hashing, which is used for distributing traffic across the inline tools in an inline tool group. The values for the hash parameter are as follows:

- **advanced**—Specifies symmetrical hashing, which is derived from the combination of packet fields based on the criteria selected for the advanced-hash algorithm. For inline bypass applications, the most common choice of criteria for the advanced-hash algorithm is the combination of source IP and destination IP addresses. This produces a hash value that sends all traffic associated with the same session to the same inline tool in the inline tool group.
- **a-srcip-b-dstip**—Specifies asymmetrical hashing, which is derived from the source IP address for side A of the network and the destination IP address for side B of the network. This produces a hash value that sends all traffic associated with the same source address residing on side A to the same inline tool in the inline tool group, regardless of destination or session.
- **b-srcip-a-dstip**—Specifies asymmetrical hashing, which is derived from the destination IP address for side A of the network and the source IP address for side B of the network. This produces a hash value that sends all traffic associated with the same source address residing on side B to the same inline tool in the inline tool group, regardless of destination or session.

11. Click **Save**.

Configure Inline Tool Series

To configure an Inline Tool Series, do the following:

1. Select **Inline Bypass > Inline Serial Tools**.
2. Click **New**.
3. In the **Alias** field, enter a name to help identify the inline tool series.
4. For **Inline Tools**, select the inline tool ports for the inline tool series.
5. Select **Enabled** to enable this configuration.
6. Select the **Failover action**.

For details about the failover actions, refer to [Inline Tool Series Local Failover Action](#)

7. Select the **Per Direction Order**.

For details about per-direction order, refer to [Inline Tool Series Per-Direction Order](#).

Configure When GigaVUE® HC Series Modules are Operationally Up

Ensure that the GigaVUE® HC Series modules are in the operationally *up* state before configuring them. Configuration changes done when a module is operationally *down* are not supported.

Also, when an inline tool or inline tool group is in the operationally *down* state, do not modify the current failover action of that inline tool or inline tool group until the tool has recovered from the failover state.

Avoid Over subscription

In general, traffic received at inline network ports is delivered to the destination ports according to the inline maps and the out-of-band maps regardless of whether the destination ports have the capacity to absorb all the traffic or not.

NOTE: When an inline network is involved in an inline map or an out-of-band map to a destination port (tool port or inline tool port), when there is temporary oversubscription, some packets arriving at the inline network port will be dropped. This can happen when the traffic path is set to bypass or monitoring.

Ensure that destination ports of maps originating from inline network ports have enough capacity to absorb the amount of traffic coming to the inline network ports.

Configure Gigamon Resiliency for Inline Protection

Gigamon Resiliency for Inline Protection (GRIP)™ is an Inline Bypass solution that connects two GigaVUE nodes together so that one node provides high availability to the other node when there is a loss of power. This redundant arrangement of two GigaVUE nodes maintains traffic monitoring by inline tools when one of the nodes is down.

GRIP makes use of the bypass protection switch relays for protected inline networks on GigaVUE-HC3, GigaVUE-HC1 nodes. The following modules are required to provide physical protection:

- bypass combo modules (BPS), for a protected pair of optical inline network ports on GigaVUE-HC3, GigaVUE-HC1.

- TAP-HC1-G10040 module, for a protected pair of copper inline network ports on GigaVUE-HC1

NOTE: GRIP is supported on GigaVUE-HC3 only when there are other modules installed in the node that can provide the stack link. The GRIP solution synchronizes the nodes through a signaling link using a stack link between two stack ports.

In the GRIP solution, two GigaVUE nodes are cabled so that traffic is guided through one GigaVUE node, acting in the primary role, while the other GigaVUE node is on standby, acting in a secondary role. If the primary node fails, the bypass protection switch relays on the modules switch the traffic over from the primary node to the secondary node. The failover or recovery process takes anywhere between 0-10 seconds depending on your configurations and devices. Contact Gigamon channel partner or Gigamon sales representatives for more information.

Using the physical protection for either copper or fiber, traffic is guided through inline tools by one of the GigaVUE nodes. The GigaVUE node with the open bypass protection switch relays is the one through which traffic flows. The traffic only flows through one GigaVUE node or the other.

To configure the GRIP solution for copper, use two TAP-HC1-G10040 modules on GigaVUE-HC1. The capacity will be 1Gb.

To configure the GRIP solution for fiber, use the following:

- two BPS-HC1-D25A24 or BPS-HC1-D25A60 or BPS-HC1-D35C60 modules on GigaVUE-HC1. The capacity will be 10Gb.
- two BPS-HC3-Q35C2G or BPS-HC3-C35C2G or BPS-HC3-C25F2G modules on GigaVUE-HC3. The capacity will be either 100Gb or 40Gb, depending on the configured port speed of the inline network port pairs

A stack port should be cabled between the two GigaVUE nodes. Also, two inline tools are needed for the GRIP solution. Refer to [Configure Gigamon Resiliency for Inline Protection](#), [Configure Gigamon Resiliency for Inline Protection](#), and [Configure Gigamon Resiliency for Inline Protection](#).

How to Cable GigaVUE Nodes

To cable two GigaVUE nodes, as shown in [Configure Gigamon Resiliency for Inline Protection](#) with the primary on the left and the secondary on the right:

- Connect the network shown at the top of [Configure Gigamon Resiliency for Inline Protection](#) to inline network port A on the primary node.
- Connect inline network port B on the primary node to inline network port A on the secondary node.

- Connect inline network port B on the secondary node to the network shown at the bottom of the [Configure Gigamon Resiliency for Inline Protection](#).
- Connect the signaling port on the primary node to the signaling port on the secondary node.

How to Handle Recovery

In the scenario in [Configure Gigamon Resiliency for Inline Protection](#), after traffic is flowing through the secondary node, at some point the primary node will come back up. The primary node will establish the configured inline traffic paths, bring the signaling link up, and open its relays. Traffic will then flow through the primary node again.

Both Nodes Go Down and Only Secondary Comes Up

In the GRIP solution, if both primary and secondary nodes are powered down or if there is a power outage causing both primary and secondary nodes to go down, powering up the secondary alone without the primary ever coming up will cause network traffic to be bypassed instead of being sent to inline tools.

It is not recommended to power up/recover only the secondary node without the primary. The recommendation is to eventually bring the primary up also.

If the primary node is prone to failures or frequent power outages, another recommendation is to change the role of the secondary node to the primary.

How to Cable GigaVUE Nodes

To cable two GigaVUE nodes, as shown in [Configure Gigamon Resiliency for Inline Protection](#) with the primary on the left and the secondary on the right:

- Connect the network shown at the top of [Configure Gigamon Resiliency for Inline Protection](#) to inline network port A on the primary node.
- Connect inline network port B on the primary node to inline network port A on the secondary node.
- Connect inline network port B on the secondary node to the network shown at the bottom of the [Configure Gigamon Resiliency for Inline Protection](#).
- Connect the signaling port on the primary node to the signaling port on the secondary node.

Configure GRIP Solution Software

To configure the GRIP solution in software:

1. On the left navigation pane, go to **Inventory**  > **Nodes**. Click on the Left Node in which the configuration needs to be done.
2. Go to **System > Ports > >Ports>All Ports**. Select the port that would act as the signaling port and click **Edit**.
3. Select **Enable for Admin**.
4. Select Type **Stack**.
5. Click **OK**.
6. Repeat Steps 1 through 5 on the right node to complete the signaling port type configuration.
7. Go to  > **Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** **Inline Flows**, and then click **Configuration Canvas** to create a new Flexible Inline Canvas.
8. In the Flexible Inline Canvas that appears, select the required device and click the '+' icon next to the **Redundancy** option.
9. Enter a name for the profile in the Alias field.
10. Click on the Signaling Port field and select the stack port configured in Step 4. Select **Primary** for Protection Role.
11. Click **OK**.
12. Repeat Steps 8 through 11 for the right node to complete the Redundancy Profile Configuration but select Protection Role as **Secondary**.

The redundancy profile specifies the following:

- **Signaling Port**—specifies the ports used to signal the state of the two GigaVUE nodes to each other. The ports provide the mechanism to detect loss of power in one of the GigaVUE nodes.
- **Protection Role**—specifies the role of the GigaVUE node, as primary, secondary, or suspended. The default is suspended. When suspended, the protection role is on hold. Changing a GigaVUE node from the primary role to the suspended role can be used to manually force the primary node down so the secondary node can become active. The suspended role is also used when performing maintenance. Refer to [Limitation for Suspended Role](#) and [How to Use Suspended Role for Maintenance](#).

The link between the signaling ports on the two GigaVUE HC Series nodes is for synchronization. When the node acting in the primary role is up, the signaling link is up, and the node acting in the secondary role sees the link as up. When the primary node loses power, the signaling link is brought down, and the secondary node sees the link as down and takes over.

The redundancy profile combines the protection role with the signaling port. The same redundancy profile is applied to the inline networks, so they have the same properties. If multiple inline networks on each GigaVUE node share the signaling link, they must be configured with the same protection role.

The primary and secondary roles on the two GigaVUE nodes do not change. That is, the role of the primary node stays the same and the role of the secondary node stays the same. The secondary always watches the state of the signaling port for whether the link is up or down.

For example, in [Configure Gigamon Resiliency for Inline Protection](#), after the primary node recovers, it will open its bypass protection switch relays. Through the signaling port, the primary node will indicate that it is ready to receive traffic by setting the link state to up. The secondary node will notice that the link is up and will close its bypass protection switch relays. After recovery, the primary node automatically goes back into service.

Limitation for Suspended Role

Though GRIP is supported in a cluster, there is a limitation when the suspended protection role is used on the standby node in the cluster. The recommendation is to either switch the standby to the leader or apply the suspended role in the redundancy profile to the leader.

Configure Synchronization

You must synchronize the configuration of the two GigaVUE nodes involved in the GRIP solution. The configuration items that must be synchronized are as follows:

- the signaling ports, as dictated by the signaling link cabling
- the inline networks, as dictated by the network path cabling between the two GigaVUE nodes
- the redundancy profiles. The redundancy profile of each GigaVUE nodes needs to have the same signaling port as well as a redundancy role that is compatible with the redundancy role on the other GigaVUE node. For example, one is configured with the primary role and one is configured with the secondary role.
- the inline tools
- the inline maps

For a configuration example of two GigaVUE-HC3 nodes, refer to [Example: Gigamon Resiliency for Inline Protection](#). In the example, the configuration is the same on both nodes, except for the protection role (primary or secondary).

Display Redundancy Control State

To display the Redundancy Control State, go to the Inline Networks page and click on the alias of the Inline Network for which you want to display the redundancy control state. The state is displayed on the Quick View under Configuration.

Table 19: Redundancy Control States

State	Description
Neutral	No redundancy profile is configured.
Suspended	The protection role is configured as suspended.
Primary Forwarding	The protection role is configured as primary. The node is acting in the primary role. Traffic flows through this node.
Secondary Bypass	The protection role is configured as secondary. The node is acting in the secondary role. Traffic bypasses this node.
Secondary Forwarding	The protection role is configured as secondary. The node is acting in the primary role due to a loss of power on the primary node. Traffic flows through this node.

How to Use Suspended Role for Maintenance

Use the suspended protection role to perform maintenance activities on the primary and secondary nodes. Maintenance activities include: bringing up a module, shutting down a module, or swapping a module.

For example, to remove a module on one of the GigaVUE-HC3 nodes (Primary node), use the following steps on that module:

1. Select **Physical > Orchestrated Flows > Inline Flows > Configuration Canvas > Redundancy**.
2. On the Redundancies page, for **Protection Role**, select **suspended**, and then click **OK**.
3. Once this is configured, the Primary node will be moved to 'Suspended' and the Secondary node will be moved to 'Secondary Forwarding' state. All the traffic will now be forwarded to the Secondary node and the Inline Tool inspection takes place.
4. Perform the maintenance activity in Primary node, like bringing up a new module, shutting down a module, swapping the modules, replacing the external inline tool.
5. Once the maintenance is done, revert the **Protection Role** in Primary Node back to **'primary'**. This will move the Redundancy Control State back to the **Primary Forwarding** and traffic will start flowing via the Primary Node.

In case of a maintenance activity (chassis, card, external inline-tool) required in Secondary node follow the below steps:

1. Set the Protection Role to '**suspended**' in the redundancy profile.
2. Once this is set, the Secondary node will be moved to '**Suspended**' and the Primary node will remain in the '**Primary Forwarding**' state and will handle the traffic.
3. Do the required maintenance activity in secondary node.
4. Once the maintenance is done, revert the Protection Role in Secondary Node back to '**secondary**'.
5. This will move the Redundancy Control State back to the "**Secondary Bypass**".

Upgrade Procedure Recommendations

There are no specific procedures for upgrading the nodes in the GRIP deployment, but we would recommend the below steps to be on the safer side. Even when the upgrading device goes into an issue state, traffic will be inspected in any one of the nodes.

- Refer the *"Supported Upgrade Path—Standalone Nodes"* to know about the order of build version in which the GigaVUE-OS needs to be upgraded.
- Save the configuration and take the backup of both Primary and Secondary node.
- First upgrade the secondary node so that traffic will get inspected in the Primary node.
- Once the Secondary node upgrade gets completed then go for the Primary node upgrade.
- When the Primary node upgrade is in progress, traffic will be inspected in the Secondary node as the failover kicks in, with the help of the signaling port.
- As soon as the Primary node upgrade is complete, the signaling port will come up, and the Primary Node will start inspecting the traffic.

Rules and Notes

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

1. The signaling port type should be a stack port, and only one port should be used.
2. All the inline components should be located in the same box within the cluster.
3. Adding the Inline Networks in the Inline Network Bundle and deploying the Solutions is recommended, which will be easy to export and import across the GRIP Nodes.
4. It is recommended to create a separate redundancy profile for each solution's sources (Inline Network) deployed from GigaVUE-FM Canvas.
5. Before deleting any Solution/Map, it is recommended to move the respective redundancy profile protection role to suspended and revert back to the respective protection role once the solution is redeployed.

6. Post-reload or Power Cycle, the signaling port link stays down when the redundancy profile is attached to the inline network and no maps have been configured. A map should be configured to bring the signaling port up. If a map exists, the signaling port will appear without any issues.
7. Link Failure Propagation is not recommended for copper ports (TAP card ports). Fabric ports support LFP only in a single path (a-to-b only) is available. In all other cases, it is best to leave LFP enabled.
8. Gigamon Resiliency for Inline Protection (GRIP™) is not supported in GigaVUE-HCT devices.
9. GRIP is not supported in other GigaVUE TA Series devices due to the absence of BPS modules.
10. Refer to [Flexible Inline Arrangements—Rules and Notes](#), which also apply to GRIP.

Limitations

- Link failure propagation is not recommended when inline network ports involve copper ports (TAP card ports) or fabric ports with only a single available path (a-to-b only). In all other cases, enabling LFP is recommended.

Troubleshoot

If any of the below issues occur, kindly follow the given steps to troubleshoot the issue.

Signaling Ports Down

1. Check if the power is proper from the optics; if not, try replacing the optics with new ones in a maintenance window
2. If the issue occurs after reloading, check if the maps are configured on the inline network to which the redundancy profile is attached.
3. If maps are not configured, deploying a map will bring up the signaling ports.
4. If a map is available and the signaling ports are still down, contact Gigamon Support for assistance.

Traffic outage in Inline Tool

1. Check if the Redundancy Control State is set to Primary Forwarding in the Primary node or Secondary Forwarding in case traffic is handled in the Secondary Node.

2. Check if the Inline Tool Flex Traffic Path is configured as “Bypass” by mistake. If so, revert to To-Inline-Tool to recover the traffic.
3. Check if the Inline Tool ports are in downstate and failover kicked into tool bypass (the default failover for Inline Tool). If so, correct the optics power, cable, and external Inline Tool faultiness.
4. If the issue persists, contact Gigamon Support for assistance.

Network Traffic Outage

1. Check if any of the Inline Tool's Flex Traffic Paths is set to “Drop” by mistake. If so, revert the Flex Traffic Path to To-Inline-Tool.
2. Check if any Inline Tool ports are down and failover kicked into tool-drop or network-drop. If so, check the optics power, cable, and External Inline Tool Faultiness and correct the same.
3. Check if the traffic hits the map and drops in any Inline Component ports using 'show port stats port-list <port-alias>.'
4. Try redeploying the Flex Inline Solution from GigaVUE-FM and check if the traffic resumes. If not, contact Gigamon Support for assistance.

FAQs

This section answers frequently asked questions when configuring the Gigamon Resiliency for Inline Protection feature.

1. How many members can we have in the GRIP configuration?

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

2. What happens if the primary node becomes unresponsive?

Even if the Primary node becomes unresponsive, traffic should not be impacted as the unresponsiveness is from Software, whereas traffic will be forwarded to the Hardware level. Also, unresponsiveness in the Primary Node will not trigger the failover switch to the Secondary.

3. Is GRIP Supported in Flexible Inline Maps?

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

4. What happens if we delete a map on the primary?

When we delete a map on the primary node, irrespective of the inline-network traffic path, the traffic is switched to the secondary node.

5. Will the configuration between the Primary and Secondary nodes get synchronized automatically?

The user needs to ensure that all the configuration changes made in the Primary node are also made in the Secondary node. Refer to Configuration Synchronization for more details.

Migrate Inline Bypass(Classic) Solution from version 6.5 to 6.12

This section explains how to migrate Inline Bypass (Classic) solutions from version 6.5.xx or earlier to version 6.12.xx.

Migration Scenarios

Based on your device model, you may need to migrate your solution or upgrade your device. Use the following table to determine the required action.

Device Model	Current Version	Migrate or Upgrade	Notes
GigaVUE-HC2	6.5.xx or earlier	Migrate the configurations to GigaVUE-HC1-Plus running on version 6.12	GigaVUE-HC2 is end of support.
GigaVUE-HC1-Plus	6.5.xx or earlier	Migrate the configurations to GigaVUE-HC1-Plus running on version 6.12	Support for Classic inline bypass solution is reintroduced in version 6.12.xx.
GigaVUE-HC1 or GigaVUE-HC3	6.5.xx or earlier	Upgrade the device	Supported upgrade path: 6.5 > 6.9 > 6.12 Refer to Supported Upgrade Path—Standalone Nodes

Migration Considerations

When migrating Inline Bypass configurations from GigaVUE-HC2 to GigaVUE-HC1-Plus, review the port differences carefully:

- Inline network that uses protected ports for Inline bypass configuration on GigaVUE-HC2 must be mapped to the available protected ports on GigaVUE-HC1-Plus.
- Inline tools that used unprotected ports on GigaVUE-HC2 must be mapped to available unprotected ports on GigaVUE-HC1-Plus.
- If breakout ports are in use, configure the same breakout mode on GigaVUE-HC1-Plus.

- The following table provides some examples of port mapping between GigaVUE-HC2 and GigaVUE-HC1-Plus.

Mapping Examples	GigaVUE-HC2 Port (Type/Role)	GigaVUE-HC1-Plus Port (Role/Type)
Inline Network (Protected) Ports pair	1/1/c1 - 1/1/c2 (protected, from BPS card)	1/2/x1 ↔ 1/2/x2 (protected ports)
Inline Tool (Unprotected) Ports	1/2/x1 (unprotected, from BPS card)	1/1/x1 (onboard fixed port)

Migrate Classic Inline Bypass Solution using GigaVUE-OS CLI

Follow these steps to migrate configurations from GigaVUE-HC2 or GigaVUE-HC1-Plus running 6.5 or earlier to GigaVUE-HC1-Plus running version 6.12:

1. Prepare the Source Configuration
 - a. Generate the running configuration.
 - b. Save the configuration as a text file and upload it to a remote server using the following command:
configuration text generate active running save <file name> upload <remote server path>
2. Edit the Configuration File
 1. Download the configuration file from the remote server and update it to match the GigaVUE-HC1-Plus device. Modify the following parameters:
 - IP address – If static, update with the new GigaVUE-HC1-Plus IP. DHCP addresses will not appear in the file.
 - Hostname – Optional.
 - Logging configurations – No changes required.
 - Local user accounts – Optional. Update username and password as needed.
 - Chassis and card configurations – Update to match the GigaVUE-HC1-Plus device.
 - Port configurations – Update to align with GigaVUE-HC1-Plus ports. See [Migration Considerations](#)
 - Inline networks, Inline tools, and map configurations – Update to reflect port configuration changes.
3. Ensure your GigaVUE-HC1-Plus is running on version 6.12.xx .
4. Apply the configuration to GigaVUE-HC1-Plus running on version 6.12.xx:
 - a. Upload the edited configuration text file to your remote server.
 - b. On the GigaVUE-HC1-Plus device, fetch the configuration file using the following command:
config text fetch <path-to-file-in-remote-server>

- c. Apply the configuration using the following command:
config text file <filename> apply fail-continue

NOTE: If mismatches occur, the system skips invalid entries and displays details of the failed configurations.

5. Verify your configurations
 Run the following **show** commands to confirm that the configurations were applied successfully.
 - a. **show inline-network**
 - b. **show inline-tool**
 - c. **show map**

Migrate Classic Inline Bypass Solution using GigaVUE-FM

Follow these steps to migrate configurations from GigaVUE-HC2 or GigaVUE-HC1-Plus running version 6.5.xx or earlier to GigaVUE-HC1-Plus running version 6.12.xx:

1. Download running configuration of your GigaVUE-HC2 or GigaVUE-HC1-Plus in version 6.5.xx to the remote server. Follow the below steps:
 - a. On the left navigation pane, click , and then select **Physical>Nodes**. Select your GigaVUE-HC2 or GigaVUE-HC1-Plus.
 - b. Go to **Settings > Config and Licenses > Configurations**. The Configurations page appears.
 - c. Select **Actions > Save as Text** to save the currently running configuration as text file.
 - d. Select the running configuration file and click **Actions > Download**.
 - e. Upload the downloaded configuration text file to your remote server.
2. Edit the Configuration File
 1. Download the configuration file from the remote server and update it to match the GigaVUE-HC1-Plus device. Modify the following parameters:
 - IP address – If static, update with the new GigaVUE-HC1-Plus IP. DHCP addresses will not appear in the file.
 - Hostname – Optional.
 - Logging configurations – No changes required.
 - Local user accounts – Optional. Update username and password as needed.
 - Chassis and card configurations – Update to match the GigaVUE-HC1-Plus hardware.

- Port configurations – Update to align with GigaVUE-HC1-Plus ports. See [Migration Considerations](#)
 - Inline networks, inline tools, and map configurations – Update to reflect port configuration changes.
3. Ensure your GigaVUE-HC1-Plus is running on version 6.12.xx.
 4. Apply the configuration to GigaVUE-HC1-Plus running on version 6.12 as follows:
 - a. On the left navigation pane, click , and then select **Physical>Nodes**.
 - b. Select the GigaVUE-HC1-Plus where you want to migrate the Inline Bypass configuration.
 - c. Select **Settings > Config and Licenses > Configurations**. The Configurations page appears.
 - d. Click **Actions > Upload**. Upload the edited configuration files from your local drive.
If you have any other map configurations in your device clear them before you apply your configuration text file.
 - e. Select the configuration text file and click **Actions > Apply**. In the Apply Configurations dialog box select *Fail Continue* and then click **Apply** to apply the configurations.

NOTE: If mismatches occur, the system skips invalid entries and displays details of the failed configurations.

5. Verify the Configurations
Confirm that the configurations were applied successfully in the following pages:
 - a. Inline Networks: On the left navigation pane, click , and then select **Physical>Nodes>System > Inline Bypass> Inline Networks**.
 - b. Inline Tools: On the left navigation pane, click , and then select **Physical>Nodes>System > Inline Bypass> Inline Tools**.
 - c. Maps: On the left navigation pane, click , and then select **Physical>Nodes>Traffic > Maps**.

Inline Bypass Solution Examples

The following sections provide examples of inline bypass solutions. The solutions are presented in an order from simple to complex. Refer to the following:

- [Example 1: Unprotected Inline Bypass with an Inline Tool Group](#)
- [Example 2: Unprotected Inline Bypass with Default Heartbeat](#)

- [Example 3: Protected Inline Bypass Using Combo Modules](#)
- [Example 4: Gigamon Resiliency for Inline Protection](#)

Example 1: Unprotected Inline Bypass with an Inline Tool Group

Example 1 is a simple, unprotected inline bypass solution. In the example, aliases are used for inline network ports (iN1 and iN2), inline tool ports (iT1 and iT2), inline network (inNet), inline tool (inTool), and inline map (inMap).

On GigaVUE-HC3, an unprotected inline bypass solution can be configured on the bypass combo module with the inline networks and inline tools on ports 1/1/x1..x16 or on ports c1..c4, or on any other module on the GigaVUE-HC3 node.

On GigaVUE-HC1-Plus, an unprotected inline bypass solution can be configured with the inline networks and inline tools on ports 1/1/x1..x16 or on ports x17..x24, or on any other module on the GigaVUE-HC1-Plus node. Refer to [Figure 52 Logical Bypass Without Bypass Combo Module](#) which shows a GigaVUE-HC1-Plus.



Figure 52 Logical Bypass Without Bypass Combo Module

On GigaVUE-HC1-Plus, an unprotected inline bypass solution can be configured on the base module, with the inline networks and inline tools on ports 1/1/x1..x12 and 1/1/g1..g4, or on the bypass combo module on ports x1..x4.

Task	Description	UI Steps
1.	Configure inline network aliases, port type (inline-network), and administratively enable inline network ports.	1. Go to System > Ports > Ports > All Ports. 2. Click Quick Port Editor. 3. Use Quick search to find the ports to configure. In

Task	Description	UI Steps
		this example, the ports are 3/1/x1 and 3/1/x2 - Set port 3/1/x1 to Type Inline-Network and select Enable. Enter iN1 for the port alias. - Set port 3/1/x2 to Type Inline-Network and select Enable. Enter iN2 for the port alias. - Make sure Enable is selected for Admin on the ports. - Click OK.
2.	Configure inline network.	- Select Inline Bypass > Inline Networks. - Click New. - In the Alias field, type InNet. - For Port A, select iN1. - For Port B, select iN2. - Click Save.
3.	Configure inline tool ports, port type (inline-tool), and administratively enable inline tool ports.	- Go to System > Ports > Ports > All Ports. - Click Quick Port Editor. - Use Quick search to find the ports to configure. In this example, the ports are 3/1/x3 and 3/1/x4 - Set port 3/1/x4 to Type Inline-Tool and select Enable. Enter iT1 for the port alias. - Set port 3/1/x4 to Type Inline-Tool and select Enable. Enter iT2 for the port alias. - Make sure Enable is selected for Admin on the ports. - Click OK.
4.	Configure inline tool, and enable it. Also enable the default heartbeat profile.	- Select Inline Bypass > Inline Tools. - Click New. - In the Alias field, type InTool. - For Port A, select iT1. - For Port B, select iT2. - Under Configuration: - Select Enabled - Select Enabled Heartbeat and set Profile to default. - Click Save.
5.	Configure map passall, from inline network to inline tool.	- Select Maps > Maps. - Click New. - In the Alias field, type InMap. - Select Inline for Type and Pass All for Subtype. - For Source, select InNet

Task	Description	UI Steps
		- For Destination, select InTool. - Click Save.
6.	Configure the path of the traffic to inline tool.	- Select Inline ByPass > Inline Networks. - Select the Inline Network InNet and click Edit - Under Configuration, set the Traffic Path field to To Inline Tool. - Click Save.

Example 2: Unprotected Inline Bypass with Default Heartbeat

Example 2 adds the default heartbeat profile to the unprotected inline bypass solution in Example 1.

Task	Description	UI Steps
1.	Configure inline network aliases, port type (inline-network), and administratively enable inline network ports.	- Go to System > Ports > Ports > All Ports. - Click Quick Port Editor. - Use Quick search to find the ports to configure. In this example, the ports are 3/1/x1 and 3/1/x2 - Set port 3/1/x1 to Type Inline-Network and select Enable. Enter iN1 for the port alias. - Set port 3/1/x2 to Type Inline-Network and select Enable. Enter iN2 for the port alias. - Make sure Enable is selected for Admin on the ports. - Click OK.
2.	Configure inline network.	- Select Inline Bypass > Inline Networks. - Click New. - In the Alias field, type InNet. - For Port A, select iN1. - For Port B, select iN2. - Click Save.
3.	Configure inline tool ports, port type (inline-tool), and administratively enable inline tool ports.	- Go to System > Ports > Ports > All Ports. - Click Quick Port Editor. - Use Quick search to find the ports to configure. In this example, the ports are 3/1/x3 and 3/1/x4 - Set port 3/1/x4 to Type Inline-Tool and select Enable.

Task	Description	UI Steps
		Enter iT1 for the port alias. 16. Set port 3/1/x4 to Type Inline-Tool and select Enable. Enter iT2 for the port alias. 17. Make sure Enable is selected for Admin on the ports. 18. Click OK.
4.	Configure default heartbeat profile.	In GigaVUE-HVUE, the default heartbeat profile is already configured. To view the profile, select Inline Bypass > Heartbeats.
5.	Configure inline tools, and enable it. Also enable the default heartbeat profile.	1. Select Inline Bypass > Inline Tools. 19. Click New. 20. In the Alias field, type InTool1. 21. For Port A, select iT1. 22. For Port B, select iT2. 23. Under Configuration: a. Select Enabled b. Select Enabled Heartbeat and set Profile to default. 24. Click Save.
6.	Configure map passall, from inline network to inline tool.	1. Select Maps > Maps. 25. Click New. 26. In the Alias field, type InMap. 27. Select Inline for Type and Pass All for Subtype. 28. For Source, select InNet 29. For Destination, select InTool. 30. Click Save.
7.	Configure the path of the traffic to inline tool.	1. Select Inline ByPass > Inline Networks. 31. Select the Inline Network InNet and click Edit 32. Under Configuration, set the Traffic Path field to To Inline Tool. 33. Click Save.

Example 3: Protected Inline Bypass Using Combo Modules

Example 3 is a protected inline bypass solution using bypass combo modules on GigaVUE-HC3. It also configures heartbeat and negative heartbeat profiles.

Protected inline networks are based on the pairs of ports associated with the physical protection switches located on the bypass combo modules. Unlike the unprotected examples, you do not need to configure inline network ports because they are created automatically. On GigaVUE-HC3, the port pairs are numbered for example: 2/2/x17 and 2/2/x18, 2/2/x19 and 2/2/x20, 2/2/x21 and 2/2/x22, 2/2/x23 and 2/2/x24.

You do not need to configure inline networks because they are also created automatically on bypass combo modules. The aliases of the default inline networks are: default_inline_net_2_2_1, default_inline_net_2_2_2, default_inline_net_2_2_3, default_inline_net_2_2_4.

On GigaVUE-HC3, protected inline bypass can be configured on the bypass combo module on ports c1..c4.

On GigaVUE-HC1, protected inline bypass can be configured on the bypass combo module. It can also be configured on the TAP-HC1-G10040 module placed in either bay 2 or bay 3, so the ports will be 1/2/g1..g8 or 1/3/g1..g8.

NOTE: The default value of the physical-bypass attribute of protected inline networks is set to enable, which means that the fibers attached to ports net-a and net-b of the inline network are optically coupled and the traffic is exchanged between end nodes without coming to the switching fabric of the GigaVUE node. As shown in Example 4, after configuring the inline tool and the map passall, the physical-bypass attribute is set to disable in order to activate the Inline Bypass solution.

Task	Description	UI Steps
1.	Configure inline tool aliases, port type (inline-tool), and administratively enable inline network ports.	- Go to System > Ports > Ports > All Ports. - Click Quick Port Editor. - Use Quick search to find the ports to configure. In this example, the ports are 2/2/x11 and 2/2/x12 - Set port 2/2/x11 to Type Inline Tool and select Enable. Enter iT1 for the port alias. - Set port 2/2/x12 to Type Inline Tool and select Enable. Enter iT2 for the port alias. - Click OK.
2.	Configure a heartbeat profile.	- Select Inline Bypass > Heartbeats. - Click New. - In the Alias field, type hb2. - For Type, select Regular. - Click Save.
3.	Configure negative heartbeat profile alias and PCAP file	- Select Inline Bypass > Heartbeats. - Click New.

Task	Description	UI Steps
		11. In the Alias field, type nhb1. 12. For Type, select Negative. 13. Click the Browse button for Custom Format and upload the pcap file; for example, hnb.pcap. 14. Click Save.
4.	Configure inline tools, and enable them. Also specify the heartbeat profile.	1. Select Inline Bypass > Inline Tools. 15. Click New. 16. In the Alias field, type InTool1. 17. For Port A, select iT1. 18. For Port B, select iT2. 19. Under Configuration: a. Select Enabled b. Select Enabled Regular Heartbeat and set Profile to hb2. c. Select Enabled Negative Heartbeat and set Negative Heartbeat Profile to nhb1. 20. Click Save. Configure the second inline tool. 1. Select Inline Bypass > Inline Tools. 21. Click New. 22. In the Alias field, type InTool2. 23. For Port A, select iT3. 24. For Port B, select iT4. 25. Under Configuration: a. Select Enabled b. Select Enabled Heartbeat and set Profile to hb_custom. 26. Click Save.
5.	Configure the inline tool group and enable it.	1. Select Inline Bypass > Inline Tool Group 27. Click New. 28. In the Alias field, type inToolGroup. 29. For Inline Tools, select InTool1 and InTool2. 30. Under Configuration, select Enabled. 31. Click Save.
6.	Configure map passall, from inline network to inline tool.	1. Select Maps > Maps. 32. Click New. 33. In the Alias field, type InMap. 34. Select Inline for Type and Pass All for Subtype.

Task	Description	UI Steps
		35. For Source, select default_inline_net_2_2_1 36. For Destination, select InTool1 37. Click Save.
7.	Configure the path of the traffic to inline tool and disable the physical bypass	1. Select Inline ByPass > Inline Networks. 38. Select the Inline Network default_inline_net_2_2_1 and click Edit 39. Under Configuration, set the Traffic Path field to To Inline Tool. 40. Make sure Physical Bypass is not selected. 41. Click Save.

Example 4: Gigamon Resiliency for Inline Protection

Example 5 is an inline bypass solution for GRIP using TAP-HC0-G100C0 modules on the GigaVUE-HC2, with copper ports.

First, configure the GigaVUE-HC2 with the primary role, then configure the GigaVUE-HC2 with the secondary role. The configuration is the same (is synchronized) on both nodes, except for step 3, in which the protection role (primary or secondary) is specified.

Note that in this example, link fail propagation (LFP) is disabled to reduce inline network recovery time after failover. When a primary to secondary failover occurs and LFP is enabled for copper inline bypass links, network service recovery may take several seconds because of Ethernet link renegotiation. Optical links failover faster and typically recover service much faster. For inline networks where only one path is available, this is a consideration. When GRIP is deployed with high availability networks where a second path is present, it is a best practice to leave LFP enabled.

You can use the Chassis page to view the chassis and modules.

Configure Primary Role GigaVUE-HC2

Task	Description	UI Steps
1.	Configure ports on the TAP-HC0-G100C0 module as passive (in	1. Go to System > Ports > Ports > All Ports . 1. Select a port of the TAP-HC0-G100C0 module.

Task	Description	UI Steps
	passive mode, relays are closed). Also configure ports, port type (inline-network).	2. Click to open the Ports page. 3. Select Passive for TapTX 4. Click Save. 5. Repeat steps 2 through 6 for each port on the TAP-HC0-G100C0 module 6. Configure Inline Network ports a. Select the port. b. Click Quick Port Editor. g. Select Inline Network for Type. Note: You can use the Chassis page to locate the position of the module in the chassis and identify port IDs.
2.	Configure stack port (for signaling port/link) and enable it.	1. Select the port and click Edit. 8. Select Enable for Admin. 9. Select Stack for Type. 10. Click OK.
3.	Create the redundancy profile by giving it a name and configuring parameters for the redundancy profile such as the signaling port and protection role (primary).	1. Select Inline Bypass > Redundancies 11. Click New. 12. Enter a name for the profile in the Alias field. For example, RP_001. 13. Click in the Signaling Port field and select the stack port configured in Task 2. 14. Select Primary for Protection Role. 15. Click Save.
4.	Configure inline network. This step associates the redundancy profile to the inline network and also disables link fail propagation on the inline network.	1. Select Inline Bypass > Inline Tools. 16. Click New. 17. Enter a name for the inline network in the Alias field. For example, IN_001. 18. Click in the Port A field and select inline network port for network A. 19. Click in the Port B field and select an inline network port f. 20. Select Bypass for Traffic Path. (This is the default setting.) 21. Make sure Link Failure Propagation option is NOT checked. (It is enabled by default.) 22. Click in the Redundancy Profile field and select the profile created in step 3. For example RP_001.
5.	Configure inline tool ports, port type (inline-tool), and administratively enable them.	1. Go to System > Ports > Ports > All Ports. 23. Select the first port (for example, 1/4/x1) to configure

Task	Description	UI Steps
		as an inline-tool port. 24. Click Quick Port Editor. 25. Select Inline Network for Type and select Enable for Admin. 26. Click OK. 27. Select the second port (for example, 1/4/x2) and repeat steps 3 through 5.
6.	Configure inline tool and failover action. then enable inline tool.	1. Select Inline Bypass > Inline Tools 28. Enter a name for the inline tool in the Alias field. For example, IT_001 29. Click in the Port A field and select the first inline tool port configured in Task 5. 30. Click in the Port B field and select the second inline tool port configured in Task 5. 31. Make sure Enable is selected. (It is enabled by default) 32. Select NetworkBypass for Failtover action. 33. Click Save.
7.	Configure map passall, from inline network to inline tool. Note: When you delete a map on the primary node, irrespective of the inline-network traffic-path, the traffic is switched to the secondary node. The port utilization must be 0% on the primary node and active on the secondary node.	1. Select Maps > Maps > Maps. 2. Click New. 3. Enter a name for the map in the Alias field. For example, INtoIT. 4. Select Regular for Type. 5. Select Pass All for Subtype. 6. Select the inline network created in Task 4 for Source. For example, IN_001. 7. Select the inline tool created in Task 6. For example, IT_001.

Configure Secondary Role GigaVUE-HC2

Task	Description	UI Steps
1.	Configure ports on the TAP-HC0-G100C0 module as passive (in passive mode, relays are closed). Also configure ports, port type (inline-network).	1. Go to System > Ports > Ports> All Ports. 8. Select a port of the TAP-HC0-G100C0 module. 9. Click to open the Ports page. 10. Select Passive for TapTX 11. Click Save. 12. Repeat steps 2 through 6 for each port on the TAP-HC0-G100C0 module 13. Configure Inline Network ports

Task	Description	UI Steps
		a. Select the port. b. Click Quick Port Editor. n. Select Inline Network for Type. Note: You can use the Chassis page to locate the position of the module of the module in the chassis and identify port IDs
2.	Configure stack port (for signaling port/link) and enable it.	1. Select the port and click Edit. 15. Select Enable for Admin. 16. Select Stack for Type. 17. Click OK.
3.	Create the redundancy profile by giving it a name and configuring parameters for the redundancy profile such as the signaling port and protection role (secondary).	1. Select Inline Bypass > Redundancies 18. Click New. 19. Enter a name for the profile in the Alias field. For example, RP_001. 20. Click in the Signaling Port field and select the stack port configured in Task 2. 21. Select Secondary for Protection Role. 22. Click Save.
4.	Configure inline network. This step associates the redundancy profile to the inline network and also disables link fail propagation on the inline network.	1. Select Inline Bypass > Inline Tools. 23. Click New. 24. Enter a name for the inline network in the Alias field. For example, IN_001. 25. Click in the Port A field and select inline network port for network A. 26. Click in the Port B field and select an inline network port f. 27. Select Bypass for Traffic Path. (This is the default setting.) 28. Make sure Link Failure Propagation option is not checked. (It is enabled by default.) 29. Click in the Redundancy Profile field and select the profile created in step 3. For example RP_001.
5.	Configure inline tool ports, port type (inline-tool), and administratively enable them.	1. Go to System > Ports > Ports > All Ports. 30. Select the first port (for example, 1/4/x1) to configure as an inline-tool port. 31. Click Quick Port Editor. 32. Select Inline Network for Type and select Enable for Admin. 33. Click OK. 34. Select the second port (for example, 1/4/x2) and

Task	Description	UI Steps
		repeat steps 3 through 5.
6.	Configure inline tool and failover action. then enable inline tool.	1. Select Inline Bypass > Inline Tools 35. Enter a name for the inline tool in the Alias field. For example, IT_001 36. Click in the Port A field and select the first inline tool port configured in Task 5. 37. Click in the Port B field and select the second inline tool port configured in Task 5. 38. Make sure Enable is selected. (It is enabled by default) 39. Select NetworkBypass for Failtover action. 40. Click Save.
7.	Configure map passall, from inline network to inline tool.	1. Select Maps > Maps > Maps. 41. Click New. 42. Enter a name for the map in the Alias field. For example, INtoIT. 43. Select Regular for Type. 44. Select Pass All for Subtype. 45. Select the inline network created in Task 4 for Source. For example, IN_001. 46. Select the inline tool created in Task 6. For example, IT_001.
8.	Configure the path of the traffic to the inline tool, disabling physical bypass on the inline network to open the relay on the node with the primary role.	1. Select Inline Bypass > Inline Networks. 47. Select the Inline Network configured in Task 4 and click Edit. 48. Select To Inline Tool for Traffic Path 49. Make sure Physical Bypass is NOT checked. 50. Click Save.

Timestamps

This chapter describes how to use the timestamp feature. The timestamp feature provides accurate time information to measure application and network performance.

Required License: Advanced Feature License

Refer to the following sections for details:

- [About Timestamps](#)
- [Using Timestamps](#)
- [Why PTP?](#)
- [Using PTP to Timestamp Packets](#)
- [Enabling Timestamp on a Port](#)
- [Viewing PTP Details](#)
- [PTP and Timestamp—Rules and Notes](#)

About Timestamps

Timestamps provide high-definition timing information for installations where packet timing is critical, particularly in high frequency trading, financial, and service provider environments. Timestamps are easily correlated across worldwide installations, improving your ability to monitor latency, Quality of Service (QoS), jitter, and end-to-end network performance across multiple instrumentation points.

This feature provides you the ability to append timestamps to ingress or egress packets. Timestamps are used in conjunction with Precision Time Protocol (PTP) for synchronization with a UTC reference clock.

You can choose to append timestamps to ingress packets, or to egress packets, or to both ingress and egress packets. [Figure 53Timestamping](#) illustrates how timestamps are added to ingress and egress packets.




Figure 53 *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	<i>"ptp" command section in the "GigaVUE-OS CLI Reference Guide"</i>
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	<i>"Related Commands" section under the "ptp" command section in the "GigaVUE-OS CLI Reference Guide"</i>
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	<i>"port" command section in the "GigaVUE-OS CLI Reference Guide"</i>
4.	View the PTP statistics for troubleshooting purposes.	Viewing PTP Details	<i>"Related Commands" section under the "ptp" command section in the "GigaVUE-OS CLI Reference Guide"</i>

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 54 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 55 PTP State Transition](#)).




Figure 55 PTP State Transition

Limitations

The following restrictions are applicable for the above feature:

- A two step PTP clock mode is not supported.
- One step GigaVUE-TA200 clock cannot synchronize time with a two step PTP Primary Source.

Synchronization of the PTP and Local System Clock

From 5.13.01 version, PTP clock will be synchronized with the local system clock to provide accurate UTC timestamping. This process takes over in the absence of a Primary source clock and when the node is acting as a source clock. In this scenario, if there is a change in the local system clock due to NTP or manual configuration the same change will be synced to the PTP clock.

Recommendation: All local clocks that are running on PTP domain should share the same time value. This can be ensured through NTP server.

Using PTP to Timestamp Packets

Refer to the following sections about how to configure and enable PTP on the network ports to timestamp packets:

- [Configuring PTP Globally](#)
- [Enabling PTP on a Network Port](#)

Configuring PTP Globally

You must configure PTP globally on a device. You can configure the PTP domain, clock mode, and priority at the device level. Based on these PTP parameters, the primary source will be selected from the clocks in a network.

To configure PTP globally:

1. From the left navigation pane, go to **Inventory > Physical > Nodes..**
2. In the Physical Nodes page, click the required cluster ID for which you want to configure PTP.
3. On the left-navigation pane, go to **System > Settings > Date and Time > PTP.**
4. Click **New** to open the PTP page.
5. In the **Alias** field, enter a unique name for the PTP template that you are configuring globally on the device. Ensure that you provide a unique name for each device in the cluster. You can also choose to leave the field blank and a default alias will be added automatically. For example, ptp_default_<box_id>.
6. From the **Box ID** drop-down list, select the box ID of the device for which you want to configure PTP.
7. In the **Domain** field, enter a PTP domain number that will be used for this clock. The valid range includes 0 and the numbers between 24 and 43.
 - Domain 0 for IEEE 1588 default profile.
 - Domain 24 - 43 for ITU-T G.8275.1 Telecom profile.
8. From the **Mode** drop-down list, select **Boundary** or **Ordinary** as the clock mode. The default is **Boundary**.
9. In the **Priority 2** field, enter a priority value that will be used to determine the primary source in a network. For example, if there are two clocks in a network that match the default criteria, the clock that has the lower priority value will be selected as the primary source. The valid range is between 0 and 255. The default is 128.

10. In the **Local Priority** field, enter a priority value that will be used to determine the PTP source when the BMC algorithm chooses more than one clock as the source. This value overrides the default criteria for the BMC selection. The clock with the lower priority value will be selected.

For example, on a particular device, the local priority of a port is compared with the local priority of the clock. Based on the following criteria, the source is determined:

- If the local priority of the port is greater than the local priority of the clock, the clock becomes the source.
- If the local priority of the port is lesser than the local priority of the clock, the port becomes the source.
- If the local priority of the port is equal to the local priority of the clock, the clock identity of the port and the clock is compared and the one that has the lower value becomes the source.

The valid range is between 1 and 255. The default is 128.

11. Click **OK**.

Enabling PTP on a Network Port

After you configure PTP globally on a device, you must enable PTP on the required network ports. You can also use the fanout ports to configure PTP. Once you have enabled PTP on a port, you cannot use the port for any other services such as maps. To configure PTP on a network port:

1. From the left navigation pane, go to **Inventory > Physical > Nodes..**
2. In the Physical Nodes page, click the required cluster ID for which you want to configure timestamp.
3. On the left-navigation pane, go to **Ports > Ports > All Ports.**
4. Select a network port on which you want to enable PTP, and then click **Edit**.
5. Scroll down to the PTP section, and then select the **Enable** check box.
6. In the **Local Priority** field, enter a priority value that will be used to determine the PTP source when the BMC algorithm chooses more than one clock as the source. This value overrides the default criteria for the BMC selection. The clock with the lower priority value will be selected.

For example, on a particular device, the local priority of a port is compared with the local priority of the clock. Based on the following criteria, the source is determined:

- If the local priority of the port is greater than the local priority of the clock, the clock becomes the source.
- If the local priority of the port is lesser than the local priority of the clock, the port becomes the source.

- If the local priority of the port is equal to the local priority of the clock, the clock identity of the port and the clock is compared and the one that has the lower value becomes the source.

The valid range is between 1 and 255. The default is 128.

7. From the **Role** drop-down list, select one of the following options:

- **Source**—The port will act as the source with which all the connected ports with receiver role will be synchronized.
- **Follower**—The port will act as the receiver and will be synchronized with the source port to which it is connected.
- **Standard**—The BMC algorithm will determine the port's role as source or receiver.

NOTE: The default is **Standard**.

8. In the **Announce Interval** field, enter the time interval, in seconds on a logarithmic scale, between PTP announce messages in the port.
 - For Domain 0, the valid range is between -2 and 4 log seconds. The default is 1.
 - For Domain 24-43, the valid range is between -3 and 4 log seconds. The default is -3.
9. In the **Delay Request Interval** field, enter the minimum interval allowed between PTP delay request messages. The valid range is between -7 and 0 log seconds.
 - For Domain 0, the default is -5.
 - For Domain 24-43, the default is -4.
10. In the **Sync Interval** field, enter the time interval between PTP synchronization messages on the port. The valid range is between -7 and 0 log seconds.
 - For Domain 0, the default is -5.
 - For Domain 24-43, the default is -4.
11. In the **VLAN** field, enter the PTP VLAN value for the port. The valid range is between 2 and 4000. The VLAN ID would be allocated based on availability of the VLAN resource.
 Hover over the help icon  to know about the available VLAN resources that can be allocated.
12. Click **OK** to save the configuration.

Enabling Timestamp on a Port

You can configure ingress timestamp, egress timestamp, or both ingress and egress timestamps on a port. Refer to [Configuring PTP Globally](#) for the ports that are supported for timestamps.

To configure timestamp on a port:

1. From the left navigation pane, go to **Inventory > Physical > Nodes..**

2. In the Physical Nodes page, click the required cluster ID for which you want to configure timestamp.
3. Go to **System > Ports > Ports > All Ports**.
4. Select a port on which you want to configure timestamp, and then click **Edit**.
5. Scroll down to the Timestamp section.
6. Select the **Enable** check box corresponding to the **Ingress** or **Egress** fields.
7. In the **Source ID** field, enter a unique ID for the ingress or egress timestamp. The packets that ingress or egress the port will be timestamped with this unique source ID. The source ID enables you to identify the port on which the timestamp is applied.
8. Click **OK** to save the configuration.

▼ Timestamp

Ingress	☒ Enable	Source ID	0	?
Egress	☒ Enable	Source ID	0	?

Figure 56 *Timestamp Configuration*

After you enable timestamp on the various ports, you can view the timestamp details for a list of ports in the Timestamp page. To view this page, go to **Settings > Date and Time > Timestamp**.

Viewing PTP Details

Following table provides information about the different PTP details that you can view:

To view:	Refer to:
The PTP configuration on a port	Viewing the PTP Configuration on a Port
The PTP clock details such as the local PTP clock, parent clock state, and the time property for a device	Viewing the PTP Clock Details
The PTP details of the source that is learned by a port	Viewing the PTP Foreign Source Details
The PTP statistics	Viewing PTP Statistics

Viewing the PTP Configuration on a Port

After you enable PTP on a network port, you can view the port's clock state and other details of PTP configuration in the PTP Port State page. To view the page, go to **Settings > Date and Time > PTP Port State**. [Figure 57 PTP 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 57 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 its datasets, hardware, and communication facilities. The port does not place any PTP messages on its communication path during this state. If one port on the boundary clock is in the Initializing state, all the other ports on the boundary clock will also be in the same state. • Listening—Only signaling messages or management messages that are responses to another management message are placed on the port's communication path. • Faulty—The fault state of the protocol. Only signaling messages or management messages that are responses to another management message are placed on the port's communication path. In a boundary clock, any activity on the port with the Faulty clock port state does

PTP Port State	Description
	not affect the other ports in the device. If any activity on the faulty port is not restricted to the port, all other ports in the device changes to the Faulty state. • Disabled—The port does not place any messages on it's communication path. In a boundary clock, the port's activity does not affect the activity in other ports that are part of the boundary clock. A port in this state discards all PTP received messages except the management messages. • Pre-master—The port acts like a source but does not place any messages on it's communication path, except for signaling or management messages. • Master—The port becomes the source. • Passive—The port does not place any messages on it's communication path. • Uncalibrated—One or more source ports are detected in the domain. The appropriate source port has been selected, and the local port is preparing to synchronize to the selected source port. This is a transient state to initiate the servo synchronization, datasets updates, and other implementation-specific activities. • Slave—The port is synchronizing to the selected source port.
Local Priority	The priority value used to determine the PTP source when the BMC algorithm chooses more than one clock as the source. This value overrides the default criteria for the BMC selection. The clock with the lower priority value will be selected. The valid range is between 1 and 255. For example, on a particular device, the local priority of a port is compared with the local priority of the clock. Based on the following criteria, 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.
Peer Delay Request Interval	The minimum interval allowed between PTP delay request messages between peer ports.
Sync Interval	The time interval between PTP synchronization messages on the port.

PTP Port State	Description
Announce Interval	The time interval, in log seconds, between PTP announce messages in the port.
Announce Receipt Timeout	The maximum number of consecutive announce messages that the port fails to receive. When the number of consecutive announce messages that the port failed to receive is greater than the value specified in this field, the port's clock state changes to source.
VLAN ID	The identifier of the VLAN interface.
Clock Port ID	The port identifier assigned by the PTP clock.
PTP Version	The PTP version that is enabled in the port.
Delay Mechanism	The mechanism used to determine the delay between a source and a receiver. Following is the delay mechanism that Gigamon supports: • End-to-End—The source and the receiver send delay requests and delay responses back and forth between themselves to measure the delay. Once the delay is measured, the receiver synchronizes the time with the source.
Peer Mean Path Delay	Reserved for future use.

Viewing the PTP Clock Details

Use the PTP Clock page to view the local PTP clock, parent clock state, and the time property for the devices for which you have configured PTP. To view this page, go to **Settings > Date and Time > PTP Clock**. [Figure 58 PTP Clock Details](#) shows the PTP clock details.

Local PTP Clock

Alias	Box ID	Clock Identity	PTP Time(sec)	Steps Removed	Mean Path Delay	Offset From Source	
ptp2	3	00:1d:ac:00:00:68:17:04	1632393577.233307057	1	0.000000130	-21386510(ns)	

< < Go to page: 1 of 1 > > Total Records: 1

Parent Clock State

*PS- Primary Source

Alias	Box ID	Clock Identity	PS Clock Identity	PS Priority1	PS Priority2	Observed Offset	
ptp2	3	00:1d:ac:00:00:66:0c:04	00:1d:ac:00:00:66:0c:04	128	50	0	

< < Go to page: 1 of 1 > > Total Records: 1

Time Property

Alias	Box ID	Current UTC Offset	Leap59	Leap61	Time Traceable	Frequency Traceab...	Time Source	PTP Time S...	
ptp2	3	36	0	0	0	0	Internal Oscillator	1	

< < Go to page: 1 of 1 > > Total Records: 1

Figure 58 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 58 PTP Clock Details](#), you can see that there is no foreign source record 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.

Port ID/P...	Box Id	Best Source	Port Iden...	Primary Source Priority1	Primary Source Priority2	Last Sy...	Last Fo...	Last Dela...	Announc...	S...	Primary Sou...	Port Mo...	Po...
1/1/c29	3	0	00:1d:ac:00:00:00	128	50	10	9	41	3	0	00:1d:ac:00:00:00	0	5

Figure 59 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 60PTP Statistics](#) shows the PTP statistics details.

☐	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 60 PTP Statistics

Following table lists the PTP statistical attributes and their descriptions:

PTP Statistical Attribute	Description
Alias	The unique name for the PTP template that you configured globally for the device.
Box ID	The box identifier for the device.
Discarded Packets	The number of packets that were discarded by the port because of mismatch in the domain or VLAN.
Rx Packets	The number of packets that were received by the port.
Tx Packets	The number of packets that were transmitted by the port.
Queue Overflow Rx Packets	The queue overflow of the received packets.
Rcpu Encap Rx Packets	The RCPU-encapsulated packets received.
Ipv4 Ptp Rx Packets	The IPv4 PTP packets received.
Ipv6 Ptp Rx Packets	The IPv6 PTP packets received.
L2 Ptp Rx Packets	The L2 PTP packets received.
Udp Ptp Rx Packets	The UDP PTP packets received.

PTP and Timestamp—Rules and Notes

Keep in mind the following rules and notes when working with PTP and timestamp:

- PTP is supported only on the GigaVUE-TA200 and GigaVUE-TA200E devices.
- Timestamp is supported only on the GigaVUE-TA200 device.
- Timestamp is not supported on circuit, stack, and hybrid ports.
- PTP supports only one-step synchronization.

- PTP requires a dedicated front-end port and can be enabled only on network ports. You cannot use the PTP-enabled port for any other GigaVUE-OS features.
- Ensure that the VLAN ID and the domain for the PTP are the same for both the nodes between which the session needs to be established. However, the VLAN IDs of the ports within a node must be unique.
- All port-related PTP intervals such as, announce, sync, and delay request must be the same between nodes.
- Transparent clocks and peer-to-peer delay request mechanism are not supported.
- Priority 1 configurations are not supported based on the ITU-T G.8275.1 standard.
- Once you have configured PTP and timestamps are being captured, it is recommended that you do not change the PTP configurations so that the timestamps are accurate.
- If you want to upgrade to release version 5.13.00 with a VLAN configured PTP enabled port, it is recommended to remove the configurations before the upgrade and re-apply as required after the upgrade.

Fabric Maps

Fabric maps simplify how you create flow mapping across multiple clusters. Using GigaVUE-FM, you can create and manage fabric maps. You can provide cross-cluster flow mapping parameters by creating a fabric map. GigaVUE-FM then allocates the required circuit ID resources, along with generating and deploying cluster specific maps and fabric paths (refer [About Circuit-ID Tunnels](#)). Fabric paths can be defined as the link between the devices through which the traffic flows. With a successfully deployed fabric map, you can save time and effort because you no longer need to replicate rules at each hop on the network to manage cross-cluster traffic.

Refer [Load Share Flow-Based Traffic to Identical Tools in a Fabric Tool Group Using Fabric Maps](#) for more detailed information.

Supported Topologies

Fabric map support is available for the following topologies:

- [Multihop Topology](#)
- [Leaf and Spine Topology](#)
- [Dual Multipath Leaf and Spine Topology](#)

Refer to the following rules and notes:

- Source port of a fabric map can be either a network port, hybrid port or a vport.

- Destination port of a fabric map can be a tool port, hybrid port, GigaStream, Fabric Port Group or a vport.

NOTE: Fabric port groups cannot be used as destination for subscriber intelligence solutions with load-balancing application.

- For topologies with leaf and spine clusters:
 - Do not configure the spine nodes as source or destination nodes. Source and destination must be configured either in the leaf nodes or in the node/clusters that are connected to the leaf nodes.
 - You can connect standalone nodes or clusters to all spine nodes of leaf-spine cluster.
 - You can connect the spine nodes to its leaf node. Spine nodes can also be connected to spine nodes of another leaf-spine cluster.
- Do not configure fabric maps in topologies that are not supported. Though GigaVUE-FM allows such configurations, it might lead to unpredictable results.


Multihop Topology

In this topology, multiple standalone nodes and/or non-leaf and spine clusters are connected using circuit GigaStreams.


Refer to the following notes:

- Source and destination nodes can be configured anywhere in the topology.
- Redundant links must not be configured between the nodes, that is, there should not be more than one path from source to destination.

The Standalone nodes or clusters can be configured as follows:




The standalone nodes or clusters can also be configured like a tree topology. Refer to the following figure.




Supported Multihop Topology

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




Unsupported Multihop Topology

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




Leaf and Spine Topology

In this topology, a standalone node or a cluster is connected to any one of the leaf nodes in a leaf-spine cluster. It can also be load balanced to multiple leaf nodes of a leaf spine cluster. Refer to the [Multi-Path Leaf and Spine](#) section for details.



Spine Source Map Mode


The Spine Source Map mode is a subset of leaf-and-spine topology that defines how Spine sources are visualized, either Separate or Combined within a leaf-and-spine fabric map.



The COMBINE setting of this topology optimizes fabric map deployments so that the cluster can use a single logical map to forward the traffic to the destination, even when traffic arrives via multiple decap circuit GigaStreams. This mode introduces a Spine Source Map Mode setting which is global in the FM. Refer to [Configure Fabric Map Settings](#) for more information on how to configure spine source map mode.

Dual Multipath Leaf and Spine Topology

In this topology, a leaf-spine cluster is connected to another leaf-spine cluster by configuring circuit GigaStreams across the spine nodes. Each spine node in a cluster is load-balanced with spine nodes in another cluster.




Supported Failovers

In a dual multipath leaf and spine cluster topology, traffic from one cluster to another cluster will be routed in case of the following failover scenarios:

- **Scenario 1:** At least one of the spine nodes in both clusters is in active state.
- **Scenario 2:** At least one circuit link connecting the spines across the cluster is active.


Scenario 1



In the above topology:

- Spine nodes 1 and 4 are down.
- Spine nodes 2 and 3 are up.
- Traffic is routed from leaf nodes on cluster 1 to leaf nodes on cluster 2 even though spine 1 and spine 4 are down.
- Traffic flows through the green links between spine 2 and spine 3.

Scenario 2



In the above topology:

- All spine nodes are active.
- Circuit GigaStream from Spine 1 to Spine 4 is down.
- Circuit GigaStream from Spine 2 to Spine 4 is down.
- Circuit GigaStream from Spine 1 to Spine 3 is up.
- Circuit GigaStream from Spine 2 to Spine 3 is up.
- Traffic from leaf nodes in cluster 1 are routed to leaf nodes in cluster 2 through circuit links from spine 1 and spine 2 to spine 3.


Thus, even if one of the links is up, traffic is routed across the leaf nodes in the clusters.

Unsupported Failovers

In a dual multipath leaf and spine cluster topology, traffic from one cluster to another cluster will not be routed in case of the following failover scenarios:

- **Scenario 1:** Circuit links connecting the spine nodes across cluster is down
- **Scenario 2:** Spine links connecting the leaf nodes to spine node is down.


Scenario 1



In the above topology:

- Traffic from Leaf 1 and Leaf 2 is load-balanced to spine node 1.
- Traffic received on Spine 1 will not be delivered to the adjacent cluster as the following circuit GigaStreams are down:
 - Circuit GigaStream from Spine 1 to Spine 3
 - Circuit GigaStream from Spine 1 to Spine 4.

Scenario 2



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

- Spine link connecting leaf 1 to spine 1
- Spine link connecting leaf 1 to spine 2

Fabric Maps Prerequisites

To deploy fabric maps, you must complete the following prerequisites:

- Designate inter-cluster connection ports as circuit ports. Refer to [Configure Ports](#)
- Configure a regular circuit GigaStream with the required circuit ports. Refer to [Configure Regular GigaStream](#).
- Create links between clusters. Refer to [Create Links between Clusters](#).



Note: Keep in mind the following notes about circuit-id tunnels when configuring circuit ports:

Circuit-ID tunnels support the following:

- From Port in Regular Map
- GigaStream
- Tool port in regular map and collector map



- Port filter
- Port groups

Circuit-id tunnels does not support the following. However, you are not restricted from configuring these resources as circuit ports, but the functionality does not work as expected.

- Ingress ports with VLAN tag
- Port Pair
- Tool Mirror
- Inline Tool
- Pass-all/collector with from port as circuit-port

For more information about circuit-id tunnels, refer to [Circuit Tunnels](#).

Using GigaVUE-FM, you can create connections between devices using manual topology links, or you can have the devices discovered through the Gigamon Discovery (GDP). After you create physical links between devices, the device configuration needs to be added into GigaVUE-FM, so that Fabric Maps can use it. Once the circuit and GigaStream ports are created and connected to the devices, you can start creating Fabric Maps. See [Enable Gigamon Discovery on Chassis](#) for more information on GDP.

Notes on Circuit Ports and Circuit GigaStream in Fabric Maps

As mentioned in the pre-requisites section, Circuit ports and Circuit GigaStreams are required for creating Fabric Maps. Refer to the following notes:

- You cannot specify the Circuit Ports and Circuit GigaStreams that are to be used in a fabric map. You can only define the Circuit Ports and Circuit GigaStreams.
- Circuit Ports and the Circuit GigaStream for the fabric path are selected as follows:
 - Circuit GigaStreams at the first clusters are selected randomly for efficient load balancing.
 - For hop beyond second hops, previously used Circuit Gigastreams may be selected if 'Resource Sharing' is set to enabled mode. This allows to reuse the circuit tunnels previously created. Refer to [Configure Fabric Map Settings](#).
- For Circuit GigaStreams on the Circuit links that connect two clusters, it is recommended to not have more than one ingress GigaStream in any of the devices of the cluster.

- Do not configure redundant Circuit GigaStream between non leaf-spine clusters as this may lead to various issues such as traffic duplication, source port configuration and other such issues.

Consider the following leaf-spine topology. Fabric map is created between the two leaf nodes and the spine node. The second Circuit GigaStream in the leaf node may be used when the circuit ports are of different speeds.




Figure 61 Circuit GigaStreams in Fabric Maps

Notes for Circuit ID

Keep in mind the following points while configuring the Circuit ID:

- Circuit ID is used internally to pass traffic from one cluster to another. You can configure Circuit ID ranges from 2 to 4000, or set your own custom range, with each topology able to reuse ID ranges (e.g., 2-513) across multiple topologies.
- Circuit ID allocation is managed globally, beginning from the lower limit of the defined range. This ensures that IDs are allocated efficiently across different topologies. It is allocated automatically.
- To enable the Circuit ID optimal allocation feature in your currently deployed maps, it is necessary to delete the existing maps. After deletion, create a new map that includes the optimal allocation feature, as it cannot be applied to the existing maps directly.
- Each cluster can use up to 512 Circuit IDs from the global pool (default range is 2-4000), allowing for 512 unique destination sets per cluster or standalone node. Multiple fabric maps can also be created, each with 512 unique destination sets.
- Circuit IDs are allocated on a per-hop and per-direction basis.
- The Circuit ID optimal allocation is not applicable for second-level maps with Fabric Port Group destination.

- The circuit tunnel is not shared for first-level maps even in Shared mode.

Create Links between Clusters

You can either connect the devices manually or you can have them automatically discovered by Gigamon Discovery Protocol (GDP). If a user's ports are already physically connected, then the link between those ports will be displayed if GDP is enabled at both ends and the physical links are up. GDP must be enabled in both the source and destination ports to have the devices displayed in the topology map.

- [Create Manual Links](#)
- [Create Gigamon Discovery Protocol \(GDP\) Based Links](#)

Create Manual Links

To create links manually, do the following.

1. On the left navigation pane, go to **Inventory**  select **Physical >Topology**.
2. Select **Add Link(s)** option from the **Add** tab drop-down menu on the Topology menu bar.
3. Select the **Source** device and the circuit port which you already created.
4. Select the **Destination** device and the circuit port which you already created.
5. Click **Submit** to connect the devices.




Figure 62 Connecting Devices

- 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:

- From the left navigation pane, go to **Inventory > Physical > Nodes..**
- Select the node (Source or Destination).
- Select **Ports**. On the Ports screen, select the circuit port.
- Click **Edit** from the top menu bar. The Port screen is displayed for the device you selected.




Figure 63 Gigamon Discovery Protocol

- Enable Gigamon GDP. Click **OK** to submit configuration.
- You must then enable GDP on the device chassis as follows:
 - Select **> Chassis**. The chassis screen is displayed showing the Gigamon hardware.
 - Click the **List View** icon located at the top of the page. Details about the chassis is displayed in a list view format.
 - Select **Enable Gigamon Discovery** from the **Actions** menu at the top of the screen.

Box ID 2 - GigaVUE-HC3

Box: ID 2 - GigaVUE-HC3

PROPERTIES

Box ID	Chassis Id/Serial ...	Hardware Type	mode	Gigamon Discovery	Hardw
2	J5227c10000	HC3-Chassis	default	Disabled	1.0

Go to page: 1 of 1 Total Records: 1

CARDS

Slot Id	Hardwar...	Configur...	Health S...	Operatio...	Fabric H...	Filter Te...	Power R...	Power P...	Serial Nu...	Al...
1	PRT-HC...	✓	✓ Slot i...	✓ Up	N/A	None	60	1	10-0...	10-0...

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

This completes the prerequisites steps for creation of Fabric Maps. Refer to section [Create Fabric Maps](#) for details on creating Fabric Maps.

Create Fabric Maps

To create Fabric Maps:

1. Go to **Traffic**  and under **Physical > Orchestrated Flows > Fabric Maps**, click **Create**.

The following table describes the parameters displayed in the fabric map list view.

Field	Description
Alias	Alias name of the fabric map. Note: GigaVUE-FM does not allow you to configure fabric maps alias to start with <i>_FM_</i> , <i>_FM-</i> , <i>FMAuto-</i> . These prefixes are used in auto-generated fabric map alias.
Enabled	Indicates if fabric map is enabled – Yes/No
Source	Number of source ports associated with fabric map
Destination	Number of destination ports associated with fabric map
Rules	Rules for fabric map
GigaSMART Alias	Alias name of the GigaSMART operation

Field	Description
Type	Port type
Status	Health status of the fabric map

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

Create Map

Basic Info

Alias

FabricMap_Reg

DESCRIPTION (optional)

Enable

Resource Allocation

SHARED

Source

Port Editor

Search by Node

10.115.46.135

Select one or more ports...

Rules

Pass all traffic that doesn't match the rule criteria

Quick Editor

Import

Add Rule

GigaSMART Operation

10.115.46.135

Destination

Port Editor

Tool Finder

Search by Node

10.115.46.135

Select one or more ports...

Tags

Create Tag

TagKey

Values

Cancel

Create

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

Field	Description
Basic Info	
Alias	Alias name of the fabric map
Description	Fabric map description
Enable	Toggle option that lets you toggle between: Enable: Traffic passes through the fabric map that is created. Disable: Traffic is temporarily prevented from passing through fabric map if 'Disable' is selected. Note: You cannot disable the second-level maps (option is disabled in the UI). Alternatively, you must disable the corresponding first-level map.
Resource Allocation	Displays whether it uses shared or non-shared circuit id resources.
Type	Map type. Options are: - Regular - First level

Field	Description
	Second level
Subtype	Map subtype. Options are: - By Rule - Pass All - Collector
Source	Source port. Use the Search by Node toggle bar to toggle between the following two options - Enter the port alias/port id (Respective Cluster Id is displayed for selecting the required ports), or - Specify the source node and source ports. Use the +/- icon to add new nodes or remove source ports. Note: You cannot use ports already used as source ports in the orchestrated configuration.
Rules	Enable Non-Matching Rules for map type 'Regular'. Use the toggle bar to either Pass all traffic or not. The following options are available: Quick Editor: To select the port number or numbers to add for a pass or drop rule or both. Import: To import a map template Add a Rule: To add a description and select one or more conditions. Use this option to add rules one at a time. Click the +/- icon to add or delete rules and conditions. - You can also associate tags to the map rules. For detailed information about map rules, refer to Map Rules. You can configure Inner Header qualifiers and MPLS Header qualifiers for GigaVUE-TA400 device. Refer to Inner Header and MPLS Header Filtering Note: When you associate a tag to a rule in Fabric map, GigaVUE-FM associates the same tag to the internally created flow maps. The visualizations displayed in the Map Rule Statistics dashboard is based on the tags associated with the internal flow maps. Refer to the Add Tags to Map Rules section for more details.
GigaSMART Operation	Select a GigaSMART node and operation.
Destination	Destination port. You can either: - Enter the port alias/port id, or - Specify the destination node and destination ports. Use the +/- icon to add new nodes or remove destination ports. Use the toggle bar to toggle between these two options. Click the Tool Finder option to search for the available destination ports.

Field	Description
	You can also use Fabric Port Group as destination for the second level maps to allow load-balanced traffic. You must enter the name of the FPG created. (For details, refer to Fabric Port Group section). Note: Starting in software version 5.7, you can also select the Null Port option for the second level maps if you do not want to send the traffic to any physical port. Refer to the <i>"Create Application Filtering Intelligence by Selecting Applications from Dashboard"</i> 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.

- 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:

- Select the required fabric map.
- Click the copy icon on the top.

A copy of the selected map is created.

NOTE: In GigaVUE-OS nodes, when using an API script, you may experience a discrepancy in the GigaVUE-FM GET API response for a brief period. This can occur when the script is updating some device components via the GigaVUE-FM API while configuration synchronization is also in progress. In such cases, the configuration synchronization may overwrite the latest data (updated during the GigaVUE-FM API call) with older data (synced during the configuration refresh) in the database. GigaVUE-FM will automatically update to reflect the latest data after subsequent configuration synchronization.

Edit and Delete Fabric Maps

You can edit and delete the flow maps . Refer to the following section for details:

- [Edit Fabric Maps](#)
- [Delete Fabric Maps](#)

Edit Fabric Maps

To edit Fabric Maps:

1. Select a fabric map. Each fabric map has a summary page.
 - The top half of the summary page displays the fabric map component interactions and traffic flow.
 - The bottom half of the screen displays the following tabs:
 - **Summary:** Summary of ports and maps
 - **Rules:** List of rules associated with the fabric map
 - **Statistics:** Fabric Map statistics.
2. Click on any component in the fabric map to display more details about the component.

NOTE: If the *Show Auto Generated* maps option is enabled, then the Fabric Maps page lists the auto generated Fabric Maps that have the prefix as *FmAuto-*, *_FM_*, *_FM-*.


Edit Fabric Map Component

1. Click the **FabricMapcomponent** icon. The fabric map component summary screen appears.
2. Click the option menu and select **edit**.
3. Edit fabric map parameters.
4. Click **Save** to update the fabric map components.

NOTE:

- You cannot edit or delete a map if another person is working on the same map. The error message is available in the audit log page. For information on the audit log page, refer to .
- You can edit Fabric Maps with Fabric Port Group configured as destination in your second level maps. You can replace the FPG with another FPG or replace it with tool ports, as required. Refer to the following screenshot.

- In a cluster, the same source ports are often used across multiple fabric maps. When you need to delete the source ports in one fabric map within the cluster, you must delete the corresponding source ports from all related fabric maps to avoid any overlapping of sources. For example, if a cluster has two source ports—1/1/x1 and 1/1/x2—used across several fabric maps, and you decide to remove source port 1/1/x1 from one fabric map, GigaVUE-FM will also remove source port 1/1/x2 from that fabric map. This automatic removal helps to prevent source overlapping within the cluster.



Delete Fabric Maps

1. Select the fabric map you want to delete from the Fabric Maps summary page.
2. Click the **Trash Can** icon.

A confirmation message appears to confirm whether to delete the selected Fabric Maps.

3. In the message, do any one of the following as per your requirement:
 - Type **DELETE** if you have selected multiple maps, and click **OK**.
 - Type **DELETE ALL** if you have selected all the maps, and click **OK**.

You can view the error message in the audit logs description field for all the failed attempts. For more information on the audit log, refer to [Dashboard](#).

NOTE: You cannot select and delete the auto-generated Fabric Maps, and you cannot select the auto-generated fabric maps using the select-all option.

After the fabric map is deleted you will see an acknowledgment on the screen. Refresh the fabric map summary page to confirm the deletion.

Prioritize Fabric Maps

You can prioritize the traffic flowing through the flow maps. The priority of the fabric map is implemented through cluster level map priorities and can be adjusted at the cluster level.

To prioritize the traffic flow through flow maps:

1. On the left navigation pane, go to **Inventory**  and select **Physical > Fabric Maps**.
2. Select a fabric map and click **View Details** either from the **Task** drop-down menu or from the fabric map selected.
3. Scroll down to view the current priority setting of the flow maps.
4. Click **Change FabricMap Priority**.
5. Use the drag and drop icon to change the order of priority.

Priority SAVE CANCEL

▼ 10.60.80.61

Order	Related Maps	Status	
1	Fabric_map_with_HC1source_x3_TA10tool_x5	● Map is healthy	
2	Fabric_map_with_HC1source_x3_and_TA10tool_x48	● Map is healthy	

 Drag and Drop Icon

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

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

Configure Fabric Map Settings

Share Circuit ID Resource

You can share the circuit id resources for the Fabric Maps created for the same destination. To do this:

1. Select the **Settings** icon on the top navigation bar to open the Share Resources dialog.
2. Enter the **Circuit Id Range** (Min and Max values). Use the +/- icon to increment/decrement the circuit id range.
3. Click **Save** to save the settings.

The Share Resources option is **enabled** by default. Use the toggle bar to disable this option.

When you change the option from shared to non-shared, a note message appears. Click **OK** to proceed and type **Yes** in the confirmation message to change from Shared to Non-Shared. You can use this option only if you who have Read and Write privileges in both the Traffic Control Management and GigaVUE-FM Security Management.

Spine Source Map Mode

You can configure this global setting in FM that controls how cluster-local maps are created in a leaf-and-spine cluster when the source is at the spine (i.e., a decap circuit GigaStream in a leaf-spine cluster). This setting is applicable for creation of the first map that uses the spine mode. Refer to [Supported Topologies](#) topic for more information on topology of Spine Source Map mode.

The following image shows the Spine Source Map Mode in the SETTINGS dialog.

SETTINGS

Share Resources ⓘ

☒
 Share resources across Fabric Maps.

Circuit ID Range

to

Spine Source Map Mode ⓘ

☒ SEPARATE
 ☐ COMBINE

When a leaf-and-spine cluster acts as a middle or destination cluster in a fabric map, its cluster-level map source can be one or more circuit GigaStreams.

If these circuit GigaStreams are on the cluster's spine nodes, you can create the map in the leaf-and-spine cluster in one of two ways:

- **COMBINE** - A single map that combines the circuit GigaStreams from all spine nodes
- **SEPARATE** - Multiple maps, with one map per spine node's circuit GigaStream. This is the default mode.

NOTE: Use SEPARATE mode when the map destination is on nodes that are connected to all spine nodes.

Use COMBINE mode to optimize fabric map deployments so that the cluster can use a single logical map to forward the traffic to the destination, even when traffic arrives via multiple decap circuit GigaStream on the spine nodes.

Limitations of Spine Source Map Mode

- If existing maps are in SEPARATE mode, the future maps from the same circuit GigaStream continue to be SEPARATE even if the global mode is switched to COMBINE.
- If existing maps are in COMBINE mode, the future maps from the same circuit GigaStream to be COMBINE even if the global mode is switched to SEPARATE.
- To truly change mode for the circuit GigaStream, you must delete the existing maps and recreate them after changing the global setting.

Fabric Maps Statistics

GigaVUE-FM provides the ability to view detailed information about flow maps configured between the devices including packets received and transmitted by the following ports:

- Network ports
- Tool ports
- Hybrid ports
- Circuit ports
- Inline-network ports
- Inline-tool ports
- GigaSMART engine ports

Display Fabric Map Statistics

Use the fabric map statistics page to view the statistics associated with the port types. If there are packet drops in the ports, you can use the statistics page to investigate the cause of the packet drops. Fabric map statistical data is generated from the rules you specify when you create a fabric map.

Display Fabric Map Details

To display the fabric map details:

1. On the left navigation pane, go to **Inventory**  and select **Physical > Fabric Maps**.
2. Click on a **Fabric Map** from the main page. The fabric map list view is displayed.
3. Click the **View Statistics** from the options menu. The statistics screen is displayed. Using this screen, you can view fabric map traffic data rates intervals.




Figure 64 Fabric Maps Statistics

NOTE: Statistics and data rates can be displayed per hour, day, week, or month.

- 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.

- To use the filtering functionality, click the **Filter** icon. The Filter quick view dialog is displayed.
- 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 within a cluster.

Refer to the following sections for more details:


- [Fabric Port Group](#)
- [Rules and Notes](#)
- [Create Fabric Port Group](#)
- [Configure Fabric Map Using Fabric Port Group](#)

How Fabric Port Groups Work?

A Fabric Port Group may be defined as a logical collection of tool ports that reside on different clusters. Each of the tool port in the Fabric Port Group receives load-balanced traffic.

Consider a GigaVUE-FM instance with the following clusters configured with tool ports:

- Cluster 1: Configured with Port 1 (1/1/x1), Port 2 (1/1/x2)
- Cluster 2: Configured with Port 3 (1/1/x3), Port 4 (1/1/x4)



FM1: First Level: Source: 1/4/x1 Destination: vPort of Cluster 1
FM2: Second Level: Source: vPort of Cluster 1, Destination: FPG

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

- Creates cluster specific GigaSMART port groups with unique global circuit IDs. These circuit IDs route the traffic tagged by the source vport to the GigaSMART port groups of the destination cluster.
- Creates circuit tunnels with matching circuit ids along the fabric path. Based on the circuit id rules, the maps generated at the cluster level route the traffic to the next hop towards the final cluster level GigaSMART port groups.

NOTE: If any of the ports in the fabric port group goes down, the traffic from the port is redistributed to other healthy ports in the fabric port group.

Rules and Notes

Keep in mind the following rules and notes when you work with the Fabric Port Group feature:


- You can create a FPG using ports of types: Hybrid or Tool. Ports in a single FPG must be of the same port type.

- You can use FPG as destination ports in second-level maps that use load balanced GigaSMART operation in Fabric Maps. You cannot use FPG in the first-level, RegularbyRule, PassAll, or the collector maps.
- You cannot delete or edit an FPG that is deployed in a fabric map. You must first delete the fabric map that deploys the FPG and then delete/edit the FPG.
- You cannot use the individual tool ports of an FPG in another FPG.
- You cannot change the port type and mode of a tool port that is used in the FPG.
- You must not use the tool ports of a FPG in device or cluster-level configurations such as GigaStream, Port Pair, Cluster-level port group (though GigaVUE-FM does not restrict this).
- You cannot use two different load balancing GigaSMART operations for the same FPG.

Create Fabric Port Group

To create a Fabric Port Group:

1. Go to **Traffic**  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

- 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.

- 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

- 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**.

Edit Fabric Port Group

Basic Info

Alias: FPG1

Description (optional): Fabric Port Group with 2 ports

Smart Load Balancing: ☒

Ports * **Port Editor** ☒ Search by Node

10.115.94.53	Select one or more ports...	⊖
T 1/4x4		
30		
10.81.8.20	Select one or more ports...	⊕ ⊖
T 1/1x12		
50		

Tags **Create Tag**

TagKey	Values
	⊕ ⊖

Cancel Save

Configure Fabric Map Using Fabric Port Group

To configure the Fabric Port Group in Fabric Maps perform the following steps:

Task	Description	UI Steps
1	Create a Fabric Port Group with alias FPG1.	Refer to Create Fabric Port Group
2	Create a first level fabric map.	Refer to Create Fabric Maps for the steps. Map Alias: FM_Firstlevel

Task	Description	UI Steps
		Map Type: First Level Source Port(s): Select a single port or multiple source ports from different clusters in a given topology. Destination Port(s): Select virtual port 'vport1' for the destination. Rule: Add the required rules.
3	Create a second level fabric map.	Map Alias: FM_Secondlevel Map Type: Second Level Source Port(s): Select the virtual port 'vport1' (defined as destination in Task 2). Destination Port(s): To select the FPG1 created in Task 1: Enter the first three characters of the FPG created to view the FPG for selection. Rule: Add the required rules.
4	Click Create .	First-level and second-level fabric maps are created, with the FPG1 deployed in the second-level fabric map.

Troubleshooting


You may face situations when fabric fails because of a configuration error or you need to investigate an unhealthy GigaStream or port associated with fabric map. Fabric map maintains the configuration status and health state.

The configuration statuses are SUCCESS, FAILED, and PARTIAL_SUCCESS. You can investigate issues related to flow maps and once issues are corrected, you can run resubmit to push changes. Health state is consolidated from generated components healthStates.

Fabric Maps health states consist of the cluster level maps, (ports), and circuit tunnels (ports and links). After configuration and health state issues are corrected you can sync the Fabric configStatus and healthStates.

To use the troubleshooting feature,

1. On the left navigation pane, go to **Inventory**  under **Physical > Fabric Maps**.
2. Select a **Fabric Map** to troubleshoot.
3. From the **Task** drop-down menu, Select **View Details**.



The Details screen displays fabric map component details.

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




Figure 65 Troubleshooting

- 6. Select the **> arrow** next to the fabric map to display the map details.
Use this screen to investigate configuration issues related to your fabric map.

Fabric Map Configuration Status

The following table helps to troubleshoot a failed or partially successful fabric map configuration.

Configuration Status	Definition	Suggested Action
SUCCESS	Fabric map was accepted at GigaVUE-FM level and traffic flow deployed correctly.	No action required
FAILED	Fabric map was not configured correctly and not accepted or deployed at GigaVUE-FM level.	- Go to the topology page and ensure there is a link between the cluster. - Make sure the circuit GigaStreams are successfully created. - If the clusters are no longer connected, connect the clusters and reapply the configuration. - If the circuit GigaStreams are configured correctly, make the necessary changes and reapply the configuration.
PENDING	- This means no fabric path exists between 2 nodes.or - A failed operation on the device and further deployment of cluster level maps are not allowed in that device.	Create a fabric path and reapply fabric map.
PARTIAL_SUCCESS	This means a portion of the fabric map was accepted at GigaVUE-FM level, but devices cannot provision the generated components.	For details, see How to Troubleshoot Partial Success Errors

Fabric Map Health State

The following table helps to troubleshoot the health state of your fabric map.

Health State	Definition	Suggested Action
Healthy	All fabric map components (nodes, ports, GigaStreams) are all healthy.	No action required.
Unhealthy	An unhealthy state indicates that a port is down or packets dropped.	- Make sure no ports are down and no packets dropped. - Correct any configuration connection issues, and re-sync the fabric map.

How to Troubleshoot Partial Success Errors

If your fabric map displays a red status light on the list view page, this means a portion of the fabric map was accepted at GigaVUE-FM level, but devices cannot provision the generated components.

The error message associated with fabric map appears when you hover over the fabric map status column. You can use the error message to help you troubleshoot and identify the components that are in conflict or mis-configured.

Example 1:

In the following example, the status error message indicates that the ports are mis-configured. You can use the following workflow as a best practice for investigating and correcting map configuration problems.

1. Hover over the fabric map status column displaying a red status light to view the error message. The error message provides information on the possible issues present with fabric map components.




Figure 66 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 67 Fabric Maps Details

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




Figure 68

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.

- Click **Edit** from the options menu to edit the fabric map.

Figure 69 Editing Fabric Maps

- Select the **Source Ports** component.




Figure 70 Source Port Component

- The Source Ports screen appear.
- Edit the source ports with correct information and save your configuration.
- 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:

- On the left navigation pane, go to **Inventory**  under **Physical > Fabric Maps**.
- Click the **Config Audit** link at the top of the screen.



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

If the config audit is not successful, check the fabric map configuration error message(s) for help in resolving issues.

Limitations of Fabric Maps

The following are the limitations of Fabric Maps:

- Circuit ports, circuit GigaStreams and manual topology links (or GDP must be enabled on circuit ports to discover the connected clusters) must be configured before creating fabric map.
- In first-level type fabric map, one vport is required for the destination. Fabric Maps can have multiple tool/hybrid ports from multiple clusters in the destinations in addition to the vport, but the map type remains as the first-level map type.
- In multi-hop scenarios, it is not recommended to create a single level fabric map with GSOP. You must create a vport associated with GSOP, and then create the first and second level maps with the vport to perform the GSOP operations.
- Fabric map configurations are saved in GigaVUE-FM and not in the individual devices. Therefore, if you change the GigaVUE-FM instance that manages the devices, then fabric map configurations cannot be rediscovered from the devices. In this scenario, the GigaVUE-FM configuration that has been saved and backed-up from the original GigaVUE-FM instance must be restored to the new GigaVUE-FM instance to continue.
- To avoid FAILED or PENDING state, or no traffic stats found during deployment of flow maps, you must ensure that the associated devices are connected to GigaVUE-FM and also interconnected for traffic paths.
- Fabric Map or a map with an encapsulation circuit-id cannot be configured from a source port if that source port is used by a passall-map. If you try to configure this from the CLI, the following error message is displayed: **Cannot have encap on a map with a map-passall or port-pair on the source ports.**
- To enable the Circuit ID optimal allocation feature in your currently deployed maps, it is necessary to delete the existing maps. After deletion, create a new map that includes the optimal allocation feature, as it cannot be applied to the existing maps directly.
- A single circuit GigaStream can source and decap traffic from up to 1024 Fabric maps. Multiple decap circuits GigaStreams may exist within a cluster.
- Fabric maps do not support ingress VLAN port tagging. Circuit-ID tunnels used for inter-cluster connections in fabric maps do not support ingress ports with VLAN tags; configuring ingress VLAN tagging on these ports may result in unexpected behavior and is not supported.
- Circuit IDs are allocated on a per-hop and per-direction basis.
- The Circuit ID optimal allocation is not applicable for second-level maps with Fabric Port Group destination.
- The circuit tunnel is not shared for first-level maps even in Shared mode.

- When you create a regular fabric map with GSOP, the GSOP can be placed anywhere along the path from the source to the destination cluster or node, including within the source or destination cluster itself. However, it must not be positioned beyond either the source or destination cluster.
- If existing maps are in SEPARATE mode, the future maps from the same circuit GigaStream continue to be SEPARATE even if the global mode is switched to COMBINE.
- If existing maps are in COMBINE mode, the future maps from the same circuit GigaStream to be COMBINE even if the global mode is switched to SEPARATE.
- To truly change mode for the circuit GigaStream, you must delete the existing maps and recreate them after changing the global setting.

Backup and Restore Fabric Maps and Orchestrated Configurations

You must backup the devices and GigaVUE-FM at the same time. Ensure that the devices are not undergoing configuration edits during backup or restore. For more information, refer to the *“Restore Devices and GigaVUE-FM for Traffic Management Solutions”* section in the *“GigaVUE Administration Guide”*.

Orchestrated Configurations


This chapter provides an overview of the configurations that can be automated from GigaVUE-FM and describes how to use the **Orchestrated Flows> Out-of-Band Flows** option in the GigaVUE-FM GUI to define traffic flow through out-of-band visibility fabric.

Featured Content:

- [About Intent Based Orchestrated Configurations](#)
- [Benefits of Orchestrated Configurations](#)
- [Supported Topologies](#)
- [Orchestrated Configuration: Examples](#)
- [Rules and Notes for Orchestrated Configurations](#)
- [How to Create a Policy](#)
- [Drop Rules](#)
- [Import and Export Orchestrated Policies](#)
- [Egress Filters for Additional Filtering Capabilities](#)
- [Glossary](#)

About Intent Based Orchestrated Configurations

Intent Based Orchestration (IBO) is a new approach that enables you to automate the configuration tasks in GigaVUE-FM, thereby reducing the complexity and the manual labor involved in those tasks. IBO configurations (or 'Orchestrated Configurations') leverage the intelligence of GigaVUE-FM to detect the intention of the user, and accordingly translates the intentions into capabilities for configuring end-to-end flow maps.



With Orchestrated Configurations, you can:

- Create **policies**, which are user-defined prescriptions for what to do with the traffic. The policies are translated internally into device-level maps that have the information required for the flow of traffic.

- Use filters and advanced mapping rules in policies, such as the following:
 - L2-L4 filters
 - Advanced filters such as GTP, ERSPAN, VXLAN, VNTAG, and MPLS
 - **NOTE:** Refer to the GigaVUE-FM User's Guide and GigaVUE-OS CLI Reference Guide for detailed information about the map rules and filters.
- Use packet transformation options along with the filters mentioned above to deliver the traffic to the desired tools. Supported packet transformation options include:
 - De-duplication
 - Masking
 - Slicing
 - Header Stripping
 - Header Addition
 - **NOTE:** Tunneling and flow-based operations (GTP, NetFlow, Application Filtering Intelligence) will be supported in future releases.

Benefits of Orchestrated Configurations

Orchestrated configurations provides the following advantages:

Priority-free Policies

The device-level flow maps created using the standard GigaVUE-FM GUI operations are priority-based maps. As a result some of the traffic may not reach the destination tools due to priority. With orchestrated configuration, you can create priority-free policies. These priority free policies are translated into internal device-level maps that direct the traffic to the configured tools. Refer to [Priority-Free Policies](#) section for detailed information.

Overlapping Sources and Overlapping Rules

Using the standard GigaVUE-FM operations (CLI and GigaVUE-FM), you cannot create maps with overlapping source ports. However, with orchestrated configuration you can create policies with overlapping sources and also policies with overlapping rules. Refer to [Overlapping Sources in Policies](#) for detailed information.

Usability and Ease of Deployment

Without the Orchestration option, if you have to guide the tapped-and-aggregated traffic to the available traffic-monitoring and analysis tools, you must perform the following manual configurations using the standard GigaVUE-FM GUI operations:

- Create and group engine ports.
- Bind engine ports with the required application profiles to create GigaSMART Operations.
- Incorporate GigaSMART operations into flow maps.
- Use hybrid ports, if required.

Using Orchestrated Configurations you can overcome the above restrictions by simply specifying your intention for the traffic and the software intelligently creates the required device-level maps. You need not perform any of the required manual configurations mentioned above.

Orchestrated Configuration: Examples

This section provides examples for orchestrated configurations:


- [Priority-Free Policies](#)
- [Overlapping Sources in Policies](#)

Priority-Free Policies

In this example traffic is tapped and sent through network port sources N1 and N2 to tools T1 and T2. The intention is to send IP traffic of specific source addresses (10.10.10.17 and 10.10.10.125) to tool T1 and VLAN traffic of a specific VLAN value to tool T2.

If you use the standard approach from the GigaVUE-FM GUI to create maps, then you have to configure the following device-level maps:


- Map from source ports N1, N2 with a rule to send traffic with IP source 10.10.10.17 and 10.10.10.125 to tool T1.
- Map from source ports N1, N2 with a rule to send VLAN 101 traffic to tool T2.



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

With Orchestrated configuration, you only have to specify your intention for the traffic in GigaVUE-FM GUI and create a policy. The software intelligently creates the following internal maps:


- Map from Sources N1, N2 to send both IP and VLAN traffic to tools T1 and T2.
- Map from Sources N1, N2 to send IP traffic to tool T1.
- Map from Sources N1, N2 to send VLAN traffic to tool T2.



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

Overlapping Rules in Policies

With priority free policies, internal maps with overlapping rules will be created. You can use the **Ignore rules overlap** option to restrict creating overlapping rules across policies, thereby saving filter resources. However, to ensure that the designated traffic reaches the destination tools in case the traffic matches multiple rules, use a comma separated list of tags to the ignore matching rules.




Overlapping Sources in Policies

In this example, traffic is tapped and sent through network port sources N1, N2, and N3 to tools T1, T2 and T3. The intention is to send:

- IP traffic from source ports N1, N2 to tool T1
- VLAN Traffic from source ports N1, N2 to tool T2
- UDP traffic from source ports N2, N3 to tool T3

You cannot create a map with overlapping sources using the standard GigaVUE-FM operations (CLI and GUI). However, with orchestrated configurations, you can create priority-free policies with overlapping sources as shown in the following figure.



Supported Topologies

Refer to [Supported Topologies](#)

Rules and Notes for Orchestrated Configurations

Refer to the following rules and notes for Orchestrated Configurations.

- The orchestrated configurations co-exist with the configurations created and maintained using the standard GigaVUE-FM interface options. However, you cannot edit the orchestrated configurations using the standard GigaVUE-FM operations and vice-e-versa.
- You can view the orchestrated policies and maps by enabling the *Show Auto Generated* maps option. GigaVUE-FM does not allow you to edit the auto generated maps (maps with prefix FMAuto/_FMAuto).
- When creating a policy, ensure to not include the following special characters: . @ \ " % * ; / , ? in the policy name.
- The priority-free policies are translated into multiple device-level map configurations, which in turn can result in exhaustion of memory resources. However, the software does not allow you to deploy a policy that may result in exhaustion of memory resources. Appropriate warning messages that describe the reason for the policy not getting deployed, will be displayed.
- Intent Based Orchestration feature is integrated with the GigaVUE-FM Tagging and RBAC infrastructure. Refer to the [Tagging and RBAC Support](#) section for detailed information.
- GigaVUE-FM triggers alarms to track the changes in the health status of the policies. Refer to the "Alarms" section in the *GigaVUE Administration Guide* for more details.
- In orchestrated configurations, you can drop the traffic instead of sending the traffic to a specified tool port. Refer to [Drop Rules in IBO](#) section for details related to drop rules.
- You can use hybrid ports in a policy as both source ports and destination ports. You must configure the hybrid ports prior to using them in a policy. Refer to [Work with Hybrid Ports](#) section for details about configuring the hybrid ports.

You can use hybrid ports for policy destination as follows:

- Tool object
- Hybrid port
- GigaStream comprising of hybrid ports

Hybrid port that are part of the destination in a policy can also be used as source in another policy.

- Ports used in a policy will not be released until you delete the policy.
- When you upgrade to software version 5.11.00, the policies created previously will not be retained in the system and will not be listed. You must create the policies again.

- Orchestration is not supported in GigaVUE-HBI device (both as a stand-alone device and also when it exists within a cluster).

Create Orchestrated Policies

Go to **Traffic**  and under **Physical > Orchestrated Flows > Out-of-Band Flows** option in the GigaVUE-FM GUI to create and view policies, go to **Traffic >**. The following tabs are available:

- **Policies:** Displays the list of configured policies. Click on a policy to view the following details:
 - Health Status
 - Deployment Status

You can deploy, undeploy, edit, delete, and create policies from this page.
- **Tools:** Displays the configured tools. You can associate new tools using the **Create New Tool** option. You can also group the tools to create a tool GigaStream. You can apply egress filters in tool ports. Refer to the Egress Filters section for more details.
- **Sources:** Displays the configured ports. You can edit the ports and create port groups from sources. Do not use fan out port as source port in policy.
 - **Filters:** Displays the available filters. You can create a new template from the Filters option. Starting in software version 5.9, you can use advanced filters for creating the filter rules.
 - **Packet Transformations:** Displays the supported packet transformation options. You can also create a template with the required Packet Transformation option.
 - **Visualize:** You can select multiple policies and click **Visualize** to compare the policies.

NOTE: You can create policies by navigating to **Actions > Create Policy** from all of the tabs listed above. You can view the health status of the ports, tools and policies when creating the policy.

Prerequisites

You must have the following licenses:

- GigaSMART Masking
- GigaSMART Packet Slicing
- GigaSMART De-duplication
- GigaSMART Header Addition
- GigaSMART Header Stripping

How to Create a Policy

Refer to the [Rules and Notes for Orchestrated Configurations](#) section before creating the policy.

To create orchestrated policies:

1. Go to **Traffic**  and under **Physical > Orchestrated Flows> Out-of-Band Flows**.
2. Click **Create Policy**. The Create Policy wizard is displayed. You can specify a name for the policy.
3. **NOTE:** You can configure the sources, tools and policies in any sequence you want.
4. Select **Sources** to select the available source ports or port groups.
5. **NOTE:** If a port (port type: network or hybrid) is already used as a source port in any of the following standard GigaVUE-FM configurations, then those ports will not be listed in the drop-down:
MapsFabric MapsApplication Intelligence 5G CUPS solution
However, 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. Within 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 within 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:
- criterial:
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

Explore Traffic Policy Interface

In the GigaVUE-FM Homepage, go to **Traffic > Traffic Policy** to view the Traffic Policy landing page as shown below.



The Traffic Policy landing page includes two main tabs, namely:

- Overview
- Resources

In the **Overview**, you can view the details of the existing maps, their health status, deployment status and you can access the maps to edit, delete, clone, deploy, and undeploy, and so on. In the **Resources**, you can create fabric port groups and can export or import tool templates.

Overview

The Overview includes the Dashboard, Map Details Grid, and Menu including actions that can be performed in the map. The following sections explain these in detail.


- Dashboard
- Map Details Grid
- Menu

Dashboard

The Dashboard provides the overview of the traffic policies and their health status. You can hide or unhide the Dashboard using the Dashboard icon  in the Menu. The Dashboard section includes the Overview section and Health Status section.


- **Overview section**

- Displays the total number of maps, number of maps based on the deployment status such as Success, Failed, Undeployed, and Others (status including Pending, Partial Success, and Damaged).
- Sorts the maps in the Map Details Grid based on your selection of items in the Overview section. For example, if you click the Success item, the maps with Success status are listed in the grid followed by other statuses. This section displays the selected item labels at the top of the Dashboard. The Map Details Grid displays the maps in the sorted order of selected items in the Overview section.



Health Status section

- Represents the overall health of all traffic policies created in the Traffic Policy interface in a color-coded horizontal bar with segments for Healthy (green), Warning (Amber), and Unhealthy (red).
- Displays the statuses with the count, when you hover over the horizontal bar as shown below.
- Sorts the traffic policies in the Map Details Grid based on the health status you select.



The following table lists the health status and its descriptions.



Color	Health Status	Description
Green	Healthy	All components in the Traffic Policy (sources, applications, destinations) are up and passing traffic as expected.
Amber	Warning	One or more components are in a partially available state. Traffic may be impacted, but some traffic can still flow.
Red	Unhealthy	One or more components are down or misconfigured, and traffic is likely not flowing as expected for this policy.


Map Details Grid

The Map Details Grid displays the maps that are created in the Traffic Policy with the tags, deployment status, health status, and so on.

The following table displays the different columns in the Map Details Grid and their description.

Grid Columns	Description
Map Alias	Name of the Traffic Policy.
Tags	Tags linked to the Traffic Policy. To know how to create, work with and add tags to maps, refer to Tags .
Source -> Destination	Displays a visual representation of the map components that are used in the Traffic Policy.  - Application  - Application Ruleset  - By Rule  - Collector  - Destination

Grid Columns	Description
	 - Pass All  - Source
Deployment Status	Displays the deployment status of the Traffic Policy. Valid values are Success, Failed, Undeployed, Pending, Partial Success, and Damaged.
Health Status	Displays the health status of the Traffic Policy. Valid values are Healthy, Warning, and Unhealthy.
Settings	Enables you to add or remove columns from the grid view. 
Select All	Enables you to select the traffic policies in the current page, select the traffic policies in all pages, or unselect all the traffic policies.  When you click the Select All Pages option, all the traffic policies in the interface are selected and Clear option is given as shown below.


Grid Columns	Description
	

The following table lists the possible deployment statuses for the Traffic Policy and suggested actions to be done for each status.

Deployment Status	Definition	Suggested Action
Success	The Traffic Policy was accepted at GigaVUE-FM level and traffic flow was deployed correctly.	No action required. Optionally verify the device-side maps and statistics if you want to confirm counters are incrementing as expected.
Failed	The Traffic Policy was not configured correctly and not accepted or deployed at GigaVUE-FM level.	View Troubleshooting details for the policy to review the error message. Check device reachability and configuration (ports, GigaSMART operations, licenses). Fix the reported issue and deploy the Traffic Policy again.
Pending	The Traffic Policy deployment request is in progress or waiting for device feedback. GigaVUE-FM has started the deployment but the final status from all devices is not yet available.	Wait for the deployment to complete. If the status remains Pending for an extended period, verify device connectivity and job status in GigaVUE-FM, then redeploy the Traffic Policy if needed. To redeploy the Traffic Policy, use Resubmit All option from the Actions menu.
Partial Success	The Traffic Policy was deployed successfully on some devices or	View Troubleshooting details for the policy to review the error message. Review per-device or per-component errors, and

Deployment Status	Definition	Suggested Action
	components, but failed on others—for example, some maps or logical components (vPorts, GigaSMART operation chains) could not be created or synchronized.	correct the failing resources (such as fixing ports, GigaSMART operation resources, or missing maps). Redeploy the Traffic Policy and confirm it moves to Success. To redeploy the Traffic Policy, use Resubmit All option from the Actions menu.
Undeployed	The Traffic Policy configuration exists in GigaVUE-FM but is not currently deployed on any devices, or it has been explicitly undeployed from the devices. No traffic is being steered by this policy.	This status applies to policies that have never been deployed which means the policy configuration is saved in GigaVUE-FM but is not applied on the device. If you want the policy to be active, select the Traffic Policy and choose Deploy from the Actions menu or from the canvas. If the policy should remain inactive, no action is required.
Damaged	GigaVUE-FM detected that the deployed state of the Traffic Policy does not fully match the intended configuration—for example, some underlying maps or components are missing or out of sync during/after deployment or configuration sync.	Use View Troubleshoot and Deployment Status to identify missing or inconsistent components. Resolve the underlying issues (for example, recreate missing maps or fix device problems) and redeploy the Traffic Policy. If the policy remains Damaged and is no longer needed, you can delete it and recreate a clean policy.

The Map Details Grid includes a Three-dots menu that enables you to edit, delete, view map details, view map statistics, and troubleshoot the Traffic Policy if there are any errors. Refer to [Manage Traffic Policy](#) for more information.




Menu

The Menu includes the following functions:

- View Migration Results
- New Map
- Actions
- Filters
- Settings

View Migration Results

Select this to view the migration status of the existing fabric maps to traffic policies.



MIGRATION ALIAS	TYPE	STATUS	TOTAL COUNT	SUCCESS COUNT	FAILURE COUNT	SKIPPED COUNT	START TIME
test-2bcc9f2dc-ff67-48d9-9a80-b03ad8247a71	App-Intel	Completed	1	1	—	—	2026-01-21 1...
test31b9f728d-297c-453f-9fd2-e750b5a331c8	App-Intel	Completed	1	1	—	—	2026-01-21 1...
FabricMap_Aliases-e1c982b5-12e4-40bb-b233-7874990db5f3	FabricMap	Completed	1	—	1	—	2026-01-19 1...
AllFabricMap2f630f83-da9a-4d85-9261-ff623e4d7800	FabricMap	Completed	1	—	1	—	2026-01-13 1...
10.114.82.204-17e0379a-5c51-42aa-bf6b-6112ea457627	FlowMap	Completed	2	1	—	1	2026-01-13 1...
FabricMap_Aliases-1e448c18-8c04-4a08-beba-93ccd420cce3	FabricMap	Completed	1	—	1	—	2026-01-12 1...
10.114.82.204-68723b17-c8c9-47bc-9d15-cb32bfe4a0ac	FlowMap	Completed	2	1	—	1	2026-01-12 1...
10.114.82.204-1ca614d0-322e-4ac4-b1ce-debb64906e42	FlowMap	Completed	2	1	—	1	2026-01-12 1...
10.114.82.204-15c17722-cb01-4b9e-9799-5c2bf904010b	FlowMap	Completed	2	1	—	1	2026-01-12 1...
10.114.82.204-7493838a-1fe4-4a47-b151-2d67c1d3cf2c	FlowMap	Completed	2	1	—	1	2026-01-12 1...
10.114.82.204-8c1df88b-0b75-493d-b887-565f55c418f7	FlowMap	Completed	2	1	—	1	2026-01-12 1...

83 migrations total


Refer to [Migrate Maps to Traffic Policy](#) for more information.

Filters


You can filter the maps based on the Alias, deployment status, health status, and tags to search for a specific map. The

To create a map filter:

1. Click **Filters** in the top right corner of the landing page. The Map Filters page appears.



2. Enter or select the required fields to filter the maps. The landing page displays the maps based on the filter criteria, and the selected filter criteria in the top left corner as shown below.



NOTE: The Alias field supports multiple entries separated by comma, and supports regex-based filtering of traffic policies.

3. Select **Clear** to reset the filter.

NOTE:

If you create a filter and log out of the GigaVUE-FM, then when you login back, the Traffic Policy landing page displays the existing maps based on the filter. If you want to filter the traffic policy of a specific cluster, then enter the cluster ID in the Source or Destination field.

Settings

You can configure the map settings to control how resources, spine source maps, configuration synchronization, and circuit ID ranges are handled for traffic flows. Refer to [Configure Map Settings](#) topic for more information.

Actions

Select the map that you want to edit, delete, clone, deploy, undeploy, and so on and click the Actions dropdown to view the applicable options for the Traffic Policy.

Config Audit

This option enables you to conduct an audit of the map configuration. This process audits the configuration including the components of the Traffic Policy. Config audit verifies that the map configuration is active, the devices are connected, and traffic is flowing through the map with no issues.

To configure audit:

1. Select the map that you want to audit.
2. Go to **Actions > Config Audit**. Once the audit is completed, a success message displays as shown below.

Resubmit All

Resubmit All is a bulk retry option used for all outstanding config operations in the unified Traffic Policy framework. This option is used after you fix any deployment issues (for example, device unreachable, missing circuit/GigaStream, validation errors), to push or retry earlier failed or pending changes to the devices without changing the policy definitions themselves.

Edit


You can edit a map and its components including the configurations of source, source ruleset, application, application ruleset and destination based on the requirement before deploying the map.

You can edit a map in two ways:

1. From the Menu, click **Actions > Edit** or

2. Select the map to edit then click the **Three-dots Menu**  > **Edit**.

NOTE: If you do not deploy the edited version of the map, and come back to the landing page, the Map Details Grid displays an edit icon in the Map Alias column as shown below.



You can always edit either the deployed policy or the policy in draft mode.

Delete

You can delete a map and its components including the configurations of source, source ruleset, application, application ruleset and destination based on the requirement before deploying the map. The delete operation removes the configuration from the device and GigaVUE-FM.

You can delete a map in two ways:

1. From the Menu, click **Actions > Delete** or
2. Select the map to delete then click the **Three-dots Menu**  > **Delete**.
3. When you select multiple maps to delete, a Delete Map dialog appears.
4. Enter Delete in the Type Delete field to delete the selected maps.

Clone

You can clone a map and its components based on the requirement.

You can clone a map in two ways:

1. From the Menu, click **Actions > Clone** or
2. Select the map to delete then click the **Three-dots Menu**  > **Clone**. The Clone Map page appears.
3. Edit and rename the map in the **Alias** field.
4. Click **Save**. The cloned map opens in the map canvas.

Deploy

To deploy a map

1. From the Menu, click **Actions > Deploy** or
2. Select the map to deploy. The selected map appears in the map canvas with the **Deploy** button.
3. Click **Deploy**. The map is deployed successfully.

Undeploy

The undeploy operation removes the configuration only from the device and the configuration exists in GigaVUE-FM.

To undeploy a map


- From the Menu, click **Actions > Undeploy** The map is undeployed successfully.

View Details


You can view the details of the Traffic Policy to view the map components, configuration, status, map type, map priority, and tags associated with the policy.

To view the details:

- Click the map alias or
 - Select the map you want to view the details, click **Three-dots menu**  **> View Details**
- The selected map details appears as shown below.



3. Click **More Info** to view the details of the map components in the policy as shown



below

The Overview tab displays the general map information, such as Alias, description, and timestamps. You can select the other tabs namely Sources, Ruleset, Applications, Application Ruleset, and Destinations to view their details.

NOTE:


In the View mode, you can select **Edit** from the **Actions** dropdown to edit the Traffic Policy if required. When you click the Map Alias to view the map, the deployed version of the map appears.

View Statistics


This option helps you to quickly verify that a Traffic Policy is actively passing traffic and to spot drops, spikes, or anomalies without going to device-level CLIs or separate tools.

To view the statistics information:

1. Select the map and click **Three-dots menu**  > **View Statistics** The Traffic Policy statistics graph appears as shown below.



2. Hover over the statistics graph to view the Data Rate/Packet rate values at any point of time.
3. You can change the X-axis and Y-axis values of statistics graph from the list as required. The statistics graph displays the Data Rate/Packet rate values at the selected time.



NOTE: You can also select the Statistics , to view the statistics information of the Traffic Policy. This is not applicable to maps with deployment status as Undeployed.

View Troubleshoot

Click **Three-dots menu**  > **View Troubleshoot** to view the troubleshooting details of the Traffic Policy. The selected traffic policy appears in the troubleshooting mode. Select the map to troubleshoot. Refer to [Troubleshoot Traffic Policy](#) for more information on troubleshooting steps.

NOTE: You can view this option for all the traffic policies with deployment status other than Undeployed.

Resources Tab

The Resources tab of the Traffic Policy includes Fabric Port Group and Tool Templates sections. You can create and configure fabric port groups, export or import tool templates in the Resources tab. The following sections explain these in detail:

- Fabric Port Groups
- Tool Templates

Fabric Port Groups

A Fabric Port Group (FPG) can be defined as a logical collection of tool ports that reside on different clusters. Each of the tool port in the Fabric Port Group receives load-balanced traffic. Thus, allowing load balancing between tool ports that reside on different clusters. However, you can also use FPG within a cluster.

To know more information on creating, managing and configuring fabric port groups for a Traffic Policy refer to [Create and Manage Fabric Port Group](#).

Tool Templates

Tool templates are pre-built profiles that define a set of applications including the metadata attributes to export to your monitoring tools. You can select the required tool template when configuring a Traffic Policy for an Application Intelligence solution, and reuse it across multiple exporters and devices.

To know more information on creating customized tool templates, refer to [Create and Manage Custom Tool Template](#) section.

Create and Manage Fabric Port Group

Create Fabric Port Group

To create fabric port groups:

1. Click **New** in the top right of the Fabric Port Group section. The New Fabric Port Group page appears.
2. **Alias** - Enter a name (FPG1) for the port group.
3. (Optional) **Description** - Enter a description for the port group.
4. **Smart Load Balancing** option is enabled by default.

NOTE: You cannot disable the Smart Load Balancing option. You can use fabric port group in fabric maps only if this option is enabled.

5. Select the required ports in the Port Editor pane to create the port group.
6. Enter the weights for each port.
7. (Optional) **Tags** - Select the required tag key and tag value combination from the existing tags to associate the FPG to tags.
8. Click **Create Tags** to create a new tag. The Create Tags pane appears.
9. Do the following in the Create Tags pane. Refer to [Tags](#) to know more about tag creation, tag types, and working with tags.
 - a. **Name** - Enter a name for the tag.
 - b. (Optional) **Description** - Enter a description for the tag.
 - c. **Multi Valued** - Turn on the toggle to create a multi-valued tag. If a tag is multi-valued, then the resources can be assigned multiple tag values.
 - d. **Tag Type** - Select the required type from the list. Valid values are Rbac and Aggregation.
 - e. **Hierarchy** - This option is enabled by default.

NOTE: If Hierarchy is enabled, for any node or cluster that is tagged, the ports and slots belonging to the node or cluster will automatically inherit the tags.

- f. **Tag Values** - Enter the tag values separated by comma.
- g. Click **Save**. The created tag appears in the New Fabric Port Group page.

10. Click **Apply**. The created fabric port group appears in the Fabric Port Group section.

Edit Fabric Port Groups

To edit the fabric port groups:

1. Select the port group that you want to edit, click **Actions > Edit**. The Edit Fabric Port Group page appears.
2. Make the required changes.
3. Click **Apply**. The fabric port group is updated successfully.

Delete Fabric Port Groups

To delete the fabric port groups:

1. Select the port group that you want to delete, click **Actions > Delete**. The confirmation dialog appears.
2. Click **Ok**. The fabric port group is deleted successfully.

NOTE: You cannot edit or delete a fabric port group that is already in use by a Traffic Policy.

Export Fabric Port Groups

To export fabric port groups:

1. To export all the port groups, click **Export > Export All > CSV/XLSX**. The fabric port groups are exported in the CSV or XLSX format.
2. To export the selected port groups, select the required port groups and click **Export > Export Selected > CSV/XLSX**. The selected fabric port groups are exported in CSV or XLSX format.

Configure Traffic Policy using Fabric Port Groups

To configure Traffic Policy using fabric port groups, create fabric port groups based on the instructions in the previous section.

To configure a GigaSMART group:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for the cluster where you want to create the virtual port. The **Overview** page of the selected node appears.

3. In the left menu, go to **System > GigaSMART > GigaSMART Groups**. The **GigaSMART Groups** page appears.
4. Click **New**. The **Add GigaSMART Group** page appears.
5. Enter a name (gs-grp) for the GigaSMART group in the **Alias** field.
6. Select a GigaSMART engine port from the **Port List**, example 1/3/e1.
7. Click **Apply**. The GigaSMART Groups page appears with the group you created.

Configure a Virtual Port

To configure a virtual port:

1. In GigaVUE-FM, go to **Inventory > Nodes**.
2. Click the Cluster ID link for which you want to create the virtual port. The **Overview** page of the selected node appears.
3. In the left menu, go to **System > GigaSMART > Virtual Ports**. The **Virtual Ports** page appears.
4. Click **New** to add a virtual port. The **New Virtual Port** page appears.
5. Enter a name (vp1) for the virtual port in **Alias** field.
6. Select a GigaSMART group to which you want to assign the virtual port from the list.
7. Click **Apply**. The **Virtual Ports** page appears with the configured virtual port.

Configure a First-Level Map:

To configure a first-level map:

1. Access Traffic Policy
 - a. In the Traffic Policy landing page, click **New Map**. The New Map pane appears.
 - b. Create a new map with **Alias** as map1.
2. Select Source.
 - a. Click **Sources**. The **Source 1** page appears.
 - b. Select the network ports 1/1/x3 and 1/1/x4 from different clusters.

NOTE: To configure or modify the port type, use the **Quick Port Editor**.

 - c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset.
 - a. Expand the **Source Ruleset** component and then drag and drop **By Rule**. The **Source Ruleset** page appears.

- b. Click **Add Rules**. The **Add Ruleset** pane appears.
 - c. Turn on **Pass** toggle to pass the IPv4 traffic.
 - d. Select IP Version from the **Condition** list, and then select V1 in the **Version** field.
 - e. Click **Save**. The added rule appears in the **Source Ruleset** page.
 - f. Click **Close**. The By Rule block appears on the canvas with the configured rule.
4. Select Destination.
 - a. Click **Destinations**. The **Destination 1** page appears.
 - b. Select the configured virtual port vp1.
 - c. Click **Add**. The Destination block appears on the canvas with the selected virtual port.
5. Deploy Traffic Policy.
 - a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
6. Verify deployment.
 - a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
 - b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure a Second-Level Map:

To configure a second-level map:

1. Access Traffic Policy
 - a. In the Traffic Policy landing page, click **New Map**. The New Map pane appears.
 - b. Create a new map with **Alias** as map2.
2. Select Source
 - a. Click **Sources**. The **Source 1** page appears with the list of available ports.
 - b. Select the virtual port vp1.
 - c. Click **Add**. The Source block appears on the canvas with the selected virtual port.
3. Select Application
 - a. Expand the **Applications** component and then drag and drop **My GSOP**. The **GSOP** page appears.
 - b. Select either a standalone node or a cluster from the **Select Node** list.

- c. Click **New** from the **GSOP** field. The **New GSOP** pane appears with the selected node.
 - d. Enter a name for the GigaSMART operation in the **GSOP Alias** field.
 - e. Select the GigaSMART group (gs-grp) you created earlier.
 - f. Select Load Balancing from the **Applications** list. By default, the application you select is enabled.
 - g. In the Load Balancing section, select Stateful option, and type as GTP
 - h. Select Weighted Least Bandwidth in the **Metric** field.
 - i. Click **Save**. The My GSOP block appears on the canvas with the configured GigaSMART operation.
4. Select Application Ruleset.
 - a. Expand the **Application Ruleset** component and then drag and drop a **By Rule**. The **Application Ruleset** page appears.
 - b. Click **Add Rules**. The **Add Ruleset** pane appears.
 - c. Turn on **Pass** toggle to pass the configured traffic.
 - d. Select IP Version from the **Condition** list, and then Version 4 from the **Version** list.
 - e. Click **Save**. The **Application Ruleset** page appears with the configured rule.
 - f. Click **Close**. The By Rule block appears on the canvas with the configured application rule.
 5. Select Destination.
 - a. Click **Destinations**. The **Destination 1** page appears with the list of available ports.
 - b. Select the fabric port group FPG1.

NOTE: To configure or modify the port type, use the **Quick Port Editor**.

 - c. Click **Add**. The Destination block appears on the canvas with the selected tool ports.
 6. Deploy Traffic Policy.
 - a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
 7. Verify deployment.
 - a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
 - b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure Traffic Policy Without GigaSMART Operations

Use a regular map without GSOP when you want to forward traffic from network ports to tools without any GigaSMART processing. You filter traffic using source rules only.

Use cases

- Send all traffic from SPAN or TAP ports to a monitoring tool.
- Send specific VLANs or IP subnets to a tool without changing packets.

Example


Send all HTTP and HTTPS traffic from a set of SPAN ports to a packet recorder tool.

Applicable Map Components

The following table lists the map components you must configure for this example.

Map Components	Sample Configuration
Source	Network (SPAN) ports - 1/1/x1 – 1/1/x4
Source Ruleset	Match TCP destination ports 80 and 443
Application Ruleset	N/A
Applications	N/A
Destination	Tool port - 2/1/x8

Traffic Policy for a Regular Map (Without GSOP)



Steps:

1. Access Traffic Policy.
 - a. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.

- b. Click **New Map** in the top right corner of the landing page. The New Map dialog appears.
 - c. In the **New Map** dialog, do the following:
 - **Alias** - Enter map1.
 - (Optional) **Description** - Enter the description for the Traffic Policy.
 - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
 - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
2. Select Source.
 - a. Click **Sources**. The Source 1 page appears.
 - b. Select the network ports 1/1/x1 - 1/1/x4.

NOTE: To configure or modify the port type, use the Quick Port Editor.

 - c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select Source Ruleset
 - a. Expand the Source Ruleset component and then drag and drop **By Rule** component. The **Source Ruleset** page appears.
 - b. Click **Add Rules**. The **Add Ruleset** page appears.
 - c. Turn on **Pass** toggle to pass the HTTP traffic.
 - d. Select **Port Destination** from the **Condition** list.
 - e. Enter **80** in the **Min** field.
 - f. Click **+** icon next to the Rule Summary to add an OR rule.
 - g. Turn on **Pass** toggle to pass the HTTPS traffic.
 - h. Select **Port Destination** from the **Condition** list.
 - i. Enter 443 in the **Min** field.
 - j. Click **Save**. The configured rules appear in the **Source Ruleset** page.
 - k. Click **Close**. The By Rule block appears on the canvas with the configured rules.
4. Select Destination.
 - a. Click **Destinations**. The Destination 1 page appears with the list of available ports.
 - b. Select the tool port 2/1/x8.

NOTE: To configure or modify the port type, use the Quick Port Editor.

- c. Click **Add**. The Destination block appears on the canvas with the selected tool port.
5. Deploy Traffic Policy.
 - a. On the Traffic Policy canvas, click **Deploy** to implement the map configuration in the device.
6. Verify deployment.
 - a. **Status:** In the Traffic Policy landing page, confirm that the map status is deployed and healthy.
 - b. (Optional) **Implemented Device:** View the generated map and confirm that the traffic statistics is incrementing as expected.

Configure VLAN Manipulation

VLAN manipulation can be configured based on map rules and maps:

- **Rule-Based**
- **Map-Based**

Refer to the [VLAN Manipulation](#).

NOTE: If you have configured both map level and rule level VLAN manipulation functionality in the same map, then rule-based configuration takes priority.

NOTE: From 6.8 Release, VLAN manipulation add is supported in standalone mode. From 6.10 Release, VLAN manipulation add is supported in Cluster mode.

NOTE: Starting from 6.13 Release, VLAN manipulation Modify, VLAN manipulation Strip is supported only in standalone mode. It is not supported in cluster mode.

Map-Level VLAN Manipulation

The below VLAN Manipulation operations can be configured on map-level:

- **Map-Level Add VLAN Manipulation**
- **Map-Level Modify VLAN Manipulation**
- **Map-Level Strip VLAN Manipulation**

Refer to the [VLAN Manipulation](#).

To configure VLAN manipulation on map-level,

1. Access Traffic Policy.
 - a. In GigaVUE-FM, go to **Traffic > Traffic Policy**.

- b. Select either a standalone node or a cluster from the **Select Node** list.
 - c. Click **New Map** in the top right corner of the landing page. The **New Map** dialog appears.
 - d. In the **New Map** dialog, do the following:
 - **Alias** - Enter map1.
 - (Optional) **Description** - Enter the description for the Traffic Policy.
 - (Optional) **Enabled** - Clear this option, if you want to disable the map. By default, this option is selected.
 - (Optional) **Tags** - Select a tag key and value for the Traffic Policy, to support policy-level analytics.
2. Select Source.
- a. Click **Sources**. The Source 1 page appears.
 - b. Select the network ports 1/1/x1 and 1/1/x2.
- NOTE:** To configure or modify the port type, use the **Quick Port Editor**.
- c. Click **Add**. The Source block appears on the canvas with the selected network port.
3. Select **Source Ruleset**.
- a. Expand the Source Ruleset component and then drag and drop **By Rule/Collector** component. The **Source Ruleset** page appears.
 - b. Click  near the **Source Ruleset**. The **Map Configuration** page appears.
 - c. If you want to add map-level VLAN manipulation configuration, go to the **VLAN** section, and select one of the below options.
 - Select **Add**, from the **VLAN Action** list.
 - Enter the **VLAN ID** value between 1 and 4095.
 - In the **Tag Protocol Id** drop-down list, select the **TPID** value for the VLAN Tag. The default value is 0x8100, but you can also select the other supported values 0x9100 and 0x88a8.
 - Select **Modify**, from the **VLAN Action** list.
 - Enter the **VLAN ID** value between 1 and 4095.
 - In the **Tag Protocol Id** drop-down list, select the **TPID** value for the VLAN Tag. The default value is 0x8100, but you can also select the other supported values 0x9100 and 0x88a8.
 - Select **Strip**, from the **VLAN Action** list.
4. Click **Save** to complete the map level VLAN manipulation configuration.
5. Click **Add Rule(s)**. The **Add Ruleset** page appears.

- d. For example, select **IPv4 Source** from the **Conditions** list.
 - e. Enter the **IPv4 Address** and **Cidr** or **Net Mask** values.
 - f. If you want to add rule level VLAN manipulation configuration, click  **Settings**.
 - In the Rule Configuration dialog, in the VLAN section, select **VLAN Manipulation Add/VLAN Manipulation Modify** from the VLAN Action list.
 - For Add/Modify you need to specify the VLAN ID and TPID.
 - Click **Save**.
 - g. After adding the rule conditions, click **Save** in the Add Ruleset page.
 - h. Click **Close**. The By Rule block appears on the canvas with the configured source ruleset.
6. Select Destination.
 - a. Click **Destination**. The Destination 1 page appears with the list of available ports.
 - b. Select the tool port 2/1/x10.
- NOTE:** To configure or modify the port type, use the **Quick Port Editor**.
- c. Click **Add**. The Destination block appears on the canvas with the selected tool port.
7. Deploy Traffic Policy.
 - a. On the Traffic Policy canvas, click **Deploy** to deploy the Traffic Policy in the device.
 8. Verify deployment.
 - a. Status: In the Traffic Policy landing page, verify that the Traffic Policy shows Deployment Status as Success and Health Status as Healthy.
 - b. (Optional) Implemented Device: View the generated map and confirm that the traffic statistics is incrementing as expected.

Migrate Maps to Traffic Policy

Traffic Policies provide a unified way to configure all physical out-of-band maps (Flow Map, Fabric Map, App Intelligence solutions) through a single User Interface.

You can migrate existing Flow Maps, Fabric Maps, and App Intelligence solutions into Traffic Policies without recreating them manually. After migration, a given configuration exists as a Traffic Policy map. The underlying device configuration is not modified during migration; traffic forwarding continues to work as it is until you change the configuration using Traffic Policies.

Dry Run before Migration

Dry Run lets you validate whether a selected Traffic Intelligence , Application Intelligence or Subscriber Intelligence solutions can be migrated successfully

NOTE: Running a *Dry Run* in Traffic Policy does not migrate or change your existing configuration. It only reports, for each map, whether migration would succeed or fail and the reason for failure.

When you start a dry run for Traffic Intelligence , Application Intelligence or Subscriber Intelligence pages, you can choose:

- Dry Run All- simulate migration and record results for all records.
- Dry Run Selected – simulate migration and record results for selected records.

After Dry Run completes, go to **Traffic**  **>Physical>Traffic Policy > View Migration Results** and review the run:

- Ensure the Status is Completed.
- Check that the Failed and Skipped counts are 0, or that any skipped entries are acceptable.
- If there are failures, use the failure details shown in the Migration Result to correct configuration or data issues, or contact Gigamon Support with the Migration Result information.

We recommend running a Dry Run before performing an actual migration. This helps detect configuration issues in advance and increases confidence that the migration will succeed. This approach helps you avoid migration failures and unexpected errors during the actual migration.

NOTE: After migration, the selected maps are represented as Traffic Policies and can no longer be viewed or edited from the legacy Flow Maps, Fabric Maps, or Application Intelligence pages. If multiple maps share the same alias across clusters (Flow Map vs Flow Map) or between Flow Maps and Fabric Maps, the migrated Traffic Policy alias is adjusted by prefixing *ClusterId_CL_* to make it unique. If there is a collision between a migrated Flow Map, Fabric Map, or Application Intelligence item and an existing

Traffic Policy alias, the migrated Traffic Policy alias is adjusted by appending `_TF_xxx`. If the final alias would exceed the supported length, that map is not migrated and appears under Failed in *View Migration Results*, with the reason shown in the details.

Migrating Traffic Intelligence and Subscriber Intelligence Solutions

Traffic Policy migration affects solution families differently:

- Traffic Intelligence solutions (for example, GigaSMART APF, de-duplication, masking, slicing, tunneling, NetFlow) are built on classic Flow Maps and Fabric Maps. You migrate these by migrating the underlying Flow Maps and Fabric Maps described in this section.
- Subscriber Intelligence solutions (for example, GTP and CUPS correlation, FlowVUE, SIP/RTP correlation, 4G/5G correlation) ,if you have Fabric Maps, that use Subscriber Intelligence applications, you can migrate those Fabric Maps using the standard Fabric Map migration flow.

When you migrate Flow Maps or Fabric Maps to Traffic Policy:

- The underlying device configuration is preserved; only the control plane representation moves to Traffic Policy. Existing Traffic Intelligence behavior continues unchanged.
- After migration, configuration and troubleshooting must be done from the Traffic Policy page; the entries in the legacy Flow Maps and Fabric Maps pages are no longer intended for user edits
- The corresponding entries on the legacy Flow Maps and Fabric Maps pages are treated as internal artifacts and are no longer intended for user edits.

Migrating Flow Maps

You can migrate existing Flow Map into Traffic Policies, which preserves their behavior while moving configuration and troubleshooting to the Traffic Policy page.

NOTE: Before migrating Flow Maps on a device, disconnect the corresponding node or cluster in GigaVUE-FM (**Inventory > Nodes > Actions > Disconnect Node**). GigaVUE-FM still completes any in-progress sync, but new configuration changes are not pushed while the node is disconnected.

1. Go to  > **Physical > flow maps**.
2. Click on Migrate button and either:
 - Migrate All- migrates all records
 - Migrate Selected-migrates only the selected records.
3. Click on OK to migrate your flow maps.

4. To view the progress of the migration use the *View Migration Results* page:
 - Success count matches the number of maps you intended to migrate.
 - Failed count is 0.
 - Skipped entries are acceptable if they correspond to auto-created / internal maps

After migration, reconnect the cluster/standalone so normal configuration sync resumes. The Flow Maps that were migrated will have their configuration represented in Traffic Policies.

NOTE: After you migrate a Flow Map to Traffic Policy, that Flow Map no longer appears in the legacy Flow Maps list and is treated as an internal artifact. You cannot view or edit it from the legacy Flow Maps page; manage the configuration only from the Traffic Policy page

Migrating Fabric Maps

You can migrate existing Fabric Maps into Traffic Policies, which preserves their behavior while moving configuration and troubleshooting to the Traffic Policy page.

5. Go to  > **Physical > Orchestrated Flows > Fabric Maps.**
6. Click on Migrate button and either:
 - Migrate All- migrates all records
 - Migrate Selected-migrates only the selected records.
7. Click on OK to migrate your Fabric Maps.
8. To view the progress of the migration use the *View Migration Results* page:
 - Success count matches the number of maps you intended to migrate.
 - Failed count is 0.
 - Skipped entries are acceptable if they correspond to auto-created / internal maps.

After the Migration your Fabric Map configuration can now only be managed via the Traffic Policy page. The migrated Fabric Map will not be editable from the Fabric Map page.

NOTE: After you migrate a Fabric Map to Traffic Policy, that Fabric Map no longer appears in the legacy Fabric Maps list and is treated as an internal artifact. You cannot view or edit it from the legacy Fabric Maps page; manage the configuration only from the Traffic Policy page

Migrating Application Intelligence Solutions

Migration converts existing Application Intelligence solutions (on the old App Intel page) into unified Traffic Policies, preserving behavior but moving management to the Traffic Policy page.

To Migrate Application Intelligence solutions do the following:

1. Go to **Traffic**  **> Solutions > Application Intelligence.**
2. Select the  Actions option of the session that you wish to migrate.
3. Click on Dry Run if you wish to see how the solution would be migrated before the actual migration.
4. Click on Migrate option to migrate your session to Traffic Policy
5. To view the progress of the migration use the *View Migration Results* page:
 - o Success count matches the number of maps you intended to migrate.
 - o Failed count is 0.
 - o Skipped entries are acceptable if they correspond to auto-created / internal maps
 - o Check that all expected sources, App Intel applications, and destinations are present.

NOTE: After you migrate a Fabric Map to Traffic Policy, that Fabric Map no longer appears in the legacy Fabric Maps list and is treated as an internal artifact. You cannot view or edit it from the legacy Fabric Maps page; manage the configuration only from the Traffic Policy page

Force rollback is not supported for Application Intelligence solutions. If you want to roll back the App Intel solution, you must contact customer support.

View Migration Results

Use the *View Migration Results* page to check the outcome of Dry Run and Migration actions.

1. In GigaVUE-FM, go to **Traffic > Traffic Policy.**
2. Click View Migration Results.
3. In the list, locate the run by its Start Time and Map Type (Flow Map, Fabric Map, or Application Intelligence solution).
4. Verify that:
 - o Success count matches the number of maps you intended to migrate.
 - o Failed count is 0.
 - o Skipped entries are acceptable if they correspond to auto-created / internal maps.

You can use the filters at the top of the page to narrow the view by:

- Map type – Show only Flow Maps, Fabric Maps, or Application Intelligence solutions.
- Status – Show only runs that are In Progress, Completed, Failed, or Skipped.


For each successful migration, confirm that a corresponding Traffic Policy exists and that its sources, applications, and destinations match the original configuration.

Migration converts the configuration stored in GigaVUE-FM for Flow Maps, Fabric Maps, and Application Intelligence solutions into Traffic Policies. Migrated configurations must be managed only from the Traffic Policies page. The corresponding entries in the legacy Flow Map / Fabric Map / Application Intelligence pages are treated as internal artifacts and are no longer intended for user edits.

Configure Map Settings

You can configure the map settings to control how resources, spine source maps, configuration synchronization, and circuit ID ranges are handled for traffic flows.

The following image shows the Maps Settings dialog in Traffic Policy.



The image shows a 'Maps Settings' dialog box with a title bar containing a close button (X), the title 'Maps Settings', and 'Cancel' and 'Save' buttons. The dialog is divided into several sections:

- Shared Resources**: Contains a help icon and a text block explaining that sharing circuit ID space across traffic flows with the same destination can reduce packet drops. Below this is a checked checkbox labeled 'Share Resources across Traffic Flows'.
- Spine Source Map Mode**: Contains a help icon and a text block explaining that in SEPARATE mode, a map is created for each spine node, while in COMBINE mode, a single map aggregates all spine nodes' sources. Below this are two radio buttons: 'SEPARATE' (selected) and 'COMBINE'.
- Configuration Sync**: Contains a help icon and a text block explaining that enabling this option automatically turns device-generated flow maps into traffic flows. Below this is an unchecked checkbox labeled 'Configuration Sync'.
- Circuit ID Range**: A section with a downward arrow icon. It contains two input fields: the first has the value '2' and the second has the value '4000', separated by the word 'to'. To the right of the second input field are two circular buttons with '+' and '-' signs.

Share Circuit ID Resource

You can share the circuit ID resources for the flow maps created for the same destination. To do this:

1. Select the **Settings** icon on the top right corner of the Traffic Policy landing page to open the **Map Settings** dialog.
2. Select **Share Resources Across Traffic Flows** option to share the circuit ID resources.
3. Enter the **Circuit ID Range** (Min and Max values). Use the +/- icon to add or remove the circuit ID ranges.
4. Click **Save** to save the settings.

When you change the option from shared to non-shared, a note message appears. Enter Yes in the confirmation message and click **OK** to change from Shared to Non-Shared. You can use this option only if you who have Read and Write privileges in both the Traffic Control Management and GigaVUE-FM Security Management.

Spine Source Map Mode

You can configure this global setting in FM that controls how cluster-local maps are created in a leaf-and-spine cluster when the source is at the spine (i.e., a decap circuit GigaStream in a leaf-spine cluster). This setting is applicable for creation of the first map that uses the spine mode. Refer to [Supported Topologies](#) topic for more information on topology of Spine Source Map mode.

When a leaf-and-spine cluster acts as a middle or destination cluster in a fabric map, its cluster-level map source can be one or more circuit GigaStreams. If these circuit GigaStreams are on the cluster's spine nodes, you can create the map in the leaf-and-spine cluster in one of two ways:

- **COMBINE** - A single map that combines the circuit GigaStreams from all spine nodes
- **SEPARATE** - Multiple maps, with one map per spine node's circuit GigaStream. This is the default mode.

NOTE: Use SEPARATE mode when the map destination is on nodes that are connected to all spine nodes.

Use COMBINE mode to optimize fabric map deployments so that the cluster can use a single logical map to forward the traffic to the destination, even when traffic arrives via multiple decap circuit GigaStream on the spine nodes.

Limitations of Spine Source Map Mode

- If existing maps are in SEPARATE mode, the future maps from the same circuit GigaStream continue to be SEPARATE even if the global mode is switched to COMBINE.
- If existing maps are in COMBINE mode, the future maps from the same circuit GigaStream to be COMBINE even if the global mode is switched to SEPARATE.
- To truly change mode for the circuit GigaStream, you must delete the existing maps and recreate them after changing the global setting.

Configuration Sync

You can enable this option to automatically turn device-generated flow maps into Traffic Policy.

Troubleshoot Traffic Policy Issues

This section describes how to use the troubleshooting capabilities in GigaVUE-FM to diagnose and resolve issues with Traffic Policy used for Traffic Intelligence, Application Intelligence, and Subscriber Intelligence solutions.

The Troubleshooting view for a deployed Traffic Policy provides:

- A consolidated health view of Traffic Policies and their components (sources, applications, destinations).
- Near real-time (NRT) statistics for supported map types, including:
 - PassAll and ByRule maps
 - Collector maps
 - Inline first-level and second-level maps
 - Maps used in Subscriber Intelligence topologies (e.g., GTP, 5G user/control plane).
- A canvas based visualization of the policy topology, with health indicators for each logical block and links to additional details.

NOTE: Draft or undeployed policies do not show the Troubleshoot option.

Use this view whenever a Traffic Policy does not behave as expected (for example, no traffic, degraded health, or migration related issues).

NOTE: If a Traffic Policy deployment fails and you later see “**No changes have been made**” when you click Deploy without changing the policy, this is expected. After fixing the underlying issue, either wait for the system to retry automatically or use *Dashboard > Actions > Resubmit All* to trigger a new attempt.

View the Troubleshooting Page

To open troubleshooting for a specific **Traffic Policy** :

- In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
- The first view in Troubleshoot is a topology canvas for the selected Traffic Policy:
 - For most Flow Map and Application Intelligence policies, the topology typically spans a single cluster.
 - For Fabric Map-style policies, the topology can span multiple clusters, showing how clusters and maps are connected.
- Locate the required Traffic Policy in the list view.
- Click on **View Troubleshoot** to view the topology in detailed manner.
- Confirm that the map topology, health indicators, and NRT statistics are visible for the selected policy.
- Under the **View Info** section the following details will be listed
 - Source Ports - The Port IDs, Alias, Type and Status.
 - Destination Ports- The Port IDs, Alias, Type and Status.
 - Auto-refresh: when enabled, NRT stats refresh automatically; when disabled, the last snapshot is frozen.
 - Clear Stats: clears existing counters so you can observe fresh statistics.
- You can also use the **View Map Info** page to view a detailed information on all components.

Troubleshooting Statistics and Health Indicators

The Troubleshooting view uses color-coded health states and messages to show the status of the policy and its components.

Policy Level Health

The overall policy health indicates whether any component in the Traffic Policy is unhealthy:

- **Healthy (green):** All monitored components (sources, applications, destinations) are up and passing traffic as expected.
- **Warning (yellow/amber):** A component is in a degraded or partially available state. Traffic may be impacted.

- Critical (red): One or more components are down or misconfigured, and traffic is likely not flowing as expected.

The policy health aggregates underlying component health from sources, destinations, and application elements (for example, GSOP-based GigaSMART chains).

Component Level Health

Each logical block on the Traffic Policy canvas has its own health state:

- Source and Destination components- Shows link state and configuration issues for network ports, tool ports, port groups, GigaStream, vPorts, and inline components. If a component's underlying device ports are down, the component is shown with a critical health state and a corresponding message.
- Application components (GSOP / App chains)- The health of GigaSMART applications (such as Dedup, Masking, APF, AMI/AFI, etc.) is reflected either:
 - On the canvas (for top-level health of the application block), and/or in the Applications tab, which provides additional metadata and health details for GSOP elements.
 - If application engines are down, the health state is marked critical, and the corresponding reason string includes GSOP metadata and vPort.
- Depending on the map type, the view shows:
 - Packet and Byte counters for:
 - PassAll maps
 - ByRule maps (rule-based VLAN/IP match conditions)
 - Collector maps
 - Inline maps (first-level, second-level, and transit maps)
 - Rule-level counters (for example, per pass rule or drop rule).

FAQ Traffic Policy

Why don't I see traffic in the dashboard after creating the Application Intelligence solution?

After you create an Application Intelligence solution, traffic should begin appearing in the dashboard within 5 to 7 minutes. If you do not see any traffic, first verify the creation time of the solution to ensure enough time has passed for traffic ingestion to start.

To view the creation time:

1. In GigaVUE-FM, go to **Traffic > Traffic Policy**. The Traffic Policy landing page appears.
2. Select the required Traffic Policy you have configured for the Application Intelligence solution. The Traffic Policy canvas appears with the map components deployed.

3. Click **More Info**. The Traffic Policy details appear. You can view the creation time here.

The screenshot shows the 'Map-AMI' configuration page with tabs for OVERVIEW, SOURCES, RULESET, APPLICATION RULESET, and DESTINATIONS. The 'General Info' section is expanded, displaying the following details:

Map Alias	Map-AMI
Type	Intent
Traffic Distribute	☒
Tags	—
Description	No Description added for this map...
Created At	2026-01-02 10:51:58
Last Modified At	2026-01-02 11:19:29

If more than 7 minutes have elapsed and the dashboard still shows no traffic, the issue may be related to your environment or configuration. In that case, contact Customer Support for assistance.

What can be migrated?

Traffic Policy migration supports Traffic Intelligence, Subscriber Intelligence (legacy Flow Map-based solutions, user-created Fabric Map-based solutions) and Application Intelligence solutions. Internally generated Fabric Maps that belong to other solutions (for example, FM-auto maps) are treated as solution artifacts and are not migrated.

What solutions cannot be migrated?

Flexible Inline and Intent-Based Orchestration (IBO) solutions cannot be migrated to Traffic Policy.

Where can I see the migration results?

Use **Traffic > Traffic Policy > View Migration Results**. You can also open [View Migration Results](#) from the legacy Flow Maps or Fabric Maps pages where the migration was triggered.

What should I check in the migration results?

Verify that Total and Success counts match your expectations. If there are any Skipped or Failed maps, click the run and then the Skipped or Failed count to see which maps were affected and why.

What happens after migration?

After migration, the configuration is represented as Traffic Policies. Any further edits to those maps must be done from the Traffic Policy page. The corresponding entries on the legacy Flow Maps / Fabric Maps / App Intel pages are treated as internal artifacts and are no longer intended for user edits.

Will migration affect traffic on the nodes?

No. Migration does not change the existing data-plane configuration on the devices; it only changes how the configuration is managed in GigaVUE-FM.

What happens to statistics collected before migration?

Existing traffic trending statistics continue to work as they are. New dashboards and views created for Traffic Policy will show statistics starting from the time of migration.

What should I do if there is a failure in migration?

Open the migration result, click the Failed count, and review the reason shown for each map. After collecting a sysdump and the failure details, contact Gigamon Support for assistance.

Why do I see “Unable to acquire lock for cluster <clustername>” during migration even after disconnecting the node from FM?

This usually means a configuration sync was already in progress when you disconnected the node. The Disconnect option does not interrupt in-progress communication between GigaVUE-FM and the node. Wait 5–10 minutes and try again.

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 VUE Community](#)

Documentation

©This table lists all the guides provided for GigaVUE software and hardware. The first row provides an All-Documents Zip file that contains all the guides in the set for the release.

NOTE: In the online documentation, view [What's New](#) to access quick links to topics for each of the new features in this Release; view [Documentation Downloads](#) to download all PDFs.

Table 1: Documentation Set for Gigamon Products

GigaVUE 6.14 Hardware and Software Guides
DID YOU KNOW? If you keep all PDFs for a release in common folder, you can easily search across the doc set by opening one of the files in Acrobat and choosing Edit > Advanced Search from the menu. This opens an interface that allows you to select a directory and search across all PDFs in a folder.
Hardware how to unpack, assemble, rackmount, connect, and initially configure ports the respective GigaVUE devices; reference information and specifications for the respective GigaVUE devices
GigaVUE-HC1 Hardware Installation Guide
GigaVUE-HC3 Hardware Installation Guide
GigaVUE-HC1-Plus Hardware Installation Guide
GigaVUE-HCT Hardware Installation Guide
GigaVUE-TA25 Hardware Installation Guide
GigaVUE-TA25E Hardware Installation Guide
GigaVUE-TA100 Hardware Installation Guide

GigaVUE 6.14 Hardware and Software Guides
GigaVUE-TA200 Hardware Installation Guide
GigaVUE-TA200E Hardware Installation Guide
GigaVUE-TA400 Hardware Installation Guide
GigaVUE-TA400E Hardware Installation Guide
GigaVUE-OS Installation Guide for DELL S4112F-ON
G-TAP A Series 2 Installation Guide
GigaVUE M Series Hardware Installation Guide
GigaVUE-FM Hardware Appliances Guide
Software Installation and Upgrade Guides
GigaVUE-FM Installation, Migration, and Upgrade Guide
GigaVUE-OS Upgrade Guide
GigaVUE V Series Migration Guide
Fabric Management and Administration Guides
GigaVUE Administration Guide covers both GigaVUE-OS and GigaVUE-FM
GigaVUE Fabric Management Guide how to install, deploy, and operate GigaVUE-FM; how to configure GigaSMART operations; covers both GigaVUE-FM and GigaVUE-OS features
GigaVUE Application Intelligence Solutions Guide
GigaVUE Inline Solutions Guide(NEW) (previously included in the GigaVUE Fabric Management Guide)
Cloud Guides how to configure the GigaVUE Cloud Suite components and set up traffic monitoring sessions for the cloud platforms
GigaVUE V Series Applications Guide
GigaVUE Cloud Suite Deployment Guide - AWS
GigaVUE Cloud Suite Deployment Guide - Azure
GigaVUE Cloud Suite Deployment Guide - OpenStack
GigaVUE Cloud Suite Deployment Guide - Nutanix
GigaVUE Cloud Suite Deployment Guide - VMware (ESXi)
GigaVUE Cloud Suite Deployment Guide - VMware (NSX-T)

GigaVUE 6.14 Hardware and Software Guides

GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration

Universal Cloud TAP - Container Deployment Guide

Gigamon Containerized Broker Deployment Guide

GigaVUE Cloud Suite Deployment Guide - AWS Secret Regions

GigaVUE Cloud Suite Deployment Guide - Azure Secret Regions

Reference Guides

GigaVUE-OS CLI Reference Guide

library of GigaVUE-OS CLI (Command Line Interface) commands used to configure and operate GigaVUE HC Series and GigaVUE TA Series devices

GigaVUE-OS Security Hardening Guide

GigaVUE Firewall and Security Guide

GigaVUE Licensing Guide

GigaVUE-OS Cabling Quick Reference Guide

guidelines for the different types of cables used to connect Gigamon devices

GigaVUE-OS Compatibility and Interoperability Matrix

compatibility information and interoperability requirements for Gigamon devices

GigaVUE-FM REST API Reference in GigaVUE-FM User's Guide

samples uses of the GigaVUE-FM Application Program Interfaces (APIs)

Factory Reset Guidelines for GigaVUE-FM and GigaVUE-OS Devices

Sanitization guidelines for GigaVUE Fabric Management Guide and GigaVUE-OS devices.

Release Notes

GigaVUE-OS, GigaVUE-FM, GigaVUE-VM, G-TAP A Series, and GigaVUE Cloud Suite Release Notes

new features, resolved issues, and known issues in this release ;
important notes regarding installing and upgrading to this release

Note: Release Notes are not included in the online documentation.

Note: Registered Customers can log in to [My Gigamon](#) to download the Software and Release Notes from the Software and Docs page on to [My Gigamon](#). Refer to [How to Download Software and Release Notes from My Gigamon](#).

In-Product Help

GigaVUE-FM Online Help

how to install, deploy, and operate GigaVUE-FM.

How to Download Software and Release Notes from My Gigamon

Registered Customers can download software and corresponding Release Notes documents from the **Software & Release Notes** page on to [My Gigamon](#). Use the My Gigamon Software & Docs page to download:

- Gigamon Software installation and upgrade images,
- Release Notes for Gigamon Software, or
- Older versions of PDFs (pre-v5.7).

To download release-specific software, release notes, or older PDFs:

1. Log in to [My Gigamon](#).
2. Click on the **Software & Release Notes** link.
3. Use the **Product** and **Release** filters to find documentation for the current release. For example, select Product: "GigaVUE-FM" and Release: "5.6," enter "pdf" in the search box, and then click **GO** to view all PDF documentation for GigaVUE-FM 5.6.xx.

NOTE: My Gigamon is available to registered customers only. Newer documentation PDFs, with the exception of release notes, are all available through the publicly available online documentation.

Documentation Feedback

We are continuously improving our documentation to make it more accessible while maintaining accuracy and ease of use. Your feedback helps us to improve. To provide feedback and report issues in our documentation, send an email to:

documentationfeedback@gigamon.com

Please provide the following information in the email to help us identify and resolve the issue. Copy and paste this form into your email, complete it as able, and send. We will respond as soon as possible.

Documentation Feedback Form		
About You	Your Name	
	Your Role	
	Your Company	

For Online Topics	Online doc link	(URL for where the issue is)
	Topic Heading	(if it's a long topic, please provide the heading of the section where the issue is)
For PDF Topics	Document Title	(shown on the cover page or in page header)
	Product Version	(shown on the cover page)
	Document Version	(shown on the cover page)
	Chapter Heading	(shown in footer)
	PDF page #	(shown in footer)
How can we improve?	Describe the issue	Describe the error or issue in the documentation. (If it helps, attach an image to show the issue.)
	How can we improve the content? Be as specific as possible.	
	Any other comments?	

Contact Technical Support

For information about Technical Support: Go to **Settings**  **> Support > Contact Support** in GigaVUE-FM.

You can also refer to <https://www.gigamon.com/support-and-services/contact-support> for Technical Support hours and contact information.

Email Technical Support at support@gigamon.com.

Contact Sales

Use the following information to contact Gigamon channel partner or Gigamon sales representatives.

Telephone: +1.408.831.4025

Sales: inside.sales@gigamon.com

Partners: www.gigamon.com/partners.html

Premium Support

Email Gigamon at inside.sales@gigamon.com for information on purchasing 24x7 Premium Support. Premium Support entitles you to round-the-clock phone support with a dedicated Support Engineer every day of the week.

The VÜE Community

The [VÜE Community](#) is a technical site where Gigamon users, partners, security and network professionals and Gigamon employees come together to share knowledge and expertise, ask questions, build their network and learn about best practices for Gigamon products.

Visit the VÜE Community site to:

- Find knowledge base articles and documentation
- Ask and answer questions and learn best practices from other members.
- Join special-interest groups to have focused collaboration around a technology, use-case, vertical market or beta release
- Take online learning lessons and tutorials to broaden your knowledge of Gigamon products.
- Open support tickets (Customers only)
- Download the latest product updates and documentation (Customers only)

The VÜE Community is a great way to get answers fast, learn from experts and collaborate directly with other members around your areas of interest.

Register today at community.gigamon.com

Questions? Contact our Community team at community@gigamon.com.

Glossary

D

decrypt list

need to decrypt (formerly blacklist)

decryptlist

need to decrypt - CLI Command (formerly blacklist)

drop list

selective forwarding - drop (formerly blacklist)

F

forward list

selective forwarding - forward (formerly whitelist)

L

leader

leader in clustering node relationship (formerly master)

M

member node

follower in clustering node relationship (formerly slave or non-master)

N

no-decrypt list

no need to decrypt (formerly whitelist)

nodecryptlist

no need to decrypt- CLI Command (formerly whitelist)

P

primary source

root timing; transmits sync info to clocks in its network segment (formerly grandmaster)

R

receiver

follower in a bidirectional clock relationship (formerly slave)

S

source

leader in a bidirectional clock relationship (formerly master)