Gigamon®

GigaVUE Traffic Intelligence Guide

GigaVUE-FM,GigaVUE HC Series

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
6.14.01/6.14.02	2.0	09/07/2026	Unified Record type was introduced. For more details refer to Application Metadata Intelligence Capabilities.
6.14	1.0	02/25/2026	The original release of this document with 6.14.00 GA.

Contents

GigaVUE Traffic Intelligence Guide **1**
Change Notes 3
Contents 4
GigaSMART-Traffic-Intelligence **17**
Traffic Intelligence 17
GigaSMART Adaptive Packet Filtering (APF) 18
Feature Overview 18
Implement APF Through the UI 19
Content-based Filtering 20
Encapsulation Awareness 24
Pattern Matching on Gen 2 GigaSMART modules 25
Cross-Packet Pattern Matching 31
Pattern Matching on Gen 3 GigaSMART modules 33
Regex Masking on Gen 3 GigaSMART modules 35
Feature Parity with Gen 2 GigaSMART module and Gen 3 GigaSMART module 37
Map Statistics 38
Display APF Statistics 38
Adaptive Packet Filtering Examples 38
GigaSMART Load Balancing 80
Stateful Load Balancing 80
Examples 85
Stateless Load Balancing 90
Examples 93
Enhanced Load Balancing 99
GigaSMART De-Duplication 101
Feature Overview 102
De-duplication Configuration Steps 103
Configure GigaSMART Parameters for Packet De-duplication 103
Example – GigaSMART De-duplication 104
Display De-duplication Statistics 106
GigaSMART Traffic Performance Enhancement 106
Flow Masking 106
GigaSMART Encapsulated Traffic Performance Enhancement 107
GigaSMART Header Stripping 109
Header Stripping Illustrated 114

Example – Stripping MPLS Headers and Adding a VLAN ID 114
Example – Stripping Layer 3 GRE IP Encapsulated Packets 115
Example – Stripping Layer 2 GRE Ethernet Encapsulated Packets . 116
Display Header Stripping Statistics 116
Example – FM6000 Timestamping 116
Generic Header Stripping 118
Custom 120
PPPoE 123
Cisco LISP 124
L2 MPLS 125
VXLAN 127
TRILL 128
Avaya SPB 129
GigaSMART Header Addition 131
Required License: Header Stripping 131
GigaSMART Masking 133
Examples – GigaSMART Masking 135
GigaSMART NetFlow Generation 136
NetFlow Generation Components 139
Network Ports 140
Map(s) 140
GigaSMART Group 141
NetFlow Records 141
NetFlow Monitors 142
Sampled NetFlow Data 142
NetFlow Exporters 143
IP Interface with Tool Ports 143
Enhancements to NetFlow 143
Display Exporter Statistics 166
Display Monitor Statistics 166
Display IP Interfaces Statistics 167
NetFlow Generation Configuration Modification and Removal 167
Modify NetFlow Generation Monitor Configuration 167
Modify NetFlow Generation Record Configuration 168
Remove NetFlow Generation Configuration 168
V5 Fixed Record Template 169
NetFlow Generation Match/Key and Collect/Non-Key Elements 170
Match/Key Syntax 170
Collect/Non-Key Syntax 174
NetFlow Cacheless export using TCP protocol 182
Rules and Notes 182
Configure Netflow Generation 183
GigaSMART Packet Slicing 191
Example – GigaSMART Packet Slicing 192

Display Slicing Statistics ... 193
GigaSMART Advanced Flow Slicing ... 193
Required License: Advanced Flow Slicing ... 193
Example-GigaSMART Advanced Slicing ... 195
Display Slicing Statistics ... 195
GigaVUE Inline Solutions ... 195
GV-ISSL ... 196
Inline Network and Inline Network Ports ... 197
Physical Bypass Parameter ... 200
Resiliency Features ... 201
Network Port Link Status Propagation Parameter ... 201
Heartbeat Support Between GigaVUE Nodes ... 201
Heartbeats in a Multi-Layer Inline Topology ... 202
Display Current State of Inline Bypass Solution ... 203
IP Interface ... 208
Inline Network Link Aggregation Group (LAG) ... 208
Inline Network Link Aggregation Group— Rules and Notes ... 209
Inline Network Bundle ... 209
Inline Tool Ports and Inline Tools ... 210
Inline Tool Recovery Modes ... 211
Inline Tool Group ... 212
Inline Bypass Restriction ... 214
Forwarding ... 215
Failover ... 215
Out-of-Band and Inline Tools ... 216
Service Chaining of Decrypted Traffic ... 217
Packet Flows ... 217
Modules Matrix ... 219
Inline Single VLAN Tag ... 219
Flexible Inline Maps ... 220
Types of Flexible Inline Maps ... 221
Unprotected Inline Network ... 222
Resilient Inline Arrangement ... 223
Resilient Inline Arrangement— Classic ... 224
Resilient Inline Arrangement With Single VLAN Tag ... 224
Inter-broker Pathway (IB-P) ... 225
Resilient Inline Arrangement—Rules and Notes ... 226
Hardware Security Modules (HSM) ... 226
Why HSM ... 227
How HSM Solution Works ... 227

HSM - Supported Platforms 228
Internet Content Adaptation Protocol (ICAP) 228
How ICAP Solution works 229
ICAP - Supported Platforms 230
ICAP - Rules, Notes, and Limitations 230
ICAP - Supported GigaSMART Engine Ports 231
Gigamon Resiliency for Inline Protection 231
GRIP in Classic Inline Bypass 232
GRIP in Flexible Inline (Redundancy) 232
GRIP Scenarios 232
Traffic Flows Through Node with Primary Role 233
233
Traffic Flows Through Node with Secondary Role after Primary is Lost 233
Both Nodes Go Down and Only Secondary Comes Up 234
Both Nodes Fail; No Traffic Monitoring 234
Both Nodes are in Suspended State 235
How to Handle Recovery 235
How to Cable GigaVUE-FM,GigaVUE HC Series Nodes 236
Redundancy Profile 236
Rules and Notes 236
Limitations 237
Inline TLS/SSL Decryption 237
Why Decrypt TLS/SSL Traffic 237
Inline vs. Passive Decryption 238
How Inline TLS/SSL Decryption Works 238
Mutual TLS (mTLS) with Inline TLS/SSL Decryption 238
Privacy and Sensitive Data Handling 239
What Applications Use TLS/SSL 240
240
Supportability and Compatibility for Inline TLS/SSL Decryption 240
Supported Platforms 240
Mutual TLS (mTLS) Support and Limitations 241
Supported platforms for mTLS 241
Limitations of mTLS functionality 241
GigaSMART Licensing 241
Inline Bypass Requirements 242
Port Requirements 243
GigaSMART Compatibility 243
Supported Ciphers 243
Post-Quantum Cryptography (PQC) Cipher Support 246

Supported Algorithms 246
Supported Key File Types 247
Limitations 247
Attributes of Inline TLS/SSL Decryption Solution 247
TLS/SSL Sessions 248
TLS/SSL Handshake 248
TLS/SSL Handshake Steps: 248
TCP Transition States during TLS/SSL Session 249
TLS/SSL Session Resumption 249
TLS/SSL Session Search 249
StartTLS and HTTP CONNECT 249
Behavior with StartTLS 250
TLS/SSL Keys and Certificates 250
Key Store 250
Trust Store 251
Client Trust Store (for mTLS) 251
Certificate Validation 251
Types of Certificate Validation Errors 252
Client Authentication 252
Client Authentication with mTLS enabled 252
Re-Signed Certificates 253
Certificate Revocation Status 253
Methods to Check Certificate Revocation Status 254
Policy Profile 254
Types of Policy Rules 255
Scale and Flexibility 255
Policy Evaluation 255
Network Phase 255
SNI Phase 256
Certificate Validation 256
Certificate Phase 257
Policy Profile Options 257
Inline TLS/SSL Decryption Port Map 257
Enable or Disable Tool Bypass 258
High Availability Active Standby 258
Inline Network Group Multiple Entry 259
Inline TLS/SSL – Multiple Entry Behavior 260
Tool Early Engage 260
One-Arm Mode 261
Failover Support 262
One-Arm Mode — Rules and Notes 262
Tool Early Inspect 263

Inline TLS/SSL L3 Tool NAT/PAT Support ... 264
HTTP 2.0 Downgrade ... 264
Decryption Port Mapping ... 265
Cache Server Certificate Timeout ... 265
Cache Persistence ... 266
GigaSMART Overload Bypass ... 266
CPU Overload Threshold ... 267
Inline TLS/SSL Monitor Mode ... 268
Inline TLS/SSL Traffic Filtering ... 268
No-decrypt Listing Policy ... 269
Decrypt Listing Policy ... 269
No-Decrypt/Decrypt List Policy — Rules and Notes ... 269
IP Address Subnet with Longest Prefix Match(LPM) ... 270
URL Categorization ... 270
URL Category Look-ups and Caching ... 271
Inline SSL URL categories ... 272
Proxy Server Profile for URL Categorization and Certificate Revocation status ... 278
Inline TLS/SSL Decryption Deployments ... 279
Inbound Deployment ... 279
Inbound Deployment with Client Certificate Authentication (mTLS) ... 279
Outbound Deployment ... 280
Basic Deployments for Inline TLS/SSL Decryption Solution ... 281
Inline TLS/SSL Decryption Solution with Layer 2 Transparent Tools ... 281
Rules and Notes ... 283
Inline TLS/SSL Decryption Solution with Layer 3 Tools (NAT-PAT) ... 284
Architecture of an L3 Tool based Inline TLS/SSL Decryption Solution ... 285
Rules and Notes ... 285
Differences Between Layer 2 and Layer 3 Tools Deployments ... 287
Advanced Features for Inline TLS/SSL Decryption Solution ... 287
Inline TLS/SSL Decryption Solution with RIA ... 287
Symmetric Traffic in RIA ... 289
Asymmetric Traffic in RIA ... 289
VLAN Tagging Behavior for Decrypted Traffic ... 291
VLAN Tagging Behavior for Non-Decrypted Traffic ... 291

Entrust nShield and Thales- Luna HSM for TLS/SSL Decryption for iSSL **292**
HSM - Supported Solutions 293
HSM Configuration Rules, Notes, and Limitations 293
Inline TLS/SSL Decryption Solution with ICAP Client **294**
ICAP - Limitations 295
TLS/SSL Terminology and Acronyms **296**
What is a Flexible Inline Arrangement 297
Flexible Inline Maps **299**
Types of Flexible Inline Maps 300
Supported Platforms, License, and Software Version **301**
Flexible Inline Solution Supported in Clustered Nodes 301
Limitations 302
Flexible Inline TLS/SSL Decryption Solution—Rules and Notes **303**
Configure Software Constructs Using Flexible Inline Arrangement Canvas **304**
Configure Inline Network Ports and an Inline Network **304**
Configure IP Interface **305**
Configure Inline Network LAG **306**
Configure Inline Network Bundle **308**
Configure Inline Tool Ports and Inline Tools **309**
Configure Inline Tool Group **311**
Configure Inline Single VLAN Tag **314**
Configure Flexible Inline Maps **316**
Configure Resilient Inline Arrangement Solution **317**
Create Inter-broker Pathway 317
Configure Resilient Inline Arrangement 319
Configure ICAP Client **321**
ICAP Client—Field References 323
Configure GRIP Solution **324**
Configure Synchronization 324
Example: Gigamon Resiliency for Inline Protection 325
Redundancy Control State 329
How to Use Suspended Role for Maintenance 329
How to Upgrade GigaVUE Nodes in GRIP Deployment 330
Troubleshoot 331

Signaling Ports Down ... 331
Traffic outage in Inline Tool ... 331
Network Traffic Outage ... 331
Configure TLS/SSL Decryption Solutions ... 332
Inline SSL App—Field References ... 332
Configure Inline TLS/SSL Decryption Solution with Layer 2 Tools ... 339
Prerequisites ... 339
Access Flexible Inline Canvas ... 339
Configure Inline Network ... 339
Configure Inline Tool ... 340
Create an Inline SSL APP ... 341
Deploy Inline TLS/SSL Decryption Solution ... 342
Verify the Solution ... 342
What to Do Next ... 342
Configure Inline TLS/SSL Decryption Solution with Layer 3 Tools ... 343
Prerequisites ... 343
Access Flexible Inline Canvas ... 343
Configure Inline Network ... 343
Configure Inline Network Bundle ... 344
Configure Inline Tool ... 344
Create an Inline SSL APP ... 345
Deploy Inline TLS/SSL Decryption Solution ... 346
Verify the Solution ... 346
What to Do Next ... 347
Configure Inbound mTLS for Inline TLS/SSL Decryption ... 347
Prerequisites ... 347
Create a Client Trust store ... 347
Access Flexible Inline Canvas ... 347
Create an Inline SSL APP ... 347
Deploy the solution ... 349
Verify the Solution ... 349
What to Do Next ... 349
Configure Inline TLS/SSL Decryption Solution with RIA ... 350
Prerequisites ... 350
Access Flexible Inline Canvas ... 350
Configure Inline Network ... 350
Configure Inline Tool ... 351
Create Inter-broker Pathway ... 352
Configure Resilient Inline Arrangement ... 352

Create an Inline SSL APP 354
Deploy Inline TLS/SSL Decryption Solution 355
Verify the Solution 356
What to Do Next 356
Configure Thales-Luna HSM for TLS/SSL Decryption 356
Prerequsites 356
Access Flexible Inline Canvas 357
Configure Inline Network 357
Configure Inline Tool 357
Create HSM Group 358
Add Thales -Luna HSM Appliance to the HSM Group 359
Create an Inline SSL APP and Attach the HSM Group 360
Deploy the Inline SSL Solution 362
Register GigaSMART Engine on Thales Luna HSM Server 363
Verify the Solution 363
What to Do Next 364
Modify an HSM Deployment for Inline TLS/SSL Decryption 364
Configure ICAP Client for Inline TLS/SSL Decryption Solution 365
Prerequistes 365
Access Flexible Inline Canvas 365
Configure Inline Network 365
Configure Inline Tool 366
Configure IP Interface 367
Configure ICAP Client 367
Create an Inline SSL APP 368
Deploy the Inline SSL Solution 369
Verify the Solution 370
What to Do Next 371
View Inline Solution Status and Statistics 371
View Inline Solution Status 371
Rules and Notes 372
View the Forwarding States of Inline Networks 373
View Inline TLS/SSL Session Statistics 379
View Inline TLS/SSL Monitor Statistics 380
View Inline TLS/SSL Certificate Statistics 382
View HSM Statistics 383
View ICAP Statistics 384
View Inline TLS/SSL Dashboards 386
Access TLS/SSL Dashboards 386

Basic Dashboards 387
Advanced Dashboards 390
System Requirements to configure Inline TLS/SSL Advanced Dashboards 390
Rules and Notes 391
Configure Advanced dashboard 391
View Advanced Dashboards 392
393
Import and Export Flexible Inline Solution — Rules and Notes 394
Import and Export a Flexible Inline Solution 394
Backup and Restore Flexible Inline Flows 395
Troubleshoot Inline TLS/SSL Decryption Solution Issues 395
Example: Troubleshoot Traffic Issues Between Side A and Side B 396
Inline Bypass Solutions 397
Introduction to Inline Bypass Solutions 398
Capabilities of Inline Bypass Solutions 398
Logical Bypass and Physical Bypass 399
Simple and Complex Inline Bypass Solutions 401
Typical Configuration 402
Distribution to Multiple Inline Tools 402
Inline Tools in a Series 403
Multiple Inline Networks 404
Inline Flow Mapping® 405
Send Traffic to Out-of-Band Tools 406
Types of Inline Networks 406
Unprotected Inline Network 406
Inline Network Groups 407
Inline Tool Sharing 408
Configurable VLAN Tagging 409
Add VLAN Tag 409
Tools in Bridge Mode 409
Inline Tools 410
Inline Tool Failover Action 411
Inline Tool Failover Action with Inline Flow Mapping® 412
Inline Tool Recovery Mode 412
Inline Tool Sharing Mode 414
How to Use SNMP Polling to Obtain Inline Tool State 415
Heartbeats 415
Rules and Notes 416

Heartbeat Profiles ... 416
Standard Heartbeat ... 418
Standard or Custom Heartbeat Packet ... 418
Detect Inline Tool Failure ... 418
Negative Heartbeat Profiles ... 419
Heartbeat Support Between GigaVUE Nodes ... 420
Troubleshoot Heartbeat Failures ... 422
Configure a Heartbeat Profile ... 422
View Heartbeat Profile Statistics ... 423
Heartbeat Status after System Reload ... 424
Inline Tool Groups ... 424
Inline Tool Group Failover Action ... 427
Inline Tool Group Spare Inline Tool ... 428
Symmetrical and Asymmetrical Hashing ... 428
Asymmetrical Hashing Restrictions ... 430
Resilient Weighted Hashing ... 431
Inline Serial Tools ... 431
Inline Serial Tools Global Failover Action ... 434
Inline Tool Series Local Failover Action ... 436
Inline Tool Series Per-Direction Order ... 438
Associate Inline Networks with Inline Tools Using Inline Maps ... 440
Inline Map Passall ... 441
Inline Map ... 441
Inline Map Shared Collector ... 442
Inline Maps to Individual Members of an Inline Tool Group ... 442
Map, Inline Tool, and Inline Tool Group Configuration Restrictions ... 444
Inline Tool Failures and Failover Actions ... 445
Maps That May Lead to Selective Traffic Drops ... 446
Out-of-Band (OOB) Map ... 446
Symmetric and Asymmetric Maps ... 448
Configure Inline Bypass ... 448
Configuration Step Details ... 449
Configure Inline Network Ports ... 449
Configure Inline Network (Unprotected) ... 450
Configure Inline Network Group ... 450
Create Heartbeat Profile ... 451
Create Negative Heartbeat Profile ... 452
Configure Inline Tool Ports ... 453
Configure Inline Tool ... 453
Create Inline Tool Group ... 454

Configure Inline Tool Series 455
Configure When GigaVUE-FM,GigaVUE HC Series® HC Series Modules are Operationally Up 456
Avoid Over subscription 456
Inline Bypass Solution Examples 457
Example 1: Unprotected Inline Bypass with an Inline Tool Group 457
Example 2: Unprotected Inline Bypass with Default Heartbeat 459
Example 3: Protected Inline Bypass Using Combo Modules 461
GigaSMART TLS/SSL Decryption for Out-of-Band Tools 464
GigaSMART Passive TLS/SSL Decryption 464
Set Up Key Store Certificate Management for Passive TLS/SSL 472
Upload TLS/SSL Private Keys 472
Delete TLS/SSL Key 473
Display Key Store Details for Passive TLS/SSL 473
Prerequisites 475
Delete TLS/SSL Service 475
Notes about Private Keys and Passwords 475
Example 1: TLS/SSL Decryption with a Regular Map 478
Example 2: TLS/SSL Decryption with De-duplication 480
Other Usage Examples 480
Entrust nShield and Thales-Luna Network HSM for TLS/SSL Decryption for Out-of-Band Tools(Passive) 480
Configure HSM Group 483
Configure Set Key Handler 485
Configure Passive TLS/SSL Network Access 486
Use RFS to Manage Encrypted Keys 487
Add RFS to GigaVUE-OS Device 488
Map Encrypted Keys with Servers 488
Configure a GigaSMART Group 489
Create a GigaSMART Operation (GSOP) 490
Create TLS/SSL Keychain Password 490
Upload TLS/SSL Private Keys 492
Configure TLS/SSL Service 493
Configure Maps 494
PCAPng Application 496
Tunneling Operations 497
GigaSMART IP Encapsulation/Decapsulation (GigaSMART Tunnel) 497
Required License for IP Decapsulation: Base (GigaVUE-HC3), Tunneling (GigaVUE-HC1)Required License for IP Encapsulation: Advanced Tunneling (GigaVUE-HC3), Tunneling (GigaVUE-HC1) . 497
GigaSMART IP Encapsulation (GigaSMART Tunnel) 509
Required License for IP Encapsulation: Advanced Tunneling (GigaVUE-HC3), Tunneling (GigaVUE-HC1) 509

GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation 510
Required License for L2GRE Decapsulation: Base (GigaVUE-HC3), Tunneling (GigaVUE-HC1)Required License for L2GRE Encapsulation: Advanced Tunneling (GigaVUE-HC3), Tunneling (GigaVUE-HC1) 510
Display L2GRE IPV4/IPv6 Tunnel Statistics 515
Display L2GRE Tunnel Encapsulation Statistics 515
Display L2GRE Tunnel Decapsulation Statistics 516
GigaSMART VXLAN Tunnel Encapsulation 530
Required License for VXLAN Encapsulation: Base (GigaVUE-HC1 Gen 3), Tunneling (GigaVUE-HC3 Gen 3) 530
Display VXLAN IPV4/IPv6 Tunnel Statistics 533
Display VXLAN Tunnel Encapsulation Statistics 533
IP Fragmentation and Reassembly on L2GRE and GMIP Tunnels .. 535
Required Licenses: IP Fragmentation on L2GRE and GMIP Tunnels-Advanced Tunneling license on GigaVUE-HC3. Referred to as “Tunneling license” on GigaVUE-HC1. Reassembly on L2GRE and GMIP Tunnels -Base (GigaVUE-HC3, GigaVUE-HC1-Plus), Tunneling (GigaVUE-HC1). 535
IP Fragmentation on Encapsulation 536
IP Reassembly on Decapsulation 536
Notes and Considerations 536
GigaSMART ERSPAN Tunnel Decapsulation 537
Required License for ERSPAN Decapsulation: Advanced Tunneling (GigaVUE-HC1 Gen3, GigaVUE-HC3, GigaVUE-HC1-Plus, and GigaVUE-HCT Gen 3), Tunneling (GigaVUE-HC1 and GigaVUE-HCT Gen 3). 537
Display ERSPAN Statistics 545
GigaSMART VXLAN Tunnel Decapsulation 546
Required License for VXLAN Decapsulation: Base (GigaVUE-HC3), Tunneling (GigaVUE-HC1) 546
GigaSMART Custom Tunnel Decapsulation 550
Required Licenses for Custom Decapsulation: Base (GigaVUE-HC3), Tunneling (GigaVUE-HC1)and Header Stripping 550
GigaSMART TCP tunnel 554
Required License- Advanced Tunneling license on GigaVUE-HC3. Referred to as “Tunneling license” on GigaVUE-HC1. 554
Secure Tunnels 555
Additional Sources of Information 561
Documentation 561
Documentation Feedback 564
Contact Technical Support 565
Contact Sales 566
The VÜE Community 566

Glossary **567**

GigaSMART-Traffic-Intelligence

Traffic Intelligence

Designed to serve a wide variety of network operations environments, these GigaSMART operations provide complete network visibility to increase the efficiency of network performance tools.

Traffic Intelligence	
Adaptive Packet Filtering	▪ GigaSMART Adaptive Packet Filtering (APF)
Advanced Load Balancing	▪ GigaSMART Load Balancing
De-duplication	▪ GigaSMART De-duplication
Flow Masking	▪ GigaSMART MPLS Traffic Performance Enhancement
Header Stripping	▪ GigaSMART Header Addition ▪ GigaSMART Header Stripping
Masking	▪ GigaSMART Masking
NetFlow Generation	▪ GigaSMART NetFlow Generation
Slicing	▪ GigaSMART Advanced Flow Slicing ▪ GigaSMART Packet Slicing
PCAPng Application	• PCAPng Application
Tunneling	▪ GigaSMART Custom Tunnel Decapsulation ▪ GigaSMART ERSPAN Tunnel Decapsulation ▪ GigaSMART IP Encapsulation (GigaSMART Tunnel) ▪ GigaSMART IP Encapsulation/Decapsulation (GigaSMART Tunnel) ▪ GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation ▪ IP Fragmentation and Reassembly on L2GRE and GMIP Tunnels ▪ GigaSMART VXLAN Tunnel Decapsulation ▪ GigaSMART TCP Tunnel ▪ GRE-In-UDP Tunnel Decapsulation

GigaSMART Adaptive Packet Filtering (APF)

Required License: Adaptive Packet Filtering

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

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

Feature Overview

Adaptive Packet Filtering (APF) provides filtering on specific encapsulation protocol parameters. Additionally, it has the ability to look beyond the encapsulation protocol parameters into the original (encapsulated) data packet, to filter on source and destination IP or Layer 4 port numbers. APF offers the ability to look for content anywhere in the data packet and make intelligent filtering and forwarding decisions.

APF filters packet-by-packet, but does not have the concept of sessions. For Application Session Filtering (ASF) and packet buffering on ASF, refer to Application Session Filtering with Buffering. Adaptive Packet Filtering includes fragmentation awareness whereby all IP fragments associated with the filtered data packet are always forwarded allowing a complete view of the traffic stream for accurate analytics. This is applicable when APF is combined with ASF.

APF operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to Groups of GigaSMART Engine Ports for details.

In APF second level maps, a maximum of twenty-five (25) maps can be attached to a virtual port (vPort). Each map can contain up to 25 GigaSMART rules. The higher map limit applies only when the vPort is configured solely with APF (no other GigaSMART applications or operations are enabled); if any other GigaSMART application or operation is enabled on the vPort, the per vPort limit for APF second level maps remains 5.

NOTE: When you configure a large number of APF second level maps per vPort (for example, approaching the 25-map limit), additional APF processing can impact overall system performance.

Adaptive Packet Filtering (APF) goes deeper into packets to search for a condition, then filter and forward packets to tools, as follows:

- *Content-based Filtering*
- *Encapsulation Awareness*
- *Pattern Matching on Gen 2 GigaSMART modules*
- *Pattern Matching on Gen 3 GigaSMART modules*

Implement APF Through the UI

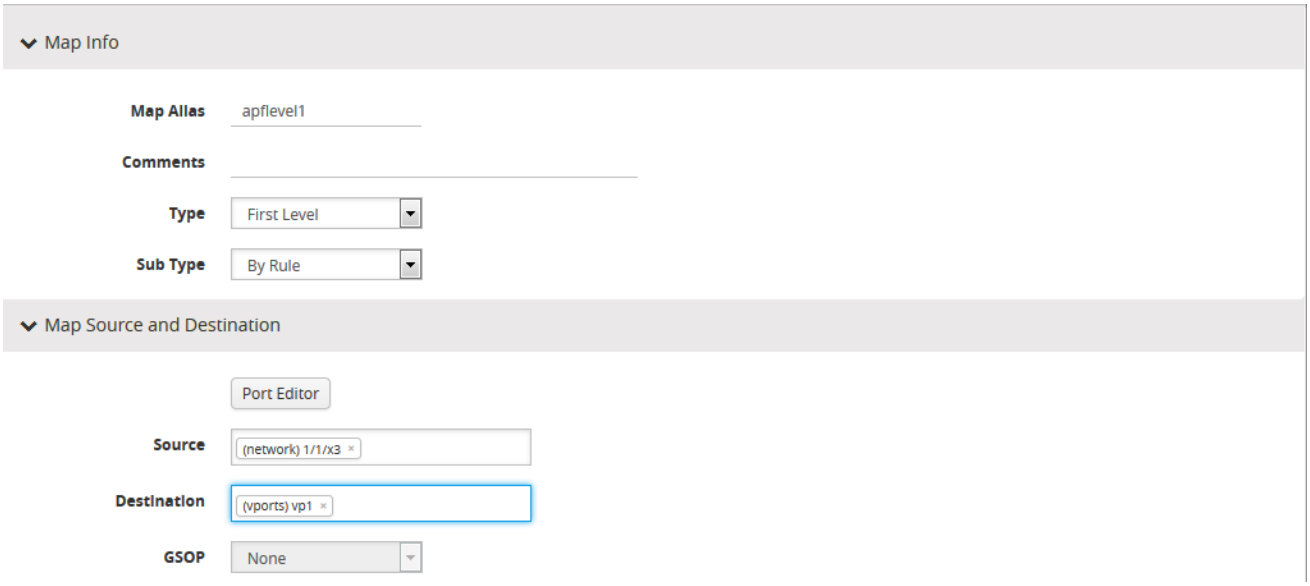
To create vports through the UI and implement APF, do the following:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**.
2. From the left navigation pane, go to **System > GigaSMART >GigaSMART Groups > GigaSMART Groups**, and then click **New**.
3. On the GigaSMART Group page, select an available engine ports in the Port List field to associate group with one of the available engine ports.
   You can associate the GigaSMART Group with one or multiple eports. For APF, no GigaSMART parameters are required unless combined with other gsops.
4. From the device view, select **GigaSMART > Virtual Ports**, and then click **New**.
5. On the Virtual Ports page, enter an alias and select the GigaSMART groups created in Step1, and then click **Save**.
6. To enable the APF operation, do the following:

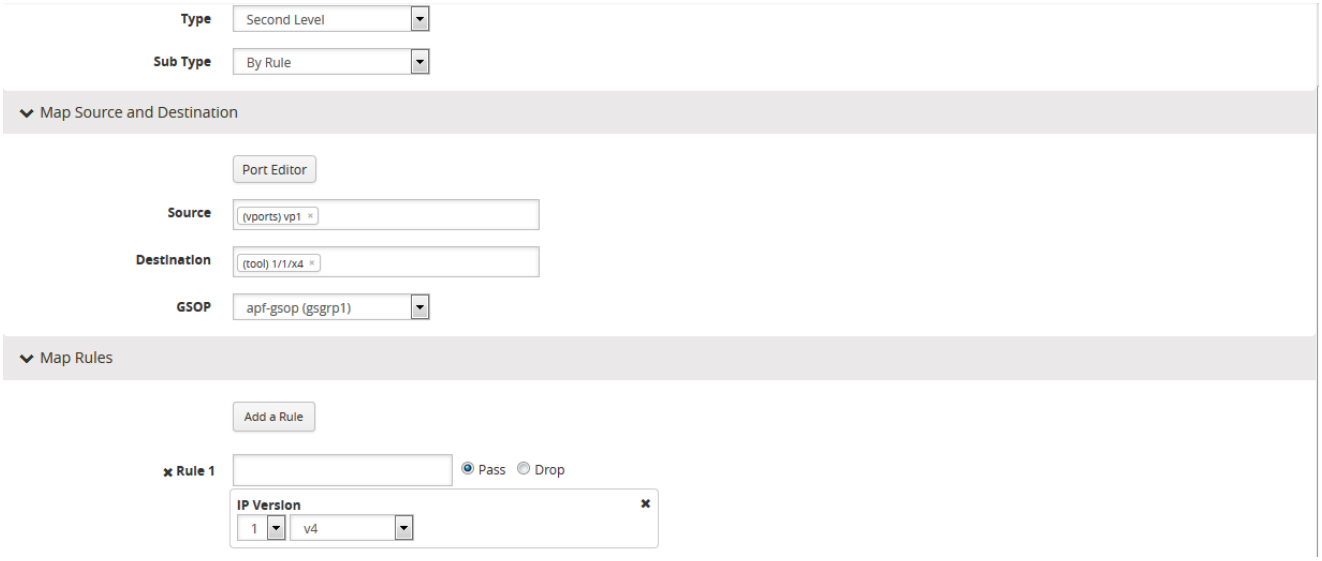
a. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.
b. On the GigaSMART Operations page, enter an alias in the **Alias** field
c. In the **GigaSMART Groups** drop-down list, select the GigaSMART group from step 1.
d. From the **GigaSMART Operations (GSOP)** list, select **Adaptive Packet Filtering** and select **Enabled**.
e. Click **Save**.

Once APF is enabled, maps can be created that the APF and the virtual port.

6. Create the first level map with virtual port created in step 2 as the destination and without applying a GigaSMART Operation.



7. Create a second level map with the APF GigaSMART operation, the virtual port as the source, and a rule. The following figure shows an example.



This completes the process to create an APF GigaSMART operation and corresponding rules. To learn more about the rules applicable for APF, see following sections.

Content-based Filtering

Content-based filtering is based on packet contents beyond Layer 2, 3, and 4 headers. The following four groups of attributes of rules in a map support content-based filtering.

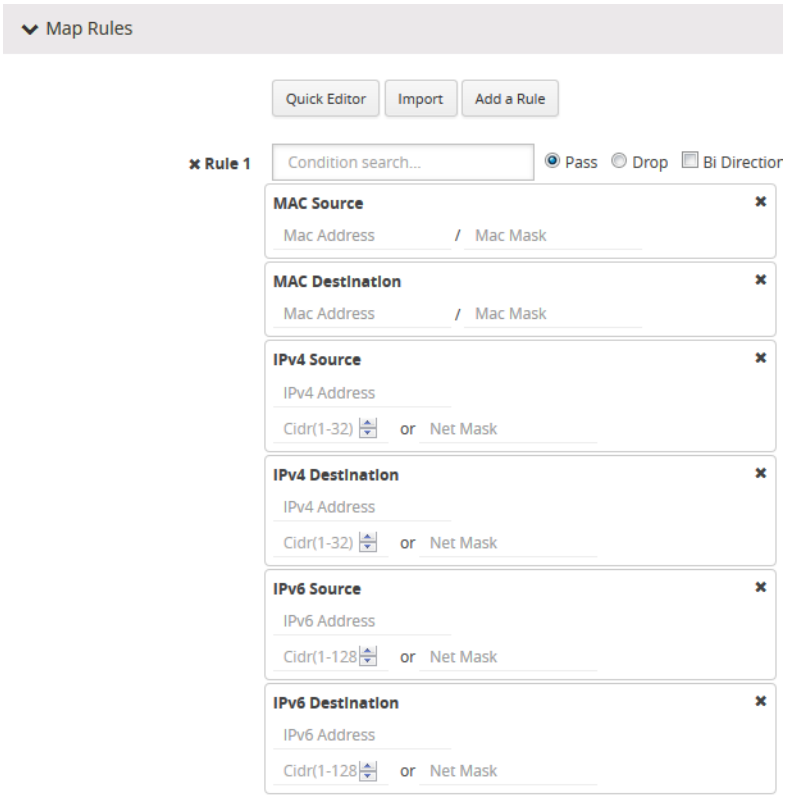
The first group of attributes has the following format:

```
<attribute> <address> <cidr>|<mask>
```

The first group of attributes that use this format are as follows:

- ipv4 src and dst
- ipv6 src and dst
- mac src and dst

The following figure shows the attributes as displayed in the UI.



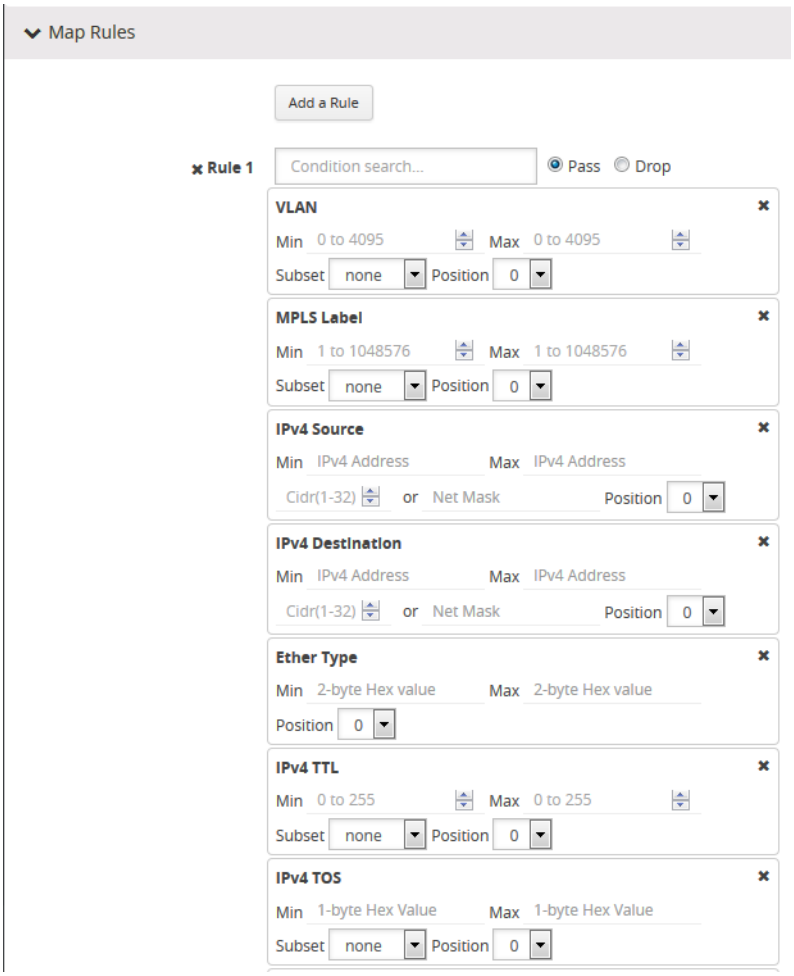
The second group of attributes has the following format:

```
<attribute> min <value> max <value> subset <odd|even|none> pos
```

The second group of attributes that use this format are as follows:

- vlan id
- mpls label
- l4port src and dst
- ethertype
- ipv4 ttl, tosval, and protocol
- ipv6 flow-label
- vntag dvifid, svifid, and viflistid

The following figure shows the attributes as displayed in the UI.



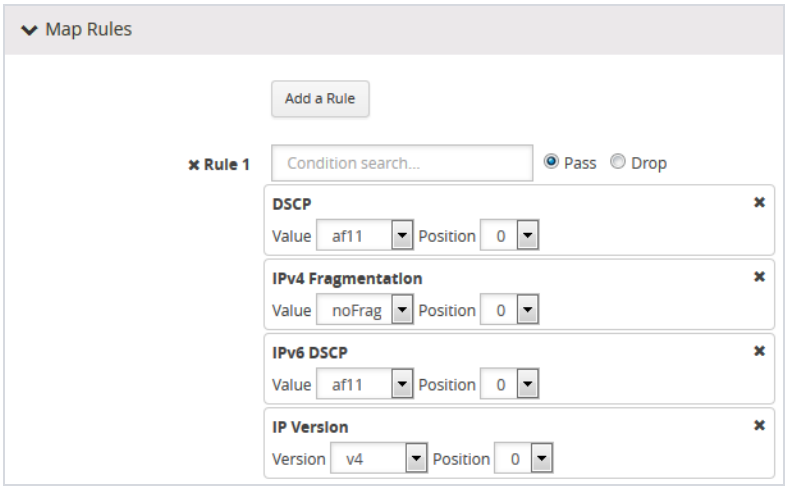
The third group of attributes has the following format: <value> <position>

```
<attribute> value <value> pos <0|1|..|n>
```

The third group of attributes that use this format are as follows:

- ipv4 dscp and frag
- ipv6 dscp
- ipver

The following figure show the attributes as displayed in the UI.



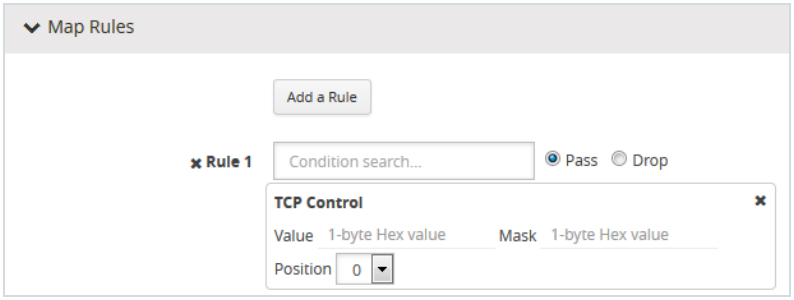
The fourth group of attributes has the following format:

```
<attribute> value <value> mask <mask> pos <0|1|2|3>
```

The fourth group of attribute that uses this format is as follows:

- tcp ctl

The following figure shows the attribute as displayed in the UI.



The maximum occurrences of each attribute supported are as follows:

Attribute	Maximum Occurrences
Attributes in IPv4 header	3
Attributes in IPv6 header	3
Attributes in MAC header	3
VLAN ID	4
MPLS label	4
Attributes in L4port	3
Ethertype	6

Attribute	Maximum Occurrences
Attributes in VNTag header	3
Attributes in TCP header	3
IP ver	3

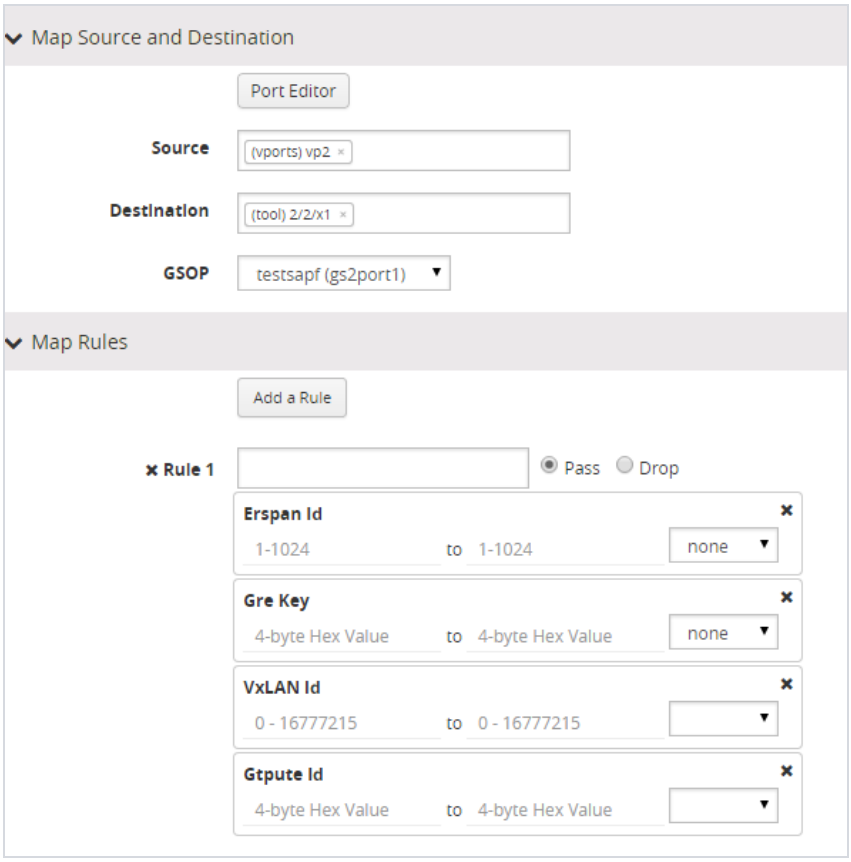
NOTE: In Generation 2 cards, for fragmented packets, only the head fragment gets processed and forwarded based on the filtering criteria while the subsequent fragments are dropped.

Encapsulation Awareness

Encapsulation awareness offers filtering across advanced encapsulation headers, including GTP tunnel ID, VXLAN ID, ERSPAN ID, and GRE key.

The following attributes of rules in a map support encapsulation awareness:

1. Enter a GTP tunnel identifier as a four-byte hex value, either a range or a single value.
2. Enter a VXLAN ID as a three-byte hex value, either a range or a single value.
3. Enter an ERSPAN ID as a decimal value from 1~1024, either a range or a single value using the corresponding arguments.
4. Enter a GRE key as a four-byte hex value, either a range or a single value.



Pattern Matching on Gen 2 GigaSMART modules

Use APF to create pattern matching filters in which the pattern is a particular sequence of data bytes at a variable or fixed offset from the start of a packet. Thus you can filter on any data patterns within a packet.

Pattern matching identifies content based on patterns in any part of the packet, including the payload. Patterns can be a static string at a user configured offset or a subset of Perl Compatible Regular Expression (PCRE) at a variable offset.

The Pattern Match attribute in a map rule supports pattern matching.

Multiple pattern matches are supported. A map can have multiple gsrules, each rule can have a pattern matching expression, and a single packet can match multiple rules.

The Pattern Match attribute in a map rule is shown in Figure 1Use Pattern Match Under Maps for Pattern Matching.

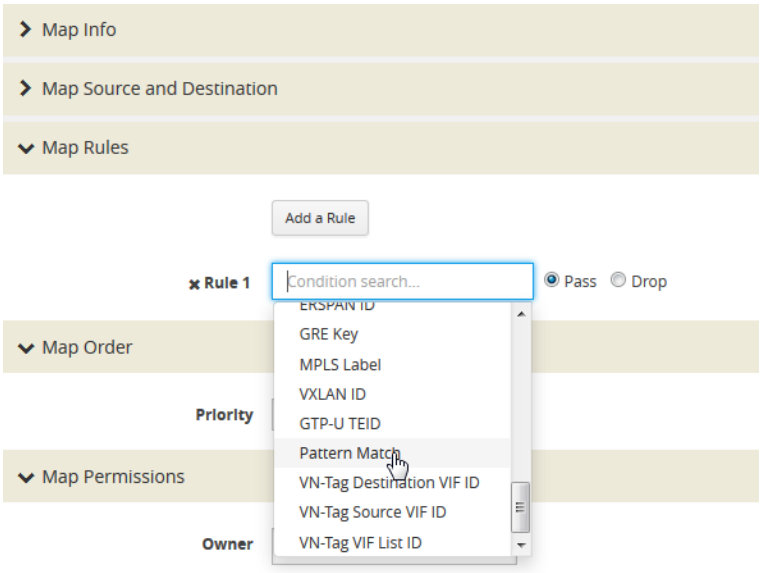


Figure 1 *Use Pattern Match Under Maps for Pattern Matching*

After selecting pattern matching for the rule, you can enter a Perl-compatible regular expression or a string to be used as a filter when pattern matching. For example to pass all packets including the string www.gigamon.com select **string** as type for the pattern match as shown in Figure 2Pattern Match with Type String.

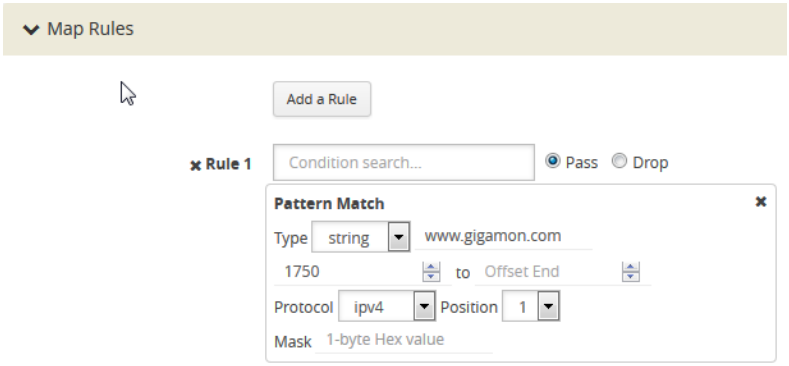


Figure 2 *Pattern Match with Type String*

To pass packets that match any phone number in the nnn-nnn-nnnn format, select **regex** for the pattern match type and enter the following regular expression in the value field: \d{3}-\d{3}-\d{4} as shown in Figure 3Pattern Match with Type RegEx.

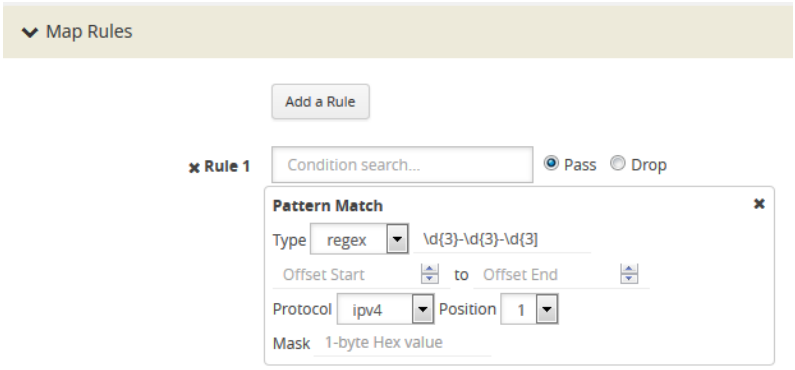


Figure 3 *Pattern Match with Type RegEx*

The offset is a value or range from 0 to 1750. The offset indicates where the pattern under search is located, specify, a value to indicate that the pattern has to start at that offset in the packet in order to be considered a match. Specify a range (beginning and ending) to indicate that the pattern can be anywhere in the packet in that range.

The optional protocol argument of the Pattern Match specifies that the matching will start after the protocol header specified in the command (IPv4, IPv6, TCP, or UDP). Pos 1 or 2 indicates the position. For example, position 2 indicates that matching is to start after the second protocol header. The offset and start and end values are also counted after the protocol header.

For example, to mask an SSL client hello packet pattern starting from the first position after the TCP header with an offset of 0 (located right after the TCP header), you define the pattern match rule as shown in Figure 4Pattern Match for SSL Client Hello Packet.

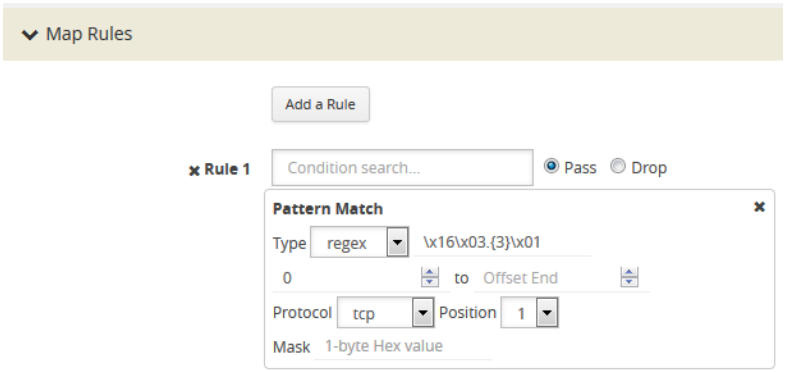


Figure 4 *Pattern Match for SSL Client Hello Packet*

Pattern Matching for IMSI (International Mobile Subscriber Identity)

An IMSI is a unique number usually fifteen digits, identifying a mobile subscriber. The IMSI field in GTP is 8 bytes and is encoded in Telephony Binary Coded Decimal (TBCD) format.

You can create a APF rule to match a specific IMSI. The APF RegEx gsrule needs to match the IMSI in TBCD format instead of the decimal format.

The procedure to convert the IMSI to TBCD hexadecimal format is as follows:

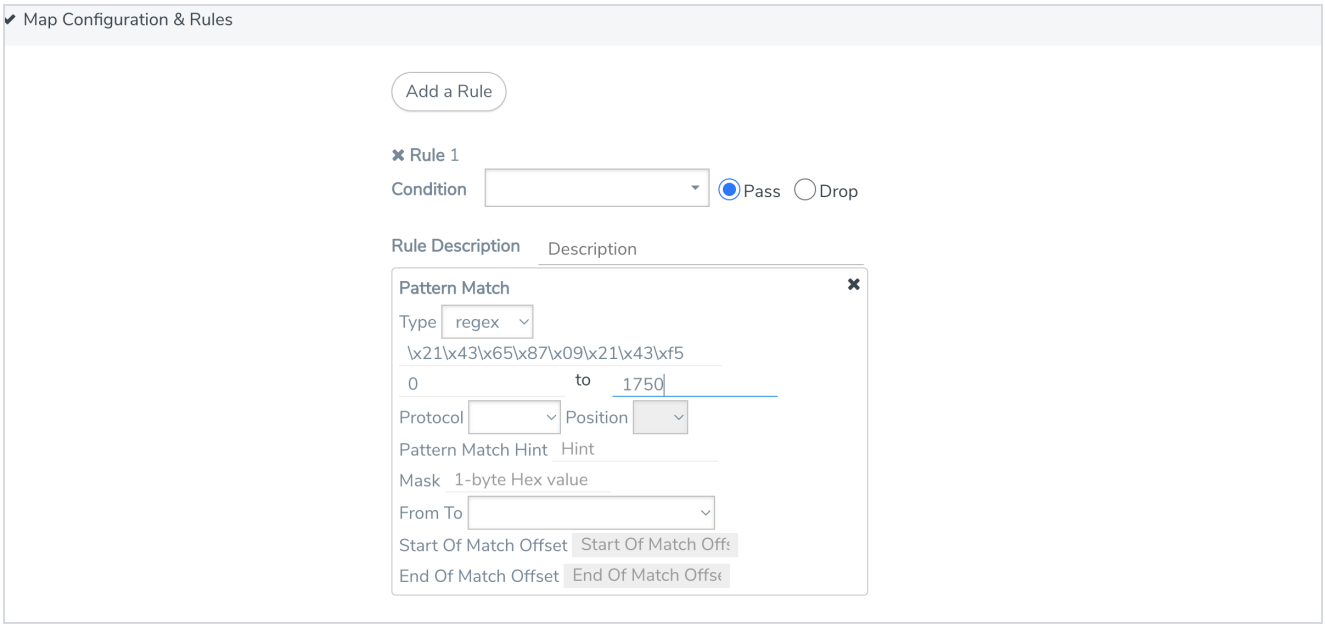
1. If the IMSI length is less than 16 digits, pad the remaining digits with F. For example, 15 digits IMSI 123456789012345 becomes 123456789012345F.

2. Group each pair of digits as 12, 34, 56, 78, 90, 12, 34, 5F.

3. Swap the most significant digit and the least significant digit of each pair.



The end result, 21436587092143F5, is the TBCD equivalent of the IMSI.



Masking with Pattern Matching

APF allows masking when there is a match through pattern matching. Use masking with pattern matching to mask out a specific portion of a packet due to security reasons or to hide sensitive information in packets.

Multiple pattern matches are supported in a map. If there is masking associated with a rule and a packet matches multiple rules, the masking action is enforced for all the matching rules in the map.

The mask specifies that the matched pattern in the gsrule will be masked with the pattern specified in the 1-byte masking pattern.

The pattern specified in the gsrule will be overwritten. The overwritten length is the length of the matching pattern specified in either a string or a RegEx pmatch. Use the 1-byte to overwrite the original pattern match pattern. If there are multiple matches in the packet, up to 10 matches will be masked.

For example, to find Social Security numbers in the format xxx-xx-xxxx, between offset 40 and 80 and replace them with zeros, create a map with a pass rule in a Second Level byRule map with the regular expression \d{3}-?\d{2}-?{4} and a mask with a 1-byte masking value of 0 as shown in Figure 5Map Rule with RegEx for Masking SSNs.

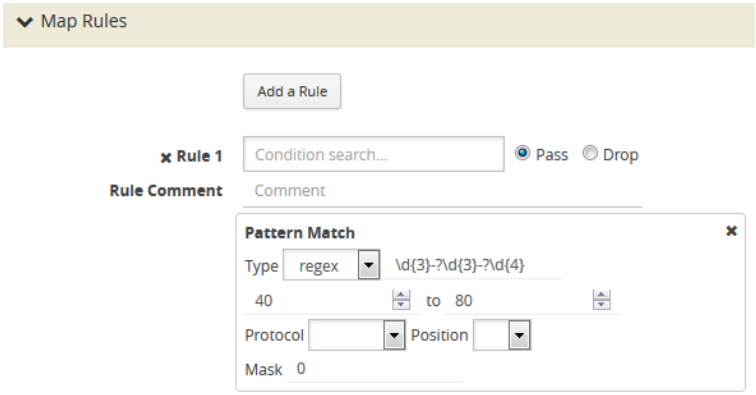


Figure 5 *Map Rule with RegEx for Masking SSNs*

Pattern Matching Hint

To optimize APF pattern matching performance in second level maps with gsrules, you can optionally use a pattern matching hint.

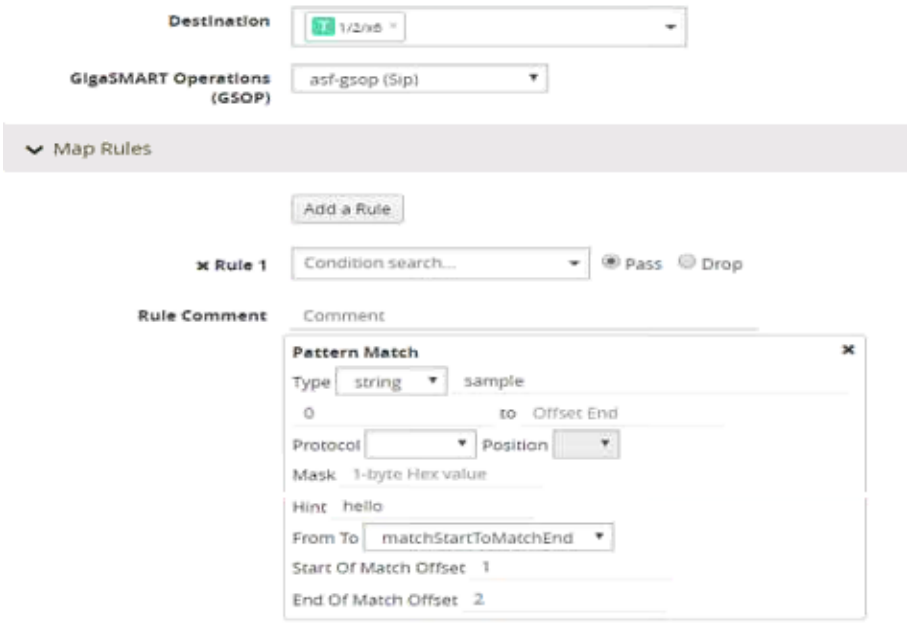


Figure 6 *Pattern Match with Hint*

The addition of the hint leads to two levels of filtering. First, the packet is subjected to a check for the simpler match comprising “gamon|GIM”. If a match is found, a second level check for a match in the complete RegEx, “a[gG]igamon|aGIMO\\s[a-f]\\d{4}”, is performed.

A hint must be selected so that all the packets that are expected to match the actual RegEx must have that string in them, otherwise the first level check will not be cleared. The hint in the example, “gamon|GIM”, was selected because a packet containing either “gamon” or “GIM” in it is a potential match to the actual RegEx.

Best Practices of Pattern Matching Hint

The pattern matching hint is optional and, to optimize performance, it should be specified for all gsrules in a map. In that map, its usage is all or none, meaning you cannot have a mix of gsrules with some having the pattern matching hint and others not. However, if there are two maps, one map can have gsrules that include the pattern matching hint, while the other map can have gsrules that do not.

The use of the pattern matching hint improves performance in complex RegEx patterns involving “lookbehind” and “lookahead” constructs of PCRE syntax. Using them in conjunction with maps with simple patterns, such as fixed length string, is not advisable as it might lead to performance degradation in some cases. Since the RegEx rule set is

limitless, there are no specific rules in which the degradation happens. A best practice is to try out both options, with and without the pattern matching hint, to find out what works best.

The rule of thumb while constructing the pattern matching hint is to keep it as simple as possible. Also, it must be a subset of the configured RegEx pattern. First, try out a 3 to 6 character-wide hint. If that does not provide the necessary scale, you can make the hint wider and more specific to prevent false positives. A maximum length of 63 bytes is supported.

Cross-Packet Pattern Matching

Cross-packet pattern matching refers to a scenario where a match initiates in one packet and ends in a subsequent packet. Staring with Gigamon software release 5.4 this feature enhancement extends the support for GSOP cross packet pattern spanning two packets.

Cross packet matching applies to connection oriented exchanges only and available for 5-tuple flows. Cross packet matching scan will be performed on frames with the following header encapsulations:

- IPv4/TCP, IPv4/UDP
- IPv6/TCP, IPv6/UDP
- IPv4/IPv6/TCP, IPv4/IPV6/UDP
- IPV6/IPv4/TCP, IPV6/IPv4/UDC

Every packet of a flow is subjected to pattern matching scan starting with the inner most L4 payload section. For example, 5-tuple TCP session with nested TCP layer will position the scan starting from start of innermost TCP payload to the end of frame. Bi-directional flow maintains match context for each direction separately and this feature supports up to 1Million flows.

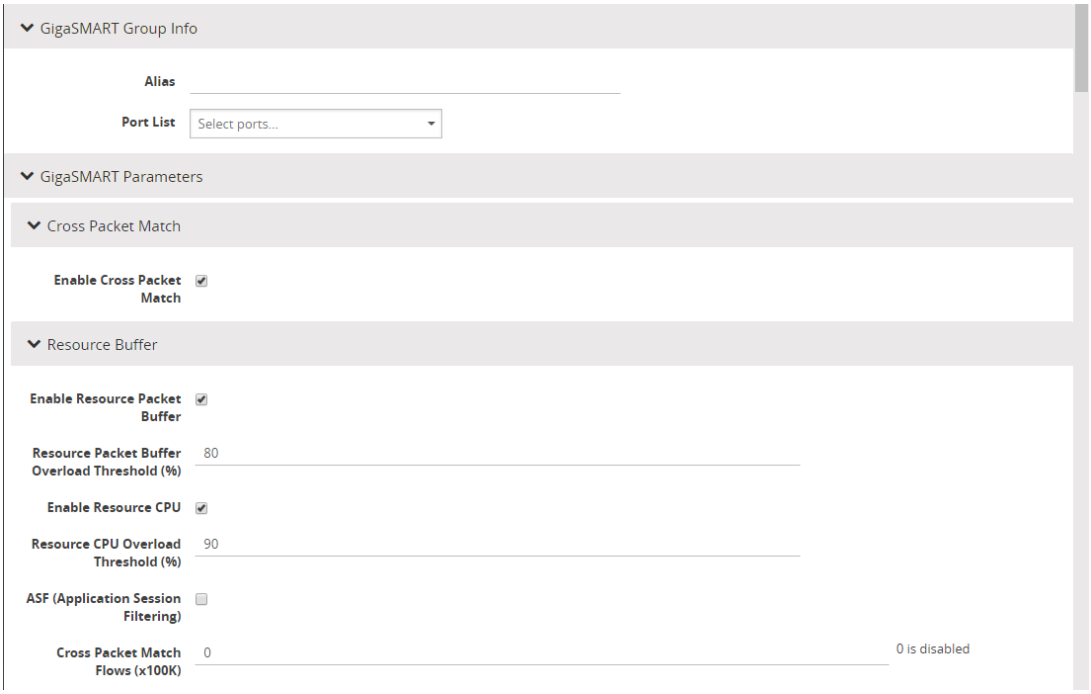
The figure below illustrates the Cross-packet pattern matching concept where the pattern search "**abcdef**" spans two packets.



Enable/Disable Cross-packet Matching

You can enable or disable Cross-packet pattern matching from the GigaSMART GSOP operation.

1. Select a Physical Node.
2. **GigaSMART** > **GigaSMART > GigaSMART Groups**.
3. Click **New**. The GigaSMART Group parameter page displays.



4. Click the **Enable Cross-packet Match** check box to enable.
5. Enter a range from 1 to 10 for the **Cross Packet Match Flows** parameter. Each unit is 100K bi-directional flows.
6. Click **OK**.

NOTE: When disabling this functionality you will be notified that change will be effective only after chassis or GigaSMART card reboot.

Disable Cross Packet Matching

1. Repeat Steps 1 through 3 from the “Enabling Cross Packet Matching” task.
2. Uncheck the **Enable Cross-packet Match** check box to disable this functionality**.**
3. Click **OK**.

View Cross-packet Matching

1. Select a Physical Node.
2. **GigaSMART** > **GigaSMART > GigaSMART Groups**.

3. Select a Group.
4. Click **Edit**. The GigaSMART Group parameters including cross pattern match details pane displays.

Pattern Matching on Gen 3 GigaSMART modules

In Gen 3 GigaSMART module, regular expression (RegEx) pattern matching is introduced to extend Adaptive Packet Filtering (APF) capabilities. This feature enables deeper analysis of packet contents, enhancing security and traffic management functionalities by going beyond basic header filtering.

Gen 3 GigaSMART rules enable advanced traffic filtering with support for up to 25 GigaSMART rules per second level map and up to 25 APF second level maps per vPort. RegEx Pattern Matching can be combined with other filtering attributes in a single rule. Both pass and drop actions are supported with RegEx.

Pattern matching operates within the first 1750 bytes of a packet and begins from protocol-specific locations. For Layer 3 protocols (IPv4 and IPv6) and Layer 4 protocols (TCP and UDP) the pattern matching starts after the last protocol header. For all other protocols it begins at the start of the packet. If Layer 3 and Layer 4 inner headers are present in the packets, the pattern matching starts after the last inner protocol header. The Character limit is 127 for configuring the RegEx pattern. RegEx pattern can be either regular expression or string. Once a RegEx rule is configured, it cannot be modified. You must delete the rule and reconfigure it as a new rule.

It supports ASF for session-based filter and fragment packet forwarding. ASF can buffer up to 20 packets for pattern matching before forwarding traffic to the tools.

Multiple pattern matches are supported. It is recommended to configure multiple pattern in different GS rule within the same map instead of configuring multiple patterns in a single GS rule. To achieve optimal performance, each rule can be configured with a different and unique pattern, especially when multiple filtering rules exist. Each rule can have a pattern matching expression, and a single packet can match multiple rules. For Example, when you configure regex pattern “testN” in one rule and another regex pattern “Next” in another rule, then payload “testNext” can match both rules.

The differences between the Gen 2 and Gen 3 map configurations for Pattern Match RegEx are outlined in the "Gen 2 Map and Gen 3 Map Configuration for Pattern Match RegEx" section in the GigaVUE-OS CLI Reference Guide.

The Pattern Match attribute in a map rule is shown in Figure 1Use Pattern Match Under Maps for Pattern Matching.

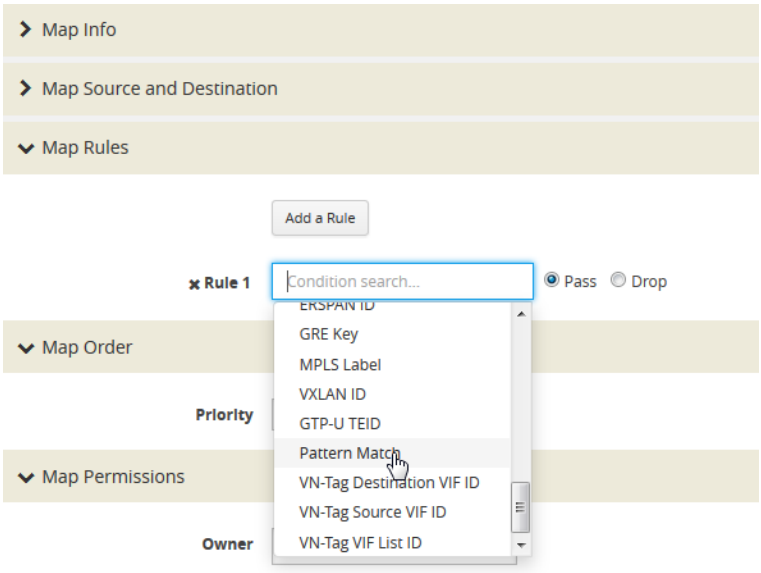


Figure 7 *Use Pattern Match Under Maps for Pattern Matching*

After selecting pattern matching for the rule, you can enter a Perl-compatible regular expression or a string to be used as a filter when pattern matching. For example to pass all packets including the string www.gigamon.com select **regex** as type for the pattern match as shown in Figure 2Pattern Match with Type String.

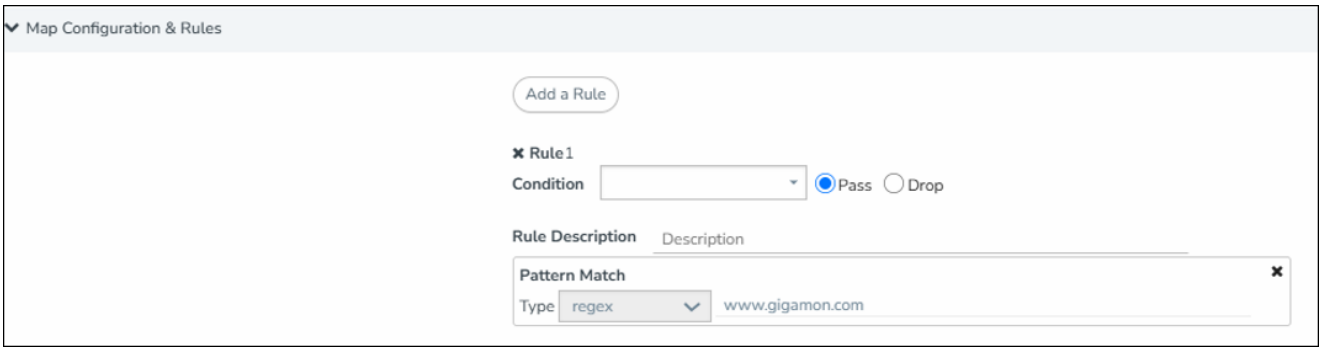


Figure 8 *Pattern Match with Type RegEx*

To pass packets that match any phone number in the nnn-nnn-nnnn format, select **regex** for the pattern match type and enter the following regular expression in the value field: \d{3}-\d{3}-\d{4} as shown in Figure 3Pattern Match with Type RegEx.

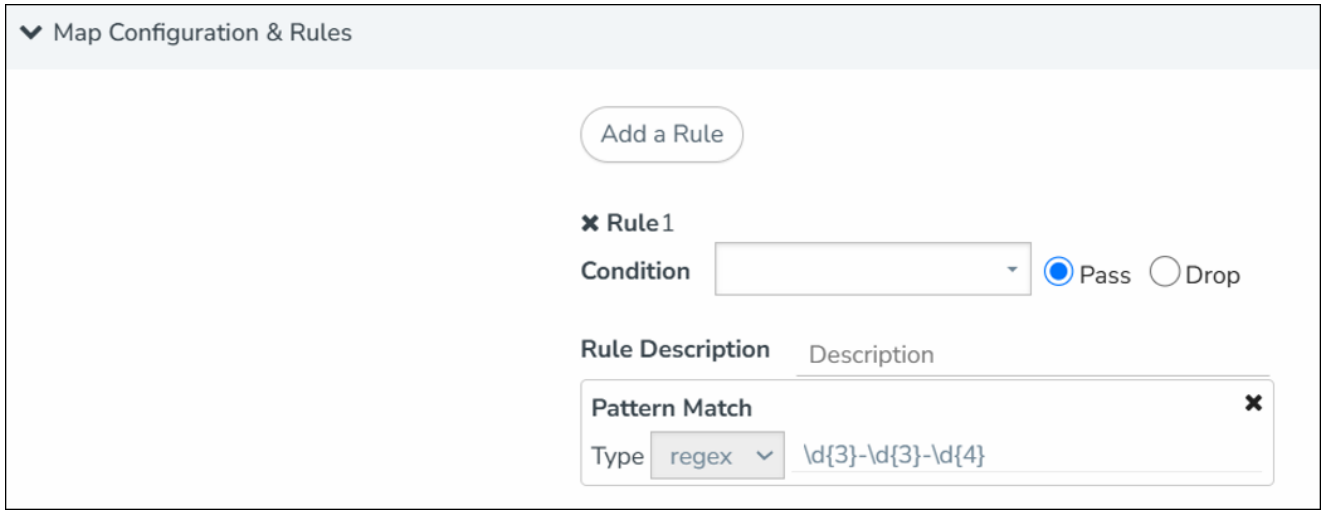


Figure 9 *Pattern Match with Type RegEx*

Regex Masking on Gen 3 GigaSMART modules

On Gen 3 GigaSMART modules, APF supports regex pattern matching with masking to mask sensitive data in the packet payload before traffic is sent to tools. Typical use cases include masking personally identifiable information (PII) such as Social Security numbers, credit card numbers, account IDs, or user IDs to help support privacy and compliance requirements.

Masking is configured on APF GigaSMART rules that use Pattern Match with the RegEx type. When a packet matches the regex, APF overwrites the matched region or a sub-range relative to the start or end of the match , with a 1 byte mask value. This allows you to:

- Mask an entire matched value (for example, a full SSN or credit card number).
- Mask only the value part of a key-value pair (for example, the bytes following user_id=) while leaving the key visible.

When ASF is configured together with APF regex-based masking, the system must mask only the first matching packet in a session. All subsequent packets in the same session must be sent to the tool without masking (if applicable).

Configuring a Regex Masking on Gen 3 GigaSMART modules

To configure Regex Masking on Gen3 GigaSMART modules :

1. From the left navigation pane, go to **Inventory > Physical > Nodes**.
2. Under **Traffic > Maps > Maps.**
3. Create or (edit) a second-level map that uses an APF GSOP on a Gen 3 GigaSMART module. Refer to Implement APF Through the UI
4. Under the **Map Configuration & Rules**

a. Set Action to **Pass**.

b. Set Condition to **Pattern Match**.

c. Set Type to **RegEx**.

5. In the Pattern Match area, configure:

   a. Value: a regex that matches the sensitive data.

   b. Range: an appropriate offset or begin to end range where the pattern can appear (for example, 0..1750).

6. In the **Masking** section (visible only for RegEx on Gen 3 nodes), fill in:

   a. Mask byte: A 1-byte mask value (for example, 0x2a to mask with *).

   b. From – Choose where masking should start relative to the regex match and set an offset:

   - Start of match: offset 0–1749
   - End of match: offset 0–1748

   c. To – Choose how far masking should extend and set the allowed range:

   - End of match: offset –1749 to +1749
   - End of packet: masks from the From position to the end of the packet.
   - Length:
     - 1–1750 bytes when From is Start of match
     - 1–1749 bytes when From is End of match.

The masking range is defined relative to the regex match—for example, from the start to the end of the match, or from the end of the match for a fixed length, so that you can either mask the entire matched value or only the value portion of a key–value pair (for example, the bytes after user_id=) while leaving the key visible.

> **NOTE:** Masking is applied only for pass rules. Drop rules can still use regex to decide whether to drop a packet, but they do not perform masking.

On Gen 3, 25 APF second-level maps per vPort are supported only when all of those maps use APF (with or without regex masking). If any other GigaSMART application is configured on that vPort (alone or in combination with APF), a maximum of five second-level maps is supported for that vPort.

Unsupported PCRE RegEx patterns

The following PCRE regex constructs are not supported by the Gen 3 APF regex engine. These limitations apply to both regex-based pattern matching and regex-based masking on Gen 3 GigaSMART modules.

- Continuation escape
- Collating elements
- Equivalence classes
- Line endings escape
- Character class subtraction
- Character class symmetric difference
- Character class complement
- Possessive quantifiers
- Backreferences
- Line comments
- Branch reset
- Lookahead
- Lookbehind
- Non-backtracking expressions
- Recursion
- Conditional expressions
- Subroutines
- Callouts
- Backtracking control verbs

Feature Parity with Gen 2 GigaSMART module and Gen 3 GigaSMART module

Pattern Matching Features	Gen 2	Gen 3
Pattern Matching	✓	✓
Masking with Pattern Matching	✓	✓

Pattern Matching Features	Gen 2	Gen 3
Cross-Packet Pattern Matching	✓	✗
Pattern Matching Hint	✓	✗

NOTE: On Gen 3, masking with pattern matching is supported only for APF gsrules that use Pattern Match with the RegEx type.

Map Statistics

Go to **Map > Statistics** to display counts of the rules that actually matched in a map. A single packet can match one or more rules. For example, if a single packet matches multiple rules in an APF map, all matching rules will have that packet counted against them and the overall map status pass counter will show 1.

Display APF Statistics

Refer to APF Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

Adaptive Packet Filtering Examples

The following are APF examples:

- Identify Social Security Numbers in User-Level Transactions
- Mask Social Security Numbers
- Filter on Fiber Channel over Ethernet (FCOE) Traffic
- Multi-Encapsulation Filtering
- Filter on Subscriber Device IP (User-Endpoint IP or UE-IP)
- Filter on Inner Layer 2-4 Parameters for Unrecognized Headers
- GTP Tunnel ID-Based Filtering
- ERSPAN Tunneling
- Distribute Traffic Based on Inner IP Addresses and Inner TCP Port Values
- MPLS Label Based Filtering
- Combine APF with GigaSMART Operations
- Conditional Header Stripping
- Facilitate Overlapping Rules

Identify Social Security Numbers in User-Level Transactions

The following example looks for packets containing Social Security Numbers in an incoming traffic stream using pattern matching. Once a match is detected, the packets are forwarded to a monitoring tool for additional analysis.

Task	Description	UI Steps
1	Configure one network and two tool ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3. Select Tool for port 1/1/x4 and port 1/1/x1. **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsgrp1 in the Alias field. **4.** Click Save.
3	Configure the GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **2.** Click New. **3.** Type gsfil in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Select APF from the GigaSMART Operations (GSOP) list. **6.** Select Enable. **7.** Click Save.

Task	Description	UI Steps
4	Create a virtual port.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map to forward traffic from network port 1/1/x3 to virtual port vp1.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the port 1/1/x3 for the Source. ▪ Select the virtual port vp1 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IP Version. **d.** Select v4 for Version. **5.** Click Save.
6	Create a second level map to forward traffic from the virtual port vp1 to GigaSMART with pattern matching.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map2 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x1 for the Destination. **4.** To add a rule for Gen 2 GigaSMART modules. Follow the below steps. ▪ Click Add a Rule. ▪ Select Pass. ▪ Select Pattern Matching. ▪ Select regex for Type and enter the value \d{3}-?\d{2}-?\d{4}. ▪ Set the Offset Start to 40. ▪ Set the Offset End to 8. **5.** To add a rule for Gen 3 GigaSMART modules. Follow the below steps. ▪ Click Add a Rule.

Task	Description	UI Steps
		▪ Select Pass. ▪ Select Pattern Matching. ▪ Select regex for Type and enter the value \d{3}-?\d{2}-?\d{4}. **6.** Click Save.

Mask Social Security Numbers

In the following pattern matching example, IPv4 packets contain Social Security Numbers (SSNs) in the format xxx-xx-xxxx. If the SSNs are between offset 40 and 80, they will be replaced with zeros. This example is not applicable to Gen 3 GigaSMART module.

Task	Description	UI Steps
1	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsgrp1 in the Alias field. **4.** Click Save.
2	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter gsTraffic in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Click Save.

Task	Description	UI Steps
3	Create a first level map to direct traffic from network port 1/1/x1 to virtual port gsTraffic.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the port 1/1/x3 for the Source. ▪ Select the virtual port gsTraffic for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IP Version. **d.** Select v4 for Version. **5.** Click Save.
4	Configure the GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations. **2.** Click New. **3.** Select gsgrp1 from the GigaSMART Groups list. **4.** Select Adaptive Packet Filtering from the GigaSMART Operations (GSOP) list. **5.** Select Enable. **6.** Click Save.
5	Create a second level map to direct traffic from the virtual port gsTraffic to GigaSMART.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map2 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x6 for the Destination. ▪ Select gsop1 from the GSOP list. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Pattern Matching. **d.** Select regex for Type and enter the value \d{3}-?\d{2}-?\d{4} **e.** Set the Offset Start to 40. **f.** Set the Offset End to 80

Task	Description	UI Steps
		g. Enter 0 for Mask. 8. Click Save.

Mask Passwords and Account IDs (Gen 3 Regex Masking)

In this example, Adaptive Packet Filtering (APF) on Gen 3 GigaSMART modules uses regex pattern matching with masking to protect sensitive fields—such as passwords or alphanumeric account IDs—embedded in application payloads. The regex defines what to match, and the From/To masking window controls which portion of the match is overwritten in the packet before it is sent to tools.

Example 1 – Mask the entire password value

Many applications send credentials in key–value form, such as password=P@ssw0rd!123, in URLs, form bodies, or headers. In this use case, APF regex masking hides the entire password value, while leaving the password= key visible for troubleshooting.

- Regex pattern
  - password=[^&\s]{5,64}
  - Matches strings that start with password= followed by 5–64 characters that are neither a space nor &.
    - Example matches:
    - password=P@ssw0rd!123
    - password=abc123XYZ!
- Masking window
  - From: start-of-match 9
    - password= is 9 characters.
    - Skips the key and starts masking at the first character of the password value.
  - To: end-of-match 0
    - Masks from that point to the end of the regex match, so the full password value is overwritten.
- Effect on payload
  - Original: password=P@ssw0rd!123
  - Masked (mask byte 0x2a, *): password=**********

For example, you can configure a GigaSMART rule similar to:

gsrule add pass pmatch mask 0x2a from start-of-match 9 to end-of-match 0 RegEx "password=[^&\s]{5,64}" 0..1750

Example 2 – Partially mask account IDs, keeping the last four characters

Some environments require partially masking identifiers, allowing operators to see only the tail of an ID (for example, for audit or troubleshooting) while hiding most of the value. In this use case, APF regex masking hides the leading portion of an account ID and preserves the last four characters.

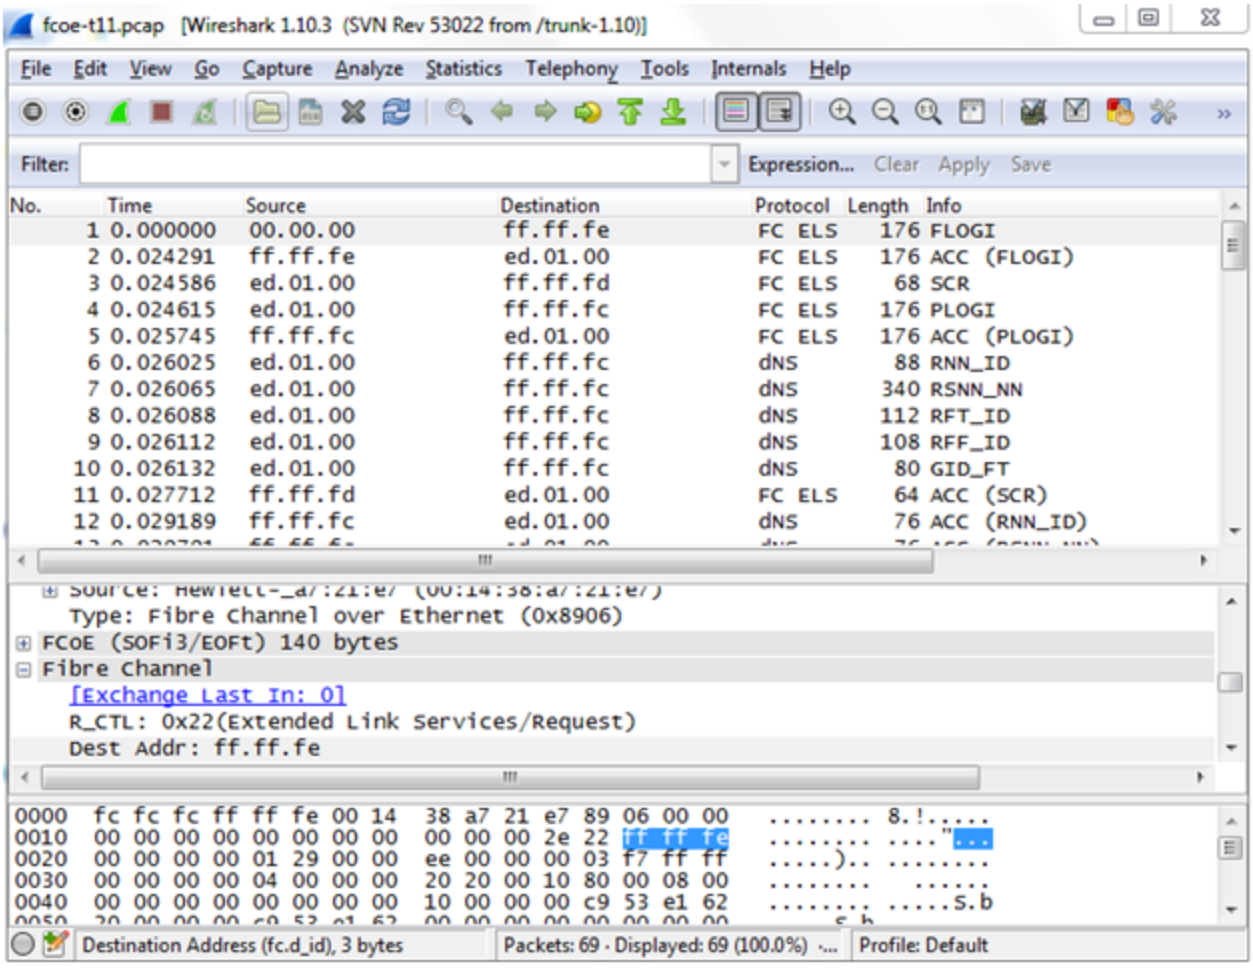
- Regex pattern for accountId=[A-Za-z0-9]{8,20}
  - Matches strings that start with accountId= followed by 8–20 alphanumeric characters. Example match: accountId=AB12CD34EF56
- Masking window
  - From: start-of-match 10
    - accountId= is 10 characters.
    - Skips the key and starts masking at the first character of the account ID value.
  - To: end-of-match -4
    - Stops masking 4 bytes before the end of the match, leaving the last four characters visible.
- Effect on payload
  - Original: accountId=AB12CD34EF56
  - Masked (mask byte 0x2b): accountId=********EF56

A corresponding GigaSMART rule can look like:

```
gsrule add pass pmatch mask 0x2b from start-of-match 10 to end-of-match-4 RegEx "accountId=[A-Za-z0-9]{8,20}" 0..1750
```

Filter on Fiber Channel over Ethernet (FCOE) Traffic

The flexibility offered by regular expression-based filters can be used as an infrastructure to classify traffic streams with protocol headers that are typically unsupported on traditional TAP/SPAN aggregation devices. In this example, regular expression-based filters are used for filtering on the source address in a Fiber Channel header.



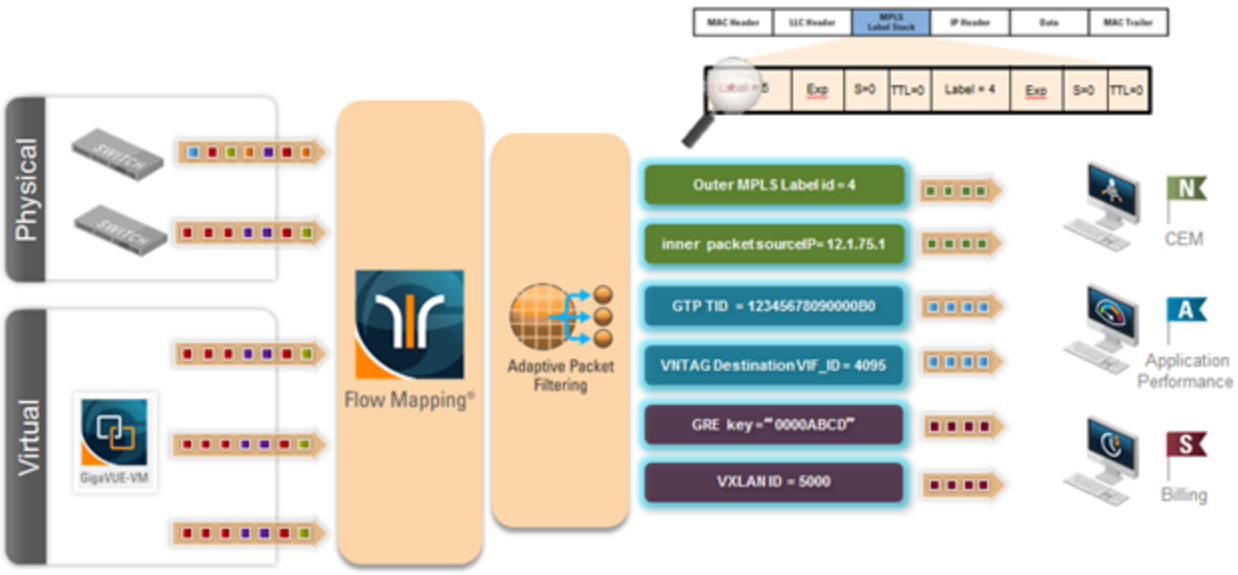
Task	Description	UI Steps
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3. Select Tool for port 1/1/x4 and port 1/1/x1. **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsgrp1 in the Alias field. **4.** Click Save.
3	Configure the GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation.

Task	Description	UI Steps
		2. Click New. **3.** Select gsfil from the GigaSMART Groups list. **4.** Select Adaptive Packet Filtering from the GigaSMART Operations list. **5.** Select Enable. **6.** Click Save.
4	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter gsTraffic in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Click Save.

Task	Description	UI Steps
5	Create a first level map to forward FCOE traffic to the virtual port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter to_vp in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the port 1/1/x3 for the Source. ▪ Select the virtual port vp1 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select EtherType **d.** Enter 8906 in the Value field. **5.** Click Save.
6	Create a second level map to filter on regular expression, using a string match to the destination address in the FCOE packet.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map2 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x1 for the Destination. ▪ Select gsfil from the GSOP list. **4.** To add a rule for Gen 2 GigaSMART modules. Follow the below steps. ▪ Click Add a Rule. ▪ Select Pass. ▪ Select Pattern Match. ▪ Select string for Type and enter txff\xff\xfe. ▪ Set the Offset Start to 0. ▪ Set the Offset End to 29. **5.** To add a rule for Gen 3 GigaSMART modules. Follow the below steps. ▪ Click Add a Rule. ▪ Select Pass. ▪ Select Pattern Match. ▪ Select regex for Type and enter txff\xff\xfe. **6.** Click Save.

Multi-Encapsulation Filtering

In order to complement the mobility brought about by the virtualized server infrastructure, network virtualization overlays like VXLAN, VNTag, NVGRE are being designed and implemented in Data Centers and Enterprise environment. Across Service Provider environments, huge volumes of traffic are being tunneled over GTP. Until now, the Gigamon Deep Observability Pipeline provided the option of stripping out these headers, thus providing visibility to monitoring tools that do not understand these overlays and encapsulation protocol. With APF, this capability is further enhanced where operators now have the option of making forwarding decisions based on the encapsulation and inner packet contents.



With encapsulation awareness enabled by APF, operators have multiple options to act on the packet including the flexibility to:

- Filter on encapsulation header parameters, Layer 2 – 4 parameters in the outer or inner headers (up to 5 layers of encapsulation) in any combination. For example:
  - Forward traffic specific to a subset of VXLAN IDs to one or more monitoring tools.
  - Distribute traffic based on MPLS label values across one or more monitoring tools.
- In combination with Header Stripping:
  - Implement “conditional” header-stripping, based on encapsulation header parameters or inner/outer packet contents, as follows:

    – Forward a subset of traffic “as-is” to monitoring tools that need these encapsulations for analysis.

– Alternatively, strip out the outer headers/encapsulations and distribute traffic to monitoring tools that do not require these outer headers for analysis.

- Since APF is implemented as a second level map, operators can also implement overlapping rules where:
  - A copy of the traffic can be distributed across a group of monitoring tools.
  - A refined subset from the same incoming stream is distributed across a different set of tools.

Filter on Subscriber Device IP (User-Endpoint IP or UE-IP)

Encapsulation awareness enabled by APF allows mobile operators to filter on Layer 2 – 4 header parameters found in an encapsulated packet.

This allows operators to filter and forward traffic specific to a mobile subscriber device or a group of subscriber devices, identified by their IP address (User-Endpoint IP) to one or more monitoring tools.

In this example, we are:

- Identifying and forwarding traffic from / to a UE-IP of 1.1.1.1 to a monitoring tool connected to 1/1/x1
- Identifying and forwarding traffic from / to a UE-IP of 1.1.1.2 to a different monitoring tool connected to tool port 1/1/x4

In many cases, the GTP control sessions are low-volume and are useful in providing some level of visibility in to the quality of experience of the subscribers. To this end, operators prefer to replicate the control sessions across all the monitoring tools, while filtering and forwarding a subset of the user-plane sessions to a subset of monitoring tools. The following example also illustrates configuration commands, leveraging the patented flow-mapping technology to replicate the GTP control sessions across all the monitoring tools involved in the traffic analysis.

Task	Description	UI Steps
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3. Select Tool for port 1/1/x4 and port 1/1/x1. **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and	**1.** From the device view, select GigaSMART >

Task	Description	UI Steps
	associate it with a GigaSMART engine port.	GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Select engine port 1/1/e1 in the Port List field. **5.** Click Save.
3	Configure the GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations. **2.** Click New. **3.** Type gsfil in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list. **5.** Select Adaptive Packet Filtering from the GigaSMART Operations (GSOP) list. **6.** Select Enable. **7.** Click Save.
4	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map to forward GTP-u traffic to the virtual port. **NOTE:** In the rule, 2152 is GTP-u traffic.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter to_vp in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the port 1/1/x3 for the Source. ▪ Select the virtual port vp1 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Port Source **d.** Enter 2152 for the port value. **5.** Click Save.
6	Create a first level map to forward GTP-c traffic to the tools. **Note**: In the rule, 2123 is GTP-c traffic.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type to_tool in the Alias field. ▪ Select Regular for Type.

Task	Description	UI Steps
		▪ Select By Rule for Subtype. ▪ Select the port 1/1/x3 for the Source. ▪ Select port 1/1/x1 and port 1/1/x4 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Port Source **d.** Enter 2123 for the port value. **5.** Click Save.
7	Create a second level map to filter on source and destination IP (bi-directional).	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x1 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Source. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Add a Rule 2. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Click Save.
8	Create another second level map to filter on source and destination IP (bi-directional).	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype.

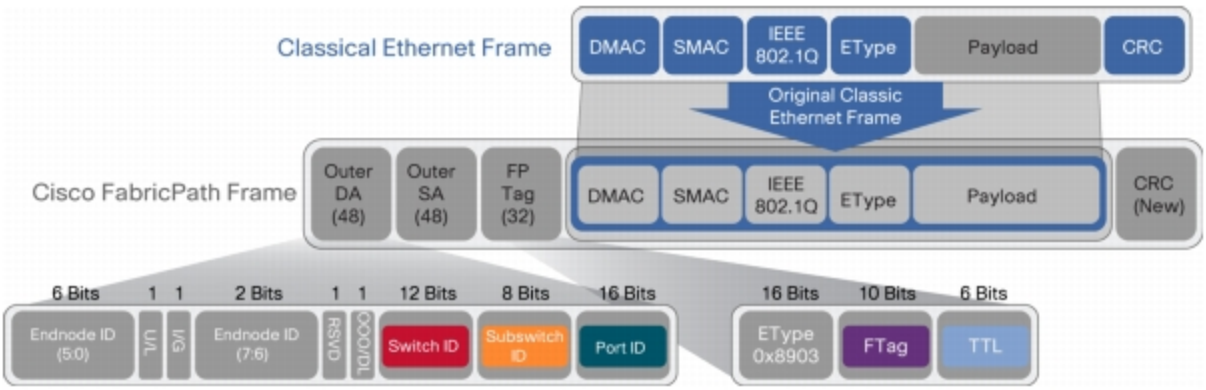
Task	Description	UI Steps
		▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x4 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Source. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Add a Rule 2. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Click Save.

Filter on Inner Layer 2-4 Parameters for Unrecognized Headers

The flexibility of encapsulation awareness enables filtering on encapsulated contents even if APF does not recognize the outer encapsulation header. The following example illustrates a packet encapsulated in Fabric Path headers. Fabric Path headers (as shown in the figure) are mac-in-mac headers that are currently not recognized by APF. However operators can still filter and forward traffic flows based on Layer 2 – 4 parameters found in the encapsulated packets.

In this example, we are:

- Identifying and forwarding traffic from/to ip 1.1.1.1 in the inner / original packet to monitoring tool connected to tool port 1/1/x1
- Identifying and forwarding traffic from/to ip 1.1.1.2 in the inner / original packet to monitoring tool connected to tool port 1/1/x4



Task	Description	UI Steps
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3. Select Tool for port 1/1/x4 and port 1/1/x1. **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Select engine port 1/1/e1 in the Port List field. **5.** Click Save.
3	Configure the GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **2.** Click New. **3.** Type **4.** Select gsfil from the GigaSMART Groups list. **5.** Select Adaptive Packet Filtering from the GigaSMART Operations (GSOP) list. **6.** Select Enable. **7.** Click Save.
4	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Click Save.

Task	Description	UI Steps
5	Create a first level map to forward fabric path packets to the virtual port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter to_vp in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the port 1/1/x3 for the Source. ▪ Select the virtual port vp1 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select EtherType **d.** Enter 8903 in the Value field. **5.** Click Save.
6	Create a second level map to filter on source and destination IP (bi-directional).	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x1 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Source. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 1. **7.** Add a Rule 2. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 1. **7.** Click Save.
7	Create another second level map to	**1.** Select Maps > Maps > Maps.

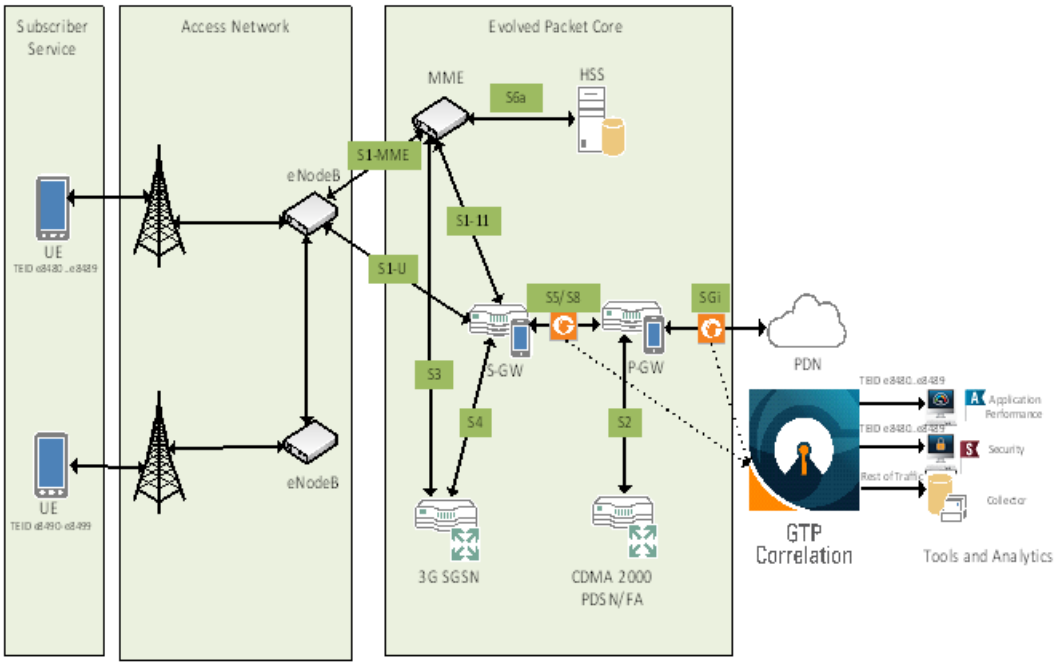
Task	Description	UI Steps
	filter on source and destination IP (bi-directional).	**2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x4 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Source. **d.** Enter 1.1.1.2 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 1. **7.** Add a Rule 2. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination. **d.** Enter 1.1.1.2 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 1. **7.** Click Save.

GTP Tunnel ID-Based Filtering

The following example demonstrates filtering and forwarding traffic based on tunnel IDs included as part of the GTP user-plane messages. It also illustrates the concept of a shared collector to which traffic not matching any of the configured filters can be optionally sent. GTP control sessions are forwarded to all the monitoring tools leveraging the power of flow mapping by filtering on Layer-4 UDP port 2123.

For GTP-u:

- Filter and forward teid ranges 0x001e8480..0x001e8489 to a monitoring tool
- Filter and forward teid ranges 0x001e8490..0x001e8499 to another monitoring tool
- Forward the rest of the traffic to a shared collector



Task	Description	UI Steps
1	Configure one network and three tool type of ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x9. Select Tool for the port s 1/1/x13, 1/1/x14, and 1/1/x15. **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Select engine port 1/1/e1 in the Port List field. **5.** Click Save.
3	Configure the GigaSMART operation and assign it to the GigaSMART group. Packets processed by this operation are evaluated using Adaptive Packet Filtering (APF) rules.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **2.** Click New. **3.** Type gsfil in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list. **5.** Select Adaptive Packet Filtering from the GigaSMART (GSOP) Operations list. **6.** Select Enable. **7.** Click Save.

Task	Description	UI Steps
4	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map that directs GTP-u traffic from physical network port/s to the virtual port created in the previous step. **NOTE:** In the rule, 2152 is GTP-u traffic.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter to_vp in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the port 1/3/x9 for the Source. ▪ Select the virtual port vp1 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select EtherType **d.** Enter 8903 in the Value field. **5.** Click Save.
6	Create a first level map that directs GTP-u traffic from physical network port/s to the tool ports. **NOTE:** In the rule, 2123 is GTP-c traffic.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter ctrl_to_tool in the Alias field. ▪ Select Regular for Type. ▪ Select By Rule for Subtype. ▪ Select the port 1/3/x9 for the Source. ▪ Select the port 1/3/x13 and port 1/3/x15 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Port Source. **d.** Enter **2123** for the port value. **5.** Click Save.
7	Create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, and matches tunnel IDs specified by the gsrule.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type m1 in the Alias field.

Task	Description	UI Steps
		▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the port 1/3/x15 for the Source. ▪ Select the virtual port vp1 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select GTP-U TEID. **d.** Enter 0x001e8480 for Min and 0x001e8489 for Max. **e.** Set Subset to none. **6.** Click Save.
8	Create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, and matches tunnel IDs specified by the gsrule.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type m2 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the port 1/3/x15 for the Source. ▪ Select the virtual port vp1 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select GTP-U TEID. **d.** Enter 0x001e8490 for Min and 0x001e8499 for Max. **e.** Set Subset to none. **6.** Click Save.
9	Add a shared collector for any unmatched data and send it to the third tool port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type scoll in the Alias field. ▪ Select Second Level for Type. ▪ Select Collector for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the port 1/3/x14 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Click Save.

ERSPAN Tunneling

In this example, APF is used to filter packets based on ERSPAN ID. The ERSPAN header is not removed from the packet.

A second level map is configured in the example. A virtual port feeds traffic to the second level map. APF filters the packets and forwards those that match the filter criteria in the map.

Task	Description	UI Steps
1	Configure a tool type of port.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Select Tool for a port. For example, port 1/1/g1. **4.** Select Enable. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsgp2 in the Alias field. **4.** Select an engine port 1/3/e1 in the Port List field. For example, 1/3/e2 **5.** Click Save.
3	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp in the Alias field. **4.** Select gsgrp2 from the GigaSMART Groups list. **5.** Click Save.

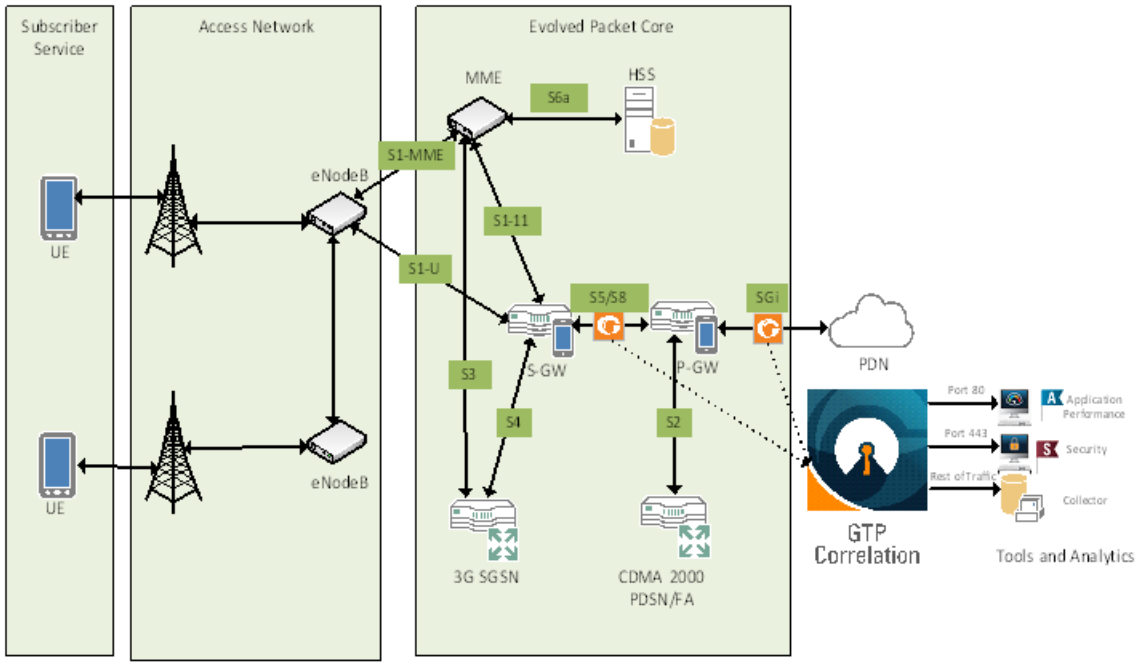
Task	Description	UI Steps
4	Configure the GigaSMART operation and assign it to the GigaSMART group.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations. **2.** Click New. **3.** Type er2 in the Alias field. **4.** Select gsgp2 from the GigaSMART Groups list. **5.** Select Adaptive Packet Filtering from the GigaSMART Operations list. **6.** Select Enable. **7.** Click Save.
5	Create a first level map.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type test1a in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the network port for the Source. For example, 1/1/g3. ▪ Select the virtual port vp for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select MAC Source **d.** Enter the address 0000.0000.0000 for Min and the address 0000.0000.0000 for Max. **5.** Click Save.
6	Create a second level map.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type test1b in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the network port for the Source. For example, 1/1/g3. ▪ Select the virtual port vp for the Destination. ▪ Select er2 from the GSOP list. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass.

Task	Description	UI Steps
		c. Select MAC Source d. Enter the address 0000.0000.0000 for Min and the address 0000.0000.0000 for Max. 5. Click Save.

Distribute Traffic Based on Inner IP Addresses and Inner TCP Port Values

In the following example, traffic is distributed based on inner IP addresses and inner TCP port values as follows:

- Packets from VLAN 20 with GTP inner IP 65.128.7.21 and 98.43.132.70, inner TCP port 80 is forwarded to one tool port
- Packets from VLAN 20 with GTP inner IP 65.128.7.21 and 98.43.132.70, inner TCP port 443 is forwarded to a second tool port
- All packets not matching these rules is forwarded to a third tool port



Task	Description	UI Steps
1	Configure one network and three tool type of ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and three tool ports. For example, select Network for port 1/1/x1. Select Tool for the port s 1/1/x10, 1/1/x11, and 1/1/x12. **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsgrp1 in the Alias field. **4.** Select an engine port in the Port List field. For example, 1/1/e1. **5.** Click Save.
3	Configure the GigaSMART operation and assign it to the GigaSMART group. Packets processed by this operation are evaluated using Adaptive Packet Filtering (APF) rules.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **2.** Click New. **3.** Type g1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Select APF from the GigaSMART Operations (GSOP) list. **6.** Select Enable. **7.** Click Save.
4	Configure a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter gsTraffic in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map that directs traffic from the physical network port to the virtual port created in the previous step.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type map1 in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the network port for the Source. For example, 1/1/x1 ▪ Select the virtual port gsTraffic for the Destination. **4.** Add a rule with three conditions. **a.** Click Add a Rule.

Task	Description	UI Steps
		b. Select Pass. **c.** Select VLAN and enter 20 for Min. **d.** Select IPv4 Protocol and select UDP for Value. **e.** Select Port Destination and enter 2152 for the port value **6.** Click Save.
6	Create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, matches the rules, and sends the traffic to one tool port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type map2 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port gsTraffic for the Source. ▪ Select the port 1/1/x10 for the Destination. ▪ Select g1 from the GSOP list. **4.** Add a rule with three conditions. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination then enter 65.128.721 for the IP address and 255.255.255.255 for the Net Mask. Set position to 2. **d.** Select IPv4 Protocol and set the Potion to 2. **e.** Select Port Destination and enter 80 for the port value and select 2 for Position. **6.** Add a rule with three conditions. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination then enter 98.43.132.70 for the IP address and 255.255.255.255 for the Net Mask. Set Position to 2. **d.** Select IPv4 Protocol and set the Position to 2. **e.** Select Port Destination and enter 80 for the port value and select 2 for Position. **6.** Click Save.
7	Create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, matches the rules, and sends the traffic to another tool port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type map3 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port gsTraffic for the Source.

Task	Description	UI Steps
		▪ Select the port 1/1/x10 for the Destination. ▪ Select g1 from the GSOP list. **4.** Add a rule with three rule conditions. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination then enter 65.128.721 for the IP address and 255.255.255.255 for the Net Mask. Set Position to 2. **d.** Select IPv4 Protocol. Set Position to 2 **e.** Select Port Destination and enter 443 for the port value and select 2 for Position. **6.** Add another rule with three rule conditions. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination then enter 98.43.132.70 for the IP address and 255.255.255.255 for the Net Mask. Set position to 2. **d.** Select IPv4 Protocol. Set position to 2. **e.** Select Port Destination and enter 443 for the port value and set Position to 2. **6.** Click Save.
8	Add a shared collector for any unmatched data and send it to the third tool port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type mapclin the Alias field. ▪ Select Second Level for Type. ▪ Select Collector for Subtype. ▪ Select the virtual port gsTraffic for the Source. ▪ Select the port 1/1/x12 for the Destination. **4.** Click Save.

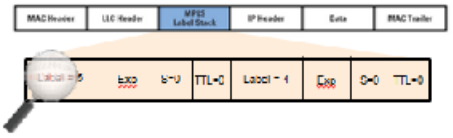
MPLS Label Based Filtering

Multiprotocol Label Switching (MPLS) is a mechanism in high-performance telecommunications networks that directs data from one network node to the next based on short path labels rather than long network addresses, avoiding complex lookups in a routing table. The labels identify virtual links (paths) between distant nodes rather than endpoints.

MPLS is a scalable, protocol-independent transport. In an MPLS network, data packets are assigned labels. Packet-forwarding decisions are made solely on the contents of this label, without the need to examine the packet itself. This allows one to create end-to-end circuits across any type of transport medium, using any protocol.

However in the context of Gigamon Deep Observability Pipeline nodes, traffic flows encapsulated in MPLS labels cannot be filtered and forwarded. With the wide-scale adoption of MPLS as a technology across enterprise and service provider environments, the ability to classify traffic flows based on MPLS labels would be a huge value add to granularly control the flow of traffic to the monitoring tools. APF can be leveraged to filter and forward traffic flows based on MPLS label values. MPLS can stack multiple labels to form tunnels within tunnels. The flexibility of APF facilitates traffic classifications across up to 5 levels of MPLS label stacks in addition to the capability to filter and forward based on Layer 2-4 parameters found in the encapsulated packet. The following example illustrates filtering and forwarding traffic based on MPLS labels, as follows:

- Filter and forward traffic flows specific to mpls label = 4 at the second level in the MPLS label stack to tool 1
- Filter and forward traffic flows specific to mpls label = 3 at the first level in the MPLS label stack to tool 2



Step	Description	Command
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3 and select Tool for the port s 1/1/x4 and 1/1/x1 **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Select an engine port in the Port List field. For example, 1/1/e1 **5.** Click Save.
3	Configure the GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > Operations.

Step	Description	Command
		2. Click New. **3.** Type gsfil in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list. **5.** Select APF from the GigaSMART Operations list. **6.** Select Enable. **7.** Click Save.
4	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp1 in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list. **5.** Click Save.

Step	Description	Command
5	Create a first level map to forward traffic to the virtual port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type to_vp in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the network port for the Source. For example, 1/1/x3 ▪ Select the virtual port vp1 for the Destination. **4.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IP Version and set Version to v4. **4.** Add Rule 2. **a.** Click Add a Rule. **b.** Select Pass and Bi Directional. **c.** Select MAC Source and enter 00:00:00:00:00:00 for the address. **d.** Set Version to v4. **5.** Click Save.
6	Create another second level map to filter on MPLS label.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the port 1/1/x1 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select MPLS Label. **d.** Set the value to 4 and the Position to 1 **5.** Click Save.
7	Create another second level map to filter on MPLS label.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type map2 in the Alias field. ▪ Select Second Level for Type.

Step	Description	Command
		▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the port 1/1/x4 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select MPLS Label. **d.** Set the value to 3 and the Position to 1 **5.** Click Save.

Combine APF with GigaSMART Operations

APF can also be combined with other GigaSMART functions including Header Stripping, Packet Slicing or Masking, De-duplication and FlowVUE. This provides network administrators and operators to perform a second layer of filtering in combination with the GigaSMART tool optimization and packet manipulation operations.

In the following example, operators can distribute traffic to monitoring tools based on decapsulated contents, more specifically, after Header stripping VXLAN:

- Identifying and forwarding traffic from/to ip 1.1.1.1 from the decapsulated packets to monitoring tool connected to tool port 1/1/x1
- Identifying and forwarding traffic from/to ip 1.1.1.2 in the decapsulated packets to monitoring tool connected to tool port 1/1/x4

NOTE: This can be applied to any protocol that is supported through header-stripping, for example:

- GTP, VXLAN, ISL, MPLS, MPLS+VLAN, VLAN, VN-Tag, fabric-path.
- This is also supported for Gigamon tunnel decapsulation.

Step	Description	Command
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3 and select Tool for the port s 1/1/x4 and 1/1/x1 **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Select an engine port in the Port List field. For example, 1/1/e1 **5.** Click Save.
3	Configure the GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations. **2.** Click New. **3.** Type gsfil_vxlanhs in the Alias field. **4.** Select gsg1 from the GS Groups list. **5.** Select Adaptive Packet Filtering from the GigaSMART Operations (GSOP) list and Enable. **6.** Select Strip Header from the GigaSMART Operations (GSOP) list and select VXLAN. **7.** Click Save.
4	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp1 in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map to forward VXLAN traffic to the virtual port. VXLAN accepts destination UDP ports 8472 and 4789. Starting in software version 4.5.01, VXLAN also accepts destination UDP port 48879.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type to_vp in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the network port for the Source. For example, 1/1/x3 ▪ Select the virtual port vp1 for the Destination.

Step	Description	Command
		4. Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Port Source and set the port value to 8472. **4.** Click Save.
6	Create a second level map to filter on source and destination IP (bi-directional).	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x1 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Source. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Add a Rule 2. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Click Save.
8	Create another second level map to filter on source and destination IP (bi-directional).	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x1 for the Destination. ▪ Select gsfil from the GSOP list.

Step	Description	Command
		4. Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Source. **d.** Enter 1.1.1.2 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Add a Rule 2. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination. **d.** Enter 1.1.1.2 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Click Save.

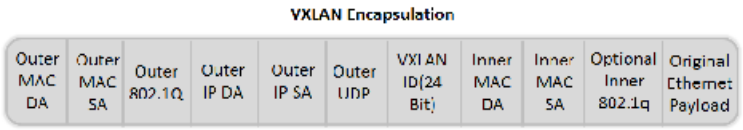
Conditional Header Stripping

Another use-case that can be addressed leveraging the flexibility of APF would be the capability to header strip packets based on specific contents found across the packet including the inner packet contents. Since the APF rules are enforced before any other GigaSMART operation, operators can filter based on encapsulation protocol values and /or encapsulated (original) packet contents and apply conditional Header Stripping operations.

The following example shows how an end-user can filter and strip out outer VXLAN headers for a subset of the traffic based on inner IP addresses, while sending the rest of the traffic “as-is” to monitoring tools that need the VXLAN headers for traffic analysis, as follows.

- Identifying and forwarding traffic from/to ip 1.1.1.1 in the inner / encapsulated packets to monitoring tool connected to tool port 1/1/x1 *after* Header Stripping VXLAN.
- Identifying and forwarding traffic from/to ip 1.1.1.2 in the inner / encapsulated packets to monitoring tool connected to tool port 1/1/x4 *without* stripping the VXLAN header.

NOTE: This can be applied to any GigaSMART operation. While this example shows filtering based on inner packet contents, conditional SMART operations can be applied by filtering on encapsulation headers as well.



NOTE: This can be applied to any protocol that is supported through Header Stripping. GTP, VXLAN, ISL, MPLS, MPLS+VLAN, VLAN, VN-Tag, and fabric-path are all supported, as is Gigamon tunnel decapsulation.

Step	Description	Command
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3 and select Tool for the port s 1/1/x4 and 1/1/x1 **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups (GSOP) > GigaSMART Groups. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Select an engine port in the Port List field. For example, 1/1/e1. **5.** Click Save.
3	Configure the GigaSMART operations.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations and create two GigaSMART Operations. **2.** Create the first operation. **a.** Click New. **b.** Type gsfil_vxlanhs in the Alias field. **c.** Select gsg1 from the GigaSMART Groups list. **d.** Select Adaptive Packet Filtering from the GigaSMART Operations list and Enable. **e.** Select Strip Header from the GigaSMART Operations list and select VXLAN. **f.** Click Save. **7.** Create second first operation. **a.** Click New. **b.** Type gsfil apf in the Alias field. **c.** Select gsg1 from the GigaSMART Groups list. **d.** Select Adaptive Packet Filtering from the GS Operations (GSOP) list and Enable. **e.** Click Save.

Step	Description	Command
4	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp1 in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list. **5.** Click Save.
5	Create a first level map to forward VXLAN traffic to the virtual port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type to_vp in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the network port for the Source. For example, 1/1/x3 ▪ Select the virtual port vp1 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Port Source and set the port value to 8472. **4.** Click Save.
6	Create a second level map to filter on source and destination IP (bi-directional), using first GigaSMART operation.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x1 for the Destination. ▪ Select **gsfil_vxlanhs** from the GSOP list. **4.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Source. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Add a Rule 2.

Step	Description	Command
		a. Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Click Save.
7	Create another second level map to filter on source and destination IP (bi-directional), using second GigaSMART operation.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Enter map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the tool port 1/1/x4 for the Destination. ▪ Select gsfil from the GSOP list. **4.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Source. **d.** Enter 1.1.1.2 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Add a Rule 2. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination. **d.** Enter 1.1.1.2 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **7.** Click Save.

Facilitate Overlapping Rules

Because APF is implemented as a second level map operation, APF can also be leveraged for implementing basic overlapping rules. For the same incoming input stream, a copy of the traffic can be sent out to a group of monitoring tools while a refined

subset of the traffic stream can be sent to a different set of monitoring tools. Typically overlapping rules would be implemented by combining APF with the patented Flow Mapping® technology.

Note that Role-Based Access control in the case of APF is applied at the gsgroup / **e** port.

In the following example, for the same input stream:

- HTTP traffic is identified and distributed to a monitoring tool connected to tool port 1/1/x1.
- At the same time, the same stream of HTTP packets are being sent out after slicing unwanted packet contents to a different monitoring tool connected to tool port 1/1/x4.

Step	Description	Command
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3 and select Tool for the port s 1/1/x4 and 1/1/x1 **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Select an engine port in the Port List field. For example, 1/1/e1. **5.** Click Save.
3	Configure the GigaSMART operations.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOPS) > Operations and create two GigaSMART Operations. **2.** Create the first operation. **a.** Click New. **b.** Type gsfil in the Alias field. **c.** Select gsg1 from the GigaSMART Groups list. **d.** Select APF from the GigaSMART Operations list and Enable. **e.** Click Save. **6.** Create second operation. **a.** Click New. **b.** Type gsfil_slice in the Alias field. **c.** Select gsg1 from the GigaSMART Groups list.

Step	Description	Command
		d. Select APF from the GigaSMART Operations (GSOP) list and Enable. **e.** Select Slicing from the GigaSMART Operations (GSOP) list and select None. Set Offset to 150. **6.** Click Save.
4	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp1 in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list. **5.** Click Save.

Step	Description	Command
5	Create a first level map to forward traffic to the virtual port. Port 1/1/x1 and vertual ort vp1 are sent destination port 80 traffic, which is HTTP.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type to_vp in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the network port 1/1/x3 for the Source. ▪ Select the virtual port vp1 and the tool port 1/1/x1 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Port Source and set the port value to 2152. **4.** Click Save.
6	Create a second level map to filter on HTTP traffic and slice it.	**1.** Click New. **2.** Configure the map. ▪ Type map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the virtual tool port 1/1/x4 for the Destination. ▪ Select gsfil_slice form the GSOP list. **3.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IP Version **d.** Set Version to v4 **e.** Set Position to 1 **6.** Click Save.
7	Create another second level map for the rest of the traffic.	**1.** Click New. **2.** Configure the map. ▪ Type map2 in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the virtual port 1/1/x1 for the Destination. ▪ Select gsfil from the GSOP list. **3.** Add a rule.

Step	Description	Command
		a. Click Add a Rule. **b.** Select Pass. **c.** Select IP Version. **d.** Set Version to v4. **e.** Set Postilion to 4 **6.** Click Save.

In the following example, for the same traffic stream, TCP traffic is sent to one monitoring tool while forwarding a subset of TCP flows specific to HTTP to another monitoring tool connected to tool port 1/1/x4.

Step	Description	Command
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3 and select Tool for the port s 1/1/x4 and 1/1/x1 **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Click New. **3.** Type gsg1 in the Alias field. **4.** Select an engine port in the Port List field. For example, 1/1/e1. **5.** Click Save.
3	Configure the GigaSMART operations.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOPS) > GigaSMART Operation and create two GigaSMART Operations. **2.** Click New. **3.** Type gsfil in the Alias field. **4.** Select gsg1 from the GS Groups list. **5.** Select Adaptive Packet Filtering from the GS Operations list and Enable. **6.** Click Save.
4	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual ports. **2.** Click New. **3.** Enter vp1 in the Alias field. **4.** Select gsg1 from the GigaSMART Groups list.

Step	Description	Command
		5. Click Save.
5	Create a first level map to forward TCP traffic to the virtual port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. ▪ Type to_vp in the Alias field. ▪ Select First Level for Type. ▪ Select By Rule for Subtype. ▪ Select the network port 1/1/x3 for the Source. ▪ Select the virtual port vp1 and the tool port 1/1/x4 for the Destination. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Protocol and set the value to TCP. **4.** Click Save.
6	Create a second level map to filter on HTTP traffic.	**1.** Click New. **2.** Configure the map. ▪ Type map1 in the Alias field. ▪ Select Second Level for Type. ▪ Select By Rule for Subtype. ▪ Select the virtual port vp1 for the Source. ▪ Select the virtual tool port 1/1/x1 for the Destination. ▪ Select gsfil form the GSOP list. **3.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Destination. **d.** Set Position to 2 **e.** Set the port value to 80. **6.** Add Rule 2. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Source. **d.** Set Position to 2 **e.** Set the port value to 80. **6.** Click Save.

GigaSMART Load Balancing

GigaSMART Enhanced Load Balancing require a separate license. Stateless Load Balancing is included with Base licenses. Stateful Load Balancing for GTP is included with the GTP Filtering & Correlation license.

Stateful Load Balancing for ASF is included with the Application Session Filtering (ASF) license. Stateful Load Balancing for tunnel is included with the Advanced Tunneling license (GigaVUE-HC3), and Tunneling license (GigaVUE-HC1)

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

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

Load balancing distributes GigaSMART outgoing traffic to multiple tool ports or multiple tunnel endpoint destinations. In this way, traffic processed by GigaSMART is shared.

Stateful load balancing distributes GigaSMART processed traffic to multiple tool ports or tunnel endpoints based on GigaSMART application-specific flow sessions. Stateless load balancing distributes GigaSMART processed traffic to multiple tool ports or tunnel endpoints based on hash values generated from predefined protocol fields in the packet.

Load balancing operations to tool ports can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to Groups of GigaSMART Engine Ports for details.

NOTE: You can combine the front and rear Gen 3 modules of GigaVUE-HC1-Plus. When combined, weighted load balancing is automatically applied for traffic processing.

The following sections describe the available load balancing schemes:

- GigaSMART Load Balancing
- Stateless Load Balancing
- Enhanced Load Balancing

Stateful Load Balancing

Stateful load balancing distributes GigaSMART processed traffic to multiple tool ports or tunnel endpoints based on GigaSMART application-specific flow sessions.

With stateful load balancing, packets belonging to the same flow session maintained by GigaSMART applications are forwarded to the same tool port or tunnel endpoint within a port group.

NOTE: When stateful tunnel load-balancing is configured in Generation 3 GigaSMART card (SMT-HC1-S), the load balancing is undefined until the end-points accessibility become stable.

Use the **GigaSMART Operations (GSOP)** page to configure load balancing. Specify one stateful application within a group of GigaSMART operations and specify a load balancing metric.

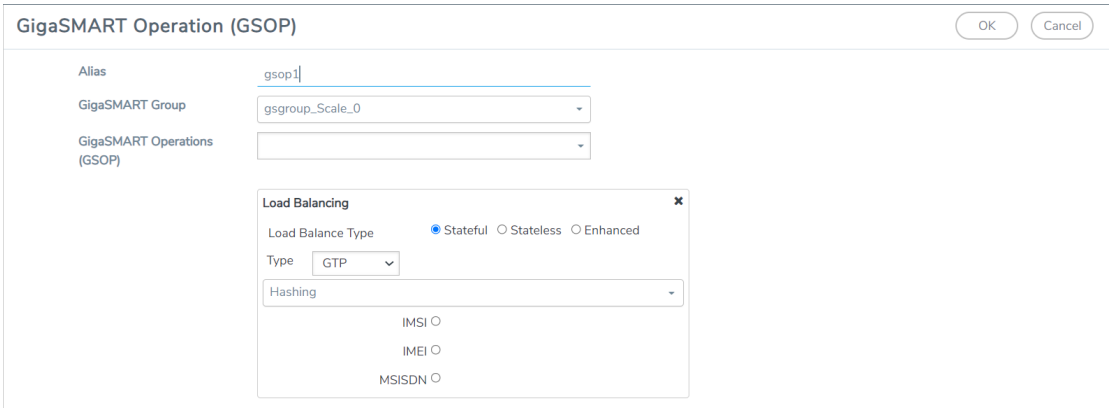
The following are the supported stateful applications:

Application	Reference
GTP	GigaSMART GTP and CUPS Correlation
Application Session Filtering (ASF)	GigaSMART Application Session Filtering ASF and Buffer ASF
Tunnel	GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation
SIP	GigaSMART SIP/RTP Correlation

To select Stateful Load Balancing, do the following:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP)**, and then click **New**.
2. Specify an alias in the **Alias** field.
3. Click in the **GigaSMART Group** field and select a GigaSMART Group.
4. Click in the **GigaSMART Operations (GSOP)** field and select **Load Balancing** from the drop-down list.
5. Select **Stateful**.
6. For **Type**, select one stateful application within a group of GigaSMART operations.

   GTP, **ASF**, **Tunnel**, and **SIP** are the supported stateful applications.
7. Specify a load balancing metric. For example, **Hashing** as shown in the following figure.
8. Click Ok.



For details on stateful load balancing, refer to the following sections:

- Stateful Load Balancing Metrics
- Hashing Key Support
- Configure Stateful Load Balancing

Stateful Load Balancing Metrics

The load balancing metrics available for stateful load balancing are described in the following table.

For weighted metrics, such as Weighted Least Bandwidth, you can either define a weight for each port such as 5,10, 25, 50, or you can use link speed: 1 for 1Gb, 10 for 10Gb, 40 for 40Gb, 100 for 100Gb. Use the **Port Groups** configuration page to select weight and use the **Port** configuration page to select link speed.

NOTE: Only the traffic from the stateful application (for example, GTP, ASF, or tunnel) is used to perform load balancing. Other traffic in the map that does not match the application's filter rule is excluded.

Metric	Description
Least Bandwidth	A tool port is selected from a port group based on the least bits per second load to the port. To compensate for bursty traffic, the history of the last 10 second bandwidth is considered on the load balancing decision. This metric is not supported for tunnel.
Weighted Least Bandwidth	A tool port is selected from a port group based on the least bits per second load to the port, as described under Least Bandwidth. This metric is not supported for tunnel.

Metric	Description
	If this metric is selected, link speed is considered in addition to the bandwidth of the port. If this metric is not selected, the weight configured for each port in the port group is considered in addition to the bandwidth of the port.
Least Packet Rate	A tool port is selected from a port group based on Least Packet Rate. To compensate for bursty traffic, the history of the last 10 second packet count is considered on the load balancing decision.
Weighted Least Packet Rate	A tool port is selected from a port group based on Least Packet Rate, as described under Least Packet Rate. With Weighted Least Packet Rate, if a port has a higher weight, it will be sent more traffic. If this metric is selected, link speed is considered with packet rate. If this metric is not selected, the weight configured for each port in the port group is considered with packet rate.
Round Robin	A tool port is selected from a port group based on round robin.
Weighted Round Robin	A tool port is selected from a port group based on least packet rate, as described under Round Robin. If this metric is selected, link speed is considered with Round Robin. If this metric is not selected, the weight configured for each port in the port group is considered with Round Robin.
Least Connection	A tool port is selected from a port group based on the current Least Connection assigned to each tool port. The port with the least number of connections assigned is selected. **Note**: The meaning of connection is defined by the application.
Weighted Least Connection	A tool port is selected from a port group based on the current Least Connection assigned to each tool port, as described under Least Connection. If this metric is selected, link speed is considered with Least Connection. If this metric is not selected, the weight configured for each port in the port group is considered with Least Connection. **Note**: The meaning of connection is defined by the application.
Least Cumulative Traffic	A tool port is selected from a port group based on the least total bytes sent to each tool port. The port with the least number of connections assigned is selected. **Note**: The meaning of connection is defined by the application.
Weighted Least Cumulative Traffic	A tool port is selected from a port group based on the least total bytes sent to each tool port, as described under Least Cumulative Traffic. If this metric is selected, link speed is considered with least cumulative traffic. If this metric is not selected, the weight configured for each port in the port group is considered with Least Cumulative Traffic.
Hashing	A tool port is selected from a port list based on hashing of data provided by the

Metric	Description
	GSOP application, which is normally extracted from the packet. The values for hashing key are: IMSI GTP key (imsi), IMEI GTP key (imei), and MSISDN GTP key (msisdn). The hashing key only applies to the GTP stateful application. Refer to *Hashing Key Support* on page 672 for details.

Hashing Key Support

The following table describes the support for GTP hashing key.

GTP Key	Hashing	(Weighted) Least Bandwidth	(Weighted) Least Packet Rate	(Weighted) Least Round Robin	(Weighted) Least Connection	(Weighted) Least Cumulative Traffic
IMSI	Supported	Supported	Supported	Supported	Supported	Supported
IMEI	Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
MSISDN	Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported

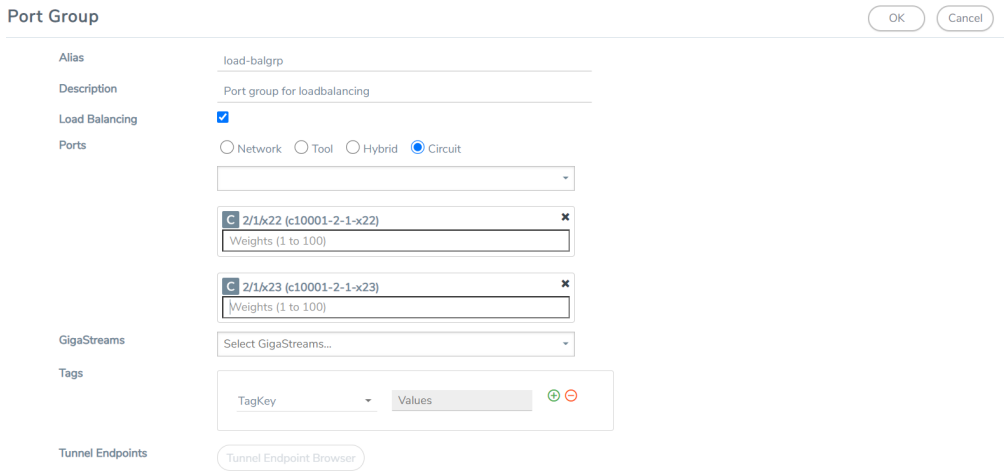
Configure Stateful Load Balancing

Use the Port Group page to specify the list of tool ports or tunnel endpoints for stateful load balancing and to enable load balancing in a port group.

To enable load balancing in a port group, do the following:

1. Select **Ports > Port Groups > All Port Groups**.
2. Click **New**.
3. In the Alias filed, enter an alias. For example, load-balgrp.
4. Select the port type.
5. If applicable, enable **Load Balancing.** Load Balancing is available only for specific port types—Tool, Hybrid, and Circuit.
6. Define the weights for each of the ports.

NOTE: Weight determines the traffic sent to a particular port. The weight of the individual ports must be less than 100. The combined value of the ports can be greater than 100, as the actual load balancing ratio is computed with individual values divided by the combined value



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

NOTES:

- Up to 50 load balancing port groups are supported.
- Ports within port groups can have different rates.

Examples

Refer to the following examples:

- Example 1: GigaSMART Stateful Load Balancing
- Example 2: GigaSMART Stateful Load Balancing

For an example of load balancing on L2GRE encapsulation tunnel, refer to GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation.

Example 1: GigaSMART Stateful Load Balancing

Example 1 configures stateful load balancing of GigaSMART GTP traffic among tool ports 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4 based on bandwidth with different weights for each port. The same subscriber (imsi) traffic will be forwarded to the same tool port. GTP-c packets are replicated to all tool ports.

Task	Description	UI Steps
1.	Create a port group, specify the tool ports for load balancing, and weights for each tool port.	**a.** Select Ports > Port Groups. **b.** Click New **c.** Type portgrp1 in the Alias field. **d.** Enable Load Balancing **e.** Click in the Ports field to select the tool ports for the group. For example, 1/1/x6,1/1/x7,1/2/x3, and 1/2/x4. **f.** Specify the weights for each port as follows: weight 1/1/x6 5 weight 1/1/x7 10 weight 1/2/x3 20 weight 1/2/x4 10 **g.** Click Save.
2.	Create a GigaSMART group and specify a port and enable replicate GTP-c packets to all tool ports in the load balancing port group.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type gsgrp1 in the Alias field. **d.** Select the engine port. For example 1/3/e1. **e.** Under the Load Balance, select Replicate GTP-c. **f.** Click Save.
3.	Create a GSOP, including GTP application and load balancing metric.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Type gsop1 in the Alias field **d.** Select gsgrp1 from the GigaSMART Groups list. **e.** Select the GigaSMART Operation Flow Filtering **f.** Select the GigaSMART Operation Load Balancing and set the operation as follows: o Select Stateful. o Select GTP for Type. o Select Weighted Least Bandwidth for the metric.
4.	Create a virtual port and associate it with the GigaSMART group.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > Virtual Ports. **b.** Click New.

Task	Description	UI Steps
		c. Type vp1 in the Alias field. **d.** Select gsgrp1 from the GigaSMART Group list. **e.** Click OK.
5.	Create an ingress (first level) map. Note the following: ▪ You can specify only one port group as part of the map tool port in the **to** statement. ▪ You can define the same load balancing port group in multiple maps, however, the load balancing metrics defined in the GSOPs on those maps have to be the same. ▪ You cannot use a shared collector map for load balancing. **Note**: In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map11 in the Alias field. **d.** Select First Level for Type and By Rule for Subtype. **e.** Select port 1/1/x1 for the Source. **f.** Select virtual port vp1 for the Destination. **g.** Add Rule 1. o Click Add a Rule o Select Pass o Select Port Destination for the condition. o Set the port value to 2123. **h.** Add Rule 2. o Click Add a Rule o Select Pass o Select Port Destination for the condition. o Set the port value to 2125. **i.** Click Save.
6.	Create a second level map.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map22 in the Alias field. **d.** Select Second Level for Type and By Rule for Flow Filter. **e.** Select virtual port vp1 for the Source. **f.** Select port group portgrp1 for the Destination. **g.** Select gsop1 from the GSOP list. **h.** Click Add a Rule o Select Pass o Select GTP IMSI for the condition. o Enter 234567* for IMSI. o Select Any for Version. **i.** Click Save.

Example 2: GigaSMART Stateful Load Balancing

Example 2 configures stateful load balancing of GigaSMART GTP traffic among tool ports 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4 based on hashing of the imei value. The same device ID (imei) traffic will be forwarded to the same tool port. GTP-c packets are replicated to all tool ports.

Task	Description	Steps
1.	Create a port group and specify the tool ports for load balancing.	**a.** Select Ports > Port Groups. **b.** Click New **c.** Type portgrp1 in the Alias field. **d.** Enable Load Balancing **e.** Click in the Ports field to select the tool ports for the group. For example, 1/1/x6,1/1/x7,1/2/x3, and 1/2/x4. **f.** Click OK.
2.	Create a GigaSMART group and specify ports and enable replicate GTP-c packets to all tool ports in the load balancing port group.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type gsgrp1 in the Alias field. **d.** Select the engine port. For example 1/3/e1. **e.** Under the Load Balance, select Replicate GTP-c. **f.** Click OK.
3.	Create a GSOP, including GTP application and load balancing metric.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations. **b.** Click New. **c.** Type gsop1 in the Alias field **d.** Select gsgrp1 from the GigaSMART Groups list. **e.** Select the GigaSMART Operation Flow Filtering **f.** Select the GigaSMART Operation Load Balancing. **g.** Configure Load Balancing as follows: o Select Stateful. o Select GTP for Type. o Select Hashing for the metric. o Select IMEI

Task	Description	Steps
		h. Click OK.
4.	Create a virtual port and associate it with the GigaSMART group.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > Virtual Ports. **b.** Click New. **c.** Type vp1 in the Alias field. **d.** Select gsgrp1 from the GigaSMART Groups list. **e.** Click OK.
5.	Create an ingress (first level) map. Note the following: ▪ You can specify only one port group as part of the map tool port in the **to** statement. ▪ You can define the same load balancing port group in multiple maps, however, the load balancing metrics defined in the GSOPs on those maps have to be the same. ▪ You cannot use a shared collector map for load balancing. **Note**: In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map11 in the Alias field. **d.** Select First Level for Type and By Rule for Subtype. **e.** Select port 1/1/x1 for the Source. **f.** Select virtual port vp1 for the Destination. **g.** Add Rule 1. o Click Add a Rule o Select Pass o Select Port Destination for the condition. o Set the port value to 2123. **h.** Add Rule 2. o Click Add a Rule o Select Pass o Select Port Destination for the condition. o Set the port value to 2125. **i.** Click Save.
6.	Create a second level map.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map22 in the Alias field. **d.** Select Second Level for Type and By Rule for Flow Filter. **e.** Select virtual port vp1 for the Source. **f.** Select port group portgrp1 for the Destination. **g.** Select gsop1 from the GSOP list. **h.** Click Add a Rule o Select Pass o Select GTP IMSI for the condition.

Task	Description	Steps
		o Enter 234567* for IMSI. o Select Any for Version. i. Click Save.

Stateless Load Balancing

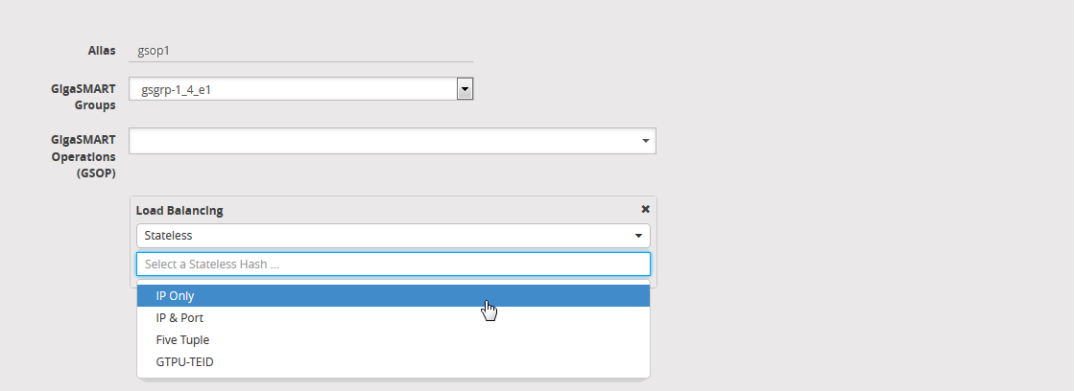
Stateless load balancing distributes GigaSMART processed traffic to multiple tools based on predefined protocol fields in the packet.

Unlike stateful load balancing, stateless load balancing can be configured together with most other GigaSMART operations or as a separate GigaSMART operation to provide more flexible traffic distribution options over what is available from a tool GigaStream. Packets processed by stateless load balancing are forwarded to a tool port within a port group.

Stateless load balancing supports packets with MPLS encapsulation and IEEE 802.1 Q-in-Q VLAN tags.

To select stateless load balancing, do the following:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) GigaSMART Operation**, and then click **New**.
2. Type an alias in the **Alias** field.
3. From the **GigaSMART Groups** drop-down list, select a GigaSMART group.
4. From the **GigaSMART Operations (GSOP)** drop-down list, select **Load Balancing**.
5. Configure Load Balancing:
   - Select **Stateless**
   - Specify a load balancing metric. For example, **IP Only** as shown in the following figure.
6. Click **Save**.



For details on stateless load balancing, refer to the following sections:

Stateless Load Balancing Metrics

The load balancing metrics available for stateless load balancing are described in the following table.

A tool port is selected from a port list based on hashing. The fields to be hashed are described in the table.

To specify a metric for stateless load balancing:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.
2. Click **New** to create a new GigaSMART Operation or **Edit** to modify an existing one.
3. Select **Load Balancing** from the **GigaSMART Operations (GSOP)** drop-down list and configure load balancing as follows.
   a. Select **Stateless**
   b. Select one of the hashing metrics: IP Only, IP & PORT, Five Tuple, or GTPU-TEID.
      Table 1: Stateless Load Balancing Metrics describes each of the metrics.
4. Click **Save**.

Table 1: Stateless Load Balancing Metrics

Metric	Description
IP Only	The source IP and destination IP addresses.
IP & Port	The source IP and destination IP addresses, and Layer 4 source port and destination port numbers.
Five	The source IP and destination IP addresses, source port and destination port numbers, and

Metric	Description
Tuple	protocol field in the IP header.
GTPU-TEID	The GTP-u tunnel identifier (ID). **Note**: There is no inner or outer field location for **`GTPU-TEID`**.
outer	The first occurrence of header or field. For example, **IP Only outer** is the first IP header in the packet, which could be IPv4 or IPv6.
inner	The second occurrence of header or field. For example, **ip-only inner** is the second IP header in the packet. The first IP header could be IPv4 or IPv6, as follows: ▪ IPv4-IPv4—**IP Only inner** is the IP addresses in the second IPv4 header ▪ IPv6-IPv4—**IP Only inner** is the IP addresses in the IPv4 header ▪ IPv4-IPv6—**IP-Only inner** is the IP addresses in the IPv6 header The supported IP encapsulation types are: IP-in-IP, VXLAN, GTP, GRE(including L2GRE), and ERSPAN.

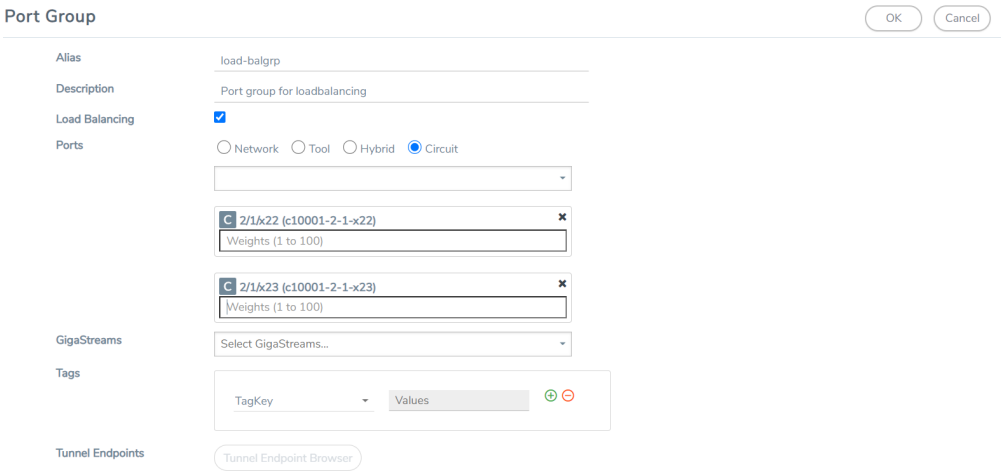
NOTE: On Gen3 GigaVUE HC Series, L2GRE traffic that uses an 8-byte GRE header (4-byte base header + 4-byte key) follows the same inner-header hashing behavior as other supported tunnels. When you select an inner stateless load balancing metric (for example, IP Only inner, IP & Port inner, or Five Tuple inner), L2GRE packets with and without a GRE key are hashed based on the inner IP and port fields. In earlier releases, L2GRE packets with a 4-byte GRE key were always hashed on the outer header, which could result in uneven distribution across GigaSMART CPU cores when many flows shared the same outer 5-tuple.

Configure Stateless Load Balancing

To configure stateless load balancing, specify the group of tool ports and enable load balancing in a port group.

1. Select **Ports > Port Groups > All Port Groups**.
2. Click **New**.
3. Type an alias in the **Alias** field. For example, load-balgrp.
4. Select the Port Type.
5. If applicable, enable **Load Balancing.** Load Balancing is available only for specific port types—Tool, Hybrid, and Circuit.
6. Define the weight for each of the ports used.

NOTE: Weight determines the traffic sent to a particular port. The weight of the individual ports must be less than 100. The combined value of the ports can be greater than 100, as the actual load balancing ratio is computed with individual values divided by the combined value



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

Notes:

- Up to 50 load balancing port groups are supported, with a maximum of 16 ports for each group.
- Ports within port groups must be on the same chassis.
- Ports within port groups can have different rates.

Examples

Refer to the following examples:

- Example 1: GigaSMART Stateless Load Balancing
- Example 2: GigaSMART Stateless Load Balancing
- Example 3: GigaSMART Stateless Load Balancing

For an example of load balancing on L2GRE encapsulation tunnel, refer to Example 3 – GigaSMART L2GRE Tunnel Encap Stateless LB.

Example 1: GigaSMART Stateless Load Balancing

Example 1 configures stateless load balancing of traffic among tool ports 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4 after slicing the packet to an offset of 70 bytes.

Task	Description	Steps
1.	Create a port group and specify the tool ports for load balancing.	**a.** Select Ports > Port Groups. **b.** Click New **c.** Type portgrp1 in the Alias field. **d.** Click in the Ports field to select the tool ports for the group. For example, 1/1/x6,1/1/x7,1/2/x3, and 1/2/x4. **e.** Enable Load Balancing **f.** Click OK.
2.	Configure a GigaSMART group and associate it with GigaSMART engine ports.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type gsgrp1 in the Alias field. **d.** Select the engine ports. For example 1/3/e1 and 1/3/e2. **e.** Click OK.
3.	Create a GSOP, with load balancing.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > Operations. **b.** Click New. **c.** Type lbiponlyouter in the Alias field **d.** Select gsgrp1 from the GigaSMART Groups list. **e.** Select the GigaSMART Operation Load Balancing. **f.** Configure Load Balancing as follows: o Select Stateless o Select IP Only Outer for the hash metric **g.** Click OK.
4.	Create a first level map. Note the following: ▪ You can specify only one port group as part of the map tool port in the **to** statement. ▪ You can define the same load balancing port group in multiple maps, however, the load balancing metrics defined in the GSOPs on those maps have to be the same.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map1 in the Alias field. **d.** Select First Level for Regular and By Rule for Subtype. **e.** Select port 1/1/x1 for the Source. **f.** Select port group portgrp1 for the Destination. **g.** Select lbiponlyouter from the GSOP list. **h.** Click Add a Rule o Select Pass

Task	Description	Steps
	▪ You cannot use a shared collector map for load balancing.	o Select IP Version for the condition. o Select v4 for Version. **i.** Click Save.

Example 2: GigaSMART Stateless Load Balancing

Example 2 configures stateless load balancing of GTP traffic among tool ports 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4. Data packets with the same GTP-u tunnel ID will be forwarded to the same tool port.

Task	Description	Steps
1.	Create a port group and specify the tool ports for load balancing.	**a.** Select Ports > Port Groups. **b.** Click New **c.** Type portgrp1 in the Alias field. **d.** Enable Load Balancing **e.** Click in the Ports field to select the tool ports for the group. For example, 1/1/x6,1/1/x7,1/2/x3, and 1/2/x4. **f.** Click OK.
2.	Configure a GigaSMART group and associate it with GigaSMART engine ports.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type gsgrp1 in the Alias field. **d.** Select the engine ports. For example 1/3/e1 and 1/3/e2. **e.** Click OK.
3.	Create a GSOP, including load balancing metric.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations. **b.** Click New. **c.** Type gsop1 in the Alias field **d.** Select gsgrp1 from the GigaSMART Groups list. **e.** Select the GigaSMART Operation Load Balancing.

Task	Description	Steps
		f. Configure Load Balancing as follows: o Select Stateless o Select GTPU-TEID for the hash metric **g.** Click OK.
4.	Create first level maps. Note the following: ▪ You can specify only one port group as part of the map tool port in the **to** statement. ▪ You can define the same load balancing port group in multiple maps, however, the load balancing metrics defined in the GSOPs on those maps have to be the same. ▪ You cannot use a shared collector map for load balancing.	Create the first map: **a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map1 in the Alias field. **d.** Select First Level for Regular and By Rule for Subtype. **e.** Select port 1/1/x1 for the Source. **f.** Select port group portgrp1 for the Destination. **g.** Click Add a Rule, and then select Pass **h.** Select IPv4 Protocol for the first condition. **i.** Select UDP for Value **j.** Select Port Destination for the second condition. **k.** Enter 2123 for the port value. **l.** Click OK. Create the second map: **a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map1 in the Alias field. **d.** Select First Level for Regular and By Rule for Subtype. **e.** Select port 1/1/x1 for the Source. **f.** Select port group portgrp1 for the Destination. **g.** Select gsop1 from the GSOP list. **h.** Click Add a Rule, and then select Pass **i.** Select IPv4 Protocol for the first condition. **j.** Select UDP for Value **k.** Select Port Destination for the second condition. **l.** Enter 2152 for the port value. **m.** Click OK.

Example 3: GigaSMART Stateless Load Balancing

Example 3 configures stateless load balancing of HTTP on GTP traffic among tool ports 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4. Data packets with the same inner IP will be forwarded to the same tool port.

Task	Description	Steps
1.	Create a port group and specify the tool ports for load balancing.	**a.** Select Ports > Port Groups. **b.** Click New **c.** Type portgrp1 in the Alias field. **d.** Enable Load Balancing **e.** Click in the Ports field to select the tool ports for the group. For example, 1/1/x6,1/1/x7,1/2/x3, and 1/2/x4. **f.** Click OK.
2.	Configure a GigaSMART group and associate it with GigaSMART engine ports.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type gsgrp1 in the Alias field. **d.** Select the engine ports. For example 1/3/e1 and 1/3/e2. **e.** Click OK.
3.	Create a GSOP, including load balancing metric.	**f.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **g.** Click New. **h.** Type gsop1 in the Alias field **i.** Select gsgrp1 from the GigaSMART Groups list. **j.** Select the GigaSMART Operation Load Balancing. o Select Stateless o Select IP Only Inner for the hash metric **k.** Click OK.
4.	Create first level maps. Note the following: o You can specify only one port group as part of the map tool port in the **to** statement.	Create the first map: **a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map1 in the Alias field. **d.** Select Regular for Type and By Rule for

Task	Description	Steps
	o You can define the same load balancing port group in multiple maps, however, the load balancing metrics defined in the GSOPs on those maps have to be the same. o You cannot use a shared collector map for load balancing.	Subtype. **e.** Select port 1/1/x1 for the Source. **f.** Select port group portgrp1 for the Destination. **g.** Click Add a Rule, and then select Pass **h.** Select IPv4 Protocol for the first condition. **i.** Select UDP for Value **j.** Select Port Destination for the second condition. **k.** Enter 2123 for the port value. **l.** Click Save. Create the second map: **a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map1 in the Alias field. **d.** Select First Level for Regular and By Rule for Subtype. **e.** Select port 1/1/x1 for the Source. **f.** Select port vp1 for the Destination. **g.** Select gsop1 from the GSOP list. **h.** Click Add a Rule, and then select Pass **i.** Select IPv4 Protocol for the first condition. **j.** Select UDP for Value **k.** Select Port Destination for the second condition. **l.** Enter 2152 for the port value. **m.** Click Save.
5.	Create the second level map.	Create the third map: **a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map22 in the Alias field. **d.** Select Second Level for Type and By Rule for Subtype. **e.** Select virtual vp1 for the Source. **f.** Select port group portgrp1 for the Destination. **g.** Select gsop1 from the GSOP list. **h.** Click Add a Rule, and then select Pass

Task	Description	Steps
		i. Select IPv4 Destination for the first condition. j. Enter 80 for the destination value. k. Select 2 for position. l. Click Save.

Enhanced Load Balancing

Required License: Enhanced Load Balancing

GigaSMART Enhanced Load Balancing supports evenly distributed traffic among multiple tool ports based on one or more user defined fields. When a tool port fails, the traffic is redistributed just for that tool port to other member tool ports. When the failed tool port recovers, the traffic that was redistributed is restored to the recovered tool port. Traffic across other member tool ports remain undisturbed during this process.

Traffic Handling and Load Balancing Distribution

- Non GTP traffic and (subsequent) fragmented packets to be load balanced to all tool ports based on outer IP.
- Rebalance traffic when the following occurs:
  - Tool port goes up or down
  - Tool port group membership changes
  - When a tool port fails redistribute the traffic just for that tool port and when the failed tool port recovers restore traffic just for that tool port

Enhanced Load Balancing Metrics

GigaSMART provides configuration support for a new enhanced load balancing (enhanced-lb) app. The enhanced-lb app allows users to define the fields used for load balancing.

Enhanced Load Balancing supports the following hashing metrics:

- inner IP address
- outer IP address
- inner L4 port

- outer L4 port
- GPRS Tunnel Endpoint Identifier (TEID)

Configure Enhanced Load Balancing

To configure enhanced load balancing, do the following:

1. Select >**Physical**.
2. Click on a node you want apply enhance load balancing.
3. Select **GigaSMART>Enhanced Load Balancing**. The enhance load balance screen displays.
4. Click **New**. The enhanced load balancing screen appears.
5. Enter an Alias.
6. Select the Hash Field Name. The following options are available:

   - IP

   - L4 Port

   - GTP-U TEID
7. Select Hash Field Location. Options:

   - inner

   - outer

   NOTE: Use the “+” or “-” icons to add and delete hash fields.

Hash Field Name	Hash Field Location
IP	**Inner**
IP	**Outer**
L4 Port	**Outer**
L4 Port	**Inner**
GTP-U TEID	**None**

You can configure up to five (5) different hash fields and location.

8. Click **OK**.

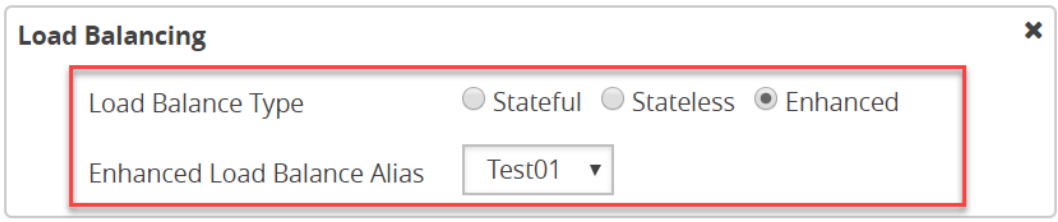
NOTE: To view details about the enhanced load balancing parameters, click the **Alias**. The detail/edit dialog box displays.

Configure GigaSMART Operation (GSOP)

1. Select > **Physical**.
2. Select **GigaSMART > GSOP**.
3. Click **New**.
4. Enter **Alias.**
5. Select a GSOP Group.
6. Select a GSOP type**: Load Balancing**.
7. Select GSOP type: Enhance load balancing. The Load Balancing dialog appears.



8. Select **Enhanced**.
9. Select an Enhanced Load Balance Alias from the drop-down. This is the enhanced load balance you previously created.



10. Click **OK**. The enhanced load balance GSOP is now available on the GSOP page.

GigaSMART De-Duplication

Required License: De-duplication

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

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

GigaSMART de-duplication detects duplicates of the following types:

- IPv4 packets
- IPv6 packets
- non-IP packets (including non-IPv4 and non-IPv6 packets)

Duplicates are packets in which the fields (including the headers and payload) are the same, with the exception of some fields such as Time-to-Live (TTL). For example, if two packets are identical except for TTL, they will be counted as duplicates.

Duplicate packets are common in network analysis environments where both the ingress and egress data paths are sent to a single output (for example, as a result of a SPAN operation on a switch). They can also appear when packets are gathered from multiple collection points along a path. GigaSMART de-duplication lets you eliminate these packets, only forwarding a packet once and thus reducing the processing load on your tools.

Feature Overview

There are two actions that can be specified for handling the duplicate packets detected:

- drop, which drops the duplicate packets
- count, which counts the duplicate packets, but does not drop them

A time interval can be configured within which an identical packet will be considered a duplicate. The greater the interval over which traffic can be checked for duplicates, the higher the accuracy of the de-duplication detection and subsequent elimination.

For example, if two of the same packets are seen in the specified time interval, the packets will be detected as duplicates. If one packet is seen in the time interval and another packet is seen in a later time interval, the packets will not be detected as duplicates.

For IPv4 and IPv6 packets, to determine if a packet is considered to be a duplicate, parts of the IP headers (Layer 3 and Layer 4), as well as part of the payload are compared.

For non-IP packets, a packet is considered to be a duplicate if it is identical.

Keep in mind the following when configuring GigaSMART de-duplication:

Feature	Description
Layer 2 Retransmissions Not Removed	Valid Layer 2 retransmissions are part of normal network behavior and are not removed by the de-duplication feature. Layer 2 retransmissions will show differences in the IP Window ID field.
Encapsulated Duplicates Not Removed	If the same packet is seen once with encapsulation (for example, GRE) and once without encapsulation, the GigaSMART will not detect and remove the duplicate.
No NAT or PAT	Packets tapped on opposite sides of a NAT or PAT boundary will differ in the Network layer and will not be detected as duplicates.

Feature	Description
MPLS and VLAN Tags	De-duplication properly parses VLAN and MPLS tags to get to the IP headers.
VN-Tag Packets	VN-Tag packets are treated as non-IP packets. User Header Stripping to strip VN-Tag to get to the IP headers for de-duplication. Refer to GigaSMART Header Stripping.
GigaSMART Engine Ports	De-duplication operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to Groups of GigaSMART Engine Ports for details.

De-duplication Configuration Steps

To configure de-duplication, use the following steps:

- Configure GigaSMART parameters on a specified GigaSMART group.
- Configure a GigaSMART operation.
- Configure a map that will use the de-duplication GigaSMART operation. This ties de-duplication to rules defined in a flow map, which applies the GigaSMART operation to specific traffic flows.

These steps are detailed in GigaSMART De-Duplication.

Configure GigaSMART Parameters for Packet De-duplication

Use the **de-dupe** section under GigaSMART Parameters on the GigaSMART Groups configuration page to configure options for GigaSMART de-duplication operations. The following table describes the configuring parameters for de-duplication on a specified GigaSMART group:

Parameter	Description
Action	Specifies whether duplicate packets are to be counted or dropped as follows: o Count– GigaSMART counts the duplicate packets, but does not drop them. o Drop– GigaSMART drops the duplicate packets. The default is drop.
IP Tclass **IP TOS** **TCP Sequence** **VLAN**	These options are useful when applying de-duplication operations to packets in a NAT environment. Different NAT implementations can change certain packet header fields (for example, the TCP sequence number). If you want to be able to detect duplicates without requiring that these fields match (ToS field, TCP sequence number, VLAN ID), you can

Parameter	Description
	disable the corresponding option. o **IP Tclass** – Ignore or include IPv6 traffic class. Use for IPv6. The default is include. o **IP TOS** – Ignore or include the IP ToS bits when detecting duplicates. Use for IPv4. The default is include. o **TCP Sequence** – Ignore or include the TCP Sequence number when detecting duplicates. The default is include. o **VLAN** – Ignore or include the VLAN ID when detecting duplicates. The default is ignore. Include means the field will be included when GigaSMART compares packets. Ignore means the field will be ignored when GigaSMART compares packets.
Timer <Value: 10-500000 µs>	Configures the time interval within which an identical packet will be considered a duplicate. The greater the interval over which traffic can be checked for duplicates, the higher the accuracy of the de-duplication detection and subsequent elimination. The default is 50,000µs. For example, if two same packets are seen in the specified time interval, the packets will be detected as duplicates. If one packet is seen in the time interval and another packet is seen in a later time interval, the packets will not be detected as duplicates. **NOTE:** Retransmissions are not counted as duplicates.

Example – GigaSMART De-duplication

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

This example shows the configuration steps for a de-duplication operation in which the GigaSMART application drops duplicate packets.

Task	Description	UI Steps
1.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**a.** From the device view, select GigaSMART > GigaSMART Groups.

Task	Description	UI Steps
		b. Click New. **c.** Type an alias in the Alias field. For example, gs2port1. **d.** Click in Port List field and select an engine port. For example, 2/1/e1. **e.** Go to Task 2.
2.	Configure parameters on the GigaSMART group.	**a.** Under the De-duplication section on the GigaSMART Group configuration page, set the parameters as the follows: Action: drop IP Tclass: Include IP TOS: Ignore TCP Sequence: Ignore Vlan: Ignore Timer (us) 500000 **b.** Click Save.
3.	Configure the GigaSMART operation for de-duplication and assign it to the GigaSMART group.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Type a name for the operation in the alias field. For example, testdedup. **d.** Select the GigaSMART Group create in task 1. **e.** Select De-duplication from the GigaSMART Operations (GSOP) list and select Enable. **f.** Click Save.
4.	Create a map.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type an alias in the Map Alias field that will help you identify this map. For example, testingdedup **d.** Select Regular and By Rule for the map type and subtype. **e.** Specify the network and tool ports in the Source and Destination fields. For example, 2/2/x4 and 2/2/x6 for Source and 2/2/x9 for Destination. **f.** From the GigaSMART Operation (GSOP) drop-down list, select the GigaSMART operation configured in Task 3. For example, testdedup in this example. **g.** Click Add a Rule under Map Rules and create

Task	Description	UI Steps
		the following rule: Select Pass, then select Bi Directional, and then select Port Source from the drop-down list and set the Min to 0 and Max to 443. **h.** Click Save.

Display De-duplication Statistics

To display the statistics for de-duplication in a cluster environment, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics**. The de-duplication statistics will be in the row labeled De-duplication in the GS Operations column.

Refer to De-duplication Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

GigaSMART Traffic Performance Enhancement

This section explains about :

- Flow Masking
- GigaSMART Traffic Performance Enhancement

Flow Masking

Flow Masking helps to identify flows, by selecting the portion of the packet for flow identification. The flow mask consists of an offset and a length, in bytes. Use the offset to specify the number of bytes from the beginning of the packet to the start of the mask within the packet. Use the length to specify the number of bytes, following the offset, to mask within the packet. The length identifies the traffic flow. Both the offset and the length are variable; however, the offset plus the length cannot be greater than 112 bytes. A default mask is provided with an offset of 14 bytes and a length of 28 bytes.

The GigaSMART encapsulated traffic performance is enhanced by using the feature Flow Masking.

For more information about encapsulated traffic performance, refer to GigaSMART Traffic Performance Enhancement.

GigaSMART Encapsulated Traffic Performance Enhancement

GigaSMART Traffic Performance Enhancement does not require a separate license.

The GigaSMART traffic performance enhancement provides a method to improve GigaSMART packet processing for MPLS traffic and other traffic having Layer 2 encapsulation, such as L2GRE or VNTag. This type of traffic has a header in the packet between the MAC address and the IP address. Figure 10MPLS Header Between MAC and IP Address in Packet shows the MPLS example.



Figure 10 *MPLS Header Between MAC and IP Address in Packet*

The GigaSMART processor is able to identify IP flows if there is only the MAC address and no other header in the packet before the IP address or if the only header before the IP address is VLAN. Without this enhancement, the GigaSMART processor cannot identify IP flows if there are MPLS headers or Layer 2 encapsulation other than VLAN before the IP address.

Most of theGigaSMART features require packets belonging to the same flow connection to be processed in the correct order. The processor uses the IP addresses and protocol to identify a traffic flow. When packets belonging to the same flow are identified, GigaSMART process these packets in a sequential manner. GigaSMART process all the encapsulated packets one at a time sequentially even though they belong to different traffic flows (different IP address pair).

Performance is impacted if the GigaSMART processor cannot use the IP source and destination (ipsrc, ipdst) to identify flows. This enhancement provides another method to identify flows. Using a flow mask, you select the portion of the packet for flow identification.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure the MPLS traffic performance enhancement, do the following:

1. From the device view, select **GigaSMART > GigaSMART Group > GigaSMART Groups**.
2. Click **New** to create a new GigaSMART group or **Edit** to modify an existing one.
3. Go to Flow Mask under GigaSMART Parameters and select **Enable**.

4. Enter the offset and length in the **Offset (bytes)** and **Length (bytes)** fields.

   If you do not enter any values in these fields, the default offset and length is used.

5. Click **Save**.

Refer to the following sections for examples:

- Flow Masking Example 1
- Flow Masking Example 2

Flow Masking Example 1

In Example 1 packets are expected to have two MPLS labels before the IP header, and no VLAN tag between the MAC and MPLS headers. IP addresses will be used to identify the flows.

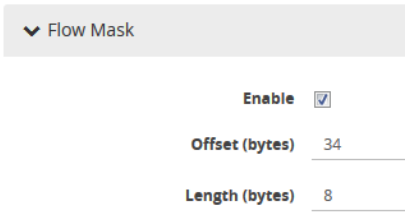
The offset will be the sum of the following: 14 bytes for the Ethernet header+ 8 bytes for the MPLS headers +12 bytes offset from the beginning of the IP header = 34 bytes.

The length will be the sum of the following: 4 bytes for ipsrc + 4 bytes for ipdst = 8 bytes.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure Example 1, do the following:

1. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.
2. Click **New** to create a new GigaSMART Group or **Edit** to modify an exiting one.
3. Under GigaSMART Parameters, go to Flow Mask and select **Enable**.
4. Enter 34 in the **Offset (bytes)** field and 8 in the **Length (bytes)** field as shown in the following figure



5. Click **Save**.

Flow Masking Example 2

In Example 2, packets are expected to have one VLAN tag and two MPLS labels before the IP header. IP addresses will be used to identify the flows.

The offset will be the sum of the following: 14 bytes for the Ethernet header + 4 bytes for the VLAN tag + 8 bytes for the MPLS headers +12 bytes offset from the beginning of the IP header = 38 bytes.

The length will be the sum of the following: 4 bytes for ipsrc + 4 bytes for ipdst = 8 bytes.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure Example 2, do the following:

To configure Example 1, do the following:

1. From the device view, select **GigaSMART > GS Groups > GS Groups**.
2. Click **New** to create a new GigaSMART Group or **Edit** to modify an exiting one.
3. Click **Enable** under GS Params Flow Mask.
4. Enter 38 in the **Offset (bytes)** field and 8 in the **Length (bytes)** field as shown in the following figure.



5. Click Save.

GigaSMART Header Stripping

Required License: Header Stripping

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

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

GigaSMART operations with a **Strip Header** component can identify and remove headers from tagged packets or headers and trailers from tunneled (encapsulated) packets.

The following types of packets can be stripped:

- **Header Stripping** – Remove headers from ERSPAN, MPLS, MPLS+VLAN, VLAN, VN-Tag, VXLAN, GRE, tagged packets, Cisco FabricPath Headers, or FM6000 timestamps before they are sent to tool ports. This feature is handy when working with tools that either cannot recognize these headers or have to engage in additional processing to adjust for them. Figure 11Strip Header GigaSMART Operation Configured shows the GigaSMART Operations page for configuring Header Stripping.
- **Tunnel Stripping** – Remove both the header and trailer of ISL or GTP-encapsulated packets, preserving the packet within for analysis. This is handy when sending data to tools that cannot parse the ISL or GTP tunnel information and analyze the packets within. Figure 12Tunnel Stripping GigaSMART Operation Configured shows an example of the GS Operations page configured for tunnel stripping.

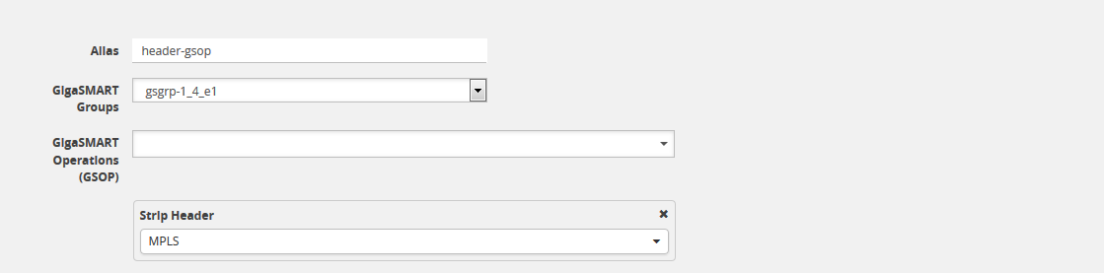


Figure 11 *Strip Header GigaSMART Operation Configured*

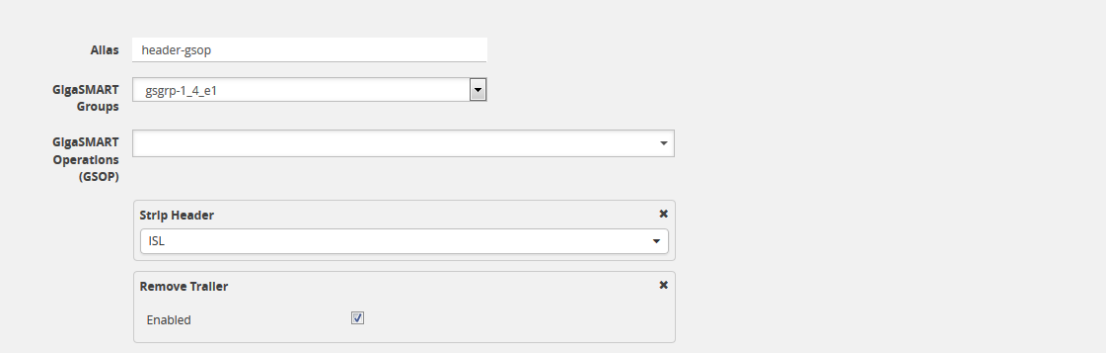


Figure 12 *Tunnel Stripping GigaSMART Operation Configured*

You can also use the **Strip Header** feature in tandem with the **Add VLAN** component to differentiate stripped packets from non-stripped packets. This is particularly useful when seeing stripped/non-stripped packets on common IP ranges (10.x.x.x; 192.168.x.x). Refer to the following table for more information.

On GigaVUE-TA400 both L3 MPLS outer Header Stripping and L2 MPLS outer Header Stripping of traffic carrying Pseudowire MPLS control word (PWMCW) or Pseudowire associated channel header (PWACH) are supported. L2 MPLS Header Stripping is not supported on traffic without PWMCW or PWACH. E

Keep in mind the following when configuring GigaSMART operations with a **Strip Header** component:

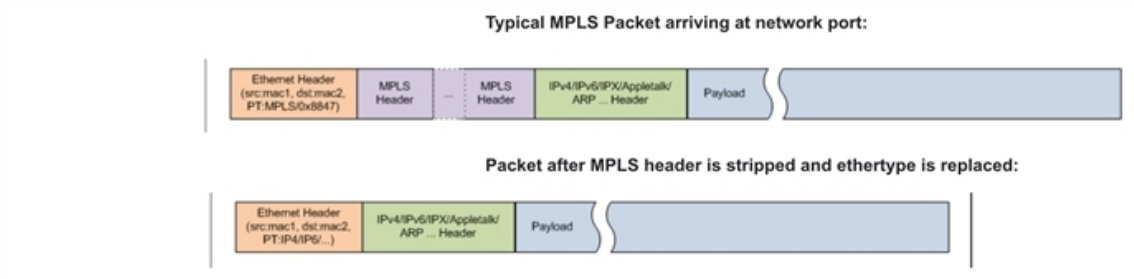
Summary	Description
ERSPAN Header Stripping	The ERSPAN header can be stripped. Specify an ERSPAN flow ID, from 0 to 1023. Use this option to strip ERSPAN Type II and Type III headers. A flow ID of zero is a wildcard value that matches all flow IDs.
Cisco FabricPath Header Stripping	The Cisco FabricPath headers can be stripped. The ability to decapsulate all packets with Cisco FabricPath headers; that is, all packets matching a destination switch ID and source switch ID. Also apply filters based on outer src/dst switch ID or ability to filter based on inner packet parameters with or without decapsulating the packet. The **Fabric Switch ID Source** and **Fabric Switch ID Destination** attributes are mandatory. Enter a value from 0 to 4095 (<0~(2^12-1)>) for a 12-bit switch ID. Enter 0 to strip all switch IDs.
FM6000 Timestamp Header Stripping	Packets entering GigaSMART from other devices may contain FM6000 timestamps. FM6000 is an Intel chip used for timestamping. The FM6000 timestamp can be stripped or it can be converted to UTC and appended to one of two Gigamon timestamping trailer formats. FM6000 has a hardware timestamp in the packet. For GigaSMART, the hardware timestamp needs to be translated into UTC time. An FM6000 device sends time mapping information in separate control packets called keyframes, which enable the UTC timestamp to be calculated. The calculated UTC timestamp can then be appended to the packets as a trailer. There are three timestamp formats: **None**, or **GigaSMART**, and **X12-TS** (for PRT-H00-X12TS). If the timestamp format is none, the FM6000 timestamp is stripped from the packet. If the timestamp format is **GigaSMART** or **X12-TS**, the FM6000 timestamp is stripped, converted to UTC, and a trailer containing the UTC timestamp is appended to the packets. The **GigaSMART** timestamp is added to the GigaSMART trailer. For the format of the GigaSMART trailer, refer to GigaSMART Trailer Reference. The **X12-TS** timestamp is added to the PRT-H00-X12TS trailer. For the format of the PRT-H00-X12TS trailer, refer to *GigaVUE-OS CLI Reference Guide*. **NOTES:** • FM6000 timestamp Header Stripping only supports an FM6000 timestamp that is present in the packet before the Frame Check Sequence (FCS). Packets containing an FM6000 timestamp that overwrites the existing FCS are discarded. Packets with a bad FCS are also discarded. • Keyframes must only be sent to GigaSMART from one FM6000 device. If there are multiple FM6000 devices, their clocks must be synchronized using PTP. • The GigaVUE-FM,GigaVUE HC Series node maintains the keyframe database per GigaSMART operation (gsop). There is a one-to-one mapping of the keyframe database to each map associated with an FM6000 gsop. Ten (10) maps with an FM6000 gsop are supported per GigaSMART engine. • The keyframe rate on the FM6000 device is important. If GigaSMART does not see the keyframe, the time that is converted and appended to the header will be a default start time (1969). In certain conditions, the keyframe rate on the FM6000 device might need to be increased so that GigaSMART does not miss seeing this frame.

Summary	Description
	• In H-Series Generation3 cards, FM6000 timestamping is not supported. For an FM6000 timestamping example, refer to Example – FM6000 Timestamping.
GRE Header Stripping	By specifying a GigaSMART Operation with a GSOP type of **Strip Header GRE**, the GigaSMART can strip GRE headers. It will automatically strip either Layer 3 or Layer 2 headers depending on the incoming packet. **Layer 3** – The GigaSMART can strip the outer IPv4/IPv6 delivery header and the GRE header to expose the encapsulated packet. Only IPv4 as the delivery protocol is supported in first and second generation cards. IPv4 and IPv6 address are supported in SMT-HC3-C08, SMT-HC1-S (third generation) cards. Any packet inside the GRE tunnel will be exposed, including IPv6 payloads. For an example, refer to Example – Stripping Layer 3 GRE IP Encapsulated Packets. **Layer 2** – The GigaSMART can strip GRE MPLS encapsulated and GRE Ethernet encapsulated packets, as follows: • GRE MPLS encapsulation – strip outer Ethernet header, outer IP header, GRE header, and MPLS header. • GRE Ethernet encapsulation (Transparent Ethernet Bridging) – strip outer Ethernet header, outer IP header, and GRE header. For an example, refer to Example – Stripping Layer 2 GRE Ethernet Encapsulated Packets. **Note**: The first-level **GRE Header Stripping** is supported on Gen 2 and Gen 3 devices.
Maximum MPLS Label Stack	The GigaSMART can strip MPLS headers up to a depth of seven labels.
Supported VLAN Types	The GigaSMART can strip both 802.1Q and Q-in-Q VLAN headers. Refer to How to Handle Q-in-Q Packets in Maps.
VXLAN Stripping	GigaSMART can strip VXLAN (Virtual eXtensible Local Area Network) headers. You can strip either matching VXLAN headers or all VXLAN headers. Select Strip Header from the GigaSMART Operation, select VXLAN for the protocol and use the following value in the Vxlan Id field: **0~(2^24-1)**. The VXLAN header is 8 bytes long with a 3-byte VXLAN Network Identifier (VNI) field. The VNI field is matched with the configured value and if it matches, the outer header (L2+IP+UDP+VXLAN) is stripped and the inner frame is sent to the tool. Specify a value of 0 to strip all outer VXLAN headers. **Note**: When processing packets with multiple encapsulation layers – for example, an ERSPAN-tunneled packet with a VXLAN tag – a VXLAN header-stripping operation strips all the way to the end of the VXLAN layer instead of just the VXLAN tag.
ISL Tunnel Stripping	ISL tunnel stripping removes the 26-byte header and the 4-byte FCS trailer associated with Cisco ISL VLAN encapsulation. **IMPORTANT:** Make sure the packets processed by a GigaSMART operation with a **Strip Header ISL** component are all using ISL encapsulation. GigaSMART operations do not distinguish between packets using ISL and packets that do not – it strips the requisite bytes from all packets it processes.
GTP Tunnel Stripping	GTP tunnel stripping removes the header and trailer for GTP-u packets inside

Summary	Description
	the GTP tunnel between the SGSN and GGSN interfaces in a 3G network, and between the eNodeB (eNb) and the SGW and between the SGW and the PGW in an LTE network. The SGSN and GGSN interfaces are also referred to as the Gn (or Gp) interface. The interface between eNb and SGW is referred to as S1U. The user plane interface between SGW and PGW is referred to as S5-U/S8-U. Both use GTPv1. Both GTPv1 and GTPv0 are supported for stripping. GTP-c control packets are not stripped. GTP[1] (also referred to as "GTP-Prime") is not supported for stripping.
Ethertype Replaced	After the VLAN/MPLS headers are stripped, the original ethertype carried in the Layer Two header is no longer valid. The GigaSMART replaces the ethertype field differently for MPLS and VLAN packets: **Ethertype Replacement for VLAN Packets** VLAN-tagged packets carry the original value for the ethertype field immediately after the VLAN tag. After the four-byte VLAN header is stripped, GigaSMART simply sets the ethertype field in the Layer 2 header to the value that was originally present in the packet past the VLAN tag. **Ethertype Replacement for MPLS Packets** Unlike VLAN-tagged packets, the Layer 3 protocol type is not carried in the packet for an MPLS packet – instead, it is applied by an egress router. To handle this, the GigaSMART examines the byte following the MPLS header to determine whether the packet is IPv4/IPv6 and takes the following actions: o **IPv4** – The GigaSMART replaces the ethertype from the MPLS packet with the IPv4 ethertype (0x0800) o **IPv6** – The GigaSMART replaces the ethertype from the MPLS packet with the IPv6 Ethertype (0x86DD). o **Non-IPv4/IPv6** – The GigaSMART passes the packet to destination tool ports without stripping the header. MPLS Header Stripping is only supported for IPv4/IPv6 packets.
CRCs Recalculated	The GigaVUE-FM,GigaVUE HC Series H Series node automatically recalculates and applies correct CRC checksums based on the new packet length after the header is stripped.
Viewing Statistics	From the device view, select **GigaSMART > GigaSMART Operations (GSOPS) > Statistics** to see statistics related to ongoing Header Stripping operations. Refer to View GigaSMART Statistics for more information.
Combine with Other Components	You can combine the **Strip Header** component with other GigaSMART components in a single operation. Refer to How to Combine GigaSMART Operations for details on the combinations of GigaSMART operations. Refer to Order of GigaSMART Operations for information on the order in which components of a single GigaSMART operation are applied.
GigaSMART Engine Ports	Header stripping operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to Groups of GigaSMART Engine Ports for details.
Generic	Use Generic Header Stripping to remove any arbitrary header from a packet by specifying the offset and the length of the header. For information about Generic Header Stripping, refer to GigaSMART Header Stripping.

Header Stripping Illustrated

The following figures illustrate how the header-stripping operations work – you can see the original MPLS packet with its label stack intact, followed by a stripped packet with a recalculated CRC and a new ethertype field.



Example – Stripping MPLS Headers and Adding a VLAN ID

The example shown in Figure 13Strip Header GigaSMART Operation for Stripping MPLS Headers illustrates a simple GigaSMART operation named **HeaderStrip** configured to strip MPLS headers and add a VLAN tag of 200. The operation is assigned to the GigaSMART group with the alias of gsGrp1.

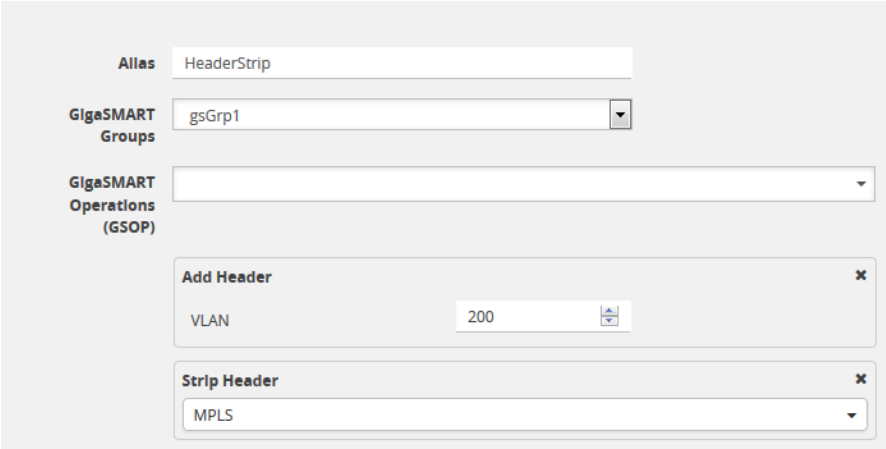


Figure 13 *Strip Header GigaSMART Operation for Stripping MPLS Headers*

The example shown in Figure 14Strip Header GigaSMART Operation for Stripping GTP Tunnel Informationillustrates a simple GigaSMART Header Stripping operation named **gtp-strip** configured to strip GTP tunnel information. The operation is performed by the GigaSMART group with the alias of GS1.

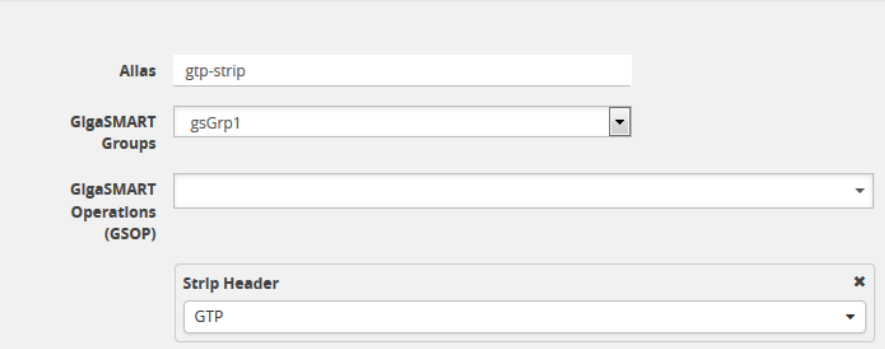
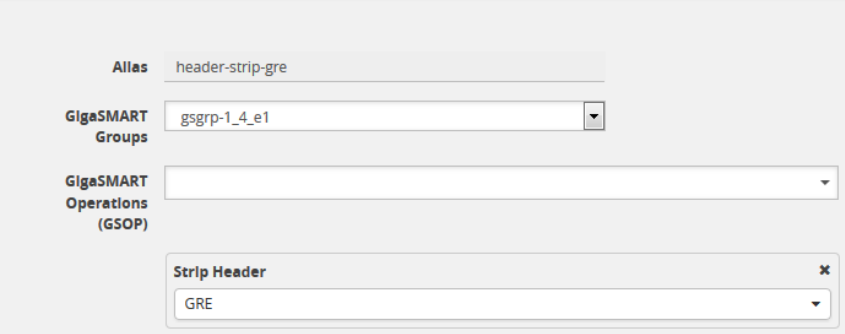


Figure 14 *Strip Header GigaSMART Operation for Stripping GTP Tunnel Information*

Example – Stripping Layer 3 GRE IP Encapsulated Packets

Use the configuration shown in the following figure to strip Layer 3 GRE IP encapsulated packets.

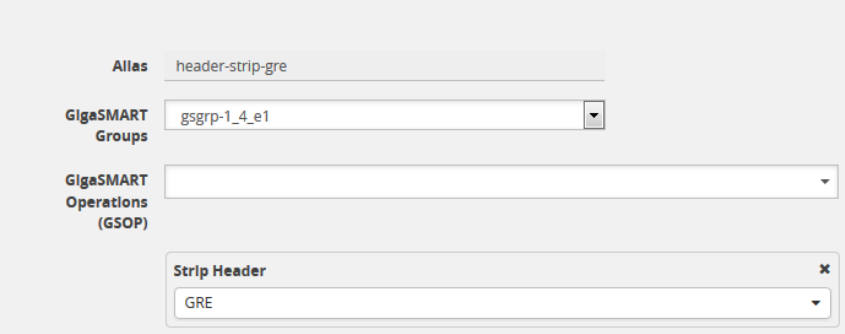


The following figure shows L3 GRE IP encapsulation before and after stripping.



Example – Stripping Layer 2 GRE Ethernet Encapsulated Packets

Use the configuration shown in the following figure to strip Layer 2 GRE Ethernet encapsulated (Transparent Ethernet Bridging) packets.



The following figure shows L2GRE Ethernet encapsulation before and after stripping.



Display Header Stripping Statistics

To display Header Stripping statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics**.

Refer to Header Stripping Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

Example – FM6000 Timestamping

Figure 15Strip Packets with FM6000 Timestamp shows an example GigaSMART Operation configured to strip packets containing the FM6000 timestamp:

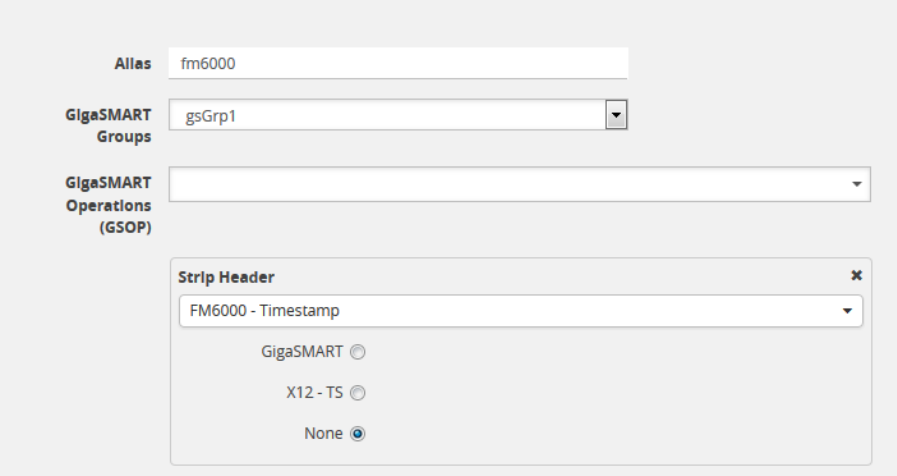


Figure 15 *Strip Packets with FM6000 Timestamp*

GigaSMART Operation with a Static Offset is an example GigaSMART Operation configured to convert packets containing the FM6000 timestamp to UTC and append the UTC timestamp to the Gigamon trailer:

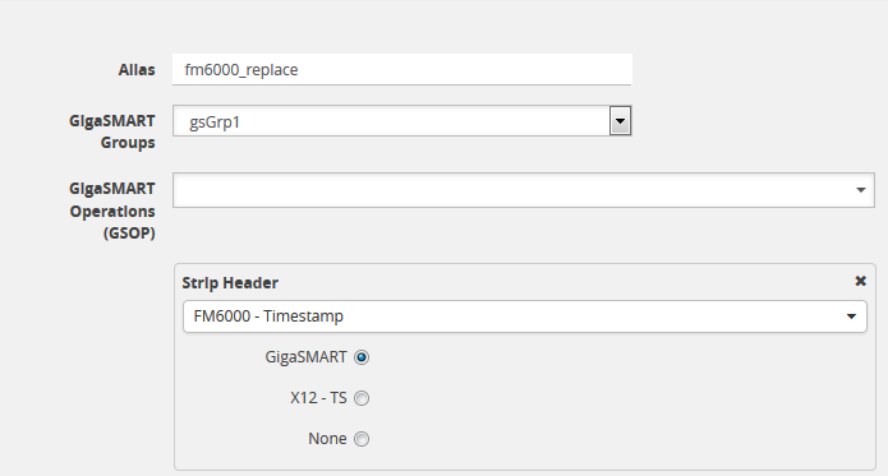


Figure 16 *Strip Packets and Append UTC Timestamp to the Gigamon Trailer*

Figure 17Map with Strip Header GigaSMART Operation is an example map using the strip header GigaSMART operation in Figure 16Strip Packets and Append UTC Timestamp to the Gigamon Trailer.

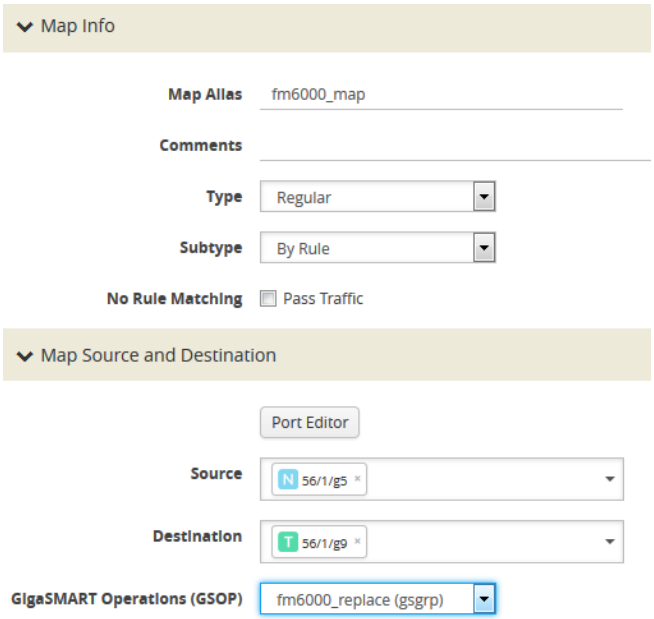


Figure 17 *Map with Strip Header GigaSMART Operation*

NOTE: There is one-to-one mapping between the GigaSMART Operation (gsop) and the map.

If there are multiple devices, each device can be configured with a different timestamp format. To configure this, use a different gsop and a different map for each device. For example, for packets arriving from FM6000 device1, configure a gsop for FM6000 device1 and associate it with map1. For packets arriving from FM6000 device2, configure a gsop for FM6000 device2 and associate it with map2.

All the maps can send all the packets to the same tool port.

Generic Header Stripping

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

GigaSMART operations use **Generic** Header Stripping to remove any arbitrary headers from a packet. The headers are stripped based on the offset and the length of the header.

To perform the generic Header Stripping operation:

1. Click **GigaSMART** on the left navigation pane. The GigaSMART Operation (GSOP) page is displayed.

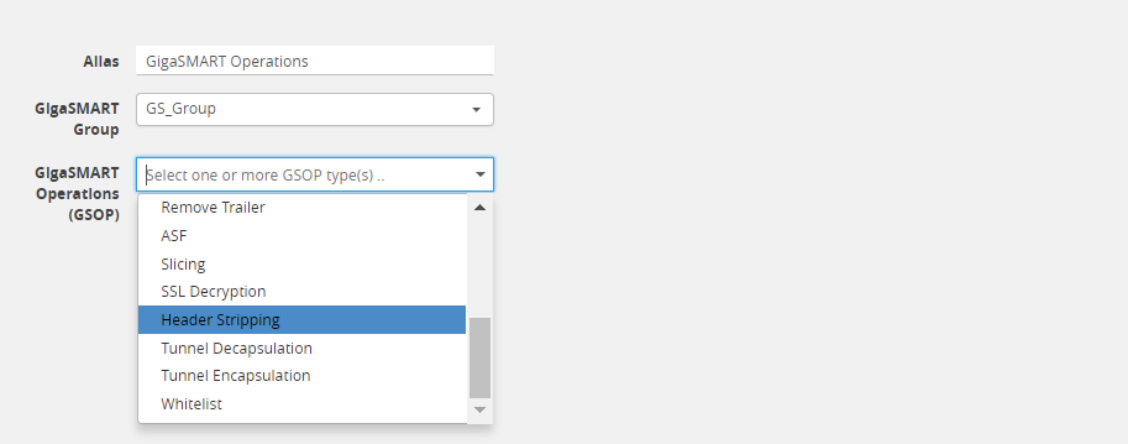


Figure 18 *GigaSMART Header Stripping*

2. In the **Alias** field, enter a name.
3. From the **GigaSMART group** drop-down list, select a GigaSMART group.
4. From the GigaSMART Operations (GSOP) drop-down list, select **Header Stripping**. A **Header Stripping** drop-down list is displayed.

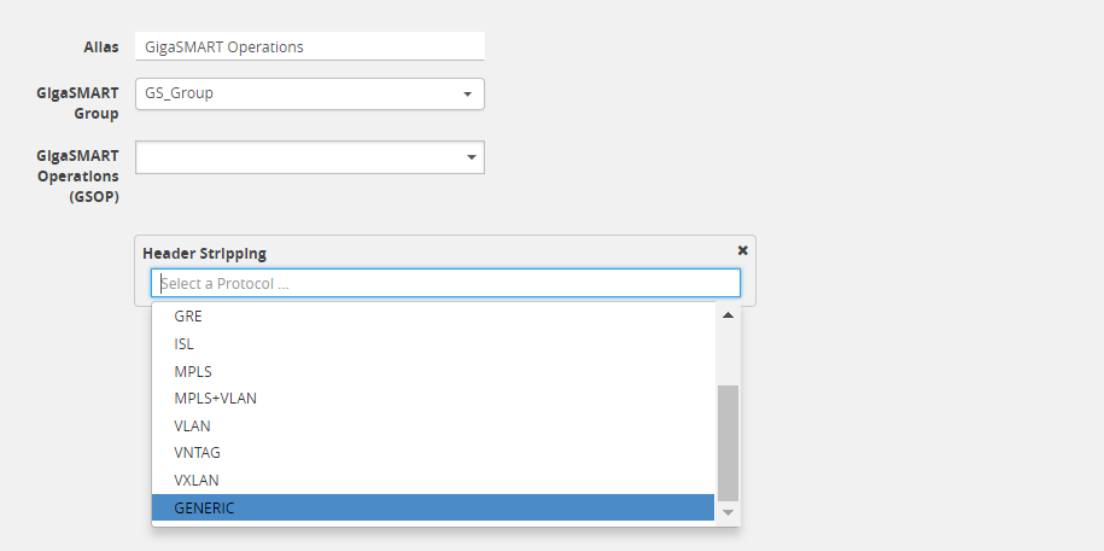


Figure 19 *GigaSMART Generic Header Stripping*

5. From the **Header Stripping** drop-down list, select **Generic**.

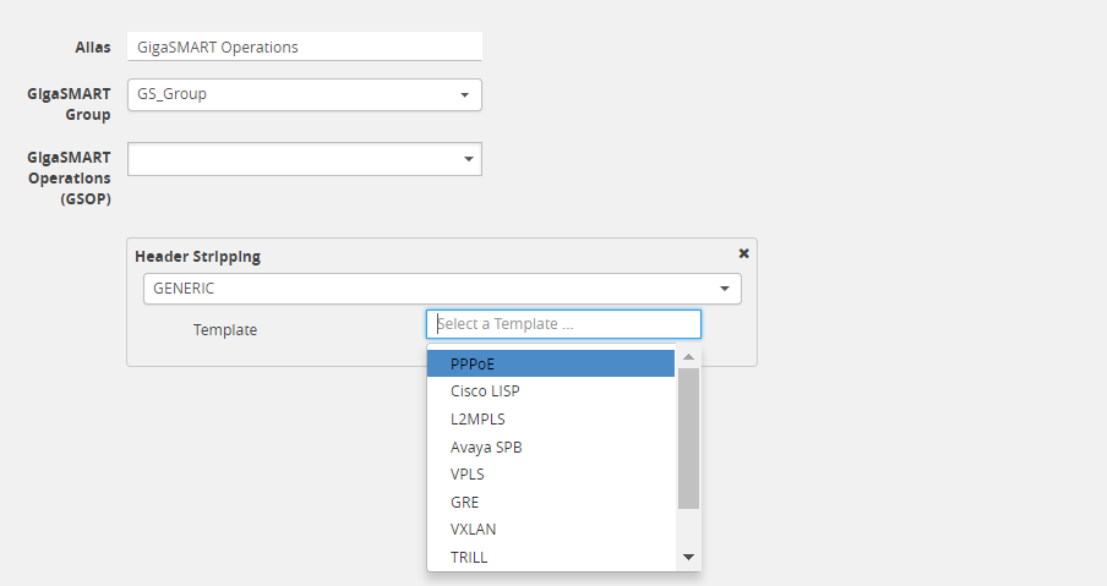


Figure 20 *Generic Header Stripping Template*

6. From the **Template** drop-down list, select one of the following options:

- Custom
- PPPoE
- Cisco LISP
- L2MPLS
- Avaya SPB
- VPLS
- GRE
- VXLAN
- TRILL
- Brocade VCS

Custom

A custom template lets you strip any arbitrary headers from a packet.

To strip any arbitrary header from a packet:

1. From the **Template** drop-down list, select **Custom.**

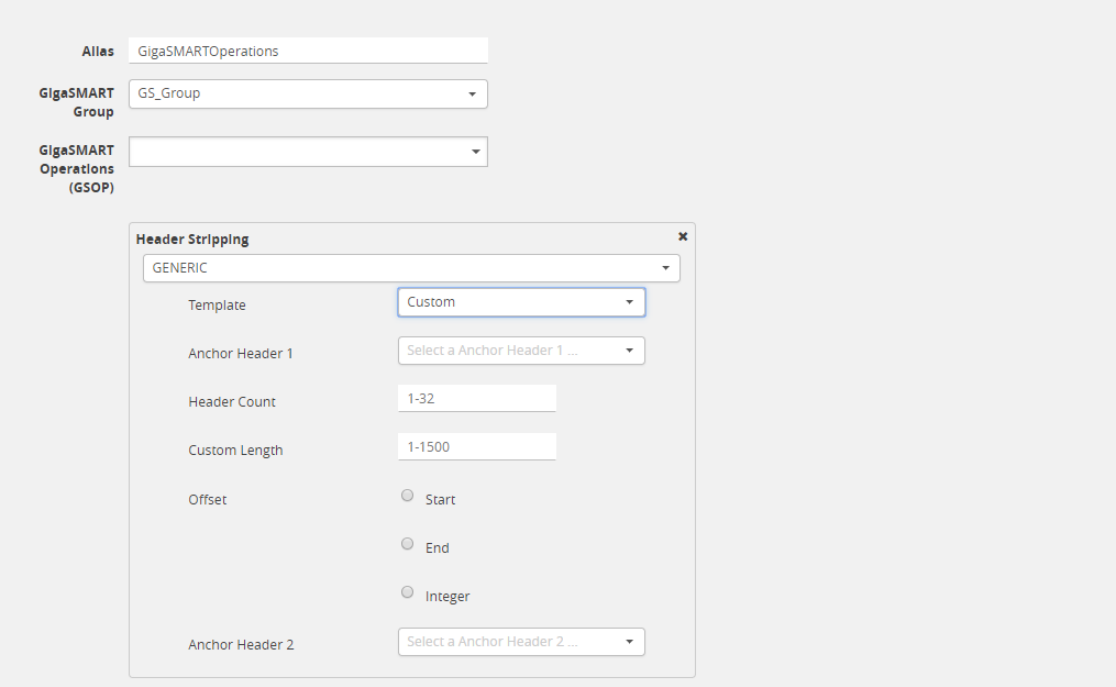


Figure 21 *Generic Header Stripping - Custom*

2. Select the following options to determine the headers to be stripped:

Component	Description
Anchor Header 1	Specifies the protocol from where GigaSMART should start stripping the header. The following values can be specified: o None o Eth o VLAN o MPLS o IPv4, and IPv6 The value **None** starts the Header Stripping operation from the start of the packet. If **None** is selected for Anchor Header 1, the Anchor Header 2 must also be set to **None.** The offset must not be set to **end**.
Offset	Specifies exactly from which end of Anchor Header 1 the stripping operation should start. You can specify the offset in terms of the following: **Start**—the Header Stripping operation starts from the left end of the Anchor Header 1. **End**—the Header Stripping operation starts from the right end of the Anchor Header 1. **Integer**—the Header Stripping operation starts from the specified integer offset of the Anchor Header 1. The integer value varies depending on the Anchor Header 1 specified.
Header Count	Specifies how many headers from the offset GigaSMART should remove. This is

Component	Description
	applicable when the packet headers are of known type. The known headers are as follows: • Ethernet • VLAN • MPLS • IPv4, and IPv6. It is important to note that, if **start** is specified for offset, the Anchor Header 1 is already counted for stripping. So, the value specified in the Header Count excludes the Anchor Header 1. If **None** is specified for Anchor Header 1, the Anchor Header 1 is not counted for stripping. So, the Header Count counts the Anchor Header 1 for stripping operation.
Custom Length	Specifies how many bytes of packet GigaSMART should remove. If the packet headers are unknown, the custom length of the unknown header can be specified to strip the packets. A combination of Header Count and Custom Length can also be used to strip the known and unknown headers. If Custom Length is specified, do not select **Any** for Anchor Header 2.
Anchor Header 2	Specifies the protocol that should become the next header after the stripping operation is complete. The following values can be specified: • None • Eth • VLAN • MPLS • IPv4 • IPv6 • TCP, UDP, and Any The value **Any** indicates that the next possible header can be any one of the options displayed for Anchor Header 2. The value **None** indicates that it is not necessary to mention the Anchor Header 2.

NOTE: Generic Header Stripping cannot strip unknown headers with variable length.

NOTE: On GigaVUE-TA400 both L3 MPLS outer Header Stripping and L2 MPLS outer Header Stripping of traffic carrying PWMCW (Pseudowire MPLS control word) or Pseudowire associated channel header (PWACH) are supported. L2 MPLS Header Stripping is not supported on traffic without PWMCW or PWACH.

PPPoE

The PPPoE template lets you strip the PPPoE encapsulated packets from the packet structure. Figure 22PPPoE Encapsulated Packets illustrates a red outline around the frame that needs to be striped.



Figure 22 *PPPoE Encapsulated Packets*

To strip the PPoE encapsulated packets:

1. From the **Template** drop-down list, select **PPPoE**. Refer to Figure 23Generic Header Stripping - PPPoE.

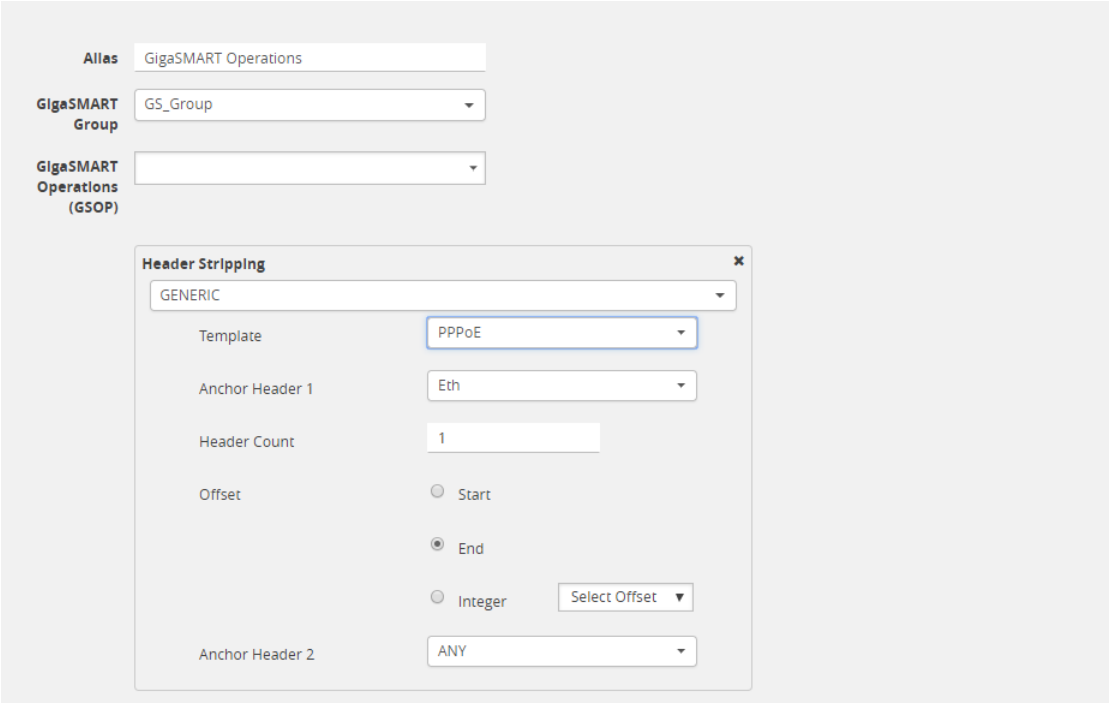


Figure 23 *Generic Header Stripping - PPPoE*

By default, the following values are selected:

Field	Value	Description
Anchor Header 1 **Offset**	Eth End	Starts the Header Stripping operation from the right end of the Ethernet header.
Header Count	1	Strips the header next to the Ethernet Header.

Field	Value	Description
Anchor Header 2	Any	Updates a valid protocol as the Anchor Header 2 in the packet. In this case, any IPv4 or IPv6 protocol can become the Anchor Header 2.

2. Click **OK**. The Header Stripping operation is displayed in the GigaSMART Operations (GSOP) page.

Cisco LISP

Cisco LISP is used to carry original IP packets to support multi-homing. In this example, the IPv4 outer header, UDP header, and LISP header are stripped from the Cisco LISP header format. The LISP header is considered as an unknown header.

Figure 24Cisco LISP Encapsulated Packets illustrates a red outline around the frame that needs to be striped.



Figure 24 *Cisco LISP Encapsulated Packets*

To strip the Cisco LISP encapsulated packets:

1. From the **Template** drop-down list, select **Cisco LISP**.

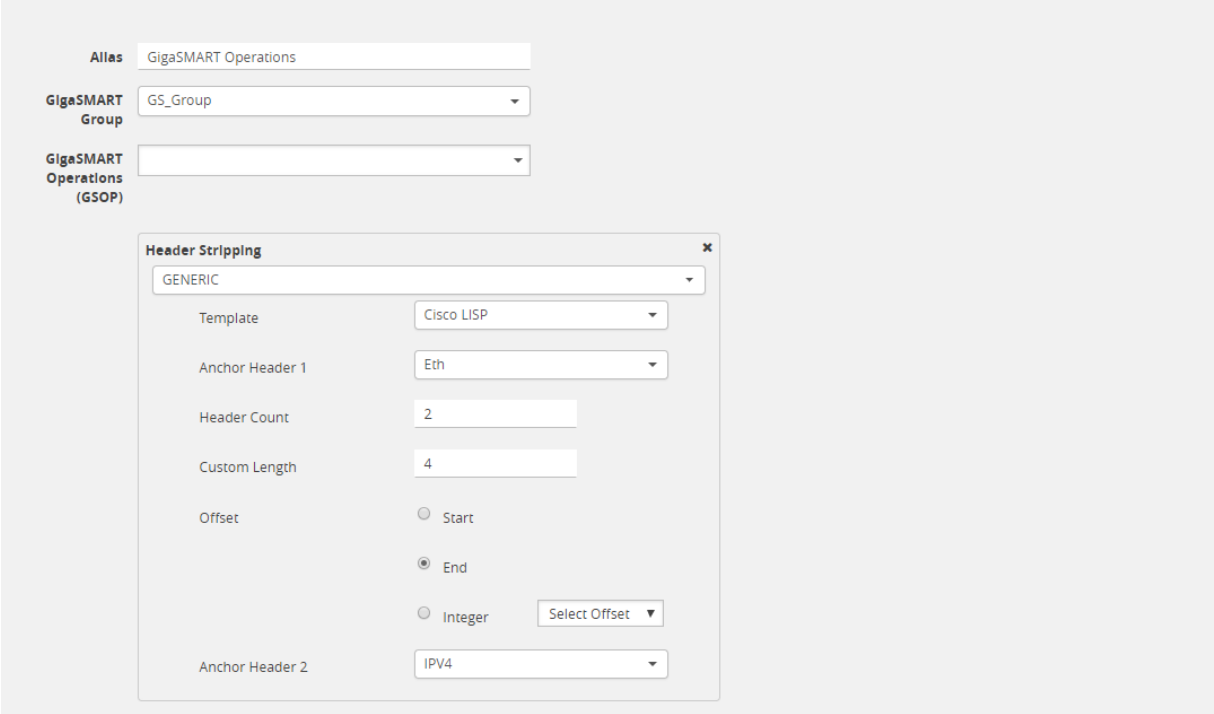


Figure 25 *Generic Header Stripping - Cisco LISP*

By default, the following values are selected:

Field	Value	Description
Anchor Header 1 **Offset**	Eth End	Starts the Header Stripping operation from the right end of the Ethernet header.
Header Count	2	Strips the next two headers— IPv4 Outer Header and UDP from the packet.
Custom Length	8	Strips 8 bytes of the unknown packet header. LISP is an unknown header.
Anchor Header 2	IPv4	Updates IPv4 protocol as the Anchor Header 2 in the packet.

2. Click **OK**. The Header Stripping operation is displayed in the GigaSMART Operations (GSOP) page.

L2 MPLS

The L2 MPLS packet, also known as VPLS, encapsulates Ethernet packets in the MPLS label stack. In this example, the outer Ethernet header and MPLS [PW Label] are stripped from the L2 MPLS encapsulated packets.

Figure 26L2 MPLS Encapsulated Packets illustrates a red outline around the frame that needs to be striped.



Figure 26 *L2 MPLS Encapsulated Packets*

To strip the outer MAC header from the L2 MPLS encapsulated packets:

1. From the **Template** drop-down list, select **L2MPLS**.

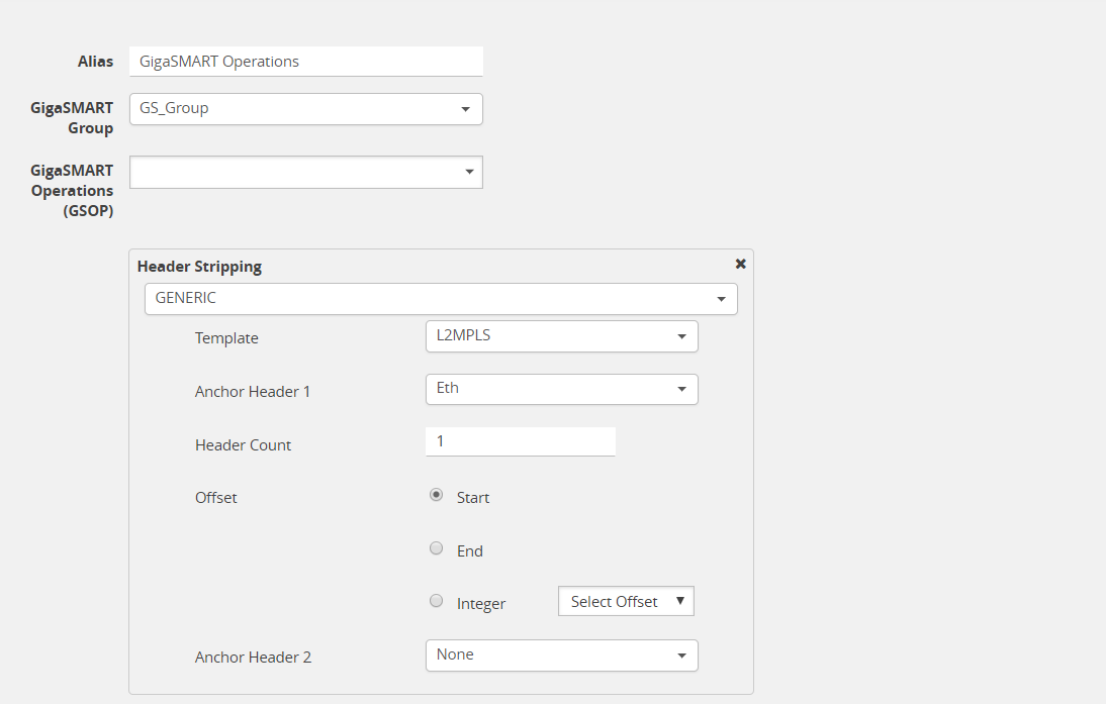


Figure 27 *Generic Header Stripping - L2 MPLS*

By default, the following values are selected:

Field	Value	Description
Anchor Header 1 **Offset**	None Start	Starts the Header Stripping operation from the start of the Ethernet header.
Header Count	2	Strips the first and the second header from the packet. The outer Ethernet header and MPLS [PW label] packet header are both removed. As Anchor Header 1 is set to none, the Header Count counts the first header for stripping.
Anchor Header 2	None	Signifies that there is no need to specify the Anchor Header 2. In this case, the IPv4 protocol forms the first header of the packet after the stripping operation is complete.

2. Click **OK**. The Header Stripping operation is displayed in the GigaSMART Operations (GSOP) page.

VXLAN

VXLAN encapsulates Ethernet packets in IP using VXLAN header. In this example, the outer Ethernet header, outer IP header, outer UDP header, and VXLAN Header are stripped from the VXLAN encapsulated packets.

Figure 28VXLAN Encapsulated Packets illustrates a red outline around the frame that needs to be striped.



Figure 28 *VXLAN Encapsulated Packets*

To strip the outer Ethernet frame from the VXLAN encapsulated packets:

1. From the **Template** drop-down list, select **VXLAN**.

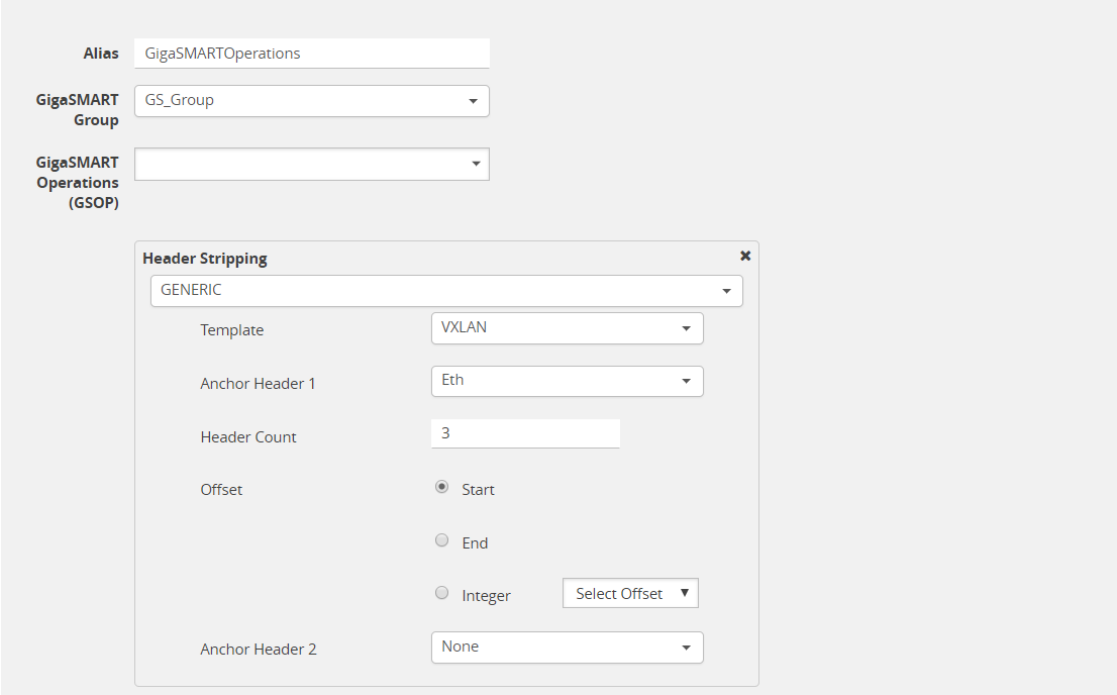


Figure 29 *Generic Header Stripping - VXLAN*

By default, the following values are selected:

Field	Value	Description
Anchor Header 1 **Offset**	Eth Start	Starts the Header Stripping operation from the start of the Ethernet header.
Header Count	3	Strips the next three headers—outer IP header,

Field	Value	Description
		outer UDP header, and VXLAN header.
Anchor Header 2	None	Signifies that there is no need to specify the Anchor Header 2. In this case, the IPv4 protocol forms the first header of the packet. **Note**: When the Anchor Header 1 is set to **None**, the Anchor Header 2 must also be set to **None**.

2. Click **OK**. The Header Stripping operation is displayed in the GigaSMART Operations (GSOP) page.

TRILL

TRILL encapsulates Ethernet packets in Ethernet frame to provide L2 layer routing in data centers. In this example, consider TRILL as an unknown header. This TRILL frame is stripped with the inner Ethernet header from the encapsulated packets. The combined length of TRILL header (6 bytes) and inner Ethernet header (14 bytes) is 20 bytes.

Figure 30TRILL Encapsulated Packets illustrates a red outline around the frame that needs to be striped.



Figure 30 *TRILL Encapsulated Packets*

To strip TRILL from the encapsulated packets:

1. From the **Template** drop-down list, select **TRILL**.

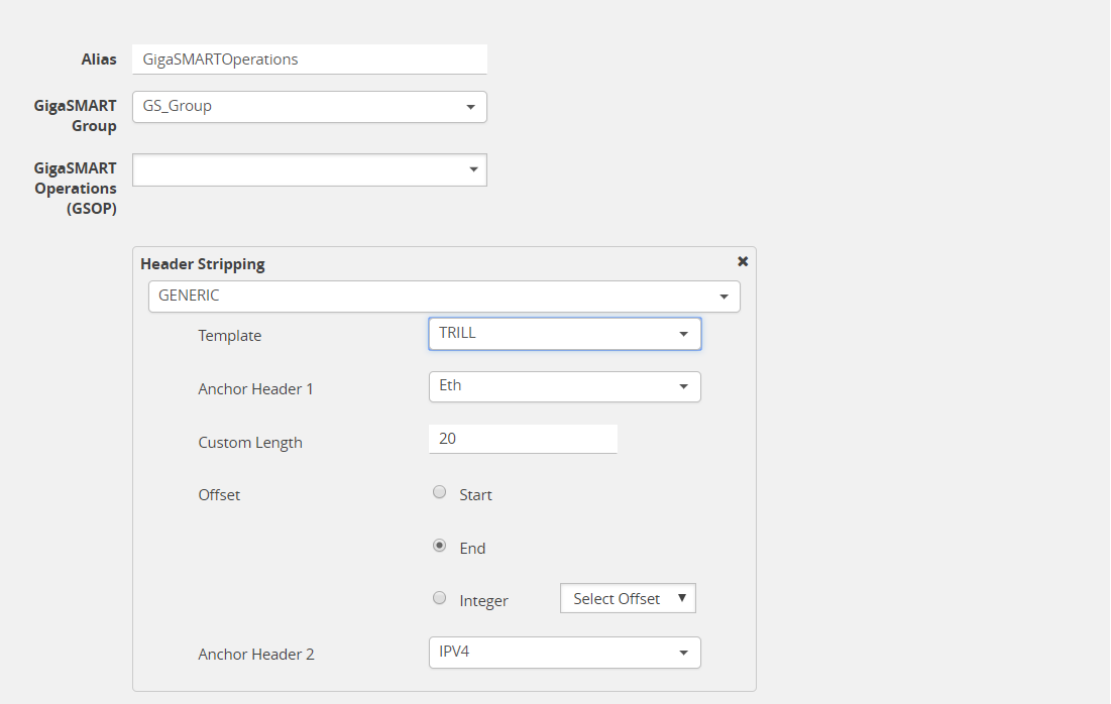


Figure 31 *Generic Header Stripping - TRILL*

By default, the following values are selected:

Field	Value	Description
Anchor Header 1 **Offset**	TRILL End	Starts the Header Stripping operation from the right end of the outer Ethernet header.
Custom Length	20	Strips 20 bytes of unknown header from the packets. In this case, the TRILL and the inner Ethernet headers are stripped.
Anchor Header 2	IPv4	Updates IPv4 protocol as the second header in the packet.

2. Click **OK**. The Header Stripping operation is displayed in the GigaSMART Operations (GSOP) page.

Avaya SPB

Avaya SPB (802.1ah) fabric encapsulates Ethernet packets using MAC-In-MAC headers. In this example, the outer Ethernet header and ITAG are removed from the packet structure.

Figure 32Avaya SPB Encapsulated Packets illustrates a red outline around the frame that needs to be striped.

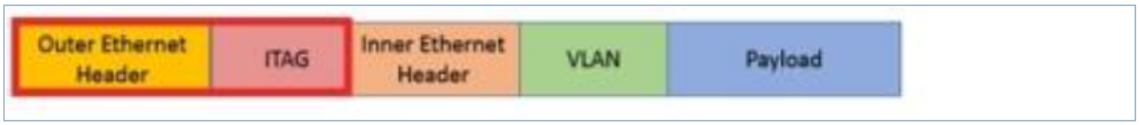


Figure 32 *Avaya SPB Encapsulated Packets*

To strip the outer Ethernet headers from the Avaya SPB encapsulated packets:

1. From the **Template** drop-down list, select **Avaya SPB**.

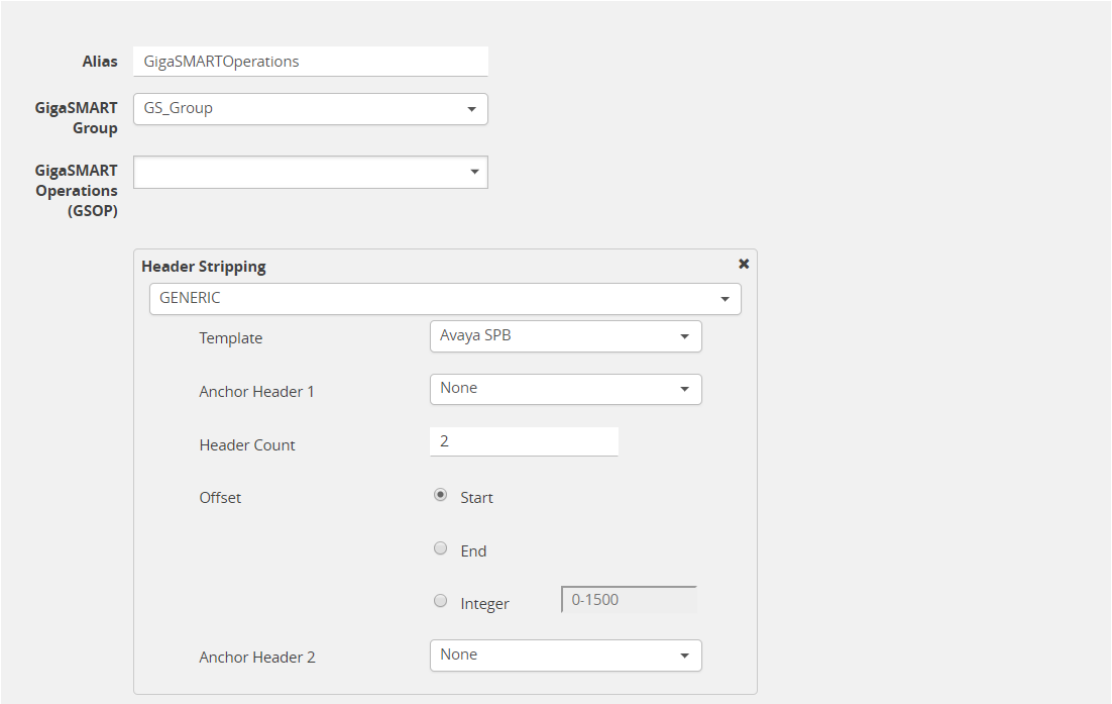


Figure 33 *Generic Header Stripping - Avaya SPB*

By default, the following values are selected:

Field	Value	Description
Anchor Header 1 **Offset**	None Start	Starts the Header Stripping operation from the left end of the outer Ethernet header.
Header Count	2	Strips the outer Ethernet and ITAG headers from the packet.
Anchor Header 2	None	Signifies that it is not necessary to specify the next header. The inner Ethernet header becomes the first header after the stripping operation is complete.

2. Click **OK**. The Header Stripping operation is displayed in the GigaSMART Operations (GSOP) page.

You can also strip the ITAG, inner Ethernet header, and VLAN from the packet structure.

Figure 34Avaya SPB Encapsulated Packets illustrates a red outline around the frame that needs to be striped.



Figure 34 *Avaya SPB Encapsulated Packets*

To strip the inner Ethernet headers from the encapsulated packets:

1. Select the following values to strip the inner Ethernet headers from the encapsulated packets:

Field	Value	Description
Anchor Header 1 Offset	Eth End	Starts the Header Stripping operation from the right end of the outer Ethernet header.
Header Count	3	Strips the ITAG, inner Ethernet, and VLAN headers from the packet.
Anchor Header 2	Any	Indicates that any valid protocol available after the Header Stripping operation can become the next header in the packet.

2. Click **OK**. The Header Stripping operation is displayed in the GigaSMART Operations (GSOP) page.

GigaSMART Header Addition

Required License: Header Stripping

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

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

GigaSMART operations with an **Add Header** selected can add VLAN tags to packets. This operation is useful in the following situations:

- Differentiating stripped packets from non-stripped packets on common IP ranges (for example,10.x.x.x; 192.168.x.x).
- Removing an arbitrary-length MPLS label stack and replacing it with a single, predictable, four-byte VLAN tag between the source address and ethertype field in the Layer 2 header. Many tools that are unable to parse the arbitrary length of an MPLS label stack can work with a predictable VLAN tag.

Keep in mind the following when configuring GigaSMART operations with an **Add Header** component:

Add VLAN Tag	You can combine **Strip Header** with **VLAN add** to help identify packets with stripped headers. This approach lets you remove an arbitrary-length MPLS label stack and replace it with a single, predictable, four-byte VLAN tag between the source address and ethertype field in the Layer 2 header. Many tools that are unable to parse the arbitrary length of an MPLS label stack can work with a predictable VLAN tag. **Packet Modifications for add_vlan** The **Add Header** operation makes the following modifications to a packet: o **TPID –** 0x8100 (802.1Q VLAN) or 0x88A8 and 0x9100 (Q-in-Q). The two-byte ethertype originally present in the Ethernet header is moved past the new VLAN header to identify the original Layer 3 header. o **CFI** – 0 o **Priority** – 0 o **VLAN ID** – User-provided value in the **VLAN** field of an **Add Header** GigaSMART Operation. Refer to How to Handle Q-in-Q Packets in Maps for TPID.
CRCs Recalculated	The GigaVUE-FM,GigaVUE HC Series H Series node automatically recalculates and applies correct CRC checksums based on the new packet length after the header is stripped.
Viewing Statistics	Use **GS Operations Statistics** page to see statistics related to ongoing GigaSMART operations. Refer to View GigaSMART Statistics for more information.
Combine with Other Components	You can combine the **Add Header** component with other GigaSMART components in a single operation. Refer to How to Combine GigaSMART Operations for details on the combinations of GigaSMART operations. Refer to Order of GigaSMART Operations for information on the order in which components of a single GigaSMART operation are applied.
GigaSMART Engine Ports	Header addition operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to Groups of GigaSMART Engine Ports for details.

GigaSMART Masking

Required License: Base

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

GigaSMART operations with **Masking** selected write over a specific portion of a packet with a specified one-byte pattern. Masking operations consist of an **offset**, **length**, and **pattern** as shown in Figure 35GigaSMART Operations Page with Masking Selected.

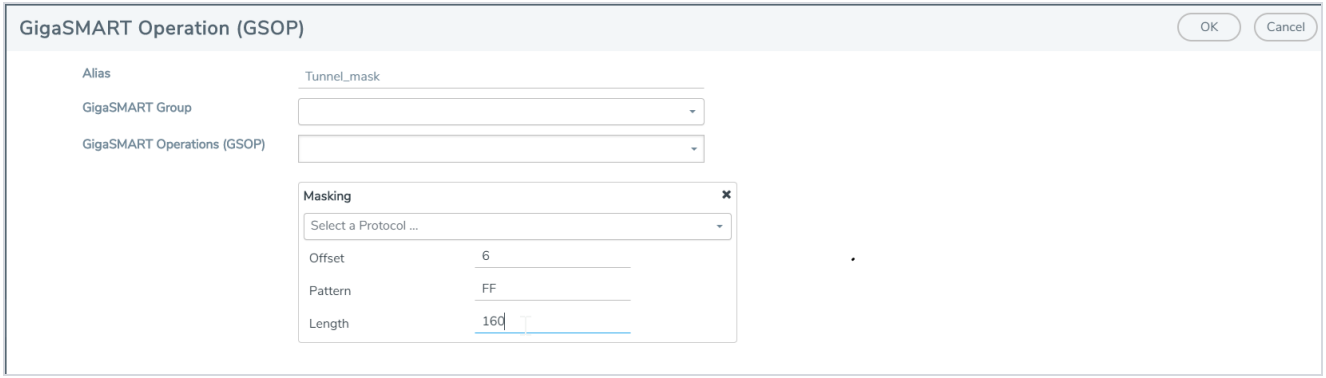


Figure 35 *GigaSMART Operations Page with Masking Selected*

The following table describes the fields.

Component	Description
Offset	Specifies **where** GigaSMART should start masking data with the supplied pattern. You can specify this in terms of either a static offset from the start of the packet or a relative offset from a particular protocol layer. This lets you automatically compensate for variable length headers, specifying a mask target in terms of a particular packet header.
Length	Specifies **how much** of the packet GigaSMART should mask. The specified one-byte pattern can be repeated to mask from 1-9600 bytes.
Pattern	Specifies **what** pattern GigaSMART should use to mask the specified portion of the packet. You can specify a one-byte hex pattern (for example, 0xFF).

Refer to the following when configuring GigaSMART operations with a **Masking** component:

Feature	Description
Protocol	The following are the protocols that you can select for from the protocol drop-down list: o **IPV4** – Mask starting a specified number of bytes after the IPv4 header.

Feature	Description
	○ **IPV6** – Mask starting a specified number of bytes after the IPv6 header. ○ **UDP** – Mask starting a specified number of bytes after the UDP header. ○ **TCP** – Mask starting a specified number of bytes after the TCP header. ○ **FTP**– Identify using TCP port 20. Mask payloads using offset from the TCP header. ○ **https** – Identify using TCP port 443. Mask payloads using offset from the TCP header. ○ **SSH** – Identify using TCP port 22. Mask payloads using offset from the TCP header. The GigaSMART-HC0 module can provide masking for GTP tunnels, provided the user payloads are unencrypted. Both GTPv1 and GTPv2 are supported – GTP' (GTP prime) is not supported. Keep in mind that only GTP-u (user plane packets) are masked. Control plane packets (GTP-c) are left unmodified. ○ **GTP** – Mask starting a specified number of bytes after the outer GTP header. ○ **GTP-IPV4** – Mask starting a specified number of bytes after the IPv4 header inside the encapsulating GTP packet. ○ **GTP-UDP** – Mask starting a specified number of bytes after the UDP header inside the encapsulating GTP packet. ○ **GTP-TCP** – Mask starting a specified number of bytes after the TCP header inside the encapsulating GTP packet.
Masking Offset and Length	You can specify either a **relative** offset or a **static** offset for the masking pattern: • **Static** offsets begin masking a specific number of bytes from the start of the packet. Choose a static offset by setting **Protocol** to **None** and supplying an **Offset** from **<0~9000>** bytes. Zero (0) indicates the start of the Ethernet frame. • **Relative** offsets begin masking a specified number of bytes from the end of a particular header – IPv4, IPv6, and so on. Choose a relative offset by selecting any of the following values for the **protocol** argument and supplying an offset from the specified protocol header of **<1~9000>** bytes: **Note**: You can only mask one contiguous portion of a packet.
Recalculated CRC	GigaSMART automatically calculates a new Ethernet CRC based on the masked packet's length and data, and uses it to replace the existing CRC. This way, analysis tools do not report CRC errors for masked packets. **Note**: IP or UDP checksum is not recalculated if masking is done on the existing IP or UDP checksum.
GigaSMART Trailer	Masking operations can optionally include the GigaSMART Trailer. If you do elect to include the GigaSMART Trailer, it will include the original packet length. Refer to How to Use GigaSMART Trailers for details in the GigaSMART Operations Reference Guide.
In Combination with Slicing	Masking operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to Groups of GigaSMART Engine Ports for details in the GigaSMART Operations Reference Guide.

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

Examples – GigaSMART Masking

The example shown in Figure 36GigaSMART Masking Operation creates a GigaSMART masking operation named **Tunnel_mask**. This example starts masking six bytes after the end of the TCP layer in the GTP-encapsulated packet and continues for 150 bytes, writing over the existing data with an FF pattern.

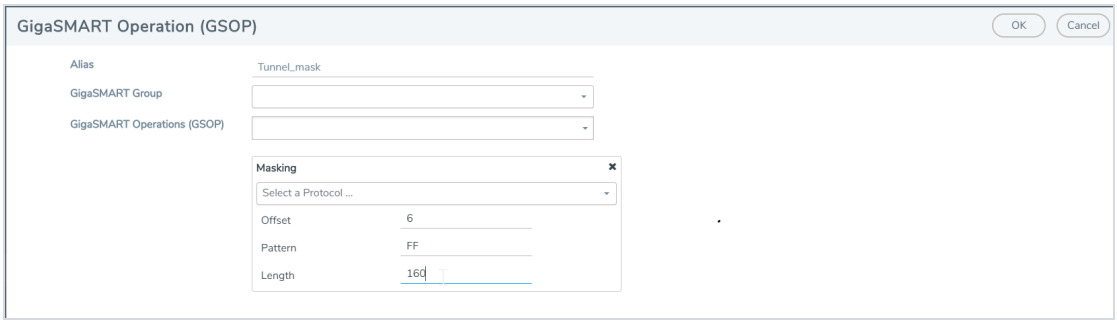


Figure 36 *GigaSMART Masking Operation*

This example shown in Figure 37GigaSMART Operation with a Static Offset creates a GigaSMART masking operation named **Mask_FIX**. This example uses a static masking offset of 148 bytes and continues for the next 81 bytes, writing over the existing data with an **FF** pattern. This GigaSMART operation is assigned to the GigaSMART group with the alias of gs2port2.

This example simulates how to mask a FIX (Financial Information eXchange) packet so that generic information is preserved at the start and end of the FIX data portion of the packet while private information within is masked. This example does not include the optional GigaSMART Trailer.

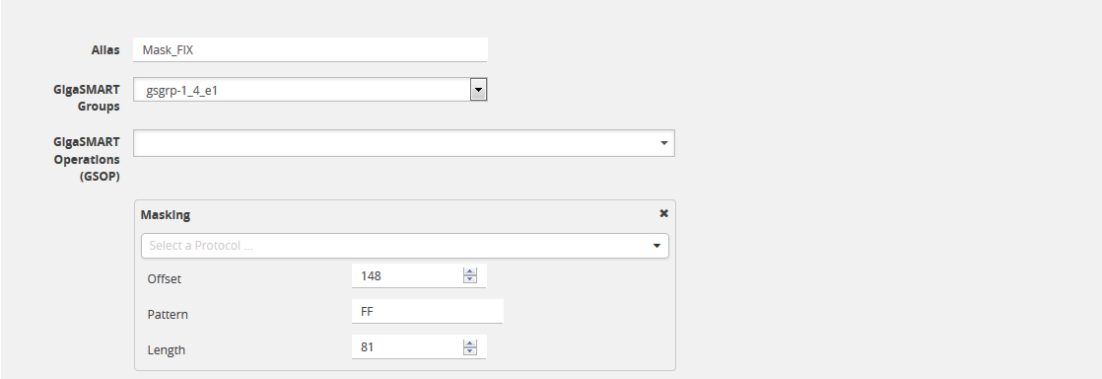


Figure 37 *GigaSMART Operation with a Static Offset*

Display Masking Statistics

To display masking statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics**. The statistics for masking will be in the row labeled Masking in the GS Operations column.

Refer to Masking Statistics Definitions for descriptions of the masking statistics as well as to GigaSMART Operations Statistics Definitions.

GigaSMART NetFlow Generation

Required License: NetFlow GenerationRequired License for NetFlow with Second Level Maps: Adaptive Packet Filtering (APF)

Supported Devices : GigaVUE-HC1 Gen 2, GigaVUE-HC3 Gen 2.

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

NetFlow Generation is a simple and effective way to increase visibility into traffic flows and usage patterns across systems. The flow-generated data can be used to build relationships and usage patterns between nodes on the network. Routers and switches that support NetFlow can collect IP traffic statistics to be exported as NetFlow records.

However, the processor and memory load of enabling NetFlow can cause service degradation and affect their ability to pass traffic without introducing latency and packet drops. Due to this processing overhead, sampled NetFlow is implemented in most of the high-end routers. Sampling in every “N” packets for NetFlow processing can severely limit the visibility needed to monitor flows.

The advanced capabilities of GigaSMART® technology can be leveraged to summarize and generate unsampled NetFlow statistics from incoming traffic streams. Offloading NetFlow Generation to an out-of-band solution like the Gigamon Deep Observability Pipeline completely eliminates the risk of using core production network resources in generating this data. Combined with the flexibility offered by Gigamon patented Flow Mapping® technology, operators can pick and choose from which flows to generate NetFlow statistics, while at the same time sending the original packets to other monitoring tools.

Support for NetFlow versions 5 and 9 and IP Information Export (IPFIX), as well as CEF, enables seamless integration with standards-based collectors. NetFlow records can also be exported to multiple collectors concurrently, providing a single flow source for business-critical management applications such as security, billing, and capacity planning. Exported flows can also be filtered so that collectors only receive the specific records relevant to them.

NOTE: Legacy NetFlow supports only one NetFlow version (v5, v9 or IPFIX) record and NetFlow exporter format version per engine unless the exporter format is CEF.

Gigamon has also extended IPFIX to include URL information, providing insight into HTTP and SIP traffic. Other enterprise extensions for IPFIX are HTTP, DNS, and SSL certificates, which provide metadata that can be used for security analysis.

Additionally, Gigamon's Deep Observability Pipeline architecture is the first in the industry to summarize flow statistics as well as to provide the flexibility of aggregating, replicating, filtering, and forwarding raw traffic streams to monitoring tools for detailed troubleshooting and analytics.

The Gigamon Deep Observability Pipeline establishes a scalable framework to deliver pervasive flow-level visibility across enterprises, data centers, and service provider environments to accurately design, engineer, optimize, and manage their network infrastructure.

NOTE: NetFlow Generation exports records using IPv4. IPv6 is not supported.

GigaSMART operations with a NetFlow component can be assigned to multiple GigaSMART groups or GigaSMART groups consisting of multiple GigaSMART engine ports.

NetFlow/IPFIX Generation is a pillar of the GigaSECURE Security Delivery Platform.

NetFlow Generation is displayed in Figure 38NetFlow Generation Gigamon Solution.

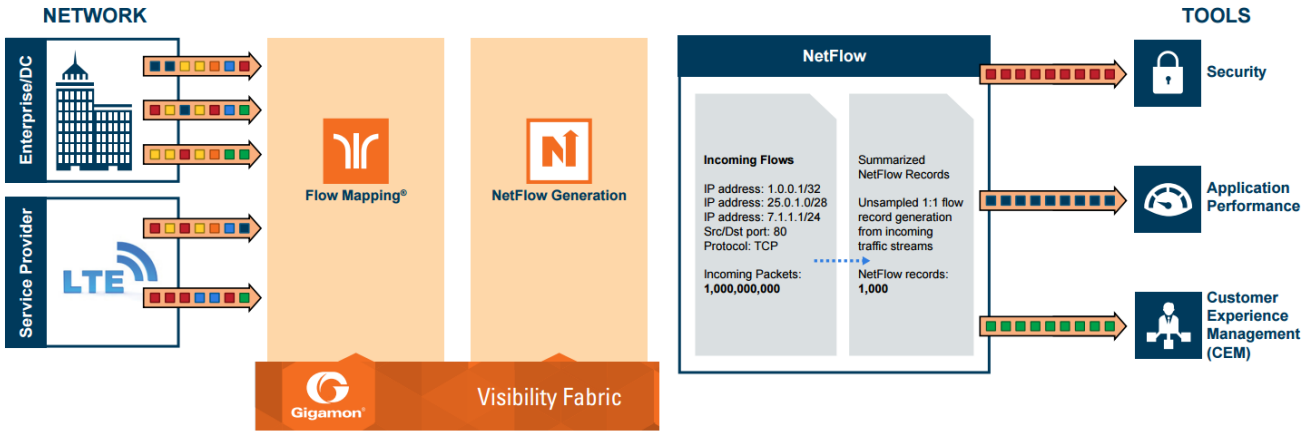


Figure 38 *NetFlow Generation Gigamon Solution*

In Figure 38NetFlow Generation Gigamon Solution, incoming packets from network(s) enter the Gigamon Deep Observability Pipeline and are directed by maps to NetFlow.

NetFlow examines the incoming packets and converts the packets of choice into flows records. Specific flows are then forwarded to specific tools, such as Security, Application Performance, and Customer Experience Management (CEM) tools.

Active Timeout

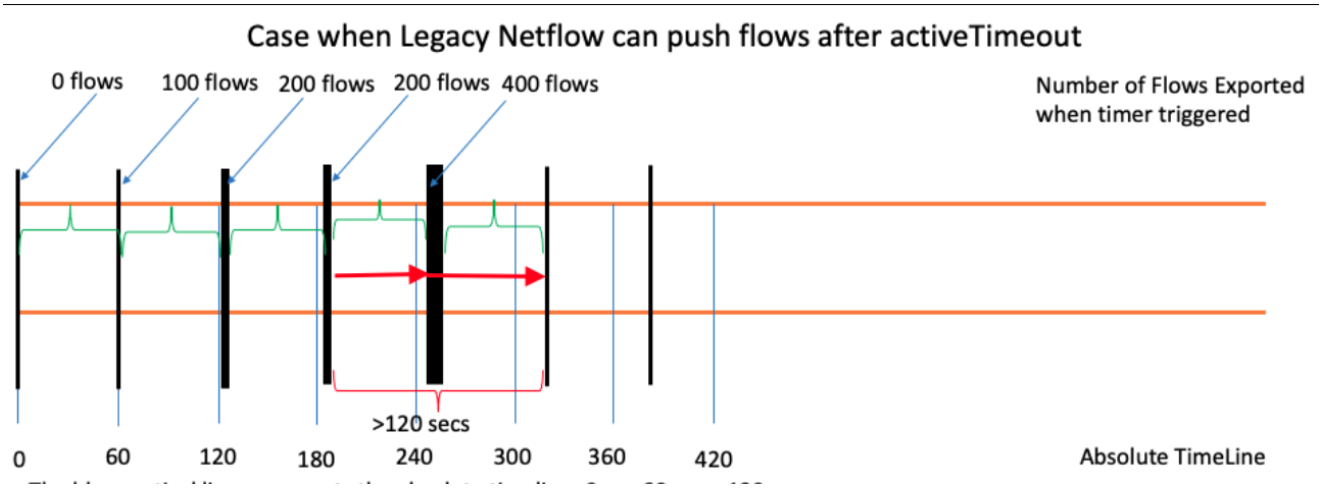
When a flow is active (GigaSMART engine receives packet) and sends packets for more than delta seconds, it hits an active timeout of delta seconds.

Inactive Timeout

When a flow stops sending packets for more than delta seconds, it hits inactive timeout of delta.

However, there can be some flows which are not exported at the end of active/inactive timeout.

The following diagram shows an example when Legacy NetFlow can push flows after active timeout.



In the above diagram, the blue vertical line represents the absolute time. The black vertical line represents the time at which GigaSMART engine starts exporting data after inactive time out. The green curly brackets represent the active timeout. The width of the black line represents the time taken by the exporter to push all the existing flows.

After a black line ends, it takes 60 seconds for another black line to start. This 60 second gap is represented by the green curly brackets. To ensure that each flow could be sent only once, there is only one vertical black line at any point of time.

In the above example, the red arrow represent a new flow in the network. When the flow starts and reaches the first black line, the GigaSMART engine calculates if the flow has reached the active time out. Here, since the flow has not reached the active Time out, the GigaSMART engine does not export the data for the flow.

The time consumed by the export process depends on the number of the flows that are being exported. The export process restarts after active timeout (the second black line). When the flow reaches the second black line, GigaSMART engine exports data for that flow.

NOTE: The difference between the time taken at which data is first exported for the flow and the flow start time is greater than the active timeout.

Hence, it is possible for the GigaSMART engine to rarely consume more time to export data of active flows.

NOTE: For inactive timeout also the Legacy NetFlow pushes the flows in same manner as explained in the example.

NetFlow Generation Components

NetFlow Generation collects IP traffic statistics on all interfaces where a NetFlow Monitor is enabled. It then gathers the statistics of the traffic flows and exports the NetFlow records to at least one NetFlow collector (typically a device that performs the actual traffic analysis based on the information from the NetFlow records).

Figure 39NetFlow Generation Components shows the NetFlow Generation components.

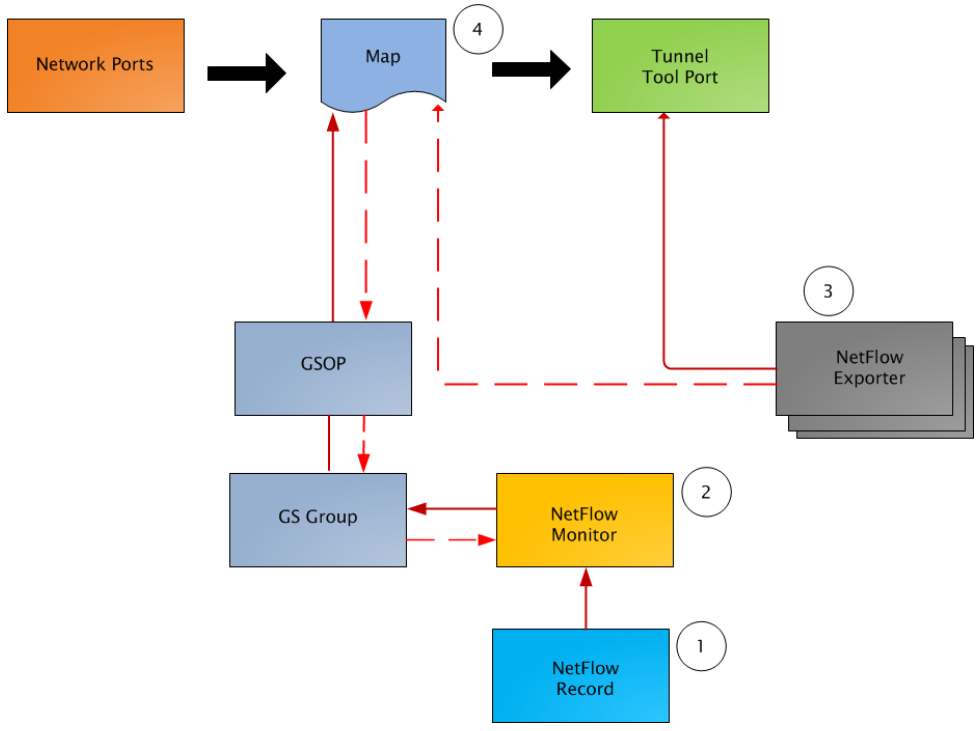


Figure 39 *NetFlow Generation Components*

GigaSMART NetFlow Generation illustrates the NetFlow Generation and how its components are associated. The NetFlow Generation associates its components in the following order:

1. One or more Records are associated to the Monitor.
2. The Monitor is associated to the GigaSMART group.
3. The Exporter is associated to the IP interface with tool port.
4. The map will eventually bind to the Exporter, Record, and Monitor.

NOTE: The dotted line from the map represents the interaction between the NetFlow Generation components.

Refer to *Example 1: NetFlow Generation Configuration* on page 631 for an example configuration of the following components.

Network Ports

NetFlow operates on the network flow. The incoming traffic on the network ports contains inputs such as, source and destination IP addresses, source and destination ports, interfaces, and so on. The network ports provide traffic to maps.

Map(s)

Traffic is received and acted upon according to maps. Maps determine what traffic is forwarded to NetFlow. Through map configuration, you add rules to filter the packets that need to go to NetFlow, and associate the map to the IP interface with tool port to specify where to send the filtered traffic.

Starting in software version 4.3.01, NetFlow supports both first level and second level maps. First level maps contain flow mapping rules to filter traffic that is needed by NetFlow and then send the filtered traffic to the IP interface with tool ports.

Second level maps are used for configuring filtering rules enabled through Adaptive Packet Filtering (APF). A virtual port is configured that directs traffic to the second level map. After the APF rules are applied, the filtered traffic that is needed by NetFlow is sent to the IP interface tool ports.

For examples of first level maps, refer to GigaSMART NetFlow Generation and GigaSMART NetFlow Generation.

For examples of first and second level maps, refer to GigaSMART NetFlow Generation and GigaSMART NetFlow Generation.

GigaSMART Group

The GigaSMART group specifies the GigaSMART engine to use, such as 8/1/e1 or 8/1/e2.

GSOP

The GigaSMART operation enables NetFlow. If a second level map is configured, the GigaSMART operation directs traffic to APF first, and then to NetFlow.

NetFlow Records

A NetFlow record contains key elements that specify what to match in the flow, such as all packets with the same source and destination port, or anything that comes in on a particular interface. A flow record also contains non-key elements that specify what information to collect for the flow, such as when the flow started or the number of bytes in the flow.

For NetFlow-v5, the fields in the flow record are fixed. For details, refer to *V5 Fixed Record Template* on page 657.

For NetFlow-v9 and IPFIX, you configure the fields, and thus create a record template. You specify how the fields are organized and in what order. The template is sent to the collector, so the collector knows what fields to expect in a NetFlow record. The template is sent periodically.

Starting in software version 4.6, multiple records are supported. An increased number of records allows more NetFlow data to be exported.

The maximum number of records is five. For all five records, each record must have the same match fields but differing collect fields. The same match fields will define the flows being considered. The different collect fields will define multiple templates sent to the NetFlow servers.

Starting in software version 5.1 for IPFIX and software version 5.2 for v5 and v9, a mix of IPv4 and IPv6 collect fields (IPv4 source/destination and IPv6 source/destination) are not supported in one record. Instead, create two records, one for IPv4 collect fields and one for IPv6 collect fields. When the IPv4/IPv6 collect fields are in separate records, an exporter will only send out records with non-blank elements.

NetFlow Monitors

Monitors store the NetFlow records associated with them. The configuration of a monitor includes the definition of the cache that specifies the records that you want to store, as well as timeouts associated with the cache. The cache can contain up to 4 million entries.

There can be a maximum of two monitors on a GigaSMART line card or module, one associated with each **e** port.

Starting in software version 4.6, up to five records can be added to the monitor. This results in the creation of five templates. For all five records, each record must have the same match fields but differing collect fields.

Sampled NetFlow Data

NetFlow data can be sampled. Sampling reduces the amount of ingress traffic sent to NetFlow for processing, which reduces the load on external collectors.

A NetFlow monitor can have multiple records with different sampling rates. The records are only updated with packets at the rate specified.

The following types of sampling are available: single-rate or multi-rate, as well as no sampling.

Sampling is enabled and disabled on the NetFlow monitor, across all flows. When sampling is enabled, you define the sampling rate by specifying a number for 1 in N, where N is the packet count.

For single-rate, the number can be from 10 to 16000. For multi-rate, the number can be from 1 to 16000. Single-rate applies to all records, whereas multi-rate applies to any record.

> **NOTE:** In a single-rate sampling type, all the NetFlow records are sampled in the same rate. In multi-rate sampling type, the sampling rate of the NetFlow records differ according to the settings defined in the individual records.

For example, if sampling is 1 in 1024, 1 packet in 1024 will be selected for NetFlow. The default is 1 in 1, which means no sampling.

> **NOTE:** The sampling mode in this release is deterministic. The selection of the packet is not random. Deterministic sampling means that if the rate is 1 in 1024, after 1023 packets, the 1024th packet is selected, while packets 1 to 1023 are ignored.

NetFlow Exporters

NetFlow records are sent to exporters. Each exporter is associated with one external collector. Records can be exported to both IPv4 and IPv6 destination. Either IPv4 or IPv6 destination address can be configured in an exporter. There can be up to six exporters that send flow records to up to six external collectors. The six destinations are per GigaSMART engine.

The configuration of an exporter includes the IP address of a collector, the transport protocol and destination port, and the template refresh interval, which specifies the frequency of when the record template is sent to the collector.

Starting in software version 5.1, an option is added to assign different exporters to different records. Instead of records being sent to all exporters, you can add an exporter to a record, which defines the exporter to which the record is sent.

IP Interface with Tool Ports

NetFlow exporters are associated with IP interface, since exporters route both records and templates to collectors in the network.

NOTE: It is expected that the gateway specified in the IP interface configuration does Layer 3 routing. However, when the IP interface and the collector's IP address are in the same subnet, the following applies:

- Configure the IP interface's gateway IP address to the same as the collector's IP address.
- Configure the IP interface's subnet mask.
- The maximum number of exporters supported per GigaSMART group is six.

Enhancements to NetFlow

In addition to the NetFlow components, there are also the following enhancements:

- *Exporter Filtering*
- *Remote Interface IDs*
- *NetFlow Option Templates*
- *IPFIX Extension: HTTP Response Code*
- *IPFIX Extension: Packet URL*
- *IPFIX Extension: User Agent*
- *IPFIX Extension: Domain Name Service (DNS)*
- *IPFIX Extension: SSL Metadata*

- *SNMP Packet Support on IP Interfaces with Tool Ports*
- *NetFlow Format Support on Exporters*

Exporter Filtering

Not all collectors are interested in all kinds of packets. On each exporter, you can configure pass filters to filter the records transmitted to a collector. Thus, you can send a subset of records to a collector, such as the flow records for UDP packets or for packets coming in on a particular port.

Filtering is based on criteria, such as ports or IP addresses. For example, you can filter on different interfaces, such as single port (1/1/x1) or a contiguous range of ports (1/1/x1..x4). Note that you can only filter the criteria or a subset of the criteria that you configured for the match fields in the record.

NOTE: If no filters are configured, all records are sent to the collectors.

The exporter pass filters are as follows:

- Input interface
- IPv4 and IPv6 DiffServ Code Point (DSCP)
- IPv4 and IPv6 source address
- IPv4 and IPv6 destination address
- IPv4 protocol
- IPv4 Type of Service (TOS)
- IPv6 flow label
- L4 source and destination port
- MAC source and destination address
- VLAN ID

Take into account the following considerations:

- an exporter can have up to 5 filter rules
- each rule can have up to 4 attributes
- input interface can only be specified once per filter
- other attributes can be specified multiple time in a rule
- two rules cannot be identical

For an example of exporter filtering, refer to GigaSMART NetFlow Generation.

Remote Interface IDs

Interface ID, ingress as well as egress, can be configured as match and collect fields. Interface IDs can be local or remote. If you are interested in the interface ID on which a packet arrives, you need the port number of the node sending the packet. To get that information, you can use the LLDP/CDP discovery protocols that talk to neighbors to fetch the remote interface ID.

Discovery has to be either enabled or disabled on all the ports in a map. If discovery is enabled, the remote interface ID is sent in the NetFlow data record, as learned through LLDP/CDP.

To configure port discovery with NetFlow, enable discovery on the port or ports that are specified in the **Source** field of the associated map.

Note the following:

- You cannot modify discovery once the map is defined.
- Local port IDs are unique across a cluster. Remote IDs might not be unique. With port discovery enabled, there is a possibility of port ID collisions.

NOTE: If port discovery is not enabled, the local port ID is sent in the NetFlow data record.

When port discovery is enabled, the sending of the remote ID requires the collaboration of the end nodes. NetFlow expects an integer for port ID. If end nodes send an alphanumeric string, MAC address, or IP address (non-integers), that cannot be translated in an integer, NetFlow interprets them as either 0xFFFF or 0xFFFFFFFF.

When port discovery is enabled and ingress LLDP/CDP packets contain interface IDs that cannot be translated into an integer, use the collect field **interface input name** in the flow record definition. Using an interface name will send meaningful information about a network port to help identify the port to which the flow record refers.

To access GigaSMART within GigaVUE-FM,GigaVUE HC Series, access a device that has been added to GigaVUE-FM,GigaVUE HC Series from the GigaVUE-FM,GigaVUE HC Series interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To use the collect field interface input name, do the following:

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Records**.
2. Click **New** to create a new Record or select the recored and click Edit to change an existing record.

3. Click in the field **Non-Key Fields (Collect)** and select **Interface** from the list.
4. Select **Input** and then **Input Name**. Specify an input **Width** as shown in Figure 40Collect Field Interface input Name.
5. Click **Save**.

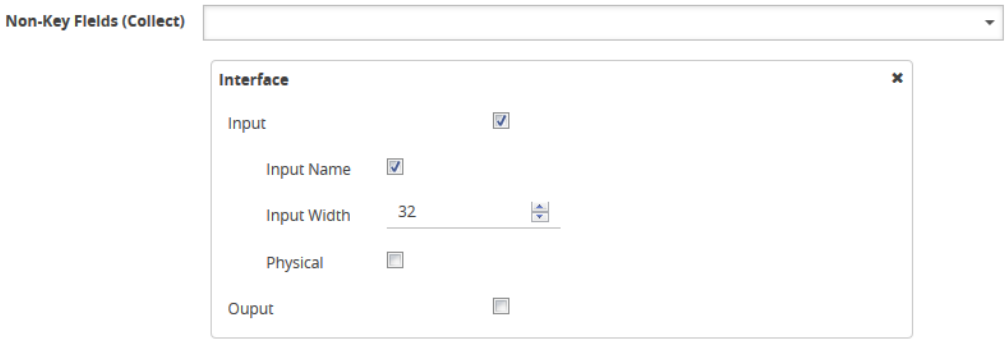


Figure 40 *Collect Field Interface input Name*

NetFlow Option Templates

For NetFlow-v9 and IPFIX, each exporter periodically sends option templates and option data records. There are two supported option templates, as follows:

- Interface ID to name mapping template and data record
- Exporter statistics template and data record

The option template for interface ID to name mapping contains an interface ID and name pair. Instead of a local port ID, the actual port number is available. For NetFlow-v9, the name field has a fixed length of 32 bytes. Names shorter than 32 bytes will be padded, while names longer than 32 bytes will be truncated. For IPFIX, the name field is of variable length.

NOTE: When port discovery is disabled for the port or ports specified in the **Source** field of the associated map, the interface option data record sends the interface ID to name mapping. But when discovery is enabled, interface option data records are not sent.

Each exporter sends out statistics, based on the standards. The exporter statistics option template includes information such as the exported flow record count, the exported message total count, and the exported octet total count.

By default, the transmission of option templates from the exporter is always enabled. The frequency of the transmission can be configured using the **Template Refresh interval** field in the NetFlow Exporter configuration page. To open the configuration page, select **GigaSMART > NetFlow / IPFIX Generation > Exporters** and click **New**.

IPFIX Extension: HTTP Response Code

For IPFIX only, use the collect field **Private PEN HTTP Response Code** in the flow record definition for capturing any packet with an HTTP response code embedded in it. This is a private information element extension, specific to Gigamon. The only valid private enterprise name (pen) is gigamon.

The HTTP response code information is captured from the packet and reported in the NetFlow record. The captured HTTP response code will be from the first packet that has HTTP/1 at the start of the HTTP header.

The field length in the flow record is a fixed length of 2 bytes. The range of response code values is from 100 to 599, as follows:

- 100-199 (informational)
- 200-299 (success related)
- 300-399 (redirection)
- 400-499 (client requests)
- 500-599 (server related)

If there is no HTTP response code in the flow, a zero value will be reported.

NOTE: In releases prior to software version 5.2, **HTTP Response Code** was directly under **Private PEN**. Starting in software version 5.2, there is a new **HTTP** section with **Response Code** under it. For backwards compatibility, both are supported.

To access GigaSMART within GigaVUE-FM,GigaVUE HC Series, access a device that has been added to GigaVUE-FM,GigaVUE HC Series from the GigaVUE-FM,GigaVUE HC Series interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure the NetFlow record for capturing any packet with an HTTP response code embedded in it, do the following:

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Records**.
2. Click **New**. Refer to Figure 41NetFlow Record Configuration for HTTP Response Code.

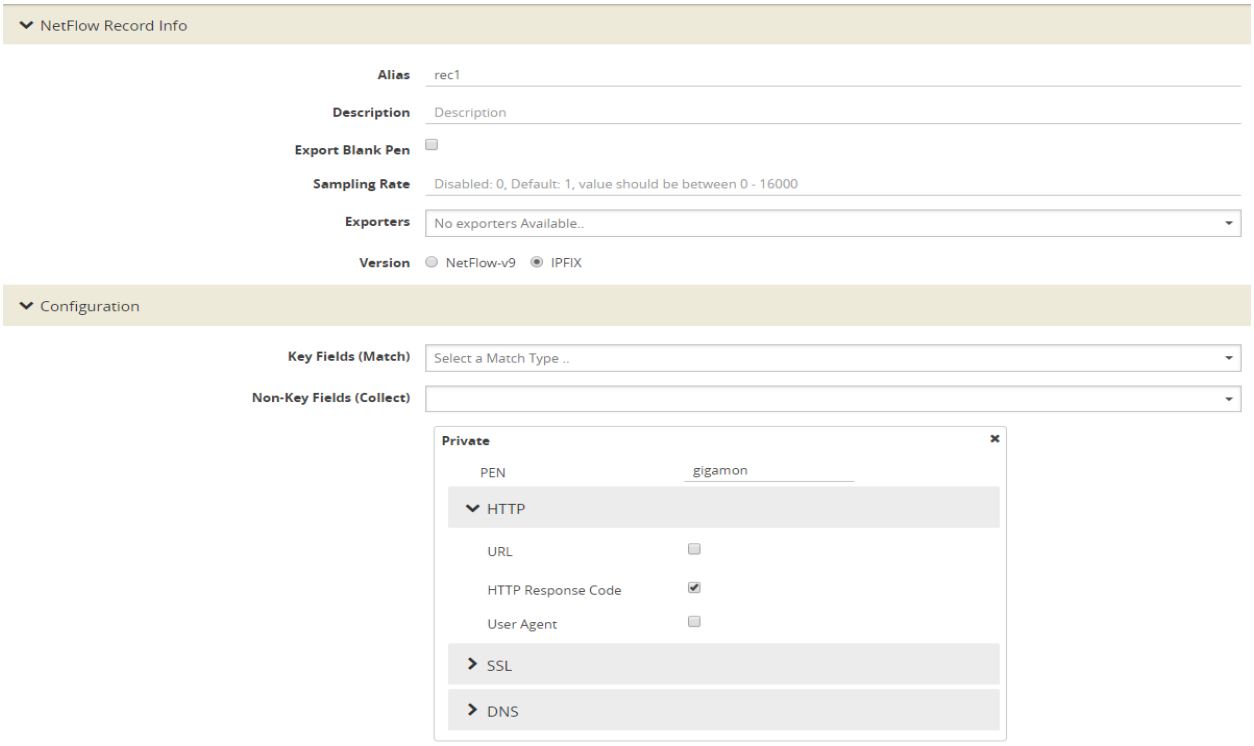


Figure 41 *NetFlow Record Configuration for HTTP Response Code*

3. For Version, select **IPFIX**.
4. Click the **Non-Key Fields (Collect)** drop-down list and select **Private**.
5. In the Private non-key field, do the following:
   - Set **PEN** to gigamon. (This is the default.)
   - Select **HTTP Response Code**.
6. Click **Save**.

IPFIX Extension: Packet URL

For IPFIX only, use the collect field **Private PEN HTTP URL** in the flow record definition for capturing any packet with a URL embedded in it. This is a private information element extension, specific to Gigamon. The only valid private enterprise name (pen) is gigamon.

The following URL information is captured from the packet and reported in the NetFlow record:

- HTTP: GET, POST, PUT, DELETE, and HEAD method types
- SIP: INVITE, ACK, BYE, REGISTER, OPTIONS, and CANCEL request types

The captured URL will be from the first packet that contains a URL. If there are additional URLs in subsequent packets in the flow, they will be ignored. If there is no URL in the flow, a zero length will be reported.

NOTE: The URL will always appear as the last element in a template, no matter the order in which the collect fields were configured.

NOTE: In releases prior to software version 5.2, **URL** was directly under **Private PEN**. Starting in software version 5.2, there is a new **HTTP** section with **URL** under it. For backwards compatibility, both are supported.

In Figure 41NetFlow Record Configuration for HTTP Response Code in the Private non-key field, select URL and enter an optional width.

IPFIX Extension: User Agent

For IPFIX only, use the collect field **Private PEN HTTP User Agent** in the flow record definition for capturing any packet with a user agent in the HTTP request header to gather information about clients user agents.

In general, the HTTP request is sent from the browser to the web application, so **User Agent** is filled in by the browser. As such, different browsers fill in this field with different values.

The maximum user agent length that is allowed in the data record is 250 bytes. The default is 150 bytes. Use the width parameter to specify a user agent length of up to 250 bytes.

In Figure 41NetFlow Record Configuration for HTTP Response Code in the Private non-key field, select User Agent and enter an optional width.

IPFIX Extension: Domain Name Service (DNS)

For IPFIX only, use the non-key or collect field **Private PEN DNS** in the flow record definition for capturing any packet with Domain Name Service (DNS) parameters embedded in it. This is a private information element, specific to Gigamon. The only valid private enterprise name (pen) is gigamon.

A domain name service translates host names into IP addresses. DNS has been exploited by attackers. Use this NetFlow enterprise element to gather metadata to help protect against security threats.

When certain DNS parameters are configured, the corresponding values for those collect parameters can be displayed in hexadecimal format or in text format. The following DNS parameters can display their values as text when the text version of that parameter is used:

- Additional Class, Additional Class Text
- Additional Type, Additional Type Text
- Authority Class, Authority Class Text
- Authority Type, Authority Type Text
- Query Class, Query Class Text
- Query Type, Query Type Text
- Response Class, Response Class Text
- Response IPv4 Address, Response IPv4 Address Text
- Response IPv6 Address, Response IPv6 Address Text
- Response Type. Response Type Text

For example, if the DNS **query-type** parameter collects a hexadecimal value of **0x1**, the **query-type-text** parameter collects the text string A, which refers to the IP address of the host.

The DNS parameters are captured from the packet and reported in the NetFlow record. Refer to Display Exporter Statistics and NetFlow Exporter Statistics Definitions.

Handle Blank Records for IPFIX

In the NetFlow record, the collect fields may contain one the following:

- Only private enterprise elements such as SSL, HTTP, or DNS
- Only non-private enterprise elements such as source IP address
- Both private and non-private elements

If all the collect fields contain only the private enterprise elements, and if during run-time, the records are blank or empty, they will not be added to NetFlow, however they will be counted in the exporter statistics as Empty Records Not Added.

If the collect fields contain both private and non-private enterprise elements, and if during run-time, the private enterprise elements are blank or empty, the records can be exported to the collector.

Configure DNS Record for IPFIX

To access GigaSMART within GigaVUE-FM,GigaVUE HC Series, access a device that has been added to GigaVUE-FM,GigaVUE HC Series from the GigaVUE-FM,GigaVUE HC Series interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure the NetFlow record for capturing any packet with (DNS) parameters embedded in it, do the following:

1. Go to **GigaSMART > NetFlow / IPFIX Generation > Records**.
2. Click **New**. Refer to Figure 42NetFlow Record Configuration for DNS.

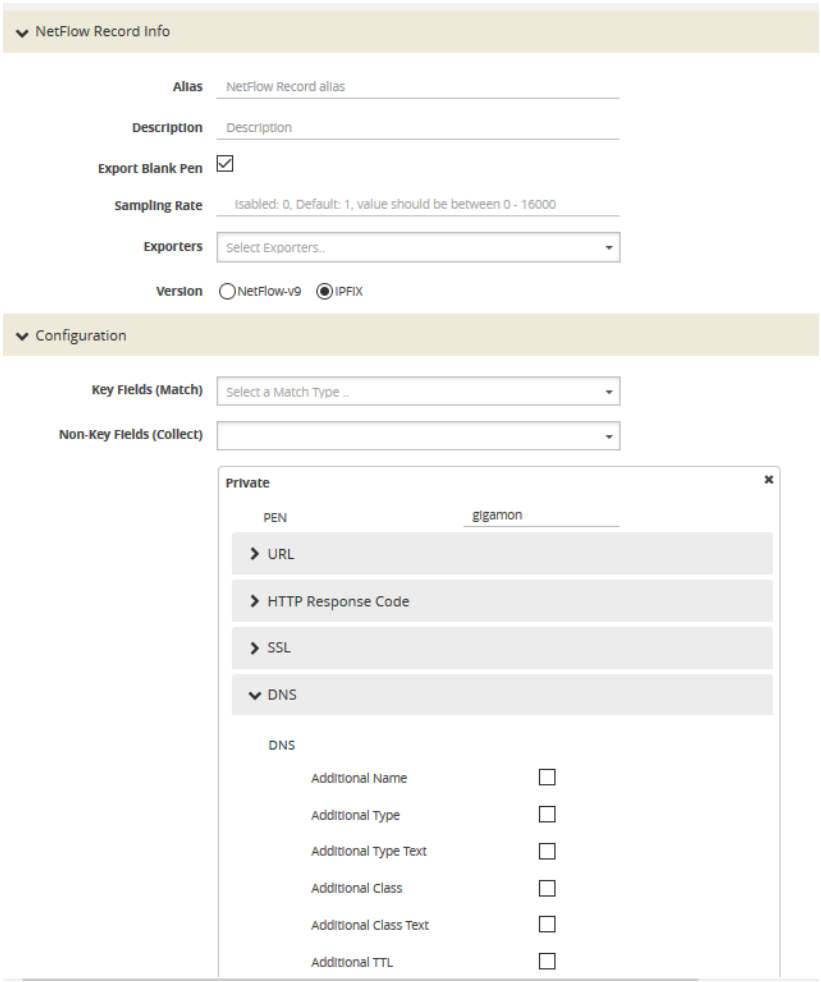


Figure 42 *NetFlow Record Configuration for DNS*

3. In the **Alias** field, enter a name.

4. To export the blank pen records, select **Export Blank Pen**.
5. For Version, select **IPFIX**.
6. Click the **Non-Key fields (Collect)** drop-down list and select **Private**.
7. In the **PEN** field, enter **gigamon**. (This is the default.)
8. Click **DNS** and select the parameters. The Number of Collects field is displayed for some DNS parameters. Refer to Table 2: DNS Parameters .
9. In the **Number of Collects** field, specify the number of instances of elements to collect from the DNS request. The value ranges from 1 to 10. The default value is 1.
10. Click **Save**.

Table 2: DNS Parameters

Method	For more information:
Additional Name	The domain name in the additional records section.
Additional Type	The additional type containing one of the RR type code.
Additional Type Text	The text string that corresponds to the hexadecimal value of the additional type containing one of the RR type code.
Additional Class	The additional class containing one of the RR class code.
Additional Class Text	The text string that corresponds to the hexadecimal value of the additional class containing one of the RR class code.
Additional TTL	The time-to-live (TTL), which is the time interval in seconds that the record is cached in the additional records section.
Additional RData	The content that describes the resource in the additional records section.
Additional RData Length	The length of the rdata field in the additional records section.
AN Count	The number of resource records in the answer section.
AR Count	The number of resource records in the additional records section.
Authority Name	The domain name in the authority section.

Method	For more information:
Authority Type	The authority type containing one of the RR type code.
Authority Type Text	The text string that corresponds to the hexadecimal value of the authority type containing one of the RR type code.
Authority Class	The authority class containing one of the RR class code.
Authority Class Text	The text string that corresponds to the hexadecimal value of the authority class containing one of the RR class code.
Authority TTL	The time-to-live (TTL), which is the time interval in seconds that the record is cached in the authority section.
Authority RData	The content that describes the resource in the authority section. The format of the rdata field varies according to the type and class of the resource record.
Authority RData Length	The length of the rdata field in the authority section.
Bits Count	The variable length of a bit map. The bit map must be a multiple of 8 bits long. For example: "/QR=1/AA=0/TC=0/RD=1/RA=1/AD=0/CD=0/Z=0", where /QR is the query (0) or a response (1), /AA is the authoritative answer, /TC is the truncation, /RD is the recursion desired, /RA is the recursion available, /AD is the authentic data, /CD is the checking disabled, and /Z is the reserved for future use.
Identifier	The identifier (Transaction ID) generated by the device that creates the DNS query and is copied by the server into the response so it can be used by that device to match that query to the corresponding reply received from the DNS server.
NS Count	The number of the name server (NS) resource records in the authority records section.
Op Code	The query type.
Qd Count	The number of entries in the question section.
Query Class	The query format containing one of the RR class codes.
Query Class Text	The text string that corresponds to the hexadecimal

Method	For more information:
	value of the query format containing one of the RR class codes.
Query Name	The domain name requested in the query. The maximum name length is 64 bytes. If the name is longer, it will be truncated.
Query Type	The query format containing one of the RR type codes.
Query Type Text	The text string that corresponds to the hexadecimal value of the query format containing one of the RR type codes.
Response Code	The type of the response.
Response Class	The response format containing one of the RR class codes.
Response Class Text	The text string that corresponds to the hexadecimal value of the response format containing one of the RR class codes.
Response Name	The domain name in the response. The maximum name length is 64 bytes. If the name is longer, it will be truncated.
Response Type	The query type specified in the response.
Response Type Text	The text string that corresponds to the hexadecimal value of the query type specified in the response.
Response RData Length	The length of the rdata field in the response data field.
Response RData	The content that describes the resource in the response data field. The format of the rdata field varies according to the type and class of the resource record.
Response-TTL	The time-to-live (TTL), which is the time interval in seconds that the record is cached.
Response iPv4 Address	The IPv4 address in the response if the response type host and class are Internet/IPv4.
Response iPv4 Address Text	The text string that corresponds to the hexadecimal value of the IPv4 address in the response if the response type host and class are Internet/IPv4. The format is dotted decimal.
Response IPv6 Address	The IPv6 address in the response if the response

Method	For more information:
	type host and class are Internet/IPv6.
Response iPv6 Address Text	The text string that corresponds to the hexadecimal value of the IPv6 address in the response if the response type host and class are Internet/IPv6. The format is dotted decimal.

IPFIX Extension: SSL Metadata

For IPFIX only, use the collect field **Private PEN ssl** in the flow record definition for capturing any packet with Secure Sockets Layer (SSL) or server metadata embedded in it, such as common name. This is a private information element extension, specific to Gigamon. The only valid private enterprise name (pen) is gigamon.

Examining the parameters associated with the SSL certificate or the SSL server provides visibility into SSL flows in the network and helps detect malicious activity. For example, checking the issuer might reveal an unknown self-signed certificate or a certificate signed by a questionable Certificate Authority (CA). Checking the certificate validity dates might reveal an expired certificate.

When NetFlow collects SSL certificate metadata, it makes use of the GigaSMART SSL application, described in GigaSMART Passive SSL Decryption. The data is routed to the SSL application first and then to NetFlow. If de-duplication is also enabled, the data is routed from de-duplication to SSL, and then to NetFlow. The SSL application does not decrypt the data.

NOTE: Only the NetFlow Generation license is needed for NetFlow to collect SSL certificate metadata.

When certain SSL certificate parameters are configured, the corresponding values for those collect parameters can be displayed in hexadecimal format or in text format. The following SSL certificate parameters can display their values as text when the text version of that parameter is used:

- Serial Number, Serial Number Text
- Signature Algorithm, Signature Algorithm Text
- Subject Algorithm, Subject Algorithm Text
- Valid Not After, Valid Not After Text
- Valid Not Before, Valid Not Before Text

To access GigaSMART within GigaVUE-FM,GigaVUE HC Series, access a device that has been added to GigaVUE-FM,GigaVUE HC Series from the GigaVUE-FM,GigaVUE HC Series interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure a record for SSL Certificate or SSL Server, do the following:

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Records**.
2. Click **New**.
3. For Version, select **IPFIX**.
4. Click in the **Non-Key fields (Collect)** field.
5. Select **Private** from the drop-down list.
6. In the Private non-key field, do the following:

- Set **PEN** to gigamon.
- Under SSL, select the SSL Certificate and SSL Server parameters and specify a width value in bytes. For example, **Certificate Issuer Common Name** and **Certificate Subject Common Name** as shown in Figure 43NetFlow Record Configuration for SSL.

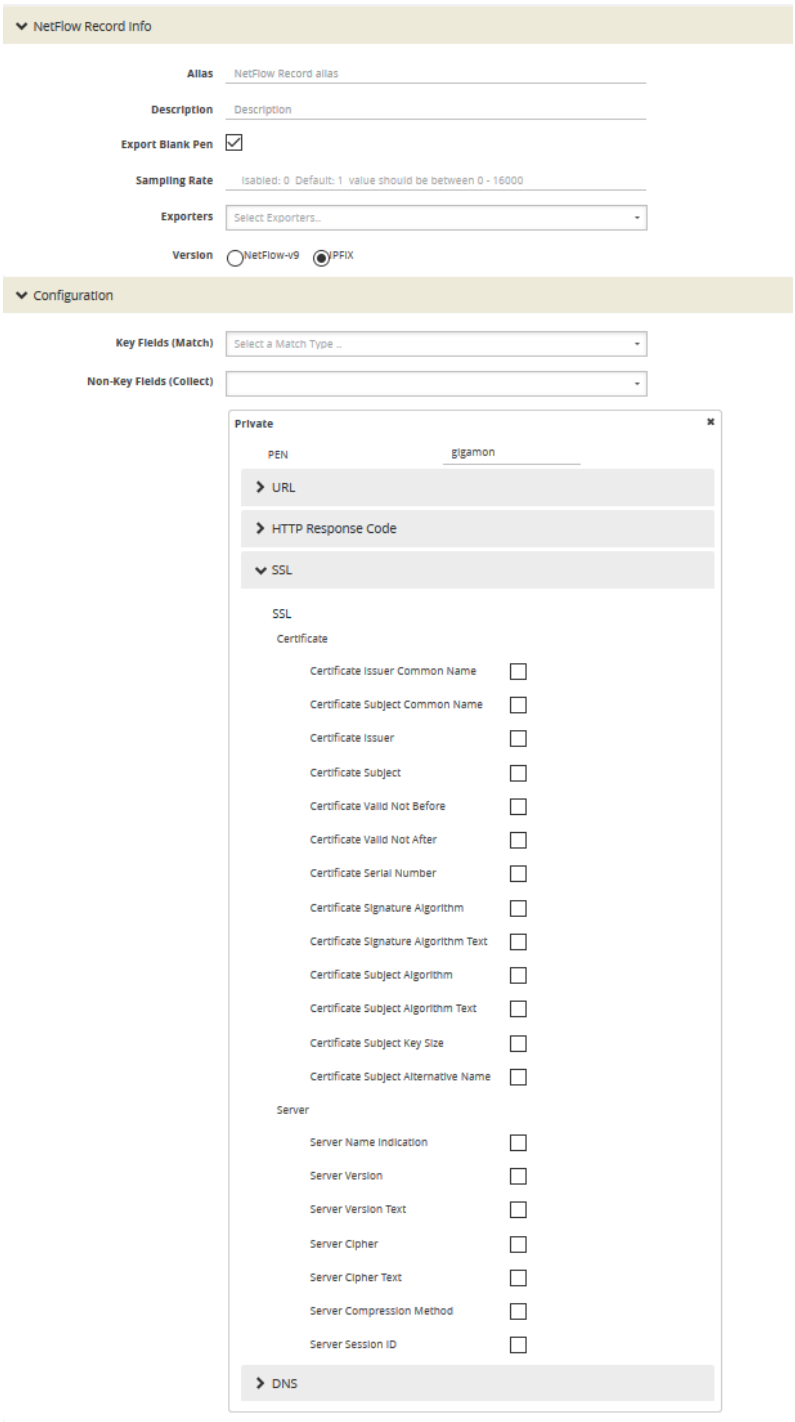


Figure 43 *NetFlow Record Configuration for SSL*

SSL Certificate Parameters

When certain SSL server parameters are configured, the corresponding values for those collect parameters can be displayed in hexadecimal format or in text format. The following SSL server parameters can display their values as text when the text version of that parameter is used:

- Cipher, Cipher Text
- Version, Version Text

For example, if the **ssl server cipher** parameter collects a hexadecimal value of **C027**, the **ssl server cipher-text** parameter collects the following text string:

TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256

The parameters supported for the SSL certificate are as follows:

- **Certificate Issuer**—the certificate issuer, which identifies the entity that has signed and issued the certificate. For example: "/C=US/ST=Arizona/L=Scottsdale/O=MyCo2.com, Inc./OU=http://certs.myco2.com/repository//CN=MyCo2 Secure Certificate Authority", where /C is the country name, /ST is the state or province, /L is the locality name, /O is the organization name, /OU is the organizational unit name, and /CN is the common name.
- **Certificate Issuer Common Name**—the certificate issuer common name, which is a subset of **Issuer**.
- **Certificate Subject**—the certificate subject, which identifies the entity associated with the public key stored in the subject public key. The **Certificate Subject** has the same fields as the **CertificateIssuer**.
- **Certificate Subject Common Name**—the certificate subject common name, which is a subset of **Subject**.
- **Certificate Subject Alternative Name**—the subject alternative name, which allows identities to be bound to the subject of the certificate. This parameter is useful to detect if the certificate claims to sign for anything else and to detect anomalies such as certificates that claim to sign for a wildcard (*). The first subject alternative name present in the certificate is collected.
- **Certificate Valid Not Before** and **Certificate Valid Not After**—the date on which the certificate validity period begins and ends. The certificate validity period is the time interval during which the CA warrants that it will maintain information about the status of the certificate. It is expressed in universal time. The format is YYMMDDHHMMSSZ, where Z is Zulu time (GMT).
- **Certificate Valid Not Before Text** and **Certificate Valid Not After Text**—the text string that corresponds to the hexadecimal value of the date on which the certificate validity period begins and ends. The format is MMM DD HH:MM:SS YYYY GMT.

- **Certificate Serial Number**—the unique number for each certificate issued by a given CA. The issuer name and serial number identify a unique certificate. This parameter is useful to detect any certificate changes or substitutions.
- **Certificate Serial Number Text**—the text string that corresponds to the hexadecimal value of the unique number for each certificate issued by a given CA.
- **Certificate Signature Algorithm**—the identifier for the cryptographic algorithm used by the CA to sign the certificate, defined in ASN.1 format. This parameter is useful to detect servers that are not compliant with an organization's cryptographic standards.
- **Certificate Signature Algorithm Text**—the text string that corresponds to the hexadecimal value of the identifier for the cryptographic algorithm used by the CA to sign the certificate, defined in ASN.1 format. This parameter is useful to detect servers that are not compliant with an organization's cryptographic standards.
- **Certificate Subject Algorithm**—the subject public key algorithm used, defined in ASN.1 format, such as RSA or DSA.
- **Certificate Subject Algorithm Text**—The text string that corresponds to the hexadecimal value of the subject public key algorithm used, defined in ASN.1 format, such as RSA or DSA.
- **Certificate Subject Key Size**—the subject public key size.

Optionally, on the **issuer**, **Certificate Issuer Common Name**, **Certificate Subject**, **Certificate Subject Common Name**, and **Certificate Subject Alternative Name** parameters, you can indicate the width of the field in bytes.

SSL Server Parameters

The parameters supported for the SSL server are as follows:

- **Server Name Indication**—the extension to the Transport Layer Security (TLS) protocol by which a client indicates the host name to which it is attempting to connect at the start of the handshaking process.
- **Server Version**—the version of SSL, including the major and minor version.
- **Server Version Text**—the text string that corresponds to the hexadecimal value of the identifier for the version of SSL, including the major and minor version.
- **Server Cipher**—the cipher that the server agreed to use for that session.
- **Server Cipher Text**—the text string that corresponds to the hexadecimal value of the identifier for the cipher that the server agreed to use for that session.
- **Server Compression Method**—the server compression method, which is typically not set (in other words, NULL). This parameter is useful to detect attacks that use compression.
- **Server Session ID**—the session identifier, generated by a server, which identifies a particular session. This parameter is useful to detect a session restart.

Optionally, on the **Server Name Indication** parameter, you can indicated the width of the field in bytes.

Restrict Ports for NetFlow SSL Sessions

SSL metadata is collected by sending all traffic to the SSL module. The SSL module accepts all IPv4 TCP packets and attempts to find SSL sessions. During the process of finding these sessions, the metadata required by NetFlow is extracted.

To improve the throughput of SSL metadata extraction for NetFlow, the TCP ports can be restricted. Reducing the TCP packets inspected by limiting the TCP ports inspected reduces the amount of packets sent to the SSL module.

Configure the monitor to scan specific ports for SSL. Options are available to scan all ports, a list of up to 10 ports, or well-known ports.

The following are the well-known ports:

- MAP_SSL_PORT 993
- POP3_SSL_PORT 995
- SMTP_SSL_PORT 465
- LDAP_SSL_PORT 636
- NNTP_SSL_PORT 563
- HTTP_SSL_PORT 443

Configure SSL Certificate and Server Parameters

To access GigaSMART within GigaVUE-FM,GigaVUE HC Series, access a device that has been added to GigaVUE-FM,GigaVUE HC Series from the GigaVUE-FM,GigaVUE HC Series interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure the SSL certificate and server parameters to collect, do the following in the UI, for example:

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Records**.
2. Click **New**.
3. Enter an alias in the **Alias** field to identify this record. For example, ipfixrec.
4. Select **IPFIX**.
5. From the **Non-Key Field (Collect)** list, select **Private**.
6. Set **PEN** to gigamon. (This is the default.)

7. Under SSL, select any of the following:
    - Certificate issuer Common Name Width 30
    - Certificate Subject Common Name Width 40
    - Certificate issuer Width 150
    - Certificate Subject Width 120
    - Certificate Valid Not Before
    - Certificate Valid Not After
    - Certificate Serial Number
    - Certificate Signature Algorithm
    - Certificate Signature Algorithm Text
    - Certificate Subject Algorithm
    - Certificate Subject Algorithm Text
    - Certificate Subject Key Size
    - Certificate Subject Alternative Name
    - Server Name Indication Width 64
    - Server Version
    - Server Version Text
    - Server Cipher
    - Server Cipher Text
    - Server Compression Method
    - Server Session ID
8. Click **Save**.

Best Practices for Collecting SSL Metadata

When collecting SSL certificate metadata, the match conditions must be configured so that the NetFlow sessions match the SSL sessions. To do this, configure the following in the NetFlow record for IPv4 flows:

1. From the device view, select **GigaSMART >**NetFlow Record / IPFIX Generation.
2. Select the record ipfixrec and click Edit. (This is the record configured in the previous section Configure SSL Certificate and Server Parameters.)
3. From the **Key Fields (Match)** list, select **IPv4**.
4. Select **Protocol**.
5. Under **Source** select **Address**.
6. Under **Destination** select **Address**.
7. From the **Key Fields (Match)** list, select **Transport** and then select the following:

- **Source Port**
- **Destination Port**

8. Click **OK**.

Or, configure the following in the NetFlow record for IPv6 flows:

1. From the device view, select **GigaSMART >**NetFlow Record / IPFIX Generation.
2. Select the record ipfixrec and click Edit. (This is the record configured in the previous section Configure SSL Certificate and Server Parameters.)
3. From the **Key Fields (Match)** list, select **IPv6**.
4. Select **Protocol**.
5. Under **Source** select **Address**.
6. Under **Destination** select **Address**.
7. From the **Key Fields (Match)** list, select **Transport** and then select the following:
   - **Source Port**
   - **Destination Port**
8. Click **OK**.

Or, configure the following in the NetFlow record for a mix of IPv4 and IPv6 flows:

1. From the device view, select **GigaSMART >**NetFlow Record / IPFIX Generation.
2. Select the record ipfixrec and click Edit. (This is the record configured in the previous section Configure SSL Certificate and Server Parameters.)
3. From the **Key Fields (Match)** list, select **IPv4**.
4. Select **Protocol**.
5. Under **Source** select **Address**.
6. Under **Destination** select **Address**.
7. From the **Key Fields (Match)** list, select **IPv6**.
8. Select **Protocol**.
9. Under **Source** select **Address**.
10. Under **Destination** select **Address**.
11. From the **Key Fields (Match)** list, select **Transport** and then select the following:
    - **Source Port**
    - **Destination Port**
12. Click **OK**.

Refer to NetFlow Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

Notes and Considerations for SSL Certificate Metadata

Refer to the following notes and considerations for SSL certificate metadata:

- When the GigaSMART SSL application is gathering SSL certificate metadata for NetFlow, it is not able to decrypt SSL packets at the same time.
- Using the SSL application plus NetFlow to collect SSL certificate metadata consumes GigaSMART resources, resulting in less memory for other GigaSMART applications.
- To retrieve SSL certificate metadata, there must be a valid SSL session in which most of the packets that form the session are received, as follows:
  - For the SSL certificate parameters, all packets up to at least the certificate packet must be received.
  - For the SSL server parameters except **SeverName Indication**, all packets up to the server hello packet must be received.
  - For the **ServerName Indication** SSL server parameter, all packets up to the client hello packet must be received.

SNMP Packet Support on IP Interfaces with Tool Ports

SNMP packet support on IP interfaces with tool ports processes SNMP packets arriving on IP interfaces with tool ports and forwards them to GigaSMART so that external NetFlow collectors can integrate with GigaSMARTNetFlow Generation.

External collectors need to recognize GigaSMARTNetFlow Generation as a valid NetFlow interface. The validation can be manual or automatic, as follows:

- with manual validation, configuration on the collector side must provide the details of GigaSMARTNetFlow Generation
- with automatic validation, the recipient collector initiates an SNMP query requesting relevant information and GigaSMARTNetFlow Generation responds to it with the required information

SNMP packet support on IP interfaces with tool ports provides automatic validation. Starting in software version 4.5, GigaSMARTNetFlow Generation processes SNMP packets arriving on IP interfaces with tool ports.

NetFlow records are sent to exporters through IP interfaces with tool ports. Each NetFlow exporter is associated with one external collector. Up to six exporters can send flow records to up to six external collectors.

An IP interface with tool port can have multiple exporters. An external collector can listen to multiple exporters.

To listen to SNMP packets from external collectors, enable SNMP under the NetFlow exporter. The following are the steps to allow external collectors to send SNMP packets to the IP interface with tool port:

1. From the device view, select **GigaSMART > NetFlow > Exporters**.
2. Click **New** to create a new exporter or Edit to configure an existing exporter.
3. Select **SNMP** under the SNMP section.

These steps enables SNMP on the default port, which is port number 161.

> **NOTE:** Only the default SNMP port is supported for packets arriving on the IP interface. If the incoming request uses a non-default SNMP port, they will be dropped at the IP interface.

To disable listening for SNMP packets by a specified NetFlow exporter, uncheck the SNMP check box.

By default, listening to SNMP packets from external collectors is disabled.

NetFlow Format Support on Exporters

NetFlow Exporters support versions IPFIX, v5, and v9. Starting in software version 5.3, the Common Event Format (CEF) version 23 is also supported. CEF is a standard format used by event collection/correlation Security Information and Event Management (SIEM) vendors. SIEMs such as Arcsight, Splunk, and QRadar accept CEF format. By supporting CEF, NetFlow metadata can integrate with and use a variety of SIEMs.

CEF is a logging format that uses the syslog message as a transport mechanism, meaning that the CEF message (header and payload) is included within the syslog message. The transport protocol that is supported is UDP and the default port number is 514.

Metadata that is generated by NetFlow can be exported in the supported formats to one or more collectors. Each exporter must have the same export type (v5, v9, IPFIX, or CEF). One CEF message is sent out per record per flow.

Also, starting in software version 5.3, IP fragmentation is supported. CEF does not allow a message to be split over multiple CEF payloads. Since CEF messages are verbose, they can be larger than the MTU.

To support CEF messages that exceed the MTU, a single IP datagram containing a CEF message will be broken up into multiple packets of smaller sizes. The reassembly of the datagram will occur at the receiving end (at the SIEMs).

For details on the CEF message format, refer to

CEF Message Format

An example of the CEF message format is as follows:

```
Fri Feb 23 02:25:37 2018 9/3/e1
CEF:23|Gigamon|metadata|5.3.00|4|metadatageneration|6| src=68.94.156.1
GigamonMdataDnsAdditionalType=41GigamonMdataDnsAdditionalTypeText=OPT
```

In the example CEF message, there is a syslog header, a CEF header, and an extension that contains the CEF payload. The fields are delimited with a vertical bar (|).

The syslog header contains the following:

- timestamp—Fri Feb 23 02:25:37 2018
- host name identifier—9/3/e1

NOTE: The host name identifier has the format <box ID>/<slot ID>/<engine ID>.For example, 9/3/e1 means 9 is the box ID, 3 is the slot ID, and e1 is the engine ID.

The CEF header contains the following:

- version—CEF:23
- device vendor—Gigamon
- device product—metadata
- device version—5.3.00
- signature identifier—4
- name—metadata generation
- severity—6

The CEF extension contains key-value pairs delimited with a space. In the example CEF message, the following is the CEF payload, in plaintext:

- src=68.94.156.1
- GigamonMdataDnsAdditionalType=41
- GigamonMdataDnsAdditionalTypeText=OPT

The CEF standard specifies key-value pairs. There are some predefined standardkeys, for example, src is a predefined key for source IP address.

For keys that are not predefined in the CEF standard, such as the NetFlow metadata elements in the CEF extension, there are custom-defined keys. Custom-defined keys have the following format:

- <VendorNameProductNameExplanatoryKeyName>

For example, GigamonMdataDnsAdditionalTypeText, is a custom-defined key that contains the following:

- VendorName—Gigamon
- ProductName—Mdata
- ExplanatoryKeyName—DnsAdditionalTypeText

Another example of the CEF format is the following SSL record:

```
Thu Mar 1 08:21:28 2018 1/1/e1 CEF:23|Gigamon|metadata|5.3.00|4|metadata
generation|6|GigamonMdataSslIssuerName=DigiCert SHA2 High Assurance S
dpt=54839 GigamonMdataSslValidNotBefore=3137303130363030303030305a
GigamonMdataSslSerialNo=0118ee3c2167b99e1b718c6eadb8fb4d00000000
GigamonMdataSslValidNotAfter=3230303131353132303030305a
GigamonMdataSslCertSigAlgo=2a864886f70d01010b
GigamonMdataSslCertSubAlgo=2a864886f70d010101
GigamonMdataSslCertSubKeySize=270 GigamonMdataSslServerVersion=771
GigamonMdataSslCertSubAltName=*.stickyadstv.com
GigamonMdataSslServerCompressionMethod=192 GigamonMdataSslServerCipher=49199
GigamonMdataSslServerVersionText=TLSv1.2 GigamonMdataSslServerSessionId=63
GigamonMdataSslIssuer=2f433d55532f4f3d44696769436572742049648413220486967682041737375726 16e636
e6469676963
6572742e636f6d2f434e3d44696769436572742053484132204869676820417373757261 6e636
52053657276 6572204341 GigamonMdataSslCertSubCommonName=*.stickyadstv.com
GigamonMdataSslSub=2f433d55532f53543d4e657720596f726b2f4c3d4e657720596f726b2f
4f3d4672656
```

5776865656c204d6564696120496e632f4f553d46726565776865656c2f434e3d2a2e737469636b796164737 4762e636f6d dst=10.50.22.59 src=38.106.34.118 spt=443

Display Exporter Statistics

To display exporter statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics** and open the GigaSMART Statistics Quick View to view the NetFlow Statistics.

Refer to NetFlow Exporter Statistics Definitions for descriptions of the statistics.

Display Monitor Statistics

To display exporter statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics** and open the GigaSMART Statistics Quick View to view the NetFlow Statistics.

Refer to NetFlow Monitor Statistics Definitions for descriptions of these statistics.

Display IP Interfaces Statistics

To display IP interfaces statistics, select **Ports > IP Interfaces > Statistics** and look for the IP interface ID in the statistics table.

Refer to IP Interfaces Statistics Definitions for descriptions of these statistics.

NetFlow Generation Configuration Modification and Removal

There may be instances where a NetFlow Generation configuration may require alteration by modifying a NetFlow Generation Monitor Configuration or a NetFlow Generation Record Configuration. It may further require that the configuration be removed entirely. In such instances, refer to the following.

Modify NetFlow Generation Monitor Configuration

This example shows the modification of a NetFlow Generation Monitor configuration.

1. Unlink the monitor from GigaSMART Parameters.
    a. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.
    b. Select the GigaSMART group to modify.
    c. Click **Edit**.
    d. Under NetFlow, select **None** in the **Monitor** field.
    e. Click **Save**.
2. Modify the monitor parameters.
    a. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Monitors**.
    b. Select the Monitor to modify.
    c. Click **Edit**.
    d. Under Config, modify the monitor parameters.
    e. Select the record from the **Record(s)** list to re-add it to the monitor.
3. Re-add the monitor to GigaSMART Parameters for the changes to take affect.
    a. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.
    b. Select the GigaSMART group to modify.

c. Click **Edit**.

d. Under NetFlow, select the monitor in the **Monitor** field.

e. Click **Save**.

Modify NetFlow Generation Record Configuration

This example shows the modification of a NetFlow Generation Record configuration.

1. Unlink the monitor from gsparams.
   a. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.
   b. Select the GigaSMART group to modify.
   c. Click **Edit**.
   d. Under NetFlow, select **None** in the **Monitor** field.
   e. Click **Save**.
2. Modify the record bound to the monitor.
   a. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Records**.
   b. Select the record to modify.
   c. Click **Edit**.
   d. Modify the record configuration.
3. Re-add the monitor to the GigaSMART Parameters for changes in record to take affect.
   a. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.
   b. Select the GigaSMART group to modify.
   c. Click **Edit**.
   d. Under NetFlow, select the monitor in the **Monitor** field.
   e. Click **Save**.

Remove NetFlow Generation Configuration

Use the following steps to remove a NetFlow Generation Configuration:

1. Remove the NetFlow parameter from the GigaSMART Group.
   a. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.
   b. Select the GigaSMART Group.
   c. Click **Edit**.
   d. Under NetFlow, select **None** in the **Monitor** field.
2. Delete the Maps.
   a. Select **Maps > Maps > Maps**.
   b. Select Table View.
   c. Select the Maps.
   d. Click **Delete**.
3. Delete the IP interface.
   a. Select **Ports > IP Interfaces**.
   b. Select the port.
   c. Click **Delete**.
4. Delete the monitor, records, and exporter
   a. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Monitors**.
   b. Select the monitor, and then click **Delete**.
   c. Select **Records**
   d. Select the record, and then click **Delete**
   e. Select **Records**.
   f. Select the record, and then click **Delete**.

V5 Fixed Record Template

NetFlow v5 records have a template of fixed fields that cannot be edited. The template contains Match/Key and Collect/Non-Key elements. It has an alias of **predefined_netflow_v5_record**.

To display the template, select **GigaSMART > NetFlow / IPFIX Generation > Records** and click on **predefined_netflow_v5_record** to display the Record Quick View shown in Figure 44NetFlow Record predefined_netflow_v5_record.

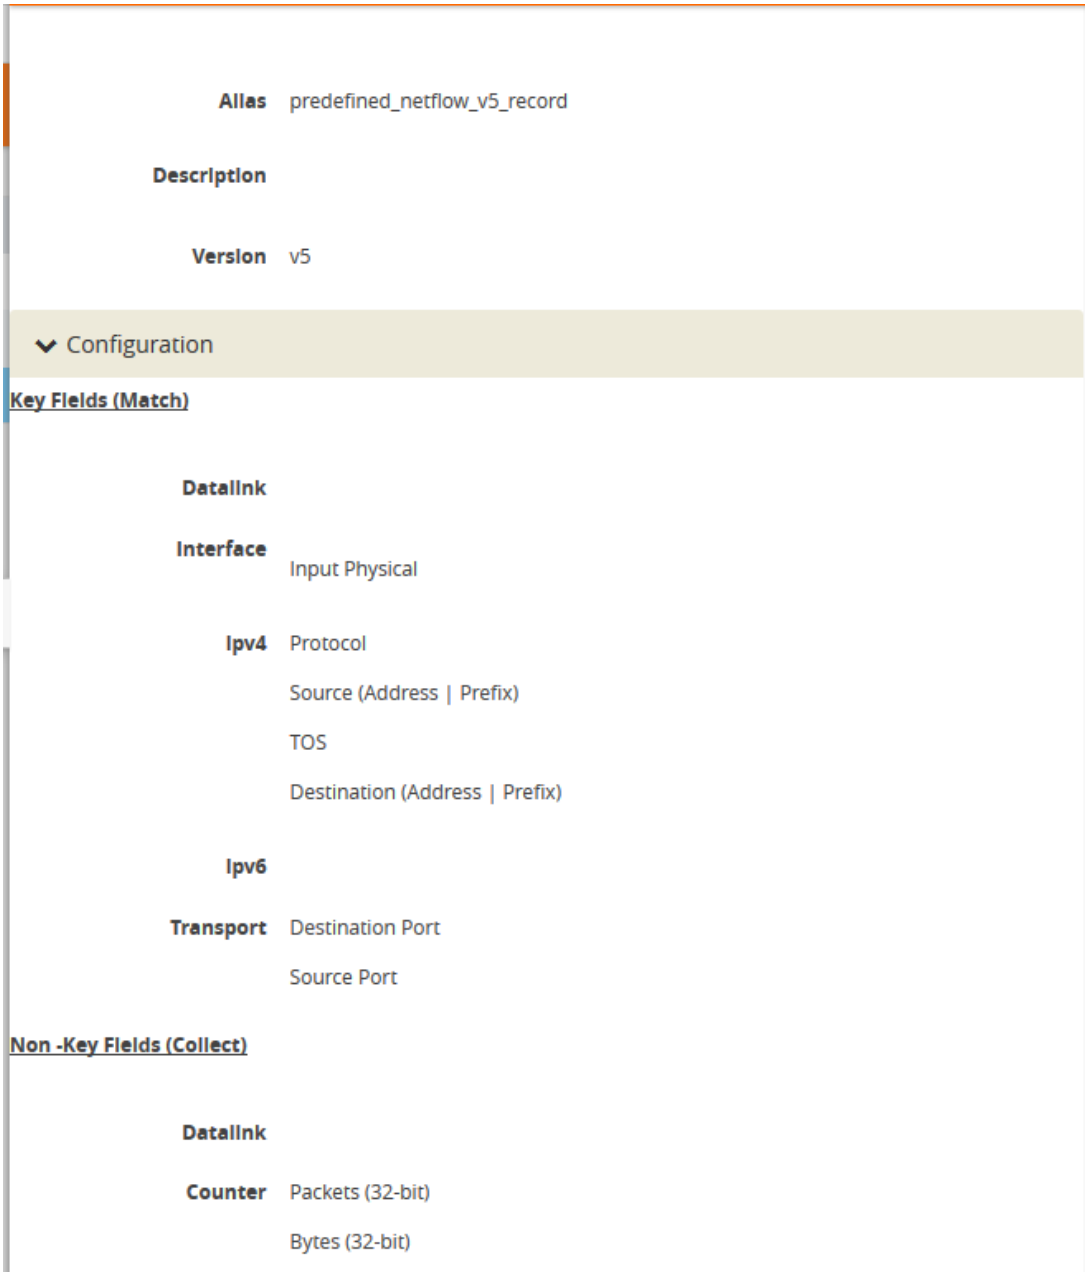


Figure 44 *NetFlow Record predefined_netflow_v5_record*

NetFlow Generation Match/Key and Collect/Non-Key Elements

NetFlow v9 and IPFIX records allow the user to configure Match/Key and Collect/Non-Key elements.

Match/Key Syntax

NetFlow v9 and IPFIX records allow the you to configure Match/Key elements.

NOTE: NetFlow v9 does not support Match/Key elements whose ID on the specified link is greater than 128. For additional information, refer to the following:

http://www.iana.org/assignments/ipfix/ipfix.xhtml

To configure the Match/Key elements, click in the Key Fields (Match) field in the NetFlow Record configuration page and select the match type.

The supported combinations of Match/Key elements are outlined in the following table:

Match Type	Parameters			Description
Data Link	**Source Mac**			Supported for v9 and IPFIX.
	Destination			Supported for v9 and IPFIX.
	VLAN			Supported for v9 and IPFIX.
Interface	**Input physical**	**Physical Width-2** **Physical Width-4**		Supported for v9 and IPFIX. for width, the only supported values are 2 or 4.
IPv4	**Destination**	**Address**		Configures the IPv4 destination address as a key field. Supported for v9 and IPFIX.
		Prefix	**<netmask | mask_ length>**	Configures a prefix for the IPv4 destination address as a key field. Supported for v9 and IPFIX.
	DSCP			Supported only for IPFIX.
	Fragmentation Flags			Supported only for IPFIX.
	Fragmentation ID			Supported for v9 and IPFIX.
	Fragmentation Offset			Supported for v9 and IPFIX.
	Header Length			Supported only for IPFIX.

Match Type	Parameters			Description
	Option Map			Supported only for IPFIX.
	Precedence			Supported only for IPFIX.
	Protocol			Supported for v9 and IPFIX.
	Section	**Header Size**	**<size>**	Configures the number of bytes of raw data starting at the IPv4 header, to use as a key field. The range is from 1 to 128. Supported only for IPFIX.
		Payload Size	**<size>**	Configures the number of bytes of raw data starting at the IPv4 payload, to use as a key field. The range is from 1 to 128. Supported only for IPFIX.
	Source	**Address**		Supported for v9 and IPFIX.
		Prefix	**<netmask | mask_ length>**	Supported for v9 and IPFIX.
	TOS			Supported only for IPFIX.
	Total Length	**maximum | minimum**		Supported only for IPFIX.
	TTL			Supported only for IPFIX.
	Version			Supported for v9 and IPFIX.
IPv6	**Destination**	**Address**		Supported for v9 and IPFIX.
		Prefix	**<netmask | mask_ length>**	Supported only for IPFIX.
	DSCP			Supported only for IPFIX.
	Extension Map			Supported for v9 and IPFIX.
	Flow Label			Supported for v9 and IPFIX.
	Fragmentation Flags			Supported only for IPFIX.
	Fragmentation ID			Supported for v9 and IPFIX.

Match Type	Parameters			Description
	Fragmentation Offset			Supported for v9 and IPFIX.
	Hop Limit			Supported only for IPFIX.
	Length	**Header**		Supported only for IPFIX.
		Payload		Supported only for IPFIX.
		Total		Supported only for IPFIX.
	Next Header			Supported only for IPFIX.
	payload-length			Supported only for IPFIX.
	Precedence			Supported only for IPFIX.
	Protocol			Supported for v9 and IPFIX.
	Section	**Header Size**	**<size>**	Supported only for IPFIX. The range is from 1 to 128.
		Payload Size	**<size>**	Supported only for IPFIX. The range is from 1 to 128.
	Source	**Address**		Supported for v9 and IPFIX.
		Prefix	**<netmask | mask_ length>**	Supported only for IPFIX.
	Traffic Class			Supported for v9 and IPFIX.
	Version			Supported for v9 and IPFIX.
Transport	**Destination Port**			Supported for v9 and IPFIX.
	ICMP	**IPv4**	**Code**	Supported only for IPFIX.
			Type	Supported only for IPFIX.
		IPv6	**Code**	Supported only for IPFIX.
			Type	Supported only for IPFIX.
	Source Port			Supported for v9 and IPFIX.
	TCP	**ACK Number**		Supported only for IPFIX.
		Destination Port		Supported only for IPFIX.
		Flags <enable	**[ACK] | [CWR] |**	Supported only for v9 and

Match Type	Parameters			Description
		| disable>	**[ECE] | [FIN] | [PSH] | [RST] | [SYN] | [URG]**	IPFIX.
		Header Length		Supported only for IPFIX.
		Sequence Number		Supported only for IPFIX.
		Source Port		Supported only for IPFIX.
		Urgent Pointer		Supported only for IPFIX.
		window-size		Supported only for IPFIX.
	UDP	**Destination Port**		Supported only for IPFIX.
		Message Length		Supported only for IPFIX.
		Source Port		Supported only for IPFIX.

Collect/Non-Key Syntax

NetFlow v9 and IPFIX records allow the user to configure Collect/Non-Key elements.

The number of Collect/Non-Key elements in a record can be up to 32. Each Collect/Non-Key element has a size. The accumulated size of the Collect/Non-Key elements in the record cannot exceed 1024 bytes. The supported Collect/Non-Key elements is determined either by the maximum number of elements in a record (32) or by the maximum size (1024 bytes), whichever is reached first.

NOTE: NetFlow v9 does not support Collect/Non-Key elements whose ID on the specified link is greater than 128. For additional information, refer to the following:

http://www.iana.org/assignments/ipfix/ipfix.xhtml

To configure the Collect/Non-Key elements, click in the **Non-Key Fields (Collect)** field in the NetFlow Record configuration page and select the match type.

The supported combinations of Collect/Non-Key elements are outlined in the following table:

Collect Type	Parameters	Size		Description
Counter	**Bytes**	**32** **64**		Supported for v9 and IPFIX.
	Packets	**32** **64**		Supported for v9 and IPFIX.
Datalink	**Source**			Supported for v9 and IPFIX.
	Mac Destination			Supported for v9 and IPFIX.
	VLAN			Supported for v9 and IPFIX.
Flow	**End Reason**			Supported only for IPFIX.
Interface	**Input Name**	**Input Width**	**[width]**	Supported for v9 and IPFIX. for width, the range is from 1 to 32.
	Physical	**Physical Width-2** **Physical Width-4**		Supported for v9 and IPFIX. For width, the only supported values are 2 or 4.
	Output	**Physical Width-2** **Physical Width-4**		Supported for v9 and IPFIX. For width, the only supported values are 2 or 4.
IPv4	**Destination**	**Address**		Configures the IPv4 destination address as a non-key field. Supported for v9 and IPFIX.
		Prefix	**<netmask | mask_ length>**	Supported for v9 and IPFIX.
	DSCP			Supported only for IPFIX.
	Fragmentation Flags			Supported only for IPFIX.
	Fragmentation ID			Supported for v9 and IPFIX.

Collect Type	Parameters	Size		Description
	Offset			Supported for v9 and IPFIX.
	Header Length			Supported only for IPFIX.
	Option Map			Supported only for IPFIX.
	Precedence			Supported only for IPFIX.
	Protocol			Supported for v9 and IPFIX.
	Section	**Header Size**	**<size>**	Configures the number of bytes of raw data starting at the IPv4 header, to use as a key field. The range is from 1 to 128. Supported for v9 and IPFIX.
		Payload Size	**<size>**	Configures the number of bytes of raw data starting at the IPv4 payload to use as a key field. The range is from 1 to 128. Supported for v9 and IPFIX.
	Source	**Address**		Supported for v9 and IPFIX.
		Prefix	**<netmask | mask_ length>**	Configures a prefix for the IPv4 destination address as a non-key field. Supported for v9 and IPFIX.
	TOS			Supported only for IPFIX.
	Total Length	**[maximum]**		Supported only for IPFIX.
		[minimum]		Supported only for IPFIX.
	TTL			Supported only for IPFIX.
	Version			Supported for v9 and IPFIX.
IPv6	**Destination**	**Address**		Supported for v9 and IPFIX.
		Prefix	**<netmask | mask_ length>**	Supported only for IPFIX.

Collect Type	Parameters	Size		Description
	DSCP			Supported only for IPFIX.
	Extension Map			Supported for v9 and IPFIX.
	Flow Label			Supported for v9 and IPFIX.
	Fragmentation Flags			Supported only for IPFIX.
	Fragmentation ID			Supported for v9 and IPFIX.
	Fragmentation Offset			Supported for v9 and IPFIX.
	Hop Limit	**[maximum]**		Supported only for IPFIX.
		[minimum]		Supported only for IPFIX.
	Length	**Header**		Supported for v9 and IPFIX.
		Payload		Supported only for IPFIX.
		Total	**[maximum]**	Supported only for IPFIX.
			[minimum]	Supported only for IPFIX.
	Next Header			Supported only for IPFIX.
	Precedence			Supported only for IPFIX.
	Protocol			Supported for v9 and IPFIX.
	Section	**Header Size**	**<size>**	Supported only for IPFIX. The range is from 1 to 128.
		Payload Size	**<size>**	Supported only for IPFIX. The range is from 1 to 128.
	Source	**Address**		Supported for v9 and IPFIX.
		Prefix	**<netmask | mask_ length>**	Configures a prefix for the IPv4 destination address as a non-key field. Supported only for IPFIX.
	Traffic Class			Supported for v9 and

Collect Type	Parameters	Size		Description
				IPFIX.
	Version			Supported for v9 and IPFIX.
Private	**PEN <pen name>**	**DNS**	**<additional-class [number-of-collects <1-10>] | additional-class-text [number-of-collects <1-10>] |additional-name [number-of-collects <1-10>] | additional-rd-length [number-of-collects <1-10>] | additional-rdata [number-of-collects <1-10> | width <1-128>] | additional-ttl [number-of-collects <1-10>] | additional-type [number-of-collects <1-10>] | additional-type-text [number-of-collects <1-10>] | an-count | ar-count |authority-class [number-of-collects <1-10>] | authority-class-text [number-of-collects <1-10>] | authority-name [number-of-collects <1-10>] | authority-rd-length [number-of-collects <1-10>] | authority-rdata [number-of-collects <1-10> |**	Supported only for IPFIX.

Collect Type	Parameters	Size		Description
			width <1-128>] | authority-ttl [number-of-collects <1-10>] | authority-type [number-of-collects <1-10>] | authority-type-text [number-of-collects <1-10>] | bits | identifier | ns-count | op-code | qd-count | query-class [number-of-collects <1-10>] | query-class-text [number-of-collects <1-10>] | query-name [number-of-collects <1-10>] | query-type [number-of-collects <1-10>] |	
Private (continued)	**PEN <pen name>**	**DNS**	**query-type-text [number-of-collects <1-10>] | response-class [number-of-collects <1-10>] | response-class-text [number-of-collects <1-10>] | response-code | response-ipv4-addr [number-of-collects <1-10>] |response-ipv4-addr-text [number-of-collects <1-10>] | response-ipv6-addr [number-of-collects <1-10>] |**	Supported only for IPFIX.

Collect Type	Parameters	Size		Description
			response-ipv6-addr-text [number-of-collects <1-10>] |response-name [number-of-collects <1-10>] | response-rd-length [number-of-collects <1-10>] | response-rdata [number-of-collects <1-10> | width <1-128>] | response-ttl [number-of-collects <1-10>] | response-type [number-of-collects <1-10>] |response-type-text [number-of-collects <1-10>]>	
Private	**PEN <pen name>**	**HTTP**	**Response Code**	Supported only for IPFIX.
Private	**PEN <pen name>**	**HTTP**	**URL**	Supported only for IPFIX. For width, the range is from 1 to 250.
Private	**PEN <pen name>**	**HTTP**	**User Agent**	Supported only for IPFIX. For width, the range is from 1 to 250.
Private	**PEN <pen name>**	**SSL Certificate**	**<Issuer [width] | Issuer Common Name [width] | Serial Number | Serial Number Text | Signature Algorithm |Signature Algorithm Text | Subject [width] | Subject Algorithm | Subject Algorithm Text | Subject Alternative Name [width] | Subject Common Name |**	Supported only for IPFIX. For width of Issuer and Subject, the range is from 1 to 250. For width of Issuer Common Name, Subject Alternative Name, and Subject Common Name, the range is from 1 to 64.

Collect Type	Parameters	Size		Description
			[width] | Subject Key Size | Valid Not After | Valid Not After Text | Valid Not Before | Valid Not Before Text>	
Private	**PEN <pen name>**	**SSL Server**	**<Cipher | Cipher Text | Compression Method | Name Indication [width] | Session ID | Version | Version Text>**	Supported only for IPFIX. For width, the range is from 1 to 64.
Private	**PEN <pen name>**	**URL**	**[width]**	Supported only for IPFIX. For width, the range is from 1 to 250.
timestamp	**Sys-uptime First**			Supported for v9 and IPFIX.
	Sys-uptime First Last			Supported for v9 and IPFIX.
transport	**Destination Port**			Supported for v9 and IPFIX.
	ICMP	**IPv4 Code**		Supported only for IPFIX.
		IPv4 Code Type		Supported only for IPFIX.
		ipv6 Code		Supported only for IPFIX.
		ipv6 Type		Supported only for IPFIX.
	Source Port			Supported for v9 and IPFIX.
	TCP Flags	**[ACK] | [CWR] | [ECE] | [FIN] | [PSH] | [RST] | [SYN] | [URG]**		Supported for v9 and IPFIX.
	TCP	**ACK Number**		Supported only for IPFIX.
		Destination Port		Supported only for IPFIX.
		Header		Supported only for IPFIX.

Collect Type	Parameters	Size		Description
		Length		
		Sequence Number		Supported only for IPFIX.
		Source Port		Supported only for IPFIX.
		Urgent Pointer		Supported only for IPFIX.
		Window Size		Supported only for IPFIX.
	UDP	**Destination Port**		Supported only for IPFIX.
		Message Length		Supported only for IPFIX.
		Source Port		Supported only for IPFIX.

NetFlow Cacheless export using TCP protocol

Classic NetFlow operates in a cacheless mode and exports packets using TCP protocol. NetFlow monitor enables you to set the cache timeout to zero and makes the NetFlow to operate in a cacheless mode. Each packet received by the NetFlow is treated as a complete flow and the metadata extracted from the packet are immediately exported without caching the flow.

The two main advantages of a cacheless mode are as follows:

- The metadata are directly sent to the export buffer as they are not copied to cache.
- The cacheless mode provides more record space.

When NetFlow operates in a cacheless mode, the NetFlow sends metadata for every incoming packet, hence the number of flows are equal to the number of incoming packets.

For NetFlow to export packets in TCP, you should Configure Apps Exporter, which is also used by 3GPP CUPS and Tunneling.

Rules and Notes

- You must define at least 2 exporters, such as **apps exporter**, **apps netflow exporter**.
- You must provide same name for the exporters. The name must be followed by a number in the range from 1 to 9.

- You can add up to 64 apps exporters to one NetFlow Engine. Flows is load balanced across these exporters based on the inner IP address and the direction of the flow.
- You can export only IPv4 using TCP through the exporter.
- The maximum number of bytes exported for the inner payload data is 9600.
- If a record has any inner collects and the packet does not have any inner fields , the record will not be collected.
- In a cluster environment, when there is a NetFlow configuration, the backup-and-restore functionality does not work as expected. The restore operation does not push the configuration to the GigaSMART engine of non-master nodes. It is recommended to copy and paste the configuration backup file instead of running the “config text apply” operation.

For more information about the commands, refer to GigaVUE-OS CLI Reference Guide.

Configure Netflow Generation

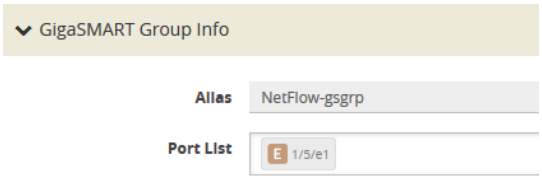
The following are the step for setting up a typical NetFlow Generation configuration with GigaVUE-FM,GigaVUE HC Series:

- Step 1: Configure a GigaSMART Group
- Step 2: Configure the NetFlow Exporter
- Step 3: Configure an IP Interface
- Step 4: Configure the NetFlow Record
- Step 5: Configure the NetFlow Monitor
- Step 6: Add the NetFlow Monitor to GigaSMART Group
- Step 7: Configure the GigaSMART Operation
- Step 8: Configure Mapping Rules to Filter Packets

Step 1: Configure a GigaSMART Group

Configure a GigaSMART Group using the following steps. you will use this GigaSMART Group in Step 6: Add the NetFlow Monitor to GigaSMART Group, where you assign a NetFlow Monitor to the group.

1. From the device view, select **GigaSMART > GigaSMART Groups**.
2. Click **New** to create a new GigaSMART Group.
3. Enter an alias to help identify this GigaSMART group. For example, Netflow-gsgrp
4. Select an engine port (the **e** port references the GigaSMART line card or module) Your GigaSMART group should look similar to the example shown in the following figure.



5. Click **Save**.

Notes:

- The GigaSMART Group can contain multiple GigaSMART engine ports.
- Only one NetFlow Generation Monitor can be configured per GigaSMART Group.
- Once a GigaSMART group is created, the Alias and port List values are no longer modifiable. You will need to recreate the group to change these values.

Step 2: Configure the NetFlow Exporter

Configure one or more NetFlow Generation Exporters. There can be up to six NetFlow Generation Exporters for each NetFlow Generation Monitor.

To access GigaSMART within GigaVUE-FM,GigaVUE HC Series, access a device that has been added to GigaVUE-FM,GigaVUE HC Series from the GigaVUE-FM,GigaVUE HC Series interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure the NetFlow Exporter, do the following:

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Exporters**.
2. Click **New**. The NetFlow Exporters page appears.
3. On the NetFlow Exporter page, enter the information for the exporter. Table 3: NetFlow ExporterConfiguration Fields describes the fields.

   NOTE: The NetFlow version must be configured with the same version of the Exporter and the Record. If no version is specified, version 9 is the default.

4. Under the **Filters** section, click **Add a Rule** to create a filter for the exporter.
5. Click **Save**.

Table 3: NetFlow ExporterConfiguration Fields

Field	Description
Alias	The alias name for the NetFlow Exporter.
Description	An optional description of the NetFlow record.
Format	The format is either NetFlow or CEF.
Version	The version is either NetFlow-v9, NetFlow-v5, or IPFIX.
Template Refresh Interval	After each template-refresh-interval, the record template is sent to the collector. Also, the option template is sent.
SNMP	Enables SNMP packet support on IP interfaces associated with the NetFlow Exporter.
Transport Protocol	The UDP port of the collector. This value cannot be changed.
IP Version	IP Version of the destination IP. You can select IPv4 or IPv6. Default is set as v4.
Destination IP	The IP address of the NetFlow/IPFIX collector. Default is set as 0.0.0.0.
Destination Port	Port for the destination IP. Default is set as 2055.
DSCP	The DSCP priority of the packet. Default is set as 0.
TTL	The Time to Live of the packet. Default is set as 64.

Step 3: Configure an IP Interface

In this step, you identify the collector port and configure it as a tool port, where the NetFlow collector will be connected, and then configure an IP interface. The steps are as follows:

1. Select the port to use and configure it as a tool port.
    a. Select **Ports > Ports > All Ports**.
    b. Click the **Quick Port Editor** button to open the Quick Port Editor.
    c. In the Quick Port Editor select the port to use for the IP interface, provide an alias to help identify the port (for example, NetFlow_Tunnel_Port), select **Tool** for the port type, and select **Enable**.
    d. Click **OK**.
2. Select **Ports > IP Interfaces**.
3. Click **New**.
4. On the IP Interface page, do the following:
    a. In the **Alias** and **Description** fields, enter a name and description for the IP interface.
    b. From the **Port** drop-down list, select the tool port that you configured in Step 1.

c. Select the type of IP interface as either **IPv4** or **IPv6**.

d. Enter the **IP Address**, **IP Mask**, **Gateway** address, and **MTU** value.

e. From the **GigaSMART Group** drop-down list, select the GigaSMART group you created in Step 1: Configure a GigaSMART Group.

f. From the Exporters drop-down list, select the NetFlow exporter you created in Step 2: Configure the NetFlow Exporter.

Step 4: Configure the NetFlow Record

Configure a NetFlow Generation Record that has the following:

- **match** parameters that identify unique flows
- **collect** parameters that identify fields you want to collect for the unique flows

To access GigaSMART within GigaVUE-FM,GigaVUE HC Series, access a device that has been added to GigaVUE-FM,GigaVUE HC Series from the GigaVUE-FM,GigaVUE HC Series interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure the NetFlow Record, do the following:

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Records**.
2. Click **New**. The NetFlow record page shown in Figure 45NetFlow Record Page displays.

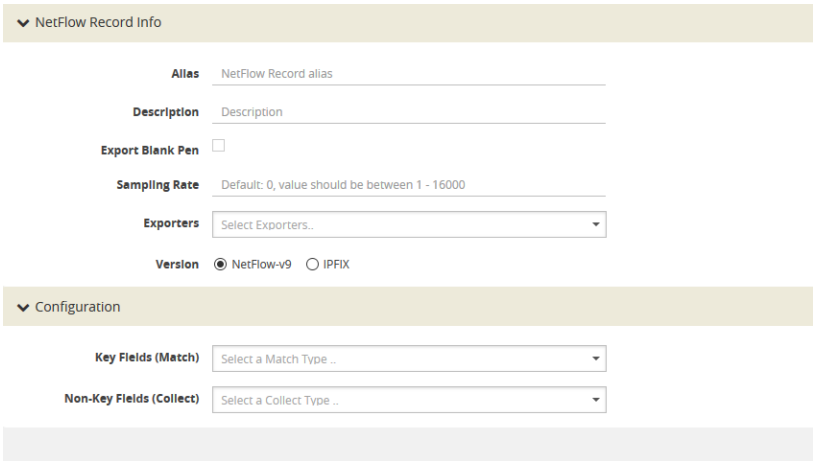


Figure 45 *NetFlow Record Page*

3. On the NetFlow Record page, do the following:

Table 4: NetFlow Record Configuration Fields

Field	Description
Alias	The alias name that would help identify the NetFlow record.
Description	An optional description of the NetFlow record.
Sampling Rate	The **Sampling Rate** is **multi-rate** only, and is specified as 1 in N, where N is the packet count. The packet count can be a number from 1 to 16000. Refer to Configure Netflow Generation. The **Sampling Rate** is disabled by default.
Exporter	Select the Exporter that you want from the **Exporters** menu.
Version	The version is either NetFlow-v9 or IPFIX. The NetFlow version must be configured with the same version of the Exporter and the Record. **NetFlow-v9** is the default. NetFlow-v9 and **IPFIX** let you configure Match/Key and Collect/Non-Key elements. Make sure that you configure the NetFlow version prior to configuring the match and collect parameters because the subsequent parameters depend on the NetFlow version configured.
Key Fields (Match)	The parameters that identify unique flows. The available Match/Key fields are based on the configured NetFlow version.
Non-Key Fields (Collect)	The parameters that identify what you want to collect for the unique flows. The number of Collect/Non-Key elements in a record can be up to 32. From the drop down select the following option: Exporter- You can collect IPv4 or IPv6 switch management interface address by including additional collects in their configuration. You can select any of the following options: **IPv4 address**—Adds new switch or router management IPv4 address. **IPv6 address**—Adds new switch or router management IPv6 address. For details about the match and collect parameters, refer to Configure Netflow Generation

Step 5: Configure the NetFlow Monitor

Configure a NetFlow Generation Monitor and associate the NetFlow Generation Record to the specified NetFlow Generation Monitor by doing the following:

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Monitors**.
2. Click **New**. The Monitors page displays.
3. On the Monitors page, do the following:
   a. Enter an **Alias** to identify the monitor.
   b. Enter a **Description** (optional).
   c. Configure the **Cache** parameters. Refer to Table 5: NetFlow Monitor Parameters.

d. Configure the **Sampling** parameters. Refer to Table 5: NetFlow Monitor Parameters.

e. Select the **Record** that you configured in Step 4: Configure the NetFlow Record.

4. Click **Save**.

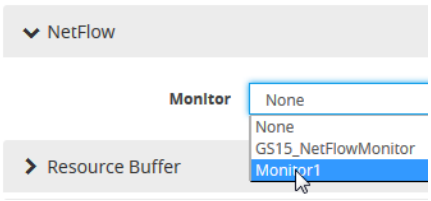
Table 5: NetFlow Monitor Parameters

Parameter	Description
Cache Type	Set as Normal.
Cache Timeout Active	Despite the flow being active, it is "flushed out" to the Exporter after this timeout, which is set in seconds.
Cache Timeout Inactive	Inactive flows are "flushed out" to the Exporter after this timeout, which is set in seconds.
Cache Timeout Event	Applies to the TCP flow. The flow is "flushed out" to the Exporter after detecting a FIN or RST.
Mode	Select the sampling mode that you want: • No sampling • Multi rate • Single rate
Single Sampling Rate	Refer to Configure Netflow Generation.

Step 6: Add the NetFlow Monitor to GigaSMART Group

Return to the GigaSMART Group configured in Step 1: Configure a GigaSMART Groupand set the NetFlow Monitor to the monitor created in Step 2: Configure the NetFlow Exporter.

1. From the device view, select **GigaSMART > GigaSMART Groups> GigaSMART Groups**.
2. Select the GigaSMART Group configured inStep 1: Configure a GigaSMART Group, and then click **Edit**.
3. Under GigaSMART Parameters, go to NetFlow. Click in the **Monitor** field and select the NetFlow monitor configured in Step 5: Configure the NetFlow Monitor as shown in the following figure.



4. Click **Save**.

Step 7: Configure the GigaSMART Operation

Define a GigaSMART operation to enable NetFlow Generation. If combining NetFlow with APF or De-duplication GSOPs, make sure that you select both operations when creating the GigaSMART Operation.

To access GigaSMART within GigaVUE-FM,GigaVUE HC Series, access a device that has been added to GigaVUE-FM,GigaVUE HC Series from the GigaVUE-FM,GigaVUE HC Series interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To configure the GigaSMART Operation, do the following:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP**).
2. Click **New**. The GigaSMART Operations (GSOP) page displays. (Refer to Figure 46GigaSMART Operation (GSOP) Page.)
3. On the GigaSMART Operations page, do the following:
   a. In the **Alias** field, enter a alias to help identify this gsop.
   b. In the **GigaSMART Groups** field, select the gsop configured in Step 1: Configure a GigaSMART Group.

      aIn the **GigaSMART Operations (GSOP)** field, select **NetFlow**. The NetFlowGigaSMART Operation is enabled by default as shown in Figure 46GigaSMART Operation (GSOP) Page.

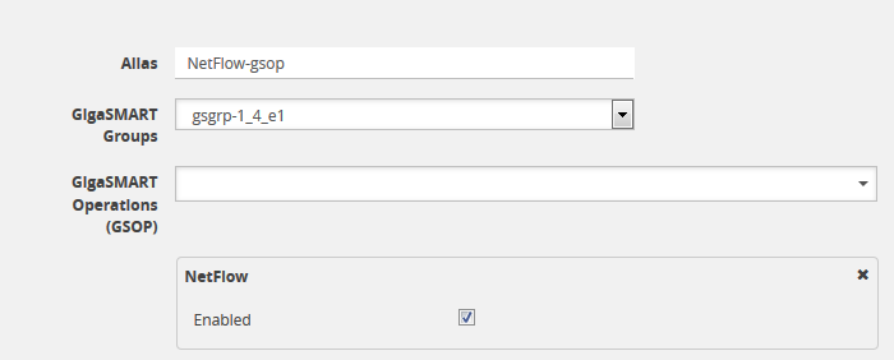


Figure 46 *GigaSMART Operation (GSOP) Page*

4. Click **Save**.

Step 8: Configure Mapping Rules to Filter Packets

To add Flow Mapping® rules to filter packets that are needed to run NetFlow, configure a map and associate the map to the IP interface with tool port.

For more detailed information about Flow Mapping®, refer to About Flow Mapping® and Manage Maps.

Notes:

- For a single NetFlowGigaSMART Operation, make sure that you create a Regular By Rule map. When combining with APF or De-duplication, use First Level or Single Level map types.
- Make sure that the other combining GigaSMART Operations are configured before creating maps using NetFlow.
- When combining NetFlow with APF or De-duplication, create virtual ports to use with the second level maps.
- The destination tool port must be the IP interface with tool port identified in Step 3: Configure an IP Interface

For second level maps, you will need to create virtual ports. To create virtual ports, do the following:

1. From the device view, select **GigaSMART > Virtual Ports**.
2. Click **New**. The Virtual Ports page displays.
3. Enter an alias in the **Alias** field to identify the virtual port.
4. In the **GigaSMART Groups** field, select the GigaSMART Group configured in Step 1: Configure a GigaSMART Group.
5. Click **Save**.

To configure mapping rules to filter packets, do the following:

1. Select **Maps > Maps > Maps**.
2. Click **New** to create a new map.
3. On the New Map page, do the following:
    a. Enter an alias in the **Alias** field and select the map **Type** and **Subtype**.
    b. Specify **Source** and **Destination** ports.
    c. In the **GigaSMART Operations (GSOP) field**, select the GigaSMART Operation configured in Step 7: Configure the GigaSMART Operation.
    d. Click **Add a Rule** to add the rules needed for the map.

4. Click **Save**.

GigaSMART Packet Slicing

Required License: Base

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

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

GigaSMART operations with a **Slicing** selected truncate packets after either a specified header/layer and offset (a **relative** offset) or at a specific offset. Slicing operations are typically configured to preserve specific packet header information, allowing effective network analysis without the overhead of storing full packet data.

Packets can have multiple variable-length headers, depending on where they are captured, the different devices that have attached their own headers along the way, and the protocols in use (for example, IPv4 versus IPv6). Because of this, slicing operations with a hard-coded offset will not typically provide consistent results.

To address this, the GigaSMART lets you configure packet slicing using **relative offsets** – a particular number of bytes after a specific packet header (IPv4, IPv6, UDP, and so on). The GigaSMART parses through Layer 4 (TCP/UDP) to identify the headers in use, slicing based on the variable offset identified for a particular header instead of a hard-coded number of bytes.

Keep in mind the following when configuring GigaSMART operations with a **Slicing** component:

Feature	Description
Protocol	The following are the protocols that you can select for from the protocol drop-down list: o **IPV4** – Slice starting a specified number of bytes after the IPv4 header. o **IPV6** – Slice starting a specified number of bytes after the IPv6 header. o **UDP** – Slice starting a specified number of bytes after the UDP header. o **TCP** – Slice starting a specified number of bytes after the TCP header. o **FTP** – Identify using TCP port 20 and slice payloads using offset from the TCP header. o **HTTPS** – Identify using TCP port 443. Slice encrypted payloads using offset from the TCP header. o **SSH** – Identify using TCP port 22. Slice encrypted payloads using offset from the TCP header.

Feature	Description
	The GigaSMART can provide slicing for GTP tunnels, provided the user payloads are unencrypted. Both GTPv1 and GTPv2 are supported – GTP' (GTP prime) is not supported. Keep in mind that only GTP-u (user plane packets) are sliced. Control plane packets (GTP-c) are left unmodified because of their importance for analysis. o **GTP** – Slice starting a specified number of bytes after the outer GTP header. o **GTP-IPV4** – Slice starting a specified number of bytes after the IPv4 header inside the encapsulating GTP packet. o **GTP-UDP** – Slice starting a specified number of bytes after the UDP header inside the encapsulating GTP packet. o **GTP-TCP** – Slice starting a specified number of bytes after the TCP header inside the encapsulating GTP packet.
Slicing Offsets	You can specify either a **relative** offset or a **static** offset for the start of the packet slice: • **Static** offsets begin slicing a specific number of bytes from the start of the packet. Choose a static offset by setting **protocol** to **none** and supplying an **offset** from **<64~9000>** bytes. • **Relative** offsets begin slicing a specified number of bytes from the end of a particular header – IPv4, IPv6, and so on. Choose a relative offset by selecting any of the values listed for the **protocol** argument, along with an offset of **<4~9000>** bytes from the end of the specified layer:
Recalculated CRC	GigaSMART packet slicing automatically calculates and appends a new four-byte Ethernet CRC based on the sliced packet's length and data and uses it to replace the existing CRC. This way, analysis tools do not report CRC errors for sliced packets. **Note**: The minimum relative offset is 4 bytes to allow the recalculated CRC to be added. The packet is sliced at the relative offset, and then the recalculated 4 bytes CRC is added to the sliced packet.
GigaSMART Trailer	Slicing operations can optionally include the GigaSMART Trailer. If you do elect to include the GigaSMART Trailer, it will include the original packet length before slicing. **Note**: Refer to How to Use GigaSMART Trailers in GigaSMART Operations Guide for details on when the GigaSMART Trailer is required for a GigaSMART Operation as well as the information found in it.
In Combination with Masking	Slicing operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to Groups of GigaSMART Engine Ports in GigaSMART Operations Guide for details.

Example – GigaSMART Packet Slicing

This example creates a GigaSMART slicing operation named **IPv6_Headers**. This operation truncates all packet data starting four bytes after the IPv6 header. The sliced packet would include the DLC, IPv6, and TCP headers, which are often enough for analysis needs.

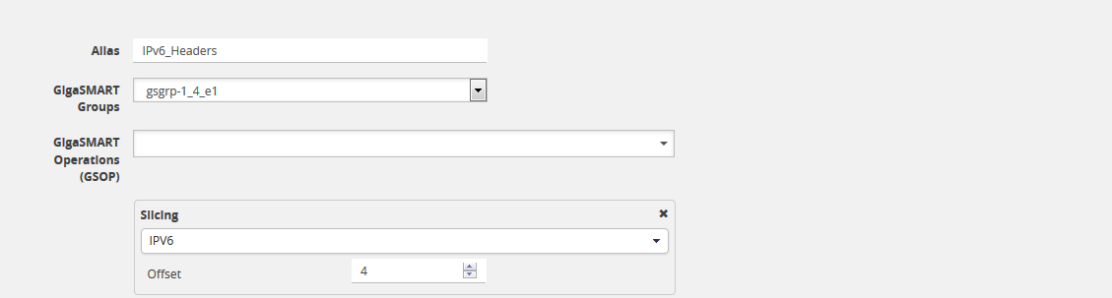


Figure 47 *GigaSMART Operations (GSOP) Page with Slicing Selected*

Display Slicing Statistics

To display slicing statistics, select **GigaSMART > GigaSMART Operations > Statistics**. The statistics for slicing will be in the row labeled Slicing in the GS Operations column.

Refer to Slicing Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

GigaSMART Advanced Flow Slicing

Required License: Advanced Flow Slicing

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

Advanced Flow Slicing (AFS) allows you to slice traffic on multiple network protocols, each with different offset. In Advanced Flow Slicing, you can configure multiple protocols within a profile with rules for slicing of a packet which in turn reduces the number of GigaSMART Operations (GSOP) and the maps to be configured. In this feature, the slicing of packets occurs after the packet count has reached a configured value thereby preventing the slicing of control information and other important data of the networking protocols.

Limitations

- Supports up to five enhanced profiles.
- Supports up to ten protocol rules for each profile.
- Slicing profile cannot be edited after attaching it to GSOP.
- Supports up to 80 million sessions per GigaSMART group.
- When enhanced slicing or AFS profile is configured with 7-tuple, the protocol rule must be configured with "flow-session inner" only.

NOTE: Only 5-tuple hashing is supported for Gen2. The 7-tuple hashing is supported for Gen3.

NOTE: It is recommended not to edit an Advanced Flow Slicing (AFS) or the enhanced-slicing profile once created. For any changes to a profile, the user must always delete and recreate with new configurations.

Create Advanced Flow Slicing Profile

To create an advanced flow slicing profile, follow these steps:

1. Go to **Physical > Physical Nodes**.
2. Click the required cluster ID.
3. On the left navigation pane, go to **System** > **GigaSMART> Advanced Flow Slicing**. The Advanced Flow Slicing page appears.
4. Click **New**.
   - Enter an allias.
   - Enter the Maximum Sessions. In a profile, you can configure from 4 to 80 million sessions.
   - Select the Protocol from the drop-down list.
   - Choose the protocol type as Inner or Outer.
   - Enter the value for Offset. The value specifies the number of bytes that should be sliced after the protocol header. The value ranges from 64 to 9000, if there is no protocol selected. The value ranges from 0 to 9000 when other protocol is selected.
   - Select the Flow Session from the drop-down list. For a profile with flow-session defined, a session is created when a packet is received and when there is no existing flow-session for that flow.

     Slicing or dropping starts on the next packet of a session after the number of packets reaches the specified value in the packet count.
   - Choose either to Slice or Drop the session. The default action is Slice.

     The slicing or drop occurs for the first matching rule in a profile.
   - Enter the timeout sessions from 10 to 300 seconds. The default value is 30 seconds.
   - Enter the Skip Packet Count value to slice or drop the packet after it reaches the given count value.

   Each profile must have atleast one protocol field. You can add a maximum of 10 rules in a profile. Click **+** to add new rules in a profile.
5. Click **OK**.

Slicing or drop cannot be performed in the following conditions:

- When the traffic does not match any rule in the profile.
- When a packet does not contain the configured inner or outer IP and l4-port in a defined flow-session.

> **NOTE:** It is recommended not to edit an Advanced Flow Slicing (AFS) or the enhanced slicing profile once created. For any changes to a profile, you must always delete the existing profile and recreate the profile with the new configurations.

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

Example-GigaSMART Advanced Slicing

This example shown in creates a GigaSMART Advanced slicing operation named **eslice-gtpu**. In this operation, slicing occurs at HTTPS (port 443) over GTPu-TCP and GTPu-UDP traffic after 10 packets for each TLS connection. This operation truncates all packet data starting 40 bytes after the TCP header. The slicing starts after TCP header gets encapsulated by GTP.

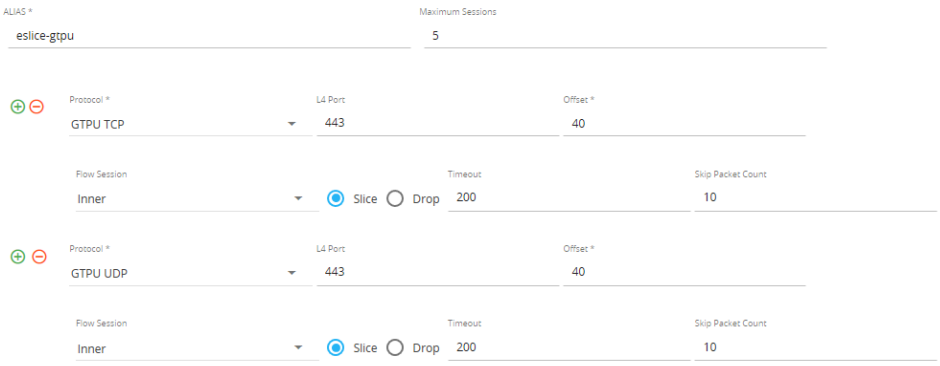


Figure 48 *Advanced Flow Slicing*

Display Slicing Statistics

To display Advanced Slicing statistics, select **System** > **GigaSMART> Advanced Flow Slicing**> **Statistics**. You can view the statistics of the Advanced Flow Slicing.

GigaVUE Inline Solutions

The Gigamon Inline Solutions Guide describes how to plan, configure, monitor, and troubleshoot inline deployments using GigaVUE nodes and related software constructs.

It is intended for network and security administrators who design and operate inline inspection environments, and provides task-oriented procedures, conceptual background, and reference information needed to deploy and maintain Gigamon inline solutions effectively.

GV-ISSL

GigaVUE Inline Solutions

A GigaVUE inline solution refers to network configuration where a GigaVUE node is placed directly in the data path ("inline") between two network segments, allowing it to capture and analyze live network traffic while simultaneously directing traffic to security tools like firewalls or intrusion prevention systems (IPS) for inspection, all while maintaining network connectivity and minimizing disruption; essentially acting as a "traffic broker" for inline security tools with fail-safe bypass capabilities in case the security appliance fails.

The various solutions are:

- Resilient Inline Arrangement - Resilient inline arrangement is a method of configuring and deploying inline threat prevention tools for dual-path, redundant network architectures
- Hardware Security Modules (HSM) - Hardware Security Modules (HSMs) with its Inline SSL (iSSL) decryption solution enhances the security and management of cryptographic keys used for decrypting SSL/TLS traffic.
- Internet Content Adaptation Protocol (ICAP) - The ICAP protocol, integrates with DLP systems for enhanced security by allowing inline inspection of decrypted TLS/SSL traffic.
- Gigamon Resiliency for Inline Protection - GRIP is an Inline Bypass solution that connects two GigaVUE-FM,GigaVUE HC Series 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-FM,GigaVUE HC Series nodes maintains traffic monitoring by inline tools when one of the nodes is down.
- Inline TLS/SSL Decryption - Inline TLS/SSL decryption detects, intercepts, filters, decrypts, analyzes, and re-encrypts encrypted traffic for security purposes. The various Inline TLS/SSL Decryption solutions are as follows:
  - Inline TLS/SSL Decryption Solution with Layer 2 Transparent Tools
  - Inline TLS/SSL Decryption Solution with Layer 3 Tools (NAT-PAT)
  - Inline TLS/SSL Decryption Solution with RIA
  - Entrust nShield and Thales- Luna HSM for TLS/SSL Decryption for iSSL
  - Inline TLS/SSL Decryption Solution with ICAP Client

Software Constructs for Inline Solutions

In Gigamon Inline solutions, a software construct is a configuration element that helps manage network traffic. These constructs include things like inline networks, inline tools, inline tool groups, and inline network groups. They are not physical devices; rather, they represent logical groupings of ports and tools that dictate how traffic is processed and inspected within the system.

For instance, an "inline tool" is formed by a pair of ports, along with the inline tool connected to those ports. This setup is configured on the GigaVUE HC Series node and includes specific attributes. An "inline network" consists of two network ports designed for inline monitoring.

Administrators can configure these constructs via a management interface or command-line interface (CLI). They play a crucial role in creating and managing inline bypass and other flexible arrangements within Gigamon deployments. These constructs allow for the definition of traffic flows, inspection policies, and failover behaviors, independent of the actual physical hardware in use.

Below is the list of software components used in Gigamon Inline solutions:

- Inline Network and Inline Network Ports
- IP Interface
- Inline Network Link Aggregation Group (LAG)
- Inline Network Bundle
- Inline Tool Ports and Inline Tools
- Inline Tool Group
- Inline Single VLAN Tag

Inline Network and Inline Network Ports

An inline network forms a logical connection between two inline network ports—Port A and Port B—that carry live production traffic. These ports operate at the same speed, use the same medium (fiber or copper), and reside on the same node. Traffic entering one port exits through the other, enabling packets to be inspected or modified by inline tools without disrupting service.

Inline networks can be configured as protected or unprotected, depending on whether failover protection is required:

- *Protected* inline networks provide high availability by automatically redirecting traffic through a bypass path if a tool or link fails.

- *Unprotected* inline networks do not offer automatic failover; traffic stops if a link or tool fails until manual intervention restores service.
- *Mixed* inline networks is a combination of protected and unprotected inline networks supported in Inline Bypass solutions (Classic).

When an inline network connects directly to inline tools, traffic continues to flow based on the configured failover action.

- **Drop**—All traffic through the inline network is stopped. No packets are exchanged between the inline network ports or with the inline tools. Traffic entering Port A or Port B is dropped.

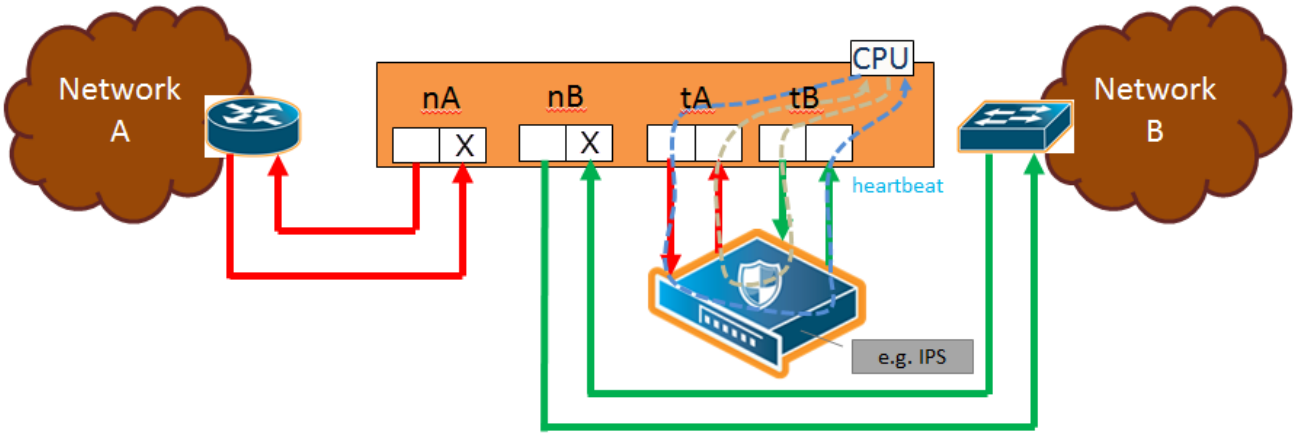


Figure 49 *Traffic Path of Drop*

- **ByPass**—Traffic is forwarded directly between the inline network ports, depending on how inline maps are configured:
  - **No maps or pass-all maps**: All traffic entering Port A is forwarded to Port B, and traffic entering Port B is forwarded to Port A through a logical bypass.
  - **Maps with drop conditions**: Only traffic that would have been forwarded in To Inline Tool mode is passed through. Packets that were dropped in that mode remain dropped here as well.

In both cases, no traffic is sent to the inline tools.

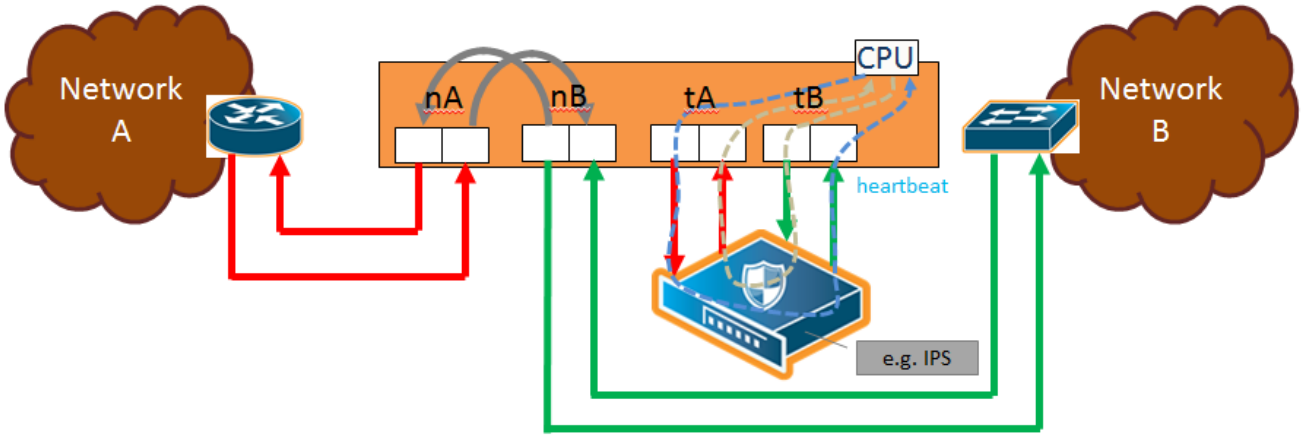


Figure 50 *Traffic Path of Bypass*

- **ByPass with Monitoring**—This mode behaves like standard bypass but also sends a copy of the traffic to the inline tools:
  - **No maps or pass-all maps**: All traffic is forwarded as in bypass mode, while a copy is sent to the inline tools for monitoring.
  - **Maps with drop conditions**: Only traffic that would not have been dropped in To Inline Tool mode is forwarded and copied for monitoring.

In both monitoring cases, no packets are received from the inline tools.

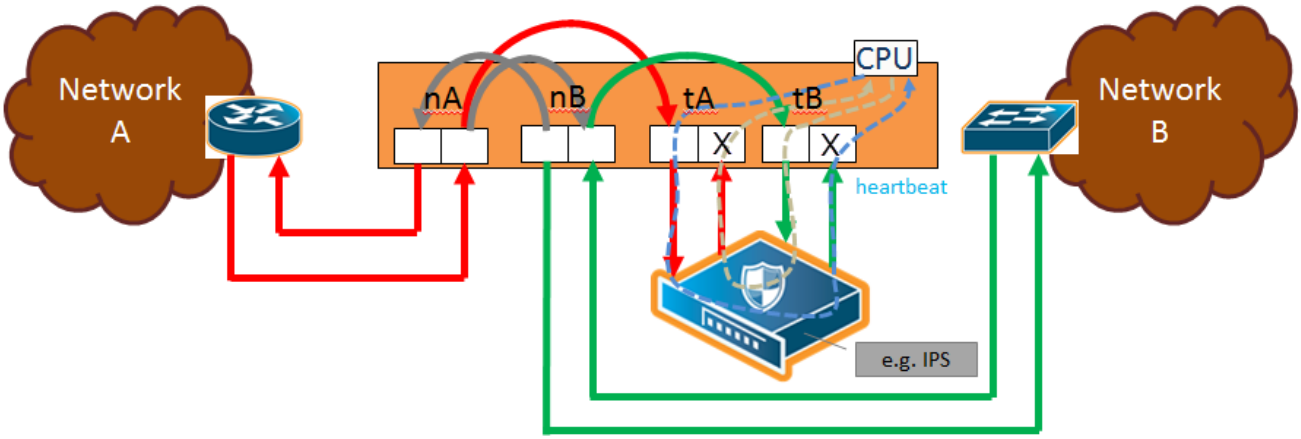


Figure 51 *Traffic Path of Monitoring*

- **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

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 49Traffic Path of Drop to Figure 52Traffic 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.

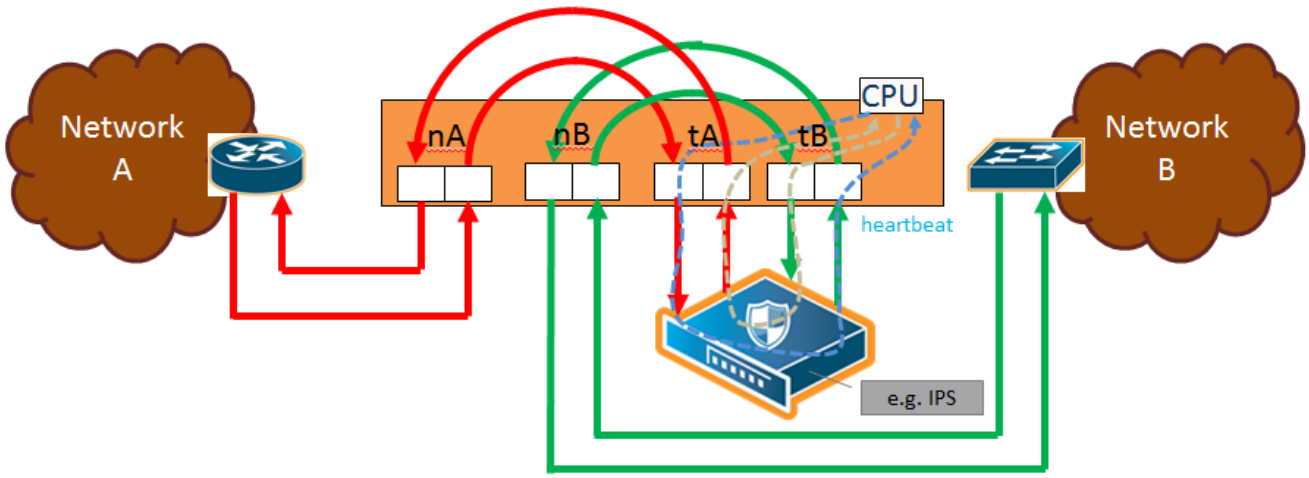


Figure 52 *Traffic Path of To-Inline-Tool*

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.

Resiliency Features

Inline Networks include resiliency mechanisms to handle failures gracefully:

Network Port Link Status Propagation Parameter

Link status propagation controls how the link state is handled for inline network ports. It is enabled by default.

- When enabled, if one side of the inline network fails, the failure is also applied to the other side.For example, if traffic stops on one port and cannot reach the inline tools, the opposite port is also marked down.
- When the original link comes back up, the opposite side is restored as well.
- The node will not forward packets to inline tools until both sides of the link are active.

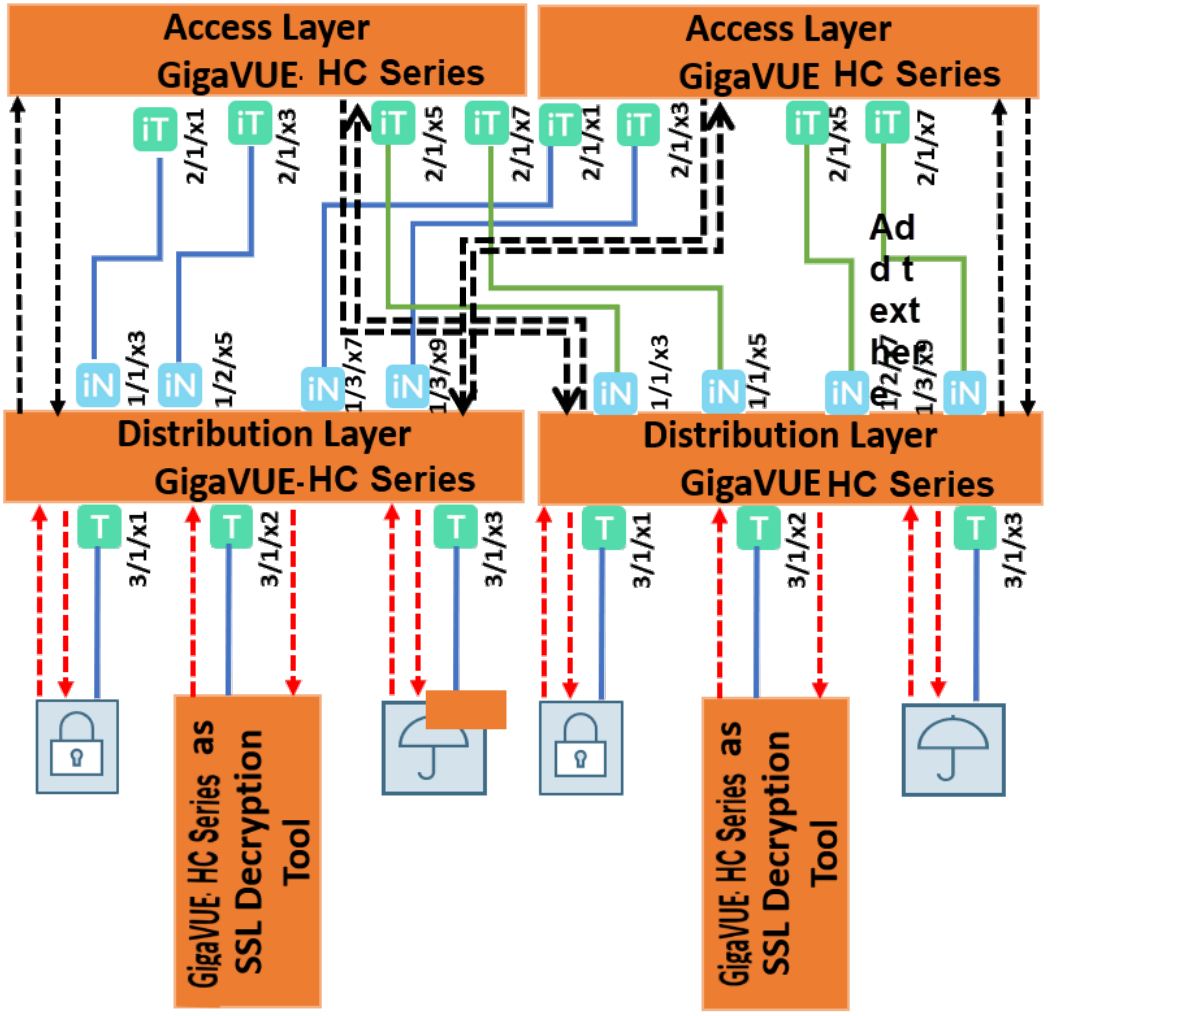
You can enable or disable this feature by selecting **Link Failure Propagation** when configuring an inline network port.

Heartbeat Support Between GigaVUE Nodes

Heartbeat packets are lightweight health-check signals exchanged between inline nodes and inline tools. They verify that traffic flows correctly through the inline path. If heartbeats stop, the system assumes a failure and automatically switches traffic to bypass mode, preventing packet loss.

When heartbeats resume, traffic again passes through the inline tools. This feature quickly detects tool or link failures and protects live traffic.

Following figure illustrates an example of a topology with GigaVUE nodes placed at three different layers.



GigaVUE HC Series means it can be either GigaVUE-HC3, GigaVUE-HC1 or GigaVUE-HC1-PLUS

Heartbeats in a Multi-Layer Inline Topology

In a multi-layer setup, heartbeat packets help maintain traffic integrity and support automatic failover.

- The access layer node captures live traffic, sends it to the tool layer for processing, and then returns it to the network.
- The distribution layer node sits between the access and tool layers, distributing and managing traffic flows.

In this example, the tool layer node performs TLS/SSL decryption. Heartbeats monitor link health between each layer.

At the access layer, ports facing the distribution layer are configured as inline tool ports. At the distribution layer, ports facing the access layer are configured as inline network ports.

Heartbeat packets flow from the inline tool port pair (access layer) to the inline network port pair (distribution layer). If the distribution layer is in a healthy forwarding state, the heartbeat is returned; otherwise, it is dropped to signal a fault. The same principle applies between the distribution and tool layers.

Here, the distribution layer ports facing the tool layer are inline tool ports, while the tool layer ports facing back are inline network ports.

Heartbeats confirm the availability of both the tool devices and their processing engines.

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 6: 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-FM,GigaVUE HC Series node. Only applies to protected inline networks.
disable	drop	any status	any status	Forwarding state—DISABLED If the inline network is configured with a traffic path of drop, no traffic is exchanged between side A network and side B network because all packets coming to the inline network ports are dropped. **Note**: If the inline network is configured with link failure propagation set to true (enabled), the status of the inline network ports will be determined by the status of the far-end ports connected to the inline network ports. If both far-end ports are up, then both inline network ports will be up. If any far-end ports are down, then both inline network ports will be down.
disable	bypass, monitoring, or to-inline-tool	at least one far-end port is down	any status	Forwarding state—DISCONNECTED When one of the inline network ports is down due to a link down caused by a far-end device, no traffic is exchanged between side A network

physical-bypass	traffic-path	Status of far-end ports connected to inline network ports	Status of inline tool side	Description
				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.
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.

physical-bypass	traffic-path	Status of far-end ports connected to inline network ports	Status of inline tool side	Description
				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.

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	at least one inline tool involved in the inline bypass solution is disabled or associated with any inline tool port that is down or associated with the heartbeat protocol in a down state	Forwarding state—FAILURE-INTRODUCED DROP No traffic is exchanged between side A network and side B network. All packets coming to the inline network ports are dropped as a result of an inline tool port failure or heartbeat failure or inline tool being disabled that led to the network-level drop due to the configured failover actions.
disable	to-inline-tool	both far-end ports are up	at least one inline tool involved in the inline bypass solution is disabled or 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.

IP Interface

An IP interface acts as a logical network interface that enables communication between devices using Internet Protocol (IP) addresses. It connects to a single IP peer within the local subnet and can reach destinations outside the subnet through the configured default gateway. Each IP interface associates an IPv4 or IPv6 address with a physical port—either a network or tool port—on a GigaVUE device.

IP interfaces are essential for GigaSMART operations such as encapsulation, decapsulation, tunnel initiation or termination, and NetFlow export. They are configured with attributes such as IP address, subnet mask, gateway, and MTU, and can be associated with GigaSMART groups and NetFlow exporters.

These interfaces handle the sending and receiving of traffic for advanced packet processing features including GigaSMART tunnels, NetFlow generation, and inline SSL decryption. Administrators can monitor interface status and performance metrics to ensure proper operation and simplify troubleshooting in visibility workflows.

NOTE: IP interfaces are supported on GigaVUE HC Series nodes but are not supported on TA Series nodes.

Inline Network Link Aggregation Group (LAG)

A Link Aggregation Group (LAG) combines multiple physical ports to act as one logical, high-bandwidth path. This approach improves throughput, distributes traffic across ports, and ensures link redundancy—so if one port fails, traffic continues through others.

When inspection or processing is required on aggregated links, the Flexible Inline Network LAG feature allows you to treat multiple inline networks as a single logical entity. Instead of configuring each inline network individually, you can manage them together.

This unified configuration makes it possible to:

- Apply consistent traffic forwarding and monitoring policies across all member links.
- Simplify load balancing and failover handling within inline deployments.
- Use the inline network LAG as a source in flexible inline maps for traffic orchestration.

Traffic entering a LAG is grouped and delivered to inline tools with the same VLAN ID. The response traffic is hashed and returned through the appropriate member links, ensuring bidirectional consistency and high availability.

Inline Network Link Aggregation Group— Rules and Notes

When configuring or using an inline network LAG, keep the following in mind:

1. Do not mix protected and unprotected inline networks or networks with different speeds within the same LAG.
2. Inline TLS/SSL Decryption is not supported with inline network LAG. Use an Inline Network Bundle instead.
3. It is strongly recommended that the first flexible inline map bypasses the LACP/PAgP protocols between the inline network peers.
4. When Bypass LACP and LLDP is enabled, all protocol packets with MAC Destination 01-80-C2-XX-XX-XX are bypassed.
5. When CDP/LLDP Bypass is enabled, neighborship discovery will not be established on the corresponding inline networks.

NOTE: Inline Network LAGs are only supported in Flexible Inline Arrangements. Classic Inline Bypass does not support the grouping of inline networks into LAGs. Classic deployments can monitor or protect individual inline networks, but they cannot aggregate them into a single logical link.

Inline Network Bundle

The Flexible Inline Network Bundle feature lets you group several inline networks into one bundle. You can then use this bundle as the source in a flexible inline map. The system automatically creates separate inline maps for each inline network in the bundle. These maps are built using the Tool Side VLAN tags and the rules you set when creating the bundle.

NOTE: This feature is only available in Flexible Inline Arrangements. It is not supported in Classic Inline Bypass.

Inline Tool Ports and Inline Tools

An **inline tool** is a device or a sequence of tools placed directly in the live traffic path between inline network ports for inspection, security, or optimization. Each inline tool is built from a pair of inline tool ports (Port A and Port B) that operate at the same speed, use the same medium (fiber or copper), and reside on the same node as the associated inline network ports. The paired ports and connected tool form a logical inline entity that processes packets passing through it.

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 an inline tool fails, traffic that would normally pass through it is redirected to the bypass path between side A and side B of the inline network. Use this mode for configurations that involve multiple inline tools associated with an inline network or inline network group using rule-based maps. In Flexible Inline, the same action applies to the failed tool instance in the flow sequence. For configurations using map pass-all, ToolBypass behaves the same as NetworkBypass.
- **NetworkBypass**:When an inline tool fails, all traffic that would normally pass through the inline network in a normal forwarding state is redirected to the bypass path. Traffic arriving at side A is forwarded to side B, and vice-versa, ensuring minimal disruption. This behavior is identical in both Classic and Flexible Inline modes.
- **ToolDrop**: When an inline tool fails, only the traffic mapped to that specific tool is dropped. Use this option when multiple inline tools share a network or tool group through rule-based or sequence maps. For map pass-all configurations, ToolDrop behaves the same as NetworkDrop.
- **NetworkDrop**:When an inline tool fails, all traffic associated with that inline network or group is dropped. Use this mode when you need to stop all packet forwarding after a failure.
- **NetworkPortForcedDown**:When an inline tool fails, the inline network ports are forced down. After a device reload, these ports become operationally up even if the tool remains down. To maintain correct behavior, manually force the inline network ports down again.

NOTE: The default failover action is **ToolBypass**. The bypass path operates between the side A and side B ports of the inline network.

When Flexible Inline Flow Mapping is used, failover actions apply to each tool within the defined sequence:

- **ToolBypass**: Traffic mapped to the failed tool is redirected to the bypass path. Other tools in the sequence continue to process their assigned traffic.

- **ToolDrop**: Traffic mapped to the failed tool is dropped, while traffic for healthy tools remains unaffected.

This model ensures that the failure of one tool does not interrupt traffic for other inline tools in the same sequence or map.

Inline Tool Recovery Modes

Inline tools detect traffic path failures between their port pairs and divert traffic as needed. After recovery, you can control how traffic returns to the inline tool using recovery modes. The Recovery Modes are as follows:

- Automatic: When the tool returns to a healthy state, traffic automatically resumes through it without user intervention. This is the default recovery mode.
- Manual: When recovery mode is set to manual, the failover action (for example, Bypass) remains active even after the tool is restored.

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 53Automatic and Manual Inline Tool Recovery from Failover for automatic and manual inline tool recovery from failover.

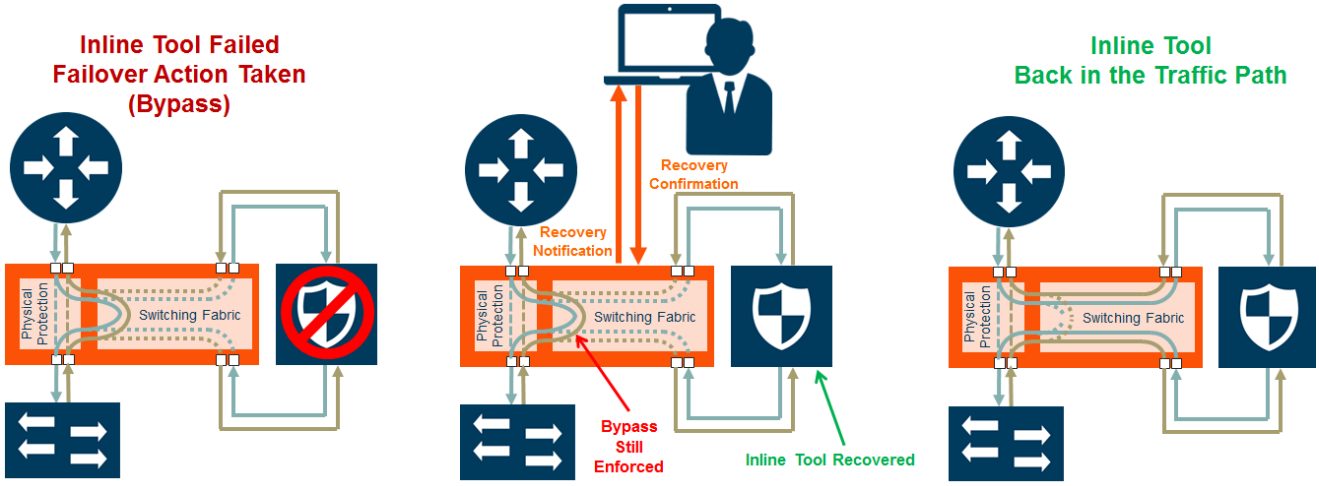


Figure 53 *Automatic and Manual Inline Tool Recovery from Failover*

The left side of Figure 53Automatic 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 53Automatic 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 53Automatic 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.

Inline Tool Group

NOTE: Resilient weighted hashing, single VLAN tag, and shared mode flexibility are supported only in Flexible Inline TLS/SSL deployments.

An inline tool group is a configuration that combines multiple inline tools into a single logical entity. Traffic is distributed among the tools using hardware-calculated hash values, and if one tool fails, the remaining tools in the group continue processing traffic. The ports of tools in a group must run at the same speed and medium and reside on the same node as the inline network ports.

In a clustered environment, the group must be configured on the cluster leader and cannot span multiple nodes. Resilient weighted hashing allows you to assign equal, relative or percentage-based weights to tools so traffic shifts appropriately when a tool fails, provided the group meets its Minimum Healthy Group Size.

If the group size falls below that threshold, traffic is handled according to the configured Failover Action: **Tool Bypass, Network Bypass, Tool Drop, Network Drop or Network Port Forced Down**. When configuring the tool group, you choose the traffic path mode (Drop, Bypass, Monitoring or To Inline Tool) and select options such as Enable, Release Spare if Possible, Minimum Healthy Group Size, Hash algorithm and spare tool assignments.

Inline tools connect to the GigaVUE-FM,GigaVUE HC Series node through inline bypass (BPS) modules, available on the GigaVUE-HC1, GigaVUE-HC3.

The inline bypass arrangements supported by inline TLS/SSL decryption are as follows:

- single inline network. The inline network ports can be protected, unprotected, or a mix of protected and unprotected.
- inline network group, consisting of multiple inline networks
- single inline tool
- inline tool group, consisting of multiple inline tools over which traffic is distributed

- inline tool group with a spare inline tool, in which the failure of one tool in the inline tool group will trigger a failover to the spare
- inline series, in which traffic is guided through inline tools in a particular order

In summary, inline TLS/SSL decryption can be deployed in any combination of inline network and inline network group with any inline tool, inline tool group, or inline series.

NOTE: The TLS/SSL sessions will fail, if a BPS card is shut down and brought back up with an Inline TLS/SSL configuration.

Inline tool ports can be configured in shared mode. When an inline tool is shared (true), the decrypted traffic will be VLAN tagged. The connected inline tool is expected to receive VLAN tagged packets instead of untagged packets. There is an extra outer VLAN tag added to the packet, which the connected inline device needs to see. When an inline tool is not shared (false), the extra VLAN tag is not added. This allows untagged traffic to be sent to the tool ports. Use false for inline tools that are not able to handle more than one VLAN tag, such as Q-in-Q tagged packets. For tagless mode, if an inline tool is involved in an inline TLS/SSL map, the inline tool cannot be used in any other classic inline map.

Refer to Figure 54Simple Inline Tool Arrangement toFigure 56Serial Inline Tool Arrangement for inline tool arrangements. Encrypted traffic is shown in solid lines, decrypted traffic is shown in dotted lines.

Figure 54Simple Inline Tool Arrangement shows a simple inline tool arrangement with one inline tool connected to the GigaVUE-FM,GigaVUE HC Series node. Traffic is decrypted and sent to the inline tool.

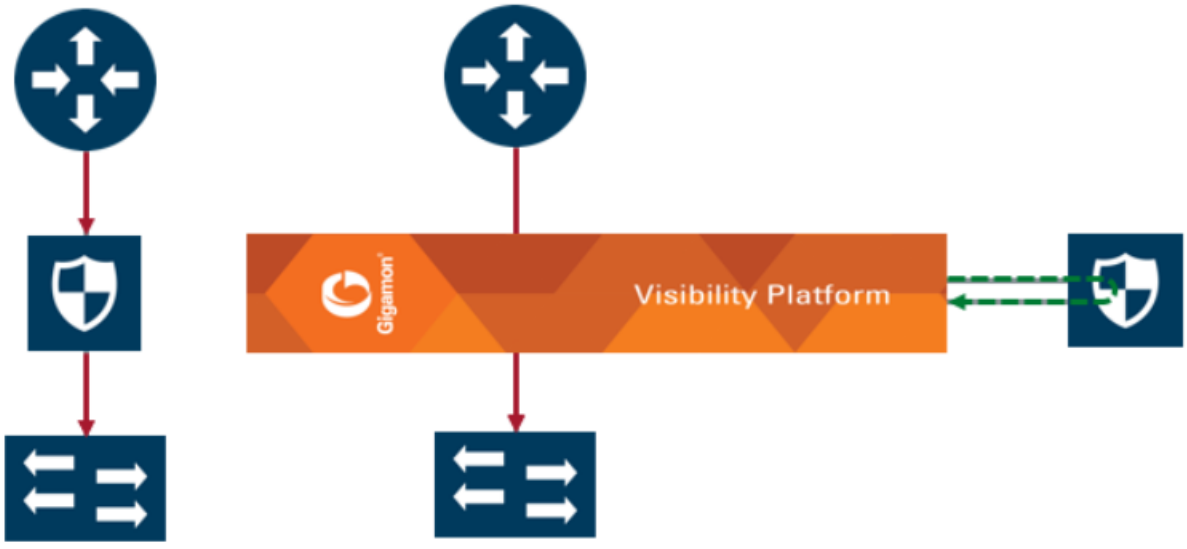


Figure 54 *Simple Inline Tool Arrangement*

Figure 55Multiple Inline Tool Arrangement shows a multiple inline tool arrangement with three inline tools connected to the GigaVUE-FM,GigaVUE HC Series node. Decrypted traffic is distributed across the tools.

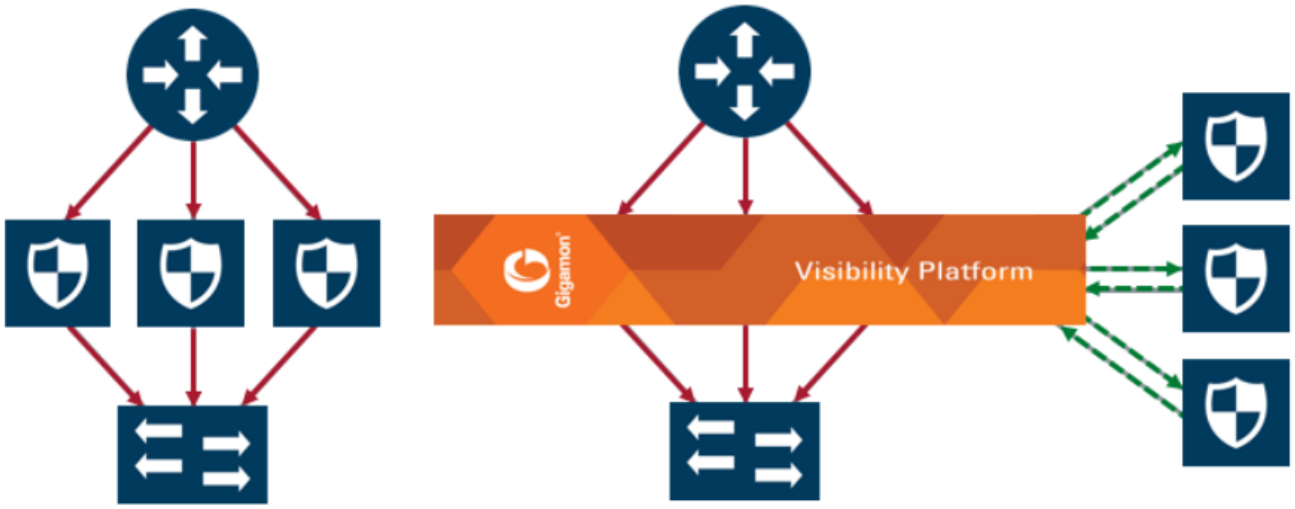


Figure 55 *Multiple Inline Tool Arrangement*

Figure 56Serial Inline Tool Arrangement shows a serial inline tool arrangement in which traffic is decrypted on the GigaVUE-FM,GigaVUE HC Series node, sent serially through the inline tool, and then re-encrypted on the GigaVUE-FM,GigaVUE HC Series node.

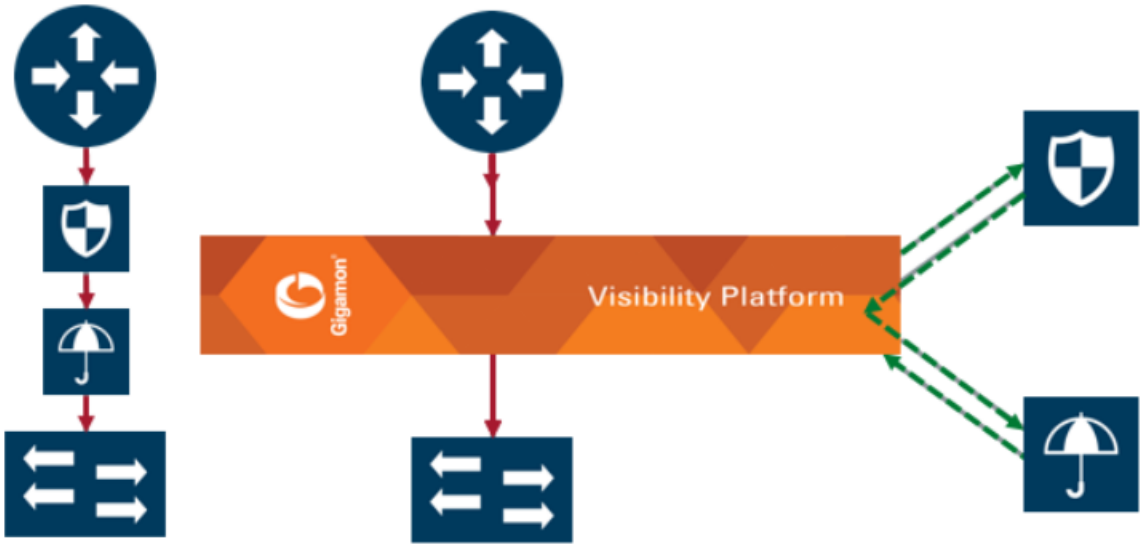


Figure 56 *Serial Inline Tool Arrangement*

Inline Bypass Restriction

The inline bypass arrangements supported by inline TLS/SSL decryption have the following restriction:

- inline network group does not support ingress VLAN tagging on the member links

Forwarding

Inline TLS/SSL decryption supports the following kinds of forwarding:

- inline forwarding—Packets can be forwarded from the inline network or inline network group to the inline tool, inline tool group, or inline series. The IPv6 traffic received from the inline network is forwarded as IPv6 traffic in tool-port. The translation from IPv6 packet in network-port to IPv4 packet in tool port will not be supported.
- inline out-of-band forwarding—Packets from inline ports can be sent to regular tool ports.
- inline bypassing—Packets can be put in loopback between two ports of an inline network.
- TLS/SSL forwarding—Packets from an inline network or inline network group can be sent to GigaSMART, then from GigaSMART to an inline tool, inline tool group, or inline series.
- GigaSMART out-of-band forwarding—Packets from GigaSMART can be copied to tool ports.

Failover

The inline bypass module detects failure either through link loss or tool heartbeat failure.

The inline bypass module supports the following failover actions:

- inline tool failover action—Specifies the failover action taken in response to a failure of an inline tool.
- inline tool group failover actions—Specifies the failover action taken in response to a failure of an inline tool group, when the number of healthy inline tools in the inline tool group (including the spare inline tool, if configured) falls below the configured minimum.
- inline tool series failover actions—Specifies the failover action taken in response to a failure of an inline tool series as a whole. An inline tool series is declared to be in a failure condition as soon as any of its member inline tools goes into a failure condition. An inline tool series recovers from a failure condition after all the member inline tools recover from their failure conditions. The failover-action attributes of the individual inline tools participating in an inline tool series are ignored. Instead, the failover-action configured for the inline tool series is respected.

The inline bypass failover actions are configurable. Refer to the *GigaVUE-OS CLI Reference Guide*for the actions and default values of the inline-tool, inline-tool-group, and inline-serial commands.

The virtual port also has configurable failover actions. Refer to the *GigaVUE-OS CLI Reference Guide* for the actions and default value of the vport command.

The GigaSMART module might also have events such as port down or card down. These failover actions are not configurable but are triggered by events. These events should not impact traffic.

Out-of-Band and Inline Tools

After decryption, traffic can be sent to multiple tools. The tools can be either inline or out-of-band.

Out-of-band tools process the decrypted packets offline. The tools are connected to the GigaVUE-FM,GigaVUE HC Series node through tool or hybrid ports, GigaStream, or port groups with tool or hybrid ports. The out-of-band tools receive a copy of the decrypted packets from the GigaSMART module. This is referred to as GigaSMART out-of-band forwarding.

Inline tools process the decrypted packets inline. Inline tools are connected to the GigaVUE-FM,GigaVUE HC Series node through inline bypass (BPS) modules.

An out-of-band tool might be an Intrusion Detection System (IDS) examining decrypted packets:

- If it detects a threat, the IDS will send a notification back, but does not have the ability to act.

An inline tool might be an Intrusion Prevention System (IPS) examining decrypted packets:

- If it does not detect a threat in the decrypted packets, the traffic comes out of the inline tool and goes back to the GigaSMART module to be re-encrypted and sent to the server.
- If it detects a threat, the IPS can act. The action depends on the tools' behavior. It can either terminate the connection or modify packets
- If the IPS terminates the connection, then GigaVUE-FM,GigaVUE HC Series node will terminate the connection between the client and the server.
- If the IPS modifies packets, then the modified packets will come out of the inline tool, go to the GigaSMART module to be re-encrypted, and sent to the server.

When an inline tool is shared, you must:

- configure the inline second level out-of-band map to forward proxy traffic from GigaSMART to the out-of-band tool port.

When an inline tool is not shared, you must configure only the inline second level out-of-band map to forward proxy and non-proxy traffic from GigaSMART to the out-of-band tool port.

NOTE: When OOB-copy is configured in a non-shared inline-tool with the tag set as “inline”, both proxy and non-proxy out-of-band traffic will be copied with the proxy map VLAN tag.

Service Chaining of Decrypted Traffic

Service chaining of decrypted traffic may be required for compliance purposes.

This is done by directing the decrypted traffic to a hybrid port and applying the required GigaSMART operations on the traffic that is looped back from the hybrid port, before forwarding the traffic to the out-of-band tool.

The GigaSMART operations must be configured on a different GigaSMART engine than the one used for inline TLS/SSL decryption.

Packet Flows

Normally, a client and server talk directly to each other, such as when you are using a browser to go to a bank website, a health care provider, or a search engine.

As shown in figure below , the GigaVUE-FM,GigaVUE HC Series node is placed in the middle between the client and the server. All traffic from the Internet goes through the GigaVUE-FM,GigaVUE HC Series node.

Incoming traffic arrives on an ingress inline network port on the inline bypass module.

SSL traffic, which is TCP traffic, is directed to the GigaSMART module. Until the traffic is processed by the GigaSMART module, it is not known if it is SSL traffic or not.

The GigaSMART module decides what traffic is bypassed, what traffic is sent to tools without decryption, and what traffic is decrypted and then sent to tools. So, there are three types of decisions as follows:

- to bypass or not
- to decrypt or not
- to send to tools or not

Figure 57Packet Flow for Inline SSL Decryption shows the flow for a configuration consisting of a single inline network, a single inline tool, and a single out-of-band tool.

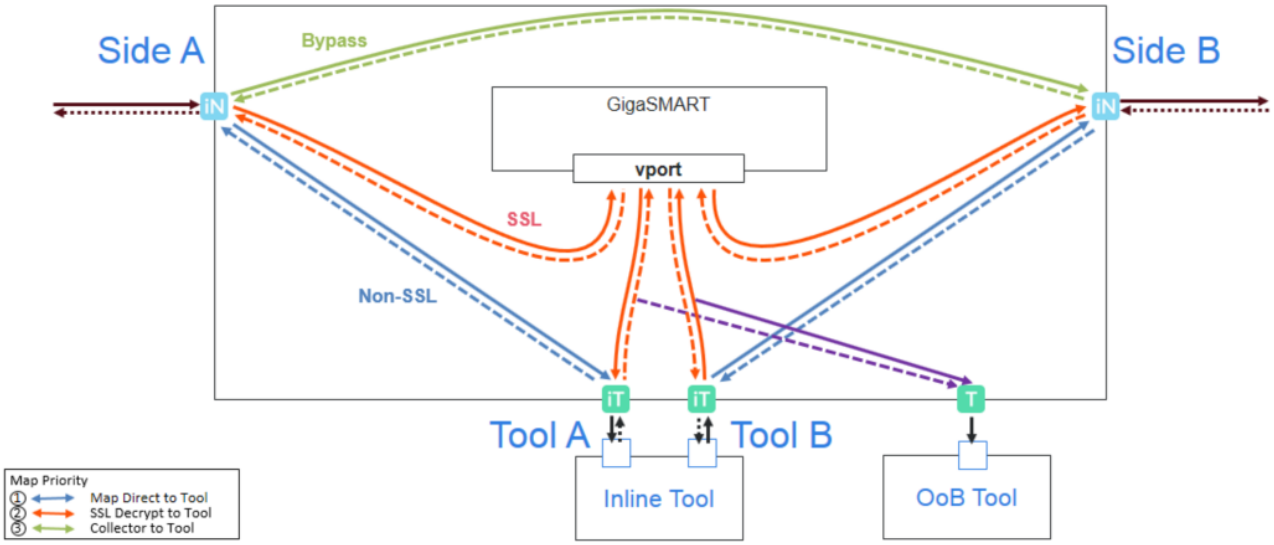


Figure 57 *Packet Flow for Inline SSL Decryption*

Traffic enters the inline SSL decryption solution at the side A inline network port on the inline bypass module.

Some traffic can be bypassed. That traffic goes from the side A inline network port to the side B inline network port on the inline bypass module as shown in the solid green line in Figure 57Packet Flow for Inline SSL Decryption. Bidirectional traffic is shown by dotted lines.

Traffic that is not encrypted (non-SSL) can be sent to tools for inspection. That traffic goes from the side A inline network port to the inline tool A port, to the inline tool, and from the inline tool to the inline tool B port and the side B inline network port on the inline bypass module without going through the GigaSMART module as shown in the solid blue lines in Figure 57Packet Flow for Inline SSL Decryption.

Traffic that is encrypted (SSL) goes from the side A inline network port on the inline bypass module to the GigaSMART module. That traffic goes to a virtual port (vport) that directs traffic to an inline SSL GigaSMART operation as shown by the red line on the left in Figure 57Packet Flow for Inline SSL Decryption. The traffic is decrypted in the GigaSMART module based on policy configuration, sent to the inline tool A port to the inline tool and from the inline tool to the inline tool B port, then back to the GigaSMART module for re-encryption as shown in the solid red lines in the center of Figure 57Packet Flow for Inline SSL Decryption. Finally, the re-encrypted traffic is sent from the GigaSMART module to the side B inline network port on the inline bypass module as shown by the red line on the right in Figure 57Packet Flow for Inline SSL Decryption.

The out-of-band tool can also receive the traffic as shown by the purple line on the right in Figure 57Packet Flow for Inline SSL Decryption.

An out-of-band map from a virtual port to a single tool port and to multiple tool ports is supported. The ports can be tool, hybrid, or GigaStream. Out-of-band maps from a vport to port groups are also supported when the ports in the group are tool or hybrid.

Modules Matrix

For traffic that is encrypted, there is a question as to whether or not it needs to be decrypted. For example, is there a policy for or against decrypting that traffic? There might be a policy, such as do not decrypt financial or health care traffic, or there might be a decrypt list that states that traffic from a particular site should always be decrypted, or there might be a no-decrypt list that states that traffic from a particular site should always be bypassed. So encrypted packets need to be filtered because some packets will not be decrypted, while others will be decrypted.

The following table is a matrix of the GigaVUE-FM,GigaVUE HC Series module used for different types of traffic.

Table 7: GigaVUE-FM,GigaVUE HC Series Modules Used for Type of Traffic

Type of Traffic	GigaSMART	Inline Bypass	GigaSMART
Bypassing GigaSMART	No	Yes - to network	No
Not encrypted (non-SSL)	Yes	Yes - to tools then to network	Yes, depending on the configuration
Encrypted (SSL), and to be decrypted	Yes	Yes - to tools	Yes - to be re-encrypted, then to network

The GigaSMART module does the decryption as well as handling policies, no-decrypt lists, and decrypt lists. The decision to decrypt or not is made in the GigaSMART module.

Inline Single VLAN Tag

During flexible inline bypass operations, network traffic sent to inline tools carries an additional VLAN tag. This tag helps distinguish return traffic from inline tools and ensures correct routing between inline networks. However, these extra tags can cause compatibility issues with certain tools. The Inline Single Tag feature resolves this by replacing the additional VLAN tag on incoming traffic. It enables mapping between network-side VLANs and tool-side VLANs, ensuring traffic reaches the inline tools with only one defined VLAN tag.

In flexible inline arrangements, VLAN tags are linked to each inline map rather than to inline network ports, as seen in classic inline bypass. A single inline network port can therefore host multiple inline maps, each with a separate VLAN tag. VLAN tags are added to traffic before it is sent to tools and removed before it returns to the network.

NOTE: The **OOB Copy** tag attribute **none** is invalid for single-tag maps; use the **original** or **as-inline** attribute instead. When configuring inline maps with a single VLAN tag, ensure that the map rules use the same VLAN tag as specified in the from parameter.

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.

Types of Inline Networks

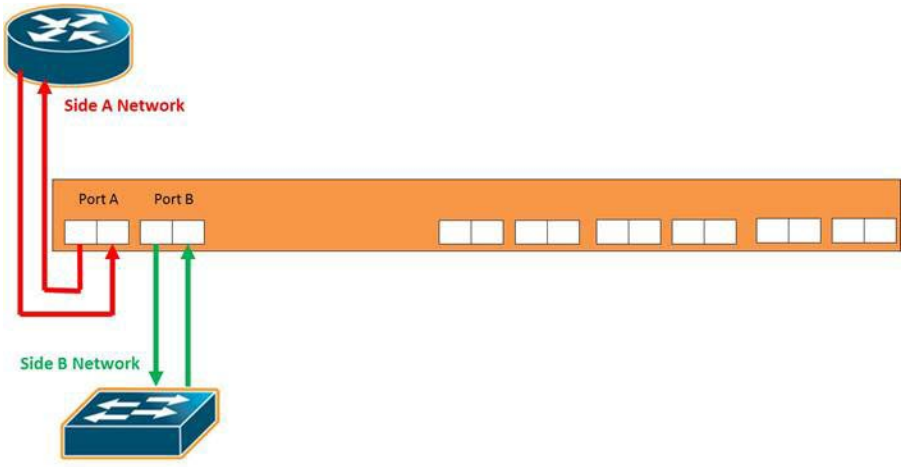
The types of inline networks are described in the following sections:

- 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-FM,GigaVUE HC Series 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.



Resilient Inline Arrangement

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

NOTE: Resilient inline arrangement is not supported in GigaVUE-HCT devices, GigaVUE-TA400, and GigaVUE-TA400E.

The following figure illustrates the resilient inline arrangement.

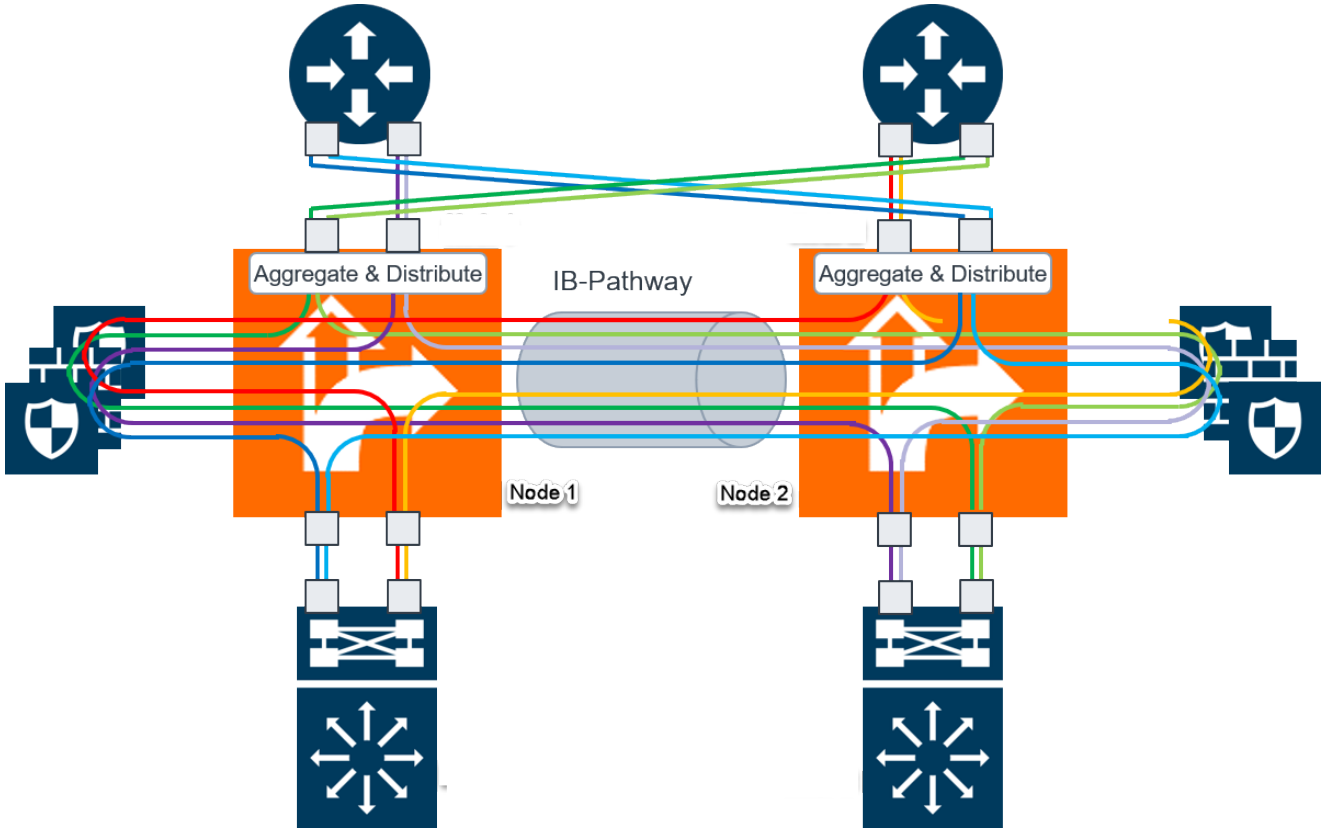


Figure 58 *Resilient Inline Arrangement*

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

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

Resilient inline arrangement is based on an aggregation and distribution principle that divides the packets received by an inline packet broker, between Node 1 and Node 2. The inline packet broker on the left, guides the Node 1 class of packets through its local tools and Node 2 class of packets through the remote tools that are reachable by a resilient inter-broker pathway. Similarly, the inline packet broker on the right, guides the Node 2 class of packets through its local tools and Node 1 class of packets through the remote tools.

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

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

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

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

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

Resilient Inline Arrangement— Classic

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

Resilient Inline Arrangement With Single VLAN Tag

You can choose to deploy a resilient inline arrangement with a single VLAN tag in which a packet received from an inline network is guided to the inline tool using a single VLAN tag, which you can configure when creating a flexible inline map. You must configure the

packet's original VLAN tag as the network side VLAN tag and provide the required tool side VLAN tag when you create the flexible inline maps. The single VLAN tag is useful when your inline tools do not support Q-in-Q VLAN tags.

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

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

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

NOTE: The **OOB Copy** tag attribute **none** is invalid for single-tag maps; use the **as-inline** attribute instead.

Inter-broker Pathway (IB-P)

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

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

- inter-broker pathway-down—the traffic is handled based on the failover action selected for the inline map configured, as follows:
  - If the failover is set to 'bypass', the traffic is passed directly between the respective inline network ports.
  - If the failover is set to 'original-map', the traffic is passed through the path that is defined by the respective original map.

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

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

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

Resilient Inline Arrangement—Rules and Notes

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

- In GEN2 GigaSMART card, a maximum of 14 VLANs will be supported for a single inline- network per GS Group. In the case of multiple inline-network ports (number of inline- network ports x number of VLANs), the number should not exceed 14 per GS Group.
- In GEN3 GigaSMART card, a maximum of 16 VLANs will be supported for a single inline- network per GS Group. In the case of multiple inline-network ports (number of inline- network ports x number of VLANs), the number should not exceed 16 per GS Group.

To configure, refer to Configure Resilient Inline Arrangement Solution.

Hardware Security Modules (HSM)

NOTE: Starting in release 6.14.00, Entrust/nCipher HSM integration is deprecated. Gen2 DSSL HSMs (nCipher v1.0) and Gen3 HSMs using nCipher 12.6 (`nshield_12_60`) are no longer supported, and upgrades are blocked if any nCipher HSM configuration is present.

A Hardware Security Module (HSM) is a dedicated, tamper-resistant device that protects cryptographic keys and performs secure key operations. In inline deployments, HSMs are used to safeguard the private keys required for TLS/SSL decryption. By offloading key storage and sensitive cryptographic tasks to an HSM, organizations ensure that decryption happens without exposing keys inside the node.

Why HSM

Most enterprise traffic today is encrypted, making TLS/SSL decryption a critical part of inline deployments. Security tools such as intrusion prevention systems (IPS) or data loss prevention (DLP) engines need decrypted traffic to analyze threats. Without an HSM, private keys must be stored directly on a Gigamon HC Series node or server, which creates risk.

With an HSM, the private keys remain inside a hardened device. The node communicates with the HSM using secure tokens, never exposing the raw keys. This means:

- Keys never leave the HSM, even if the node is compromised.
- The node can request private key operations such as decryption without handling the key itself.
- Only approved applications, such as GigaSMART® SSL/TLS decryption, can use the tokenized access.

This makes the HSM a critical part of inline solutions where compliance, regulations, or strict security rules require hardware-grade key protection.

How HSM Solution Works

In Flexible Inline Arrangements, HSMs are supported through **HSM Groups**. An HSM Group represents one or more hardware modules connected to a node. Supported HSM devices include:

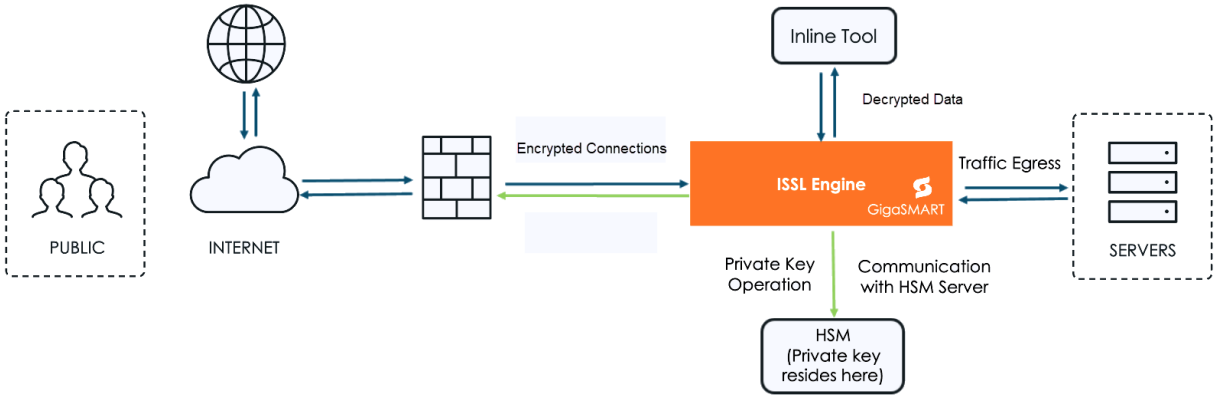
- Thales Luna Network HSMs

Each group is tied to a single vendor type; mixing vendors in the same group is not allowed. Thales Luna clustering is not supported, and IPv6 connectivity is not available. Nodes must be configured with static IPv4 addresses to communicate with HSMs.

When an HSM is in place, the node offloads TLS/SSL decryption key handling. For example, in an Inline SSL Decryption workflow, traffic passes through GigaSMART, which queries the HSM for key use. The HSM processes the request securely and returns only a tokenized response. The private key never leaves the HSM.

The following images illustrates the process flow of HSM solution

.



GigaSMART applications that commonly leverage HSMs include:

- SSL/TLS decryption for inline traffic inspection.
- Passive SSL/TLS Decryption

NOTE: Classic Inline Bypass does not support HSM integration. Classic deployments provide resiliency through protected inline networks, tool groups, and failover actions, but they do not include any mechanism for hardware-based key storage or integration with external HSM appliances.

HSM - Supported Platforms

HSM Group is supported in the following platforms:

- GigaVUE-HC1 Gen3
- GigaVUE-HC3 Gen3
- GigaVUE-HC1-Plus

To learn more about HSM configuration with Inline TLS/SSL Decryption, refer to:

- Configure Entrust nShield HSM for TLS/SSL Decryption

Internet Content Adaptation Protocol (ICAP)

The Internet Content Adaptation Protocol (ICAP) is used to connect inline deployments with external security services, such as Data Loss Prevention (DLP) systems. It provides a standard way to send decrypted HTTP traffic to a DLP-ICAP server for inspection and

policy enforcement.

In Gigamon deployments, the ICAP Client runs inside the GigaSMART® engine, where it acts like an inline tool. It receives decrypted HTTP traffic and forwards it to the ICAP server. The server inspects the content, applies any security policies, and returns the results. The ICAP Client then reinserts the inspected traffic back into the inline path.

NOTE: The ICAP Client is included as part of the TLS/SSL license SKUs.

How ICAP Solution works

In an Inline SSL/TLS Decryption (iSSL) workflow, decrypted HTTP requests and responses can be sent directly from GigaSMART to the ICAP Client. The ICAP Client then:

- Sends the decrypted traffic to the ICAP server for analysis.
- Accepts and applies any changes from the ICAP server before passing the traffic downstream.
- Works in both inbound and outbound inline traffic chains.

This process allows DLP systems to examine live, decrypted traffic in real time without breaking the inline flow.

Originally, DLP-ICAP servers had to be placed outside of the Gigamon decryption zone, which limited their ability to enforce policies on live traffic. With the ICAP Client, integration now happens directly in the inline path, providing real-time DLP enforcement.

The ICAP Client can also operate independent of Inline SSL , as long as it receives HTTP traffic (clear text) from another source.

The following figures illustrate typical **Inbound** and **Outbound** deployment scenarios where the ICAP Client is used to support DLP inspection.

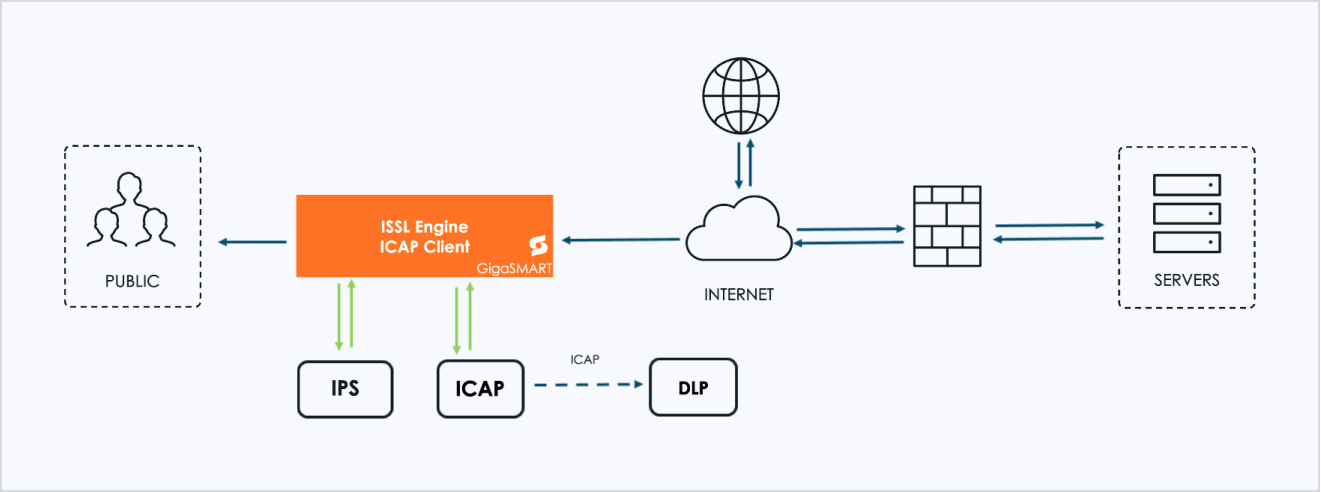


Figure 59 *Inline SSL Solution with ICAP Client support – Outbound Deployment*

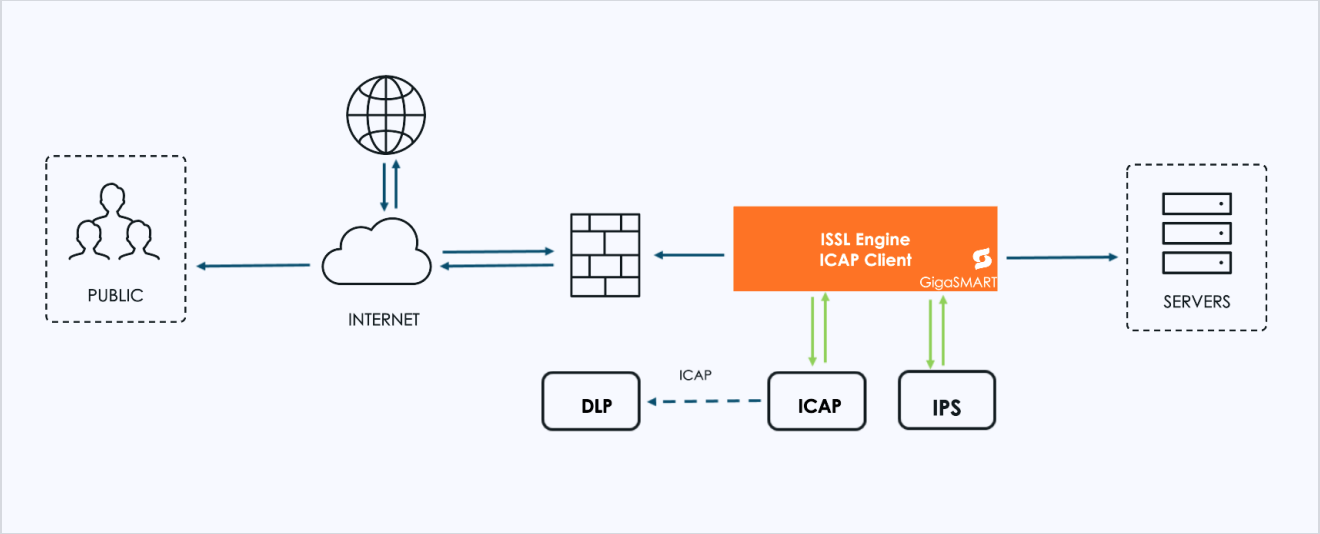


Figure 60 *Inline SSL Solution with ICAP Client support – Inbound Deployment*

ICAP - Supported Platforms

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

Platform	Card
GigaVUE-HC1 Gen3	SMT-HC1-S
GigaVUE-HC1P Gen3	SMT-HC1-S , SMT-HC1A-R
GigaVUE-HC3 Gen3	SMT-HC3-C08

ICAP - Rules, Notes, and Limitations

1. **GSOP Integration Rules**
   - The ICAP GSOP must run as a standalone operation. It cannot be chained with other GSOPs.
   - Each GSGroup can include only one ICAP GSOP.
   - ICAP is not supported in clustered environments.
2. **Protocol and Traffic Handling**
   - HTTP/2 traffic is not supported. HTTP/2 shall be downgraded to HTTP1.1 if required via a configuration option, so that ICAP inspection shall be done.
   - HTTP/1.1 pipe lining is not supported; only one HTTP request/response can be processed at a time.
   - HTTP chunk extensions are not forwarded to the ICAP server.

- HTTP trailers are forwarded as-is, but trailers within the preview length are not sent, per RFC 3507 limitations.
- Any Trailer present in HTTP messages are forwarded to the ICAP server without any modification, as the ICAP RFC does not define specific behavior for HTTP trailers.

3. **Engine Port and Traffic Flow Considerations**
   - Using multiple engine ports in an inline pair map can disrupt traffic due to asymmetric flow. For example, a SYN packet from side A may hit engine port 1 while the SYN/ACK reply from side B reaches engine port 2, causing flow inconsistency.

ICAP - Supported GigaSMART Engine Ports

Refer to the table below for information on the number of GigaSMART engine ports required for each deployment module:

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

For more information on configuring an ICAP Client, refer to Configure ICAP Client.

Gigamon Resiliency for Inline Protection

Gigamon Resiliency for Inline Protection (GRIP) is the umbrella term for node-level resiliency in inline deployments. It ensures that live traffic continues to flow and inline tools remain protected, even if one node fails.

GRIP is implemented differently depending on whether you are using **Classic Inline Bypass** or **Flexible Inline Arrangements**, but the goal is the same: maintain service continuity and protect inline inspection during node outages.

Both models rely on bypass protection switch relays on BPS modules to provide physical fail-safe behavior on protected inline networks

GRIP in Classic Inline Bypass

In Classic Inline Bypass, GRIP is a hardware-based failover mechanism that uses stack signaling and relay-protected inline links. Two nodes are paired using a stack signaling link:

- The primary node actively handles inline traffic.
- The standby node is ready to take over if the primary fails.
- Both nodes connect to the same inline network links through bypass combo modules (fiber) or copper TAP modules.

These modules include optical or electrical relays that automatically close when the primary fails, rerouting traffic to the standby node. Failover typically occurs within 0–10 seconds.

If both nodes fail, GRIP ensures that traffic still flows directly between the inline network ports, bypassing the tools so the network remains available, even if inspection is lost.

GRIP in Flexible Inline (Redundancy)

In Flexible Inline Arrangements, node-level resiliency is delivered through Redundancy, which functions as the Flex[[[Undefined variable guideVar.Guide-VS-ISSL]]]ible Inline version of GRIP. As in Classic, protected inline networks still use BPS relays for physical protection, but Flexible Inline adds policy-driven redundancy and traffic steering on top of that hardware foundation

- Both nodes can be configured in active/standby modes.
- If one node or tool path fails, traffic flows are automatically redistributed to the healthy node or tools, guided by session-aware hashing and policy rules.
- Failover is faster and more flexible.

Across both deployment models, GRIP provides:

- **Node-Level Protection** – Ensures traffic continuity even if an entire node fails.
- **High Availability** – Prevents tool disconnections and minimizes downtime.
- **Operational Flexibility** – Hardware relay-based protection in Classic; policy-driven redundancy in Flexible.
- **Seamless Recovery** – Once the failed node recovers, it can resume its role without disrupting live traffic.

GRIP Scenarios

The following scenarios illustrate how GRIP behaves in real-world deployments, showing how traffic flows through primary and secondary nodes under different conditions.

Traffic Flows Through Node with Primary Role

When both nodes are active, the primary node handles all inline traffic. Its bypass protection switch relays remain open, directing traffic to the inline tool attached to the primary node. The secondary node monitors the state of the signaling link. As long as the link is up, the secondary node remains idle with its relays closed.

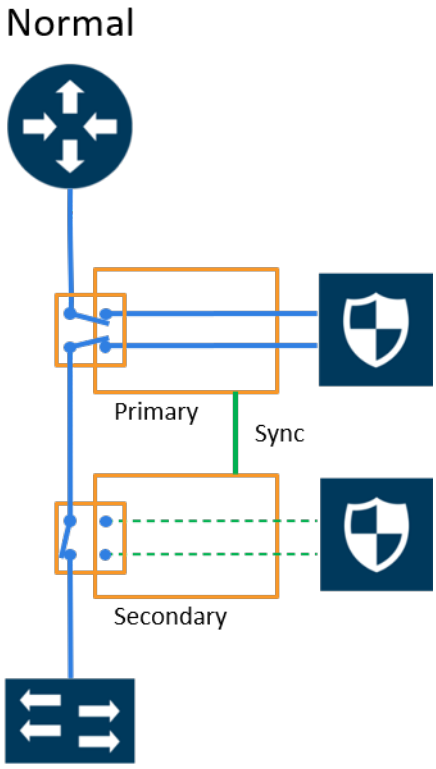


Figure 61 *Traffic Flows Through Node with Primary Role*

Traffic Flows Through Node with Secondary Role after Primary is Lost

If the primary node loses power or becomes unavailable, its relays close automatically. The signaling link informs the secondary node, which then opens its relays. Traffic immediately shifts to flow through the secondary node, keeping the inline tool path active.

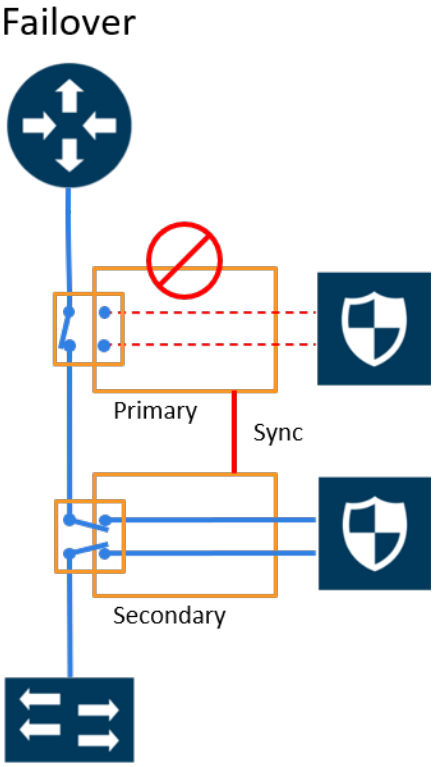


Figure 62 *Traffic Flows Through Node with Secondary Role after Primary is Lost*

Both Nodes Go Down and Only Secondary Comes Up

If both nodes fail at the same time, and only the secondary is later powered back up, traffic is not passed through the inline tools. Instead, it is bypassed across the relays. For this reason, it is recommended to always restore the primary node as well, or promote the secondary into the primary role if the original primary is unreliable.

Both Nodes Fail; No Traffic Monitoring

If both nodes remain down, the bypass relays on both close, allowing network traffic to continue to flow, but all inline tools are bypassed. The network remains up, but without inspection or monitoring.

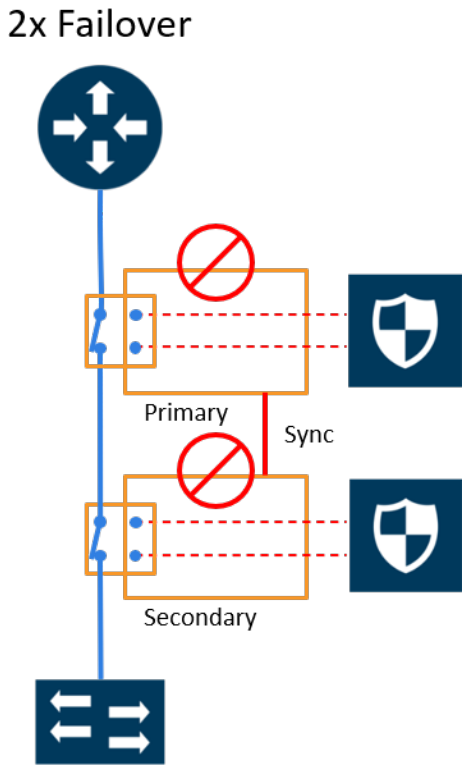


Figure 63 *Both Nodes Fail; No Traffic Monitoring*

Both Nodes are in Suspended State

In the GRIP solution, if both primary and secondary nodes are switched to suspended , Network traffic will be bypassed instead of being sent to inline tools in both the nodes. Switching redundancy profile protection role on secondary node alone from suspended to secondary will still cause network traffic to be bypassed instead of being sent to inline tools in both the nodes .

It is recommended to switch the redundancy profile protection role on primary node from suspended to primary first , then followed by switching the redundancy profile protection role on secondary node from suspended to secondary.

How to Handle Recovery

When the primary node recovers, it restores its inline traffic paths and opens its relays. The signaling link then instructs the secondary node to close its relays, returning traffic to the primary path. Recovery is automatic, and traffic resumes flowing through the primary node’s tool path.

How to Cable GigaVUE-FM,GigaVUE HC Series Nodes

To cable two GigaVUE-FM,GigaVUE HC Series nodes, as shown in Figure 61Traffic Flows Through Node with Primary Role with the primary on the left and the secondary on the right:

- Connect the network shown at the top of Figure 61Traffic Flows Through Node with Primary Role to inline network port A on the primary node.
- Connect inline network port B on the primary node to inline network port A on the secondary node.
- Connect inline network port B on the secondary node to the network shown at the bottom of the Figure 61Traffic Flows Through Node with Primary Role.
- Connect the signaling port on the primary node to the signaling port on the secondary node.

Redundancy Profile

GRIP resiliency is managed by a redundancy profile, which defines how the two nodes coordinate:

- Signaling Port – the port pair used to exchange health status between the nodes.
- Protection Role – assigns each node as Primary, Secondary, or Suspended.
  - Primary: Actively handles inline traffic.
  - Secondary: Remains on standby, takes over when the primary fails.
  - Suspended: Used for maintenance or to manually force failover.

NOTE: Cluster Limitation- GRIP is supported in clustered environments, but the Suspended role has restrictions on standby nodes. In this case, it is recommended to switch the standby into a new role or carefully apply the suspended state.

Refer to How to Use Suspended Role for Maintenance to know moreGigamon Resiliency for Inline Protection

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. Post-reload or Power Cycle, the signaling port link stays down when the redundancy profile is attached to the inline network and no maps have been configured. A map should be configured to bring the signaling port up. If a map exists, the signaling port will appear without any issues.
5. Link Failure Propagation is not recommended 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.
6. Gigamon Resiliency for Inline Protection (GRIP™) is not supported in GigaVUE-HCT devices.
7. GRIP is not supported in other GigaVUE TA Series devices due to the absence of BPS modules.
8. Refer to Flexible Inline TLS/SSL Decryption Solution—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.

Inline TLS/SSL Decryption

NOTE: In this section, Secure Sockets Layer (SSL) and Transport Layer Security (TLS) mean the same thing. The terms are used interchangeably.
Before you proceed, make sure you understand the TLS/SSL Terminology and Acronyms

Inline TLS/SSL decryption gives your security tools access to encrypted traffic. It works by decrypting packets and sending them to tools that are placed either inline (directly in the data path) or out-of-band (off the path). These tools then scan the decrypted traffic for threats, such as viruses and malware.

Unlike passive decryption, which only sends decrypted traffic to out-of-band tools that can alert but not act, inline decryption enables tools to take immediate action on threats.

Why Decrypt TLS/SSL Traffic

Most Internet traffic is now encrypted using SSL or TLS. While encryption protects data, it also makes it harder to inspect packets for threats. As a result, malware and other attacks often conceal themselves within encrypted traffic. Without decryption, these threats go unseen.

By decrypting TLS/SSL traffic:

- You can detect hidden threats across any port or application (e.g., HTTPS, email, VoIP).
- You reduce risk and gain visibility into encrypted sessions within your network.

Inline vs. Passive Decryption

Inline TLS/SSL decryption is active. It enables tools to inspect traffic in real-time and take action when a threat is detected. This is different from passive decryption, such as the existing GigaSMART SSL/TLS solution. Passive decryption only sends traffic to tools out-of-band. These tools can detect threats and alert users, but cannot stop threats directly.

Inline decryption does more. It offloads the complex decryption task so tools can focus on detecting and stopping threats faster and more effectively.

How Inline TLS/SSL Decryption Works

Inline TLS/SSL decryption performs the following key functions:

- Detects encrypted traffic across any port in your network.
- Intercepts encrypted flows between clients and servers.
- Filters traffic by policy, allowing sensitive flows (e.g., healthcare or financial data) to bypass decryption.
- Decrypts packets at a single, centralized point.
- Forwards decrypted data to one or more tools for inspection. These tools can be inline or out-of-band.
- Takes action on threats:
  - Tools can modify traffic (e.g., remove malware) or terminate sessions.
  - If modified, GigaSMART re-encrypts the packets.
  - If the session is terminated, GigaSMART ends the connection between client and server.
- Re-encrypts traffic after inspection and sends it back into the network.

Mutual TLS (mTLS) with Inline TLS/SSL Decryption

Mutual TLS (mTLS) is an extension of TLS/SSL in which both the client and the server present and validate certificates during the TLS handshake. This creates a two-way trust relationship, the client verifies that it is talking to the expected server, and the server verifies that it is talking to an authenticated client.

In a standard Inline TLS/SSL deployment, Gigamon validates the server certificate and optionally enforces certificate/revocation policies before re-signing traffic and forwarding it to inline tools for inspection. With mTLS support, Gigamon can additionally:

- Validate client certificates presented to the server.
- Enforce client-authentication policies (for example, drop connections with unknown or self-signed client certificates).
- Offload the client-auth handshake to the inline TLS/SSL solution so attached tools can focus on inspection rather than certificate processing.

In an inbound deployment with mTLS enabled, the high-level flow is as follows:

1. The client initiates a TLS connection to the protected server.
2. The GigaVUE node , acting as a Man-in-the-Middle (MitM), terminates the client-side TLS session and validates the server certificate as in a regular inbound deployment.
3. When the server requests a client certificate, the client sends its certificate to the GigaVUE node.
4. The client certificate chain is validated against a client trust store and applies client-auth policy (expired, self-signed, unknown CA, revocation status, and so on).
5. If the client is accepted, Gigamon completes the mTLS handshake on both sides, decrypts the traffic, and forwards clear-text packets to inline tools for inspection.
6. After inspection, Gigamon re-encrypts traffic and forwards it to the server, maintaining the mTLS session semantics on the wire.

NOTE: Mutual TLS (mTLS) functionality is now supported only for **Inbound and Hybrid Inline TLS/SSL Decryption** in Flexible Inline deployments. Outbound deployments continue to use one-way TLS/SSL decryption without client authentication.

Mutual TLS Inline TLS/SSL Decryption is on GigaVUE-HC Series Gen3 platforms used for Inline TLS/SSL decryption such as : GigaVUE-HC1 Gen3, GigaVUE-HC3 Gen3, GigaVUE-HC1-Plus.

The following feature combinations are not supported with Mutual TLS Inline TLS/SSL Decryption:

- Inline TLS/SSL L3 Tool NAT/PAT Support
- Entrust nShield and Thales Luna HSM for iSSL
- Post-quantum client authentication using PQC certificates (PQC cipher support is limited to server side).
- One-Arm mode and Tool Early Engage.
- Tool Early Inspect.
- Inline TLS/SSL deployments with ICAP .

Privacy and Sensitive Data Handling

Decrypted traffic may expose sensitive information, such as:

- Usernames and passwords in email
- Social security numbers in financial records

To protect user privacy and meet compliance standards, you can define policies that exclude certain traffic from decryption. This helps align with acceptable use, legal, and regulatory requirements.

What Applications Use TLS/SSL

Many applications such as email, websites, and voice calls over IP (VoIP), use TLS/SSL to secure data. Encryption ensures that sensitive data stays private while traveling over the Internet. But when data is encrypted, network tools cannot inspect it. This creates blind spots where threats can hide.

Decrypting this traffic removes those blind spots. It allows your tools to inspect data on any port or application, such as HTTPS (port 443), email, web, VoIP, FTPS, SMTP, IMAP, and POP3 (via StartTLS).

Supportability and Compatibility for Inline TLS/SSL Decryption

Refer to the following sections for details:

- Supported Platforms
- GigaSMART Licensing
-
- Port Requirements
- GigaSMART Compatibility
-

Supported Platforms

Inline TLS/SSL decryption is supported on the following platforms:

- GigaVUE-HC1
- GigaVUE-HC1-Plus
- GigaVUE-HC3

To enable decryption, both the GigaSMART module and the inline bypass module must be installed on the same node.

Mutual TLS (mTLS) Support and Limitations

Mutual TLS (mTLS) functionality is now supported only for **Inbound and Hybrid Inline TLS/SSL Decryption** in Flexible Inline deployments. Outbound deployments continue to use one-way TLS/SSL decryption without client authentication. The platforms and features that are compatible with this functionality are listed below:

Supported platforms for mTLS

Mutual TLS (mTLS) is supported only on GigaVUE-HC Series Gen 3 platforms such as

- GigaVUE-HC1 Gen3
- GigaVUE-HC3 Gen3
- GigaVUE-HC1-Plus

Limitations of mTLS functionality

- L3 Tool NAT/PAT: the Inline TLS/SSL L3 Tool NAT/PAT Support feature cannot be combined with mTLS. If NAT/PAT mode is enabled on the profile, mTLS must remain disabled.
- HSM integration: Inline TLS/SSL decryption using Entrust nShield or Thales Luna HSM cannot be used together with mTLS in the same Inline SSL application.
- PQC client authentication: Post-Quantum Cryptography (PQC) cipher support does not extend to client certificate authentication. If a server requests a PQC client certificate, the session is bypassed for client-authentication purposes.
- One-Arm mode and Tool Early Engage: One-Arm mode requires Tool Early Engage and is not supported with mTLS. Any configuration that enables One-Arm mode or Tool Early Engage must keep mTLS disabled.
- Tool Early Inspect: Tool Early Inspect changes the order of TLS handshakes and is incompatible with mTLS; these features cannot be enabled at the same time.
- ICAP client flows: Inline TLS/SSL with ICAP Client continues to operate on decrypted HTTP traffic; mTLS cannot be enforced end-to-end across the ICAP chain in this release.
- When inbound mTLS traffic is running with a large number of Server key maps or scaled configurations/rules, editing these configurations (for example, adding or removing server key map entries or updating the client trust store) can temporarily disrupt inline TLS/SSL decryption for existing connections while the new configuration is applied.

GigaSMART Licensing

Required License: Subscription based TLS/SSL Decryption license.

Inline Bypass Requirements

For physical inline bypass, install a fiber bypass (BPS) combo module. On GigaVUE-HC1, a copper TAP can also be used for physical bypass. Refer to the table for a list of supported inline bypass modules.

Model & Module Type	Description & BPS Port Pairs
GigaVUE-HC1 PLUS	Includes:
· BPS-HC1-D25A60 (HC1-Plus)	6 × SX/SR multimode inline network port pairs
· BPS-HC1-D35C60 (HC1-Plus)	6 × LX/LR single-mode inline network port pairs
GigaVUE-HC1 (Classic HC1 Chassis)	Includes:
· BPS-HC1-D25A24	2 × SX/SR multimode (50/125 µm) inline network port pairs + 4 SFP+ cages
· BPS-HC1-D25A60 (HC1-Plus)	6 × SX/SR multimode inline network port pairs
· BPS-HC1-D35C60 (HC1-Plus)	6 × LX/LR single-mode inline network port pairs
GigaVUE-HC2	Includes:
· BPS-HC0-D25A4G	4 × SX/SR (50/125 µm) multimode bypass pairs, 16 SFP/SFP+ cages
· BPS-HC0-D25B4G	4 × SX/SR (62.5/125 µm) multimode bypass pairs, 16 SFP/SFP+ cages
· BPS-HC0-D35C4G	4 × LX/LR single-mode bypass pairs, 16 SFP/SFP+ cages
· BPS-HC0-Q25A28	2 × SR4 (50/125 µm) bypass pairs, 8 SFP/SFP+ cages (40 Gb capability)
GigaVUE-HC3	Includes:
· BPS-HC3-C25F2G	2 × SR4 (40/100 Gb) BPS pairs, 16 SFP+ cages
· BPS-HC3-Q35C2G	2 × 40 Gb LR bypass pairs, 16 SFP+ cages
· BPS-HC3-C35C2G	2 × 100 Gb LR bypass pairs, 16 SFP+ cages

The following diagram shows a GigaVUE device with both the GigaSMART and inline bypass (BPS) modules installed:

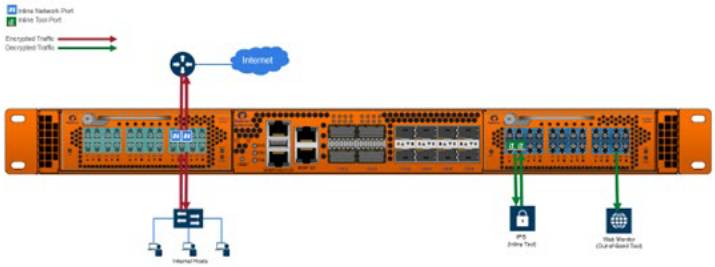


Figure 64 *GigaVUE-FM,GigaVUE HC Series Modules: GigaSMART and Inline Bypass*

- The GigaSMART module contains the SSL decryption software.
- The inline network ports are located on the inline bypass module.
- Inline and out-of-band tool ports are available on the same GigaVUE node.

Port Requirements

- For inline traffic, both inline network and inline tool ports require two links (a port pair) to handle bidirectional traffic.
- For out-of-band (offline) traffic, only one link is needed, as the traffic is not bidirectional.

GigaSMART Compatibility

Inline TLS/SSL decryption must be configured exclusively on a GigaSMART engine. It is not compatible with other GigaSMART operations, including Passive TLS/SSL decryption.

- Do not share the same GigaSMART engine with other operations when using inline TLS/SSL decryption.
- You can deploy both inbound and outbound inline TLS/SSL decryption on a single GigaSMART engine.

NOTE: On GigaVUE-HC1 nodes, Inline TLS/SSL decryption can be configured alongside other GigaSMART applications.

Supported Ciphers

Inline TLS/SSL decryption supports modern cryptographic algorithms. It supports the commonly-supported TLS 1.2 and TLS 1.3 ciphers.

Combining the following ciphers, MACs, and Key Exchange Algorithms results in many cipher suites:

- Ciphers: AES_128_CBC, AES_128_GCM, AES_256_GCM, AES_256_CBC, Camellia, Chacha20
- MAC: SHA, SHA256, SHA384, Poly1305
- Key Exchange Algorithms: RSA, DHE_RSA, ECDHE_RSA, ECDHE_ECDSA.

Diffie Hellman Ephemeral (DHE) is a key exchange protocol.

Inline TLS/SSL Decryption supports key cipher suites and exchanges without downgrading cryptography levels of the organization.

Cipher suites are a standard combination of the following:

- **bulk encryption algorithm**—Specifies how to encrypt communications, including the algorithm, key size, and the cryptographic mode used. For example, AES_128_CBC is AES with 128-bit keys in Cipher Block Chaining mode.
- **key exchange algorithm**—Specifies how both sides authenticate each other during the TLS/SSL handshake. For example, RSA.

- **message authentication code (MAC**)—Specifies the hash algorithm used to verify that communications have not been tampered with. For example, SHA.
- **pseudorandom function**—Specifies how a 384-bit master secret, which is used as a source of randomness for session keys, is generated.

NOTES:

- TLS/SSL transactions with unsupported ciphers will be bypassed/TCP proxied.
- The new TLS1.3 cipher suites are defined differently and do not specify the certificate types (RSA/DSA/ECDSA) or the key exchange mechanism (DHE/ECHDE).
- The Inline TLS/SSL session is now equipped to receive a client hello with the key exchange X25519Kyber768 and now fall back to using just X25519. This ensures the system maintains secure and functional connections, even if it cannot use the newer, quantum-resistant algorithm now.

The following key sizes are supported:

- **RSA**—2048, 3072, 4096, 8192
- **DH**—1024, 2048, 4096
- **ECC**—prime256v1, ecsecp256r1, ecsecp384r1, ecsecp521r1, secp160k1, secp160r1, secp160r2, secp192k1, secp192r1, secp224k1, secp224r1, secp256k1, secp256r1, secp384r1, secp521r1, brainpoolP256r1, brainpoolP384r1, brainpool512r1, X25519, X448

The following TLS extension is supported:

- **RFC7301**—Application-Layer Protocol Negotiation (ALPN)

The below table lists the TLS S1.3 and TLS 1.2 ciphers that support Inline TLS/SSL Decryption.

Inline TLS/SSL Supported TLS 1.3 Ciphers

Cipher Name	Encryption (Enc)	MAC
TLS_AES_256_GCM_SHA384	AES_ 256_GCM	SHA384
TLS_CHACHA20_POLY1305_SHA256	CHACHA20_ POLY1305	SHA256
TLS_AES_128_GCM_SHA256	AES_128_GCM	SHA256

Inline TLS/SSL Supported TLS 1.2 Ciphers

Cipher Name	Key Exchange (Kx)	Authentication (Au)	Encryption (Enc)	MAC
TLS_DHE_RSA_WITH_AES_128_ CBC_SHA	DHE	RSA	AES128_CBC	SHA
TLS_DHE_RSA_WITH_AES_256_	DHE	RSA	AES256_CBC	SHA

Cipher Name	Key Exchange (Kx)	Authentication (Au)	Encryption (Enc)	MAC
CBC_SHA				
TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA	DHE	RSA	CAMELLIA128	SHA
TLS_RSA_WITH_CAMELLIA_128_CBC_SHA	RSA	RSA	CAMELLIA128	SHA
TLS_RSA_WITH_CAMELLIA_256_CBC_SHA	RSA	RSA	CAMELLIA256	SHA
TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA	DHE	RSA	CAMELLIA256	SHA
TLS_DHE_RSA_WITH_AES_128_CBC_SHA256	DHE	RSA	AES128_CBC	SHA256
TLS_DHE_RSA_WITH_AES_256_CBC_SHA256	DHE	RSA	AES256_CBC	SHA256
TLS_DHE_RSA_WITH_AES_128_GCM_SHA256	DHE	RSA	AES128_GCM	SHA256
TLS_DHE_RSA_WITH_AES_256_GCM_SHA384	DHE	RSA	AES256_GCM	SHA384
TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305	ECDHE	RSA	CHACAH20	POLY1305
TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305	ECDHE	ECDSA	CHACAH20	POLY1305
TLS_DHE_RSA_WITH_CHACHA20_POLY1305	DHE	RSA	CHACAH20	POLY1305
TLS_RSA_WITH_AES_128_CBC_SHA	RSA	RSA	AES128_CBC	SHA
TLS_RSA_WITH_AES_256_CBC_SHA	RSA	RSA	AES256_CBC	SHA
TLS_RSA_WITH_AES_128_CBC_SHA256	RSA	RSA	AES128_CBC	SHA256
TLS_RSA_WITH_AES_256_CBC_SHA256	RSA	RSA	AES256_CBC	SHA256
TLS_RSA_WITH_AES_128_GCM_SHA256	RSA	RSA	AES128_GCM	SHA256
TLS_RSA_WITH_AES_256_GCM_SHA384	RSA	RSA	AES256_GCM	SHA384
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA	ECDHE	ECDSA	AES128_CBC	SHA
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA	ECDHE	ECDSA	AES256_CBC	SHA

Cipher Name	Key Exchange (Kx)	Authentication (Au)	Encryption (Enc)	MAC
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	ECDHE	RSA	AES128_CBC	SHA
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	ECDHE	RSA	AES256_CBC	SHA
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	ECDHE	ECDSA	AES128_CBC	SHA256
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384	ECDHE	ECDSA	AES256_CBC	SHA384
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	ECDHE	RSA	AES128_CBC	SHA256
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	ECDHE	RSA	AES256_CBC	SHA384
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	ECDHE	ECDSA	AES128_GCM	SHA256
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	ECDHE	ECDSA	AES256_GCM	SHA384
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	ECDHE	RSA	AES128_GCM	SHA256
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	ECDHE	RSA	AES256_GCM	SHA384

Post-Quantum Cryptography (PQC) Cipher Support

Post-Quantum Cryptography (PQC) introduces cryptographic algorithms designed to be secure against attacks from both classical and quantum computers. This update enables devices to negotiate and process SSL/TLS sessions using PQC algorithms when performing Inline SSL decryption.

Supported Algorithms

- Key Exchange (KEM):
  - ML-KEM (CRYSTALS-KYBER) in variants mlkem512, mlkem768, mlkem1024
  - Hybrid combinations (e.g., X25519_MLKEM768, SecP256r1_MLKEM768, SecP384r1MLKEM1024)
- Signature Algorithms:
  - ML-DSA (CRYSTALS-DILITHIUM) in variants mldsa44, mldsa65, mldsa87

Supported Key File Types

The supported file types for Inline SSL support for Post-Quantum Cryptography (PQC) Ciphers are as follows:

- PEM format: pq-private (for PQC private keys)
- PKCS12 format: pq-pkcs12 (for PQC key/certificate bundles)
- Certificate files: pq-certificate (for PQC certificates)

You can enable PQC ciphers via GigaVUE-OS CLI command ***apps keystore*** or select the Key type ***'PQC'*** in your Inline SSL profile.

Limitations

The main limitations of the Inline TLS/SSL support for Post-Quantum Cryptography (PQC) ciphers include:

- The FHA Inline TLS/SSL dashboard does not currently show any details related to PQC.
- This feature does not apply when NAT/PAT needs to support multiple Client Hello messages.
- The feature does not support client authentication with PQC algorithm certificates. If a server requests a client certificate, sessions will be bypassed, and PQC client certificates will not be intercepted..
- HSM Luna and NCipher do not support Post-Quantum Cryptography (PQC).
- The new PQC key type is not compatible with the GEN2 Inline SSL flex configuration.
- This feature is also incompatible with the Classic Inline Bypass configuration.

Attributes of Inline TLS/SSL Decryption Solution

This section describes the various attributes that construct an inline TLS/SSL decryption solution. You can configure these attributes in the Inline SSL App based on your deployment type. For details, refer to Configure Flexible Inline TLS/SSL Decryption Solution.

Refer to the following sections for more details:

TLS/SSL Sessions

Secure Sockets Layer (SSL) is a protocol that enables secure data transfer between a client and server. Transport Layer Security (TLS) builds on SSL and adds stronger cryptographic protection for TCP/IP communication.

A TLS/SSL session begins when the client initiates a connection. The GigaVUE node intercepts this request and negotiates a TLS/SSL session with the client. The node monitors all TCP connections and intercepts TLS/SSL sessions, while non-TCP traffic passes through unchanged.

All incoming TLS/SSL traffic terminates on the GigaVUE node. The node decrypts traffic for inbound or outbound deployments, forwards it to inline tools, and then passes it to the server.

The GigaVUE node maintains session details such as the client's IP address, allowing it to "reconnect" the client and server after decryption. The node also negotiates a new TLS/SSL session with the server. TLS/SSL traffic can be sent to tools inline or out-of-band.

TLS/SSL Handshake

TLS/SSL encryption secures traffic between client and server. Decryption uses keys to decode this traffic. TLS/SSL protocols define a series of handshake messages that set up and tear down a secure connection.

During the handshake, the client and server use Public Key Infrastructure (PKI) to exchange encryption keys for data transfer.

TLS/SSL Handshake Steps:

1. The client sends a Client Hello message with its supported TLS/SSL version, cipher suites, and a random number.

2. The server replies with a Server Hello message. It selects a cipher suite, sends its supported version, a random number, and its digital certificate.

3. The client verifies the server's certificate.

4. Using the random numbers, the client creates a pre_master_secret, encrypts it with the server's public key, and sends it.

5. The server verifies the client certificate if requested.

6. The client sends a Client Finished message, encrypted with the shared secret key, confirming its part of the handshake.

7. The server responds with a Server Finished message, also encrypted, confirming its part of the handshake.

8. The session is now established, and both sides use the shared secret key for encrypted communication.

TCP Transition States during TLS/SSL Session

The TCP handshake moves through different states during a TLS/SSL session:

- (Na:SYN RCV:Nb:SYN_SENT:INIT:INIT): Client SYN received; TCP handshake starts on the server side. If the server is busy or unreachable, the session resets.
- (Nb:SYN_RCV: EST:INIT:INIT): TCP handshake is complete on the server side.
- (Na:EST:Nb:EST:Na:INIT:Nb:INIT): TCP handshake complete between client and server; session may be decrypted depending on policy.
- (Na:EST:Nb:EST:Ta:SYN_SENT:Tb:INIT): Decryption decision made; tool-side TCP handshake begins.
- (EST:EST:Ta:SYN_SENT:Tb:SYN_RCV): Tool-side TCP handshake continues.
- (EST:EST:EST:EST): TCP handshake successfully established

TLS/SSL Session Resumption

TLS/SSL sessions can be resumed to improve performance. TLS/SSL session resumption speeds up the TLS/SSL handshake.

Once a session has been established, the keys are saved so a session can be resumed efficiently later. The resumed TLS/SSL handshake has fewer steps.

Session identifier-based resumption is supported. The GigaVUE-FM,GigaVUE HC Series node maintains the session identifier data in the cache. Session ticket-based resumption is not supported. By default, resumption is enabled.

TLS/SSL Session Search

You can search for an existing TLS/SSL session by hostname. The system matches the input against the Server Name Indication (SNI) or the certificate subject name of current sessions.

StartTLS and HTTP CONNECT

When the system detects a CLIENT HELLO packet, it switches to TLS/SSL mode.

- StartTLS: Allows protocols like SMTP, IMAP, and POP3 to begin in plaintext, then upgrade to TLS/SSL on the same port. Up to 20 ports can be monitored for StartTLS traffic.

- HTTP CONNECT: Used by explicit proxies to establish an HTTP session, then upgrade to TLS. Detection is automatic.

Both methods add security by upgrading existing plain text protocols.

Behavior with StartTLS

- If StartTLS is used, Inline TLS/SSL decryption works if the decision to decrypt is made at the certificate phase.
- If the CLIENT HELLO packet does not have SNI, policy rules are applied during the certificate phase.

Explicit proxy traffic also applies policy rules during the certificate phase.

NOTE: StartTLS must be enabled for decryption of sessions that use explicit proxy.

TLS/SSL Keys and Certificates

The TLS/SSL protocol enables secure data transfer between a server and a client. Both ends use cryptographic keys to decode the transmission and certificates to establish trust.

A TLS/SSL certificate is a digital file that contains a public key, host information, and a digital signature from a Certificate Authority (CA). This certificate ensures that communication between two endpoints can be trusted.

The Inline TLS/SSL decryption solution acts as a Break-and-Inspect system.

- Outbound traffic – The Man-in-the-Middle (MitM) generates server certificates on the fly, signed by the installed Signing CA.
- Inbound traffic – Certificate generation is not required, but the server's private key and certificate chain must be provided to the MitM.

The inline TLS/SSL decryption solution uses two main components – Key Store and Trust Store.

Key Store

The key store holds private keys and certificate-key pairs. It can store up to 1000 key pairs. Keys can be used for:

- **Inbound deployment** – For decrypting traffic using the server's private key.
- **Outbound deployment** – For re-signing and re-encrypting traffic using the Signing CA.

The key store supports encrypted or password-protected PEM files to fetch or download private keys. It also supports ECDSA keys for both inbound and outbound deployments.

Trust Store

The trust store is used to validate server certificates. Gigamon only trusts certificates that are linked to a CA in the trust store. In other words, the certificate chain must connect back to one of the trusted root CAs.

Gigamon provides a default trust store, based on Mozilla's trusted CAs, which is updated periodically. You can replace it with your own trust store in PEM format, append new CA certificates, or delete specific ones. You can also search for a certificate using its SHA1 fingerprint.

NOTE: Whenever you modify the trust store, you must reconfigure the inline TLS/SSL profile for the changes to take effect.

Client Trust Store (for mTLS)

Inline TLS/SSL decryption with mTLS introduces a **Client Trust Store** dedicated to validating client certificates. When a server requests client authentication, GigVUE node validates the client certificate chain against this trust store and applies the configured policy (for example, whether to allow expired, self-signed, or unknown-CA certificates) to decide whether to complete or terminate the mTLS handshake. The client trust store is managed separately from the server trust store and is used only for inbound and hybrid mTLS deployments. Changes to a client trust store take effect only for Inline SSL profiles that already reference it, and only after those profiles are updated or redeployed.

NOTE: When a device that previously used inbound mTLS is later re-purposed for outbound Inline TLS/SSL decryption, remove all configured client trust store and client key store entries before you deploy outbound Inline SSL profiles. Otherwise, outbound sessions may be bypassed instead of decrypted.

Certificate Validation

Gigamon validates server certificates before re-signing them. This prevents untrusted certificates from being incorrectly accepted. The validation process includes:

- **Certificate expiration date and validity period** – The certificate must not be expired and must fall within the valid date range.
- **Certificate issued by trusted CA** – The GigaVUE node maintains a list of trusted root CAs. This list determines the certificates that the client will accept. The trust store acts as a trust anchor during certificate validation. The GigaVUE node validates that each incoming certificate chain is trusted by one of the certificates in the trust store.
- **Server name** – The hostname in the certificate must match the Server Name Indication (SNI). If the client does not send SNI, this check is skipped.

- **Certificate revocation check** – The system checks the certificate's status using OCSP or CRL. This requires internet access. The certificate revocation check determines the revocation status of the server certificate.

Certificates that pass validation are signed by the primary MitM CA. If a certificate fails validation but a security exception is configured, the secondary MitM CA signs the certificate instead.

If the certificate is self-signed, the re-signed certificate will also be self-signed. Typically, client applications only trust the primary MitM CA, which helps surface certificate errors to end users so they can decide whether to reject the connection.

If no exception is configured, connections with invalid certificates are dropped.

Types of Certificate Validation Errors

The types of Certificate Validation errors are as follows:

- **Expired** – The certificate validity period has ended.
- **Self-signed** – The certificate issuer and subject are the same.
- **Unknown CA** – The CA is not in the trust store or cannot be validated. (You may download and import the root certificate if you trust the site.)
- **Invalid** – The certificate does not match the SNI or contains incorrect signature/field information.

For expired, unknown CA, and invalid certificates, Gigamon re-signs the certificate with the secondary MitM CA. This allows users to either accept or reject the connection.

Client Authentication

A server can request a client certificate after sending the Server Hello message. The client responds with its certificate, and the server validates it. Gigamon supports client authentication for both inbound and outbound traffic. If client authentication is detected during the handshake, the connection is bypassed.

Client Authentication with mTLS enabled

When mTLS is enabled for an Inline SSL profile and a server requests a client certificate:

1. The client sends its certificate to Gigamon as part of the mTLS handshake.
2. GigavUE node validates the client certificate chain against the client trust store and applies the configured client-auth policy (expired, self-signed, unknown CA, revocation status).
3. If the client certificate is accepted:

- GigaVUE node completes the handshake on both client and server sides.

- The session is decrypted and forwarded to inline tools as in a normal inbound deployment.

4. If the client certificate is rejected according to policy:

- The TLS session is terminated with an appropriate error toward the client.
- No traffic is forwarded to the server or inline tools for that session

Re-Signed Certificates

As a Man-in-the-Middle (MitM), Gigamon re-signs certificates.

- Fields copied from original certificate:
  - subject name
  - certificate validity
  - subject alternative name
- Fields removed from the certificate:
  - authority information access
  - certificate policies
  - CRL distribution points
  - SCT list
- Fields set in the re-signed certificate:
  - certificate type - v3
  - issuer
  - version
  - public key
  - random serial number
  - signature algorithm/hash
  - thumbprint
  - v3 extensions:
    - basicConstraints CA—True

NOTE: As this is a CA certificate, basicConstraints is set to True. For leaf certificates, basicConstraints is set to False.

    - keyUsage—digitalSignature and keyEncipherment
    - extendedkeyUsage—serverAuth
    - subjectKeyIdentifier—hash
    - authorityKeyIdentifier—keyid,issuer:always

Certificate Revocation Status

When Gigamon re-issues certificates for outbound traffic, clients automatically trust them because the MitM CA is already installed in the browser's trust store. This prevents clients from directly checking the revocation status of the original certificate.

Instead, the GigaVUE node performs revocation checks before generating replacement certificates.

- Soft fail mode – Traffic continues even if the revocation status is unknown.
- Hard fail mode – Traffic is blocked until the certificate's revocation status is confirmed.

If a certificate is confirmed as revoked, all future TLS/SSL connections to that server are dropped. By default, revocation checks are disabled.

Methods to Check Certificate Revocation Status

To check the certificate revocation status utilize one of the following method:

- **Certificate Revocation List (CRL)** – An online database of revoked certificates published by each CA. The list contains the serial number of the revoked certificates and the reasons for the revocation. Any revoked certificate should not be trusted even if the signatures are valid. The CRL location is included in the "CRL distribution points" X.509 extension of the certificate.
- **Online Certificate Status Protocol (OCSP)** – An Internet protocol used for obtaining the revocation status of an X.509 certificate. A real-time check can be performed using the OCSP link in the "Authority Information Access" X.509 extension.

If both CRL and OCSP are enabled, OCSP is attempted first, followed by CRL. Both methods require internet connectivity.

Policy Profile

A Policy Profile defines how traffic is processed during Inline SSL/TLS decryption. It determines which traffic is decrypted and which remains encrypted, based on a set of match rules and handling options. Each rule includes a match condition and an action—either decrypt or no-decrypt. For example, administrators can decrypt all general traffic while excluding sensitive categories such as financial transactions.

A Policy Profile is made up of several components:

- **Policy Rules**: Match traffic conditions and apply either decrypt or no-decrypt actions.
- **Default Action**: Specifies the action taken when no rule matches.
- **URL Cache Miss Action**: Defines what to do when a rule based on URL category cannot be evaluated because category data is missing.
- **Certificate Handling**: Controls how certificates are validated for decrypted traffic. Administrators can allow or block expired, invalid, self-signed, or unknown-CA certificates and enable or disable certificate-revocation checks.
- **Tool Handling:** Determines whether decrypted or non-decrypted traffic is forwarded to inline tools for further analysis.

Types of Policy Rules

Policy rules can be created based on one or more attributes such as:

- URL category (for example, finance, healthcare, or social media)
- Hostname or domain name
- Server certificate issuer
- Source or destination IP address
- Source or destination port number
- VLAN identifier

Scale and Flexibility

A single policy profile can include up to 2,048 rules, allowing fine-grained control over decryption policies and precise traffic handling across complex network environments.

Policy Evaluation

Policy evaluation defines how SSL/TLS decryption policies are applied at different stages of a connection. Each phase uses specific inputs—such as IP addresses, ports, host names, and certificates—to decide whether traffic is decrypted or left encrypted.

The below image describes the policy validation flow.

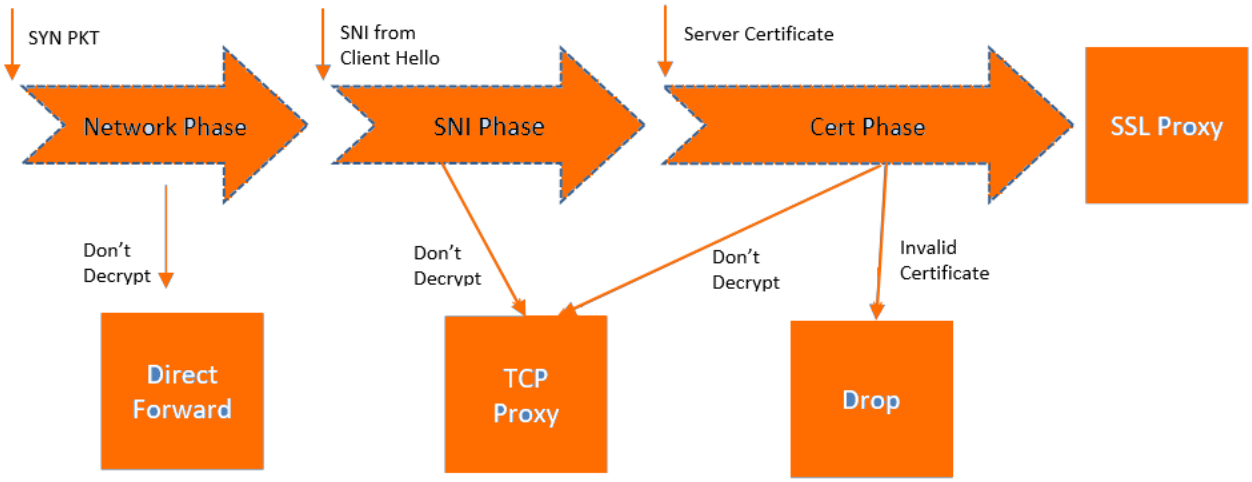


Figure 65 *Policy Validation Flow*

Network Phase

The Network Phase is evaluated on the TCP SYN packet.

- Inputs: source and destination IP addresses, source and destination ports, and VLAN identifiers.
- Evaluation Order: no-decrypt rules are evaluated before decrypt rules. Decrypt rules are checked in the following sequence: source IP, destination IP, source port, destination port, and VLAN.
- If a no-decrypt rule matches, packets are processed as non-proxy traffic using the bypass VLAN path (Direct Forward). This decision remains final for the lifetime of that TCP connection.
- If traffic is to be decrypted, processing continues to later phases. A decrypt verdict from the Network Phase can be overridden by a more specific rule in a subsequent phase.

For connections marked for decryption:

- If the Client Hello contains SNI, policy is further evaluated in the SNI and Certificate phases.
- If no SNI is present, only the Certificate Phase is evaluated.
- For HTTPS proxy connections, hostname information from the CONNECT request is also used during evaluation.

SNI Phase

The SNI Phase evaluates the Server Name Indication (SNI) field from the Client Hello.

- Input: the hostname obtained from SNI.
- Evaluation Order: no-decrypt (domain) → decrypt (domain) → no-decrypt (category) → decrypt (category).
- Wildcard domains (for example, *.example.com) can be used to match all sub-domains, while plain domains (for example, example.com) match only the exact host.
- If traffic is not decrypted, it is handled through the TCP Proxy path. For traffic that is decrypted, the plaintext version is sent through the configured tool path according to the decrypt tool-bypass settings.

When URL category-based rules are configured and a URL cache miss occurs, behavior depends on the URL Cache Miss Action:

- decrypt or no-decrypt — applied immediately.
- defer — causes a temporary delay.

For compliance reasons, no-decrypt is the recommended cache-miss action.

Certificate Validation

When traffic is marked for decryption (from the SNI phase or when SNI is absent), the system validates the server certificate using the configured trust store. Checks include:

- certificate expiry, hostname mismatch, and self-signed status;

- optional revocation status if enabled.

If the certificate is valid, a new server certificate is issued using the primary MitM CA. If validation fails, the connection is dropped unless a security exception is configured. When a secondary MitM CA is defined, it may be used to issue replacement certificates in such cases. The primary MitM CA is optional for inbound deployments.

Certificate Phase

The Certificate Phase is evaluated for all connections after certificate validation.

- Inputs: certificate issuer and subject name. The Common Name (CN) attribute is extracted from the certificate subject and used for policy matching.
- Evaluation Order: rules based on certificate subject (CN) first, then issuer-based rules using the CN and DN of the issuer.
- If traffic is not decrypted, it is processed as TCP Proxy traffic—the server SSL session is reset and the Client Hello is resent.

For HTTPS proxy connections, the Certificate Phase also evaluates policies using the hostname from the CONNECT request. If required, no-decrypt can be applied during this phase for such connections.

Together, these evaluation phases ensure that decryption behavior is accurate, rule-driven, and compliant, allowing administrators to define precise control over SSL/TLS traffic handling.

Policy Profile Options

This section describes a few of the options for the Inline TLS/SSL policy profile. Refer to the following sections:

- Inline TLS/SSL Decryption Port Map
- Enable or Disable Tool Bypass
- High Availability Active Standby
- Inline Network Group Multiple Entry
- Tool Early Engage
- One-Arm Mode
- Tool Early Inspect
- Inline TLS/SSL L3 Tool NAT/PAT Support

Inline TLS/SSL Decryption Port Map

The TCP destination port for decrypted traffic sent to inline tools can be configured in the decryption profile.

If not configured, the system automatically uses the same TCP destination port as the incoming traffic.

The appliance determines the TCP port for decrypted traffic based on two priority levels:

- Priority 1 — Port Map (User-Configured):You can define a mapping between the In Port and Out Port.
  - In Port refers to the TCP destination port received from the client.
  - Out Port is the TCP port used to send decrypted traffic to inline tools.
- Priority 2 — Default Out Port: This port is used when the incoming port does not match any entry in the configured port map.

NOTE: Decryption Port Mapping and Tool Early Inspect cannot be configured at the same time.

Enable or Disable Tool Bypass

Tool bypass determines whether specific traffic types are sent to inline tools. You can enable or disable tool bypass for the following types of traffic:

- TLS/SSL decrypted traffic
- Non-decrypted SSL traffic (non-TLS/SSL TCP)
- Non-TLS/SSL traffic (non-TCP)

By default, tool bypass is disabled for all traffic types. This means that all decrypted TLS/SSL, non-decrypted TLS/SSL, and non-TLS/SSL traffic is sent to the tools for inspection.

When you enable tool bypass for a traffic type, that specific traffic is excluded from tool processing and not sent to the inline tools.

High Availability Active Standby

Inline network High Availability (HA) in Active/Standby mode supports detection of link switch overs triggered by upstream devices.

When an upstream device, such as a firewall, performs an HA failover, the system identifies the change and continues forwarding traffic through the correct inline network.

For example, in an inline TLS/SSL network group with two network port pairs (Na1/Nb1 and Na2/Nb2), incoming traffic may shift from Na1 to Na2 during an upstream failover.

The system automatically detects this transition and maintains traffic continuity by routing packets through the appropriate inline link.

By default, this feature is disabled.

NOTE: Do not enable this option if the inline TLS/SSL network group operates in an Active/Active configuration.

Inline Network Group Multiple Entry

An inline network group topology can include multiple network port pairs (for example, Na1/Nb1 and Na2/Nb2).

With multiple port pairs, traffic from a network interface may traverse the system more than once.

Intercepted traffic from one interface can reenter the system through another interface within the same network group, as shown below.

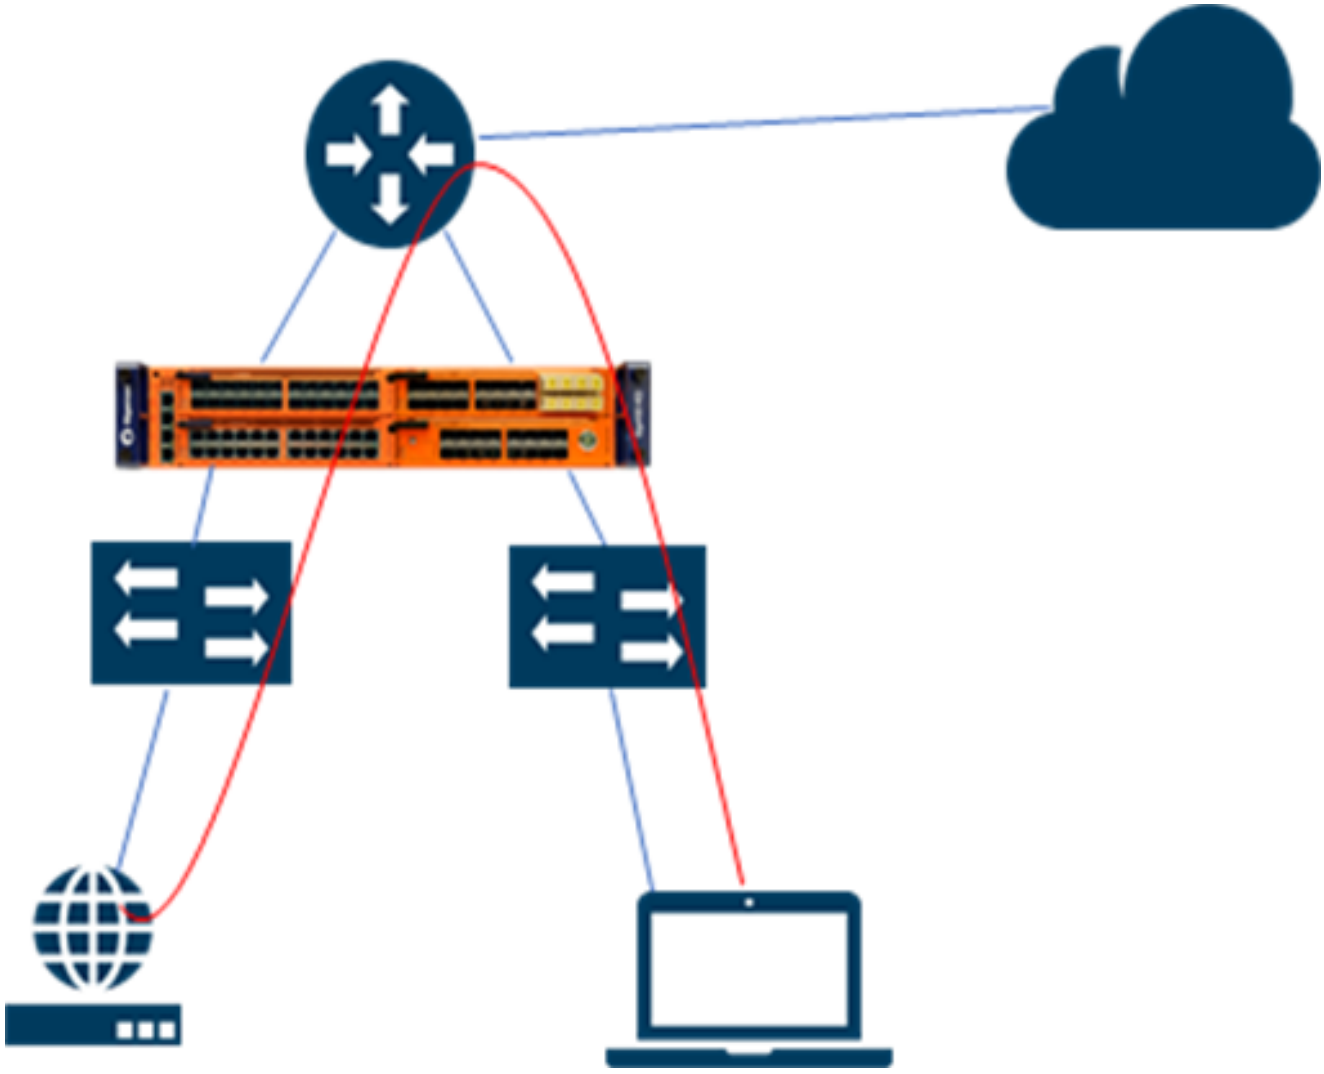


Figure 66 *Inline TLS/SSL Inline Network Group Configuration*

Inline TLS/SSL – Multiple Entry Behavior

When the inline TLS/SSL appliance is positioned between internal devices and the upstream router, traffic from internal devices to the Internet is intercepted by the appliance.

If internal devices associated with different network port pairs in the same group communicate with each other, their traffic first travels from the source device to the appliance, then to the upstream router, and finally reenters the appliance through another port pair to reach the destination device.

The appliance maintains the inline incoming interface (for example, Na1) for each connection.

If subsequent traffic from the same connection reenters through a different interface within the same group (for example, Na2), the appliance forwards the packets directly to the corresponding opposite interface (Nb2) without additional processing.

This enables seamless reentry handling for ongoing connections. Traffic that reenters through the same network port pair (for example, Nb2 → Na2) is not supported.

Similarly, any connection that uses more than the original network pair (for example, first packet on Na1/Nb1 and later through another pair) is unsupported.

All traffic for a connection must continue through the same port pair used by the first packet.

You can enable or disable Inline Network Group Multiple Entry in the profile settings. By default, this feature is disabled.

Tool Early Engage

In a Layer 3 topology, inline tools may need to modify the MAC address or VLAN IDs when forwarding client traffic back to the server.

The Tool Early Engage option allows inline tools to perform these modifications effectively.

When a client initiates a connection, the appliance first establishes the connection with the inline tool before connecting to the server.

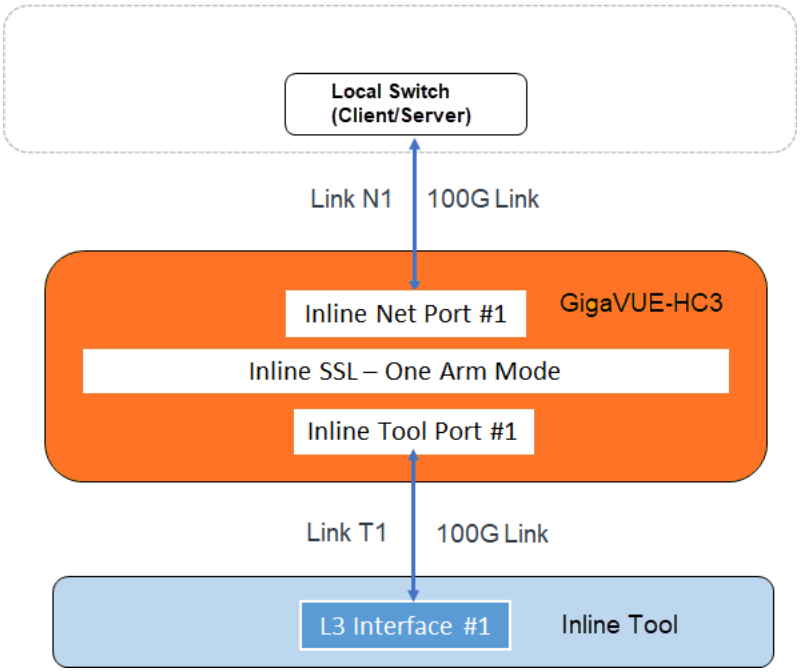
This early engagement enables the inline tool to adjust MAC addresses or VLAN IDs as needed when sending traffic back to the server.

Notes:

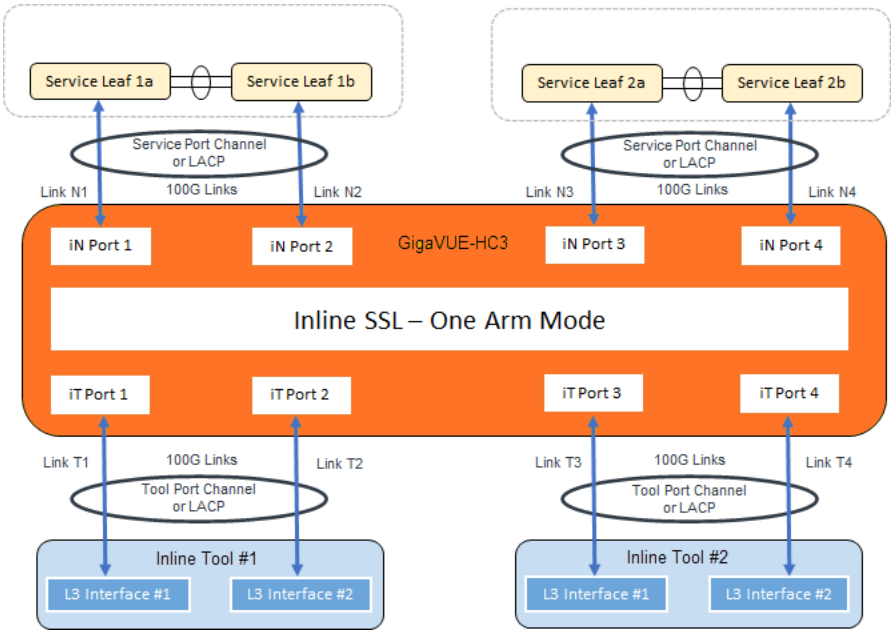
- Tool Early Engage can be enabled for a policy profile as a standalone feature, even when One-Arm mode is not used.
- Tool Early Engage and Tool Early Inspect cannot be configured together.

One-Arm Mode

In a Layer 3 topology, One-Arm mode allows both client and server traffic to traverse the same physical or logical link (for example, an aggregate port channel). You can configure this only after enabling Tool Early Engage.



For each connection between the client and server, there are two TCP sessions established between GigaSMART and the inline network and two TCP sessions established between GigaSMART and the inline tool. In the above diagram, you can see that with the One-Arm mode enabled, both TCP sessions from the inline network side arrive at GigaSMART on the same link N1. The TCP sessions from the inline tool side arrive at GigaSMART on the same link T1.



In the above figure, the inline network link and inline tool link works as a pair – (N1, T1), (N2, T2), (N3, T3), and (N4, T4). The GigaSMART sends traffic to the corresponding tool link of the received network link. Similarly, GigaSMART sends the traffic back to the server on the corresponding network link of the received inline tool link. For example, when a connection comes to GigaSMART from the inline network N1, after decryption, GigaSMART sends clear text traffic to inline tool on T1. The return traffic of the same connection arrives at GigaSMART on the inline tool link T2. GigaSMART then re-encrypts the traffic and sends it to the server on the inline network link N2.

Failover Support

If either member of a link pair fails, the paired link is also taken down to preserve flow consistency.

This behavior can be achieved by configuring **Link Aggregation Control Protocol (LACP)** or enabling **Link Failure Propagation (LFP**) between each network-tool link pair.

One-Arm Mode — Rules and Notes

Before enabling One-Arm mode for a policy profile, observe the following:

- Connect inline network interfaces to the Na side and inline tool interfaces to the Nb side.
- Tool Early Engage must be enabled.
- One-Arm mode is not supported on the Flexible Inline Decryption solution.
- To route packets through a firewall, configure the client and server gateways to use the router's IP address.
- Second-level OOB is not supported; use map-passall from the tool port instead.

- When both Tool Early Engage and One-Arm mode are enabled, MAC addresses and VLAN IDs can be changed, but IP addresses, ports, and protocols cannot.
- Inline Network Group Multiple Entry and High Availability Active/Standby features are not supported.
- Enable LACP on both network and tool sides, and LFP on the inline-network configuration.
- Bypass-tool options in the Inline TLS/SSL profile (Decrypt, No-Decrypt, Non-TLS/SSL) require another router on the network side to route packets.
- The name “one-arm” is a reserved keyword. Do not use it as an alias for any inline network, tool, or map. Rename existing aliases before upgrading to GigaVUE-OS 5.12.xx.
- Resilient Inline Arrangement (RIA) is not supported with One-Arm mode.
- If “one-arm” is configured as a tool in an inline second-level map, the VPort status displays as *Up (not Up (Normal))*.
- One-Arm mode cannot coexist with the **Tool Early Inspect** feature.

Tool Early Inspect

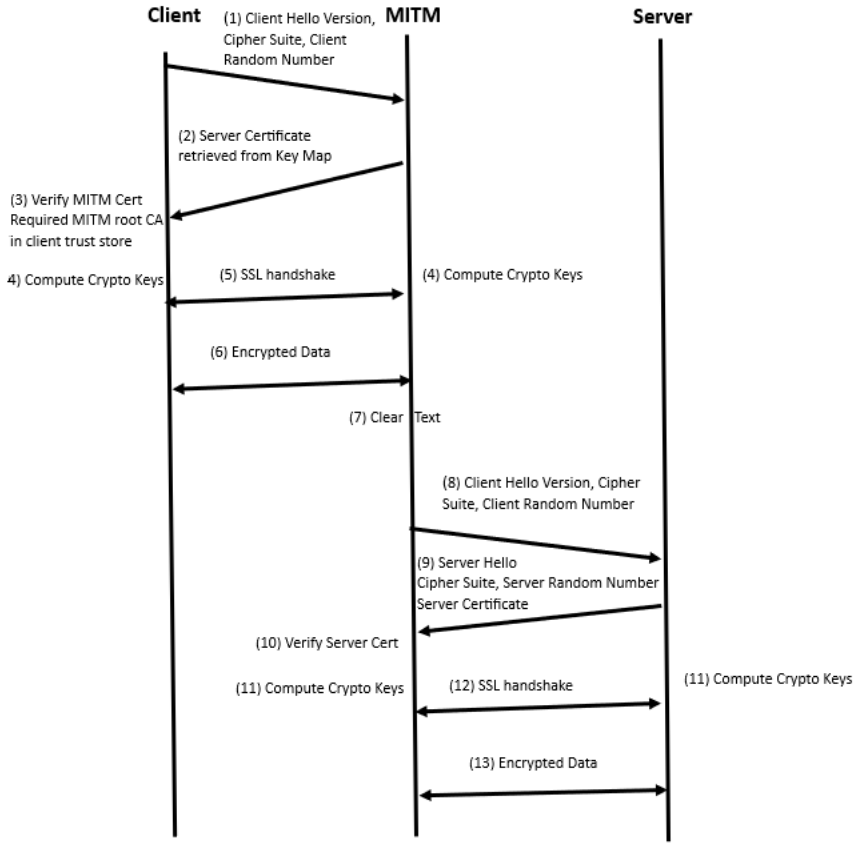
In a standard Inline TLS/SSL deployment, the node intercepts the TLS/SSL session between the client and server and forwards decrypted data to the inline tool after completing the handshake with both ends.

This means that connections rejected by the tool are still fully established to the server.With Tool Early Inspect , the client’s handshake completes first using the configured server certificate and key.

Decrypted data is then sent to the inline tool for inspection before connecting to the real server, ensuring only valid sessions proceed.

NOTE: Tool Early Inspect is supported only in inbound deployment modules. Outbound and hybrid modules are not supported.

The below diagram shows how a client handshake sequence would be once the Tool Early Inspect is enabled.



Inline TLS/SSL L3 Tool NAT/PAT Support

This feature offloads TLS decryption from Layer-3 inline tools that perform NAT/PAT (Network/Port Address Translation).

The GigaSMART engine maintains two independent sessions—one toward the client and one toward the server—to achieve this separation.

Supported Platforms

- Gen3 cards in GigaVUE-HC1 and GigaVUE-HC3
- GigaVUE-HC1-Plus

HTTP 2.0 Downgrade

When NAT/PAT mode is enabled, the HTTP 2.0 Downgrade option is enabled by default. HTTP 2.0 traffic is downgraded to HTTP 1.1 for decryption. If the downgrade option is disabled, HTTP 2.0 traffic is forwarded without decryption.

Decryption Port Mapping

In the Inline SSL application profile , you can now sent decrypted traffic to user-defined L4 ports. This feature supports the following scenarios:

- One-to-One Tool Port Address Translation: A specific clear text port for Inline SSL traffic is assigned, and after decryption, the flow is directed to this assigned port
- Many-to-One Tool Port Address Translation: Multiple incoming SSL Layer 4 ports are mapped to a specific single clear text Layer 4 port. If no specific mapping exists, decrypted traffic will be directed to the designated clear text port
- No Port Mapping- If one-to-one or default port mappings are not configured, decrypted traffic will continue to use the same original L4 port from the incoming encrypted data.

To configure this feature, enable the NAT/PAT Mode and then configure the port details in TCP Port MAP Decryption.
The port details can be configured through the *apps inline-ssl* command. Refer to GigaVUE-OS CLI Reference Guide.

Limitations

- L3 Tool port address translation does not apply to web proxy scenarios. Therefore, support for tool port address translation will not be included for web proxy traffic.
- The Start TLS port should not be configured in any port mapping settings, whether one-to-one or as a default port map.
- Any L4 port expected to receive the first data from the server must not be included in the port mapping configuration.

Cache Server Certificate Timeout

The server information is cached for performance optimization. The default time out is 30 minutes. The cache is flushed when the cache timeout value is set to zero. The cache is disabled when the timeout value is set to zero.

Refer to the following Gigamon Validated Design for more information:

- Offloading TLS Decryption for an One-Armed Inline Tool in L3 with NAT/PAT Mode
- Enabling GigaSECURE TLS Decryption to Offload SSL Inspection from Next-Generation Firewall

Limitations

- Decrypted data visible to inline tools is limited to HTTP/1.1 over TLS.
- StartTLS traffic is not decrypted.
- Tool bypass is not supported because all packets must pass through the inline tool.
- Cannot coexist with Network Group Multiple Entry, Inline Network HA, RIA, Tool Early Engage, Tool Early Inspect, or One-Arm mode.
- IPv6 is unsupported in release 6.1.00 but supported from 6.2.00 onward.

Cache Persistence

There are four in-memory caches as follows. They are not configurable.

- Re-signed certificate cache
- URL category cache
- Revocation certificate cache
- Session resumption cache

Caches are maintained for Internet lookups such as URL categorization and certificate revocation checks using OCSP or CRL for faster subsequent lookups. The cache persistence feature allows the information to be saved on the GigaVUE-FM,GigaVUE HC Series node in the control card's persistent storage so that it can be retrieved in case of reboots. This allows the GigaSMART card to start with the information learned earlier. This feature is enabled by default and can be disabled if needed.

On the **Cache Persistence** page, you can:

- Search for specific entries in the caches using the **Find Entries** option
- Clear the caches using the **Clear Store** option from the **Actions** drop-down menu
- Display a summary of the records.

GigaSMART Overload Bypass

Packet buffers, CPU, memory utilization and concurrent connections are monitored for overloaded conditions. GigaSMART goes to bypass when resource usage exceeds thresholds. Existing connections will continue to be processed by GigaSMART, but any new connections will be bypassed.

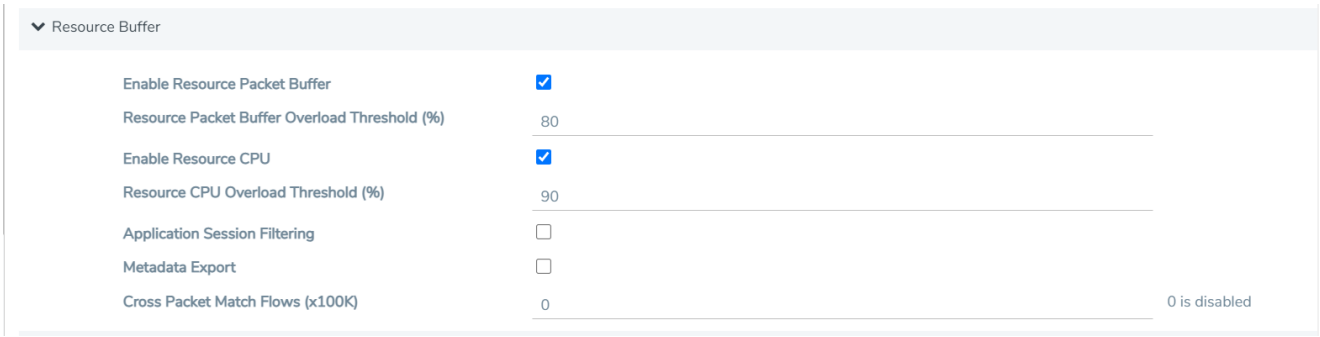
Table 8: Overload Bypass Connections and Thresholds

Criteria	GigaVUE-HC1	GigaVUE-HC3 (per GigaSMART Engine)
Maximum connections per second	1500	5000
Maximum connections	100000	200000

Criteria	GigaVUE-HC1	GigaVUE-HC3 (per GigaSMART Engine)
Resource Packet Buffer	Overload threshold for packet buffer resources for GigaSMART operations. Default is 80% (configurable)	
Resource CPU	Rising threshold for GigaSMART CPU statistics. Default is 90% (configurable)	
Heap exhaust	80%	

To configure the Packet Buffer and CPU threshold values, navigate to **GigaSMART** > **GigaSMART Operations (GSOP)** > **Resource Buffer** and configure the following:

- Resource Packet Buffer Overload Threshold (%)
- Resource CPU Overload Threshold (%)



CPU Overload Threshold

Due to sudden bursts of traffic, the GigaSMART CPU can become too busy and drop packets. However, when a system or application reaches a threshold, SSL sessions can be bypassed. When a maximum CPU is reached, incoming connections will be bypassed.

When the CPU overload threshold is set to a configured value, (for example, 90%), the lower threshold is set to two-third of the CPU overload threshold configured (in this example 60%). A mean threshold is calculated, which will be the average of the CPU overload threshold and the lower threshold (in this example 75%).

The following actions will be taken:

- If the CPU hits the overload threshold, all new SSL connections will be bypassed.
- If the CPU reduces to the mean threshold, half of the new SSL connections will be bypassed.
- If the CPU reduces further to the lower threshold, all new SSL connections will be decrypted.

If you choose connectivity-over-security, the CPU overload threshold must be set to the lower threshold value.

Inline TLS/SSL Monitor Mode

Use the inline TLS/SSL monitor mode to assist in understanding your network topology. Monitor mode provides information about the traffic going to the GigaSMART card, which can help to learn about your deployment. When monitor mode is enabled, the monitor application collects information such as TCP ports used and VLAN information about the incoming traffic.

After inline TLS/SSL decryption is configured and monitor mode is enabled, the inline TLS/SSL application does not terminate the session. Instead, the monitor application collects information and forwards packets to the tool port or network port based on the configuration of the non-TLS/SSL TCP bypass action. For any Monitor mode, you can enable or disable seamlessly without any other configuration changes.

Monitor mode is disabled by default. To enable the monitor mode, refer to Configure TLS/SSL Decryption Solutions.

For packets coming from the network port, the monitor application collects packet flow information.

From the information collected from monitor mode, you can analyze the following cases:

- duplicate TCP SYN—For a given session, the SYN messages with a different packet signature than 5tuple, for example, a different VLAN ID, indicates the packet is coming from multiple paths.
- asymmetric routing—For a given session, packets arriving from multiple network interfaces indicates a packet is coming from multiple paths.

Inline SSL Monitor mode only captures TCP information, not SSL information. However Inline SSL Persistent Monitor mode captures both TCP and SSL information.

NOTE: Monitor mode is supported for standalone nodes only, not for nodes in a cluster.

Inline TLS/SSL Traffic Filtering

Because TLS/SSL/ connections can carry sensitive data, some organizations may require the TLS/SSL connections to avoid inspection. The SSL connections that carries user data such as financial or health care information can be bypassed without inspection, based on a configured policy.

Based on the decryption policies, some connections are not decrypted and are passed through, optionally through tools, without decryption. The inline SSL decryption solution respects data privacy and supports compliance.

Inline SSL decryption provides different ways to filter traffic, as follows:

- No-decrypt lists specify traffic to always pass through. A no-decrypt list policy states that traffic from certain sites should always skip decryption.
- Decrypt lists specify traffic to always decrypt. A decrypt policy states that traffic from certain sites should always be decrypted.
- Both No-decrypt lists and Decrypt lists IP addresses, IP subnets and explicit wildcards as domain rules.
- URL Web Services categorizes the URLs by their type, such as MyBank.com is a financial institution, so as a policy, do not decrypt that traffic. This is also called URL filtering. Typically, banking and health care information are not decrypted.
- Policy rules based on network attributes, such as
  - Source IPv4 address
  - Destination IPv4 address
  - VLAN
  - L4 port

No-decrypt Listing Policy

No-decrypt lists are typically used in environments where the default is to decrypt, excepting for certain sites or classes of sites which cannot be decrypted for legal or compliance reasons. By default, traffic that is not to be decrypted is forwarded to the tools unless otherwise configured.

A no-decrypt list file can contain a maximum of 10,000 entries.

Decrypt Listing Policy

Decrypt listing is typically used at sites where specific classes of connections must be decrypted, although the default for other traffic is not to decrypt. Decrypt listed domains and host names will always be decrypted.

A decrypt list file can contain a maximum of 10,000 entries.

No-Decrypt/Decrypt List Policy — Rules and Notes

While configuring a No-Decrypt / Decrypt policy keep in mind the following rules and notes:

1. The maximum domain/hostnames support per list is 10000.
2. IP Subnets are supported from 5.13.01 version. Example, 10.10.10.0/24.
3. Special characters are not supported unless they are used to define domain names, such as *** . - @** are supported for domain names and **/** is supported if IP subnet is defined. # is supported to comment out a line. Example of a text file format would be as follows:
   - *.google.com
   - www.gigamon.com
   - gigamon.com
   - domain-registration.com.us
   - 10.10.1.1
   - 10.10.1.0/24
4. Range of IP addresses are not supported example, 10.10.10.10-20.
5. Use a newline for each entry. Adding characters such as **, ;** are not supported.
6. The domain gigamon.com as an entry matches only gigamon.com. To match all subdomains of gigamon.com on v5.9+, use *.gigamon.com.
7. If the system has large set of decrypt/no-decrypt list entries, GigaVUE-FM,GigaVUE HC Series stats page and CLI stats command does not display any output. Wait for 5 to 10 minutes after reloading to check the inline SSL show stats command in CLI and stats page in GigaVUE-FM,GigaVUE HC Series.

IP Address Subnet with Longest Prefix Match(LPM)

The No-decrypt and Decrypt database allows the user to utilize IP subnets. This allows the user to configure overlapping IP addresses, in decrypt and no-decrypt database. The decision to decrypt or no-decrypt will be based on the longest prefix match of the IP entries available in the decrypt /no-decrypt database.

The format is as follows *subnet (no space) /prefix*.Eg: 191.1.1.0/32

URL Categorization

URL categories make it convenient to apply policies on all the possible URLs matching the category and reduce the number of policy rules. Categorization is based on the hostname in the TLS Server Name Indication (SNI) or the subject name from the server certificate if there is no SNI. There are 83 categories including one for Uncategorized, which is a default category for URLs that do not match any of the other 108 categories. The categories are fixed meaning that categories cannot be added, deleted or modified.

GigaSMART ships with a local database of 1M entries and will also perform a cloud lookup for those hosts not found in the local database. The URL Web Service provides the URL categorization. The URL database is updated daily from the URL Web Service. Each update likely adds new entries and purges other entries, but always keeping the database at 1M entries.

NOTE: When a URL is not in the cache, for cloud look-ups the stack port interface on GigaSMART must be configured to provide Internet access.

URL Category Look-ups and Caching

As part of the Inline SSL processing, URL category look-ups are performed against the database. If the URL is not found in the database, then a lookup is performed against the local cache. If the URL is not found in the local cache, then an external lookup to the URL Web Services may be performed, if configured. If the URL is found in the external look-up, then it is dynamically saved in the local cache. Future look-ups may then find the URL in the local cache instead of requiring the external look-up.

NOTE:

- For TLS connections containing SNI in the Client Hello, do not perform URL category look-up in the certificate phase.
- CN based evaluation can be performed using the configuration option.

The local cache can hold up to 250k entries (in addition to the 1M entry database). The local cache works like a circular buffer – older entries are discarded to make room for newer ones if the cache is full. Each cache entry is valid for 24 hours and updated with current time stamp whenever an entry is made. If an expired entry is encountered, a new query is issued to the URL Web Services to refresh the entry in the cache. Expired entries don't get actively deleted from the cache.

While the URL Web Service is hosted on AWS, external look-ups need to occur very quickly. Gigamon provides a timeout option, up to 10 seconds for external URL category look-ups via the URL cache miss defer option.

Notes:

- URLs may get re-categorized as part of updates from the URL Web Services. This is transparent to Gigamon and customers.
- The URL category classification is fixed, and a new category cannot be added. Gigamon provides the no-decrypt list/decrypt list functionality, which can achieve the same result as creating a custom category.
- If a URL belongs to multiple categories, any no-decrypt policy would take precedence over all decrypt policies.

Inline SSL URL categories

The following are the list of Inline SSL URL categories with examples.

NOTE: Gigamon does not endorse any of the following categories, descriptions, and examples, but replicated the information from the URL Web Services. Some categories are presented without examples since they are not appropriate.

Category Name	Description and Examples
Abortion	Abortion topics, either pro-abortion and anti-abortion.
Abused Drugs	Discussion or remedies for illegal, illicit, or abused drugs such as heroin, cocaine, or other street drugs. This category includes information on the misuse of non-proscribed substances (eg. "glue sniffing"), or the misuse of prescription medications.
Adult and Pornography	Sexually explicit material for the purpose of arousing a sexual or prurient interest. Online groups, including newsgroups and forums, that are sexually explicit in nature.
Alcohol and Tobacco	Sites that provide information on, promote, or support the sale of alcoholic beverages or tobacco products and associated paraphernalia.
Auctions	Sites that support the offering and purchasing of goods between individuals as their main purpose. Does not include classified advertisements. http://ebay.co http://quibids.com
Botnets	These are URLs, typically IP addresses, which are determined to be part of a Bot network, from which network attacks are launched. Attacks may include SPAM messages, DOS, SQL injections, proxy jacking, and other unsolicited contacts.
Business and Economy	Business firms, corporate websites, business information, economics, marketing, management, and entrepreneurship. http://samsung.com http://ups.com
Content Delivery Networks	Delivery of content and data for third parties, including ads, media, files, images, and video. http://metacdn.co http://edgestream.com
Cheating	Sites that support cheating on examinations and contain such materials, including free essays, exam copies, plagiarism, etc.
Computer and Internet Info	General computer and Internet sites, technical information. SaaS sites and other URLs that deliver internet services. http://ranking.co http://system.netsuite.com
Computer and Internet Security	Computer/Internet security, security discussion groups. http://siteadvisor.co http://webroot.com

Category Name	Description and Examples
Confirmed Spam Sources	Confirmed SPAM sources.
Cult and Occult	Internet resources which include discussion of astrology, spells, curses, magical powers, satanic rituals or supernatural beings. This includes horoscope sites.
Dating	Dating websites focused on establishing personal relationships. http://eharmony.com
Dead Sites	These are dead sites that do not respond to http queries. Policy engines should usually treat these as “Uncategorized” sites. http://g00gle.com http://whitehouse.info
Dynamic Content	Domains that generate content dynamically based on arguments to their URL or other information (like geo-location) on the incoming web request. booking.com
Education Institution	Pre-school, elementary, secondary, high school, college, university, and vocational school and other educational content and information including enrollment, tuition, and syllabus. http://mit.edu http://ox.ac.uk
Entertainment and Arts	Motion pictures, videos, television, music and programming guides, books, comics, movie theatres, galleries, artists or reviews on entertainment. http://eonline.com http://warnerbros.com
Fashion and Beauty	Fashion or glamour magazines, beauty, clothes, cosmetics, style. http://visionmodels.co.uk http://genejuarez.com
Financial Services	Banking services and other types of financial information, such as loans, accountancy, actuaries, banks, mortgages, and general insurance companies. Does not include sites that offer market information, brokerage or trading services. http://firstpremierbankcards.com http://paypal.com
Gambling	Gambling or lottery web sites that invite the use of real or virtual money. Information or advice for placing wagers, participating in lotteries, gambling, or running numbers. Virtual casinos and offshore gambling ventures. Sports picks and betting pools.
Games	Playing or downloading, video games, computer games, electronic games, tips, and advice on games or how to obtain cheat codes. Also includes sites dedicated to selling board games as well as journals and magazines dedicated to game playing. http://duowan.com http://ubi.com
Government	Information on government, government agencies and government services such as taxation, public, and emergency services. Also includes sites that discuss or explain laws of various governmental entities. Includes local, county, state, and national government sites.

Category Name	Description and Examples
	http://www.nasa.gov http://premier-ministre.gouv.fr
Gross	Sites that contain material which describe or display material which would be considered foul or disgusting. Examples would include bodily fluids, injuries, gore.
Hacking	Illegal or questionable access to or the use of communications equipment/software. Development and distribution of programs that may allow compromise of networks and systems.
Hate and Racism	Sites that contain content and language in support of hate crimes and racism.
Health and Medicine	General health, fitness, well-being, including traditional and non-traditional methods and topics. Medical information on ailments, various conditions, dentistry, psychiatry, optometry, and other specialties. http://webmd.com http://missionvalleymedical.com
Home and Garden	Home issues and products, including maintenance, home safety, decor, cooking, gardening, home electronics, design, etc. http://homedepot.com http://waysidegardens.com
Hunting and Fishing	Sport hunting, gun clubs, and fishing. http://fishingworks.com http://wildlifelicense.com
Illegal	Criminal activity, copyright and intellectual property violations, etc.
Image and Video Search	Photo and image searches, online photo albums/digital photo exchange, image hosting. http://images.google.fr http://gettyimages.com
Individual Stock Advice and Tools	Promotion and facilitation of securities trading and management of investment assets. Also includes information on financial investment strategies, quotes, and news. http://stockstar.com http://morningstar.com
Internet Communications	Internet telephony, messaging, VoIP services and related businesses. http://skype.com http://www.chatib.com/
Internet Portals	Web sites that aggregate a broader set of Internet content and topics, and which typically serve as the starting point for an end user. http://yahoo.com http://qq.com
Job Search	Assistance in finding employment, and tools for locating prospective employers, or employers looking for employees. http://monster.com http://51job.com

Category Name	Description and Examples
Keyloggers and Monitoring	Downloads and discussion of software agents that track a user's keystrokes or monitor their web surfing habits.
Kids	Sites designed specifically for children and teenagers. http://www.mundogaturro.com http://www.poptropica.com
Legal	Legal websites, law firms, discussions and analysis of legal issues. http://www.pepperlaw.com http://earlcaterlaw.com
Local Information	City guides and tourist information, including restaurants, area/regional information, and local points of interest. http://downtownlittlerock.com http://sandiegorestaurants.com
Malware Sites	Malicious content including executables, drive-by infection sites, malicious scripts, viruses, trojans, and code.
Marijuana	Marijuana use, cultivation, history, culture, legal issues.
Military	Information on military branches, armed services, and military history. http://defense.gov http://www.mod.uk
Motor Vehicles	Car reviews, vehicle purchasing or sales tips, parts catalogs. Auto trading, photos, discussion of vehicles including motorcycles, boats, cars, trucks and RVs. Journals and magazines on vehicle modifications. http://www.carmax.com http://carsales.com.au
Music	Music sales, distribution, streaming, information on musical groups and performances, lyrics, and the music business. http://itunes.com http://bandcamp.com
News and Media	Current events or contemporary issues. Also includes radio stations, magazines, online newspapers, headline news sites, newswire services, personalized news services, and weather sites. http://abcnews.go.com http://newsoftheworld.co.uk
Nudity	Nude or seminude depictions of the human body. These depictions are not necessarily sexual in intent or effect but may include sites containing nude paintings or photo galleries of artistic nature.
Online Greeting Cards	Online Greeting card sites. http://123greetings.com http://greeting-cards.com
Online Personal Storage	Online storage and posting of files, music, pictures, and other data. http://box.net http://freefilehosting.net
Open HTTP Proxies	The proxy servers that are accessible by any Internet user.

Category Name	Description and Examples
P2P (Peer to Peer)	Peer to peer clients and access that includes torrents, music download and programs.
Parked Sites	Parked domains are URLs which host limited content or click-through ads which may generate revenue for the hosting entities but generally do not contain content useful to the end user. Also includes Under Construction, folders, and web server default home pages. http://000.com http://buythisdomain.com
Pay to Surf	Sites that pay users in the form of cash or prizes, for clicking on or reading specific links, email, or web pages.
Personal Sites and Blogs	Personal websites posted by individuals or groups, as well as blogs. http://blogger.com http://wordpress.org
Philosophy and Political Advocacy	Politics, philosophy, discussions, promotion of a particular viewpoint or stance in order to further a cause. http://philosophynow.org http://political.com
Phising and Other Frauds	Phishing, pharming, and other sites that pose as a reputable site, usually to harvest personal information from a user. These sites are typically quite short-lived, so examples may not last long.
Private IP Addresses	RFC 1918, Address Allocation for Private Intranets. 10.0.0.0 - 10.255.255.255 (10/8 prefix) 192.168.0.0 - 192.168.255.255 (192.168/16 prefix)
Proxy Avoid and Anonymizers	Proxy servers and other methods to gain access to URLs in any way that bypasses URL filtering or monitoring. Web-based translation sites that circumvent filtering.
Questionable	Tasteless humor, “get rich quick” sites, and sites that manipulate the user experience or client in some unusual, unexpected, or suspicious manner.
Real Estate	Information on renting, buying, or selling real estate or properties. Tips on buying or selling a home. Real estate agents, rental or relocation services, and property improvement. http://prudentialproperties.com http://realtor.com
Recreation and Hobbies	Information, associations, forums and publications on recreational pastimes such as collecting, kit airplanes, outdoor activities such as hiking, camping, rock climbing, specific arts, craft, or techniques; animal and pet related information, including breed-specifics, training, shows and humane societies. http://petloverspublications.com http://craftster.org
Reference and Research	Personal, professional, or educational reference material, including online dictionaries, maps, census, almanacs, library catalogues, genealogy, and scientific information. http://reference.com http://wikipedia.org

Category Name	Description and Examples
Religion	Conventional or unconventional religious or quasi-religious subjects as well as churches, mosques, synagogues, or other places of worship. http://therocksandiego.org http://biblesociety.ca
Search Engines	Search interfaces using key words or phrases. Returned results may include text, websites, images, videos, and files. http://google.com http://sogou.com
Sex Education	Information on reproduction, sexual development, safe sex practices, sexually transmitted diseases, sexuality, birth control, sexual development, and contraceptives. http://sexetc.org
Shareware and Freeware	Sites that contains softwares, screensavers, icons, wallpapers, utilities, ringtones including downloads that request a donation on open source projects. http://download.com http://sourceforge.net
Shopping	Department stores, retail stores, company catalogs and other sites that allow online consumer or business shopping to purchase goods and services. http://amazon.com http://groupon.com
Social Network	Social networking sites that have user communities where users interact, post messages, pictures, and otherwise communicate. http://facebook.com http://twitter.com
Society	A variety of topics, groups, and associations relevant to the general populace, broad issues that impact a variety of people, including safety, children, societies, and philanthropic groups. http://dar.org http://unicefusa.org
Spam URLs	URLs contained in SPAM.
Sports	Team or conference web sites, international, national, college, professional scores and schedules; sports-related online magazines or newsletters, fantasy sports and virtual sports leagues. http://nba.com http://schoenen-dunk.de
Spyware and Adware	Spyware or Adware sites that provide or promote information gathering or tracking that is unknown to, or without the explicit consent of, the end user or the organization, also unsolicited advertising popups and programs that may be installed on a user's computer.
Stream Media	Sales, delivery, or streaming of audio or video content, including sites that provide downloads for such viewers. http://youtube.com http://ustream.tv

Category Name	Description and Examples
Swimsuits and Intimate Apparel	Swimsuits, intimate apparel or other types of suggestive clothing.
Training and Tool	Distance education, trade schools, online courses, vocational training, software training, and skills training. http://trainingtools.com http://prezi.com
Translation	Language translation sites that allow users to see URL pages in other languages. http://translate.google.com http://microsofttranslator.com
Travel	Airlines and flight booking agencies. Travel planning, reservations, vehicle rentals, car rentals, descriptions of travel destinations, promotions for hotels or casinos. http://cheapflights.com http://expedia.com
Uncategorized	Sites that have not been categorized by URL Web Service.
Unconfirmed Spam Sources	Unconfirmed SPAM sources.
Violence	Sites that advocate violence, depictions and methods, including game/comic violence, and suicide.
Weapons	Sales, reviews, descriptions of weapons such as guns, knives, martial arts accessories.
Web Advertisements	Advertisements, media, content, and banners. http://casalemedia.com http://justwebads.com
Web Based Email	Sites offering web-based email and email clients. http://google.com/mail http://foxmail.com
Web Hosting	Free or paid hosting services for web pages and information concerning their development, publication, and promotion. http://siteground.com http://bluehost.com

Proxy Server Profile for URL Categorization and Certificate Revocation status

To ensure a stable security network you can now redirect URL look-ups and Certificate Revocation status checks to a Proxy Server Profile. This Proxy Server profile will be attached to your Inline SSL deployment . To learn more refer to *Proxy Server Configuration* section in GigaVUE-Administration Guide

Inline TLS/SSL Decryption Deployments

Inline TLS/SSL decryption can be deployed in two main ways – Inbound and Outbound.

Inbound Deployment

In an inbound deployment, the client is on the Internet, and the server is inside the enterprise network. The GigaVUE node is also deployed inside the network, on the server side. To decrypt traffic, the GigaVUE node must have access to the server's private keys. This setup enables Man-in-the-Middle (MitM) decryption. The traffic uses Diffie-Hellman or Perfect Forward Secrecy (PFS) encryption.

Refer to Diagram 1 for an example of inbound inline TLS/SSL decryption deployment.

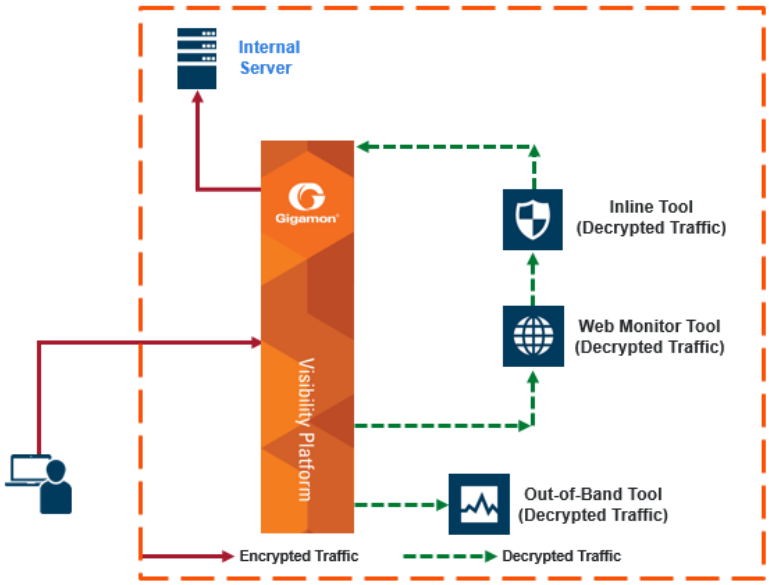


Figure 67 *Inbound Deployment of Inline TLS/SSL Decryption*

Inbound Deployment with Client Certificate Authentication (mTLS)

In a standard inbound deployment, the GigaVUE node is deployed on the server side and decrypts traffic destined for internal servers using their private keys. With mTLS enabled, the same topology can enforce client certificate authentication at the inline TLS/SSL layer.

In this deployment model:

- The client resides on the Internet or external network.
- The server and GigaVUE node reside inside the protected enterprise network.
- The server is configured to require client certificates for selected applications (for example, user-facing portals or API endpoints).
- The Inline TLS/SSL application:
  - Validates the server certificate against the server trust store.
  - Validates the client certificate against the client trust store.
  - Applies policy to accept or reject the session based on the outcome of both validations.

The resulting end-to-end connection is still an mTLS session from the client's perspective, but certificate validation and decryption are offloaded to the Inline TLS/SSL solution, which then forwards clear-text traffic to inline tools for inspection.

Outbound Deployment

In an outbound deployment, the client is inside the enterprise network, along with the GigaVUE node, which is deployed on the client side. The server is external, located on the Internet. In this case, the GigaVUE node acts as a trusted Man-in-the-Middle (MitM) but does not require access to the server's private key. It can still inspect TLS/SSL traffic.

Refer to Diagram 2 for an example of outbound inline TLS/SSL decryption deployment.

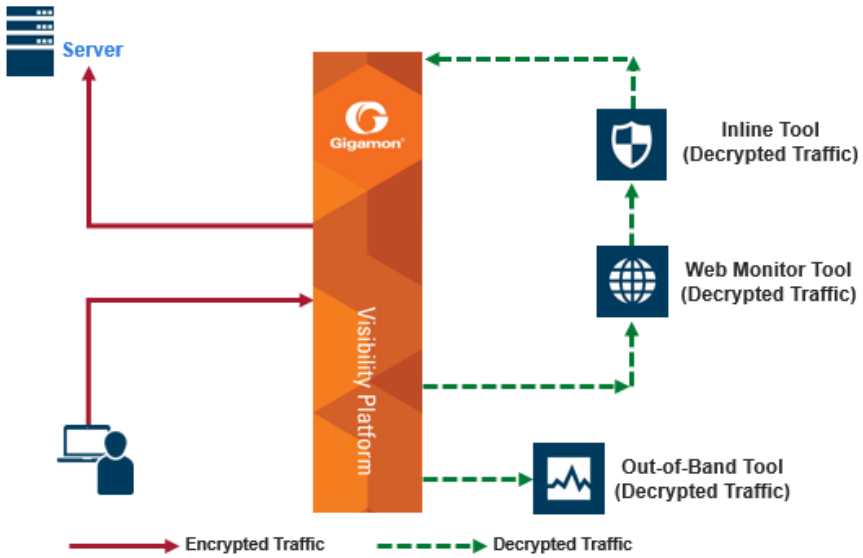


Figure 68 *Outbound Deployment of Inline TLS/SSL Decryption*

Basic Deployments for Inline TLS/SSL Decryption Solution

Choose from the following basic deployment methods to implement Inline TLS/SSL Decryption:

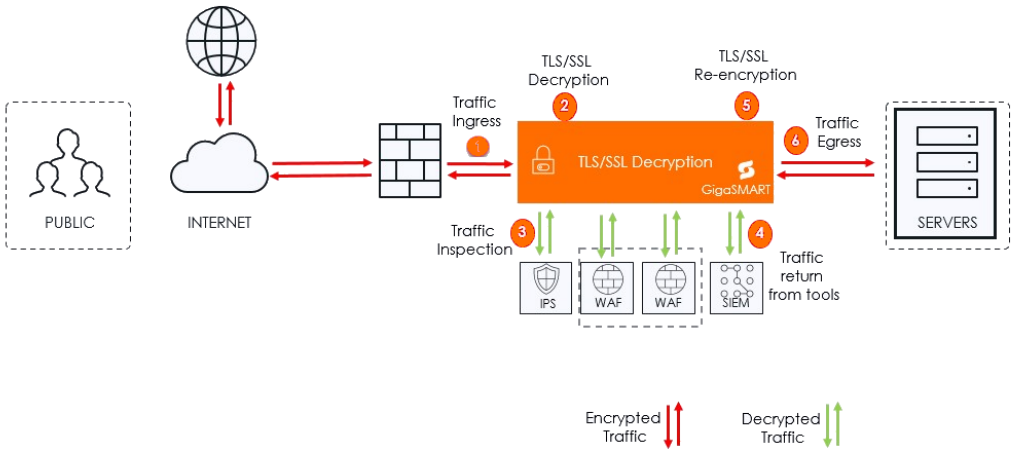
- Inline TLS/SSL Decryption Solution with Layer 2 Transparent Tools
- Inline TLS/SSL Decryption Solution with Layer 3 Tools (NAT-PAT)

Inline TLS/SSL Decryption Solution with Layer 2 Transparent Tools

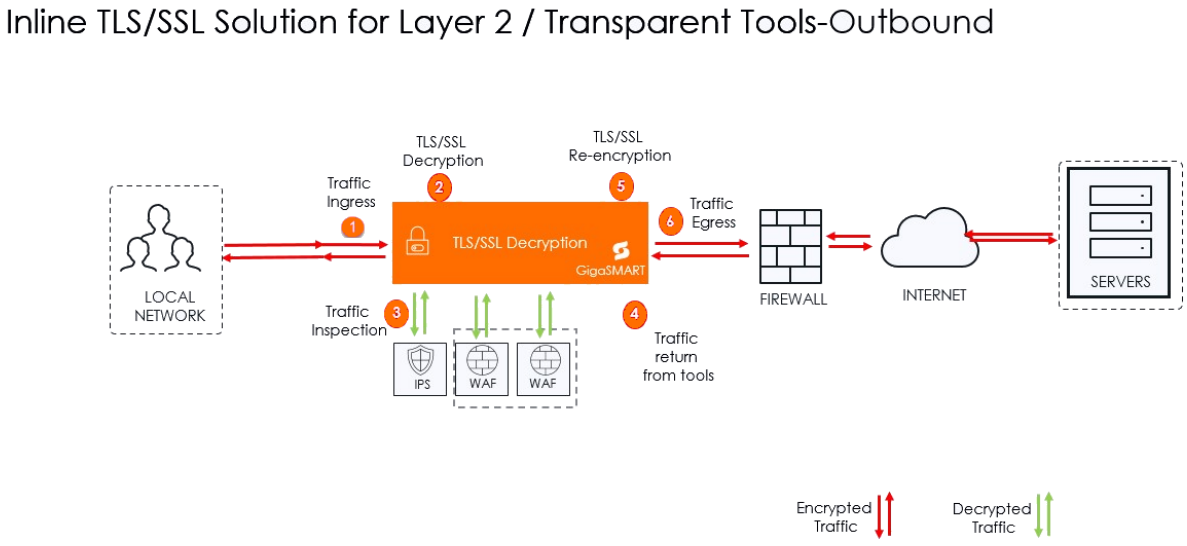
In an Inline SSL deployment, encrypted traffic passes through a sequence of components for secure decryption, inspection, and re-encryption. This deployment is designed to support Layer 2 inspection tools that analyze traffic without altering its contents.

The below shown is an inbound deployment of Inline TLS/SSL Decryption.

Inline TLS/SSL Solution for Layer 2 / Transparent Tools- Inbound



The shown is an outbound deployment of Inline TLS/SSL Decryption.



The following steps describe the end-to-end flow:

1. **Traffic Ingress**-Traffic enters the system through the Inline Network Port (Ingress). This port acts as the first point of entry for all inbound encrypted traffic.
2. **TLS/SSL Decryption**-The traffic is directed to the GigaSMART engine, where the inline TLS/SSL application is. The engine decrypts the TLS/SSL traffic, making it readable for inspection tools.
3. **Traffic Inspection**-The decrypted traffic is forwarded to inline security tools (for example, IPS, WAF, SIEM, etc). These tools operate in Layer 2 (transparent) mode, meaning they inspect but do not modify packet contents. Tools may allow, alert on, or block traffic based on inspection results, but they do not change the actual data.
4. **Traffic Return from Tools**- after inspection, the traffic is returned to the GigaSMART engine. This ensures that only inspected traffic proceeds further.
5. **SSL Re-encryption**-The GigaSMART engine re-encrypts the traffic using the original TLS/SSL parameters.
6. **Traffic Egress**-The re-encrypted traffic is sent out through the Inline Network Port (Egress) to continue to the servers.

Rules and Notes

The following table lists rules and notes that you need to be aware of while configuring your deployment.

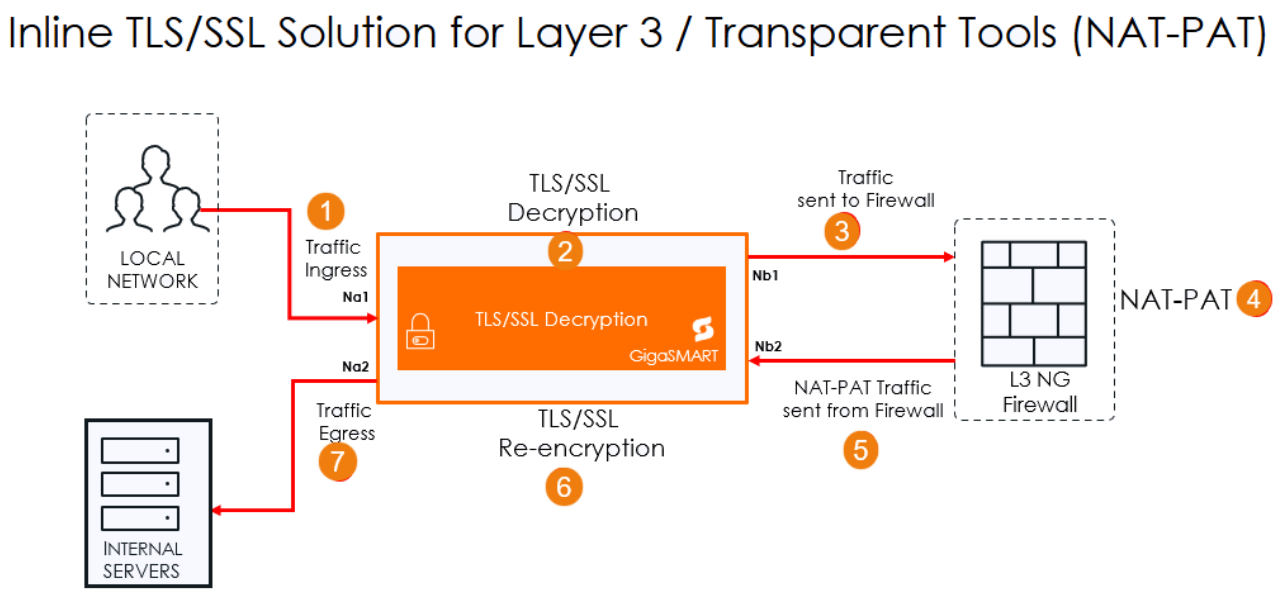
Rule / Feature	Details / Limitations
Exclusive Use	Inline tools in a flexible inline map cannot be used in classic inline or inline decryption maps. All inline networks and tools must belong to only one type of map.
Collector Maps	Only one unidirectional collector map is allowed for the same inline network. 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 (GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA25E, GigaVUE-TA25, GigaVUE-TA400, GigaVUE-TA400E)	- Physical Bypass (no BPS card) - Flexible and Resilient Inline SSL (no GigaSMART card) - GRIP (no BPS card) - ICAP (no GigaSMART card) - Classic Inline Bypass
VLAN Tagging and OOB Copy Limits (GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-HC1-Plus)	Flexible Inline Single VLAN Tag with monitoring mode may send incorrect VLAN tags. OOB copy packets may also have wrong tags. You cannot use BYPASS WITH MONITORING with MONITORING mode on the tool. OOB copy from inline network is not allowed in this mode.
Inline Map Limits — Bidirectional	GigaVUE-TA25,GigaVUE-TA25E, GigaVUE-HC1-Plus → 126 maps GigaVUE-HC1,GigaVUE-HC3 (CCv1 & CCv2), GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA400, GigaVUE-TA400E → 256 or 512 maps depending on setup.
Inline Map Limits — Unidirectional	GigaVUE-TA25,GigaVUE-TA25E, GigaVUE-HC1-Plus→ 252 mapsGigaVUE-HC1,GigaVUE-HC3 (CCv1 & CCv2),), GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA400, GigaVUE-TA400E → 512 or 1024 maps depending on setup.
Flexible Inline SSL Limits	Not supported with Inline Network LAG. Setting inline tools to “Drop” in the chain does not block Inline SSL traffic.
Filtering Limits(GigaVUE-TA400, GigaVUE-TA400E)	VLAN-based filtering in the Egress Port Filter for OOB copies is not supported. If one tool in the map is in monitoring mode, all tools must use the same mode. Asymmetric hashing (a-srcip-b-dstip and b-srcip-a-dstip) is not supported.
Protocol Pass-Through(GigaVUE-TA400, GigaVUE-TA400E)	CDP pass-through is not supported when the source is an Inline Network LAG. Bypass for LACP, CDP, and LLDP is supported.
Scaling Limits — GigaVUE-TA400, GigaVUE-TA400E	Max Inline Networks and Tools: 48 Max Inline Network LAG list: 24 Max Inline tools or tool groups per direction: 16 Max OOB copy entries per direction: 17 Max OOB copy ports per entry: 128

Inline TLS/SSL Decryption Solution with Layer 3 Tools (NAT-PAT)

A Layer 3 inline SSL decryption refers to Gigamon's deployment of its inline TLS/SSL decryption solution at the network layer (Layer 3 of the OSI model). In this mode, the GigaVUE node intercepts encrypted traffic flows between clients and servers, decrypts the SSL/TLS traffic, and subsequently delivers the decrypted packets to security tools for inspection. After the inspection, the traffic is re-encrypted and forwarded to its destination. Here are some key aspects of the solution:

- **Traffic Routing**: The solution allows for the configuration of subnet-based IP ranges at the ISSL profile level. This configuration determines whether traffic is routed through L3 tools. For example, traffic within a specific subnet range can be designated for L3 tools, which support NAT-PAT (Network Address Translation - Port Address Translation) for correlating client, server, and tool sessions.
- **Deployment Modes**: The solution can be deployed in a one-arm mode, which is suitable for enterprises looking to optimize their return on investment (ROI) by aligning with the capabilities of their inline tools.
- **Configuration and Management**: The deployment involves configuring flow maps to guide decrypted traffic to the appropriate tools, such as Next-Generation Firewalls (NGFWs). This includes setting up inline network bundles, flex maps, and inline SSL profiles.
- **Monitoring and Statistics:** The solution provides tools for monitoring traffic states, session statistics, and certificate hit counts. This helps in managing and optimizing the decryption process.

Architecture of an L3 Tool based Inline TLS/SSL Decryption Solution



The following steps describe the end-to-end flow:

1. **Traffic Ingress**-Traffic enters the system through the Inline Network Port (Ingress). This port acts as the first point of entry for all inbound encrypted traffic.
2. **TLS/SSL Decryption**-The traffic is directed to the GigaSMART engine, where the inline TLS/SSL application is. The engine decrypts the TLS/SSL traffic.
3. **Traffic Sent to Firewall**-The decrypted traffic is forwarded to the firewall.
4. **NAT-PAT**- The traffic is now either Network Address translated or Port Address translated.
5. **Traffic Sent from Firewall** :Traffic is being forwarded to the GigaSMART engine for re-encryption.
6. **SSL Re-encryption**-The GigaSMART engine re-encrypts the traffic using the TLS/SSL parameters.
7. **Traffic Egress**-The re-encrypted traffic is sent out through the Inline Network Port (Egress) to continue to the servers.

Rules and Notes

The following table lists rules and notes that you need to be aware of while configuring your deployment.

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 (GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA25E, GigaVUE-TA25, GigaVUE-TA400, GigaVUE-TA400E)	- Physical Bypass (no BPS card) - Flexible and Resilient Inline SSL (no GigaSMART card) - GRIP (no BPS card) - ICAP (no GigaSMART card) - Classic Inline Bypass
VLAN Tagging and OOB Copy Limits (GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-HC1-Plus)	Flexible Inline Single VLAN Tag with monitoring mode may send incorrect VLAN tags. OOB copy packets may also have wrong tags. You cannot use BYPASS WITH MONITORING with MONITORING mode on the tool. OOB copy from inline network is not allowed in this mode.
Inline Map Limits — Bidirectional	GigaVUE-TA25,GigaVUE-TA25E, GigaVUE-HC1-Plus → 126 maps GigaVUE-HC1,GigaVUE-HC3 (CCv1 & CCv2), GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA400, GigaVUE-TA400E → 256 or 512 maps depending on setup.
Inline Map Limits — Unidirectional	GigaVUE-TA25,GigaVUE-TA25E, GigaVUE-HC1-Plus→ 252 mapsGigaVUE-HC1,GigaVUE-HC3 (CCv1 & CCv2),), GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA400, GigaVUE-TA400E → 512 or 1024 maps depending on setup.
Flexible Inline SSL Limits	Not supported with Inline Network LAG. Setting inline tools to “Drop” in the chain does not block Inline SSL traffic.
Filtering Limits(GigaVUE-TA400, GigaVUE-TA400E)	VLAN-based filtering in the Egress Port Filter for OOB copies is not supported. If one tool in the map is in monitoring mode, all tools must use the same mode. Asymmetric hashing (a-srcip-b-dstip and b-srcip-a-dstip) is not supported.
Protocol Pass-Through(GigaVUE-TA400, GigaVUE-TA400E)	CDP pass-through is not supported when the source is an Inline Network LAG. Bypass for LACP, CDP, and LLDP is supported.
Scaling Limits — GigaVUE-TA400, GigaVUE-TA400E	Max Inline Networks and Tools: 48 Max Inline Network LAG list: 24 Max Inline tools or tool groups per direction: 16 Max OOB copy entries per direction: 17 Max OOB copy ports per entry: 128

Differences Between Layer 2 and Layer 3 Tools Deployments

Feature	L2 ISSL Decryption	L3 ISSL Decryption
Traffic Handling	This method processes traffic without the need for NAT-PAT, enabling a direct correlation between client and server sessions.	Supports NAT-PAT, with both IP and port or either of them. NAT-PAT enhances security by making it more difficult for unauthorized users to track and intercept network traffic, ultimately protecting sensitive information.
Protocol Support	Typically, it supports a wider variety of L7 protocols in OSI model without compromising performance.	The system primarily supports HTTP/1.1. Other protocols such as SMTPS and DNS-over-TLS are not supported.
Session Management	The GigaSMART maintains a single session dedicated to traffic processing.	The GigaSMART engine operates by maintaining two distinct sessions: one prior to its activation and another following its use.

Advanced Features for Inline TLS/SSL Decryption Solution

Choose from the following advanced features methods to implement Inline TLS/SSL Decryption:

- Inline TLS/SSL Decryption Solution with RIA
- Entrust nShield and Thales-Luna HSM for TLS/SSL Decryption for iSSL
- Inline TLS/SSL Decryption Solution with ICAP Client

Inline TLS/SSL Decryption Solution with RIA

GigaSMART Flex Inline Solution can now be configured in a Resilient Inline Arrangement (RIA), which is supported on all GigaVUE HC Series devices. To learn more about Resilient Inline Arrangements refer to Resilient Inline Arrangement.

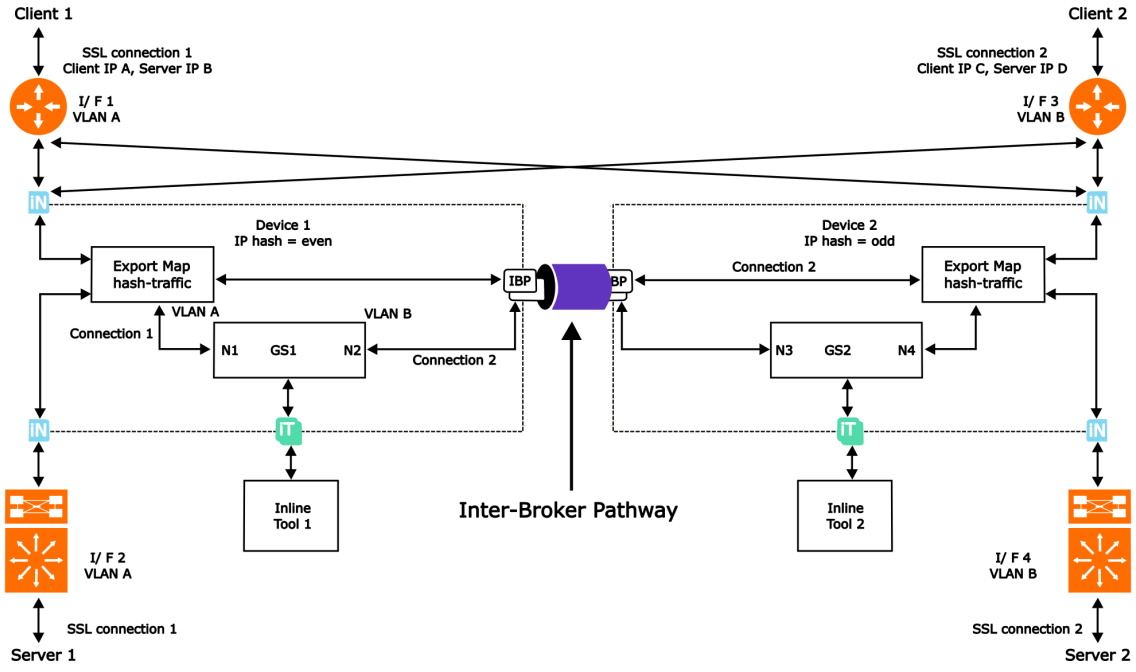
A Resilient Inline Arrangement utilizes two nodes to manage traffic in dual-path high availability environments. Both nodes process traffic simultaneously, using the source and destination IP addresses for decision-making. Traffic received from the top network interfaces is directed based on the source IP, while traffic from the bottom network

interfaces relies on the destination IP of the incoming traffic. Specifically, if the IP address ends with an even number, the traffic will be forwarded to one node; conversely, if it ends with an odd number, it will be sent to the other node.

For example:

In the Resilient Inline Arrangement described below, interface (I/F) 1 and I/F 2 are linked to node 1, while I/F 3 and I/F 4 connect to node 2. For connection 1, client 1's traffic from I/F 1, which has a VLAN A source and an IP address A, will be directed to GS1 in node 1 because the last decimal digit of IP A is even. Similarly, the traffic from server 1 of the same connection, originating from I/F 2 with VLAN A and destination IP A, will also be routed to GS1 in node 1.

For connection 2, since the last digit of IP C is even, the traffic from I/F 3 and I/F 4 will likewise be sent to GS1 in node 1. The traffic between GS1 and client 2 will occur on I/F 3 with VLAN B, requiring these packets to traverse the IBP that connects node 1 and node 2. In the same manner, the traffic between GS1 and server 2 will utilize I/F 4 after passing through the IBP.

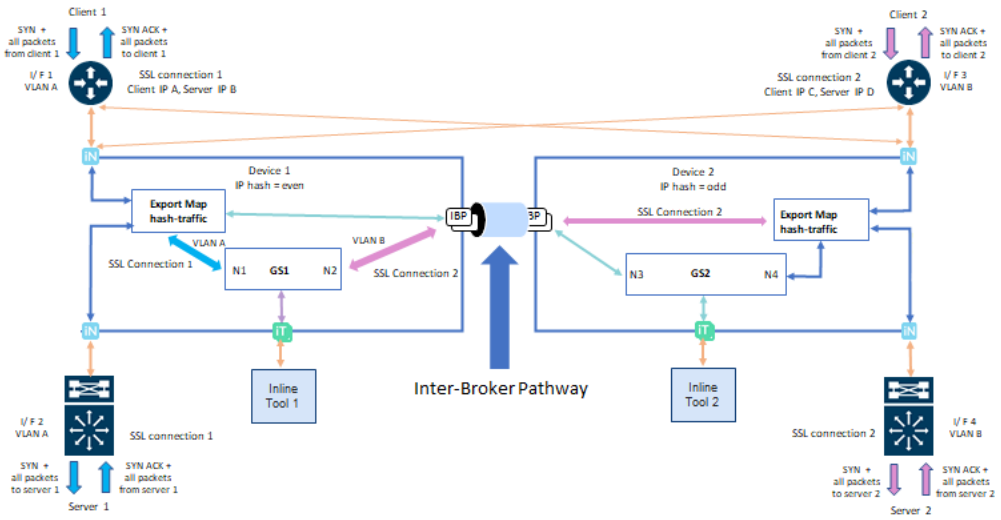


Symmetric Traffic in RIA

In a symmetric connection, the SYN packet from Client 1 arrives on Interface 1 (I/F 1) with VLAN A and is forwarded to GS1 in Node 1, as the source IP (IP A) is even. GS1 initiates a TCP connection to Server 1 by sending a SYN packet out from Interface 2 (I/F 2) with VLAN A.

Server 1 responds with a SYN-ACK from I/F 2 with VLAN A, routed back to GS1 since the destination IP (IP A) is even. Thus, all traffic between GS1 and Client 1 occurs on I/F 1, while traffic between GS1 and Server 1 uses I/F 2, connected to Node 1.

Traffic from Connection 2 on I/F 3 and I/F 4 with VLAN B is processed by GS1 in Node 1, as Client 2's IP (IP C) is also even. All traffic between GS1 and Client 2 flows through I/F 3 with VLAN B via the IBP, while traffic between GS1 and Server 2 is managed on I/F 4 with VLAN B, connected to Node 2.



Asymmetric Traffic in RIA

An asymmetric connection can occur in various scenarios:

Scenario 1

In the resilient inline arrangement below, the SYN packet for TLS/SSL connection 3 from client 3 is received on I/F 1 with VLAN A and forwarded to GS1 in device 1, as the source IP (IP A) is even. Server 3 responds with a SYN ACK from I/F 2 with VLAN A on device 1, also forwarded to GS1 since the destination IP (IP A) is even. GS1 then sends the SYN ACK back to client 3 on I/F 1 with VLAN A.

However, due to asymmetric routing or load balancing, some subsequent traffic from client 3 or server 3 related to connection 3 may arrive at I/F 3 or I/F 4 with VLAN B. This traffic is forwarded to node 1 through IBP, allowing symmetric traffic inspection. All outgoing traffic from GS1 to client 3 will be sent via I/F 1 with VLAN A, while traffic to server 3 will go through I/F 2 with VLAN A.

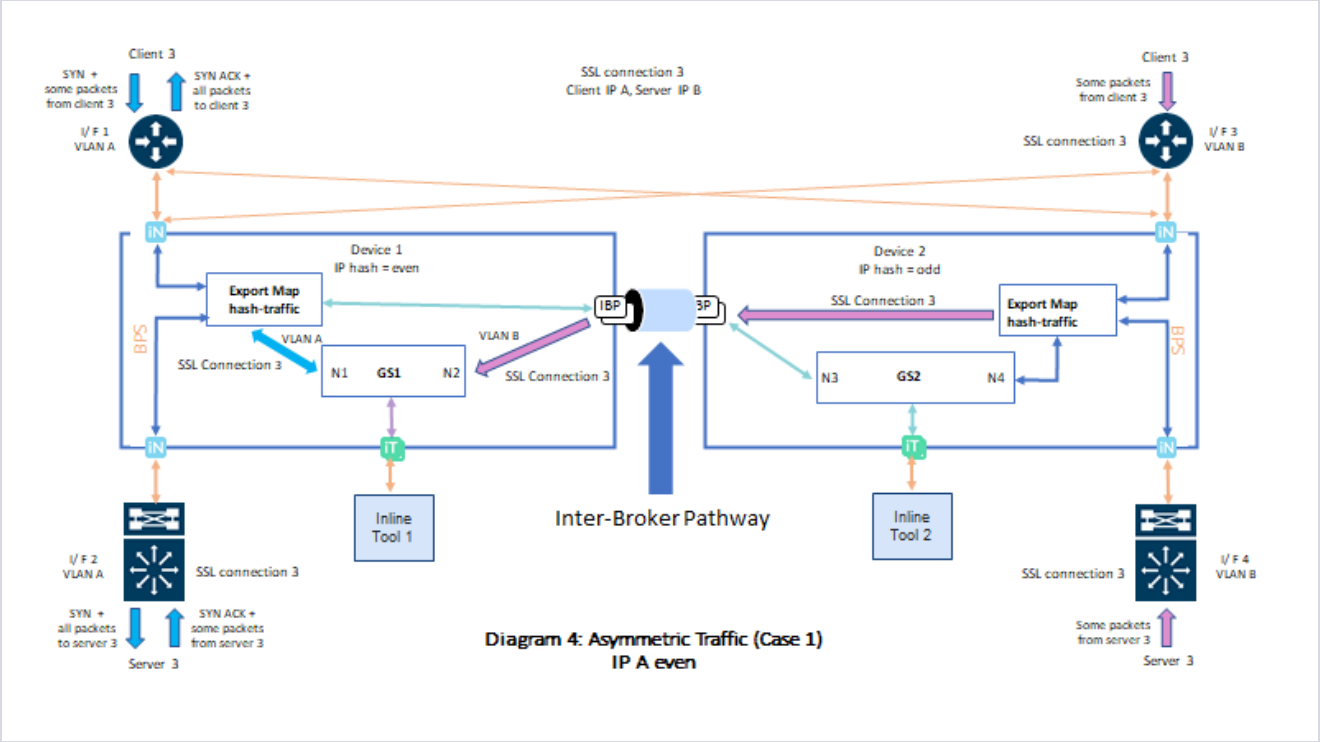


Diagram 4: Asymmetric Traffic (Case 1) IP A even

Scenario 2

Consider the following: the SYN packet for the TLS/SSL connection from client 5 arrives on interface 1 (I/F 1) with VLAN A. It is forwarded to GS1 in node 1, as the source IP (IP A) is even. GS1 initiates a TCP connection to Server 5 by sending a SYN packet from interface 2 (I/F 2) with VLAN A. Server 5 replies with a SYN ACK from interface 4 (I/F 4) with VLAN B on node 2, forwarded to GS1 through the IBP since the destination IP (IP A) is even.

Upon receiving the SYN ACK from Server 5 on I/F 4 with VLAN B in node 2, GS1 sends a SYN ACK to client 5 on interface 3 (I/F 3) with VLAN B, also on node 2. Due to asymmetric routing or load balancing, some incoming traffic from client 5 may reach I/F 3 with VLAN B on device 2, forwarded to GS1 via the IBP, as the source IP (IP A) remains even.

Additionally, some traffic from Server 5 may arrive at I/F 2 with VLAN A on node 1, also forwarded to GS1 on node 1 since the destination IP (IP A) is even. Outgoing traffic from GS1 to client 5 will go through I/F 3 with VLAN B on node 2, while traffic to Server 5 will route through I/F 2 with VLAN A on node 1, regardless of the incoming interface.

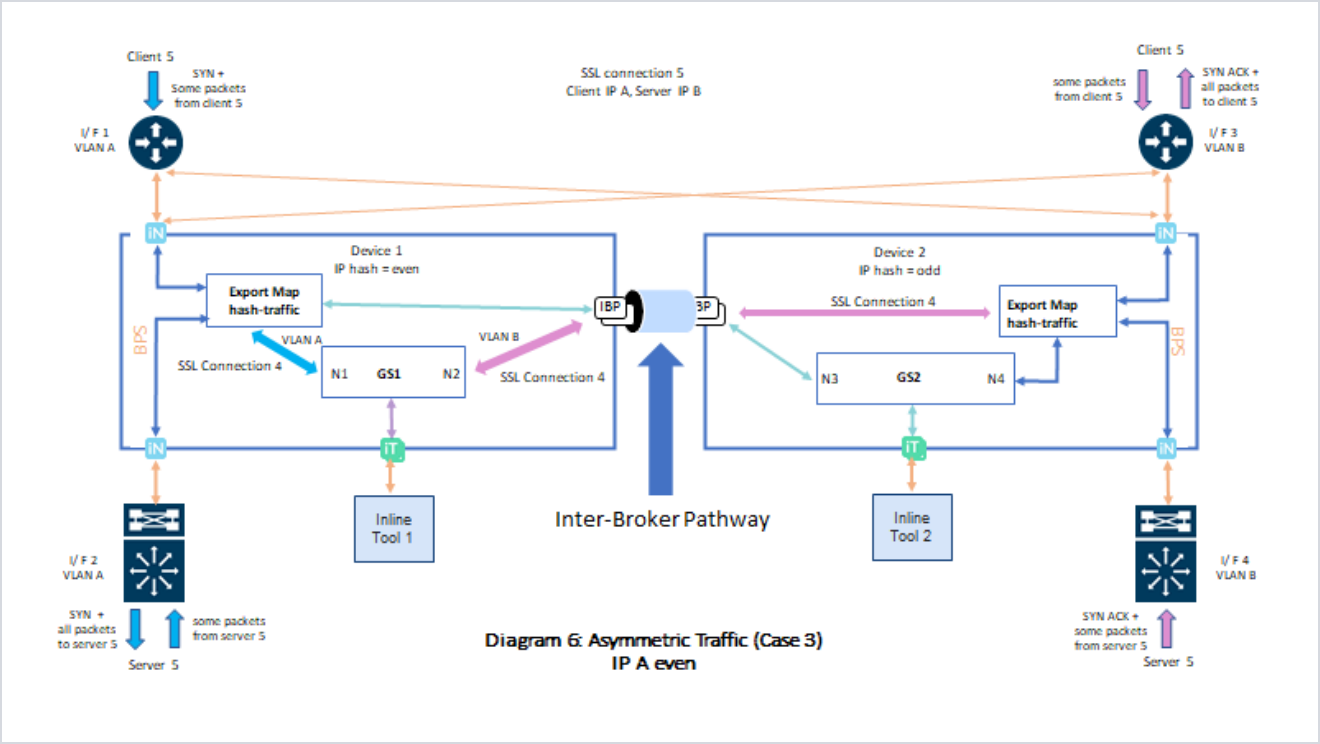


Diagram 6: Asymmetric Traffic (Case 3)
IP A even

VLAN Tagging Behavior for Decrypted Traffic

The Inline Tool positioned outside the SSL App is capable of receiving the original Map's VLAN tag for local node traffic, as well as the import map's VLAN tag for remote node traffic. In cases of asymmetric traffic, the Inline Tool will receive both the original Map's VLAN and the import map's VLAN.

Conversely, the Inline Tool located within the SSL App will consistently receive the tool tag configured in the SSL App for both local and remote node traffic.

VLAN Tagging Behavior for Non-Decrypted Traffic

The Inline Tool positioned outside the SSL App can capture the VLAN tag from the original Map or the non-proxy Map for local node traffic, as well as the VLAN tag from the import Map or import non-proxy Map for remote node traffic. In scenarios involving asymmetric traffic, the Inline Tool will receive both the VLAN tags from the original Map and the import map.

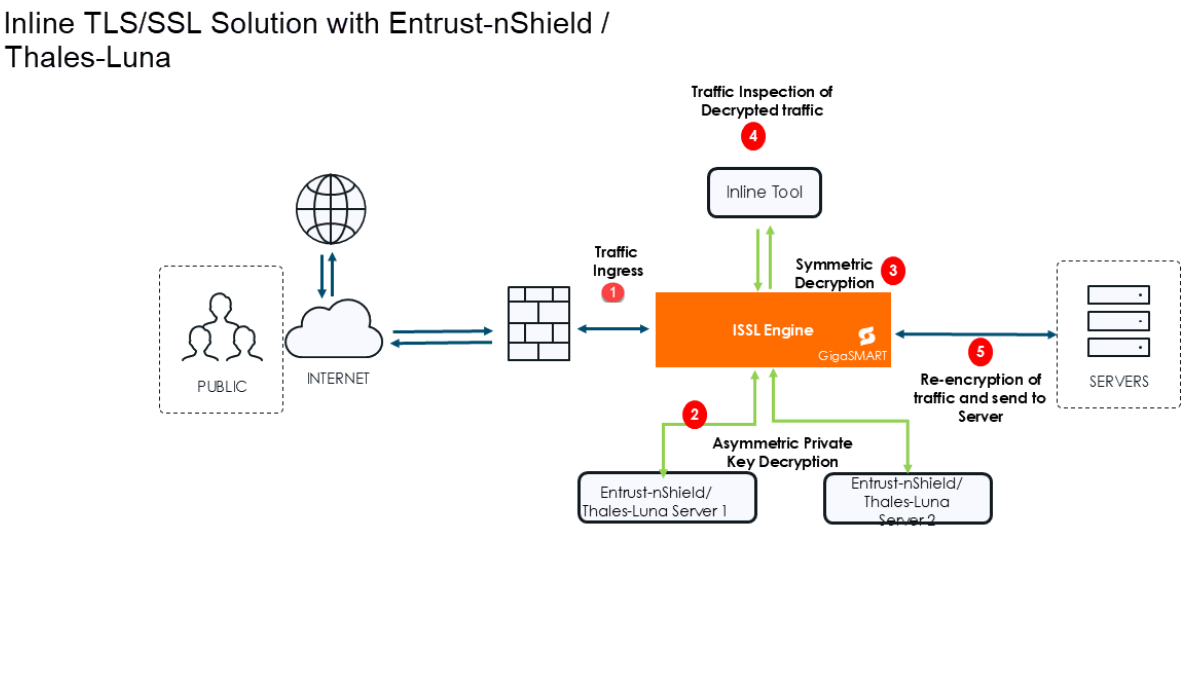
Conversely, the Inline Tool located within the SSL App will obtain the non-proxy Map's Tool tag for local node traffic and the import non-proxy Map's Tool tag for remote node traffic.

Entrust nShield and Thales-Luna HSM for TLS/SSL Decryption for iSSL

This feature enables inline TLS/SSL decryption to work with Hardware Security Modules (HSMs), enhancing the security and flexibility of cryptographic operations. HSM's are designed to:

- Securely perform cryptographic processing.
- Generate and store cryptographic keys.
- Prevent unauthorized access to sensitive data.

ISSL with Entrust nShield/Thales- Luna HSM



The following steps describe the end-to-end flow:

1. **Traffic Ingress**-Traffic enters the system through the Inline Network Port (Ingress). This port acts as the first point of entry for all inbound encrypted traffic.
2. **Asymmetric Decryption**- TLS handshake traffic which needs private key operation directs to HSM server.

3. **Symmetric Decryption**- Post private key operation done on HSM, data traffic decrypts in GigaSMART.
4. **Traffic Inspection**- The decrypted traffic is forwarded to inline security tools . These tools inspect but do not modify packet contents. Tools may allow, alert on, or block traffic based on inspection results, but they do not change the actual data.
5. **SSL Re-encryption and sent to Server** -The GigaSMART engine re-encrypts the traffic using the TLS/SSL parameters. The re-encrypted traffic is sent out to the servers.

HSM - Supported Solutions

Inline TLS/SSL decryption supports the following HSM solutions:

- Entrust nShield HSM
- Thales Luna Network HSM

HSM Configuration Rules, Notes, and Limitations

- Network & Deployment Requirements
  - Use a static IP address for the GigaSMART engine, and ensure it is registered on the HSM server.
  - DHCP is not recommended. If the IP address changes, the entire HSM setup must be deleted and reconfigured.
  - IPv6 is not supported for Thales Luna HSM server configuration when IPv6 stack port support is enabled. IPv6 traffic can be decrypted but HSM server IP and GigaSMART engine IP configuration will not support IPv6 address.
- HSM Type Usage
  - A deployment can use only one HSM type—either Entrust nShield or Thales Luna.
  - Entrust nShield and Thales Luna HSM cannot operate together within the same HSM Group or the same GigaVUE HC Series device.
- Switching Between HSM and Non-HSM Configurations: When moving from non-HSM to HSM, or between different HSM types, the GigaSMART engines and the device must be restarted before redeployment.
- Keys, Keychain, and Certificates
  - The Keychain password must be set before importing any keys. If not set, key and certificate uploads will fail.
  - Key–certificate mismatch checks cannot be performed for RSA or ECDSA keys because private keys remain on the HSM server.
  - For inbound deployments, required CA certificates must be uploaded to trust store avoid configuration errors.
  - For outbound deployments, CA certificates are optional but recommended for improved validation.
- Unsupported Deployment Combinations

    - Mixed HSM types (Entrust nShield + Thales Luna) are not supported, even if configured on different Gen3 GigaSMART cards.
    - Mixed HSM and non-HSM deployments cannot exist on the same device, including across separate Gen3 cards.
    - Mixed generation configurations are not supported:
    - Gen2 non-HSM iSSL cannot be configured with Gen3 HSM iSSL, even on different cards within the same device.
- Thales Luna-Specific Limitations
    - Thales Luna Network HSM does not support cluster, standby, or non-HA modes.

Inline TLS/SSL Decryption Solution with ICAP Client

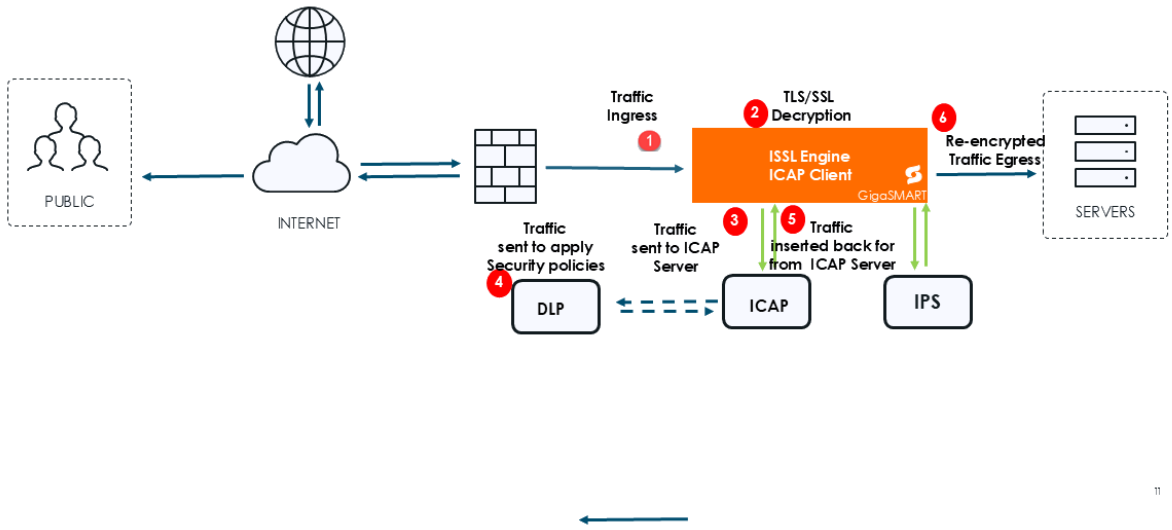
ICAP (Internet Content Adaptation Protocol) is used to inspect and modify HTTP messages by offloading them to a separate server. It allows clients to send HTTP requests and responses to an ICAP server, which can process the content and return a modified response. The message format follows HTTP/1.1, with components like the request line, headers, and body.

In a standard Inline SSL setup, the GigaVUE node decrypts SSL traffic and forwards it to an inline tool for processing. However, this setup does not support integration with DLP (Data Loss Prevention) servers running over ICAP.

By combining Inline SSL with ICAP, the decrypted traffic can now be routed through an ICAP client within the GigaSMART application. This enhancement enables direct integration with DLP ICAP servers, allowing advanced inspection and policy enforcement on decrypted content.

The following image provides an end -to end flow of Inline TLS/SSL Configuration with an ICAP Client

Inline TLS/SSL Solution with an ICAP Client



The following steps describe the end-to-end flow:

1. **Traffic Ingress**-Traffic enters the system through the Inline Network Port (Ingress). This port acts as the first point of entry for all inbound encrypted traffic.
2. **TLS/SSL Decryption**-The traffic is directed to the GigaSMART engine, where the inline TLS/SSL application is. The engine decrypts the TLS/SSL traffic, making it readable for inspection tools.
3. **Traffci sent to ICAP Client**- the decrypted traffic can now be routed through an ICAP client within the GigaSMART application.
4. **Traffic Inspection**- The traffic is now directed to DLP ICAP servers, allowing advanced inspection and policy enforcement on decrypted content.
5. **Traffic sent back for re-encryption** - The traffic is now sent from ICAP server to the GigaSMART engine.
6. **SSL Re-encryption and sent to Server** - The re-encrypted traffic is sent out to the servers.

ICAP - Limitations

- The ICAP Client cannot be deployed inline with the following iSSL configurations:
  - One-Arm iSSL
  - L3 NAT iSSL
  - RIA iSSL

- In GigaVUE-HC1-Plus, the inline network ports must reside in the same hardware slot.
- ICAP is not supported in GigaVUE-HCT chassis.

Refer to Configure ICAP Client and Configure ICAP Client for Inline TLS/SSL Decryption Solution for configuration details.

TLS/SSL Terminology and Acronyms

The below table provides definitions of TLS/SSL terminology:

Table 9: TLS/SSL Terminology

Term	Definition
Plaintext	The original, unencrypted data.
Ciphertext	The encrypted data.
Cryptography	The practice of secure communications.
Encryption	The process of turning plaintext into ciphertext.
Decryption	The process of turning ciphertext into plaintext.
Encryption algorithm	The algorithm used to perform encryption and decryption. It is also called the cipher.
Encryption key	The key used for encryption.
Decryption key	The key used for decryption.
Symmetrical encryption algorithm	The algorithm used for encryption in which the encryption key and the decryption key are identical.
Asymmetrical encryption algorithm	The algorithm used for encryption in which the encryption key and the decryption key are different.
Public key	The key used for encryption.
Private key	The key used for decryption.

The below table lists TLS/SSL acronyms:

Table 10: TLS/SSL Acronyms

Acronym	Definition
AES	Advanced Encryption Standard
CA	Certificate Authority
CBC	Cipher Block Chaining
CN	Common Name

Acronym	Definition
CRL	Certificate Revocation List
DES	Data Encryption Standard
DH, D-H	Diffie-Hellman
DSA	Digital Signature Algorithm
ECDH	Elliptic Curve Diffie-Hellman
ECDSA	Elliptic Curve Digital Signature Algorithm
FQDN	Fully Qualified Domain Name
IDS	Intrusion Detection System
IPS	Intrusion Prevention System
MAC	Message Authentication Code
MD	Message Digest
MitM	Man-in-the-Middle
OCSP	Online Certificate Status Protocol
OoB	Out-of-Band
PEM	Privacy Enhanced Mail
PFS	Perfect Forward Secrecy
PKCS12	Public Key Cryptography Standard #12
PKI	Public Key Infrastructure
RSA	Rivest-Shamir-Adleman
SHA	Secure Hash Algorithm
SSL	Secure Sockets Layer
TCP	Transmission Control Protocol
TLS	Transport Layer Security
URL	Uniform Resource Locator

Inline Solutions Configuration Using Flexible Inline Arrangement Canvas

This sections elaborates on how to configure Inline Solutions using Flexible Inline Arrangement Canvas.

What is a Flexible Inline Arrangement

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 Configurations using Inline Bypass Solutions (Classic).

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.

The figure below illustrates a complex inspection scenario that is enabled by flexible inline arrangements.

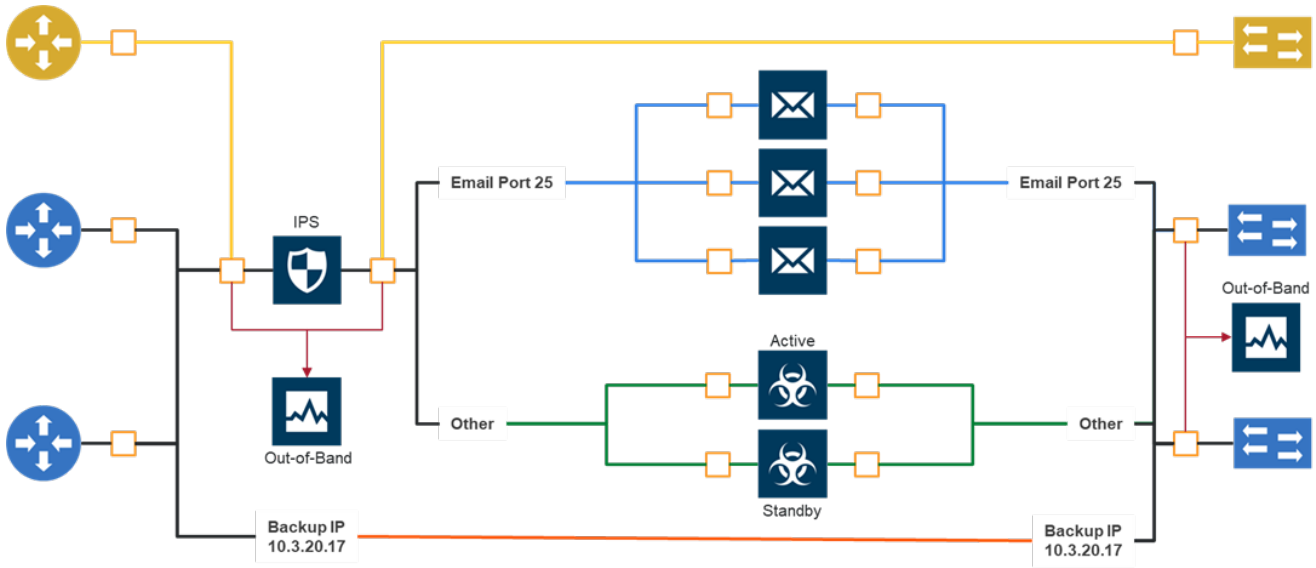


Figure 69 *Flexible Inline Arrangements Scenario*

In this example, on the left, there are three network links, all of which share the IPS. The details are as follows:

- At the top of the figure, the yellow line represent a flow of network traffic that only needs IPS inspection.
- In the middle of the figure, the Email and Other traffic represents network traffic flows that go through IPS inspection first, and then the Email traffic goes to dedicated inspection tools, here shown as three tools in an inline tool group, and the Other traffic goes to a threat protection active/standby pair.
- At the bottom of the figure, daily Backups of already-inspected traffic goes to bypass.

Although not shown graphically in Figure 69Flexible 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 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.

Supported Platforms, License, and Software Version

The below table provides you the software version, supported platforms and license that might be required to utilize flexible inline arrangements

Software Version	Supported Platforms	License	
		Monthly Subscription (Term License)	Perpetual Subscription
GigaVUE-FM and GigaVUE-OS running on software version 5.3.xx or higher support the flexible inline arrangement functionality.	GigaVUE-HC1-Plus	NA	NA
	GigaVUE-HCT	NA	NA
	GigaVUE-HC1	NA	NA
	GigaVUE-HC3 (CCv1 and CCv2)	NA	NA
	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

Flexible Inline Solution Supported in Clustered Nodes

The GGigaVUE 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	What is a Flexible Inline Arrangement
Single VLAN Tag	Configure Inline Single VLAN 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 LAG.
Hashing	Refer to Asymmetrical Hashing options in Configure Inline Tool Group.

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.
- FlexInline and FlexInline SSL App solutions don't support TAGS, so RBAC tagging is not supported.

Flexible Inline TLS/SSL Decryption Solution—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 (GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA25E, GigaVUE-TA25, GigaVUE-TA400, GigaVUE-TA400E)	- Physical Bypass (no BPS card) - Flexible and Resilient Inline SSL (no GigaSMART card) - GRIP (no BPS card) - ICAP (no GigaSMART card) - Classic Inline Bypass
VLAN Tagging and OOB Copy Limits (GigaVUE-TA25, GigaVUE-TA25E, GigaVUE-HC1-Plus)	Flexible Inline Single VLAN Tag with monitoring mode may send incorrect VLAN tags. OOB copy packets may also have wrong tags. You cannot use BYPASS WITH MONITORING with MONITORING mode on the tool. OOB copy from inline network is not allowed in this mode.
Inline Map Limits — Bidirectional	GigaVUE-TA25,GigaVUE-TA25E, GigaVUE-HC1-Plus → 126 maps GigaVUE-HC1,GigaVUE-HC3 (CCv1 & CCv2), GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA400, GigaVUE-TA400E → 256 or 512 maps depending on setup.
Inline Map Limits — Unidirectional	GigaVUE-TA25,GigaVUE-TA25E, GigaVUE-HC1-Plus→ 252 mapsGigaVUE-HC1,GigaVUE-HC3 (CCv1 & CCv2),), GigaVUE-TA200, GigaVUE-TA200E, GigaVUE-TA400, GigaVUE-TA400E → 512 or 1024 maps depending on setup.
Flexible Inline SSL Limits	Not supported with Inline Network LAG. Setting inline tools to “Drop” in the chain does not block Inline SSL traffic.
Filtering Limits(GigaVUE-TA400, GigaVUE-TA400E)	VLAN-based filtering in the Egress Port Filter for OOB copies is not supported. If one tool in the map is in monitoring mode, all tools must use the same mode. Asymmetric hashing (a-srcip-b-dstip and b-srcip-a-dstip) is not supported.
Protocol Pass-Through(GigaVUE-TA400, GigaVUE-TA400E)	CDP pass-through is not supported when the source is an Inline Network LAG. Bypass for LACP, CDP, and LLDP is supported.
Scaling Limits — GigaVUE-TA400, GigaVUE-TA400E	Max Inline Networks and Tools: 48 Max Inline Network LAG list: 24 Max Inline tools or tool groups per direction: 16 Max OOB copy entries per direction: 17 Max OOB copy ports per entry: 128

Configure Software Constructs Using Flexible Inline Arrangement Canvas

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

Refer to the following sections for details:

- Configure Inline Network Ports and an Inline Network
- Configure IP Interface
- Configure Inline Network LAG
- Configure Inline Network Bundle
- Configure Inline Tool Ports and Inline Tools
- Configure Inline Tool Group
- Configure Inline Single VLAN Tag
- Configure Flexible Inline Maps

Configure Inline Network Ports and an Inline Network

To configure inline network ports and an inline network:

1. Go to  **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed Flexible Inline Canvas, select the device where you want to configure the inline network.
3. Click the '+' icon next to the **Inline Network** to create a new entry.
4. Enter a name and description for the inline network in the **Alias** and **Description** fields.
5. Click **Port Editor**. In the **Quick Port Editor**, scroll to the inline network ports you want to configure. Select **Enable** to activate the network ports, and click **OK**.
6. Select the ports you want to configure as the inline network pair from the **Port A** and **Port B** drop-down lists.

7. Choose a traffic path from the **Traffic Path** drop-down list:
   - **Bypass**- Redirects all traffic from the inline network to the opposite-side inline network port, bypassing all inline tools and tool groups. Traffic is not decrypted during this process.
   - **Drop** - Drops all traffic originating from the inline network.
   - **Bypass with Monitoring** - Sends one copy of the traffic directly to the opposite-side inline network port and another copy to the sequence of inline tools and inline tool groups.

     > **NOTE:** If the Inline network is set to bypass with monitoring, GigaSMART Inline TLS/SSL decryption does not work.

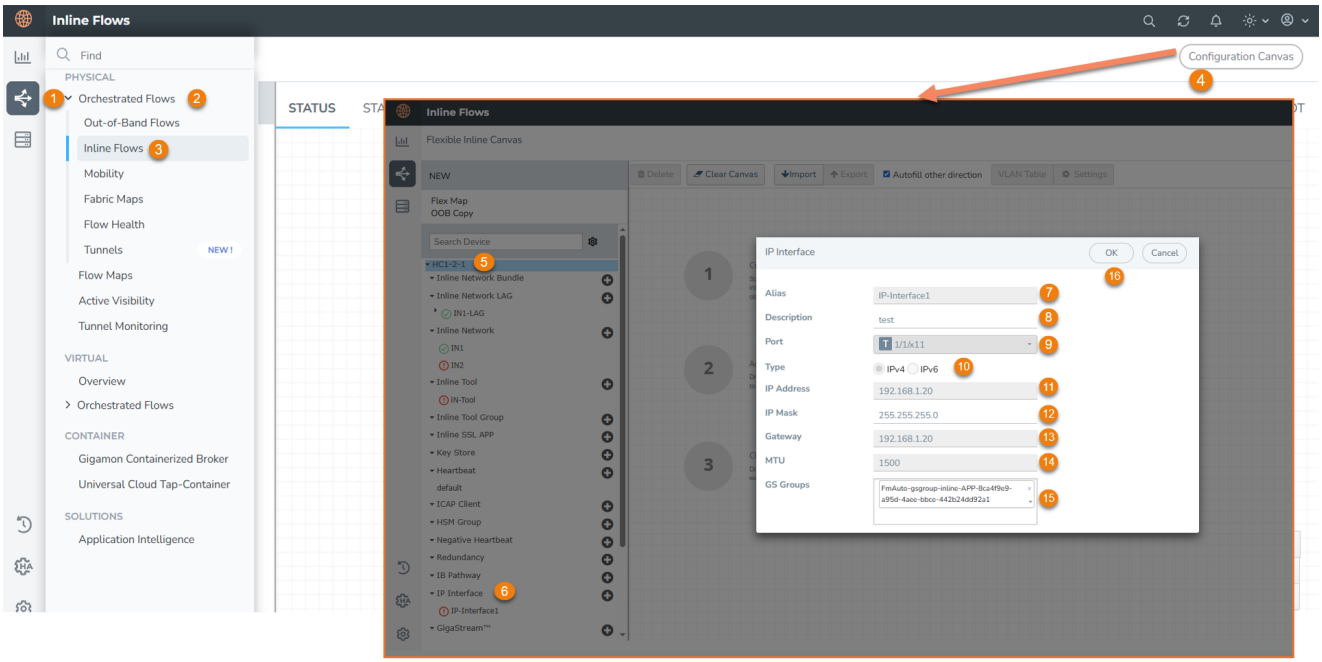
   - **To Inline Tool** - Directs all traffic through the configured sequence of inline tools and tool groups, based on their current status.
8. Select the **Link Failure Propagation** check box to propagate link failure from one side of the inline network to the other.
9. Select the **Accept Regular Heartbeat** check box to allow the inline network port pair to accept heartbeat packets from the inline tool port pair.
10. Click **OK** to save the configuration.

Configure IP Interface

To configure IP Interface:

1. Go to  **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed Flexible Inline Canvas, select the device where you want to configure the IP Interface.
3. Click the '+' icon next to the **IP Interface** to create a new entry.
4. Enter a name and description for the IP Interface in the **Alias** and **Description** fields.
5. From the **Port** drop down list, choose the port you want to configure as the IP Interface.
6. Select the **Type** of the IP Interface you want to configure.
7. Enter the **IP Address** (for example, 192.168.1.20).
8. Enter the **IP Mask** (for example, 255.255.255.0).
9. Enter the **Gateway** (for example, 192.168.1.20).

10. In the **MTU** field, specify the Maximum Transmission Unit for the port (for example, 1500).
11. From the GS Groups drop-down list, select the required GigaSMART Group you created.
12. Click **OK** to save the configuration.



Configure Inline Network LAG

Before You Begin - Configure the required inline network ports and inline networks before you set up an inline network LAG. Refer to Configure Inline Network Ports and an Inline Network for details.

To configure an inline network LAG:

1. Go to  **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed canvas, select the device where you want to configure the inline network LAG.
3. Click the '+' icon next to the **Inline Network LAG** to create a new entry.
4. Enter a name and description for the inline network in the **Alias** and **Description** fields.

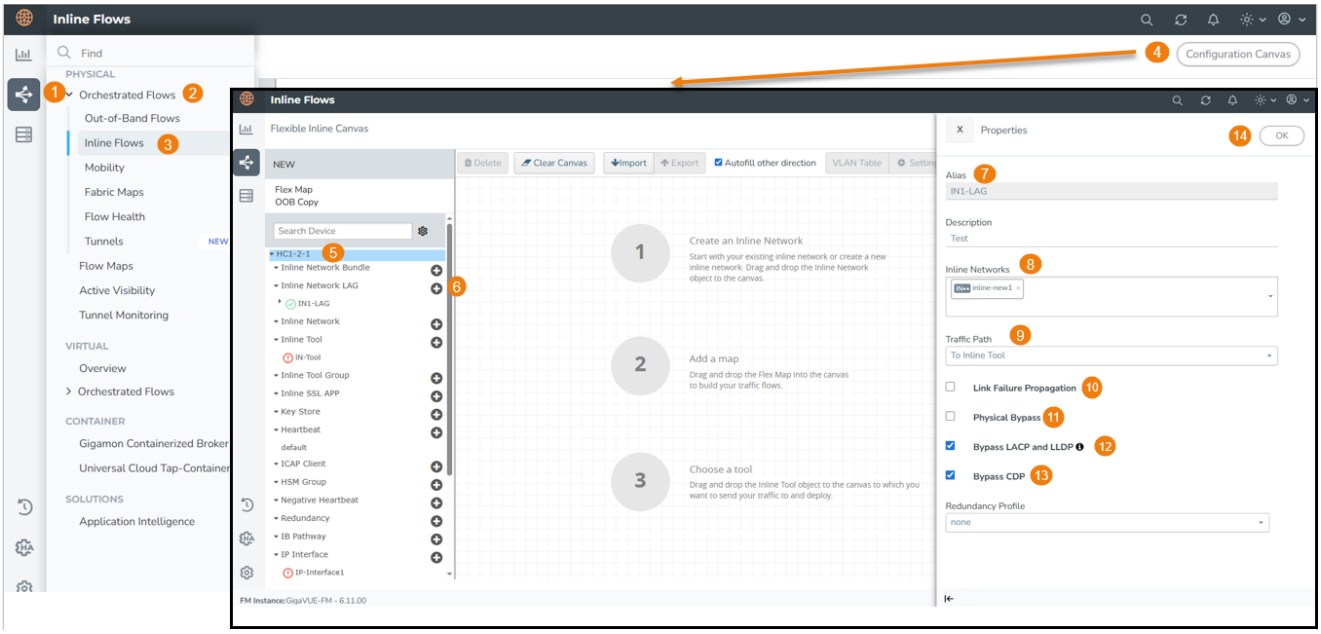
5. From the Inline Networks drop-down list, select the inline networks to include in the inline network LAG.
6. From the **Traffic Path** drop-down list, choose one of the following options:
   - **Bypass** - Forwards traffic directly between Port A and Port B of the inline network.
   - **Drop** - Drops all traffic arriving at the inline network ports.
   - **Bypass with Monitoring** - Forwards traffic as a forced bypass and sends a copy to inline tools. Ensure a traffic map is configured between the inline network and the inline tool.
   - **To Inline Tool** - Forwards traffic to the sequence of inline tools.
7. Select the **Link Failure Propagation** check box to bring down a port when its paired port goes down..
8. Select the **Physical Bypass** check box to allow traffic to flow directly between Port A and Port B of the inline network pair when a device or module powers down.
9. If the port channel uses **LACP**, select the **Bypass Link Aggregation Control Protocol** and **Link Layer Discovery Protocol** check box to maintain port channel functionality on links connected to inline network LAG ports.

   Notes:
   - Inline Network LAG needs a bypass map to handle LACP bypass.
   - When the Bypass LACP and LLDP checkbox is enabled, all protocol packets with MAC Destination **01-80-C2-XX-XX-XX** are bypassed.

10. If the port channel supports CDP, select the **Bypass CDP** (Cisco Discovery Protocol) check box to maintain CDP discovery functionality on the links that are connected to inline network LAG ports.

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

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



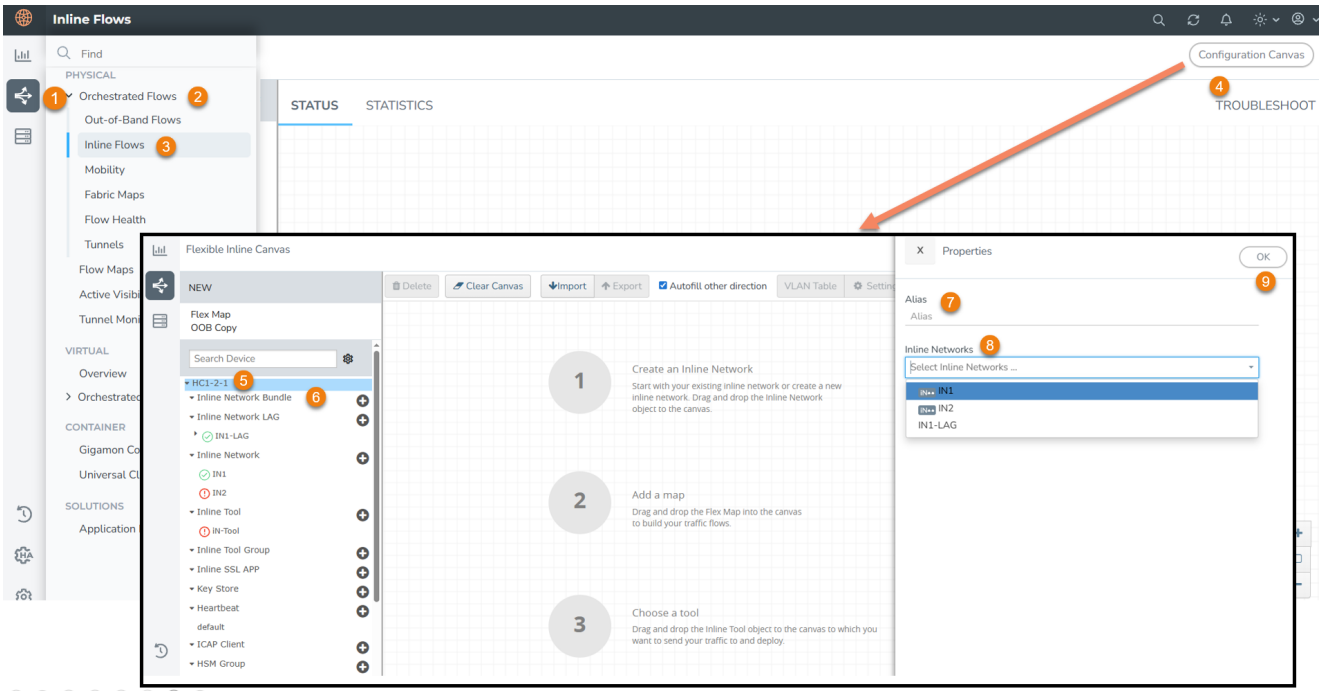
Configure Inline Network Bundle

Before You Begin - Configure the required inline network ports and inline networks before you set up an inline network bundle. Refer to Configure Inline Network Ports and an Inline Network for details.

To configure an inline network bundle:

1. Go to  **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed canvas, select the device where you want to configure the inline network bundle.
3. Click the '+' icon next to the **Inline Network Bundle** to create a new entry.
4. Enter a name for the inline network bundle in the **Alias** field.
5. From the **Inline Networks** drop-down list, select the inline networks you want to include in the bundle. Refer to Configure Inline Network Ports and an Inline Network.
6. Click **OK** to save the bundle configuration.
7. Drag and drop the inline network bundle into the canvas, and configure the required inline map.
8. Enter the **Tool Side VLAN Tag** for each inline network added to the bundle.

9. Select the **TPID** for the Tool Side VLAN Tag.
    - The default TPID is 0x8100.
    - You can also select 0x9100 or 0x88a8 from the drop-down list.
10. Add the required rules for the inline map, and click **OK** to save the configuration.



Configure Inline Tool Ports and Inline Tools

To configure the inline tool ports and inline tools:

1. Go to  **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed canvas, select the device where you want to configure the inline tool.
3. Click the '+' icon next to the **Inline Tool** to create a new entry.
4. In the **Properties** pane, enter a name and description in the **Alias** and **Description** fields.
5. From the **Type** drop-down list, select one of the following options:
    - **External** - To configure a third-party tool.
    - **GigaVUE Node** - To configure a GigaVUE node as a tool.

6. Click **Port Editor**, and in the **Quick Port Editor**, scroll to the inline tool ports you want to configure. Select **Enable** to activate the ports, and click **OK**.
7. From the **Port A** and **Port B** drop-down lists, select the inline tool ports based on the direction the inline tool expects traffic from the network.
8. Verify that the **Enabled** check box is selected.
9. From the **Failover action** drop-down list, choose one of the following options:
   - **Tool Bypass** - Redirects traffic to the next inline tool or inline network port when a failover occurs.
   - **Network Bypass** - Puts inline networks in bypass mode (traffic flows from side A to side B and vice versa).
   - **Tool Drop** - Drops traffic and redirects it to a dummy VLAN with no members.

     > **NOTE:** If the failover action is set to 'drop' for an inline tool in the Flex iSSL solution, GigaSMART drops all traffic at the vport.

   - **Network Drop** - Drops all traffic entering from side A and side B.
   - **Network Port Forced Down** - Brings down inline network ports.
10. Select the **Recovery Mode** to either **manual** or **automatic**.
11. To define additional tags on the tool side, select the **Enable** check box for the **Inline tool Sharing mode**.

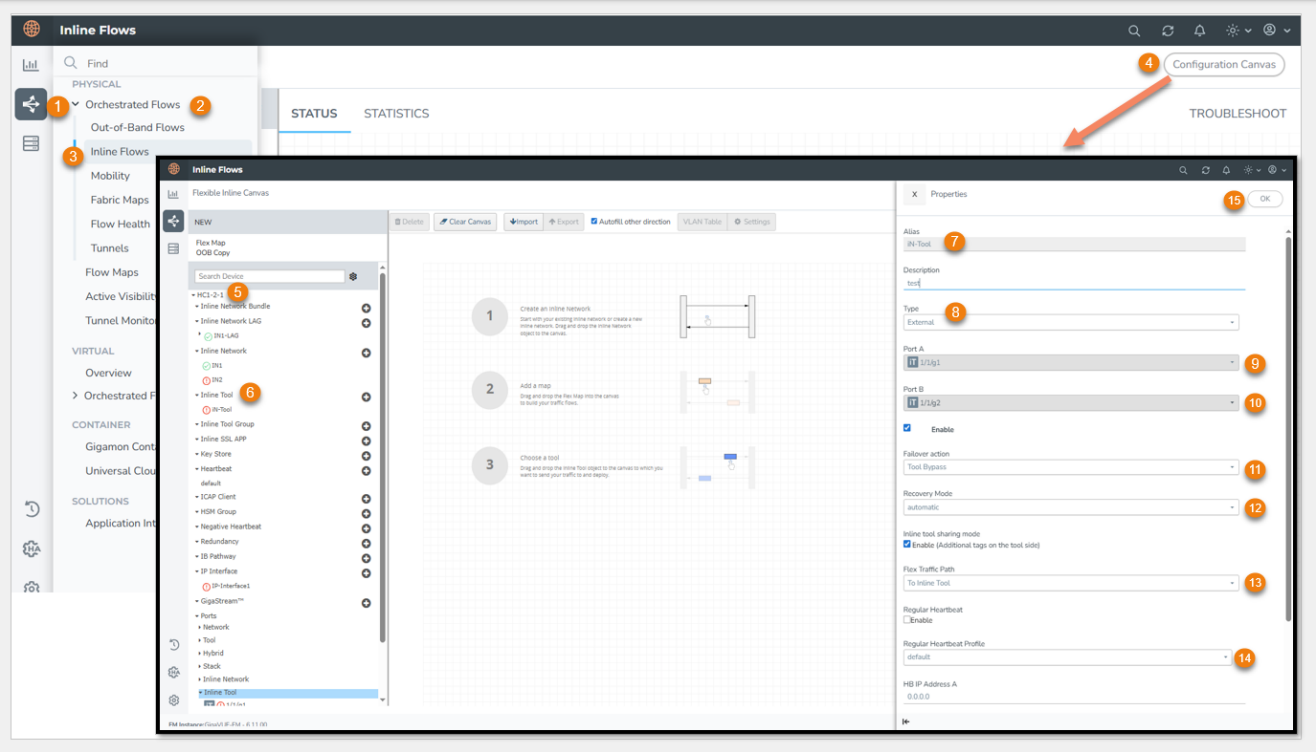
    > **NOTE:** If you disable **Inline Tool Sharing Mode**, you can use the inline tool only in one flexible inline map.

12. From the **Flex Traffic Path** drop-down list, select one of the following:
    - **Drop**- Drops traffic at the inline tool.

      > **NOTE:** If the failover action is set to 'drop' for an inline tool in the Flex iSSL solution, GigaSMART drops all traffic at the vport.

    - **Bypass** - Redirects all traffic from the inline network to the opposite-side inline network port, bypassing all inline tools and tool groups. Traffic is not decrypted during this process.
    - **Bypass with Monitoring** - Sends traffic to the inline tool and absorbs it, while forwarding a copy to the next inline tool in the sequence. Also absorbs return traffic from side B.
    - **To Inline Tool** - Directs all traffic through the configured sequence of inline tools and tool groups based on their current status, and decrypts it.
13. If required, select the **Enable** check box for **Regular Heartbeat**, and choose a suitable profile from the **Regular Heartbeat Profile** drop-down list.
14. Enter the IP addresses for side A and side B in the HB IP Address A and HB IP Address B fields, as defined in the heartbeat profile.

15. If required, select the **Enable** check box for **Negative Heartbeat**, and choose a suitable profile from the **Negative Heartbeat Profile** drop-down list.
16. Click **OK** to save the configuration.



Configure Inline Tool Group

Before You Begin - Configure the required inline tools before setting up an inline tool group. Refer to Configure Inline Tool Ports and Inline Tools.

To configure an inline tool group:

1. Go to  **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed canvas, select the device where you want to configure the inline tool group.
3. Click the '+' icon next to the **Inline Tool Group** to create a new entry.
4. In the **Properties** pane, enter a name and description in the **Alias** and **Description** fields.
5. From the **Inline Tools** drop-down list, select the required inline tools.
6. From the **Weighting** drop-down list, choose one of the following options:

- **Equal**—Distributes traffic equally across all inline tools in the inline tool group. This is the default option.
- **Relative**—Distributes traffic to the inline tools in the inline tool group based on a relative weight or ratio assigned to the respective inline tools. The valid range is 1–256.
- **Percentage**—Distributes traffic to the inline tools in the inline tool group based on assigned percentages. The valid range is 1–100.

> **NOTE:** If you select **Relative** or **Percentage**, enter the hash weights for the inline tools in the table below the **Weighting** drop-down list. Ensure that you assign a hash weight for each inline tool in the inline tool group. However, asymmetrical hashing is not supported for these options.

7. From the **Inline Spare Tool** drop-down list, select the inline tool that will receive traffic if the primary inline tool fails. To add a removed tool back to the spare list, delete it from the group and update it under the Inline Tool Group option in the canvas.

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

8. Select the **Enabled** check box to make the inline tool group available for deployment.
9. Select the **Release Spare if Possible** check box to allow the spare tool to return to standby when the primary tool recovers.
10. From the **Failover Action** drop-down list, Choose one of the following options:
    - **Tool Bypass** - Redirects traffic to the next inline tool or inline tool group when a failover occurs.
    - **Network Bypass** - Puts inline networks in bypass mode (traffic flows from side A to side B and vice versa).
    - **Tool Drop** - Drops traffic for the affected inline tool or inline tool group.
    - **Network Drop** - Drops all traffic entering from side A and side B.
    - **Network Port Forced Down** - Brings down inline network ports.
11. From the **Failover Mode** drop-down list, select **Spread** to redistribute the traffic coming from the inline network (or inline network group) to the active inline tools (excluding the failed inline tool or tools).

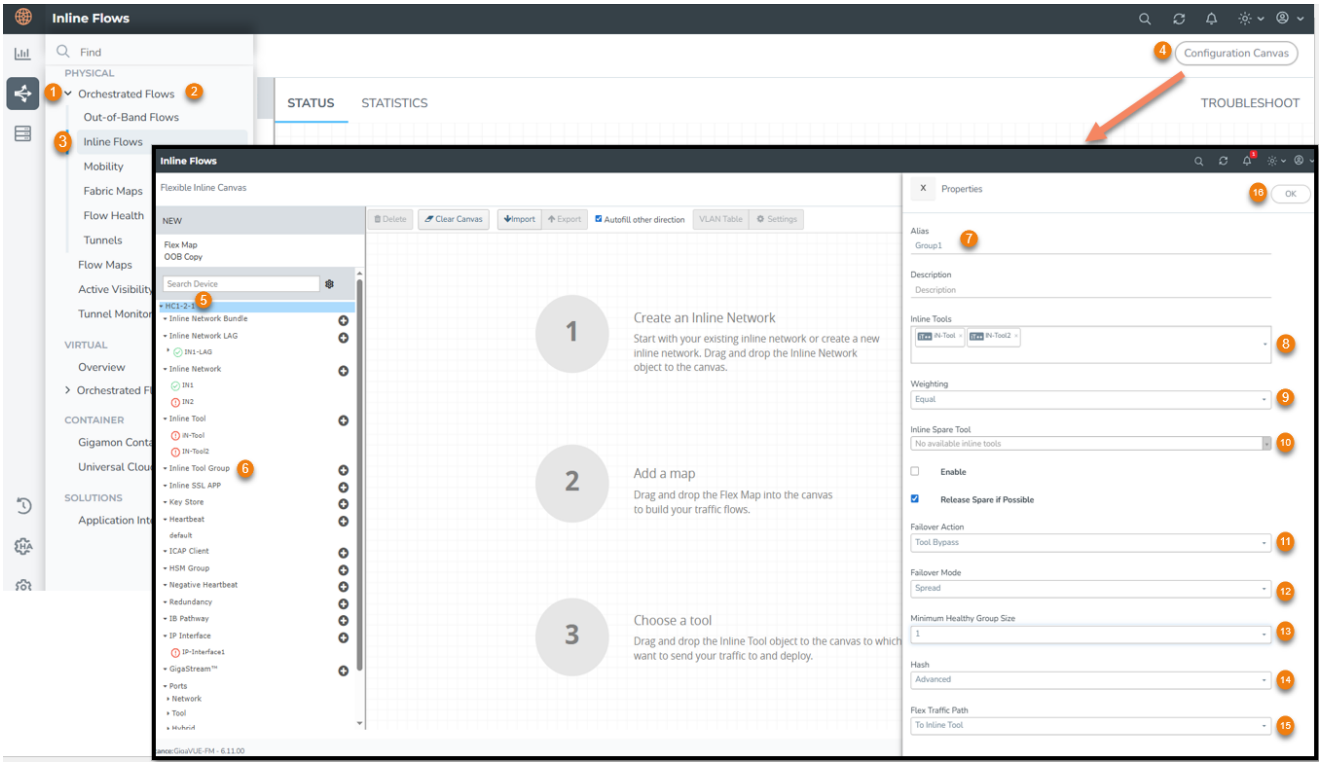
    > **NOTE:** This field is not applicable when there is only one inline tool in the tool list.

12. From the **Minimum Healthy Group Size** drop-down list, select the minimum number of inline tools (including the spare) required for the inline tool group to remain active.

13. From the **Hash** drop-down list, choose one of the following options to distribute packets across the inline tools in the group:
    - **Advanced**—Uses symmetrical hashing, derived from the combination of packet fields based on the selected advanced-hash algorithm criteria
      - The most common criteria are the combination of source IP and destination IP addresses.
      - This produces a hash value that ensures all traffic from the same session is sent to the same inline tool in the group
    - **SideA as sourceIP & SideB as destinationIP**—Uses asymmetrical hashing, derived from the source IP address on side A and the destination IP address on side B of the inline network.
      - This produces a hash value that ensures all traffic from the same source address on side A is sent to the same inline tool in the inline tool group, regardless of destination or session.
    - **SideB as sourceIP & SideA as destinationIP**—Uses asymmetrical hashing, derived from the destination IP address on side A and the source IP address on side B of the inline network.
      - This produces a hash value that ensures all traffic from the same source address on side B is sent to the same inline tool in the inline tool group, regardless of destination or session.

    > **NOTE:** This field is not applicable if you select the **Relative** or **Percentage** options in the Weighting drop-down list.

14. From the **Flex Traffic Path** drop-down list, select one of the following options:
    - **Drop**- Drops traffic at the inline tool group.
    - **Bypass** - Traffic bypasses the inline tool group. Use this option for performing maintenance on an inline tool group.
    - **Monitoring** - Sends traffic to the inline tool group and absorbs it, while forwarding a copy to the next inline tool group in the sequence. Also absorbs return traffic from side B.
    - **To Inline Tool** - Forwards traffic to the inline tool group.
15. Click **OK** to save the configuration.



Configure Inline Single VLAN Tag

During flexible Inline Bypass operations, network traffic sent to inline tools includes an extra VLAN tag. This tag helps identify traffic from inline tools and ensures correct routing to inline networks. However, many inline tools may not support extra VLAN tags. You can use Flexible Inline Single Tag to replace the extra VLAN tag on incoming traffic. This feature allows you to map incoming VLANs on the network side to outgoing VLANs on the tool side.

NOTE: The **OOB Copy** tag attribute **none** is invalid for single tag maps. Use the attribute **Original** or **as-inline** instead.

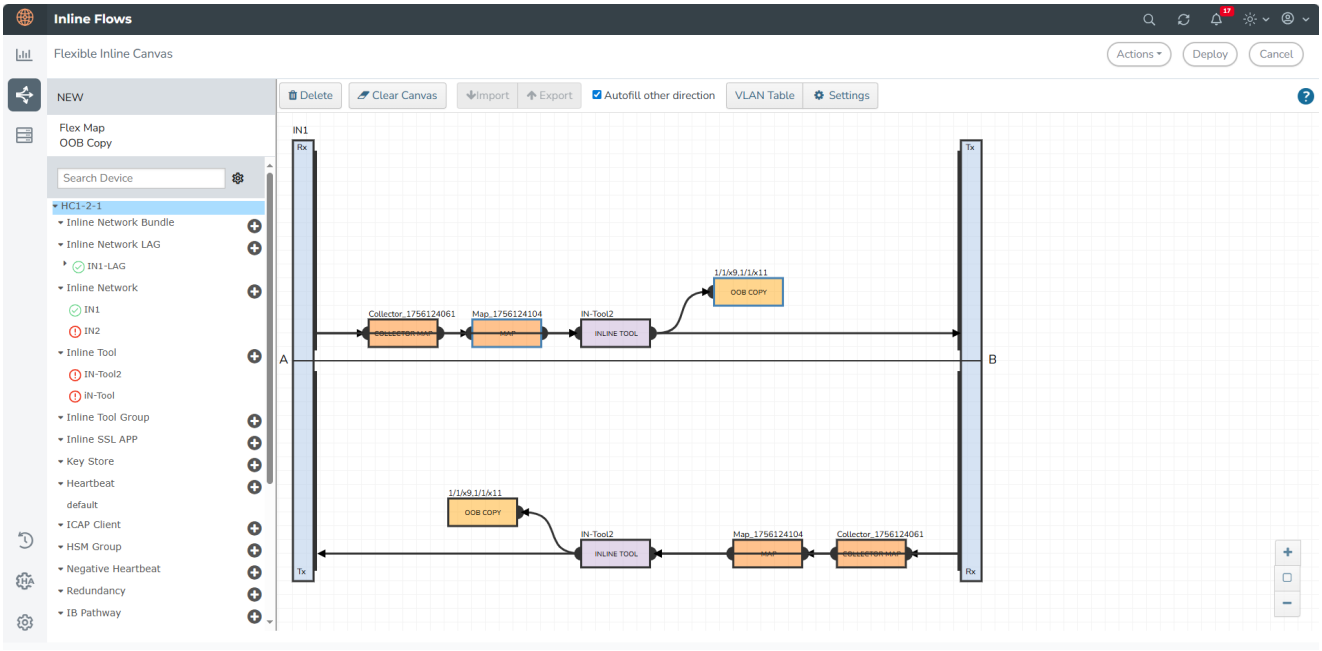
To configure an inline single tag:

1. Go to  **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed canvas, select the required inline network from the list of devices.
3. Drag and drop the inline network into the canvas.
4. Click **Settings** to open the **Settings** pane.

5. Select the **Enable** check box for **Show Single Tag Options**, and enter the VLANS expected on the inline network. Click **OK**.
6. Drag and drop a flexible inline map object into the canvas, and click the map to open the **Properties** pane.
7. Select the **Enable** check box for the **Single Tag Mode**, and enter the tool-side VLAN tags.
   a. Use the **Enable Network Side VLAN Tag** check box to bulk enable or disable VLAN tags for flexible inline maps.
   b. Select the **TPID** for the tool-side VLAN tag. The default value is 0x8100. You can also choose 0x9100 or 0x88a8 from the drop-down list. Click **OK**.

   Notes:
   - GigaVUE-FM automatically adds a VLAN qualifier to the rules. If you don't specify any rules, it adds a default rule with the VLAN qualifier.
   - You can enable or disable **Single Tag Mode** for collector maps if needed.
8. Drag and drop the required inline tools into the canvas.
9. Drag and drop the **OOB Copy** into the canvas, if required.
10. Click **Port Editor**, and in the **Quick Port Editor**, scroll to the hybrid or tool ports that you want to configure. Select **Enable** to activate the ports. Click **OK**.
11. From the **Destination Ports** drop-down list, select the required hybrid or tool ports to configure as destination ports. You can also select a hybrid or tool GigaStream. Refer to *"How to Use GigaStream" section* for details.
12. From the **VLAN Tag** drop-down list, select the required tags.
13. Click **OK** to save the configuration.

Notes:
- In GigaVUE-HC1-Plus, you cannot enable Single Tag Mode if the inline network's traffic path is in monitoring mode.
- If you enable the **Single VLAN Tag** option in a Flexible inline TLS/SSL Decryption solution, ensure that the **inline-ssl app profile** also has Single VLAN Tag configuration enabled.
- After deploying the solution, GigaVUE-FM preserves and displays any unused VLAN tags in the order they were configured.
- The VLAN Table window lists maps in the same order as shown in the configuration canvas.

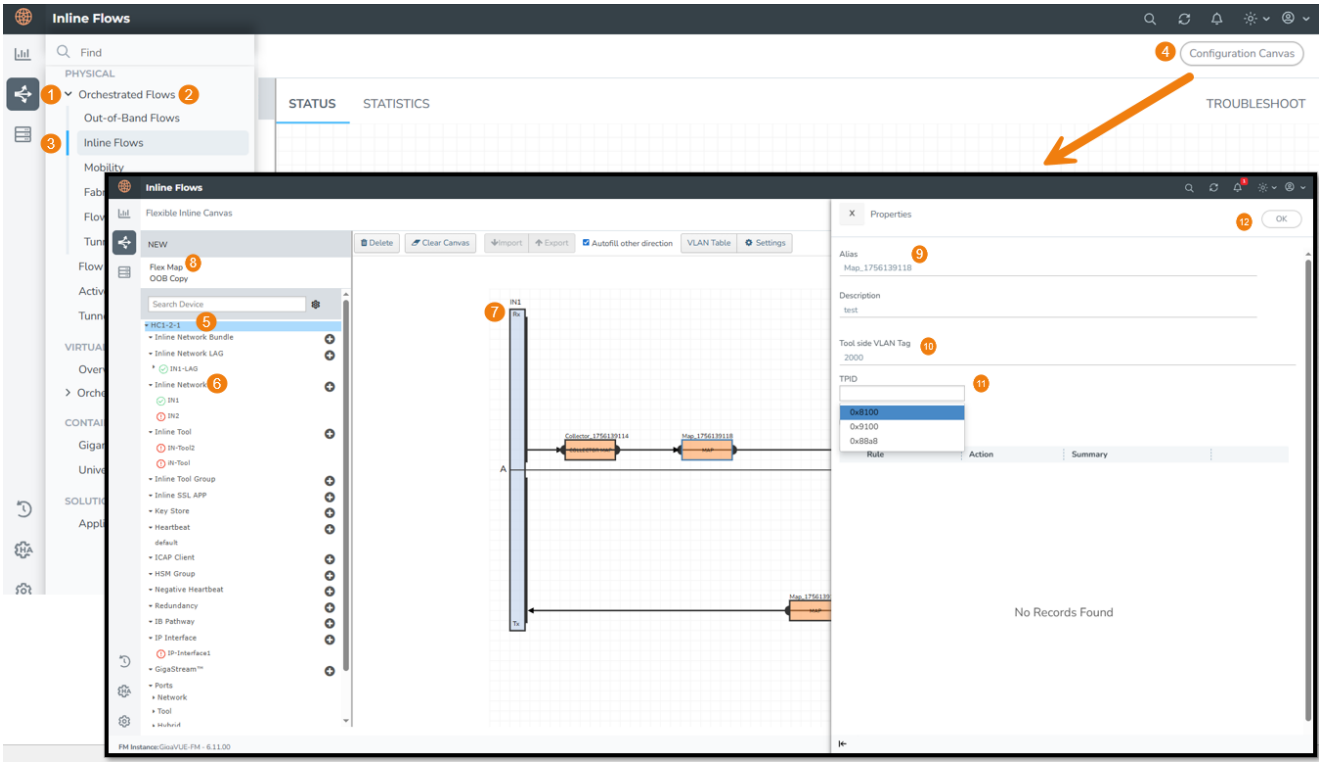


Configure Flexible Inline Maps

Before You Begin - Configure the required inline network before setting up Flexible Inline Maps. Refer to Configure Inline Network Ports and an Inline Network.

To configure Flexible Inline Maps:

1. Go to  **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed canvas, select the device where you want to configure the Flexible Inline Maps.
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, 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 TPID is 0x8100. The default TPID is 0x8100. You can also select 0x9100 or 0x88a8 from the drop-down list.
8. Add the required rules for the inline map. You can also import an existing map template..
9. Click **OK** to save the configuration.



Configure Resilient Inline Arrangement Solution

Before You Begin

1. A Resilient Inline Arrangement is performed on two GigaVUE devices .
2. Ensure that the names on both GigaVUE-FM,GigaVUE HC Series devices are identical, that is, the inline networks, inline tools, inline tool groups, out-of-band tools, and out-of-band tool GigaStreams must all have the same alias names on both the devices.
3. If you choose to use the inline network bundle, the alias of the inline network bundle on both the devices must be identical. However, the inline networks that are grouped into the bundle can have different aliases.

Once you are set with these basic prerequisites you can move forward to configure the following components:

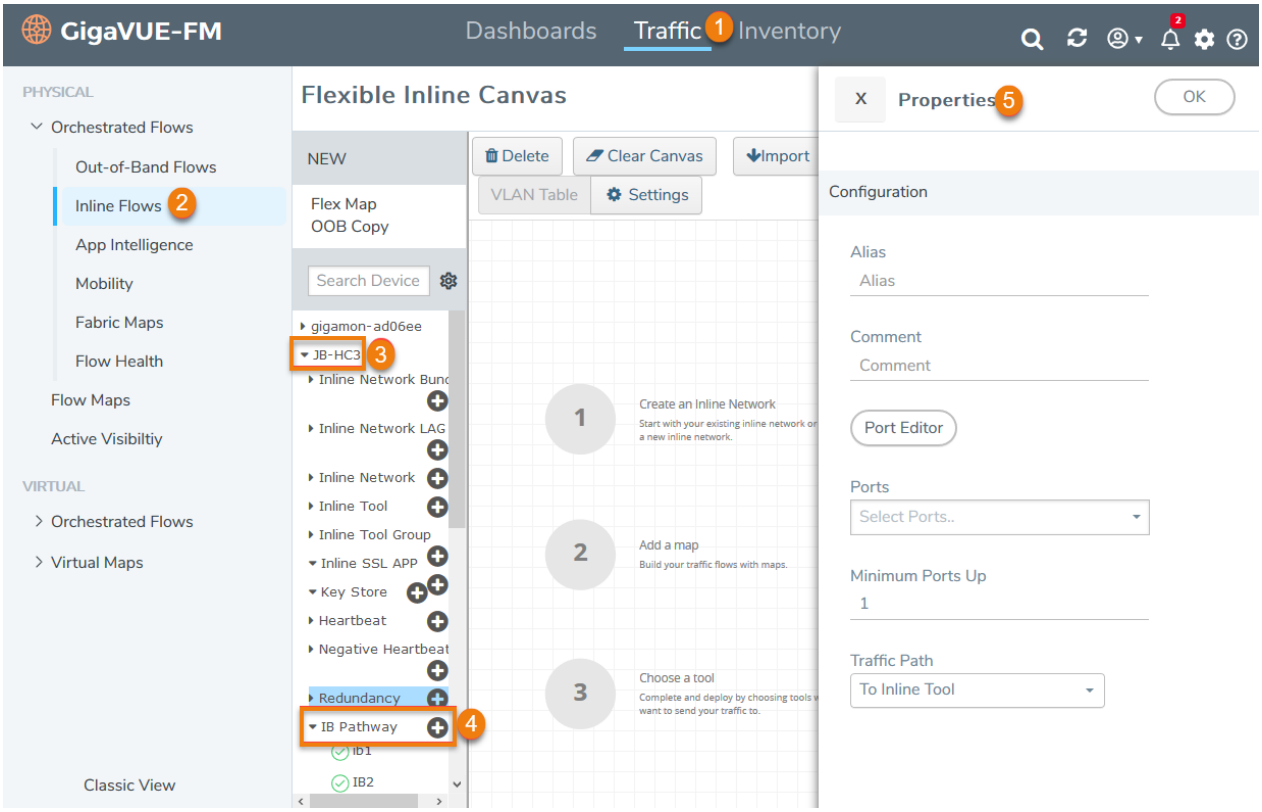
Create Inter-broker Pathway

To create a new inter-broker pathway:

1. Go to  **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed canvas, select the device where you want to create the inter-broker pathway.
3. Click the '**+**' icon next to the **IB Pathway** option to create a new inter-broker pathway.
4. In the **Properties** pane, enter a name and description in the **Alias** and **Description** fields..
5. From the **Ports** drop-down list, select the required tool ports to attach to the inter-broker pathway.

   > **NOTE:** If the required tool ports are not available, you can enable them administratively. Click **Port Editor**, scroll to the tool ports you want to configure in the **Quick Port Editor** page, select **Enable**, and click **OK**.

6. In the **Minimum Ports Up** field, enter the minimum number of tool ports that must be operational for the inter-broker pathway status to be "**Up**".
7. From the **Traffic Path** drop-down list, select one of the following options:
   - **Bypass**—Traffic bypasses the inter-broker pathway and is redirected to the next inline network port.
   - **Monitoring**—Traffic is forwarded to the sequence of inline tools in the monitoring mode.
   - **To Inline Tool**—Traffic is forwarded to the sequence of inline tools you have configured.
8. Click **OK** to save the configuration.



Configure Resilient Inline Arrangement

To configure a resilient inline arrangement:

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

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

3. Select the **Enable** check box next to **Show Resilient Inline Menu**.
4. Select the required **Node 1**, **Node 2**, **IB Pathway1**, and **IB Pathway2** for the resilient inline arrangement.
5. From the **Hashing Source** drop-down list, select one of the following options:
   - **Side A** - Hashing uses source IP or source port from Side A; destination IP or destination port from Side B.
   - **Side B** - Hashing uses source IP or source port from Side B; destination IP or destination port from Side A.
6. From the **Hashing Type** drop-down list, select one of the following options:
   - **L3 (IP Based)** - Hashing uses the IP address.

- **L4 (Port Based)** - Hashing uses the transport layer port number.

7. From the **Hashing LSB Node** drop-down list, select one of the following options:
   - **Node 1 as 0** - Traffic from IPs ending in 0 is hashed to Node 2.
   - **Node 2 as 0** - Traffic from IPs ending in 0 is hashed to Node 1.

   NOTE: This option is available only if you selected **L3 (IP Based)** in the Hashing Type field.
8. From the **Hashing Port** drop-down list, select one of the following options:
   - **Node 1 as odd** - Traffic with odd port numbers is hashed to Node 2, while traffic with even port numbers is hashed to Node 1.
   - **Node 2 as odd** - Traffic with odd port numbers is hashed to Node 1, while traffic with even port numbers is hashed to Node 2.

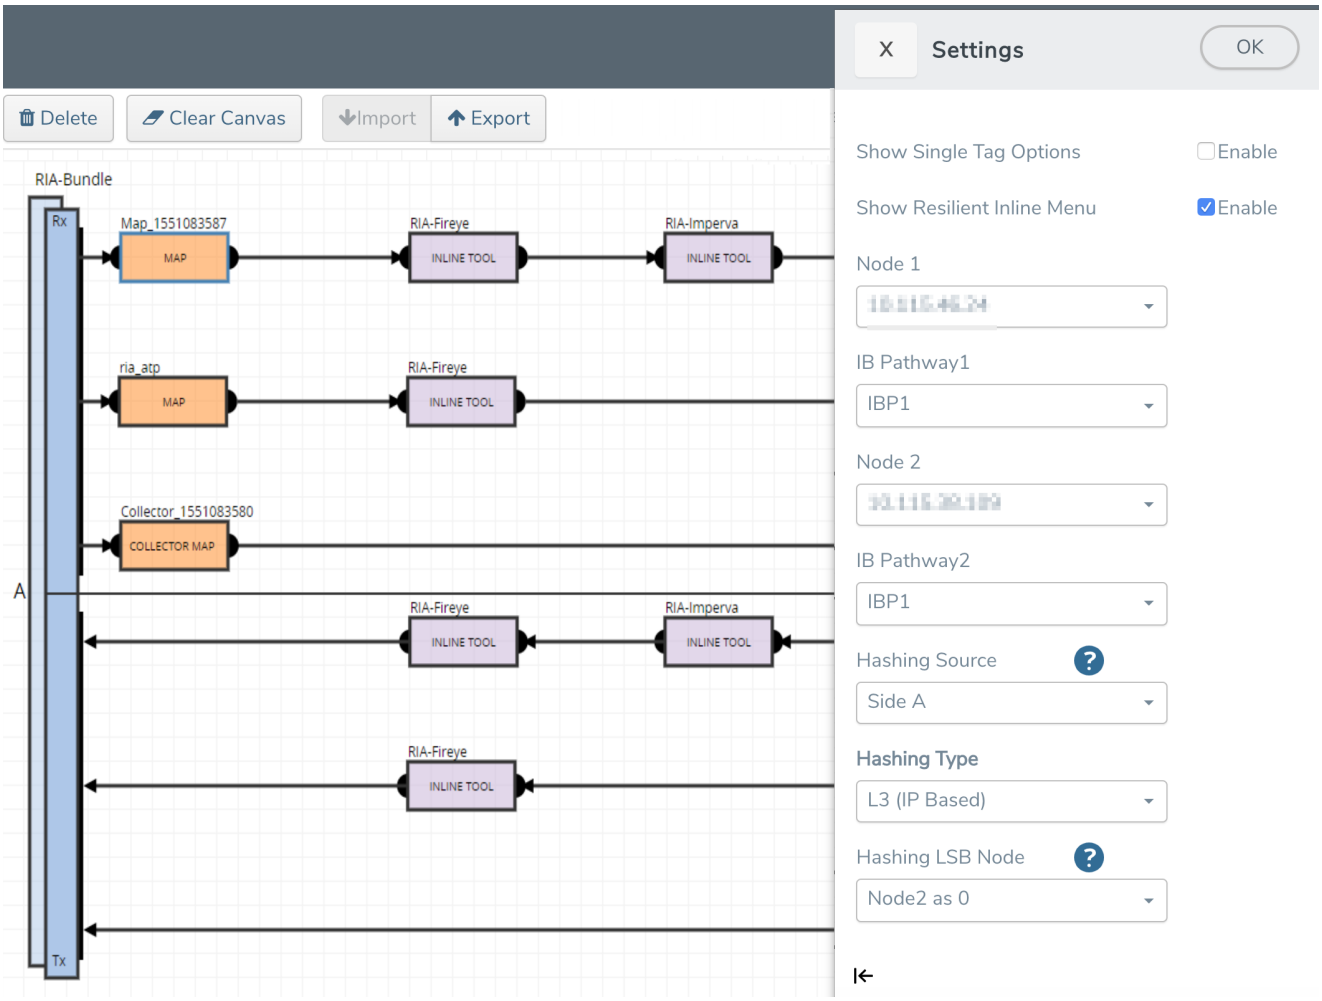
   NOTE: This field is available only if you select the **L4 (Port Based)** option in the **Hashing Type** field.
9. Click **OK** to save the settings.
10. Drag and drop the flexible inline map into the canvas. Click the map to open the Properties pane.
11. In the **Alias** and **Description** fields, enter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

    NOTE: You can choose to disable the **Single Tag Mode** for collector maps, if required.
13. Enter the **Tool Side VLAN Tag** for the inline network you are configuring.
14. Select the **TPID** for the Tool Side VLAN Tag. The default value is 0x8100. You can also choose from the supported values 0x9100 and 0x88a8 in the drop-down list.
15. From the **FlexInline Failover** drop-down list, select one of the following options:
    - **Bypass** - Traffic passes directly between the respective inline network ports.
    - **Original Map** - Traffic follows the path defined in this Flexible Inline Map.
16. Add the required rules for the inline map. Click **OK** to save the configuration.
17. Drag and drop the required **Inline Tools** or **Inline Tool Group** into the canvas.
18. If needed, drag and drop the **OOB Copy** into the canvas.
19. From the **Destination Ports** drop-down list, select the required hybrid or tool ports.
20. From the **VLAN Tag** drop-down list, select one of the following options:
    - **None** - No VLAN tag is used; traffic is routed to a different destination.

- **Original** - Uses the original VLAN tag from the packet received from the inline network.
- **As Inline** - Uses the VLAN tag configured for the Flexible Inline Map.

NOTE: The **As Inline** option is available only when you configure the Resilient Inline Arrangement with a single VLAN tag.

21. Click **Deploy**, select a traffic path and click **OK**.



Configure ICAP Client

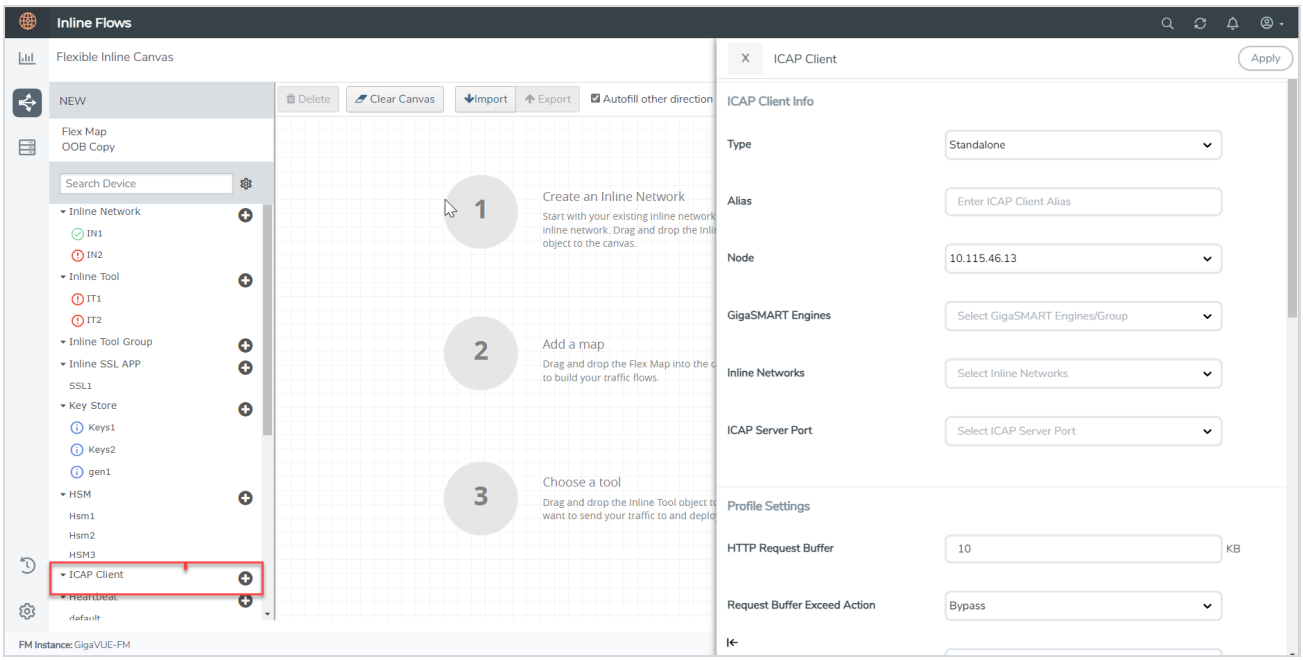
Before You Begin:

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

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

To configure the ICAP Client:

1. Go to  **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed Flexible Inline Canvas, select the device where you want to configure the ICAP Client.
3. Click the '+' icon next to the **ICAP Client** to create a new entry.
4. Enter a name and description for the inline network in the **Alias** and **Description** fields.



5. In the **ICAP Client** properties pane that appears on the right, complete the required fields in the ICAP Client Info, Profile Settings, and the Server sections. Refer to ICAP Client—Field References for more details.
6. Select one from the following:
   - If ICAP and iSSL are integrated and deployed on the same node, select the **Same Node** option.
   - If ICAP and iSSL are integrated and deployed on different node, select the **Different Node** option.
7. Select the required tool port of Inline SSL profile, which is connected to the inline network of ICAP. (This option will not appear if you select standalone as type.)
8. Click **Apply** to save the configurations.

ICAP Client—Field References

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

Field	Description
ICAP Client Info	
Type	Select the required type from the following: • If ICAP is deployed without iSSL and the traffic from devices other than Gigamon is sent to ICAP, select the **Standalone** option. • If ICAP and iSSL are integrated and deployed on the same node, select the **Same Node** option. • If ICAP and iSSL are integrated and deployed on different node, select the **Different Node** option.
Alias	Enter a unique name for the ICAP Client.
Node	Select the required node for which you want to configure the ICAP Client.
GigaSMART engines	Select the required GigaSMART engine.
Inline Networks	Select the required inline networks, which are the source for the ICAP app.
Source Ports	Select the required tool port of iSSL, which is connected to the inline network of ICAP. (This option will not appear if you select standalone as type.)
ICAP Server Port	Select the required IP interface physical port.
Profile Settings	
HTTP Request Buffer	Enter the HTTP request buffer size.
Request Buffer Exceed Action	Select the buffer action on exceeding the size: • Drop - The traffic is dropped. • Bypass - The traffic bypasses on exceeding the buffer size.
Preview Size	Enter the required preview size.
REQMOD	Enable or disable the Request Modification Mode.
RESPMOD	Enable or disable the Response Modification Mode.
Response Timeout	Enter the response timeout in seconds.
Response Timeout Action	Select the response timeout action on exceeding the timeout: • Drop - The traffic is dropped • Bypass - Traffic is bypasses on exceeding the response timeout limit.
Inactivity Timeout	Enter the session inactivity timeout value in seconds.
Port Range	Enter the ICAP client source port range for connecting to ICAP server.
Server	
Server Alias	Enter a unique name for the ICAP Server.
Description	Enter a description for the ICAP Server.
Port	Enter the L4 port number on which the ICAP server is listed.
Address	Enter the L3 IPV4/IPV6 address of the ICAP server.

Field	Description
Request Modification URL	Enter the Request Modification service URL.
Response Modification URL	Enter the Response Modification service URL.
Options URL	Enter the ICAP options URL (if necessary).

Configure GRIP Solution

To configure the GRIP solution 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** to create a new Flexible Inline Canvas.
8. In the displayed canvas, select the device where you want to configure and click the '+' icon next to the **Redundancy** to create a new entry.
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**.

Configure Synchronization

You must synchronize the configuration of the two GigaVUE-FM,GigaVUE HC Series 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-FM,GigaVUE HC Series nodes
- the redundancy profiles. The redundancy profile of each GigaVUE-FM,GigaVUE HC Series 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-FM,GigaVUE HC Series node. For example, one is configured with the primary role and one is configured with the secondary role.
- the inline tools
- the inline maps

In the example below , the configuration is the same on both nodes, except for the protection role (primary or secondary).

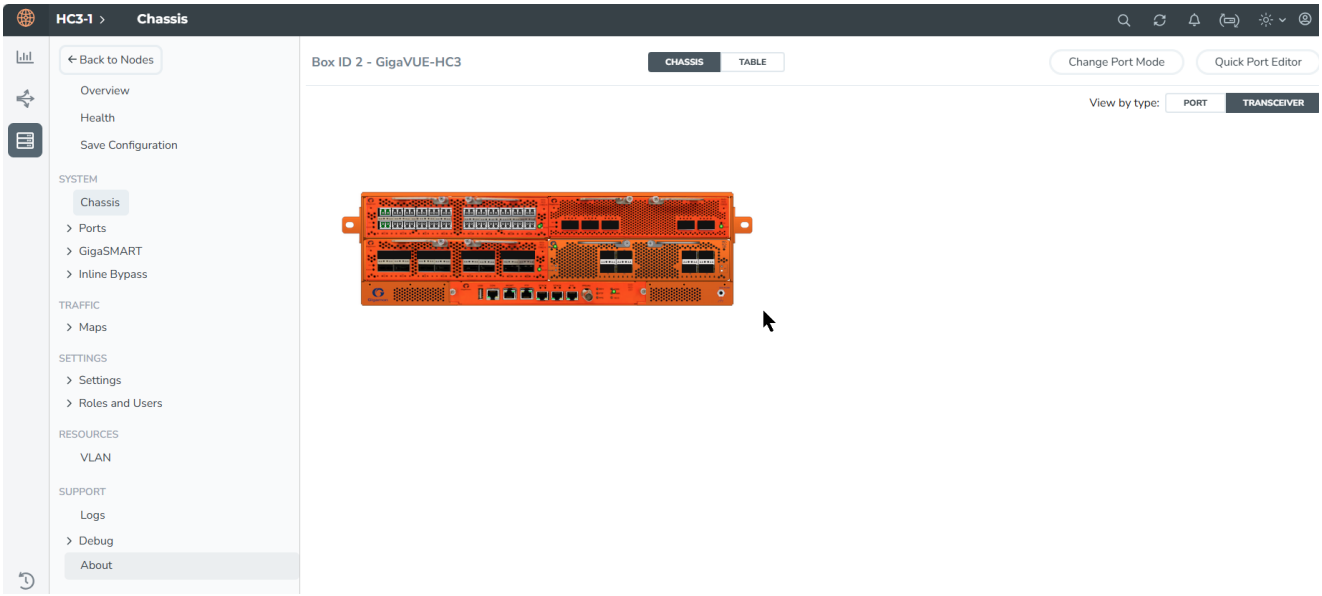
Example: Gigamon Resiliency for Inline Protection

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

First, configure the GigaVUE-HC1 with the primary role, then configure the GigaVUE-HC1 with the secondary role. The configuration is the same (is synchronized) on both nodes, except for step 3, in which the protection role (primary or secondary) is specified.

Note that in this example, link fail propagation (LFP) is disabled to reduce inline network recovery time after failover. When a primary to secondary failover occurs and LFP is enabled for copper Inline Bypass links, network service recovery may take several seconds because of Ethernet link renegotiation. Optical links failover faster and typically recover service much faster. For inline networks where only one path is available, this is a consideration. When GRIP is deployed with high availability networks where a second path is present, it is a best practice to leave LFP enabled.

You can use the Chassis page to view the chassis and modules.



Configure Primary Role GigaVUE-HC1

Task	Description	UI Steps
1.	Configure ports on the TAP-HC0-G100C0 module as passive (in passive mode, relays are closed). Also configure ports, port type (inline-network).	**1.** Go to **System > Ports > Ports > All Ports**. **2.** Select a port of the TAP-HC1-G10040 module. **3.** Click to open the Ports page. **4.** Select Passive for TapTX **5.** Click Save. **6.** Repeat steps 2 through 6 for each port on the TAP-HC1-G10040 module **7.** Configure Inline Network ports **a.** Select the port. **b.** Click Quick Port Editor. **c.** Select Inline Network for Type. **NOTE:** You can use the Chassis page to locate the position of the module in the chassis and identify port IDs.
2.	Configure stack port (for signaling port/link) and enable it.	**1.** Select the port and click Edit. **2.** Select Enable for Admin. **3.** Select Stack for Type. **4.** Click OK.
3.	Create the redundancy profile by giving it a name and configuring parameters for the redundancy profile such as the signaling port and protection role (primary).	**1.** Select Physical > Orchestrated Flows > Inline Flows > Configuration Canvas > Redundancy. **2.** Click '+' icon. **3.** Enter a name for the profile in the Alias field. For example, RP_001. **4.** Click in the Signaling Port field and select the stack port configured in Task 2.

Task	Description	UI Steps
		5. Select Primary for Protection Role. **6.** Click OK.
4.	Configure inline network. This step associates the redundancy profile to the inline network and also disables link fail propagation on the inline network.	Refer to the Configure Inline Network Ports and an Inline Networksection for configuration details.
5.	Configure inline tool ports, port type (inline-tool), and administratively enable them.	**1.** Select Physical > Orchestrated Flows > Inline Flows > Configuration Canvas > Inline Tool. **2.** Select the first port (for example, 1/4/x1) to configure as an inline-tool port. **3.** Select Inline Network for Type and select Enable for Admin. **4.** Select the second port (for example, 1/4/x2) and repeat steps 3 through 5.
6.	Configure inline tool and failover action. then enable inline tool.	Refer to the Configure Inline Network Ports and an Inline Networksection for configuration details.
7.	Configure map passall, from inline network to inline tool. **NOTE:** When you delete a map on the primary node, irrespective of the inline-network traffic-path, the traffic is switched to the secondary node. The port utilization must be 0% on the primary node and active on the secondary node.	Refer to the Configure Inline Network Bundle section for configuration details.

Configure Secondary Role GigaVUE-HC1

Task	Description	UI Steps
1	Configure ports on the TAP-HC0-G100C0 module as passive (in passive mode, relays are closed). Also configure ports, port type (inline-network).	**1.** Go to **System > Ports > Ports> All Ports**. **2.** Select a port of the TAP-HC1-G10040 module. **3.** Click to open the Ports page. **4.** Select Passive for TapTX **5.** Click Save. **6.** Repeat steps 2 through 6 for each port on the TAP-HC1-G10040 module **7.** Configure Inline Network ports **a.** Select the port. **b.** Click Quick Port Editor. **c.** Select Inline Network for Type. **NOTE:** You can use the Chassis page to locate the position of the module in the chassis and identify port IDs.
2	Configure stack port (for signaling port/link) and enable it.	**1.** Select the port and click Edit. **2.** Select Enable for Admin. **3.** Select Stack for Type. **4.** Click OK.
3	Create the redundancy profile by giving it a name and configuring parameters for the redundancy profile such as the signaling port and protection role (secondary).	**1.** Select Physical > Orchestrated Flows > Inline Flows > Configuration Canvas > Redundancy. **2.** Click '+' icon. **3.** Enter a name for the profile in the Alias field. For example, RP_001. **4.** Click in the Signaling Port field and select the stack port configured in Task 2. **5.** Select Secondary for Protection Role. **6.** Click OK.
4	Configure inline network. This step associates the redundancy profile to the inline network and also disables link fail propagation on the inline network.	Refer to the Configure Inline Network Ports and an Inline Network section for configuration details.
5	Configure inline tool ports, port type (inline-tool), and administratively enable them.	**1.** Select Physical > Orchestrated Flows > Inline Flows > Configuration Canvas > Inline Tool. **2.** Select the first port (for example, 1/4/x1) to configure as an inline-tool port. **3.** Select Inline Network for Type and select Enable for Admin. **4.** Select the second port (for example, 1/4/x2) and repeat steps 3 through 5.

Task	Description	UI Steps
6	Configure inline tool and failover action. then enable inline tool.	Refer to the Configure Inline Tool Ports and Inline Tools section for configuration details.
7	Configure map passall, from inline network to inline tool.	Refer to the Configure Inline Network Bundlesection for configuration details.
8	Configure the path of the traffic to the inline tool, disabling physical bypass on the inline network to open the relay on the node with the primary role.	Refer to the Configure Inline Network Ports and an Inline Network section for configuration details.

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.

The below table describes the different Redundancy Control States.

Table 11: 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**".

How to Upgrade GigaVUE Nodes in GRIP Deployment

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.

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.

Configure TLS/SSL Decryption Solutions

NOTE: Starting in release 6.14.00, Entrust/nCipher HSM integration is deprecated. Gen2 DSSL HSMs (nCipher v1.0) and Gen3 HSMs using nCipher 12.6 (`nshield_12_60`) are no longer supported, and upgrades are blocked if any nCipher HSM configuration is present.

The Inline TLS/SSL Decryption solution can be configured based on use cases. Refer to the following topics for various use cases:

- Configure Inline TLS/SSL Decryption Solution with Layer 2 Tools
- Configure Inline TLS/SSL Decryption Solution with Layer 3 Tools
- Configure Inline TLS/SSL Decryption Solution with RIA
- Configure Thales-Luna HSM for TLS/SSL Decryption
- Configure ICAP Client for Inline TLS/SSL Decryption Solution

Inline SSL App—Field References

These are the fields are that you will come across while configuring the **Inline SSL APP**. The table below provides a list and description of the attributes that define the Flexible Inline Decryption solution.

Field	Description
Alias	Enter a unique name for the flexible inline SSL APP.
Resilient Inline Arrangements	Enable this to configure a Resilient Inline Arrangement.
GS engines	Select the required GigaSMART engines.
TLS/SSL Monitor Mode	Select a TLS/SSL Monitor Mode from one of the following options: • Enable—When the monitor mode is enabled, the TLS/SSL decryption or encryption is off. The monitor application collects information, such as the TCP ports in use and VLAN information about the incoming traffic, and forwards the packets to the tool port or network port based on the non-TLS/SSL TCP bypass action. • Disable—This is the default value. When the monitor mode is disabled, the TLS/SSL decryption or encryption is on. Use this mode during the deployment stage. • Inline—Both monitor mode and TLS/SSL decryption or encryption are on. Use this mode to debug issues.
HSM Group	Select an HSM Group alias that you have configured from the drop-down list. Select **Disable** from the drop-down list to disable the HSM Group.

Field	Description
	Notes: • Thales-Luna Network HSM configuration is supported in Inbound, Outbound, and Hybrid deployment types.
Advanced Session Statistics	Enable this option to visualize advanced Inline SSL Session dashboards, such as Session Insights and Session Table, in the Fabric Health Analytics dashboard. The basic dashboards are available by default as you configure an Inline SSL session. Refer View Inline TLS/SSL Dashboards to know more.
Keychain Password	The keychain password must be configured before installing certificates and keys. To add or reset the Keychain Password: **a.** Click **Keychain Password**, and then choose either **Add** or **Reset**. **b.** If you choose to reset the Keychain Password, enter a password that is 8 to 30 characters long and contains at least one numerical character, one uppercase character, one lowercase character, and one special character. **c.** Select the **Auto Login** check box to enable GigaVUE-FM to unlock the keystore when the device reboots. **d.** Click **OK** to save the Keychain Password.
Add new keys	To configure a certificate-key pair: **a.** Click **Add new keys** to open the Key page. **b.** Enter a name and description for the key. **c.** Select the required **Key Type** and **File Type**. **d.** You can choose to include a Passphrase for the key when you select PEM or PKCS12 as File type if needed. **e.** When you choose Luna-HSM, enter the Key label for the key. **f.** Add the required **Private Key** and **Certificate**. **g.** Click **OK** to save the configuration.
Deployment Type	Select one of the following deployment types: • Inbound—For inbound deployments, add a new **Server Key Mapping**. Enter the domain name or IP address of the server, and then select the required **Key Pair Alias**. • Outbound—For outbound deployments, add a primary and a secondary signing Certificate Authorities (CA). • Hybrid—For hybrid deployments, add a new **Server Key Mapping**, and a primary and a secondary signing CA.
mTLS	Enables mutual TLS (client certificate authentication) for inbound and hybrid Inline TLS/SSL sessions that use this profile.
Configurations	
Default Action	Select one of the following options : ▪ Decrypt—Decrypt all the traffic that is guided into the Inline SSL APP. ▪ No Decrypt—Do not decrypt the traffic that is guided into the Inline SSL APP.
URL Cache Miss Action	Select one of the following options: • Decrypt— This option Decrypts all the traffic that is guided into the Inline SSL APP.

Field	Description
	• No Decrypt— This will not decrypt the traffic that is guided into the Inline SSL APP. • Defer—Delay the decryption until the Defer Timeout seconds provided.
Tool Fail Action	The failover action taken in response to a failure of an inline tool. Select one of the following options: • Bypass Tool—The traffic bypasses the failed inline tool. • Drop Connection—The traffic is dropped.
Tool Bypass	Select the required options: • Decrypted TLS/SSL Traffic—Bypasses the decrypted SSL traffic. • No Decrypted TLS/SSL Traffic—Bypasses the non-decrypted SSL traffic. • Non-TLS/SSL TCP Traffic—Bypasses the non-TLS/SSL, that is the TCP intercepted traffic.
High Availability	Select the check box to detect the link switchover by upstream device that is in active or standby mode. **NOTE:** Do not select this check box if the inline network links are in active state.
Network Group Multiple Entry	Select this check box to allow the traffic from different inline network to reenter GigaSMART.
Tool Early Engage	Select this check box to allow the inline tools to change the MAC address or VLAN IDs. When a connection request is received from the client, GigaSMART establishes the connection with the inline tool first, before connecting with the server. This helps the inline tools to modify the MAC address or VLAN IDs when sending the traffic back to the server.
HTTP Downgrade	HTTP 2.0 Downgrade option is enabled by default. HTTP 2.0 traffic is downgraded to HTTP 1.1 for decryption. If the downgrade option is disabled, HTTP 2.0 traffic is forwarded without decryption.
NAT/PAT mode	Enable this to perform NAT/PAT (Network/Port Address Translation) in Layer-3 inline tools .
Tool Early Inspect	Select this check box to allow the inline tool to inspect the decrypted data first before connecting to the server. This will allow the inline tool to validate the data and ensure that only valid connections are sent to the server. Notes: • You can access **Tool Early Inspect** feature from the flex Inline SSL APP only. **Tool Early Inspect** cannot co-exist with features such as RIA, NAT/PAT mode, Tool Early Engage, One-Arm, and Decryption Port Mapping. • If **Tool Early Inspect** is enabled, you can configure the connections timeout value. Connection timeout represents the time by which the tool should respond after receiving the first decrypted data. If no response is received within the configured time interval, the connections will be reset.
StartTLS Port	Enter the required SSL/TLS ports.
MTU	The Maximum Transmission Unit (MTU) is the maximum size of each packet that can be transferred as a single entity in a network connection.

Field	Description
Session Logging	
Session Logging	Select the Enable check box to log the Inline TLS/SSL session related information to a remote server.
IP Version	Select IPV4 or IPV6 as the IP Version for the Session Logging server. You can select one session logging configuration per GigaSMART group.
Remote Syslog Server IP	Enter the IP address of the remote syslog server.
Associated IP Interface	In the Associated IP interface drop-down list, select the IP interface that you assigned to the GigaSMART group.
Remote Syslog Port Number	Enter the port number of the remote syslog server.
Log Level	In the Log Level drop-down list, select the severity log level of the events that you want to send to the inline TLS/SSL session logging server.
Traffic Path	
Single VLAN Tag	Enable the check box to deploy flexible inline TLS/SSL solution with a single VLAN tag. If an inline tool is involved in an inline TLS/SSL map, the inline tool can be supported across multiple maps with different single VLAN tags. **NOTE:** Deploying a flexible Inline SSL solution with SVT is optional, and you can choose to enable or disable the Single VLAN Tag option. If you choose to enable the Single VLAN Tag option in the iSSL solution, you should also enable the Single VLAN Tag configuration in the flex map deployed in that solution. **NOTE:** If you enable the Single VLAN tag option in the Flexible Inline SSL solution, you should also enable the Single VLAN Tag configuration in the inline-ssl app profile deployed in the solution
Tool Side VLAN Tag	Enter the required tool side VLAN tag for the inline network.
TPID	Select the **TPID** for the Tool Side VLAN Tag. The default value of TPID is 0x8100. You can select the other supported values 0x9100 and 0x88a8 from the drop-down list.
Traffic Path	Select one of the following options: • Drop—Traffic is dropped at the virtual port. • Bypass—Traffic bypasses the virtual port. • Monitoring—Traffic is fed to the virtual port and absorbed, while a copy of the traffic is sent to the next inline tool in the sequence. Traffic returned from side B of the network is also absorbed at the virtual port in the monitoring mode. **NOTE:** You can select the **Monitoring** option only if you have set the **SSL Monitor Mode** to either **Enable** or **Inline**. • To Inline Tool—Traffic is forwarded to the inline tool. This is the default value.
Inline Failover Action	Select one of the following options: • Virtual port bypass—All virtual ports configured as the source of any map that triggered this failover action, will be put in the bypass mode, that is all traffic will bypass the virtual port and will be guided to the inline tool or inline tool group. • Virtual port drop—All virtual ports configured as the source of any map that triggered this failover action, will be put in the drop mode, that is all traffic will be

Field	Description
	dropped at the virtual port. • Network bypass—All inline networks configured as the source of any map involving the inline tool or inline tool group that triggered this failover action, will be put in the bypass mode, that is, all traffic coming to side A will be directed to side B and vice versa. • Network drop—All inline networks configured as the source of any map involving the inline tool or inline tool group that triggered this failover action, will be put in the drop mode, that is, all traffic coming to side A or side B will be dropped. • Network port forced down—For all inline networks configured as the source of any map involving the inline tool or inline tool group that triggered this failover action, the inline network ports will be brought down.
Security Exceptions	You can choose to either decrypt or drop the traffic for the following certificates: • Self-signed certificate • Unknown CA certificate • Invalid certificate • Expired certificate • You can also choose to configure the security exceptions for certificate revocation validation based on OCSP or CRL on an inline decryption profile. Select one of the following options: o Soft Fail—If you select this option, the client browser displays the secondary MitM certificate and the inline decryption session stats in GigaVUE-FM displays as Decrypt. o Hard Fail—If you select this option, the client browser displays the certificate from DigiCert and the inline decryption session stats in GigaVUE-FM displays as Bypass: Unknown Revocation.
No-decrypt list/Decrypt list	Select the following check boxes: • No-decrypt list—Allows traffic from certain classes such as sites, domains, host-based IP address and IP subnets (decision based on LPM) to bypass decryption. • Decrypt list—Allows traffic from certain sites, domains, host-based IP address and IP subnets (decision based on LPM) to always be decrypted. Select from the below operations that can be performed on an uploaded list: • Append _ This would add to the uploaded list. • Replace- This would remove the previously added list and add a new list. This option is supported only on Generation 3 cards. • Clear- This would completely clear the list. • Download - This would download the list that has been uploaded. If you select Append/Replace, you can enter the list using any of the following options: · Copy and Paste · Install from URL · Install from Local Directory
Policy Rules	Add the required policy rules for the inline decryption profile. Click **Add a Rule**. In the **Condition** field, Select one of the following options from the drop-down list:

Field	Description
	• Category • Domain • IPv4 Destination • IPv4 Source • IPv6 Destination • IPv6 Source • L4 Port Destination • L4 Port Source • VLAN • X509 Certificate Issuer Name Select one of the following options: • Decrypt—Decrypt all the traffic that is guided into the Inline SSL APP. • Decrypt—Do not decrypt the traffic that is guided into the Inline SSL APP.
Network Access	Network access configuration is used to get URL categorization updates To configure the network access for the GigaSMART engine ports: o Select either DHCP or IP Address as the network access configuration mode. o If you select IP Address as the mode, enter the IP Address, Netmask, Gateway, DNS, MTU, and VLAN. ▪ **DNS or Split DNS**- Configure either a default single DNS server or a Split DNS Server. If you want to attach Split DNS server profile to your Inline SSL deployment choose a Split DNS server from the drop-down. To configure a new Split DNS Profile, click on **Create new Split DNS**. o Select either Eth2 or Eth3 as the Interface. o If you want to attach a Proxy profile to your Inline SSL deployment select a **Proxy Server Profile** from the drop-down. To configure a new Proxy Server Profile, click on **Create new Proxy**. **Notes:** • The Eth3 option is available only for GigaVUE-HC3 devices. • IP Address configuration mode details should be entered when you select Luna HSM configuration from the **HSM Group** drop down. • If your Proxy Server profile is associated with an Inline SSL application, choose **'None'** in the Proxy Server profile field on the Inline SSL configuration page to disconnect the proxy server profile prior to deleting the profile. • You cannot enable Gen 2 and Gen 3 GigaSMART engine for network access simultaneously.
Decryption Port Mapping	The TCP destination port for decrypted traffic sent to inline tools can be configured as part of the inline decryption profile. Configure the required Priority 1 map, which is user configurable and Priority 2 map, which is the default out port.
Trust Store	**Server Trust Store**: The trust store contains a trusted certificate authority (CA) for server validation. You can choose to either append or replace the trust store. **Client Trust store** The trust store contains a trusted certificate authority (CA) for client validation

Field	Description
	when mTLS is enabled for this profile. You can choose to either append or replace the trust store. **NOTE:** The Client trust store is used only for inbound and hybrid mTLS sessions. If you later disable mTLS and redeploy the device for outbound only Inline TLS/SSL decryption, then it will remove all client trust store and signing certificate configurations that are attached to this profile.
TCP Settings	Configure the required TCP settings as follows: • **TCP Inactive Timeout**— TCP Inactive session timeout in minutes. • **TCP Delayed ACK**—GigaSMART Inline TLS/SSL decryption ACKs every TCP packet by default. If TCP Delayed ACK is enabled, then GigaSMART decryption will wait for 100ms or ACK every third packet – whichever comes first. • **TCP SYN Retries**—number of retries made by the MitM to initiate a session with the destination server. If a SYN/ACK response isn't received from the destination server on initial TCP SYN, GigaSMART attempts for additional number of TCP SYN Retries as defined by the user. • **TCP TIMEWAIT Timeout**— Configure the 'TCP TIMEWAIT' timeout value from 0-300 seconds. The default value is 30 seconds. The TCP connection in the TIME_WAIT state gets deleted after the timeout period.
Split-Proxy Settings	
Split-Proxy	Select the check box to enable the split proxy settings for the inline decryption solution. The TLS connection between the server and client is divided into two independent connections, and the security parameters are kept separate.
Non-PFS Ciphers (Server)	Select the check box to enable the non-PFS ciphers settings for the inline decryption solution that has the split proxy settings enabled. This setting is to indirectly force the server to use protocols that are lower than TLS 1.3 with non-PFS ciphers. This means that the ciphers with DHE/ECDHE key-exchange will not be used on the server side.
Miscellaneous (Global Settings)	
SSL/TLS Version	Select the minimum and maximum SSL/TLS version.
Connection Reset Action	Select one of the following options for the minimum SSL/TLS version: • Drop—Closes all sessions that are below the minimum SSL/TLS version specified. This ensures that the network is safe from weak TLS/SSL connections. This is the default option. • No Decrypt—Bypasses all sessions that are below the minimum SSL/TLS version specified. Select one of the following options for the maximum SSL/TLS version: o No Decrypt—Bypasses all sessions that are above the maximum SSL/TLS version specified. This is the default option. o Drop—Closes all sessions that are above the maximum SSL/TLS version specified.
Caching persistence	Select this check box to allow the information to be saved on the node in the control card's persistent storage so that it can be retrieved in case of reboots.
DHE Cipher Suites	Enable this use DHE Cipher suites.

Configure Inline TLS/SSL Decryption Solution with Layer 2 Tools

This section describes how to configure and deploy an Inline TLS/SSL Decryption solution using Layer 2 tools.

Prerequisites

- A configured and reachable GigaSMART engine on the selected device to host the Inline SSL app.
- Configure and unlock the Keychain Password in GigaVUE-FM,GigaVUE HC Series to access the keystore; optionally enable Auto Login for unattended restarts.
- Prepare required keys and certificates for your deployment direction:
  - Inbound: server private key(s) and corresponding certificate(s).
  - Outbound/Hybrid: appropriate CA/certificate chain for re-signing or policy-based decryption.
- Keep key types and file formats (PEM/PKCS12/Luna-HSM) and any passphrases ready.
- Decide the deployment type you will use (Inbound, Outbound, or Hybrid) so the Inline SSL app can be configured accordingly.
- Plan the software constructs you will create during configuration: at minimum, an Inline Network and an Inline Tool to connect your Layer 2 (transparent) inspection device(s). Ensure the required ports are identified and available.
- Ensure your connected tools can operate in Layer 2 transparent mode (inspect decrypted traffic without altering packet contents)

Access Flexible Inline Canvas

1. Go to  **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed Flexible Inline Canvas, select the device where you want to configure the TLS/SSL decryption.

Configure Inline Network

1. On the left pane, click the '+' icon next to **Inline Network** option to create a new entry.

2. Enter a name and description for the inline network in the **Alias** and **Description** fields. Then, click **Port Editor**.
3. In the **Alias** and **Description** fields, enter a name and description for the inline network.
4. Click **Port Editor**. In the **Quick Port Editor**, scroll down to the inline network ports that you want to configure. Select **Enable** to administratively enable inline network ports, and click **OK** to apply the changes.
5. From the **Port A** and **Port B** drop-down lists, select the ports that you want to configure as the inline network pair.
6. From the **Traffic Path** drop-down list, select **To Inline Tool**.
   - **To Inline Tool**—All traffic originating from the inline network is directed to the sequence of inline tools and inline tool groups and is guided through the inline tools and inline tool groups according to the current inline tool and inline tool group status.
7. Click **OK**. You have successfully created an Inline Network.

To explore additional options during configuration, refer to Inline Network Ports and Inline Network.

> **NOTE:** If there are multiple Inline Network Ports, you can configure Inline Network bundle. Refer to Configure Inline Network Bundle.

Configure Inline Tool

1. On the left pane, click the '+' icon next to the **Inline Tool** option to create a new inline tool.
2. In the **Properties** pane, in the **Alias** and **Description** fields, enter a name and description for the inline tool.
3. From the **Type** drop-down list, select one of the following options:
   - **External**- To configure a third-party tool.
   - **GigaVUE Node**- To configure a GigaVUE node as a tool.
4. Click **Port Editor**. In the **Quick Port Editor**, scroll down to the inline tool ports that you want to configure. Select **Enable** to administratively enable the inline tool ports, and then click **OK** to apply the changes.
5. From the **Port A** and **Port B** drop-down lists, select the inline tool ports according to the direction the inline tool expects traffic from the network.
6. Verify that the **Enabled** check box is selected.
7. Click **OK**. You have successfully created an Inline Network.

To explore additional options during configuration, refer to Inline Tool Ports and Inline Tools.

Create an Inline SSL APP

1. On the left pane, click the '+' icon next to the **Inline SSL APP** option.
2. Enter a name for the Inline SSL APP and select the required GigaSMART engines.
3. In the **Alias** field, enter a name for the Inline SSL APP.
4. From the **GigaSMART Engine** drop-down list, select the required engine.
5. Under **Deployment Type**, set up **Key chain Password**:
    a. Click **Keychain Password**, and choose either Add or Reset.
    b. If you choose to reset the Keychain Password, enter a password that is 8 to 30 characters long and contains at least one numerical character, one uppercase character, one lowercase character, and one special character.
    c. Select the **Auto Login** check box to enable GigaVUE-FM,GigaVUE HC Series to unlock the keystore when the device reboots.
    d. Click **OK** to save the Key chain Password.
6. Configure **Keys and Certificates**. A key in an inbound deployment can be selected only for decryption or for re-signing and re-encryption in an outbound deployment.
    a. Click **Add Keys** to open the Key page.
    b. Enter a name and description for the key.
    c. Select the required **Key Type** and **File Type**.
    d. If using PEM or PKCS12 as file type, optionally include a passphrase.
    e. You can choose to include a Passphrase for the key when you select PEM or PKCS12 as file type if required.
    f. When you choose Luna-HSM, enter the Key label for the key.
    g. Add the required **Private Key** and **Certificate**.
    h. Click **OK** to save the configuration.
7. Select the required deployment type.
    - **Outbound** - Add the configured primary and a secondary signing Certificate Authorities (CA).
    - **Inbound** - Add a new Server Key Mapping. Enter the domain name or IP address of the server, and select the required Key Pair Alias.
    - **Hybrid** - Add a new Server Key Mapping, and a primary and a secondary signing CA.
8. Click **OK** to save the configuration.

Deploy Inline TLS/SSL Decryption Solution

In the Flexible Inline Canvas:

1. Drag and drop the required inline network or inline network bundle.
2. Drag and drop the flexible inline map into the canvas. The **Properties** pane opens automatically. If configuration is needed, provide values for the following fields:
    - In the **Alias** and **Description** fields, enter the name and description of the inline map.
    - Enter the **Tool Side VLAN Tag** for the inline network for which you are configuring the map.
    - 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.
    - Add the required rules for the inline map, and click **OK** to save the configuration.
3. Drag and drop the Inline SSL APP.
4. Drag and drop the required inline tools or inline tool group.
5. Drag and drop the **OOB Copy**, if required.
6. Click **Deploy**, and choose the Traffic Path—either **Logical Bypass** or **Keep as is.** Select the option that best suits your deployment needs, and click **OK**.

Verify the Solution

In GigaVUE-FM, open the Flexible Inline Canvas: **Physical > Orchestrated Flows > Inline Flows > Configuration Canvas**, select the target device, and open the deployed solution on the canvas. Confirm the solution shows as deployed.

What to Do Next

After deployment, you can view the Monitor and Session statistics. Refer to View Inline TLS/SSL Session Statisticson Statistics. You can also view the Inline TLS/SSL Dashboards. Refer to View Inline TLS/SSL Dashboards.

Configure Inline TLS/SSL Decryption Solution with Layer 3 Tools

This section describes how to configure and deploy an Inline TLS/SSL Decryption solution using Layer 3 tools.

Prerequisites

- A configured and reachable GigaSMART engine on the selected device to host the Inline SSL app.
- Configure and unlock the Keychain Password in GigaVUE-FM to access the keystore (optionally enable Auto Login for unattended restarts).
- Prepare required keys and certificates for your selected deployment type: inbound (server private keys/certs), outbound (CA/resigning chain), or hybrid. Ensure supported file types (PEM/PKCS12/Luna-HSM) and passphrases are available.
- Decide the deployment type (Inbound, Outbound, or Hybrid) so the Inline SSL app can be configured accordingly.
- Plan the software constructs you will create: at minimum an Inline Network (and, where applicable, an Inline Network Bundle) plus an Inline Tool path for the Layer 3 security device(s); ensure the necessary ports are identified and available.
- Define routing/policy selection for Layer 3: identify the subnet-based IP ranges at the Inline SSL profile level that will be steered through the L3 tool path.
- Plan to enable NAT/PAT in the Inline SSL app when offloading TLS decryption from Layer 3 inline tools (Policy Profile > NAT/PAT).

Access Flexible Inline Canvas

1. Go to  **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed Flexible Inline Canvas, select the device where you want to configure the TLS/SSL decryption.

Configure Inline Network

1. On the left pane, click the '+' icon next to **Inline Network** option to create a new entry.

2. Enter a name and description for the inline network in the **Alias** and **Description** fields. Then, click **Port Editor**.
3. In the **Alias** and **Description** fields, enter a name and description for the inline network.
4. Click **Port Editor**. In the **Quick Port Editor**, scroll down to the inline network ports that you want to configure. Select **Enable** to administratively enable inline network ports, and click **OK** to apply the changes.
5. From the **Port A** and **Port B** drop-down lists, select the ports that you want to configure as the inline network pair.
6. From the **Traffic Path** drop-down list, select **To Inline Tool**.
   - **To Inline Tool**—All traffic originating from the inline network is directed to the sequence of inline tools and inline tool groups and is guided through the inline tools and inline tool groups according to the current inline tool and inline tool group status.
7. Click **OK**. You have successfully created an Inline Network.

To explore additional options during configuration, refer to Inline Network Ports and Inline Network.

Repeat steps 1 through 7 to create the required number of Inline Networks and group them into an Inline Network Bundle.

Configure Inline Network Bundle

1. On the left pane, click the '+' icon next to **Inline Network Bundle** option to create a new entry.
2. Enter a name and description for the inline network bundle in the **Alias** and **Description** fields.
3. From the **Inline Networks** drop-down list, select the inline networks you want to include in the bundle.
4. Click **OK** to save the bundle configuration.

Configure Inline Tool

1. On the left pane, click the '+' icon next to the **Inline Tool** option to create a new inline tool.
2. In the **Properties** pane, in the **Alias** and **Description** fields, enter a name and description for the inline tool.
3. From the **Type** drop-down list, select one of the following options:
   - **External**- To configure a third-party tool.
   - **GigaVUE Node**- To configure a GigaVUE node as a tool.

4. Click **Port Editor**. In the **Quick Port Editor**, scroll down to the inline tool ports that you want to configure. Select **Enable** to administratively enable the inline tool ports, and then click **OK** to apply the changes.
5. From the **Port A** and **Port B** drop-down lists, select the inline tool ports according to the direction the inline tool expects traffic from the network.
6. Verify that the **Enabled** check box is selected.
7. Click **OK**. You have successfully created an Inline Network.

To explore additional options during configuration, refer to Inline Tool Ports and Inline Tools.

Create an Inline SSL APP

1. On the left pane, click the '+' icon next to the **Inline SSL APP** option.
2. Enter a name for the Inline SSL APP and select the required GigaSMART engines.
3. In the **Alias** field, enter a name for the Inline SSL APP.
4. From the **GigaSMART Engine** drop-down list, select the required engine.
5. Under **Deployment Type**, set up **Key chain Password**:
    a. Click **Keychain Password**, and choose either Add or Reset.
    b. If you choose to reset the Keychain Password, enter a password that is 8 to 30 characters long and contains at least one numerical character, one uppercase character, one lowercase character, and one special character.
    c. Select the **Auto Login** check box to enable GigaVUE-FM,GigaVUE HC Series to unlock the keystore when the device reboots.
    d. Click **OK** to save the Key chain Password.
6. Configure **Keys and Certificates**. A key in an inbound deployment can be selected only for decryption or for re-signing and re-encryption in an outbound deployment.
    a. Click **Add Keys** to open the Key page.
    b. Enter a name and description for the key.
    c. Select the required **Key Type** and **File Type**.
    d. If using PEM or PKCS12 as file type, optionally include a passphrase.
    e. You can choose to include a Passphrase for the key when you select PEM or PKCS12 as file type if required.
    f. Add the required **Private Key** and **Certificate**.
    g. Click **OK** to save the configuration.
7. Select the required deployment type.

- **Outbound** - Add the configured primary and a secondary signing Certificate Authorities (CA).
- **Inbound** - Add a new Server Key Mapping. Enter the domain name or IP address of the server, and select the required Key Pair Alias.
- **Hybrid** - Add a new Server Key Mapping, and a primary and a secondary signing CA.

8. Under **Advanced > Configurations**, enable NAT/PAT (Network Address Translation/Port Address Translation) on traffic passing through Layer 3 inline tools to offload TLS decryption. Select the NAT-PAT checkbox to apply this setting.
9. Click **OK** to save the configuration. Refer toInline SSL App—Field References for any additional details when configuring the Inline SSL APP.

Deploy Inline TLS/SSL Decryption Solution

In the Flexible Inline Canvas:

1. Drag and drop the required inline network bundle.
2. Drag and drop the flexible inline map into the canvas. The **Properties** pane opens automatically. If configuration is needed, provide values for the following fields:
   - In the **Alias** and **Description** fields, enter the name and description of the inline map.
   - Enter the **Tool Side VLAN Tag** for the inline network for which you are configuring the map.
   - 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.
   - Add the required rules for the inline map, and click **OK** to save the configuration.
3. Drag and drop the Inline SSL APP.
4. Drag and drop the required inline tools or inline tool group.
5. Drag and drop the **OOB Copy**, if required.
6. Click **Deploy**, and choose the Traffic Path—either **Logical Bypass** or **Keep as is.** Select the option that best suits your deployment needs, and click **OK**.

Verify the Solution

In GigaVUE-FM, open the Flexible Inline Canvas: **Physical > Orchestrated Flows > Inline Flows > Configuration Canvas**, select the target device, and open the deployed solution on the canvas. Confirm the solution shows as deployed.

What to Do Next

After deployment, you can view the Monitor and Session statistics. Refer to View Inline TLS/SSL Session Statistics. You can also view the Inline TLS/SSL Dashboards. Refer to View Inline TLS/SSL Dashboards.

Configure Inbound mTLS for Inline TLS/SSL Decryption

This section describes how to enable client certificate authentication (mTLS) for inbound Inline TLS/SSL decryption deployments using Flexible Inline arrangements

Prerequisites

- Ensure that servers are configured to request or require client certificates for the relevant services.
- Ensure you have access to the CA certificates used to issue client certificates (to populate the client trust store).

Create a Client Trust store

1. On the left navigation pane, go to **Inventory** > **Physical Node>Nodes>Settings>Trust Store> Client Trust Store**
2. Create a new Client Trust Store:
   - Specify a unique alias for the client trust store.
   - Upload the CA certificates that issue client certificates for your environment.
3. Click on Apply to save the Client Authentication and it appears in the Client Trust store list.

Access Flexible Inline Canvas

1. Go to **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed Flexible Inline Canvas, select the device where you want to configure the TLS/SSL decryption.

Create an Inline SSL APP

1. On the left pane, click the '+' icon next to the **Inline SSL APP** option.

2. Enter a name for the Inline SSL APP and select the required GigaSMART engines.
3. In the **Alias** field, enter a name for the Inline SSL APP.
4. From the **GigaSMART Engine** drop-down list, select the required engine.
5. Under **Deployment Type**, set up **Key chain Password**:
    a. Click **Keychain Password**, and choose either Add or Reset.
    b. If you choose to reset the Keychain Password, enter a password that is 8 to 30 characters long and contains at least one numerical character, one uppercase character, one lowercase character, and one special character.
    c. Select the **Auto Login** check box to enable GigaVUE-FM,GigaVUE HC Series to unlock the keystore when the device reboots.
    d. Click **OK** to save the Key chain Password.
6. Configure **Keys and Certificates**. A key in an inbound deployment can be selected only for decryption or for re-signing and re-encryption in an outbound deployment.
    a. Click **Add Keys** to open the Key page.
    b. Enter a name and description for the key.
    c. Select the required **Key Type** and **File Type**.
    d. If using PEM or PKCS12 as file type, optionally include a passphrase.
    e. You can choose to include a Passphrase for the key when you select PEM or PKCS12 as file type if required.
    f. When you choose Luna-HSM, enter the Key label for the key.
    g. Add the required **Private Key** and **Certificate**.
    h. Click **OK** to save the configuration.
7. Select the required deployment type.
    - **Inbound** - Add a new Server Key Mapping. Enter the domain name or IP address of the server, and select the required Key Pair Alias.
8. Enable the **mTLS** toggle.
9. Under Client Trust store. Select the **Client Trust Store** or create a new one by clicking on the **Add Trust Store** button. You can configure to allow or drop sessions when:
    - The client certificate is expired.
    - The certificate chain is self-signed.
    - The certificate chain is signed by an unknown CA.
10. Click **OK** to save the configuration.

Deploy the solution

In the Flexible Inline Canvas:

1. Drag and drop the required inline network or inline network bundle.
2. Drag and drop the flexible inline map into the canvas. The **Properties** pane opens automatically. If configuration is needed, provide values for the following fields:
    - In the **Alias** and **Description** fields, enter the name and description of the inline map.
    - Enter the **Tool Side VLAN Tag** for the inline network for which you are configuring the map.
    - 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.
    - Add the required rules for the inline map, and click **OK** to save the configuration.
3. Drag and drop the Inline SSL APP.
4. Drag and drop the required inline tools or inline tool group.
5. Drag and drop the **OOB Copy**, if required.
6. Click **Deploy**.

Verify the Solution

In GigaVUE-FM, open the Flexible Inline Canvas: **Physical > Orchestrated Flows > Inline Flows > Configuration Canvas**, select the target device, and open the deployed solution on the canvas. Confirm the solution shows as deployed.

What to Do Next

After deployment, at the Monitor and Session statistics you can view if :

- the mTLS sessions are counted as decrypted inbound sessions
- Client certificate-related counters show successful validation events.

Refer to View Inline Solution Status and Statistics. You can also view the Inline TLS/SSL Dashboards. Refer to View Inline TLS/SSL Dashboards.

NOTE: If you later remove this inbound mTLS solution and deploy an outbound Inline SSL solution on the same device, first delete all client trust-store entries created for mTLS and any client key-store entries used for mTLS client authentication, and then deploy the new outbound Inline SSL application.

Configure Inline TLS/SSL Decryption Solution with RIA

This section describes how to configure and deploy Resilient Inline TLS/SSL.

Prerequisites

- Resilient Configuration requires two nodes . Ensure you have configured both devices with the required software constructs.
- Use inline-capable ports and, where applicable, inline bypass modules for fail-open behavior on power or device failure.
- Cable two inline nodes in path (client–device–server) with identical inline tool connectivity on each.
- Build an Inter-broker Pathway (IB-P) using at least one—preferably two or more aggregated—tool ports between nodes, and set a minimum-up count for IB-P status “Up.”
- For Inbound iSSL, prepare server private keys/certificates (RSA/ECDSA) with key pair aliases, define actions for unknown/invalid/self-signed certs, and set TLS policy (for example, decrypt TLS 1.2/1.3 and drop TLS 1.0/1.1).
- Ensure inline tools on both nodes are identical in count, type, speed, and capacity for consistent distribution and service behavior.
- Use only one RIA-enabled iSSL app per two-node pair.

Access Flexible Inline Canvas

1. Go to  **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed Flexible Inline Canvas, select the device where you want to configure the TLS/SSL decryption.

You should proceed to configure the following software constructs for both nodes before you deploy your solution.

Configure Inline Network

1. On the left pane, click the ‘+’ icon next to **Inline Network** option to create a new entry.
2. Enter a name and description for the inline network in the **Alias** and **Description** fields. Then, click **Port Editor**.

3. In the **Alias** and **Description** fields, enter a name and description for the inline network.
4. Click **Port Editor**. In the **Quick Port Editor**, scroll down to the inline network ports that you want to configure. Select **Enable** to administratively enable inline network ports, and click **OK** to apply the changes.
5. From the **Port A** and **Port B** drop-down lists, select the ports that you want to configure as the inline network pair.
6. From the **Traffic Path** drop-down list, select **To Inline Tool**.
   - **To Inline Tool**—All traffic originating from the inline network is directed to the sequence of inline tools and inline tool groups and is guided through the inline tools and inline tool groups according to the current inline tool and inline tool group status.
7. Click **OK**. You have successfully created an Inline Network.

To explore additional options during configuration, refer to Inline Network Ports and Inline Network.

> **NOTE:** If there are multiple Inline Network Ports, you can configure Inline Network bundle. Refer to Configure Inline Network Bundle.

Configure Inline Tool

1. On the left pane, click the '+' icon next to the **Inline Tool** option to create a new inline tool.
2. In the **Properties** pane, in the **Alias** and **Description** fields, enter a name and description for the inline tool.
3. From the **Type** drop-down list, select one of the following options:
   - **External**- To configure a third-party tool.
   - **GigaVUE Node**- To configure a GigaVUE node as a tool.
4. Click **Port Editor**. In the **Quick Port Editor**, scroll down to the inline tool ports that you want to configure. Select **Enable** to administratively enable the inline tool ports, and then click **OK** to apply the changes.
5. From the **Port A** and **Port B** drop-down lists, select the inline tool ports according to the direction the inline tool expects traffic from the network.
6. Verify that the **Enabled** check box is selected.
7. Click **OK**. You have successfully created an Inline Network.

To explore additional options during configuration, refer to Inline Tool Ports and Inline Tools.

Create Inter-broker Pathway

The inter-broker pathway (IBP) links the two RIA nodes so flows that hash to one node can traverse to the peer when needed, ensuring symmetric inspection and continuous traffic forwarding across nodes even during asymmetric paths or node/tool events.

To create a new inter-broker pathway:

1. Go to  **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed canvas, select the device where you want to create the inter-broker pathway.
3. Click the '**+**' icon next to the **IB Pathway** option to create a new inter-broker pathway.
4. In the **Properties** pane, enter a name and description in the **Alias** and **Description** fields..
5. From the **Ports** drop-down list, select the required tool ports to attach to the inter-broker pathway.

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

Configure Resilient Inline Arrangement

To configure a resilient inline arrangement:

1. Drag and drop the required inline network or inline network LAG into the flexible inline canvas, and click **Settings**.

2. In the **Settings** pane, select the **Enable** check box next to **Show Single Tag Options** to configure the resilient inline arrangement with a single VLAN tag.

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

3. Select the **Enable** check box next to **Show Resilient Inline Menu**.
4. Select the required **Node 1**, **Node 2**, **IB Pathway1**, and **IB Pathway2** for the resilient inline arrangement.
5. From the **Hashing Source** drop-down list, select one of the following options:
   - **Side A** - Hashing uses source IP or source port from Side A; destination IP or destination port from Side B.
   - **Side B** - Hashing uses source IP or source port from Side B; destination IP or destination port from Side A.
6. From the **Hashing Type** drop-down list, select one of the following options:
   - **L3 (IP Based)** - Hashing uses the IP address.
   - **L4 (Port Based)** - Hashing uses the transport layer port number.
7. From the **Hashing LSB Node** drop-down list, select one of the following options:
   - **Node 1 as 0** - Traffic from IPs ending in 0 is hashed to Node 2.
   - **Node 2 as 0** - Traffic from IPs ending in 0 is hashed to Node 1.

   NOTE: This option is available only if you selected **L3 (IP Based)** in the Hashing Type field.

8. From the **Hashing Port** drop-down list, select one of the following options:
   - **Node 1 as odd** - Traffic with odd port numbers is hashed to Node 2, while traffic with even port numbers is hashed to Node 1.
   - **Node 2 as odd** - Traffic with odd port numbers is hashed to Node 1, while traffic with even port numbers is hashed to Node 2.

   NOTE: This field is available only if you select the **L4 (Port Based)** option in the **Hashing Type** field.

9. Click **OK** to save the settings.
10. Drag and drop the flexible inline map into the canvas. Click the map to open the Properties pane.
11. In the **Alias** and **Description** fields, enter the name and description of the inline map.
12. To deploy the resilient inline arrangement with a s**ingle VLAN tag**, select the **Enable** check box next to Single Tag Mode. Refer to Configure Inline TLS/SSL Decryption Solution with RIA.

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

13. Enter the **Tool Side VLAN Tag** for the inline network you are configuring.
14. Select the **TPID** for the Tool Side VLAN Tag. The default value is 0x8100. You can also choose from the supported values 0x9100 and 0x88a8 in the drop-down list.
15. From the **FlexInline Failover** drop-down list, select one of the following options:
    - **Bypass** - Traffic passes directly between the respective inline network ports.
    - **Original Map** - Traffic follows the path defined in this Flexible Inline Map.
16. Add the required rules for the inline map. Click **OK** to save the configuration.
17. Drag and drop the required **Inline Tools** or **Inline Tool Group** into the canvas.
18. If needed, drag and drop the **OOB Copy** into the canvas.
19. From the **Destination Ports** drop-down list, select the required hybrid or tool ports.
20. From the **VLAN Tag** drop-down list, select one of the following options:
    - **None** - No VLAN tag is used; traffic is routed to a different destination.
    - **Original** - Uses the original VLAN tag from the packet received from the inline network.
    - **As Inline** - Uses the VLAN tag configured for the Flexible Inline Map.

    > **NOTE:** The **As Inline** option is available only when you configure the Resilient Inline Arrangement with a single VLAN tag.
21. Click **Deploy**, select a traffic path and click **OK**.

Create an Inline SSL APP

1. On the left pane, click the '+' icon next to the **Inline SSL APP** option.
2. Enter a name for the Inline SSL APP.
3. Enable Resilient Inline Arrangements check box.
4. Select the nodes that would be configured and their respective GigaSMART modules.
5. Select the required GigaSMART engines.
6. Under Deployment Type, set up Key chain Password:
    a. Click **Keychain Password**, and then choose either Add or Reset.
    b. If you choose to reset the Keychain Password, enter a password that is 8 to 30 characters long and contains at least one numerical character, one uppercase character, one lowercase character, and one special character.

c. Select the **Auto Login** check box to enable GigaVUE-FM,GigaVUE HC Series to unlock the keystore when the device reboots.

d. Click **OK** to save the Key chain Password.

7. Configure **Keys and Certificates**. A key in an inbound deployment can be selected only for decryption or for re-signing and re-encryption in an outbound deployment.

   a. Click **Add new keys** to open the Key page.

   b. Enter a name and description for the key.

   c. Select the required **Key Type** and **File Type**.

   d. If using PEM or PKCS12 as file type, optionally include a passphrase.

   e. You can choose to include a Passphrase for the key when you select PEM or PKCS12 as file type if required.

   f. When you choose Luna-HSM, enter the Key label for the key.

   g. Add the required **Private Key** and **Certificate**.

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

8. Select the required deployment type.
   - **Outbound** - Add the configured primary and a secondary signing Certificate Authorities (CA).
   - **Inbound** - Add a new Server Key Mapping. Enter the domain name or IP address of the server, and select the required Key Pair Alias.
   - **Hybrid** - Add a new Server Key Mapping, and a primary and a secondary signing CA.

9. The platform pushes the added keys to both nodes. To delete a key, delete it from each node individually.

10. Click **OK** to save the Inline SSL App configuration. You can configure Inline SSL App for any one of the nodes. It will be available for the second node as well.

Deploy Inline TLS/SSL Decryption Solution

In the Flexible Inline Canvas:

1. Drag and drop the required inline network or inline network bundle.
2. Drag and drop the flexible inline map into the canvas. Use the Flex Map, Inline Tools, and Inline SSL App—available on both nodes with the same alias—to configure the Flex Inline TLS/SSL maps.
3. Under the Settings option, enable the '**Show Resilient Inline Menu**' check-box and setup the Node, IB Pathway, and Hashing configurations.
4. Click **Deploy**.

Verify the Solution

- In GigaVUE-FM, open the Flexible Inline Canvas: **Physical > Orchestrated Flows > Inline Flows > Configuration Canvas**, select the target device, and open the deployed solution on the canvas. Confirm the solution shows as deployed for both nodes.
- Ensure Resilient Inline settings (node selection, hashing/IP parity) are correct and the Inter-broker Pathway (IBP) state is Up.
- Check that traffic and sessions appear on both nodes, with asymmetric flows traversing the IBP as designed.
- If Single VLAN Tag (SVT) is used, confirm it's enabled and tool-side VLANs match your policy.

What to Do Next

After deployment, you can view the Monitor and Session statistics. Refer to View Inline TLS/SSL Session Statistics Statistics. You can also view the Inline TLS/SSL Dashboards. Refer to View Inline TLS/SSL Dashboards.

Configure Thales-Luna HSM for TLS/SSL Decryption

This section describes the configuration and use of Thales-Luna HSM for TLS/SSL decryption.

Prerequsites

- From the Thales Luna HSM administrator:
  - HSM server IP address and port, and a valid partition label and partition password for the Luna partition you will use.
  - The private keys' key labels in the Luna partition (for example, for primary and secondary signing). With Luna, the private key remains in the HSM; you reference it by its unique label/handle during configuration.
  - The corresponding server certificate(s) for those key labels, and any required CA certificate(s) for trust.
- A Keychain password must be configured on the device before installing certificates and keys in the keystore.
- A static IP address to assign to the GigaSMART engine for HSM communication; static IP is recommended. When Luna HSM is selected, IP address configuration details must be entered for the Inline SSL application.

Access Flexible Inline Canvas

1. Go to  **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed Flexible Inline Canvas, select the device where you want to configure the TLS/SSL decryption.

Configure Inline Network

1. On the left pane, click the '+' icon next to **Inline Network** option to create a new entry.
2. Enter a name and description for the inline network in the **Alias** and **Description** fields. Then, click **Port Editor**.
3. In the **Alias** and **Description** fields, enter a name and description for the inline network.
4. Click **Port Editor**. In the **Quick Port Editor**, scroll down to the inline network ports that you want to configure. Select **Enable** to administratively enable inline network ports, and click **OK** to apply the changes.
5. From the **Port A** and **Port B** drop-down lists, select the ports that you want to configure as the inline network pair.
6. From the **Traffic Path** drop-down list, select **To Inline Tool**.
   - **To Inline Tool**—All traffic originating from the inline network is directed to the sequence of inline tools and inline tool groups and is guided through the inline tools and inline tool groups according to the current inline tool and inline tool group status.
7. Click **OK**. You have successfully created an Inline Network.

To explore additional options during configuration, refer to Inline Network Ports and Inline Network.

Configure Inline Tool

1. On the left pane, click the '+' icon next to the **Inline Tool** option to create a new inline tool.
2. In the **Properties** pane, in the **Alias** and **Description** fields, enter a name and description for the inline tool.
3. From the **Type** drop-down list, select one of the following options:
   - **External**- To configure a third-party tool.
   - **GigaVUE Node**- To configure a GigaVUE node as a tool.

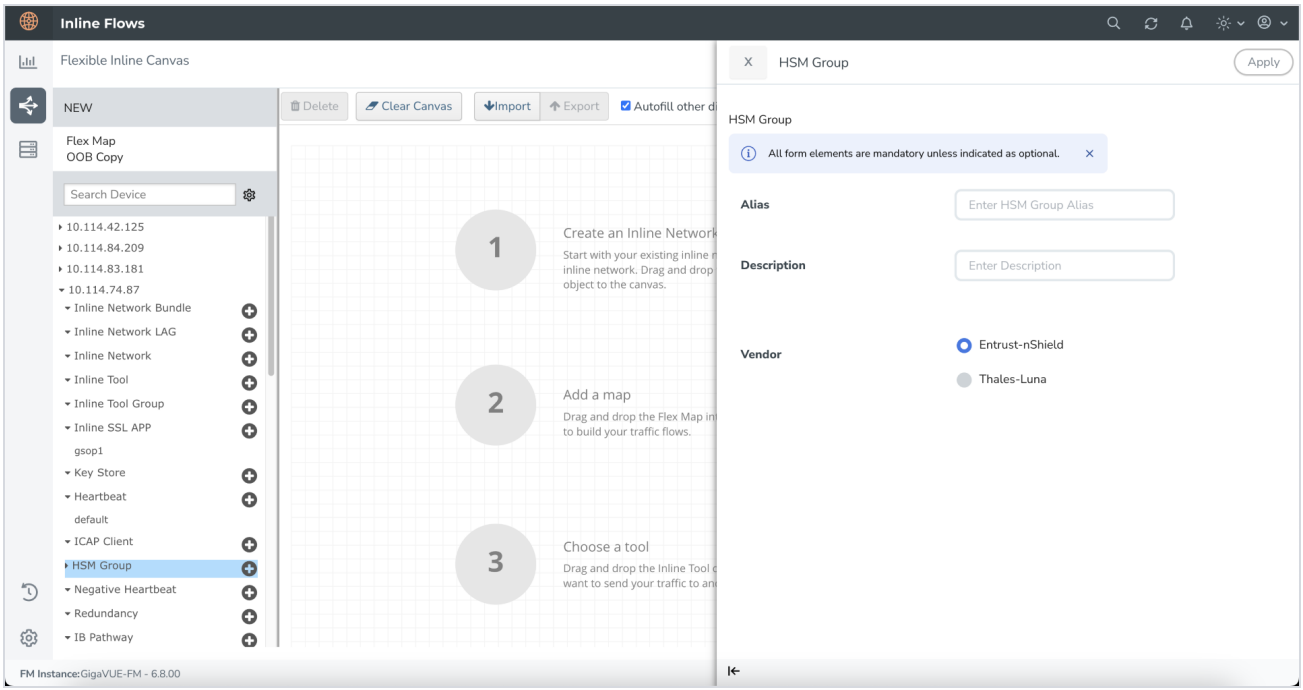
4. Click **Port Editor**. In the **Quick Port Editor**, scroll down to the inline tool ports that you want to configure. Select **Enable** to administratively enable the inline tool ports, and then click **OK** to apply the changes.
5. From the **Port A** and **Port B** drop-down lists, select the inline tool ports according to the direction the inline tool expects traffic from the network.
6. Verify that the **Enabled** check box is selected.
7. Click **OK**. You have successfully created an Inline Network.

To explore additional options during configuration, refer to Inline Tool Ports and Inline Tools.

Create HSM Group

To configure HSM Group:

1. Go to  **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed Flexible Inline Canvas, select the device where you want to configure the HSM Group.
3. Click the '+' icon next to the **HSM Group** to create a new entry.
4. Enter a name and description for the **HSM Group** in the **Alias** and **Description** fields.



5. Select the required vendor type from the available options —**Entrust-nShield** or **Thales-Luna**—to create the corresponding HSM Group.

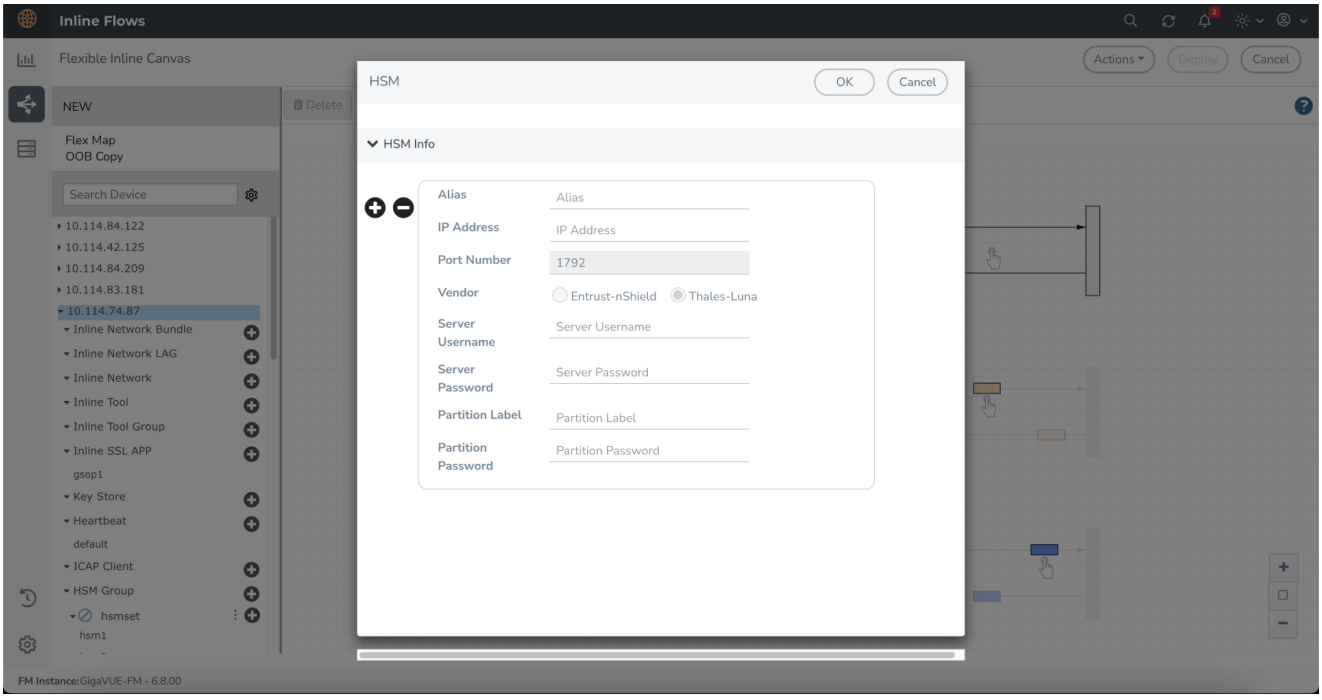
   NOTE: For Thales-Luna Network HSM, you can create a maximum of 16 HSM units per device.

6. Click **Apply** to save the configurations. All individual HSM units you create will be listed under the configured HSM Group.

Add Thales -Luna HSM Appliance to the HSM Group

To add your vendor type as Thales-Luna to your configured HSM Group follow the below steps:

1. In the Flexible Inline Canvas, select the device for which the HSM Group is configured.
2. Click the expand menu on the configured **HSM Group**.
3. Click the '+' icon next to the configured **Thales-Luna** HSM to proceed.



4. In the HSM pop-up that appears, enter a name for the HSM appliance in the **Alias** field.
5. Enter the IP address of the HSM server and **Port Number**.

Notes:

- The default port number for Thales- Luna HSM is 1792.
- It is preferable to use a static IP address to prevent the Thales Luna registration from expiring.

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

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

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

Notes

For **Thales-Luna** configuration, if the operational status of the HSM Group shows '**Registration pending**', follow the below steps to register the HSM client (GigaSMART engine) with the HSM server:

1. Register the HSM Client (GigaSMART engine's IP address) on the HSM server.

Create an Inline SSL APP and Attach the HSM Group

1. On the left pane, click the '+' icon next to the **Inline SSL APP** option.
2. Enter a name for the Inline SSL APP and select the required GigaSMART engines.
3. From the **HSM Group** drop-down list, select the configured **hsm-luna** Group alias.
4. Under Deployment Type, set up Key chain Password:
   a. Click **Keychain Password**, and then choose either Add or Reset.
   b. If you choose to reset the Keychain Password, enter a password that is 8 to 30 characters long and contains at least one numerical character, one uppercase character, one lowercase character, and one special character.
   c. Select the **Auto Login** check box to enable GigaVUE-FM,GigaVUE HC Series to unlock the keystore when the device reboots.
   d. Click **OK** to save the Key chain Password.

5. Configure **Keys and Certificates**. A key in an inbound deployment can be selected only for decryption or for re-signing and re-encryption in an outbound deployment.
   a. Click **Add Keys** to open the Key page.
   b. In the **Key Alias** and **Description** fields, enter a name and description for the SSL key.
   c. For Key Type, select either **RSA** or **ECDSA**.
   d. From the **File Type** drop-down list, select **Luna-HSM**.
   e. Choose one of the following methods to import the required key token:
      - **Install from URL**—Enter a valid directory path including the file name and enter the password to access the server.
      - **Install from Local Directory**—Browse and select the key handler file from your local directory.
   f. Choose one of the following methods to import the corresponding certificate:
      - **Install from URL**—Enter a valid directory path including the file name and enter the password to access the server.
      - **Install from Local Directory**—Browse and select the certificate file from your local directory.

   NOTE: You can obtain the key token and the corresponding certificate from your HSM administrator.

   g. Click **OK**. You have successfully created the Primary key.
   h. Repeat **steps 1-6** to configure the secondary signing certificate and private key.
6. Thales-Luna HSM configuration is supported in Inbound, Outbound, and Hybrid deployment types. Select the required deployment type. Refer to TLS/SSL Sessions for more details.
   - **Outbound** - Add the configured primary and a secondary signing Certificate Authorities (CA).
   - **Inbound** - Add a new Server Key Mapping. Enter the domain name or IP address of the server, and then select the required Key Pair Alias.
   - **Hybrid** - Add a new Server Key Mapping, and a primary and a secondary signing CA.
7. Select Outbound as your deployment type and add the configured primary and a secondary signing Certificate Authorities (CA).
8. Under **Configurations**, for Default Action select one of the following options:
   - **Decrypt**—Decrypt all the traffic that is guided into the Inline SSL APP.
   - **No Decrypt**—Do not decrypt the traffic that is guided into the Inline SSL APP.
9. For **URL Cache Miss Action**, select one of the following options

- **Decrypt**—Decrypt all the traffic that is guided into the Inline SSL APP.
- **No Decrypt**—Do not decrypt the traffic that is guided into the Inline SSL APP.
- **Defer**—Delay the decryption until the Defer Timeout seconds provided.

10. Under **Security Expectations**, choose to either decrypt or drop the traffic for the following certificates:
    - Self-signed certificate
    - Unknown CA certificate
    - Invalid certificate
    - Expired certificate
11. Under **Network Access**, select IP Address as the network access configuration mode. Enter the IP Address, Netmask, Gateway, DNS, MTU, and VLAN.

    Notes:

    - DHCP mode is not supported when configuring Thales-Luna HSM, as IP changes can break connectivity if the new address isn't registered with the HSM server
    - Ensure that IP address you are using is already registered in the HSM server.

12. Under **Trust Store**, set up Set up the Signing Certificate Authority (CA).
    - Click **Append** and choose one of the following methods:
      - **Copy and Paste** - Directly paste the Trust Store Certificate.
      - **Install from Local Directory** - Browse and select the certificate file from your local directory.
    - Click **Replace** to update the existing certificate with a new one.
13. Click **OK**. You have successfully created the Inline SSL APP.

Refer to Inline SSL App—Field References for any additional details when configuring the Inline SSL APP.

Deploy the Inline SSL Solution

1. Drag and drop the required inline network or inline network bundle into the flexible inline canvas.
2. Drag and drop the flexible inline map into the canvas.
3. In the Properties pane, click **Add Rule**, and add **Bi-directional** as the rule condition.
4. In the rule description, add protocol as TCP. Add the required rules for the inline map, and then click OK to save the configuration.
5. Drag and drop the Inline SSL APP into the canvas.

6. Drag and drop the required inline tools or inline tool group into the canvas.
7. Click **Deploy**.

Register GigaSMART Engine on Thales Luna HSM Server

After deploying the Gigamon configuration, the Luna HSM administrator must register the GigaSMART engine (client) on the Luna HSM using the GigaSMART IP you configured in Network Access.

If client registration is not completed within the required time window (10 minutes), operational status can move to a registration timeout state. Work with your HSM admin to complete registration on the Luna HSM side. Refer to Client Register - Luna Command Reference for more details.

Verify the Solution

- Confirm the solution is Deployed on the Flexible Inline Canvas **(Physical > Orchestrated Flows > Inline Flows > Configuration Canvas**) and the Inline SSL app is present in the canvas.
- Verify HSM group, connectivity, and trust
  - InGigaVUE-FM,GigaVUE HC Series, verify the HSM Group exists and is associated to the target node(s).
  - Confirm the Thales Luna HSM appliance/client is registered and reachable; check that the Luna partition/slot assignment is correct and shows a healthy/connected state.
  - Ensure you are not mixing HSM vendors on the same device (do not combine Entrust and Thales).
- Confirm that the Inline SSL app is using Thales Luna HSM-backed keys (not local files), the Keychain Password is unlocked, and the keys/certificates load without errors or HSM/keystore alerts so they are available for decryption
  - In the Inline SSL App's keys/certificates area:
    - Verify each required key/cert shows Thales Luna HSM as the source (HSM-backed), not local/PEM/PKCS12.
    - Ensure the GigaVUE-FM,GigaVUE HC Series Keychain Password is unlocked so the app can access the keystore.
  - Confirm keys/certs load without errors (no red badges or "key not found/permission denied/keystore locked" messages).
  - Check Events/Alerts for any HSM-related issues (e.g., partition login failures, HSM connectivity, permission errors). There should be no HSM/keystore alerts.

What to Do Next

After deployment, you can view the the statistics details , refer to View HSM Statistics.

Modify an HSM Deployment for Inline TLS/SSL Decryption

When modifying an HSM Decryption deployment, follow the steps below to ensure a smooth and secure transition:

1. **Switch to Bypass Mode**

   Move the Inline Network traffic path to Bypass mode to prevent disruption during the configuration change.

2. **Apply Deployment Changes**

   Make the required deployment modifications. Supported transitions include:
   - From non-HSM based decryption to Thales Luna HSM based decryption
   - From HSM-based decryption to non-HSM based decryption

3. **Reboot GigaSMART Card**

   After applying the changes, reboot the GigaSMART card to apply the new configuration.

4. **Switch to Inline Mode**

   Move the Inline Network traffic path from Bypass mode to '**To Inline Tool**' mode to resume normal operation.

Notes:

- Ensure that at least one active Luna HSM is available in the High Availability (HA) group to avoid decryption service interruption.
- If a partition fails in the Luna HSM configuration but at least one active partition remains, do not reload the device.
- The network connection between the HSM and GigaSMART must use a static IP address. Do not use DHCP, as the IP address must remain consistent for proper communication.

Configure ICAP Client for Inline TLS/SSL Decryption Solution

This section provides topics on configuring and using **ICAP Client** for Inline TLS/SSL Decryption Solution.

Prerequistes

- Allocate two separate GigaSMART engines—one dedicated to Inline SSL and one to ICAP—and if HTTP/2 downgrade is required in Inline SSL, use Gen3 devices for both engines.
- Rely on the system to add the GigaSMART Group automatically for ICAP when the port is attached to the ICAP Client, so you do not need to add it during IP Interface creation.
- Gather the ICAP server details, including IP/hostname and listening port (typically 1344), plus any required service URIs (REQMOD, RESPMOD, OPTIONS).
- Decide on the deployment type—Same Node or Different Node—and for Same Node be prepared to select the Inline SSL Inline Tool as the ICAP Client's Source port, while for Different Node you must ensure a network interconnect carries the clear-text feed from the Inline SSL node to the ICAP node's Inline Network.
- Verify compatibility constraints beforehand, because the ICAP Client cannot be deployed inline with One-Arm iSSL, L3 NAT iSSL, or RIA iSSL.

Access Flexible Inline Canvas

1. Go to  **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed Flexible Inline Canvas, select the device where you want to configure the TLS/SSL decryption.

Configure Inline Network

For an Inline SSL with ICAP deployment, create the Inline SSL constructs (one Inline Network with NA/NB legs and an Inline Tool) and a separate Inline Network for ICAP, then configure an IP Interface on a tool port with the required IP/subnet/gateway/VLAN to reach the ICAP server.

1. On the left pane, click the '+' icon next to **Inline Network** option to create a new entry.

2. Enter a name and description for the inline network in the **Alias** and **Description** fields. Then, click **Port Editor**.
3. In the **Alias** and **Description** fields, enter a name and description for the inline network.
4. Click **Port Editor**. In the **Quick Port Editor**, scroll down to the inline network ports that you want to configure. Select **Enable** to administratively enable inline network ports, and click **OK** to apply the changes.
5. From the **Port A** and **Port B** drop-down lists, select the ports that you want to configure as the inline network pair.
6. From the **Traffic Path** drop-down list, select **To Inline Tool**.
   - **To Inline Tool**—All traffic originating from the inline network is directed to the sequence of inline tools and inline tool groups and is guided through the inline tools and inline tool groups according to the current inline tool and inline tool group status.
7. Click **OK**. You have successfully created an Inline Network.

To explore additional options during configuration, refer to Inline Network Ports and Inline Network.

Configure Inline Tool

1. On the left pane, click the '+' icon next to the **Inline Tool** option to create a new inline tool.
2. In the **Properties** pane, in the **Alias** and **Description** fields, enter a name and description for the inline tool.
3. From the **Type** drop-down list, select one of the following options:
   - **External**- To configure a third-party tool.
   - **GigaVUE Node**- To configure a GigaVUE node as a tool.
4. Click **Port Editor**. In the **Quick Port Editor**, scroll down to the inline tool ports that you want to configure. Select **Enable** to administratively enable the inline tool ports, and then click **OK** to apply the changes.
5. From the **Port A** and **Port B** drop-down lists, select the inline tool ports according to the direction the inline tool expects traffic from the network.
6. Verify that the **Enabled** check box is selected.
7. Click **OK**. You have successfully created an Inline Network.

To explore additional options during configuration, refer to Inline Tool Ports and Inline Tools.

Configure IP Interface

1. On the left pane, click the '+' icon next to the **IP Interface** option to create a new IP Interface.
2. In the **Alias** and **Description** fields, enter a name and description for the IP Interface.
3. From the **Port** drop down list, select the port that is connected to Network and has reach ability to the ICAP server.
4. Select the **Type** of the IP Interface that you want to configure.
5. Enter an **IP Address**. For example, 192.168.1.20.
6. Enter an **IP Mask**. For example, 255.255.255.0.
7. Enter a **Gateway**. For example, 192.168.1.20.
8. Enter the Maximum Transmission Unit (MTU) for this port in the **MTU** field. For example, 1500.
9. Click **OK** to save the configuration.

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

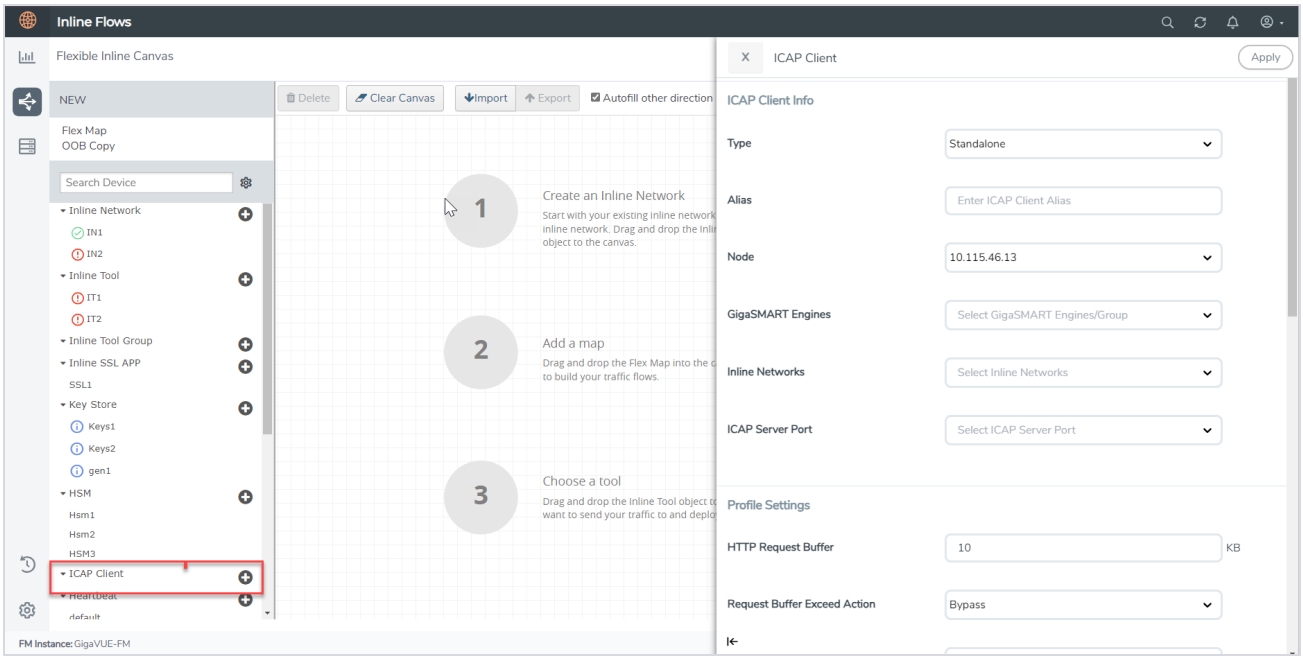
Configure ICAP Client

In an Inline TLS /SSL Decryption Solution with an **ICAP Client** the main choice for your ICAP configuration is :

- Type should be either Same Node or Different Node

To configure the ICAP Client:

1. Go to  **> Physical > Orchestrated Flows > Inline Flows > Configuration Canvas** to create a new Flexible Inline Canvas.
2. In the displayed Flexible Inline Canvas, select the device where you want to configure the ICAP Client.
3. Click the '+' icon next to the **ICAP Client** to create a new entry.
4. Enter a name and description for the inline network in the **Alias** and **Description** fields.



5. In the **ICAP Client** properties pane that appears on the right, complete the required fields in the ICAP Client Info, Profile Settings, and the Server sections. Refer to Configure ICAP Client for Inline TLS/SSL Decryption Solution for more details.
6. Select one from the following:
   - If ICAP and iSSL are integrated and deployed on the same node, select the **Same Node** option.
   - If ICAP and iSSL are integrated and deployed on different node, select the **Different Node** option.
7. Select the required tool port of Inline SSL profile, which is connected to the inline network of ICAP. (This option will not appear if you select standalone as type.)
8. Click **Apply** to save the configurations.

Create an Inline SSL APP

1. On the left pane, click the '+' icon next to the **Inline SSL APP** option.
2. Enter a name for the Inline SSL APP and select the required GigaSMART engines.
3. ICAP configuration is supported in Inbound, Outbound, and Hybrid deployment types. Select the required deployment type. Refer to TLS/SSL Sessions for more details.
   - **Outbound** - Add the configured primary and a secondary signing Certificate Authorities (CA).
   - **Inbound** - Add a new Server Key Mapping. Enter the domain name or IP address of the server, and then select the required Key Pair Alias.
   - **Hybrid** - Add a new Server Key Mapping, and a primary and a secondary signing CA.

4. Under **Configurations**, for Default Action select one of the following options:
   - **Decrypt**—Decrypt all the traffic that is guided into the Inline SSL APP.
   - **No Decrypt**—Do not decrypt the traffic that is guided into the Inline SSL APP.
5. For **URL Cache Miss Action**, select one of the following options:
   - **Decrypt**—Decrypt all the traffic that is guided into the Inline SSL APP.
   - **No Decrypt**—Do not decrypt the traffic that is guided into the Inline SSL APP.
   - **Defer**—Delay the decryption until the Defer Timeout seconds provided.
6. Under **Security Expectations**, choose to either decrypt or drop the traffic for the following certificates:
   - Self-signed certificate
   - Unknown CA certificate
   - Invalid certificate
   - Expired certificate
7. Under **Network Access**, select DHCP or IP Address as the network access configuration mode. If you select IP Address as the mode, enter the IP Address, Netmask, Gateway, DNS, MTU, and VLAN.

   NOTE: It is recommended to use static IP address option when configuring ICAP Client.
8. Click **OK**. You have successfully created the Inline SSL APP.

To explore additional options during configuration, refer to Inline SSL App—Field References.

Deploy the Inline SSL Solution

1. Drag and drop the required inline network configured for Inline TLS/SSL into the flexible inline canvas.
2. Drag and drop the Flexible Inline Map into the canvas, then in the Properties pane, click **Add Rule** and set the rule condition to **Bi-directional**.
3. In the rule description, add the protocol as TCP. Add the required rules for the inline map, and then click OK to save the configuration.
4. Drag and drop the Inline SSL APP into the canvas.
5. Drag and drop the required inline tools configured for Inline TLS/SSL into the canvas.
6. Click **Deploy**. You get two options to select your Traffic path during Deployment . It could be either **Logical Bypass** or **Keep as is.** Select an option based on your requirement and then click **OK**.

If the configuration is correct and the source ports of ICAP are properly linked to the Inline Tool port of Inline SSL, the canvas will display the inline tool indicating the ICAP client.

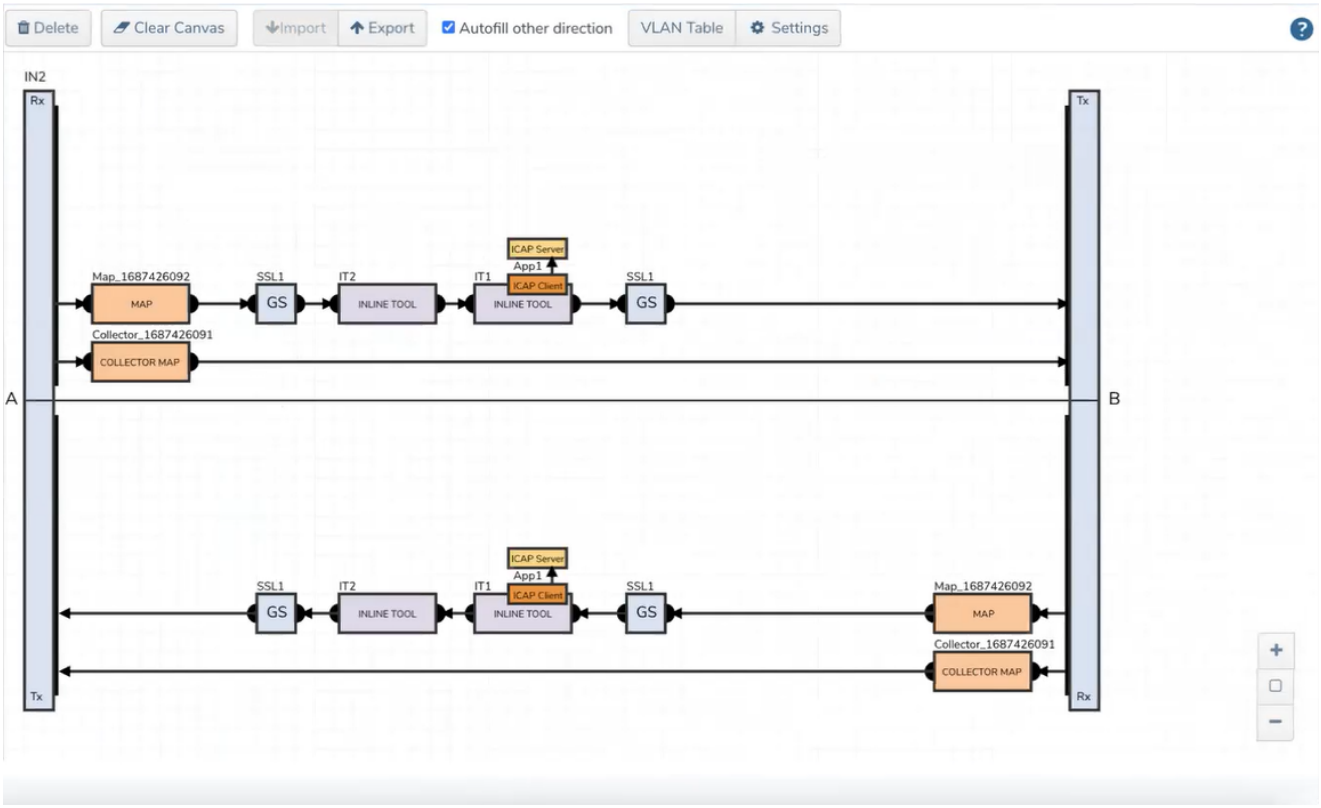


Figure 70 *ICAP Client APP—Deployed*

Verify the Solution

- Confirm the solution shows as Deployed on the Flexible Inline Canvas (**Physical > Orchestrated Flows > Inline Flows > Configuration Canvas**); ensure both Inline Networks exist: one attached to the Inline SSL app and a separate one attached to the ICAP client.
- Verify ICAP client wiring and health
  - Open the ICAP client properties and confirm:
    - The ICAP client is attached to the intended Inline Network (the ICAP-specific inline network you created).
    - Source Port/Interface points to the correct inline or tool port that reaches the ICAP server (as per your design).
  - If you created an IP Interface on a tool port to reach the ICAP server, verify:
    - IP/subnet/gateway/VLAN are correct and the interface shows Up.
    - Basic reachability to the ICAP server (no Events/Alerts indicating ARP/route failures).

What to Do Next

After deployment, you can view statistics details, refer to View Inline Solution Status and Statistics.

View Inline Solution Status and Statistics

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 inline flows has the following tabs:

- Status
- Statistics

Refer to the following sections for more details:

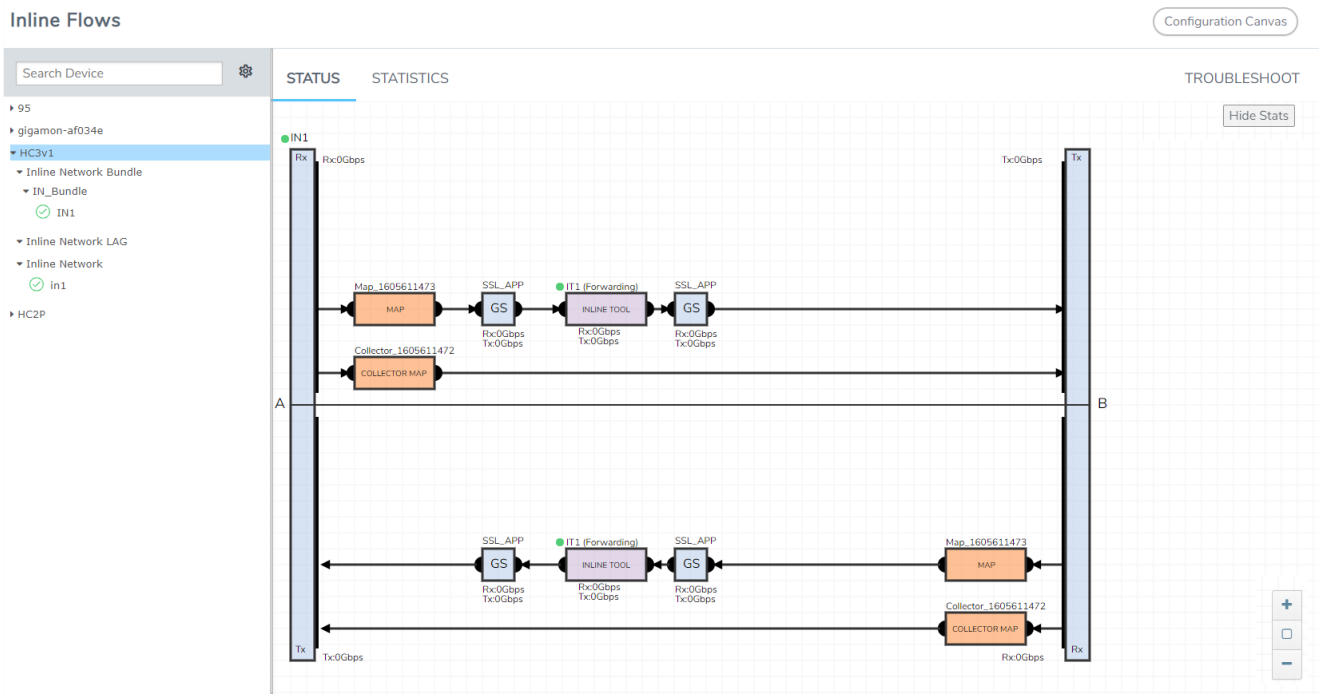
- View Inline Solution Status
- View Inline TLS/SSL Session Statistics
- View Inline TLS/SSL Monitor Statistics
- View Inline TLS/SSL Certificate Statistics
- View HSM Statistics
- View ICAP Statistics

View Inline Solution 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.

It also provides the status of the components that are part of the selected inline network. Hover over the status to view the description.

Click the **Show Stats/Hide Stats** toggle button to view the Rx/Tx rate for the components that are part of the flexible inline flow deployment. Refer to the following figure for details.



Rules and Notes

Keep in mind the following points when you are viewing the details in the Status tab:

- The ability to view the Rx/Tx rate for the components is supported only for GigaVUE-HC1, GigaVUE-HC3 and GigaVUE-HC2 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.

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.

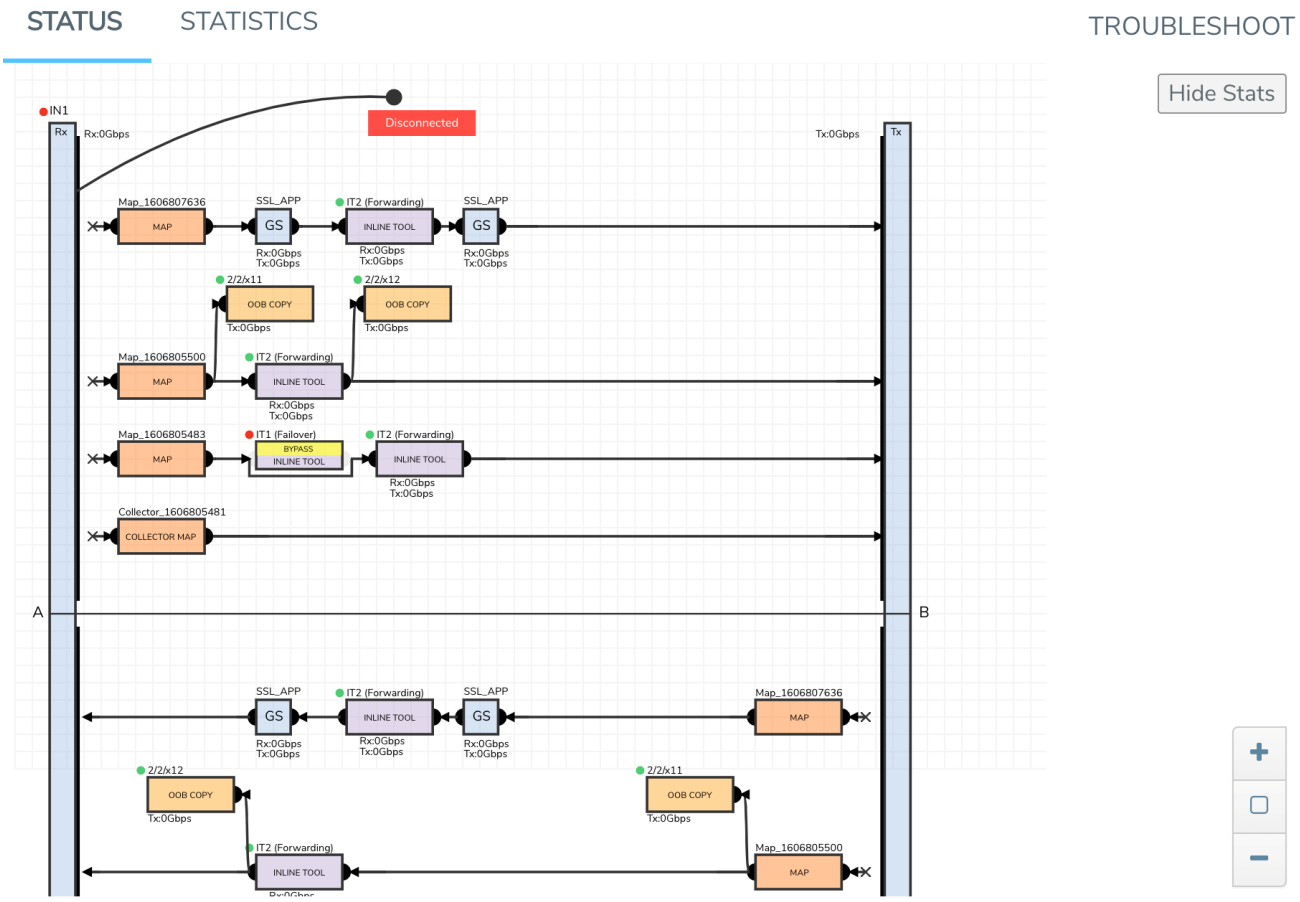


Figure 71 *Inline Network Forwarding States*

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 12: 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-FM,GigaVUE HC Series node acting as a wire or fiber)
disable	**traffic path set to drop**	any combination of far-end ports status	any combination of operational state of the inline tools or inline tool groups involved in the maps originating from the inline network	DISABLED	all traffic arriving at the inline network ports is dropped
disable	traffic path set to bypass, monitoring, or to-inline-tool	**at least one far-end port is down**	any combination of operational state of the inline tools or inline tool groups involved in the maps originating from the inline network	DISCONNECTED	No traffic is exchanged between the nodes

Inline Network Physical Bypass	Inline Network Traffic Path	Far-End Status of Links Connected to Inline Network Ports	Operational State of *Not* Forced Inline Tools and Inline Tool Groups Involved in Maps from the Inline Network	Forwarding State	Description
disable	**traffic path set to bypass**	both far-end ports are up	any combination of operational state of the inline tools or inline tool groups involved in the maps originating from the inline network	FORCED BYPASS	All traffic that matches any of the maps originating from the inline network is redirected through a logical bypass
disable	**traffic path set to monitoring**	both far-end ports are up	any combination of operational state of the inline tools or inline tool groups involved in the maps originating from the inline network	FORCED BYPASS WITH MONITORING	A copy of the traffic originating from the inline network bypasses the sequence of inline tools and inline tool groups and is re-directed to the opposite-side inline network port. Another copy of the traffic is directed to the sequence of inline tools and inline tool groups, except that no traffic of the second copy is sent to the exit port.
disable	**traffic path set to to-inline-tool**	both far-end ports are up	all inline tools involved (directly or indirectly as members of inline tool groups) in the maps originating from the inline network are in the *up* operational state	NORMAL	The traffic is guided between the source inline network port and the destination inline network port according to the status of the inline tools and inline tool groups. **NOTE:** The state of all inline tools must

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
					be *up*, including inline tools configured as spare in an inline tool group, inline tools or inline tool group members in the **a-to-b** and **b-to-a** lists configured with any traffic path other than to-inline-tool.
disable	**traffic path set to to-inline-tool**	both far-end ports are up	at least one of the inline tools or inline tool groups involved in the maps originating from the inline network configured with the traffic path parameter to-inline-tool and failover action of network-port-forced-down is in the *down* operational state	NETWORK PORTS FORCED DOWN	No traffic is exchanged between the 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	• **none of the inline tools or inline tool groups involved in the maps originating from the**	FAILURE INTRODUCED DROP	All traffic arriving at the inline network ports is dropped

Inline Network Physical Bypass	**Inline Network Traffic Path**	**Far-End Status of Links Connected to Inline Network Ports**	**Operational State of *Not* Forced Inline Tools and Inline Tool Groups Involved in Maps from the Inline Network**	**Forwarding State**	**Description**
			inline network configured with to-inline-tool and failover action network-port-forced- down is in the *down* operational state • **at least one of the inline tools or inline tool groups involved in the maps originating from the inline network configured with to-inline-tool and failover-action of network-drop is in the *down* operational state**		
disable	traffic path set to to-inline-tool	both far-end ports are up	• none of the inline tools or inline tool groups involved in the maps originating from the inline	FAILURE INTRODUCED BYPASS	All traffic that matches any of the maps originating from the inline network is redirected through a logical bypass

Inline Network Physical Bypass	Inline Network Traffic Path	Far-End Status of Links Connected to Inline Network Ports	Operational State of *Not* Forced Inline Tools and Inline Tool Groups Involved in Maps from the Inline Network	Forwarding State	Description
			network configured with to-inline-tool and failover action of network-port-forced-down or network-drop is in the *down* operational state • at least one of the inline tools or inline tool groups involved in the maps originating from the inline network configured with to-inline-tool and failover action of network-bypass is in the *down* operational state		
disable	traffic path set to to-inline-tool	both far-end ports are up	**any combination of conditions not listed for the forwarding state definitions of PHYSICAL BYPASS, DISABLED,**	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

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
			DISCONNECTED, FORCED BYPASS, FORCED BYPASS WITH MONITORING, NORMAL, NETWORK PORTS FORCED DOWN, FAILURE-INTRODUCED DROP, or FAILURE-INTRODUCED BYPASS		tool involved in flexible inline maps (directly or indirectly as a member of an inline tool group) is in the *down* operational state and there is no network-level failover action in effect, the inline network is in the ABNORMAL state.

NOTE: When the Inline Network traffic path is set to monitoring and the Inline Tool Failover action is configured as 'Network Port Force Down,' the Inline Tool will continue to receive traffic, even if it is disabled or in an operational down state. This behavior is expected.

View Inline TLS/SSL Session 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.

To display the inline TLS/SSL summary details, go to **GigaSMART > Inline SSL > Session Statistics**, view the Summary details under **Summary** tab which is displayed initially. Refer to Figure 72Inline TLS/SSL Session Statistics in GigaVUE-FM.

There are four sections: Session Statistics, Performance Statistics, Policy Statistics, and Certificate Statistics. Click **Show Summary** to view these sections. Click **Clear Session Summary** to clear all the displayed summary details.

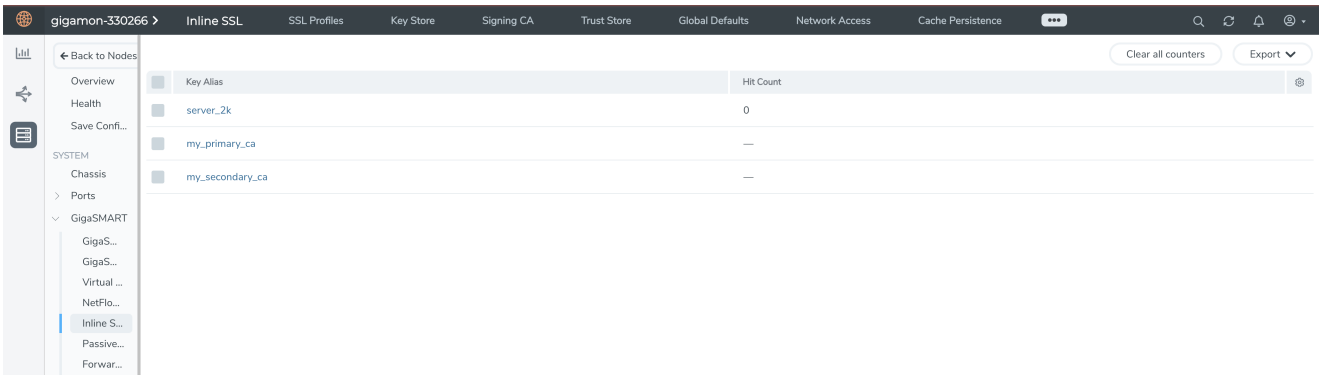


Figure 72 *Inline TLS/SSL Session Statistics in GigaVUE-FM*

To view the inline TLS/SSL session details, go to **GigaSMART > Inline SSL > Session Statistics**. Click on the **Sessions** tab. The list of available sessions will be displayed. To search and filter the session details, click **Filter** and enter an IPv4 source or destination, an L4 port source or destination, or a host name.

View Inline TLS/SSL Monitor Statistics

To display monitor statistics, go to **GigaSMART > Inline SSL > Monitor Statistics**. Select the required GigaSMART engine from the drop-down menu.

There are three sections. The first section, which has a graph for INTERFACE TRAFFIC and Interface Packet statistics, is displayed initially. To return to this display, click the small graph, TOTAL INCOMING PACKETS. Refer to Figure 73 Inline TLS/SSL Session Monitor—Interface Packet Statistics.

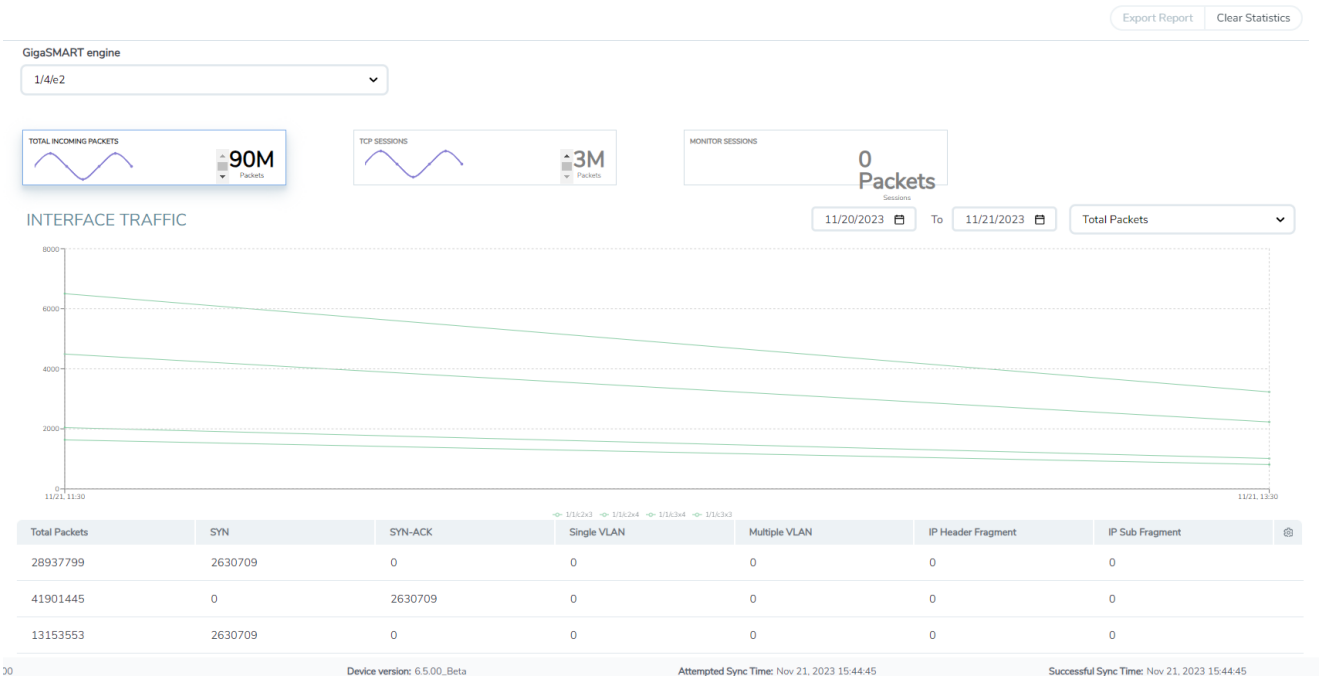


Figure 73 *Inline TLS/SSL Session Monitor—Interface Packet Statistics*

To display the graph and statistics for TCP Sessions, click the small graph, TCP SESSIONS. Refer to Figure 74 Inline TLS/SSL Session Monitor—TCP Sessions.

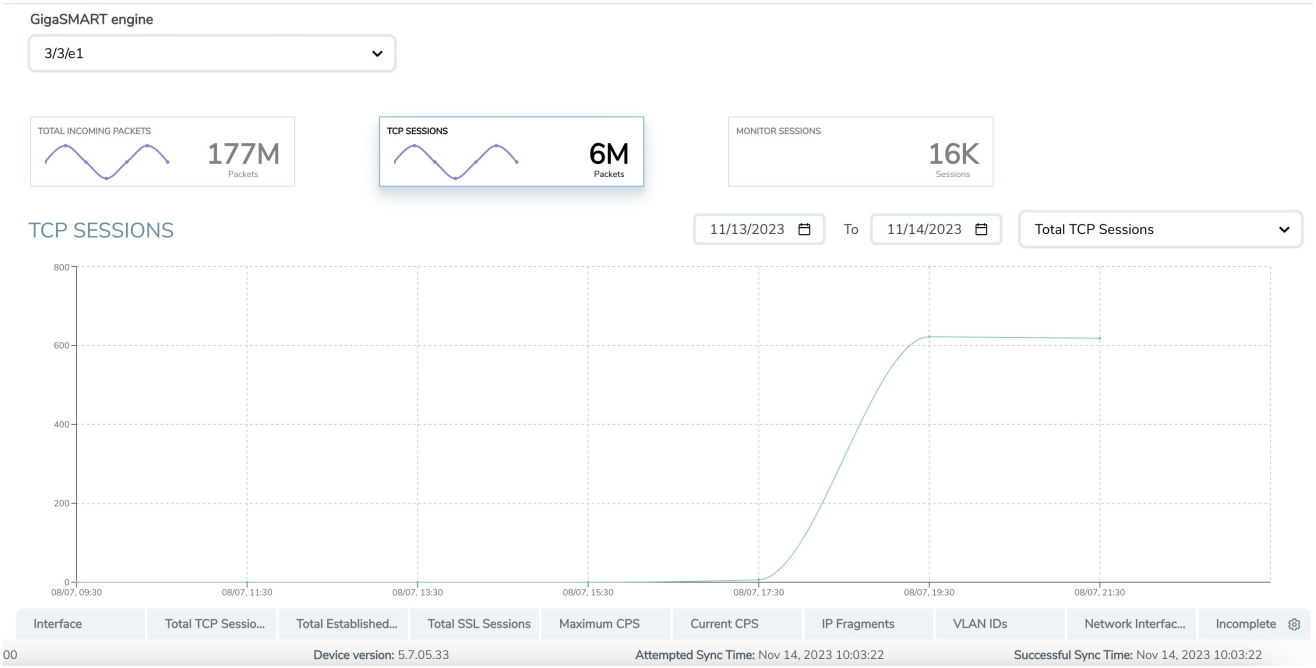


Figure 74 *Inline TLS/SSL Session Monitor—TCP Sessions*

To display Monitor Sessions, click the small graph, MONITOR SESSIONS. Refer to Figure 75 Inline TLS/SSL Session Monitor—Monitor Sessions.

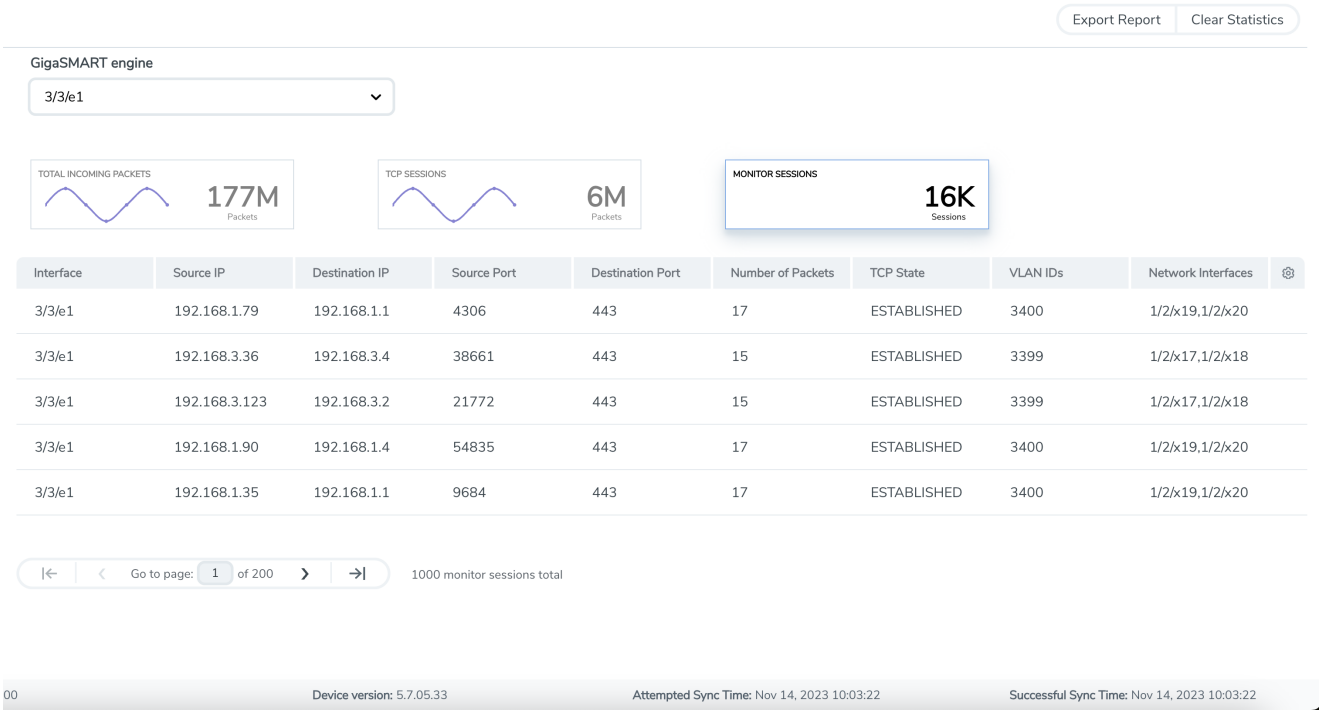


Figure 75 *Inline TLS/SSL Session Monitor—Monitor Sessions*

Click **Clear Statistics** to clear all the displayed statistics details.

View Inline TLS/SSL Certificate Statistics

To view certificate statistics, go to **GigaSMART > Inline SSL > Certificate Statistics**. The page displays the hit count of each key store certificate and allows to track whether the certificate is actively used or not. The hit count is numbered only if the policy is set for decryption. If the policy is set to no-decryption or if the deployment is outbound, the hit count will not be considered.

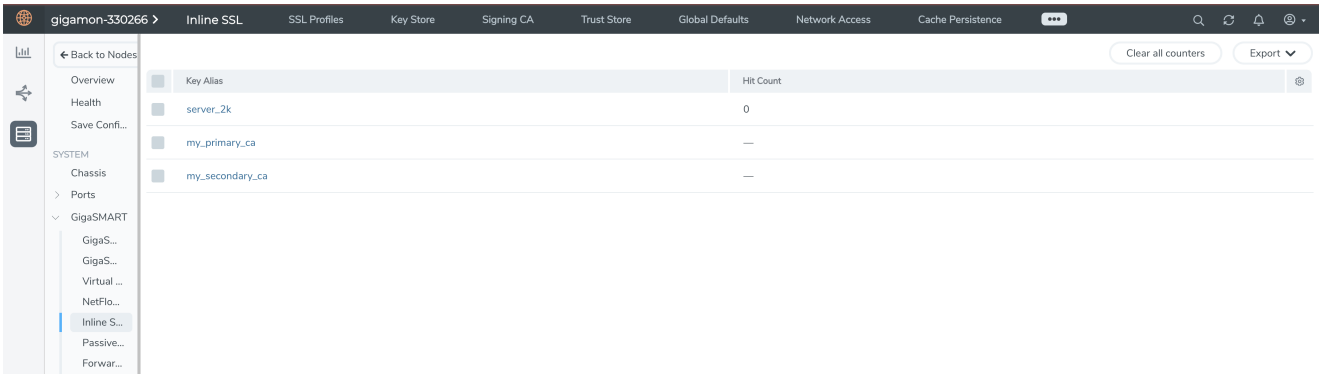
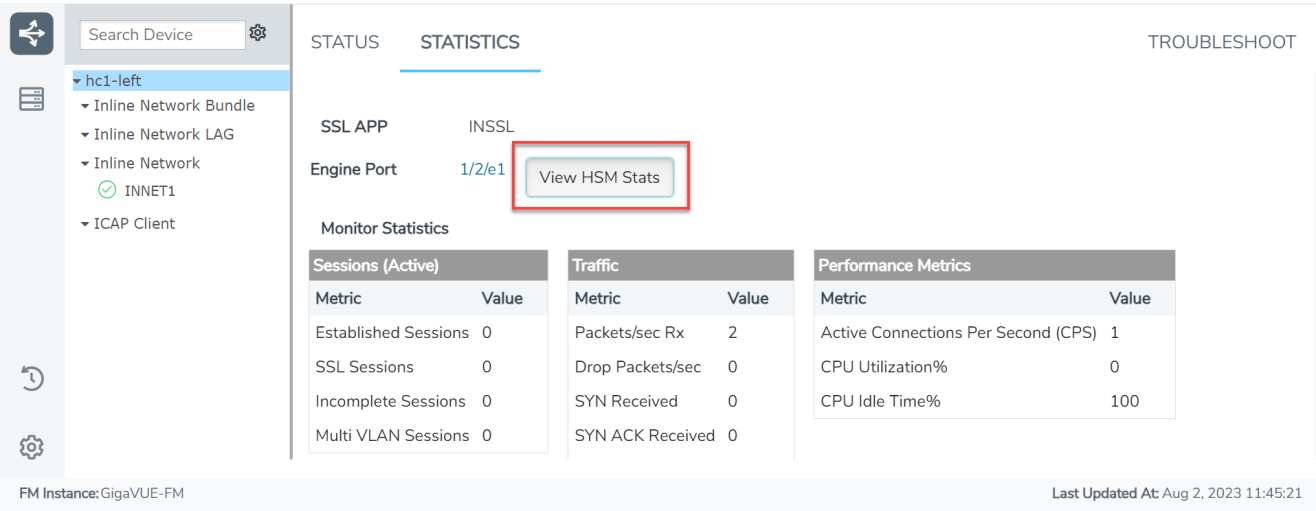


Figure 76 *Inline SSL - Certificate Statistics*

Click **Clear all Counters** to clear the hit counter of all the certificates. Click **Export** to export the available hit count details.

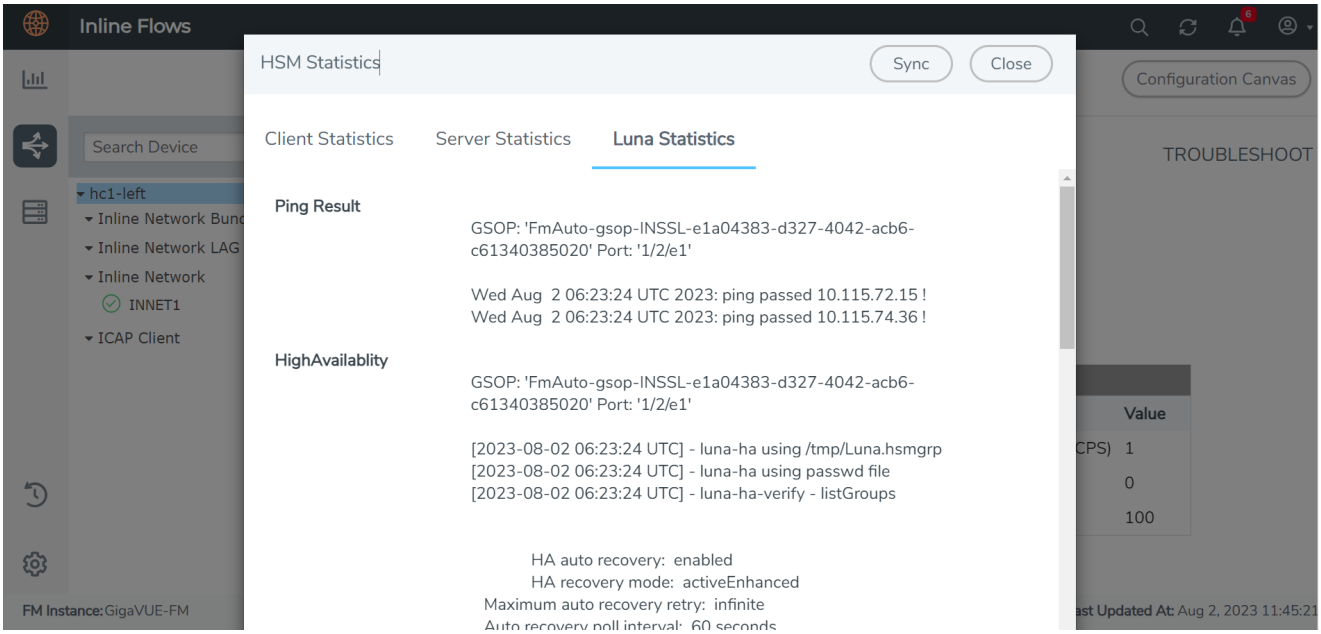
View HSM Statistics

1. In the **Inline Flows** page, select the required device for which you want to view the statistics.
2. Click **Statistics > View HSM Statistics**. You can view the following statistics details for **Entrust nShield HSM** under Client Statistics and Server Statistics:
   - Number of Request(s) Received
   - Number of Response(s) Sent
   - Metric values
   - Error details
   - Number of Request(s) Sent
   - Number of Response(s) Received
   - Delay time
   - Average Round Trip Time



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.
- **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.



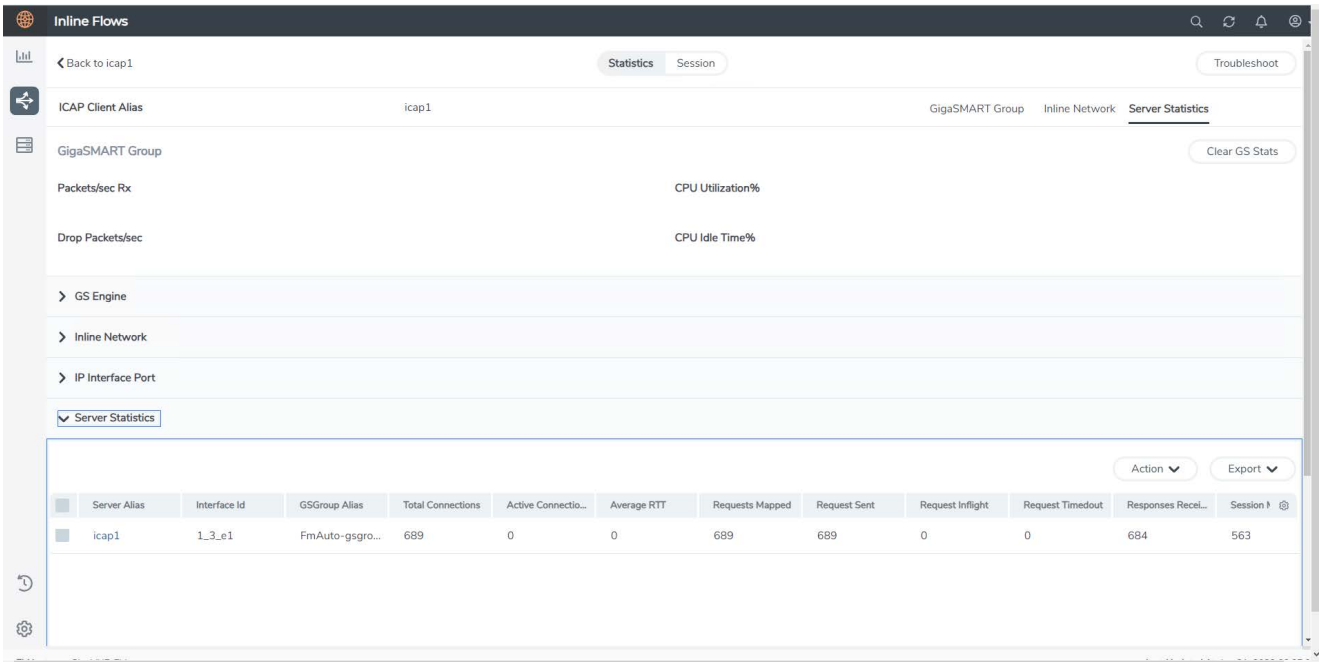
View ICAP Statistics

The Show Stats option will be displayed after the ICAP Client app is configured and deployed. The Show Stats option has the following tabs:

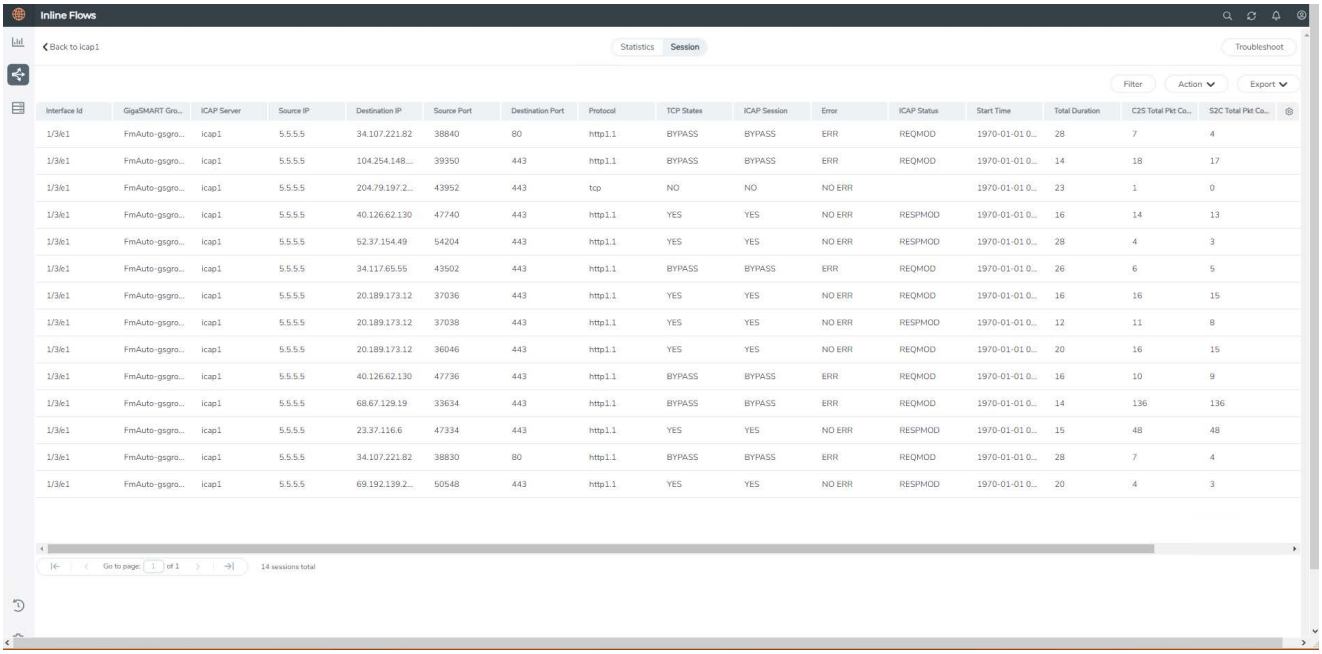
- Statistics
- Session

The **Statistics** tab provides statistical information on the GigaSMART group, GS engines, inline network, IP interface port, and server statistics. Click the drop-down menu of each component listed on the statistics page to know more about the configuration details. You can also view the statistics from the sidebar on the Inline Flows page.

The GS Engine drop-down menu provides information on the ICAP session statistics. Click **Clear GS Stats** to clear the statistics of the GS group.



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



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

Field	Description
Filter	To view the statistics of a particular IP with selected parameters.
Action	**Upload to server**: To export the log details to an external server.
Export	To export the current session details to local.

View Inline TLS/SSL Dashboards

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

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

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

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

Access TLS/SSL Dashboards

To access the dashboard:

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

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

- Basic Dashboards
- Advanced Dashboards

Basic Dashboards

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

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

The following are the basic dashboards and its visualizations:

Table 13: Session Overall Dashboard-Displays visualizations on the overall details of encrypted traffic.

Visualizations	Details
Total Intercepted Sessions	Displays overall count of intercepted sessions by the node over time period. This page does not display per engine unless specified by a filter.
Sessions Trend	Displays the trend of all Inline TLS/SSL session that has been received over a specified time period and per specified GigaSMART engines. The trend included the visualization of the Intercepted/Decrypted/Non-SSL Sessions.
Average Decryption Rate	Displays the average rate of Inline TLS/SSL sessions that have been decrypted over the specified time period..
Average CPU	Displays an average CPU utilization of all engines in Session Overall Page.
Client TLS Version Trend	Displays an overview of the incoming traffic's TLS version of the incoming data.

Visualizations	Details
Server TLS Version Trend	Provides an insight into the TLS version distribution at the server side.
Policy based Intercepted Session	Displays the trend of decryption status of Inline TLS/SSL session based policy.
Intercepted Sessions By Policy Rules	Displays the trend of Inline TLS/SSL session based on Policy rules such as; Domain ,Category, Issuer, URL Cache Miss, Network and Default. **NOTE:** The no. of sessions that gets matched to a Network Policy Rule will not be displayed in the Total Intercepted Session widget.

Table 14: Session Engine Overview Dashboard-Displays visualizations related to Inline TLS/SSL Sessions per Engine

Visualizations	Details
Sessions Rate per Engine	Displays the rate at which Inline TLS/SSL sessions are intercepted per engine.
Average Decryption Rate per Engine	Displays the average rate of sessions that got decrypted per engine.
Average CPS per Engine	Displays the average Connections per Second (CPS) performance metric per engine.
Average CPU per Engine	Displays the average CPU utilization per engine.
Engine Metric Table	Displays the Decryption rate per engine, average CPS and average CPU rate in a tabular format. The details are displayed as Host Name/Engine ID. For example; FHA-HC1 (Host Name)_1/3/e11(GSEngine ID)

Table 15: Traffic Insights Dashboard-Displays visualizations related to the traffic that is handled with a Inline TLS/SSL sessions

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

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

Table 16: Certificate Insights Dashboard-Provides mTLS-focused certificate visibility for Inline TLS/SSL sessions, showing how server and client certificates—and their revocation status

Visualizations	Details
Server Certificate Trends	Displays how server certificate usage and status evolve over time across Inline TLS/SSL sessions.
Client Certificate Trends	Displays how client certificate usage and status evolve over time for mTLS enabled Inline TLS/SSL sessions.
Server Certificate OCSP Trend	Displays the trend of OCSP revocation-check results for server certificates (for example, good, unknown, or failed) over time in mTLS sessions.

Visualizations	Details
Client Certificate OCSP Trend	Displays the trend of OCSP revocation-check results for client certificates over time in mTLS sessions.
Server Certificate CRL Trend	Displays the trend of CRL based revocation-check results for server certificates over time.
Client Certificate CRL Trend	Displays the trend of CRL based revocation-check results for client certificates over time in mTLS sessions.

Table 17: Engine Diagnostics Dashboard-Displays the certificates and SSL alerts related to a GigaSMART engine

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

Advanced Dashboards

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

Refer the following sections for more details:

- System Requirements to configure Inline TLS/SSL Advanced Dashboards
- Rules and Notes
- Configure Advanced dashboard
- View Inline TLS/SSL Dashboards

System Requirements to configure Inline TLS/SSL Advanced Dashboards

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

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

Rules and Notes

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

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

Configure Advanced dashboard

To configure advanced dashboards:

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

View Advanced Dashboards

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

- Host Name
- GigaSMART Engine ID (GSEngine ID)
- URL (Only for Session Table Dashboard)
- Source IP
- Destination IP
- URL Category (Only For Session Table Dashboard)

The following are the advanced dashboards and its visualizations:

Table 18: Session Insight Dashboard-Displays visualizations on the details of an Inline TLS/SSL session.

Visualizations	Details
Decryption Status	Displays the number of Inline TLS/SSL sessions that were decrypted and not decrypted.
SSL Mode	Displays the distribution of TLS/SSL Session modes. The modes are as follows: • TLS/SSL Outbound- : Sessions decrypted due to ISSL inbound deployment. • TLS/SSL Inbound-Sessions decrypted due to ISSL outbound deployment. • TLS/SSL Bypass- The session mode that is neither inbound or outbound. • Non-SSL - TCP sessions that are not an TLS/SSL session.
SSL State	Displays the distribution of TLS/SSL Session statuses.
Policy Match By Rules	Provides an insight into the TLS/SSL session that matches the Policy Rules. **NOTE:** The Policy Rule CATEGORY indicates the URL category.
TLS Version	Displays the TLS version of the sessions. **NOTE:** The counter “Bypass/Error” denotes sessions that were not able to determine the TLS version.
Top URLS (Max 10)	Displays the top 10 URLs that were accessed during the Inline TLS/SSL Session.
Top URL Category (10 Max)	Displays the Category of top 10 URLs accessed in Inline TLS

Visualizations	Details
	sessions **NOTE:** 'Uncategorized' signifies SNIs that could not be categorized or Non TLS sessions. **NOTE:** 'Unknown' signifies TLS Bypass and IP address based URLs.
Top Ciphers (Max 10)	Displays the top 10 Ciphers that performed the Inline TLLS/SSL Decryption.
Certificates by Type	Displays the certificates received are valid or non-valid.

Table 19: Session Table Dashboard-Displays visualizations related to Inline TLS/SSL Sessions per Engine in a tabular format.

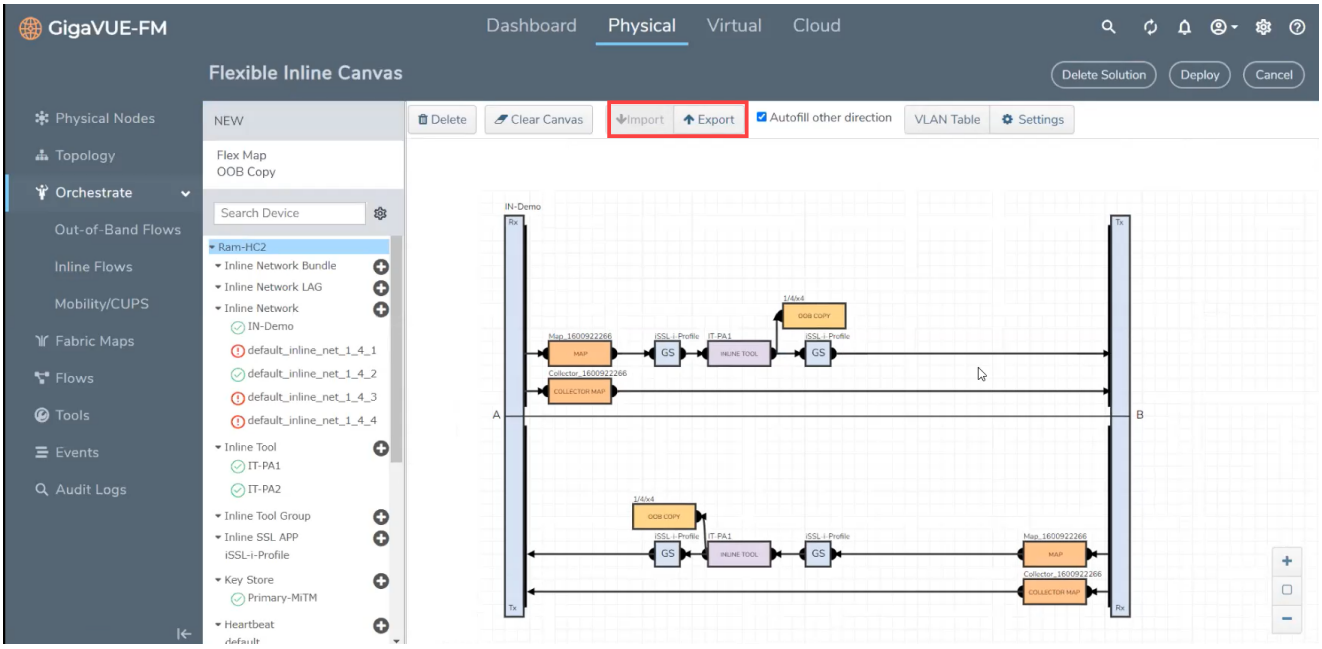
Visualizations	Dashboard
Session Debug Table	Displays the entire Sessions Debug details throughout the system that has enabled Advanced Session Statistics. Each field can be added or removed as a customized filter option by using ⊕ ⊖ button.
Session Policy Debug Table	Displays the entire Sessions Policy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

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.



Import and Export Flexible Inline Solution — 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 ,ensure you have the .YAML file ready of the flexible Inline solution:

1. Click the **Import** 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, from the Flexible Inline Canvas select the solution:

1. Click the **Export** button. The solution is downloaded as an YAML file.
2. Save the file to the required location.

Backup and Restore Flexible Inline Flows

You must backup the devices and GigaVUE-FM,GigaVUE HC Series 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,GigaVUE HC Series for Traffic Management Solutions"* section in the *"GigaVUE Administration Guide"*.

Troubleshoot Inline TLS/SSL Decryption Solution Issues

In the **Inline Flows** page, select the required device for which you want troubleshoot. Click the **Troubleshoot** button to view more on the below map configuration details for the ICAP Client app:

- Map Alias
- From
- To
- GSOP

The Troubleshoot tab provides details about the device-level maps that are created for the selected inline network. The details include the source of the map, inline tools or inline tool groups used in the A to B and B to A directions, tool side and network side VLAN tags, and OOB copies. Click the required component to view its configuration details. It also provides details about the map statistics. In case of inline tool group, you can view the list of inline tools that are associated with the inline tool group. Click the required inline tool from the list to view its configurations. Refer to the following figure for details.

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

1. From the flexible inline canvas, go to the **Status** tab to check the forwarding states of the required inline network.
2. Go to the **Statistics** tab to check the port statistics to ensure that there are no drops, errors, or discards.

> **NOTE:** Click the **Statistics** tab again to refresh the data.

3. Go to the **Troubleshoot** tab to check the required map configurations and isolate the issue.

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.

Inline Configurations using Inline Bypass Solutions (Classic)

This chapter provides the following information about inline bypass solutions:

- Inline Bypass Solutions
- Configure Inline Bypass

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.

Inline bypass is supported on GigaVUE-FM,GigaVUE HC Series® HC Series nodes: GigaVUE-HC3, GigaVUE-HC1 and GigaVUE-HC1-Plus.

This section describes inline bypass solutions. Refer to the following sections for details:

- Introduction to Inline Bypass Solutions
- Logical Bypass and Physical Bypass
- Types of Inline Networks
- Simple and Complex Inline Bypass Solutions
- Types of Inline Networks
- Inline Network Groups
- Inline Tools
- Heartbeats
- Inline Tool Groups
- Inline Serial Tools
- Associate Inline Networks with Inline Tools Using Inline Maps
- Configure Inline Bypass

NOTE: The configuration of inline bypass solutions can be complex. Follow the order of configuration outlined in Configure Inline Bypass 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 in the event of a tool failure.
- Supports out-of-band maps for selective forwarding of inline traffic to monitoring tools.
- Supports the 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-FM,GigaVUE HC Series 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-FM,GigaVUE HC Series node, such as if power is lost. On GigaVUE-FM,GigaVUE HC Series® 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-FM,GigaVUE HC Series 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 77Physical Bypass Protection and Protected Inline Network.

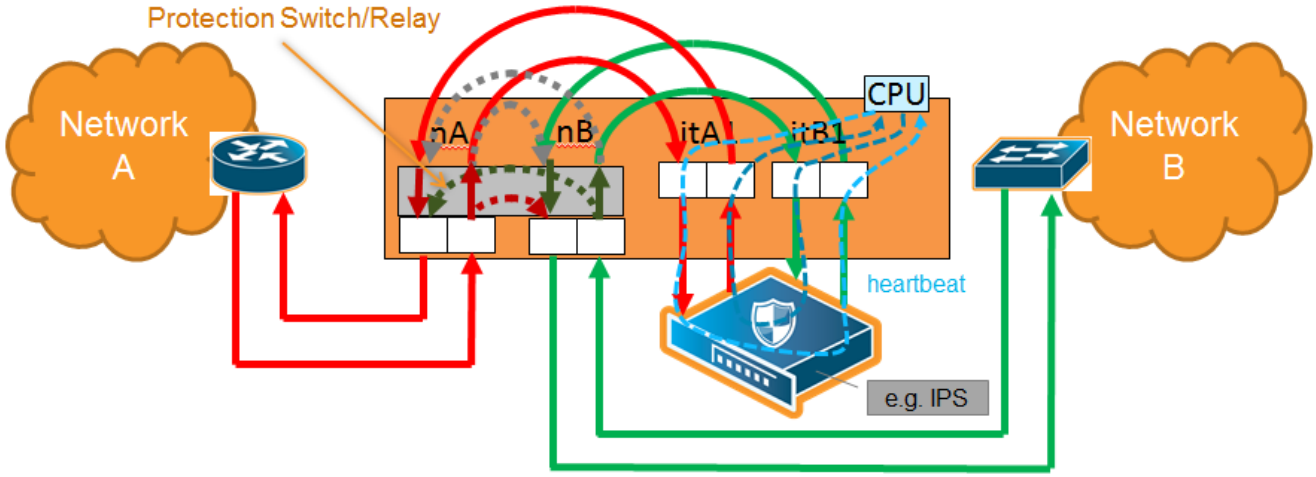


Figure 77 *Physical Bypass Protection*

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 78Simple Configuration. The physical protection switch is optional and present only when using bypass combo modules or copper TAP modules.

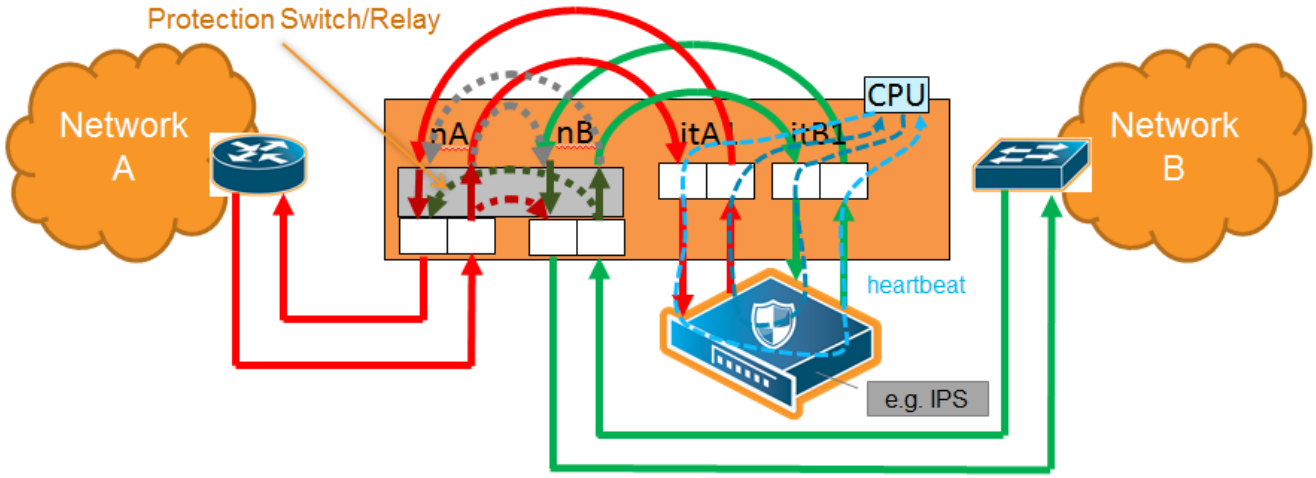


Figure 78 *Simple Configuration*

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 79Multiple Tool Arrangement.

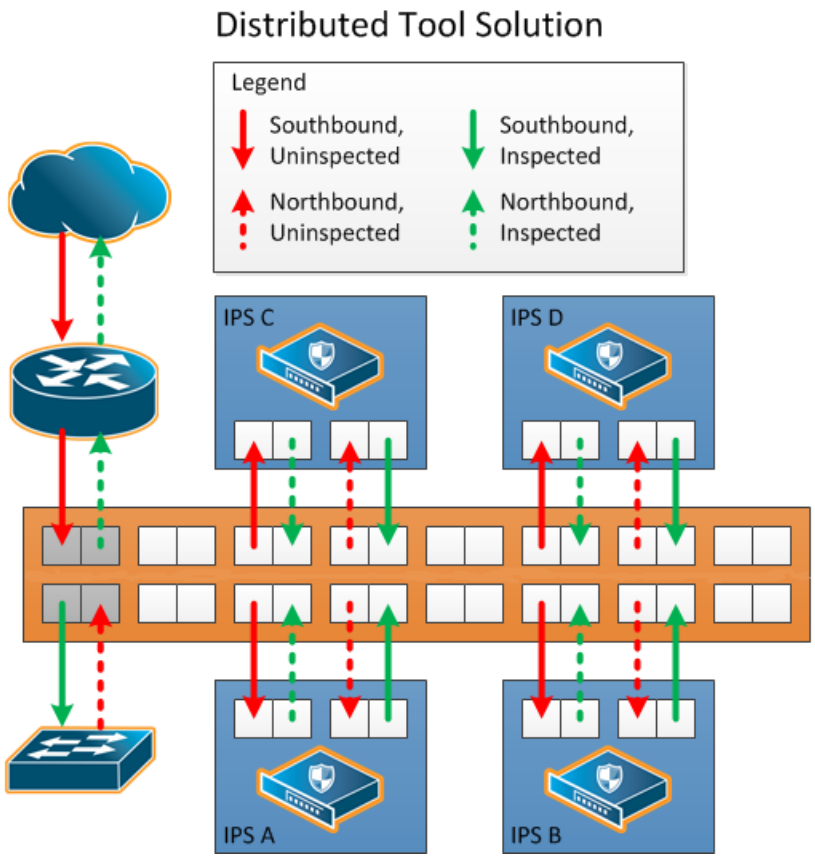


Figure 79 *Multiple Tool Arrangement*

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 80Inline Tool Series for an inline tool series. In Figure 80Inline Tool Series, traffic is only shown from A-to-B.

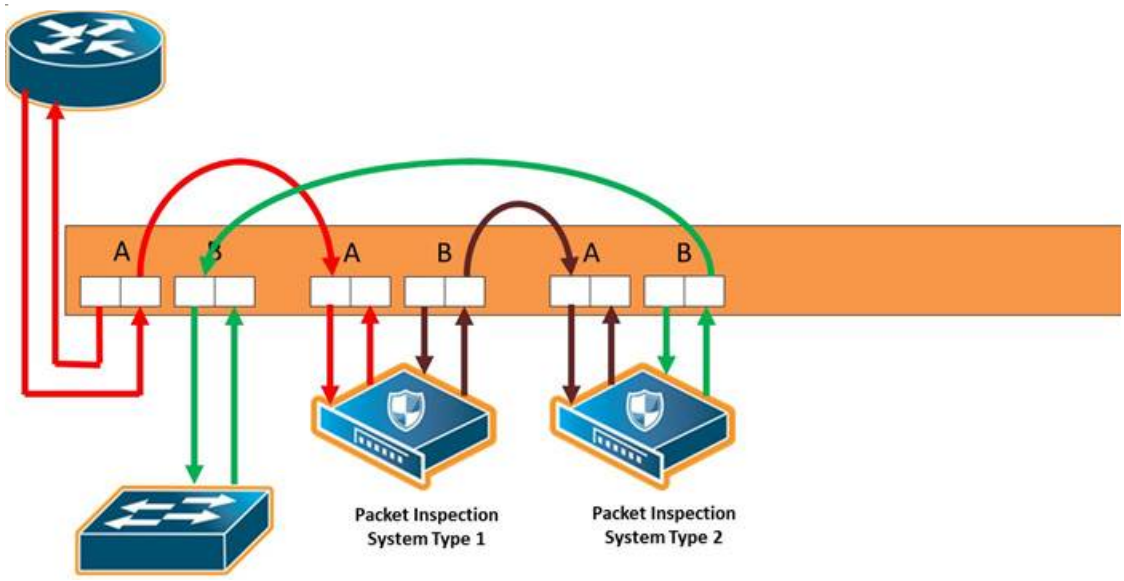


Figure 80 *Inline Tool Series*

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 81Inline Network Group or many-to-many. Traffic is guided to a particular inline network through internal VLAN ID tagging.

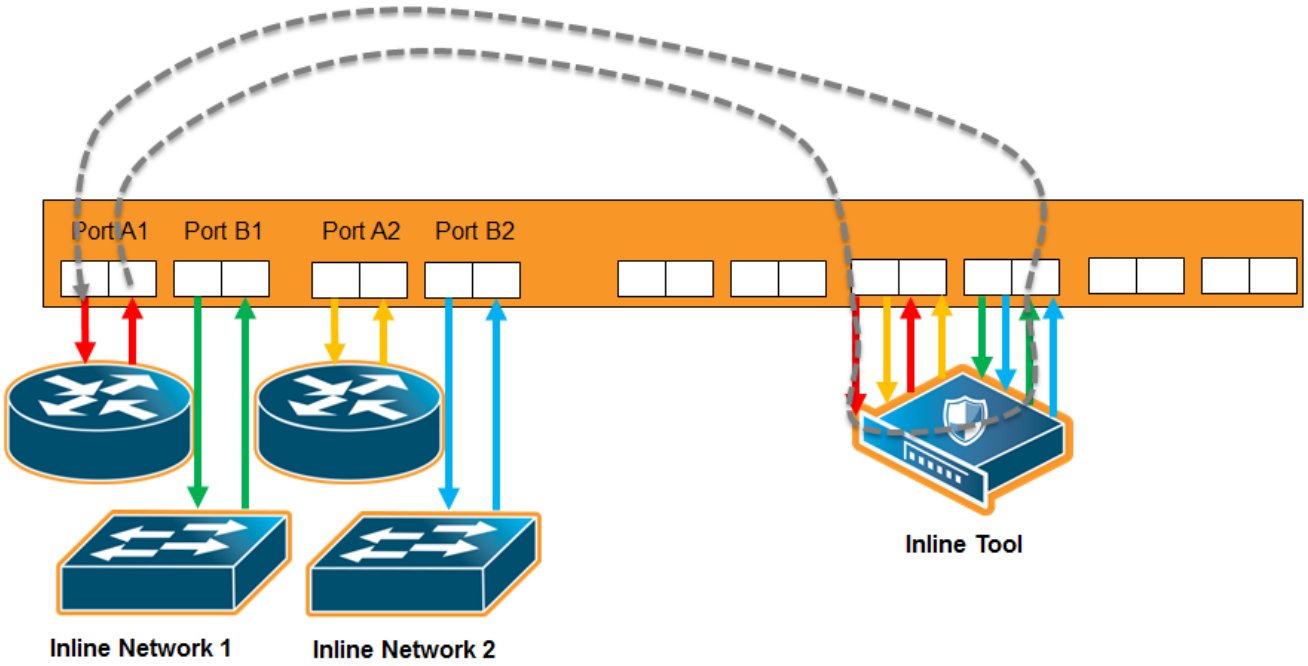


Figure 81 *Inline Network Group*

For more information on inline network groups, refer to Inline Configurations using Inline Bypass Solutions (Classic).

Inline Flow Mapping®

Some tools are optimized for particular inline traffic. With inline flow mapping, the GigaVUE-FM,GigaVUE HC Series 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 82Inline Flow Mapping® Based Traffic Distribution. In Figure 82Inline Flow Mapping® Based Traffic Distribution, traffic is only shown from A-to-B.

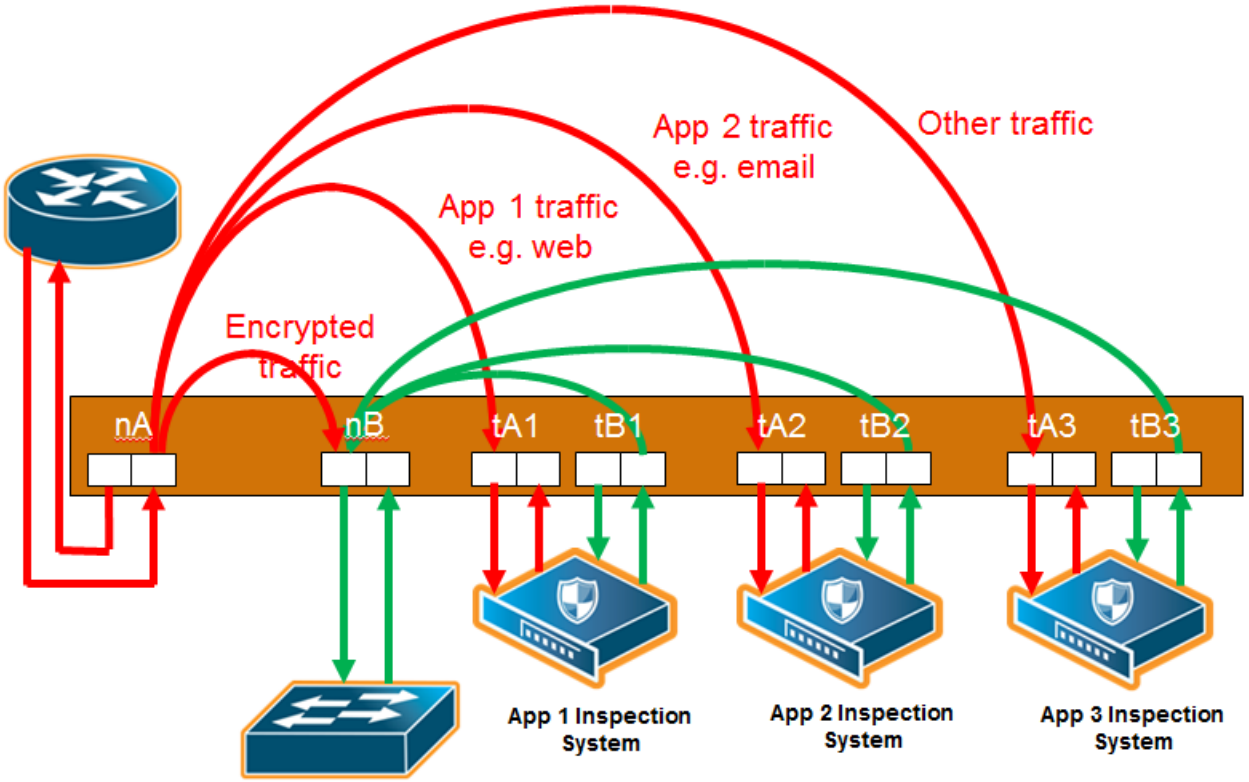


Figure 82 *Inline Flow Mapping® Based Traffic Distribution*

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 83Out-of-Band Arrangement.

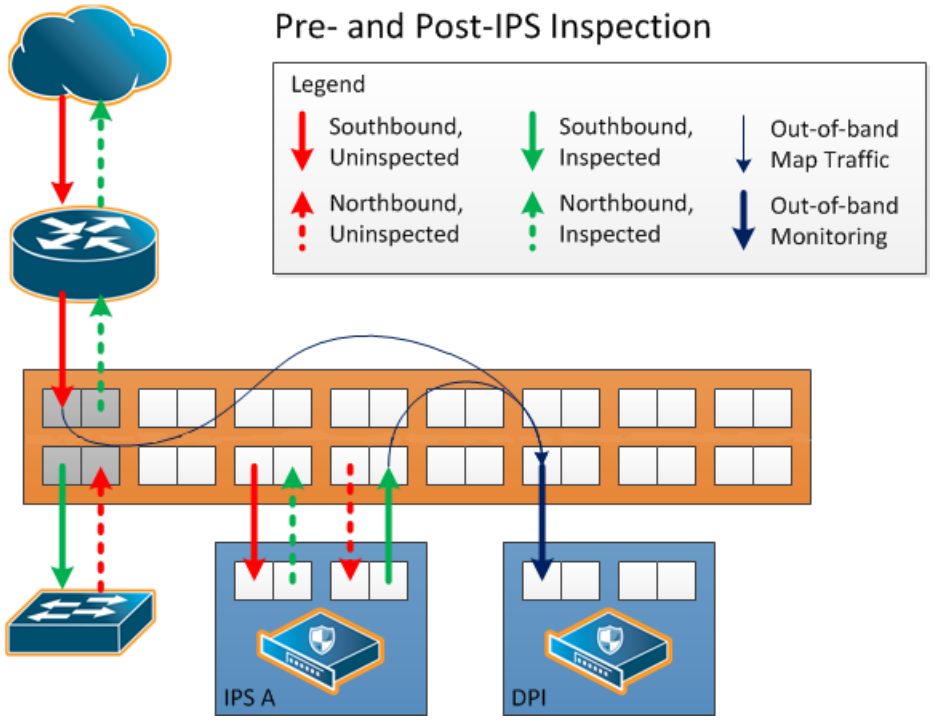


Figure 83 *Out-of-Band Arrangement*

Types of Inline Networks

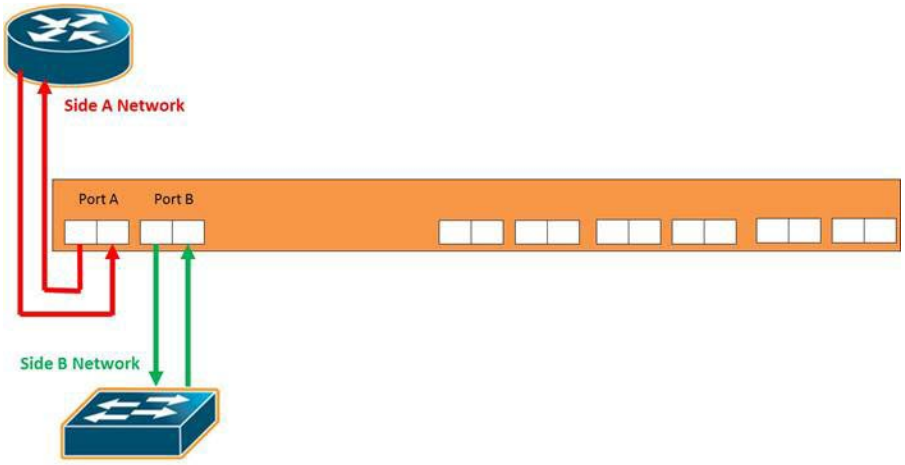
The types of inline networks are described in the following sections:

- 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-FM,GigaVUE HC Series 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.



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 84Inline Network Group Many-to-One. Many-to-many means from an inline network group to an inline tool group or inline tool series.

Figure 84 *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-FM,GigaVUE HC Series 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 85Inline Tool Inserted in Multiple NetA-NetB Links.



Figure 85 *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 84Inline 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-FM,GigaVUE HC Series® 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-FM,GigaVUE HC Series 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 86Inline Tool and Inline Tool Ports.

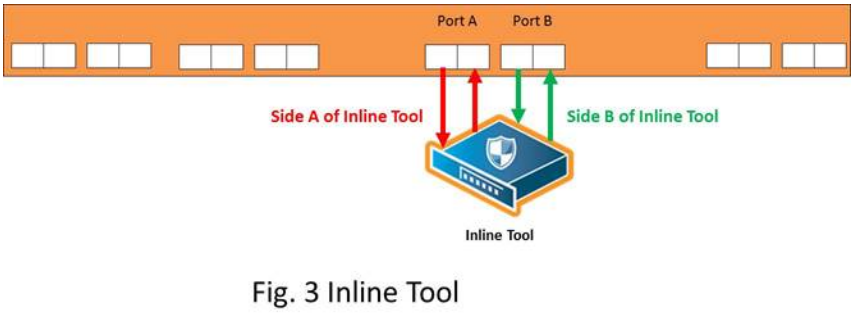


Figure 86 *Inline Tool and Inline Tool Ports*

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 87Automatic and Manual Inline Tool Recovery from Failover for automatic and manual inline tool recovery from failover.

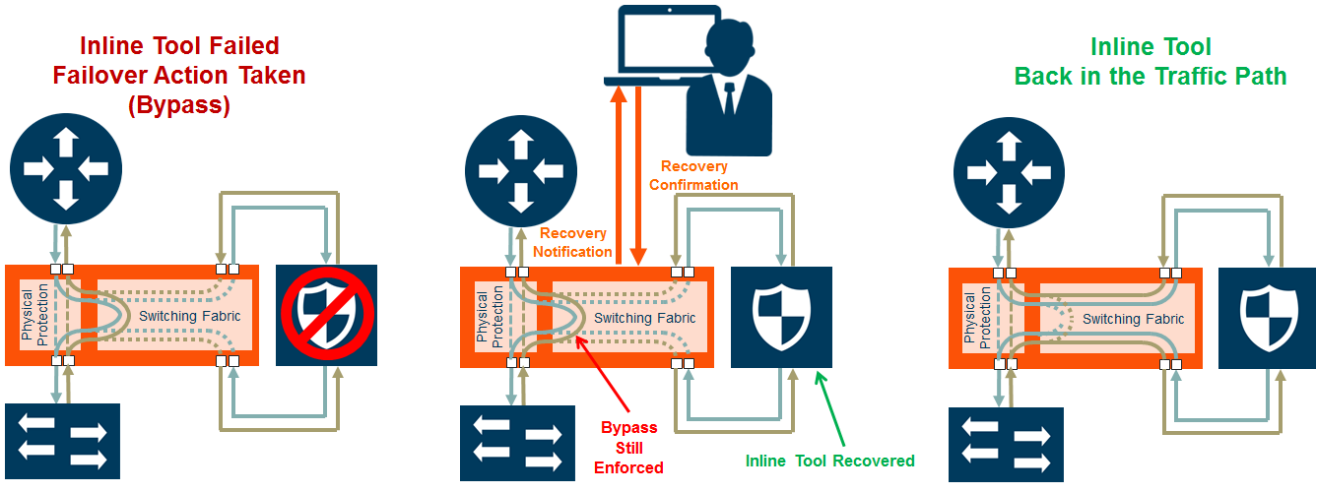


Figure 87 *Automatic and Manual Inline Tool Recovery from Failover*

The left side of Figure 87Automatic 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 87Automatic 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 87Automatic 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-FM,GigaVUE HC Series® 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 20: Aggregate Inline Tool States. They are an aggregate of the inline tool port statuses and the heartbeat status.

Table 20: 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

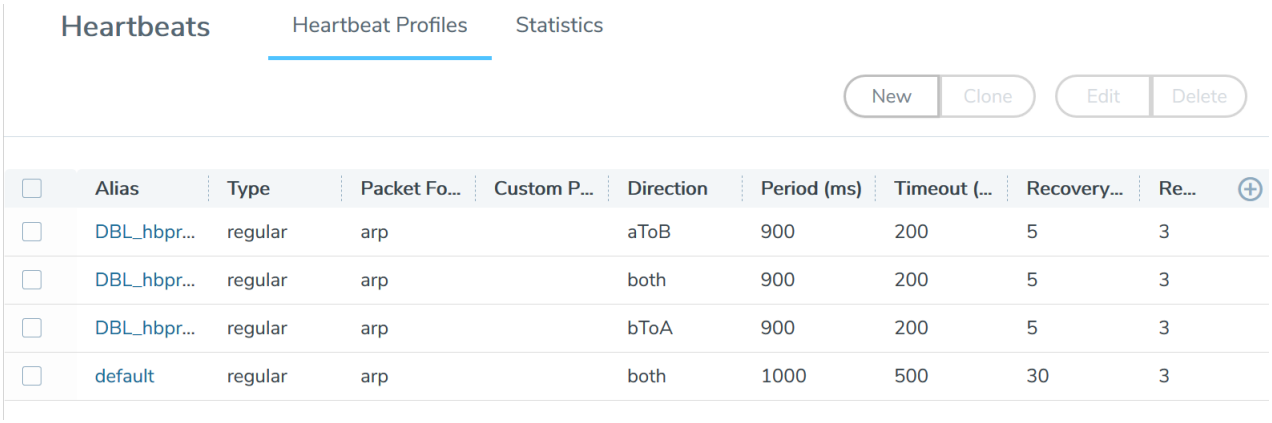
Keep in mind the following rules and notes when you configure the heartbeat mechanism between GigaVUE nodes:

- 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 88Heartbeats Page with Heartbeat Profiles.



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 88 *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-FM,GigaVUE HC Series node that passes through the inline tool to verify that it is passing traffic, even if the link is *up*. If the packet is not passed through the tool, the tool is considered to have failed and a bypass action is triggered.

Even when the tool is considered down, heartbeat packets continue to be sent so that the bypass action can be reversed when the tool is healthy again.

Heartbeats are sent bidirectionally to the inline tool.

Standard or Custom Heartbeat Packet

The format of the heartbeat packet can be the standard ARP packet or a custom packet. For a custom packet, you must provide a URL from which a PCAP file can be imported. If the PCAP file contains several packets, the first packet present in the file is taken as the heartbeat packet. The size of a custom heartbeat packet must be less than 128 bytes.

NOTE: The system will overwrite the MAC address portion of the custom heartbeat packet.

If the inline tool through which the heartbeat packets are passed is expecting IPv6 traffic exclusively, you must select a custom heartbeat packet.

Custom heartbeat packets are needed in situations in which inline tools do not reliably pass standard ARP packets. For example, if an inline tool is configured to pass only IPv6 traffic, an ICMPv6 ARP packet might be appropriate.

If a custom heartbeat packet is specified, the **Heartbeats** page displays the name of the PCAP file from which it was imported.

Detect Inline Tool Failure

The health of the inline tool is critical to the proper handling of traffic. An inline tool is determined to have failed if:

- link is lost to the tool
- inline heartbeat fails

When the tool fails in one direction, it is considered to have failed in both directions. For example, if the heartbeat stops flowing in the northbound direction, neither northbound or southbound packets are sent to the tool.

Negative Heartbeat Profiles

A negative heartbeat profile is a group of attributes applied to an inline tool to configure its negative heartbeat operation. Multiple inline tools can share a negative heartbeat profile. The content of a negative heartbeat is configurable using the same PCAP file mechanism as for a custom heartbeat packet.

Use the Add Heartbeats page shown in Heartbeats to configure a negative heartbeat profile. A negative heartbeat profile can be created by selecting **Negative** in the **Type** field. The profile will have the following settings:

- **PacketFormat** set to **Custom**
- **Direction** set to **Bi-directional**
- **Period** set to **1000** (period is specified in milliseconds)
- **Recover Time** set to **30** (recovery is set in seconds)

You must provide a valid PCAP file when Packet Format is set to Custom before the negative heartbeat profile can be applied to an inline tool.

When a negative heartbeat is configured, the system will send packets specified by the **Custom Packet Format**, in the time specified by **Period**, in the direction specified by **Direction**. The inline tool absorbs the negative heartbeat packets until the number of seconds specified by **Recovery Time** has passed. **Recovery Time** specifies the number of seconds of not receiving negative heartbeat packets in order for the inline tool to be declared healthy.

The negative heartbeat mechanism supports the maximum number of inline tools, which is 48 on the GigaVUE-HC3, 16 on the GigaVUE-HC1 , and 48 on the GigaVUE-HC1-Plus.

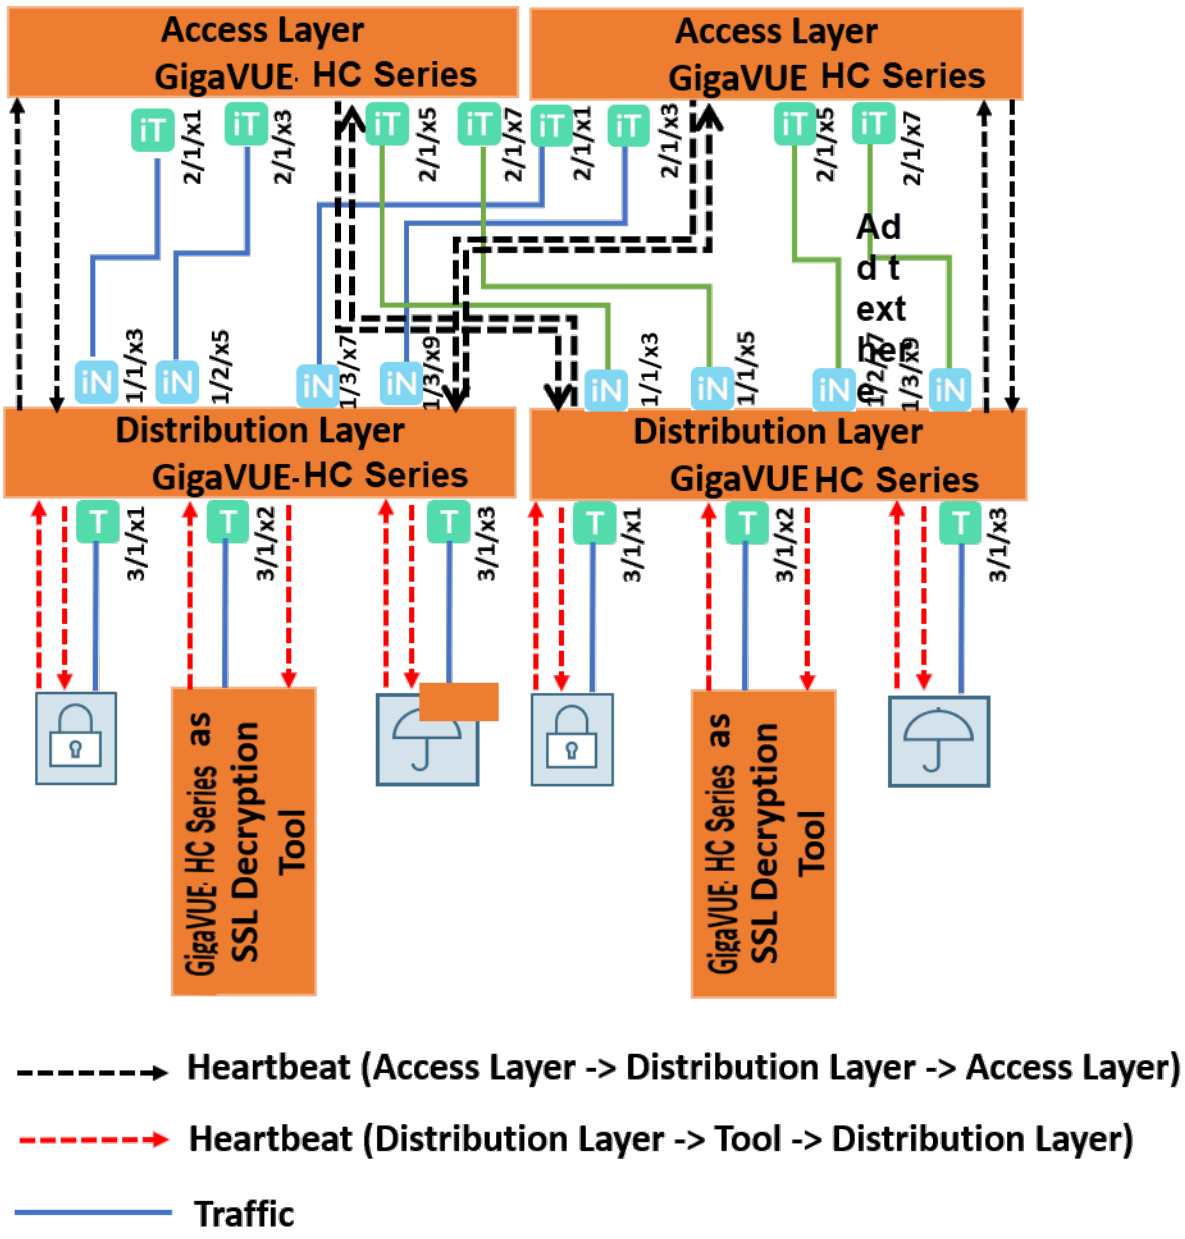
Use the Heartbeat page to display the configured negative heartbeat profiles. Use the Inline Tool page to display the negative heartbeat profile associated with an inline tool, the negative heartbeat status, and the counters of received packets in each direction. There is also a combined heartbeat status, which combines heartbeat and negative heartbeat statuses and indicates the tool health used for inline tool failover actions or SNMP polling. The combined heartbeat status is the combination of both heartbeat statuses. For example, if both are enabled and both are up, the combined status is up. If both are enabled and one is down, the combined status is down.

Heartbeat Support Between GigaVUE Nodes

The heartbeat mechanism focuses on providing extended heartbeat capability to monitor the following types of devices when the devices are connected to the inline-tool pair of ports as a tool:

- GigaVUE nodes
- GigaVUE nodes with GigaSMART operations configured

Following figure illustrates an example of a topology with GigaVUE nodes placed at three different layers.



GigaVUE HC Series means it can be either GigaVUE-HC3, GigaVUE-HC1 or GigaVUE-HC1-PLUS

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

The GigaVUE node at the distribution layer distributes the traffic from the access layer to the tool layer.

The GigaVUE node at the tool layer acts as the SSL decryption tool.

In this topology, heartbeats are essential to monitor the traffic integrity at the distribution layer and to ensure automatic failover in case of a tool failure. In the access layer device, the ports that are connected to the distribution layer device are configured as inline tool ports because they face the tool side. In the distribution layer device, the ports that are connected to the access layer device are configured as inline network ports because they face the network side of the topology. The heartbeat packets will be sent from the inline tool port pair of the access layer device to the inline network port pair of the distribution layer device. If the forwarding state of the inline network pair is normal, the heartbeat packet is sent back to the inline tool port pair of the access layer device. Else, the packet is dropped.

The heartbeat mechanism is extended to support the GigaVUE node at the distribution layer to monitor the GigaVUE node that acts as a tool at the tool layer. In the distribution layer device, the ports that are connected to the tool layer device are configured as inline tool ports. In the tool layer device, the ports that are connected to the distribution layer device are configured as inline network ports. The heartbeat packets that are sent from the distribution layer device to the tool layer device will monitor the availability of both, the tool layer device and its GigaSMART engines.

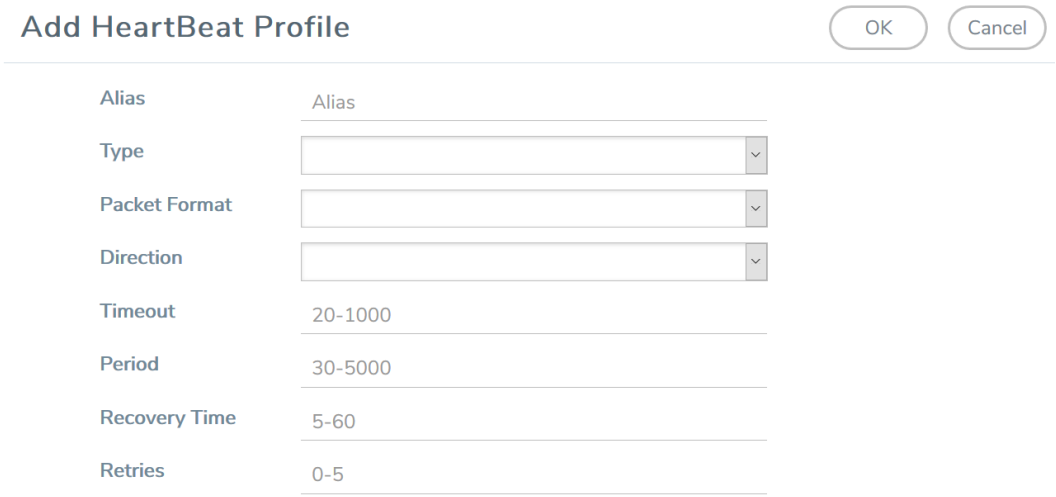
Troubleshoot Heartbeat Failures

In your topology, you may sometimes notice heartbeat packets getting dropped at the inline network or the heartbeat status of the inline tool is not up. For details about how to isolate and troubleshoot such failures, refer to the following sections:

- Heartbeat Status of Inline Tool is Not Operationally Up
- Heartbeat Packets are Dropped

Configure a Heartbeat Profile

Use the Add Heartbeats page shown in the below figure to configure a regular heartbeat profile or a negative heartbeat profile by selecting **Inline Bypass > Heartbeats > Heartbeats**, and then clicking **New**.



View Heartbeat Profile Statistics

Statistics about heartbeat profiles are displayed on the heartbeats statistics page. To open the Statistics page, select **Inline Bypass > Heartbeats > Statistics**. The page shows the following information:

Column	Description
Alias	The alias of the inline tool or inline network to which the heartbeat profile is associated with.
Inline Type	The inline type to which the heartbeat profile is associated with. The valid values are Inline Tool or Inline Network.
Heartbeat Profile	The alias of the heartbeat profile
Heartbeat Type	Indicates the type of heartbeat profile: Regular or Negative
A to B Packets	The number of packets sent/received from port A to port B of the inline tool.
B to A Packets	The number of packets sent/received from port B to port A of the inline tool.
A to B Packet Drops	The number of packets dropped from port A to port B of the inline tool or inline network.
B to A Packet Drops	The number of packets dropped from port B to port A of the inline tool or inline network.

You can clear the statistics for a specific heartbeat profile or all heartbeat profile. To clear the statistics for a heartbeat profile, select the required row, and then click **Clear**.

To clear the statistics for all heartbeat profiles, select **Clear All**, and the choose one of the following options:

- Select **Clear All Inline Network Heartbeat Stats** to clear the heartbeat statistics for all inline networks.
- Select **Clear All Inline Tools Heartbeat Stats** to clear the heartbeat statistics for all inline tools.
- Select **Clear All Negative Heartbeat Stats** to clear only the statistics for negative heartbeat profiles. The statistics for regular heartbeat profiles will remain.

Heartbeat Status after System Reload

Following a reload, there is a 5-minute delay for the system to stabilize before heartbeat packets are sent or received. During this delay, the heartbeat status is down.

Inline Tool Groups

Use the **Inline Tool Groups** configuration page to configure an inline tool group, which is an arrangement of multiple inline tools to which traffic is distributed based on hashing. In an inline tool group, traffic is shared. Each inline tool in the group receives a portion of the traffic. The distribution mechanism includes a way of dealing with failures of individual tools through traffic redistribution to the remaining healthy tools.

The inline tool ports that make up the inline tools participating in the inline tool group are always in pairs, running at the same speed, on the same medium (fiber or copper). All inline tool ports of the inline tool group must be on the same GigaVUE-HC3, or GigaVUE-HC1 and GigaVUE-HC1-Plus node, but can be on different modules on the node. The inline tool ports must also be on the same GigaVUE-HC3, or GigaVUE-HC1 and GigaVUE-HC1-Plus node as the inline network ports.

Inline tool groups can be configured as follows:

- non-redundant—multiple inline tools with no spare inline tool. Refer to Figure 89Inline Tool Group With No Spare, Non-Redundant.
- 1+1 redundancy—single inline tool with a spare inline tool. Refer to Figure 90Inline 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 91Inline Tool Group With Spare, Redundant N+1 Scenario.

With 1+1 redundancy, an inline tool is paired with a standby tool. When there is a loss of link or a heartbeat failure on an active tool, the traffic will be sent to the standby tool with no loss. In addition, if the standby tool fails, you can configure what happens to the traffic in that case, such as drop it or forward it.

With N+1 redundancy, one tool is added to a group of N distributed inline tools. When any one of the N tools fails, the traffic from that tool is sent to the standby (or spare) tool with no loss. In addition, if the standby tool fails, you can configure what happens to the traffic in that case, such a redistribute it, send it to another tool, or declare a failure on the tool group.

Refer to Figure 89Inline Tool Group With No Spare, Non-Redundant for a non-redundant inline tool group.

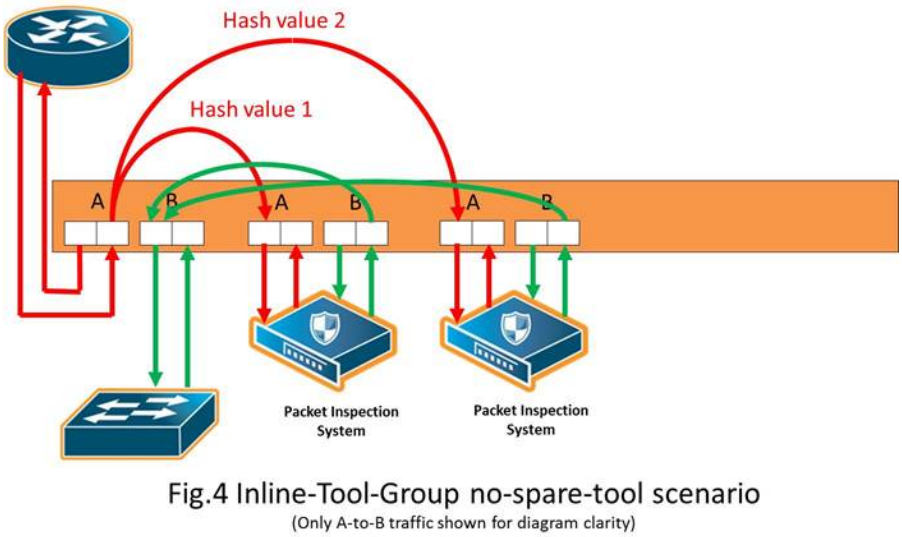


Figure 89 *Inline Tool Group With No Spare, Non-Redundant*

Refer to Figure 90Inline 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 90Inline Tool Group With Spare, Redundant 1+1 Scenario, traffic is only shown from A-to-B.

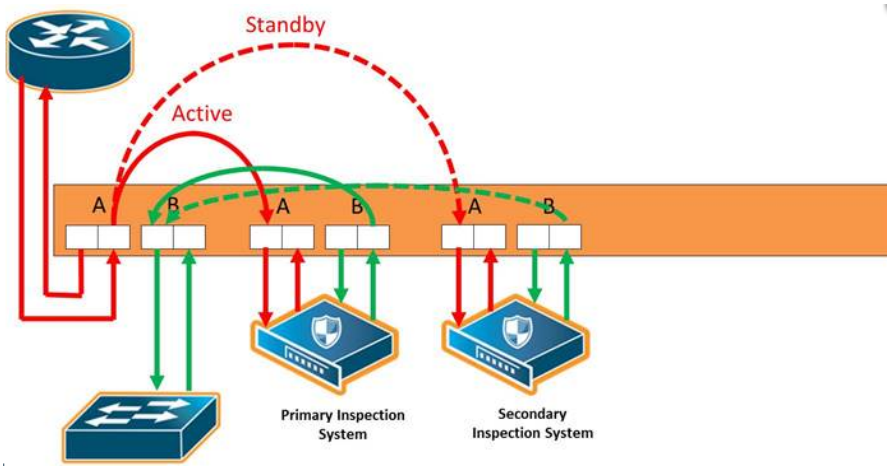


Figure 90 *Inline Tool Group With Spare, Redundant 1+1 Scenario*

Refer to Figure 91Inline 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 91Inline Tool Group With Spare, Redundant N+1 Scenario, traffic is only shown from A-to-B.

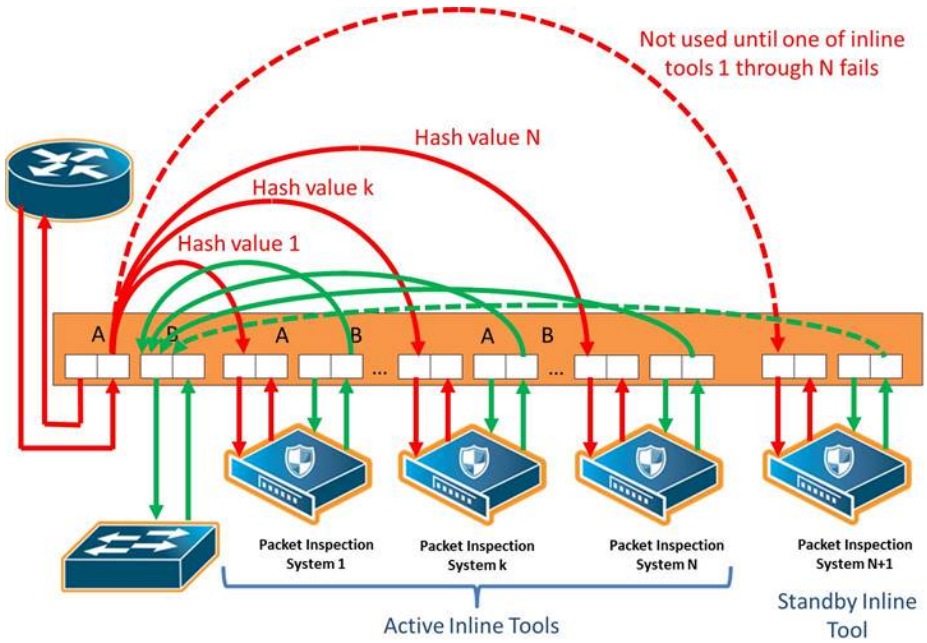


Figure 91 *Inline Tool Group With Spare, Redundant N+1 Scenario*

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-FM,GigaVUE HC Series 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 92Asymmetrical Hashing for asymmetrical hashing.

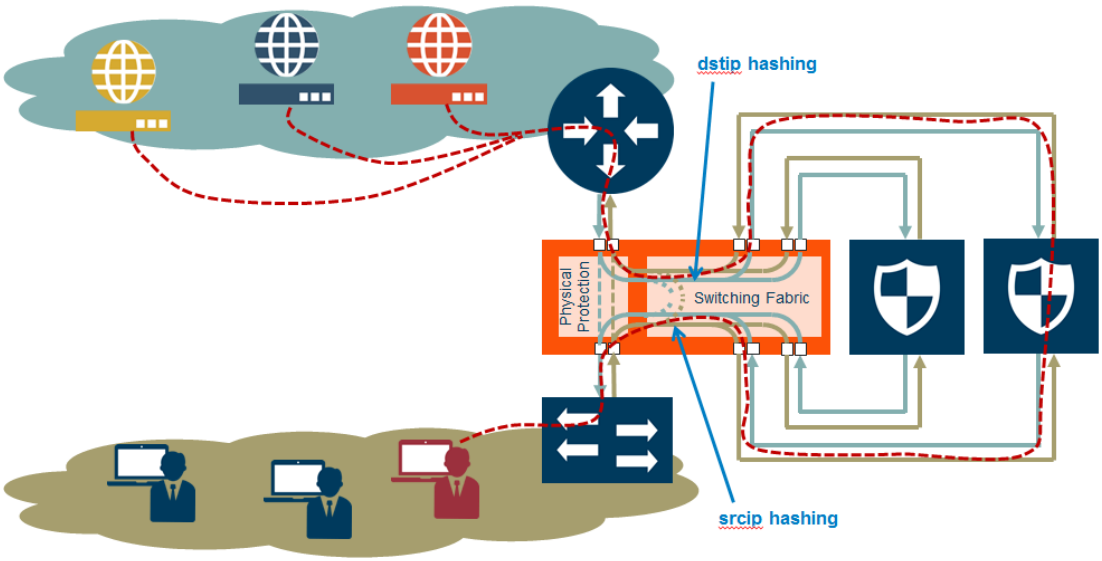


Figure 92 *Asymmetrical Hashing*

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 93Inline Serial Tools for an illustration of an inline tool series. In Figure 93Inline Serial Tools, traffic is only shown from A-to-B.

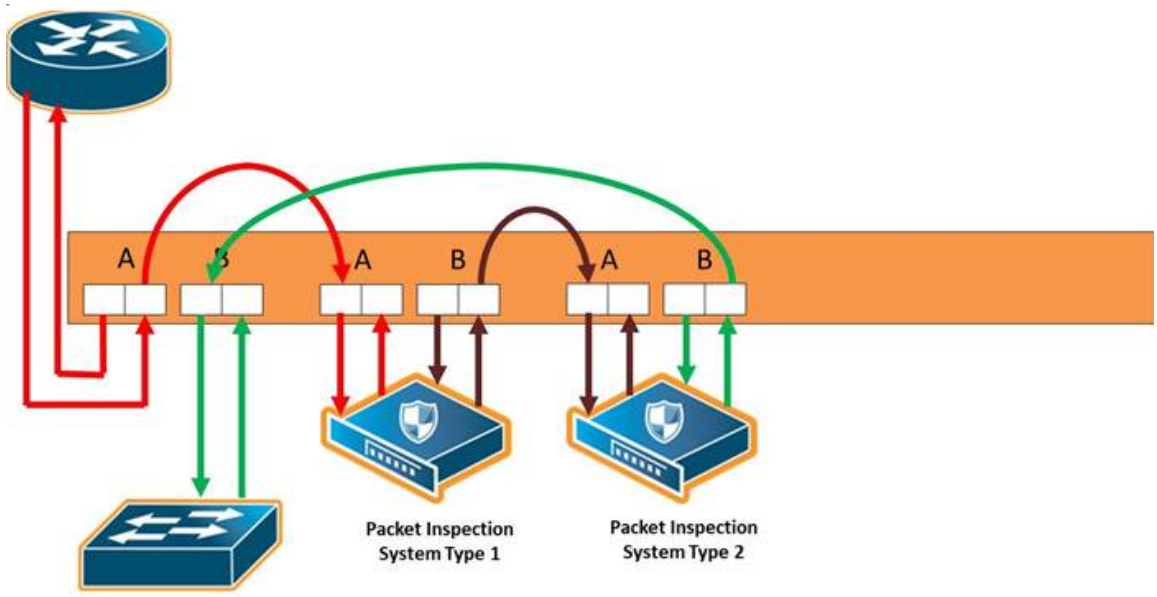


Figure 93 *Inline Serial Tools*

Refer to Figure 94Inline 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 94Inline 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.

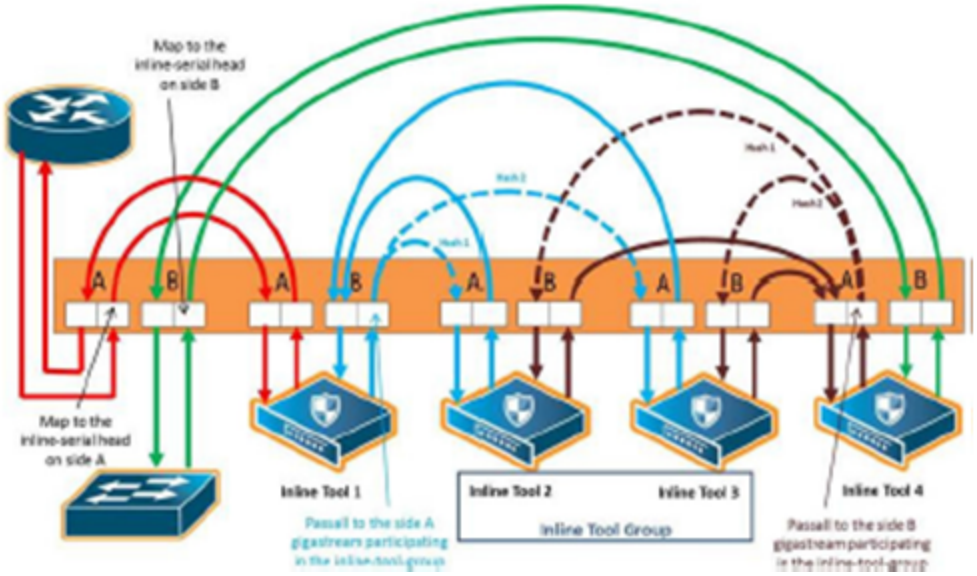


Figure 94 *Inline Tool Series, Including Inline Tool Group*

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 95Inline 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 95Inline 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 96Inline 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 97Inline 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 97Inline 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 95Inline Tool Series Global Bypass Failover Action to Figure 97Inline Tool Series Global Drop Failover Action show the global failover actions for the inline series when any individual inline tool in the series fails.

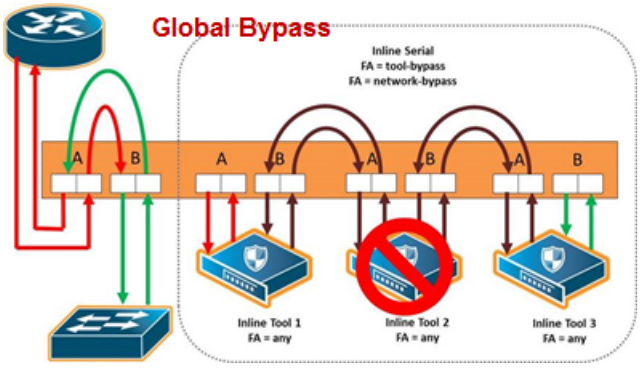


Figure 95 *Inline Tool Series Global Bypass Failover Action*

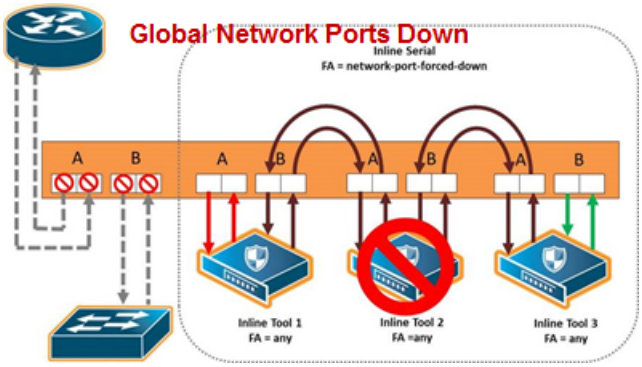


Figure 96 *Inline Tool Series Global Network Ports Forced Down Failover Action*

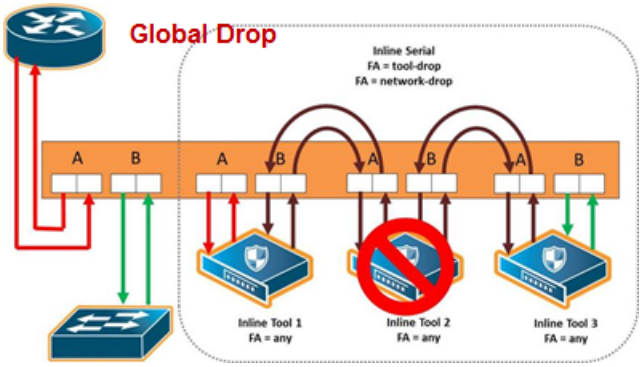


Figure 97 *Inline Tool Series Global Drop Failover Action*

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 98Inline 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 99Inline 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 100Inline Tool Series Local Network Bypass Failover Action.
- **NetworkDrop**—when the inline tool fails, traffic is dropped at the inline network ports. Refer to Figure 101Inline 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 102Inline Tool Series Local Network Ports Forced Down Failover Action.

Figure 98Inline Tool Series Local Tool Bypass Failover Action to Figure 102Inline Tool Series Local Network Ports Forced Down Failover Action show the local failover actions when an individual inline tool in a series fails

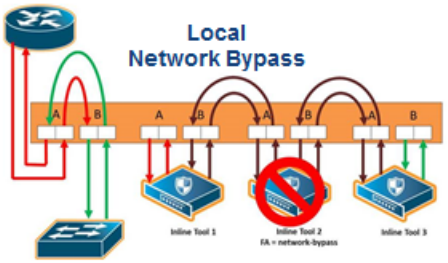


Figure 98 *Inline Tool Series Local Tool Bypass Failover Action*

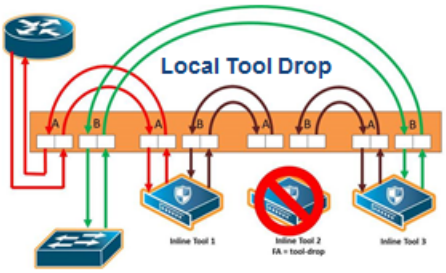


Figure 99 *Inline Tool Series Local Tool Drop Failover Action*

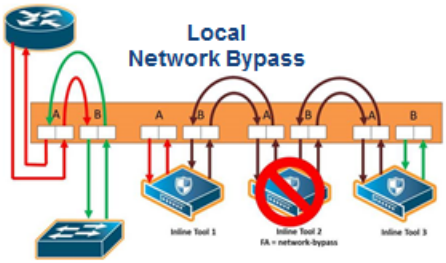


Figure 100*Inline Tool Series Local Network Bypass Failover Action*

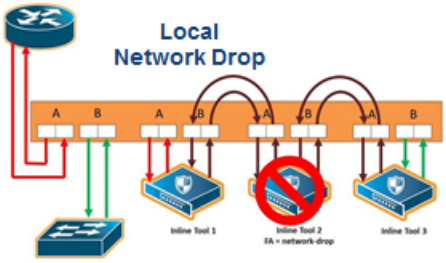


Figure 101 *Inline Tool Series Local Network Drop Failover Action*

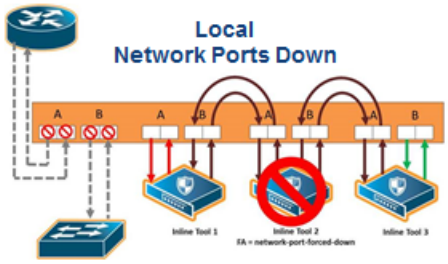


Figure 102*Inline Tool Series Local Network Ports Forced Down Failover Action*

Figure 103Inline Tool Series Local Failover of Two Tools shows the failure of two individual inline tools in a series with different configured failover actions.

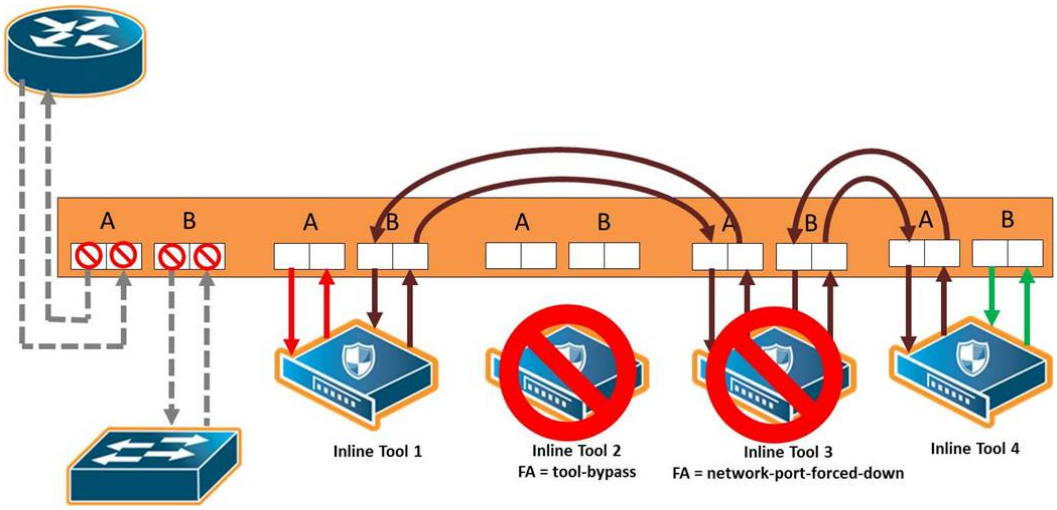


Figure 103*Inline Tool Series Local Failover of Two Tools*

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 104Inline Tool Series Per-Direction Order: Simplified Reverse and Forward shows a simplified view of the flow through the tools:

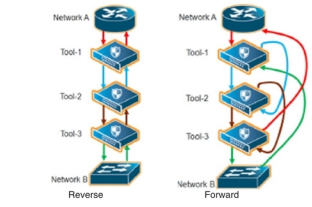


Figure 104*Inline Tool Series Per-Direction Order: Simplified Reverse and Forward*

In Figure 104Inline 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 105Inline Tool Series Per-Direction: Reversed shows the reverse direction with the tools connected to the GigaVUE-FM,GigaVUE HC Series node.

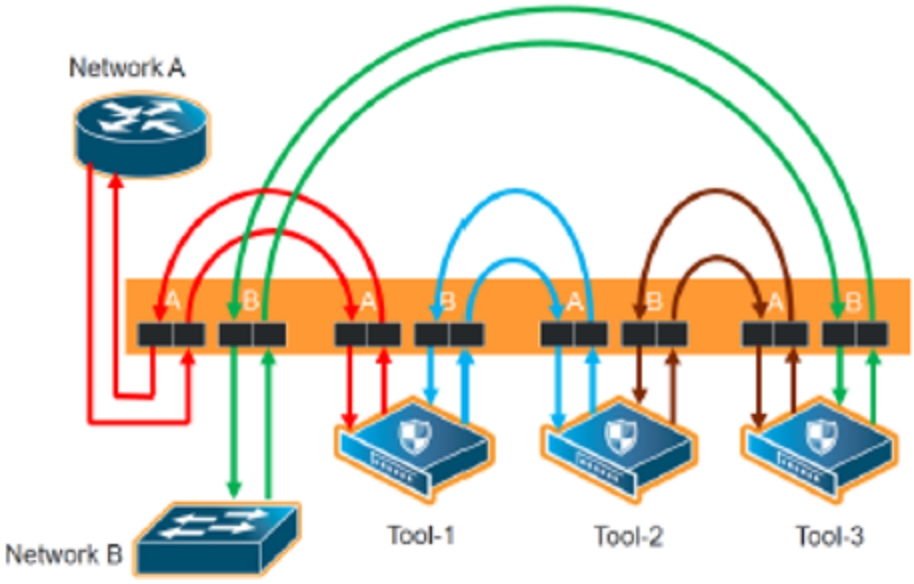


Figure 105*Inline Tool Series Per-Direction: Reversed*

Figure 105Inline Tool Series Per-Direction: Reversed shows the forward direction with the tools connected to the GigaVUE-FM,GigaVUE HC Series node.

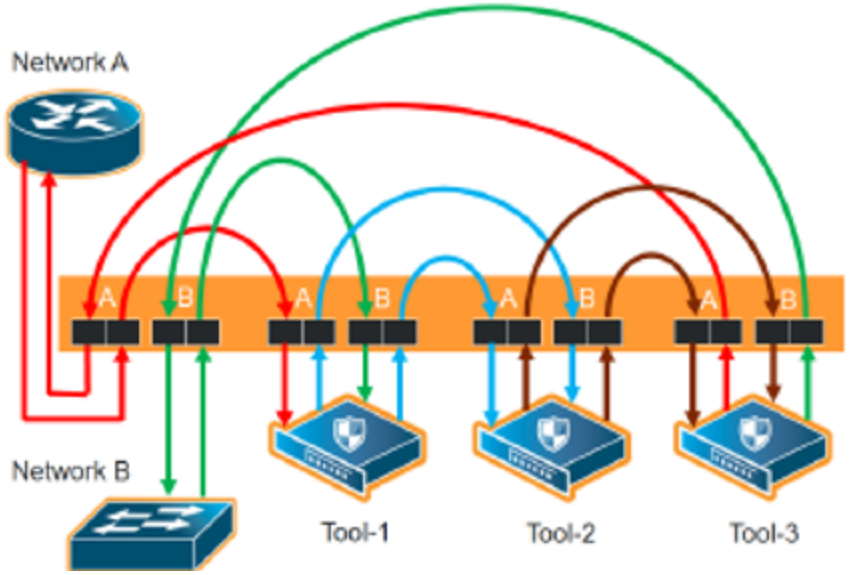


Figure 106*Inline Tool Series Per-Direction: Forward*

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-FM,GigaVUE HC Series 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-FM,GigaVUE HC Series 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 Inline Tool Groups.

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 107Rule-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 107Rule-Based Maps to Individual Tools in an Inline Tool Group, traffic is only shown from A-to-B.

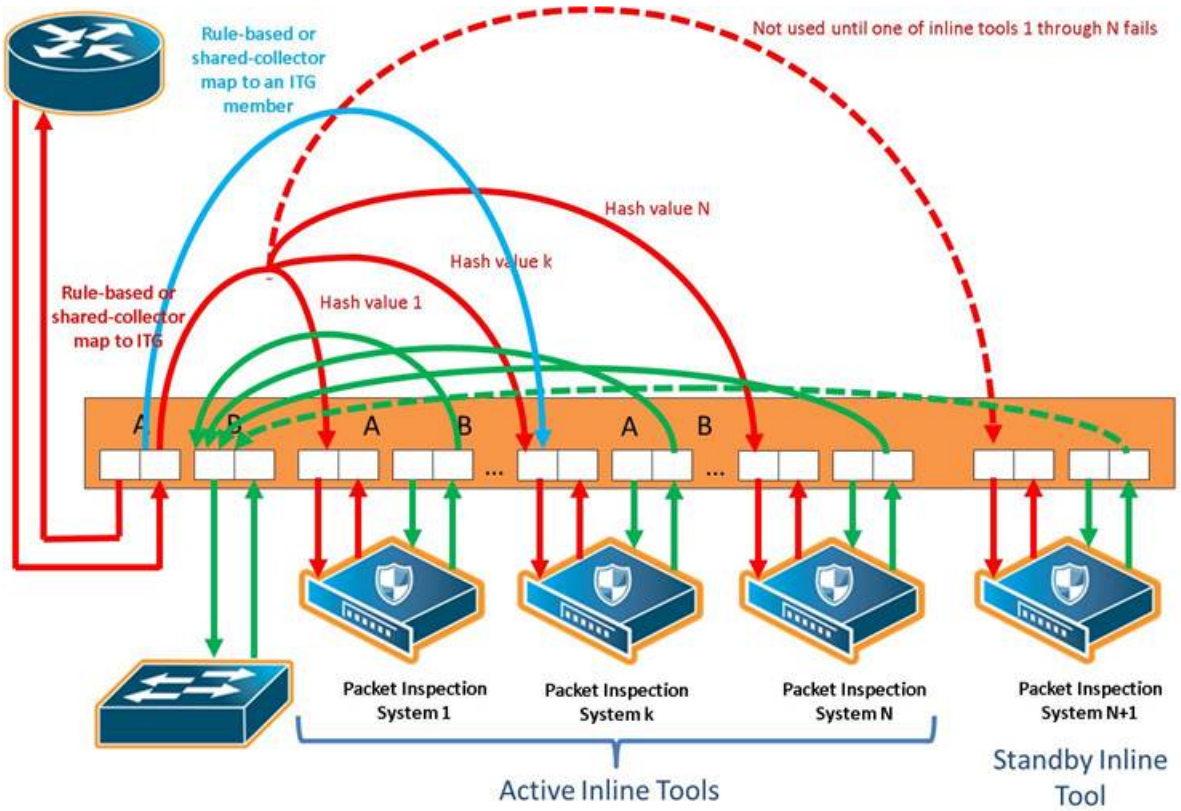


Figure 107*Rule-Based Maps to Individual Tools in an Inline Tool Group*

Map, Inline Tool, and Inline Tool Group Configuration Restrictions

The following are map, inline tool, and inline tool group configuration restrictions for the rule-based maps to the individual members of the inline tool group:

- If there is a map to the inline tool group as a whole, there cannot also be rule-based maps to the individual inline tools in the group.
- Maps to the individual inline tools in a group must be rule-based. Map passalls to the individual tools cannot be configured.
- The source of rule-based maps to the individual inline tools in a group must be the same (either the same inline network or the same inline network group). The shared collector must have the same source as well.
- If there is a map configured to the individual inline tools in a group, the inline networks must have their traffic path set to **to-inline-tool**. This applies to individual inline networks as well as to inline networks involved in an inline network group.
- For the individual inline tools in a group, the recovery mode of the individual inline tools must be configured as **automatic**. A recovery mode of **manual** cannot be configured.
- For the inline tool group, the failover action must be either **NetworkBpass** or **NetworkDrop**. A failover action of either **ToolBypass** or **ToolDrop** for the inline tool group cannot be configured.
- Only one inline shared collector map can be configured (among the set of inline maps).
- Maps must be created in a specific order. The shared collector map must be configured last. For example, if there are three inline tools in the group, configure the three maps to the individual members of the group first, then configure the shared collector map.
- Maps must be deleted in a specific order. The shared collector map must be deleted first. Then the maps to the individual members of the group can be deleted.
- Once the shared collector map is configured, any changes to the maps to the individual members of the group are restricted. Only the map rules can be edited.

  NOTE: All the rules in a map cannot be deleted. All maps must have attributes for **Source**, **Destination**, and at least one **rule** configured.

- When an inline tool group is included as a member of an inline series, inline maps to individual members of an inline tool group are not supported.

- If one of the inline tools is disabled in an inline-network-group to inline-tool-group map that has the same ingress VLAN tag on the ports of an inline-network pair then the traffic will be looped back to the same network. This behavior is seen only with passall maps. In above scenarios use rule-based maps instead of passall maps.

Inline Tool Failures and Failover Actions

An inline tool group has a failover action for the group as a whole. The failover action is taken in response to a failure when the number of healthy inline tools in the inline tool group (including the spare inline tool, if configured) falls below the configured minimum. In addition, the individual inline tools in the group have failover actions.

When there are maps to individual inline tool members of the group, an inline tool has both a group failover action and an individual failover action.

Refer to Failover Actions When an Individual Tool in an Inline Tool Group Fails for the failover actions when an individual inline tool in an inline tool group fails.

Table 21: 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.
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 108Out-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 109Out-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

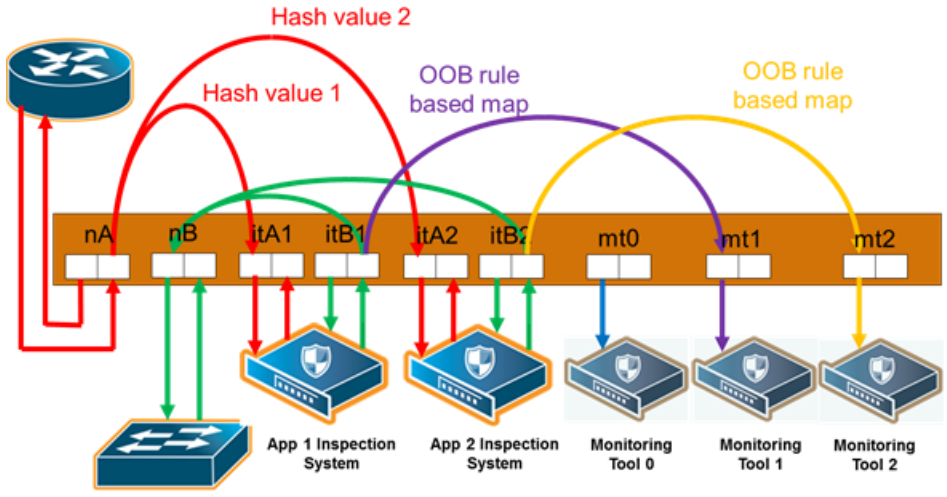


Figure 108*Out-of-Band Rule-Based Maps*

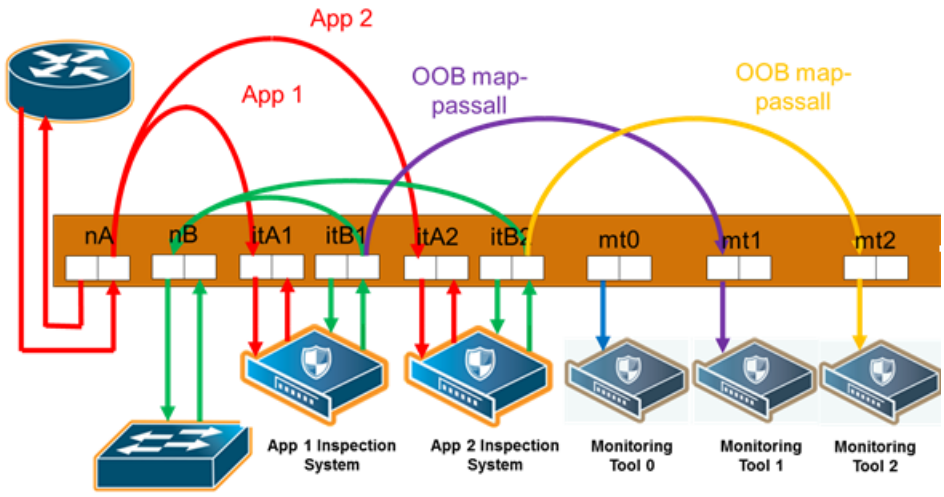


Figure 109*Out-of-Band Map Passalls*

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.

Configure Inline Bypass

The steps to configure inline constructs are only applied to GigaVUE-FM,GigaVUE HC Series® HC Series nodes. In a cluster environment, these steps are only applied to GigaVUE-FM,GigaVUE HC Series® HC Series nodes through the cluster leader.

The inline constructs must all be configured on one GigaVUE-HC3, GigaVUE-HC1-Plus, or one GigaVUE-HC1 node, not across nodes, even if the nodes are in a cluster.

In addition to the steps listed, a map passall, map, or map shared collector is part of the configuration. For more information, refer to Associate Inline Networks with Inline Tools Using Inline Maps.

For out-of-band maps, refer to Out-of-Band (OOB) Map.

For the order in which to configure the software constructs associated with inline bypass solutions, refer to Associate Inline Networks with Inline Tools Using Inline Maps.

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 110Configuration Steps for Inline Bypass Solutions.

Figure 110 *Configuration Steps for Inline Bypass Solutions*

Configuration Step Details

This section provides detailed steps for configuring inline bypass.

Configure Inline Network Ports

Use the following procedure to create inline network ports:

1. Go to **System > Ports > Ports > All Ports.**
2. Open Quick Port Editor by clicking the **Port Editor** button.
3. Use the Quick search field to Find the ports to configure.
4. In th an Alias field enter a name to help identify the inline port.
5. For Type, select **Inline Network**.

   An Inline Network (unprotected) is a software arrangement of two network-type ports allocated to facilitate access to a bidirectional link between two networks (far end network devices) that are linked to inline tool ports.
6. Click **OK**.

NOTE: Any available network-type ports on a Gigamon node can be used to form an unprotected inline network.

Configure Inline Network (Unprotected)

An Inline Network Group is an arrangement of multiple inline network ports to which traffic is distributed based on calculated hash values used by a Gigamon node.

Perform the steps to configure an inline network:

1. After configuring the Inline Network ports, select **Inline Bypass > Inline Networks**.
2. Click **New**.
3. In the **Alias** field, enter an alias for the inline network to help identify the inline network.
4. From the **Port A** drop-down list select an inline network port.
5. An inline network port is automatically selected for **Port B**. To select a different port, select one from the **Port B** drop-down list if there is more than one inline network port.
6. Select a traffic path from the **Traffic** Path drop-down list. The types of traffic paths are:
    - **Bypass** — All traffic arriving at the Port A inline network port is directly forwarded to the Port B inline network port and all traffic arriving at the Port B inline network port is directly forwarded to the Port A inline network port.
    - **Drop** — No traffic is exchanged through the inline network ports (all traffic coming to these ports is dropped).
    - **ByPass with Monitoring** — All traffic is forwarded as a forced bypass value and a copy of the traffic is also forwarded to the inline tools. A traffic map must first be configured between the inline network and the inline tool to have the traffic forwarded with no traffic taken from the inline tools.
    - **To Inline Tool** — The traffic received at the inline network ports is forwarded based on:

1. The traffic map between the inline network and the respective inline tools.

    2. The failover action attributes of the inline tools

    3. The health state of the inline tools.

7. Select the **Link Propagation** check box to enable whether the inline network link on one side of the inline network gets propagated to the other side.
8. Click **Save**.

Configure Inline Network Group

An Inline Network Group is an arrangement of multiple inline network ports to which traffic is distributed based on calculated hash values used by Gigamon node.

Perform the following steps to configure an Inline Network Group:

1. After configuring the Inline Network Ports, select **Inline Bypass > Inline Network Groups**.
2. Click **New**.
3. In the **Alias** field, enter a name for the network group.
4. From the **Inline Network** drop-down list, select the Inline Network ports.
5. Click **Save**.

Create Heartbeat Profile

The Heartbeat Profile is a data structure that contains the heartbeat attributes that are applied to an Inline Tool for configuring its heartbeat. The Create Heartbeat Profile wizard allows you to apply attribute values for the heartbeat profile. Use the following procedure:

1. Select **Inline Bypass > Heartbeats**.
2. Click **New**.
3. In the **Alias** field, enter a name for your heartbeat profile.
4. In the **Type** field, select **Regular**.

   For details about regular heartbeats, refer to Heartbeats.
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 Heartbeats. Heartbeats

     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 Heartbeats.
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 Heartbeats.
6. Use the **Direction** drop-down menu to select the direction for sending heartbeat. The directions are:
   - **A to B**—From Port A to Port B of the inline tool.
   - **B to A**—From Port B to Port A of the inline tool.
   - **Bi-directional**—Both directions.
7. In the **Period** field, enter a number in milliseconds for sending subsequent negative heartbeat packets. The acceptable range is 30 to 5000 milliseconds. The default is 1000 milliseconds.
8. In the **Recovery Time** field, enter the minimum number of seconds since the last negative heartbeat packet is received to declare that the inline tool is up.

The inline tool is up from the standpoint of the negative heartbeat if the negative heartbeats sent are *not* received. When a tool is declared down, sent heartbeats should not be received for a number of seconds in order to declare the tool as being up.

The acceptable range for the Recovery Time field is 5 to 60 seconds. The default is 30 seconds.

9. Click **Save**.

   The heartbeat profile appears in the Heartbeat Profile table.

   NOTE: Highlight the heartbeat profile and click **Edit** to modify the parameters, if needed.

Configure Inline Tool Ports

Use the following procedure to create inline tool ports:

1. Go to **System > Ports > Ports > All Ports.**
2. Open Quick Port Editor by clicking the **Port Editor** button.
3. Use the Quick search field to Find the ports to configure.
4. In th an Alias field enter a name to help identify the inline port.
5. For **Type**, select **Inline Tool**.

   An Inline Tool represents a pair of inline tool ports.
6. Click **OK**.

Configure Inline Tool

An Inline Tool represents a pair of inline tool ports. To configure an Inline Tool, do the following:

1. Select **Inline Bypass > Inline Tools**.
2. Click **New**.
3. If needed, click **Port Editor** to open the Quick Port Editor to configure the inline tool ports.
4. Select the inline tool ports for the inline tool.
   - For **Port A**, select an inline tool port configured in the previous procedure, Configure Inline Tool Ports.
   - For **Port B**, select another inline tool port configured in the previous procedure, Configure Inline Tool Ports.
5. Select **Enabled** to set the inline tool ports as enabled for inline bypass traffic.
6. For **Failover action**, select one of the following:

- **Tool Bypass** — When the inline tool fails all traffic coming to the respective inline tool is directed via the bypass path.
- **Network Bypass** — When the inline tool fails the traffic is directed to multiple inline tools associated with an inline network or inline network group using rule-based inline maps.
- **Tool Drop** — When the inline tool fails all traffic coming to the respective inline tool is dropped.
- **Network Drop**—When the inline tool fails all traffic coming to the respective inline tool is dropped.
- **Network Port Forced Down**—When the inline tool fails the inline network ports of the respective inline network are forced as "down".

7. Select the **Recovery Mode**. The recovery mode can be one of the following:
   - **automatic**—Specifies automatic recovery, which redirects traffic back to the inline tool as soon as it has recovered from all faulty conditions.
   - **manual**—Specifies manual recovery, which lets you control when to put an inline tool back into service after the tool has recovered. For example, you may want to wait for a maintenance window to return the inline tool to service.
8. Select **Enable Heartbeat** to set the heartbeat.

   If the heartbeat is enabled, do the following:

   a. From **Profile**, select the desire heartbeat profile. The available heartbeat profiles are existing profiles from other inline tools. Once a heartbeat profile is selected, its attributes are displayed.

   b. In the **HB IP Address A** field, enter the server's IP address to send the heartbeat packets.

   c. In the **HB IP Address B** field, enter the server's IP address to send a second heartbeat packet.
9. Click **Save**.

Create Inline Tool Group

An Inline Tool Group is an arrangement of multiple Inline Tools to which traffic is distributed based on calculated hash values used by Gigamon node.

To configure an Inline Tool Group, do the following:

1. Use the **Quick Port Editor** to configure the inline tool ports.
2. Select **Inline Bypass > Inline Tool Groups**.
3. Click **New**.
4. In the **Alias** field, enter a name to help identify the inline tool group.
5. For **Inline Tools**, select the inline tool ports for the inline tool group.
6. (Optional) For the **Inline Spare Tool**, select another inline tool port.

If a spare is selected, the inline tool group becomes a redundant arrangement of inline tools. When the first failure occurs in a set of active inline tools, traffic will be forwarded to the spare with no loss, thus the spare will replace the failed tool in the active set.

7. Select **Enable** to enable the inline tool group.
8. Select the **Failover Action**. The failover actions are:
   - **Tool Bypass**—When the inline tool group fails all traffic coming to the respective inline network is directed via the bypass path.
   - **Tool Drop**—When the inline tool group fails all traffic coming to the respective inline network is dropped.
   - **Network Bypass**—When the inline tool group fails all traffic coming to the respective inline network is directed to the inline tool group via the bypass path.
   - **Network Drop**—When the inline tool group fails all traffic coming to the respective inline network group is dropped.
   - **Network Port Forced Down**—When the inline tool group fails the inline network ports of the respective inline network are forced as "down".
9. For **Minimum Healthy Group Size**, select a number that represents the minimal amount of inline tools that are required to have a state of Normal.
10. Select the **Hash** for the inline tool group.

    Hashing, which is used for distributing traffic across the inline tools in an inline tool group. The values for the hash parameter are as follows:
    - **advanced**—Specifies symmetrical hashing, which is derived from the combination of packet fields based on the criteria selected for the advanced-hash algorithm. For inline bypass applications, the most common choice of criteria for the advanced-hash algorithm is the combination of source IP and destination IP addresses. This produces a hash value that sends all traffic associated with the same session to the same inline tool in the inline tool group.
    - **a-srcip-b-dstip**—Specifies asymmetrical hashing, which is derived from the source IP address for side A of the network and the destination IP address for side B of the network. This produces a hash value that sends all traffic associated with the same source address residing on side A to the same inline tool in the inline tool group, regardless of destination or session.
    - **b-srcip-a-dstip**—Specifies asymmetrical hashing, which is derived from the destination IP address for side A of the network and the source IP address for side B of the network. This produces a hash value that sends all traffic associated with the same source address residing on side B to the same inline tool in the inline tool group, regardless of destination or session.
11. Click **Save**.

Configure Inline Tool Series

To configure an Inline Tool Series, do the following:

1. Select **Inline Bypass > Inline Serial Tools**.
2. Click **New**.
3. In the **Alias** field, enter a name to help identify the inline tool series.
4. For **Inline Tools**, select the inline tool ports for the inline tool series.
5. Select **Enabled** to enable this configuration.
6. Select the **Failover action**.

   For details about the failover actions, refer to Inline Tool Series Local Failover Action
7. Select the **Per Direction Order**.

   For details about per-direction order, refer to Inline Tool Series Per-Direction Order.

Configure When GigaVUE-FM,GigaVUE HC Series® HC Series Modules are Operationally Up

Ensure that the GigaVUE-FM,GigaVUE HC Series® 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.

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 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 111Logical Bypass Without Bypass Combo Module which shows a GigaVUE-HC1-Plus.



Figure 111 *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 this example, the ports are 3/1/x1 and 3/1/x2 **4.** Set port 3/1/x1 to Type Inline-Network and select Enable. Enter iN1 for the port alias. **5.** Set port 3/1/x2 to Type Inline-Network and select Enable. Enter iN2 for the port alias. **6.** Make sure Enable is selected for Admin on the ports. **7.** Click OK.
2.	Configure inline network.	**1.** Select Inline Bypass > Inline Networks. **2.** Click New. **3.** In the Alias field, type InNet. **4.** For Port A, select iN1. **5.** For Port B, select iN2. **6.** Click Save.
3.	Configure inline tool ports, port type (inline-tool), and administratively enable inline tool ports.	**1.** Go to **System > Ports > Ports > All Ports**. **2.** Click Quick Port Editor. **3.** Use Quick search to find the ports to configure. In this example, the ports are 3/1/x3 and 3/1/x4 **4.** Set port 3/1/x4 to Type Inline-Tool and select Enable. Enter iT1 for the port alias. **5.** Set port 3/1/x4 to Type Inline-Tool and select Enable. Enter iT2 for the port alias. **6.** Make sure Enable is selected for Admin on the ports. **7.** Click OK.

Task	Description	UI Steps
4.	Configure inline tool, and enable it. Also enable the default heartbeat profile.	**1.** Select Inline Bypass > Inline Tools. **2.** Click New. **3.** In the Alias field, type InTool. **4.** For Port A, select iT1. **5.** For Port B, select Ii2. **6.** Under Configuration: **a.** Select Enabled **b.** Select Enabled Heartbeat and set Profile to default. **7.** Click Save.
5.	Configure map passall, from inline network to inline tool.	**1.** Select Maps > Maps. **2.** Click New. **3.** In the Alias field, type InMap. **4.** Select Inline for Type and Pass All for Subtype. **5.** For Source, select InNet **6.** For Destination, select InTool. **7.** Click Save.
6.	Configure the path of the traffic to inline tool.	**1.** Select Inline ByPass > Inline Networks. **2.** Select the Inline Network InNet and click Edit **3.** Under Configuration, set the Traffic Path field to To Inline Tool. **4.** Click Save.

Example 2: Unprotected Inline Bypass with Default Heartbeat

Example 2 adds the default heartbeat profile to the unprotected inline bypass solution in Example 1.

Task	Description	UI Steps
1.	Configure inline network aliases, port type (inline-network), and administratively enable inline network ports.	**5.** Go to **System > Ports > Ports > All Ports**. **6.** Click Quick Port Editor. **7.** Use Quick search to find the ports to configure. In this example, the ports are 3/1/x1 and 3/1/x2 **8.** Set port 3/1/x1 to Type Inline-Network and select Enable. Enter iN1 for the port alias. **9.** Set port 3/1/x2 to Type Inline-Network and select Enable. Enter iN2 for the port alias. **10.** Make sure Enable is selected for Admin on the ports. **11.** Click OK.
2.	Configure inline network.	**1.** Select Inline Bypass > Inline Networks. **12.** Click New. **13.** In the Alias field, type InNet. **14.** For Port A, select iN1. **15.** For Port B, select iN2. **16.** Click Save.
3.	Configure inline tool ports, port type (inline-tool), and administratively enable inline tool ports.	**1.** Go to **System > Ports > Ports > All Ports**. **17.** Click Quick Port Editor. **18.** Use Quick search to find the ports to configure. In this example, the ports are 3/1/x3 and 3/1/x4 **19.** Set port 3/1/x4 to Type Inline-Tool and select Enable. Enter iT1 for the port alias. **20.** Set port 3/1/x4 to Type Inline-Tool and select Enable. Enter iT2 for the port alias. **21.** Make sure Enable is selected for Admin on the ports. **22.** Click OK.
4.	Configure default heartbeat profile.	In GigaVUE-FM,GigaVUE HC Series-HVUE, the default heartbeat profile is already configured. To view the profile, select Inline Bypass > Heartbeats.

Task	Description	UI Steps
5.	Configure inline tools, and enable it. Also enable the default heartbeat profile.	**1.** Select Inline Bypass > Inline Tools. **23.** Click New. **24.** In the Alias field, type InTool1. **25.** For Port A, select iT1. **26.** For Port B, select iT2. **27.** Under Configuration: **a.** Select Enabled **b.** Select Enabled Heartbeat and set Profile to default. **28.** Click Save.
6.	Configure map passall, from inline network to inline tool.	**1.** Select Maps > Maps. **29.** Click New. **30.** In the Alias field, type InMap. **31.** Select Inline for Type and Pass All for Subtype. **32.** For Source, select InNet **33.** For Destination, select InTool. **34.** Click Save.
7.	Configure the path of the traffic to inline tool.	**1.** Select Inline ByPass > Inline Networks. **35.** Select the Inline Network InNet and click Edit **36.** Under Configuration, set the Traffic Path field to To Inline Tool. **37.** Click Save.

Example 3: Protected Inline Bypass Using Combo Modules

Example 3 is a protected inline bypass solution using bypass combo modules on GigaVUE-HC3. It also configures heartbeat and negative heartbeat profiles.

Protected inline networks are based on the pairs of ports associated with the physical protection switches located on the bypass combo modules. Unlike the unprotected examples, you do not need to configure inline network ports because they are created automatically. On GigaVUE-HC3, the port pairs are numbered for example: 2/2/x17 and 2/2/x18, 2/2/x19and 2/2/x20, 2/2/x21 and 2/2/x22, 2/2/x23 and 2/2/x24.

You do not need to configure inline networks because they are also created automatically on bypass combo modules. The aliases of the default inline networks are: default_inline_net_2_2_1, default_inline_net_2_2_2, default_inline_net_2_2_3, default_inline_net_2_2_4.

On GigaVUE-HC3, protected inline bypass can be configured on the bypass combo module on ports c1..c4.

On GigaVUE-HC1, protected inline bypass can be configured on the bypass combo module. It can also be configured on the TAP-HC1-G10040 module placed in either bay 2 or bay 3, so the ports will be 1/2/g1..g8 or 1/3/g1..g8.

NOTE: The default value of the physical-bypass attribute of protected inline networks is set to enable, which means that the fibers attached to ports net-a and net-b of the inline network are optically coupled and the traffic is exchanged between end nodes without coming to the switching fabric. of the GigaVUE-FM,GigaVUE HC Series node. As shown in Example 4, after configuring the inline tool and the map passall, the physical-bypass attribute is set to disable in order to activate the Inline Bypass solution.

Task	Description	UI Steps
1.	Configure inline too aliases, port type (inline-tool), and administratively enable inline network ports.	**1.** Go to **System > Ports > Ports > All Ports**. **38.** Click Quick Port Editor. **39.** Use Quick search to find the ports to configure. In this example, the ports are 2/2/x11 and 2/2/x12 **40.** Set port 2/2/x11 to Type Inline Tool and select Enable. Enter iT1 for the port alias. **41.** Set port 2/2/x12 to Type Inline Tool and select Enable. Enter iT2 for the port alias. **42.** Click OK.
2.	Configure a heartbeat profile.	**1.** Select Inline Bypass > Heartbeats. **43.** Click New. **44.** In the Alias field, type hb2. **45.** For Type, select Regular. **46.** Click Save.
3.	Configure negative heartbeat profile alias and PCAP file	**1.** Select Inline Bypass > Heartbeats. **47.** Click New. **48.** In the Alias field, type nhb1. **49.** For Type, select Negative. **50.** Click the Browse button for Custom Format and upload the pcap file; for example, hnb.pcap. **51.** Click Save.
4.	Configure inline tools, and enable them. Also specify the heartbeat profile.	**1.** Select Inline Bypass > Inline Tools. **52.** Click New. **53.** In the Alias field, type InTool1. **54.** For Port A, select iT1. **55.** For Port B, select iT2. **56.** Under Configuration:

Task	Description	UI Steps
		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. **57.** Click Save. Configure the second inline tool. **1.** Select Inline Bypass > Inline Tools. **58.** Click New. **59.** In the Alias field, type InTool2. **60.** For Port A, select iT3. **61.** For Port B, select iT4. **62.** Under Configuration: **a.** Select Enabled **b.** Select Enabled Heartbeat and set Profile to hb_custom. **63.** Click Save.
5.	Configure the inline tool group and enable it.	**1.** Select Inline Bypass > Inline Tool Group **64.** Click New. **65.** In the Alias field, type inToolGroup. **66.** For Inline Tools, select InTool1 and InTool2. **67.** Under Configuration, select Enabled. **68.** Click Save.
6.	Configure map passall, from inline network to inline tool.	**1.** Select Maps > Maps. **69.** Click New. **70.** In the Alias field, type InMap. **71.** Select Inline for Type and Pass All for Subtype. **72.** For Source, select default_inline_net_2_2_1 **73.** For Destination, select InTool1 **74.** Click Save.
7.	Configure the path of the traffic to inline tool and disable the physical bypass	**1.** Select Inline ByPass > Inline Networks. **75.** Select the Inline Network default_inline_net_2_2_1 and click Edit **76.** Under Configuration, set the Traffic Path field to To Inline Tool. **77.** Make sure Physical Bypass is not selected. **78.** Click Save.

GigaSMART TLS/SSL Decryption for Out-of-Band Tools

Required License: TLS/SSL Decryption for Out-of-Band Tools

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

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

TLS/SSL decryption for out-of-band tools is described in GigaSMART Passive TLS/SSL Decryption. It is supported on all GigaVUE-HC1, GigaVUE-HC1-Plus, and GigaVUE-HC3 devices.

GigaSMART Passive TLS/SSL Decryption

GigaVUE-FM,GigaVUE HC Series H Series nodes support Secure Sockets Layer (SSL) decryption. TLS/SSL is a cryptographic protocol that adds security to (Transmission Control Protocol) TCP/IP communications such as Web browsing and email. The protocol allows the transmission of secure data between a server and client who both have the keys to decode the transmission and the certificates to verify trust between them. Passive TLS/SSL decryption delivers decrypted traffic to out-of-band tools that can then detect threats entering the network.

Passive TLS/SSL decryption is a pillar of the GigaSECURE Security Delivery Platform. For an overview of GigaSECURE, refer to GigaSECURE Security Delivery Platform.

On GigaVUE-FM,GigaVUE HC Series H Series nodes, GigaSMART line cards or modules perform the decryption of TLS/SSL traffic. Using GigaSMART for decryption offloads the decryption function from tools and offers improved tool performance by removing this computationally intensive task. GigaSMART provides a centralized decryption point. Decrypted TLS/SSL traffic can be sent from GigaSMART to inspection tools for further analysis, for example, to look at encrypted communications or to detect malware.

Before TLS/SSL traffic is decrypted, the de-duplication GigaSMART operation can be performed. Decrypted traffic from the GigaSMART line card or module can be filtered, aggregated, and replicated and then sent to one or more monitoring tools for analysis.

Passive TLS/SSL decryption is supported on the following GigaVUE-FM,GigaVUE HC Series H Series products with GigaSMART line cards or modules installed:

- GigaVUE-HC3
- GigaVUE-HC1
- GigaVUE-HC1-Plus
- GigaVUE-HCT

Use Passive TLS/SSL decryption on the GigaSMART line card or module with passive or offline traffic. Tap the traffic to and from a server and pass it to the GigaVUE-FM,GigaVUE HC Series H Series node with the GigaSMART line card or module.

Passive TLS/SSL decryption operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to Groups of GigaSMART Engine Ports for details.

For secure storage of private keys in Passive TLS/SSL decryption scenarios, integration with a Hardware Security Module (HSM) is supported. Both Entrust nShield and Thales Luna Network HSMs are compatible and can be used based on deployment requirements. Refer to Entrust nShield HSM for TLS/SSL Decryption for Out-of-Band Tools for details.

Gigamon also offers inline TLS/SSL decryption, which inspects TLS/SSL encrypted traffic inline. Refer to Work With Inline TLS/SSL Decryption for details.

About Passive TLS/SSL Decryption

TLS/SSL encryption secures traffic between a client and a server, such as a Web server. TLS/SSL decryption uses keys to decode the traffic between the client and server.

SSL and Transport Layer Security (TLS) protocols consist of a set of messages exchanged between a client and server to set up and tear down the TLS/SSL connection between them. To set up the connection, the client and server use the Public Key Infrastructure (PKI) to exchange the bulk encryption keys needed for data transfer.

Figure 112Basic TLS/SSL Handshake shows the basic TLS/SSL handshake between a client and server to establish a session. The messages are unencrypted up to step 6 in Figure 112Basic TLS/SSL Handshake. The messages are encrypted after step 6, including the step 9 Finished message.

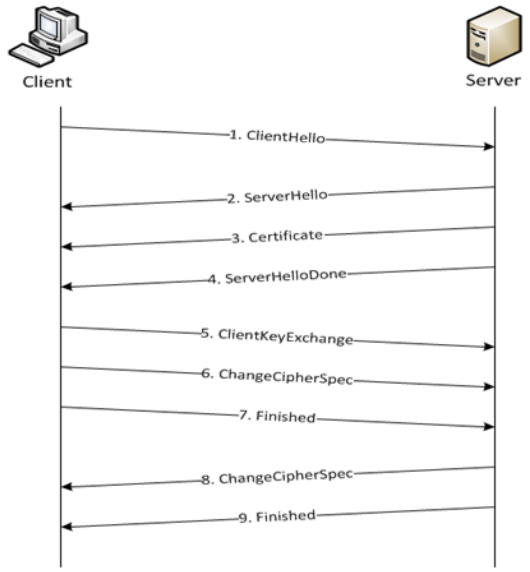


Figure 112 *Basic TLS/SSL Handshake*

Once a session has been established, the keys are saved so a session can be resumed efficiently later. Figure 113Resumed TLS/SSL Handshake shows the resumed TLS/SSL handshake, with fewer steps.

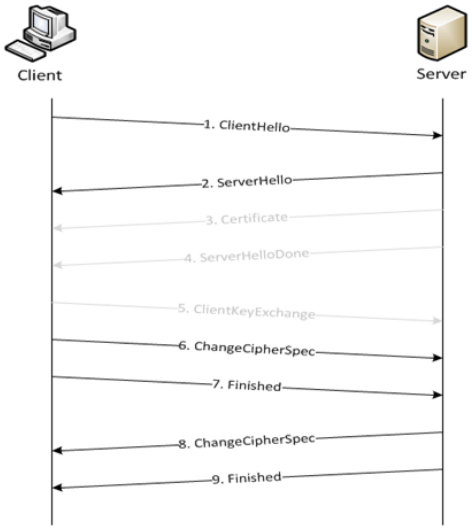


Figure 113 *Resumed TLS/SSL Handshake*

Passive TLS/SSL decryption can be deployed close to the server, as shown in Figure 114Inbound (Server Side).

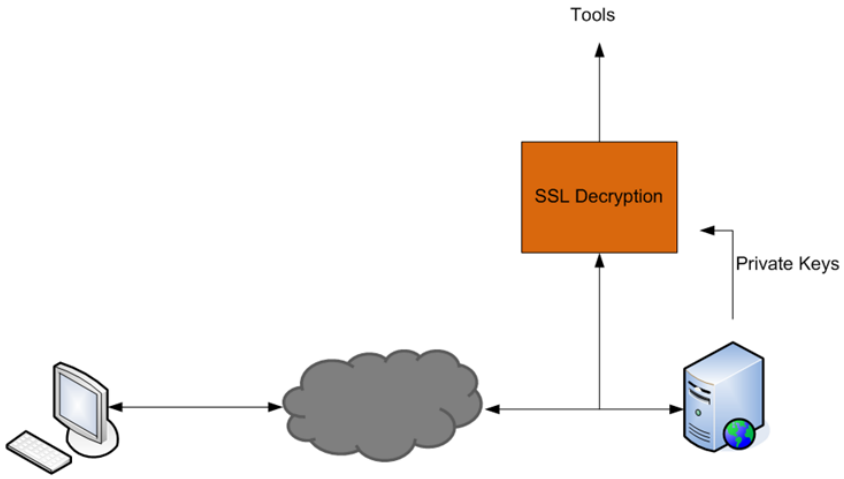


Figure 114 *Inbound (Server Side)*

Passive TLS/SSL decryption can also be deployed close to an TLS/SSL proxy, with the server in the Enterprise domain as shown in Figure 115Outbound (Client/Enterprise Side).

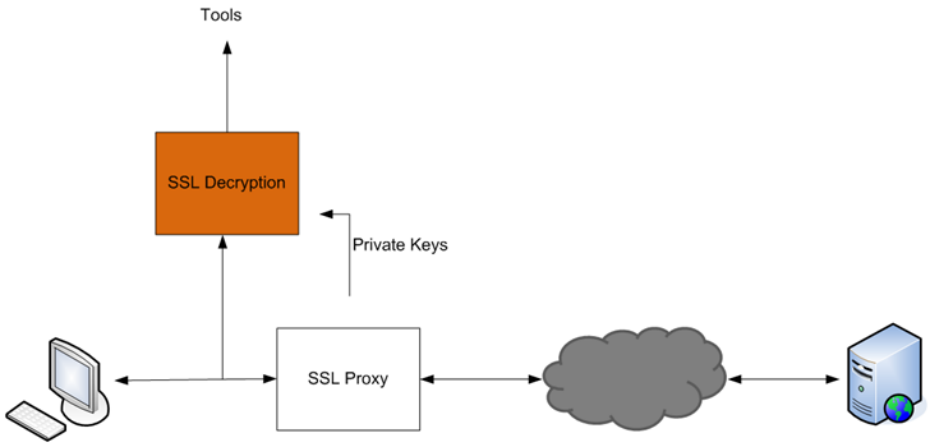


Figure 115 *Outbound (Client/Enterprise Side)*

In Figure 114Inbound (Server Side), tap traffic to the server and then send it for decryption. In Figure 115Outbound (Client/Enterprise Side), tap traffic to the proxy and then send it for decryption. You can have a deployment with either a server or a proxy, but not both.

The following sections describe Passive TLS/SSL decryption on GigaSMART:

- Supported Protocols, Algorithms, and Ciphers
- Limitations
- Create and Reset TLS/SSL Keychain Passwords
- Work with Keys and Services

Supported Protocols, Algorithms, and Ciphers

The supported protocols are as follows:

- SSL 3.0
- TLS 1.0
- TLS 1.1
- TLS 1.2

The supported authentication (Au) is as follows:

- RSA

The supported key exchange (Kx) is as follows:

- RSA

The supported encryption algorithms (Enc) are as follows:

- NULL
- RC4
- DES
- 3DES
- AES (including GCM mode)
- CAMELLIA
- SEED
- IDEA

The supported compression algorithm is as follows:

- NULL

The supported digest algorithms are as follows:

- MD5
- SHA1
- SHA2

The supported key sizes are 128, 256, 512, 1024, 2048, and 4096.

The supported TLS extensions are as follows:

- Extended Master Secret, RFC 7627
- Encrypt-then-MAC, RFC 7366

The supported ciphers are listed in Table 22: Supported Ciphers for Passive SSL decryption.

Table 22: Supported Ciphers for Passive SSL decryption

Cipher Name	Kx	Au	Enc	Bits	Mac
TLS_RSA_WITH_NULL_MD5	RSA	RSA	NULL	0	MD5
TLS_RSA_WITH_NULL_SHA	RSA	RSA	NULL	0	SHA
TLS_RSA_EXPORT_WITH_RC4_40_MD5	RSA_EXPORT	RSA_EXPORT	RC4_40	40	MD5
TLS_RSA_WITH_RC4_128_MD5	RSA	RSA	RC4_128	128	MD5
TLS_RSA_WITH_RC4_128_SHA	RSA	RSA	RC4_128	128	SHA
TLS_RSA_EXPORT_WITH_RC2_CBC_40_MD5	RSA_EXPORT	RSA_EXPORT	RC2_CBC_40	40	MD5
TLS_RSA_WITH_IDEA_CBC_SHA	RSA	RSA	IDEA_CBC	128	SHA
TLS_RSA_EXPORT_WITH_DES40_CBC_SHA	RSA_EXPORT	RSA_EXPORT	DES40_CBC	40	SHA
TLS_RSA_WITH_DES_CBC_SHA	RSA	RSA	DES_CBC	56	SHA
TLS_RSA_WITH_3DES_EDE_CBC_SHA	RSA	RSA	3DES_EDE_CBC	168	SHA
TLS_RSA_WITH_AES_128_CBC_SHA	RSA	RSA	AES_128_CBC	128	SHA
TLS_RSA_WITH_AES_256_CBC_SHA	RSA	RSA	AES_256_CBC	256	SHA
TLS_RSA_WITH_CAMELLIA_128_CBC_SHA	RSA	RSA	CAMELLIA_128_CBC	128	SHA
TLS_RSA_WITH_CAMELLIA_256_CBC_SHA	RSA	RSA	CAMELLIA_256_CBC	256	SHA
TLS_RSA_WITH_SEED_CBC_SHA	RSA	RSA	SEED_CBC	128	SHA
TLS_RSA_WITH_NULL_SHA256	RSA	RSA	NULL	0	SHA256
TLS_RSA_WITH_AES_128_CBC_SHA256	RSA	RSA	AES_128_CBC	128	SHA256
TLS_RSA_WITH_AES_256_CBC_SHA256	RSA	RSA	AES_256_CBC	256	SHA256
TLS_RSA_WITH_AES_128_GCM_SHA256	RSA	RSA	AES_128_GCM	128	SHA256
TLS_RSA_WITH_AES_256_GCM_SHA384	RSA	RSA	AES_256_GCM	256	SHA384

All algorithms used for Passive TLS/SSL decryption are FIPS 140-3 compliant.

All key URLs must point to an RSA private key stored in the PEM or PKCS12 format, as follows:

- http://keyserver.domain.com/path/keyfile.pem
- https://keyserver.domain.com/path/keyfile.pem
- ftp://keyserver.domain.com/path/keyfile.pem
- tftp://keyserver.domain.com/path/keyfile.pem
- scp://username[:password]@keyserver.domain.com/path/keyfile.pem

The supported applications are as follows:

- HTTPS
- FTPS

- SMTP, IMAP, and POP3 with StartTLS

Limitations

The limitations of Passive TLS/SSL decryption are as follows:

- Only IPv4
- Only regular maps; no virtual ports (vports)
- Only combined with the de-duplication GigaSMART operation
- Only one private key per PKCS12 file
- Only server-side authentication
- Only the protocols and ciphers listed in Supported Protocols, Algorithms, and Ciphers.

  NOTE: If an TLS/SSL session cannot be decrypted due to having a non-supported protocol or cipher and if the GigaSMARTS Parameter **SSL Decryption** has **Decrypt Fail Action** is set to **Pass to Tool Port**, the packets will be forwarded to the tool without decryption . Non-supported ciphers and protocols include SSL 2.0, Diffie-Hellman (DHE keys), Ephemeral keys, Elliptic Curves Extension, compression, and 8K key size.

Licensing

The GigaSMART license for Passive TLS/SSL decryption is installed as any other license.

There are no limits to the number of Passive TLS/SSL decryption sessions or the number of users.

Create and Reset TLS/SSL Keychain Passwords

To perform the configuration in the following section, you must have an admin level access role.

Before uploading keys or configuring TLS/SSL, you must create an SSL keychain password. The password is used to encrypt the private keys that you upload to the node.

Keychain passwords are not saved on the node. Refer to Notes about Private Keys and Passwords.

NOTE: When uploading TLS/SSL keys, make sure that you are not creating a duplicate key. Adding a duplicate key can cause errors.

To create an TLS/SSL keychain password, use the following steps:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**.
2. In the Physical Nodes page, select the node for which you want to create the Key Store password.
3. Go to **GigaSMART**>**Passive SSL** >**Key Store**.
4. Click **Actions**>**Keychain Password**.
5. Enter a password in the **Password** and **Confirm Password** fields.

   You can only configure a strong password. A strong password should include at least eight (8) or more characters (up to 64) and include the following:
   - one uppercase letter
   - one lowercase letter
   - one numerical character
   - one special character
6. Enable the **Auto login** check box to let GigaVUE-FM,GigaVUE HC Series unlock the key store when the node reboots.
7. Click **Submit**.

After keys are installed on the node, you will be prompted to enter the password after any login as well as after a node reboot, for example:

If you are a user who does not have an admin level access role, when you enter the configure terminal mode, the following message is displayed:

Password required. Please contact administrator.

If you are a user with an admin level access role, but you enter an incorrect password, the following message is displayed:

Password does not match. Please reenter the password

If an TLS/SSL keychain password is lost, it can be reset, but all existing private keys will be revoked. When there are keys installed on the node, a warning is displayed before you are prompted for the new password.

Once you have a new password, you will have to upload the keys again.

Work with Keys and Services

This section describes working with private keys as well as services. Keys must be uploaded to the GigaVUE-FM,GigaVUE HC Series H Series node using a unique alias. Services must be defined for each server destination that needs decryption. There are two types of decryption. If the service is associated with a certificate then it is called

service based decryption. If only certificate is added to the node without creating service then it is key certificate based decryption.

To perform the configuration in the following section, you must have an admin level access role.

Encrypted private keys are saved on the node. Refer to Notes about Private Keys and Passwords.

NOTE: When uploading TLS/SSL keys, make sure that you are not creating a duplicate key. Adding a duplicate key can cause errors.

Set Up Key Store Certificate Management for Passive TLS/SSL

A Key Store certificate can be setup to be auto- enabled, auto-deleted and auto-retained to a passive TLS/SSL service. The configuration can be done as follows:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**.
2. In the Physical Nodes page, select the node for which you want to configure Key Store settings.
3. Go to **GigaSMART**>**Passive SSL** >**Key Store**.
4. Click **Settings**. Configure the following settings:
   - **Auto Enable New Certificates** – This setting is not applicable for passive TLS/SSL service.
   - **Auto Delete Expired Certificates** – This option allows you to delete the expired certificates automatically. This setting will be triggered once a day at 12:00:00 UTC. Specify the number of days to retain an expired certificate in the **Number of days to retain expired certificates** field. The default value would be 30 days.
   - **Auto Delete Certificates with same entity** – This option allows you to automatically delete expired certificates that have a similar name and are associated with a passive TLS/SSL service. When you enable this option, you need to specify the maximum number of certificates to retain for the same entity in the corresponding field. This means that if there are more certificates than the specified number, the oldest one will be deleted. This helps you manage your certificates and avoid cluttering your system with unnecessary or redundant certificates.

Upload TLS/SSL Private Keys

To upload an TLS/SSL private key, do the following:

1. From the device view, select **GigaSMART >Passive SSL > Key Store** to open the Key Store page.
2. Click **New**. The SSL Key page appears.
3. In the **Add Key** page, enter the following details:
   - For **Key Alias**, enter an alias for the SSL key.
   - For **Key Upload Type**, select **PEM**, **PKCS12**,**PKCS11**,**nShield-HSM,Luna-HSM**.
   - (optional) For **Passphrase**, enter a passphrase for the key.
   - Select a **Private Key** by pasting the copied key in PEM format or installing from URL or installing from local directory.
   - Select a **Certificate** by pasting the copied key in PEM format or installing from URL or installing from local directory.

   NOTE:
   - You can either use the server based private key or the key certificate based private key to decrypt a passive TLS/SSL traffic.
   - If the Default Service is enabled in TLS/SSL Service, you cannot use the server certificate to decrypt Passive TLS/SSL traffic.
4. Click **OK**.

If you choose to use Venafi, Inc. as your Electronic Key Management System (EKMS), ensure that you have the GigaVUE-FM host address and credentials to push the generated key into the selected nodes. For details on pushing keys from Venafi, see the GigaVUE-FM/Venafi Trust Protection Platform Integration Guide.

Delete TLS/SSL Key

To delete a particular SSL private key, select the key on the TLS/SSL Keys page, and then select Delete. To delete all TLS/SSL private keys, select multiple keys.

Display Key Store Details for Passive TLS/SSL

The key store certificates added would be displayed in the Key Store page.

To access the Key Store page:

1. From the left navigation pane, go to **Inventory > Physical > Nodes**.
2. Select the node for which you want to view the key store certificate information.
3. From the left navigation pane, go to **System > GigaSMART > Passive SSL > Key Store**. The details about the key store certificates added for the selected node is displayed.

The following table describes the fields:

Component	Description
Key Alias	The alias name of the Key certificate.
Type	Defines whether the Key Store is a Certificate or a Private Key .
Health Status	The health indicator of a certificate used in traffic flow. The three major indicators with their respective color legend are as follows: Green - The key certificate is attached to a passive TLS/SSL service and the service is part of the passive TLS/SSL GSOP, which is used in a traffic map. This will also indicate a certificate which is being used as a signing CA in outbound deployment . Blue - The key certificate is not actively participating in any traffic flow. Red - The key certificate has expired .
Common name	A common name given to group the key based on domain.
Organization	Organization name that provided the key .
.**Organization Unit**	Organization unit name that provided the key.
Expiry Date	Date on which the key certificate would get expired.
Installed On	Date on which the key certificate was installed.
Description	Description or additional information about the key certificate.
Status	Status of the key certificate. The valid values are: Expiring—The key certificate is nearing the expiry date. Expired—The key certificate has expired.

You can control the key store certificates display by utilizing the filters provided.

Create TLS/SSL Service

After you have uploaded a private key, you can add a service. A service maps to a physical server, such as an HTTP server. One server can run multiple services. A service is a combination of an IP address and a server port number. Also, the key and the service must be tied together.

Prerequisites

Before creating a service, you must do the following:

- Upload a private key as described in Upload TLS/SSL Private Keys
- Create GigaSMART Group with TLS/SSL Decryption enabled.

To create a service, do the following:

1. From the device view, select **GigaSMART > Passive SSL > SSL Services**. The SSL Services page appears.
2. Click **New**.
3. On the Create SSL Service configuration page, do the following:
    - Enter an alias.
    - Click on the **Default Service** (inbuilt key will be used for decryption for any service) check box to allow GigaVUE-FM to enable default service.

NOTE: When opting for the Default service along with service based decryption, other than the particular service that is opted for service based, other connections will be decrypted using the key available on the default service.

    - Enter the information for the service: IP Address, Server Port.
    - Select the alias of SSL Key previously uploaded. For the steps, refer to Uploading TLS/SSL Private Keys.
    - Select the GigaSMART Group with TLS/SSL decryption enabled to associate with this TLS/SSL service.
4. Click **Apply**.

Delete TLS/SSL Service

To delete a particular TLS/SSL service select the service on the TLS/SSL Services page, and then select Delete. To delete all TLS/SSL services, select multiple keys.

Notes about Private Keys and Passwords

Consider the following notes about private keys and passwords:

- Encrypted private keys are stored on the node. When a private key is uploaded, it is encrypted with a password before it is stored, therefore keys are password-protected. Keychain passwords are not stored on the node.
- Because only encrypted private keys are stored on the node and because the keychain password is not stored on the node, after any node reboot you will be prompted to enter the password. Until the password is entered, Passive TLS/SSL decryption is not working.
- Key content cannot be displayed.

- Keys that are synchronized across a cluster are encrypted.

ECODES for Troubleshooting Passive TLS/SSL Decryption

Use the following table of ECODE messages to assist with troubleshooting Passive TLS/SSL decryption:

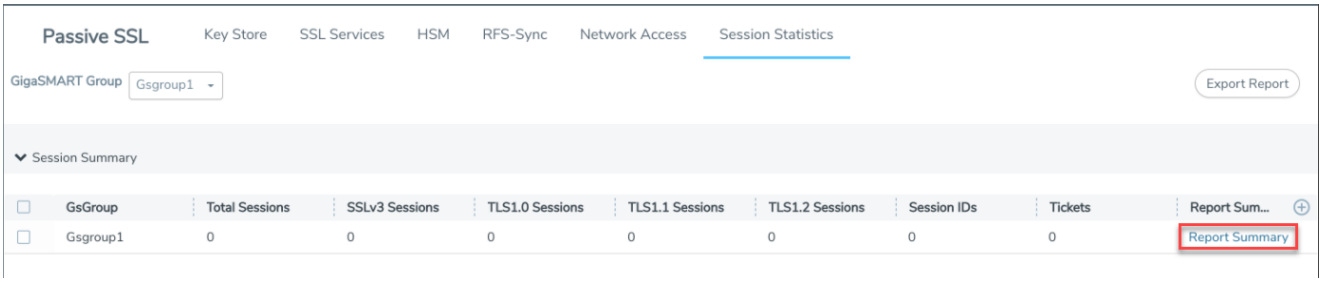
ECODE	Description
81	TCP flow errors detected. Make sure you see the complete TCP flow. Use the de-duplication GigaSMART operation with Passive TLS/SSL decryption.
103	Session limit reached. The session table has been exhausted. If the session timeout (**session-timeout**) value is large, lower it.
104	Key/ticket cache limit reached. The allocated cache entries have been used up. If the timeout (**key-cache-timeout** or **ticket-cache-timeout**) value is large, lower it.
206	No server info. A flow has been received for which service-key mapping is not defined.
213	Packets for missed TCP handshake. Packets were received for TCP flows that do not exist. If the device was just started, this should trend down quickly.
218	Unknown TLS/SSL version. An TLS/SSL handshake processing error occurred. Use the de-duplication GigaSMART operation with Passive TLS/SSL decryption.
221	Unknown TLS/SSL version. An unsupported TLS/SSLv2 handshake was seen.
222	Protocol error. An unsupported protocol version was seen.
225	Unsupported cipher. The cipher suite cannot be decrypted.
226	Pre-master secret error. Check that the private key is correct and that the session is complete.
228	Generic decryption error. Usually indicates errors in the handshake. Check that you are getting the full session from both sides.
231	Invalid MAC. Likely indicates that invalid or truncated packets have been received.
232	Session not in cache. Indicates that you are trying to decrypt a restarted session where the original negotiation was not seen. These should trend down in time, but if they do not, increase the **key-cache-timeout** value.
237	Cannot decrypt ephemeral key based encryption. One of the Ephemeral/PFS ciphersuites, usually Diffie-Hellman Ephemeral, has been seen. These are not supported.
245	Ticket not in cache. This is usually not an error. Indicates that you are trying to decrypt a restarted session where the original negotiation was not seen. These should trend down in time, but if they do not, increase the **ticket-cache-timeout** value.

View Passive TLS/SSL Decryption Flow Ops Report

GigaSMART provides support for Flow Ops reporting. You can generate the Flow Ops report for Passive TLS/SSL Decryption, view the session summary and session statistics, and export the report to a remote server.

To view the session summary and session details of the Passive TLS/SSL Decryption Flow Ops report:

1. From the device view, go to **System > GigaSMART > Passive SSL > Session Statistics**.
2. From the **GigaSMART Group** drop-down list, select the required GigaSMART group. The Session Summary and Session Details appear. For descriptions of the session statistics, refer to Flow Ops Report Statistics for Passive TLS/SSL Decryption.
3. In the Session Summary table, click the **Report Summary** link to view the graphical representation of the trending data for the Passive SSL Decryption session summary as shown in the following figures:





You can also choose to export the session summary to a remote server. Select the required row in the Session Summary table, and then click **Export Report**. The Upload Flow Ops Report page appears. Enter the remote server path and password to access the server. Click **Upload**. The details are exported to the remote server.

View Certificate Statistics

To view certificate statistics, from the device view, go to **System > GigaSMART > Passive SSL > Certificate Statistics**. The page displays all the SSL server-certificates that are attached to the GigaSMART group.

View Service Statistics

To view service statistics, from the device view, go to **System > GigaSMART > Passive SSL > Service Statistics**. The page displays all the SSL services that are associated with the GigaSMART group.

View Error Statistics

To view error statistics from the device view, go to **System > GigaSMART > Passive SSL > Error Statistics**. The page displays all the SSL error code messages that were received for the device.

Configuring Passive TLS/SSL Decryption Examples - Command Line Reference

The following sections provide examples and commands (CLI) of TLS/SSL decryption. Refer to the following:

- Example 1: TLS/SSL Decryption with a Regular Map
- Example 2: TLS/SSL Decryption with De-duplication
- Other Usage Examples

Example 1: TLS/SSL Decryption with a Regular Map

In Example 1, a regular map is configured to use with the TLS/SSL decryption GigaSMART operation.

Step	Description	Command
1.	Upload a key and create a service. Refer to *Working with Keys and Services* on page 609.	**(config) # apps ssl key alias key1 download type private-key url https://keyserver.domain.com/path/keyfile.pem** **(config) # apps ssl service alias service1 server-ip 192.168.1.1 server-port 443**
2.	Configure a GigaSMART	**(config) # gsgroup alias gsgrp1 port-list 1/1/e1**

Step	Description	Command
	group..	
3.	Specify the GigaSMART group alias.	**(config) # gsparams gsgroup gsgrp1**
4.	Specify a failover action.	**(config gsparams gsgroup gsgrp1) # ssl-decrypt decrypt-fail-action drop**
5.	Configure session timeouts, in seconds.	**(config gsparams gsgroup gsgrp1) # ssl-decrypt pending-session-timeout 60** **(config gsparams gsgroup gsgrp1) # ssl-decrypt session-timeout 300** **(config gsparams gsgroup gsgrp1) # ssl-decrypt tcp-syn-timeout 20**
6.	Configure cache timeouts, in seconds.	**(config gsparams gsgroup gsgrp1) # ssl-decrypt key-cache-timeout 9000** **(config gsparams gsgroup gsgrp1) # ssl-decrypt ticket-cache-timeout 9000**
7.	Configure a key/service mapping that maps how a key is assigned to an IP address of a server.	**(config gsparams gsgroup gsgrp1) # ssl-decrypt key-map add service service1 key key1**
8.	Enable TLS/SSL decryption.	**(config gsparams gsgroup gsgrp1) # ssl-decrypt enable**
9.	Exit the GigaSMART group configuration mode.	**(config gsparams gsgroup gsgrp1) # exit** **(config) #**
10.	Configure a GigaSMART operation for TLS/SSL decryption.	**(config) # gsop alias gdssl1 ssl-decrypt in-port any out-port auto port-list gsgrp1**

In the previous step, **gdssl1** is the alias for a GigaSMART operation, **in-port** specifies the destination port on which to listen, **out-port** specifies the destination port on which to send decrypted traffic, and **port-list** is set to the GigaSMART group alias previously configured. The **in-port** and **out-port** arguments can also be a port number between 1 and 65535.

Next, configure a traffic map, as follows:

Step	Description	Command
1.	Specify a map alias (m1) and specify the map type and subtype.	**(config) # map alias m1** **(config map alias m1) # type regular byRule**
2.	Specify the GigaSMART operation alias (gdssl1) as part of the map. This applies the associated GigaSMART	**(config map alias m1) # use gsop gdssl1**

Step	Description	Command
	functionality to packets matching a rule in the map.	
3.	Specify a map rule.	**(config map alias m1) # rule add pass ipver 4**
4.	Specify the destination for packets matching this map.	**(config map alias m1) # to 1/1/g2**
5.	Specify the source port(s) for this map.	**(config map alias m1) # from 1/1/g1**
6.	Exit the map prefix mode.	**(config map alias m1) # exit** **(config) #**
7.	Display the configuration.	**(config) # show gsop** **(config) # show map** **(config) # show gsparams**

Example 2: TLS/SSL Decryption with De-duplication

In Example 2, the configuration steps are the same except when you configure a GigaSMART operation you send the decrypted traffic to de-duplication for additional filtering, as follows:

```
(config) # gsop alias gdssl1 ssl-decrypt in-port any out-port auto dedup set port-list gsgrp1
```

Other Usage Examples

Two typical usage examples are as follows:

- Use map rules to filter on the IP address of the server and send everything to GigaSMART. Configure a GigaSMART operation to listen on the **in-port** used by the server. The GigaSMART will drop other traffic.
- Use map rules to filter on the IP address of the server and **in-port** and send specific port traffic to the GigaSMART. Configure a GigaSMART operation to listen on **in-port any**.

Entrust nShield and Thales-Luna Network HSM for TLS/SSL Decryption for Out-of-Band Tools(Passive)

Required License: Included with TLS/SSL Decryption for Out-of-Band Tools (Passive)

Starting in software version 5.3, Entrust nShield Hardware Security Module (HSM) is integrated with decryption for Out-of band Tools (Passive SSL/TLS Decryption). Hardware Security Modules (HSMs) are specialized systems that logically and physically safeguard cryptographic operations and cryptographic keys. HSMs protect sensitive data from

being stolen by providing a highly secure operation structure. HSMs are comprehensive, self-contained solutions for cryptographic processing, key generation, and key storage. The hardware and firmware (i.e., software) required for these functions are automatically included in these appliances.

The application could be a web server or a database server, but, in the case of TLS/SSL decryption for out-of-band tools, the application is GigaSMART. The application interfaces with HSM to use the keys that are stored. There must be network connectivity between HSM and the application.

Keys are added to the HSM by an administrator. When an application's key is on HSM, the HSM creates an application key token. The key token is sent to the application. When the application wants to use a key, the application sends the token to HSM, which establishes a session with HSM to use the key. In this way, the use of keys by the application is secure because only key tokens are exchanged.

You can use Remote File System (RFS), a component in the Entrust nShield HSM to store and manage encrypted keys. The RFS helps to automate the key distribution process. You can enable RFS on the GigaVUE-OS device using GigaVUE-FM so that the device can access the encrypted keys stored in RFS. You can synchronize RFS with GigaVUE-OS device to perform a bulk download of the encrypted keys.

Starting from software version 6.8, decryption for Out-of band Tools (Passive SSL/TLS Decryption) is also enhanced to include the Thales-Luna Network HSM support.

Entrust nShield HSM is supported on GigaVUE-HC1, GigaVUE-HC3,Generation 3 GigaSMART card (SMT-HC1-S) and GigaVUE-HCT.

Entrust nShield and Thales-Luna HSM for TLS/SSL Decryption for Out-of-Band Tools—Rules and Notes

Keep in mind the following rules and notes before you configure and use HSM to store and manage keys:

- GigaSMART uses keys that are already stored on the HSM. There is no key generation.
- The key token that is uploaded to GigaSMART can only be in PKCS11 format.
- Only RSA keys (private keys) are supported.
- When TLS/SSL Decryption for Out-of-Band Tools is configured with HSM, only one map and GigaSMART operations will function. If you configure with more than one map will result in undefined behavior.
- The network connectivity between the HSM and GigaSMART must use a static IP address. Do not use DHCP because the IP address needs to remain the same.

> **NOTE:** If the GigaSMART® engine is configured using DHCP, the following issues may arise:
> 1. Whenever a new DHCP IP is assigned to the GigaSMART® engine, the user must delete and re-create the ISSL App and deploy the solution.
> 2. Additionally, the user needs to register the new DHCP IP with the HSM server for client use.

- Only IPv4 addresses are supported.
- Each GigaSMART card that interfaces with the Entrust nShield HSM will use one Entrust nShield license.
- Clustering is not supported.
- Increase the HSM timeout to 5000ms when using 4K size keys for decryption.
- When uploading RSA keys, validity check for protocol mismatch cannot be performed since the private keys are available on the HSM server.
- If a HSM Decryption deployment is modified follow the below steps:
  - Move the Inline Network traffic path to bypass mode.
  - Make the desired deployment change such as:
    - From non-HSM based decryption to Thales-Luna HSM based decryption.
    - From non-HSM based decryption to Entrust nShield HSM based decryption.
    - From Entrust nShield HSM based decryption to Thales-Luna HSM based decryption.
    - From HSM based decryption to non-HSM based decryption.
  - Reboot the GigaSMART card.
  - Move the Inline Network out of bypass mode to 'To inline Tool' mode.
- There should be at least one active Luna HSM in the High Availability to ensure that the decryption is not interrupted.
- Do not attempt to reload the device, if a partition fails in the Luna HSM configuration and at least one active partition is present.

Configure HSM for TLS/SSL Decryption for Out-of-Band Tools

This section provides topics on how to configure and use HSM for TLS/SSL decryption for out-of-band tools:

Topics:

- Configure HSM Group
- Configure Set Key Handler
- Configure Passive TLS/SSL Network Access
- Use RFS to Manage Encrypted Keys
- Entrust nShield and Thales-Luna Network HSM for TLS/SSL Decryption for Out-of-Band Tools (Passive)
- Create a GigaSMART Operation (GSOP)
- Create TLS/SSL Keychain Password
- Upload TLS/SSL Private Keys
- Configure TLS/SSL Service
- Configure Maps

Configure HSM Group

To configure an HSM group ,add at least one HSM appliance by specifying an alias, a static IP address, and port number. Obtain the ESN and KNETI from your HSM administrator.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation pane of the device view on supported devices. Refer to the Access GigaSMART from GigaVUE-FM for details.

To add an HSM appliance, do the following:

1. From the device view, go to **GigaSMART**> **Passive SSL** > **HSM Groups**.
2. Click **New**.

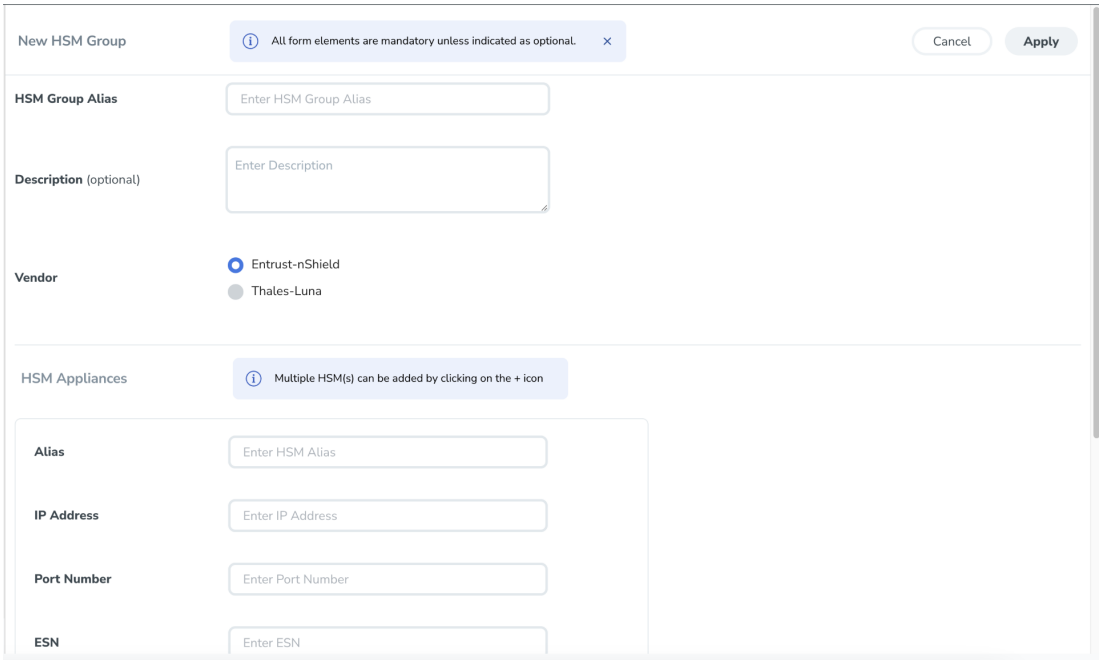


Figure 116*Adding a New HSM Appliance*

3. Enter the details in the **HSM Group Alias** and **Description** fields under the **New HSM Group** section.
4. Select the required vendor type from the options (**Entrust-nShield** or **Thales-Luna**) to create the respective HSM Group. By default, **Entrust-nShield** will be selected as the vendor type.
5. Under the **HSM Appliances** section, in the **Alias** field, enter a name for the HSM appliance.
6. Enter a valid **IP address** and **Port Number**.

NOTE: For Thales-Luna HSM Group , the port group by default is 1792.

7. Configure the following fields for the vendor type **Entrust-nShield**:
   a. Enter the **ESN** (Electronic Serial Number) and **KNETI** that you obtained from the HSM administrator.
   b. Choose one of the following methods to select the required key handler file:
      - **Install from URL**—Enter a valid directory path including the file name and enter the password to access the server.

      NOTE: SCP, SFTP, HTTP, FTP, and TFTP are the supported protocols from where you can select the key handler file.

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

8. Configure the following fields for the vendor type **Thales-Luna**:
   a. Enter the valid username and password in the Server Username and Server Password fields.
   b. Enter the valid details in the Partition Label and Partition Password fields.

Click **Apply**.

NOTE: You cannot configure multiple HSM Groups for Passive SSL with HSM solution.

Configure Set Key Handler

1. From the device view, go to **GigaSMART > Passive SSL > HSM Groups**.
2. Select the **HSM appliance** you just created.
3. Click **Configure** from the **Actions** drop down.
4. Choose one of the following methods to install the key handler file:
   - **Install from URL**—Enter a valid directory path including the file name and enter the password to access the server.

     NOTE: SCP, SFTP, HTTP, FTP, and TFTP are the supported protocols from where you can select the key handler file.

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

   NOTE: Ensure that the file name is "world"

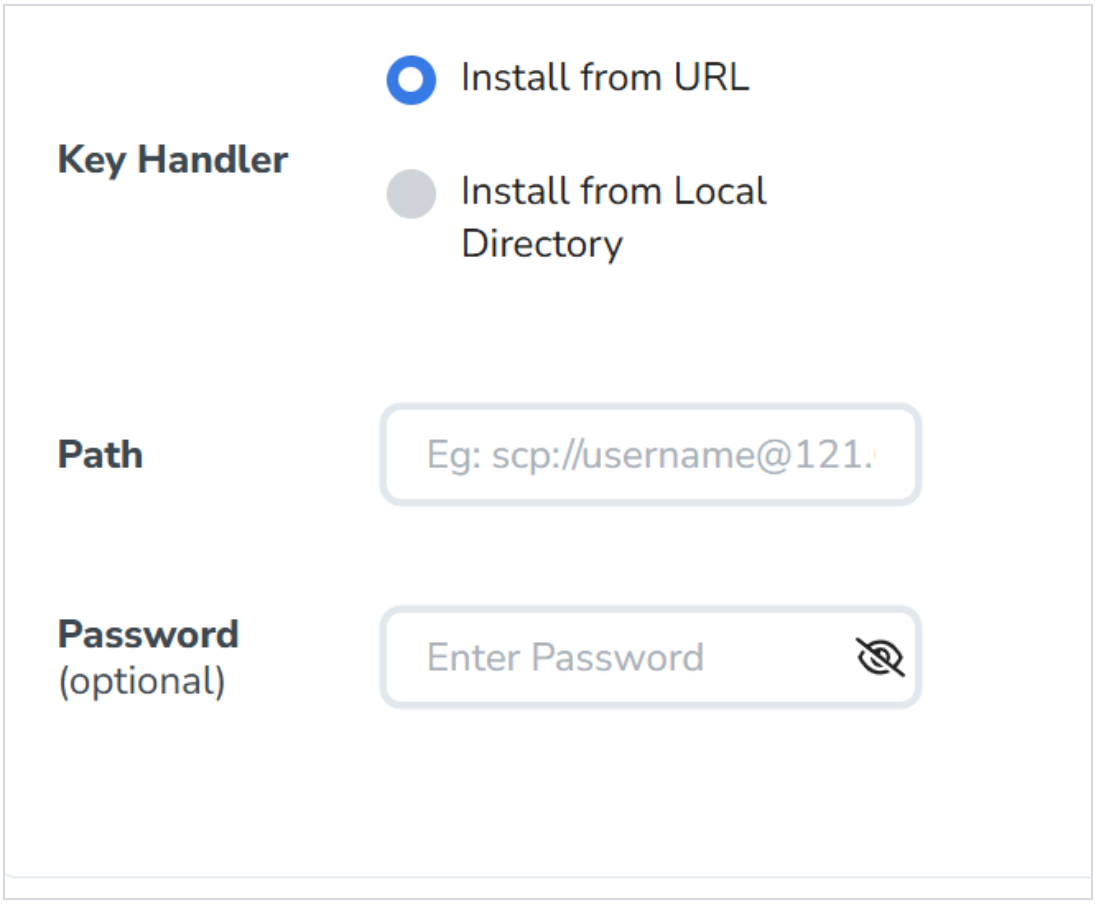


Figure 117*HSM-Configure Key Handler*

5. Click **OK**.

NOTE: Configuring Key Handler is not applicable for a Thales-Luna HSM group.

Configure Passive TLS/SSL Network Access

Each GigaSMART card requires IP address configuration for network access. To configure IP address details:

1. From the device view, go to **GigaSMART > Passive SSL > Network Access.**
2. Select the **GigaSMART appliance.**
3. Click **Edit**.
4. Select **IP Address** from the **Network Access configuration** options.

   NOTE: For the Thales Luna HSM group, it is preferable to use a static IP address to prevent the Thales Luna registration from expiring.

> **NOTE:** If the IP address of the GigaSMART engine is changed, the GigaSMART engine needs a reboot to complete the HSM registration with the new IP address.

5. Enter **IP Address, Netmask, Gateway, DNS, MTU and VLAN** parameters.
6. Select the required management interface.
7. Under the **Ping Test** section, select the GigaSMART port and enter the **IP Address / Host Name** and the **Ping Test** parameters.

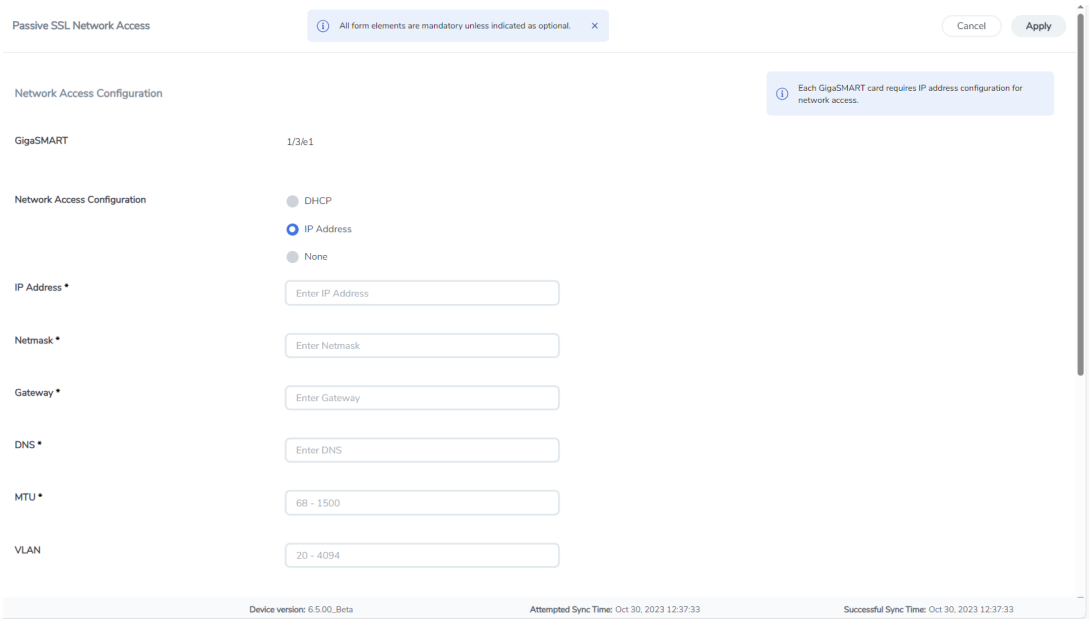


Figure 118*Passive TLS/SSL Network Access - IP Configuration*

8. Click **Apply**.

Use RFS to Manage Encrypted Keys

Use Remote File System (RFS), a component in HSM to store and manage encrypted keys. The RFS helps to automate the key distribution process. You can enable RFS on the GigaVUE-OS device using GigaVUE-FM so that the device can access the encrypted keys stored in RFS. You can synchronize RFS with GigaVUE-OS device to perform a bulk download of the encrypted keys.

> **NOTE:** Remote File System is not applicable for Thales-Luna HSM.

Refer to the following sections:

- Add RFS to GigaVUE-OS Device
- Map Encrypted Keys with Servers

Add RFS to GigaVUE-OS Device

To add and synchronize RFS to the GigaVUE-OS device:

1. From the device view, go to **GigaSMART > Passive SSL > RFS-Sync** .
2. Select your HSM Group.
3. From the RFS section, select the RFS got to **Actions>Sync** to automatically synchronize the RFS with the GigaVUE-OS device.
4. In the **Sync Period** field, enter the time interval for synchronization in hours.
5. Click **OK**
6.

The details of the RFS, such as the IP address, sync period, last sync time, next sync time, and the total number of keys stored and managed in the RFS appears in the RFS section. Click the **Show Details** link next to the **Total Keys** column to view the key token and key name mapping for the encrypted keys stored in RFS.

You can choose to modify the fields for the RFS you have added. Select the RFS click on **Actions > Edit** .

Click **Sync** to manually synchronize the RFS and the GigaVUE-OS device to fetch the latest encrypted keys from RFS at any point in time.

Map Encrypted Keys with Servers

A key name or a key token must be mapped to a server IP address. A total of up to 1000 key mappings is allowed per RFS. You can either manually map the keys with the servers or do a bulk key mapping using a text file. If a key mapping already exists, the new key mapping will be rejected. You must delete the existing key mapping to add the new mapping.

To map a key name or key token to a server IP address:

1. From the device view, go to **GigaSMART > Passive SSL > Key Mapping**.
2. In the Key Mapping section, click **Add**. The Add Key Mapping page appears.
3. Choose one of the following types to map the keys with the server IP address:
    - **Manual**—Manually map the key name or key token to the server IP address. You must keep adding the key mapping one at a time.
    - **From URL**—Create a text file with the key mappings and upload it to a server. Enter a valid directory path including the text file name and enter the password to access the server. It is recommended to use a secure protocol, such as SCP or HTTPS to access the URL.

- **From Local Directory**—Create a text file with the key mappings and save the text file in your local directory. Browse and select the text file from your local directory.

4. Click **Submit**.

Configure a GigaSMART Group

Refer to the section GigaSMART Group for more details on GigaSMART Group.

To configure a GigaSMART group for passive TLS/SSL:

1. From the device view, go to **GigaSMART > GigaSMART Groups.**
2. Click **New**.
3. In the **Alias** field, enter a name for the GigaSMART group that you are creating for Passive TLS/SSL.
4. From the **Port List** drop-down list, select the required port you want to associate with this group.
5. Scroll down to the **GigaSMART Parameters > Passive SSL> HSM Group** section of the page, and then select the required HSM Group from the drop-down list.

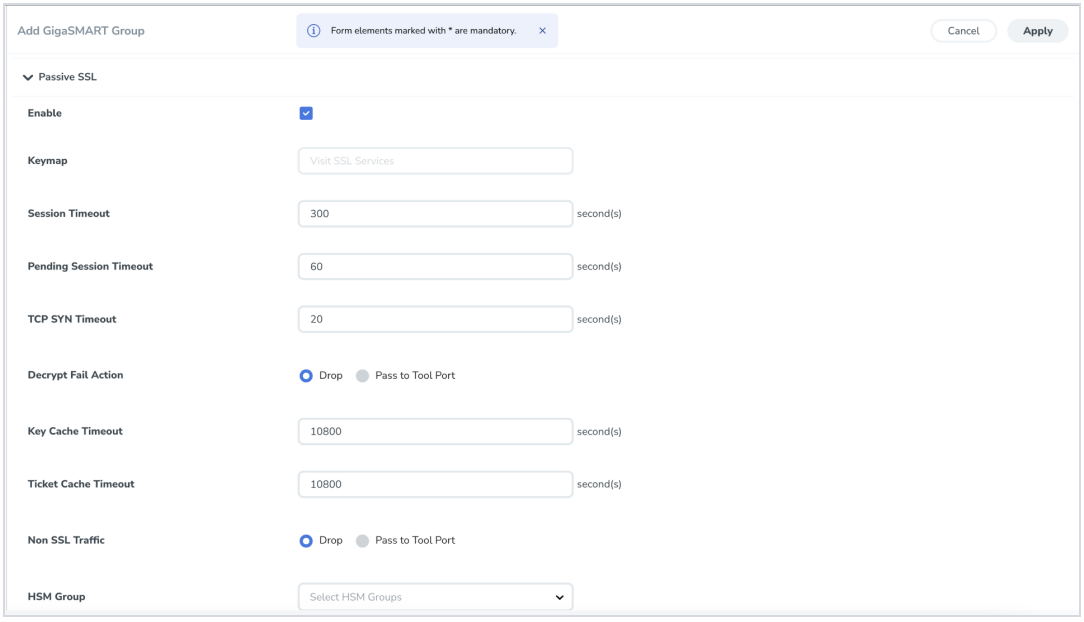


Figure 119*GigaSMART Group Setup Page*

6. Click **OK**.

Create a GigaSMART Operation (GSOP)

Refer to the section GSOP for more details on GigaSMART Operation.

To create a GigaSMART operation with an TLS/SSL Decryption component:

1. From the device view, go to **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.
2. Click **New**.
3. In the **Alias** field, enter a name for the GigaSMART operation.
4. From the **GigaSMART Group** drop-down list, select the GigaSMART group that you have created for passive TLS/SSL.
5. From the **GigaSMART Operations (GSOP)** drop-down list, select **SSL Decryption**.

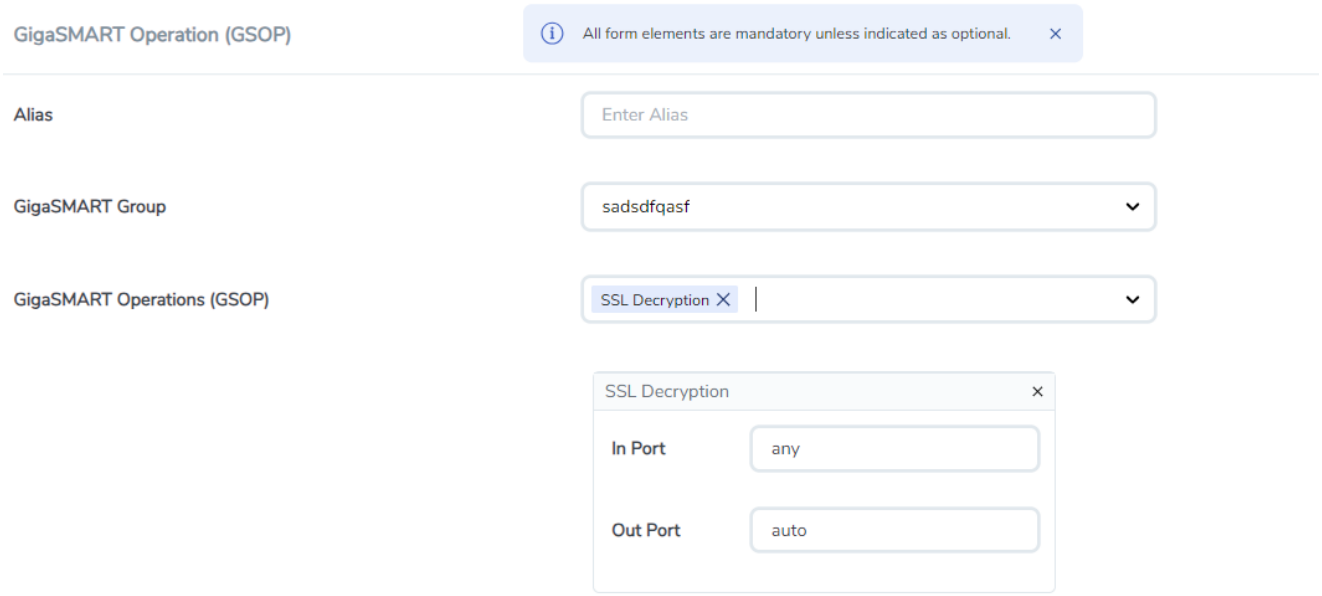


Figure 120*GigaSMART Operations - Setup Page*

6. Click **Apply**.

Create TLS/SSL Keychain Password

Before uploading keys or configuring TLS/SSL, you must create an TLS/SSL keychain password. The password is used to encrypt the private keys that you upload to the node.

NOTE: When uploading TLS/SSL keys, make sure that you are not creating a duplicate key (adding same key with different key alias). Adding a duplicate key can cause errors.

To create an TLS/SSL keychain password:

1. From the device view, go to **GigaSMART > Passive SSL > Key Store.**
2. Click **Keychain Password** from the **Actions** drop down list.

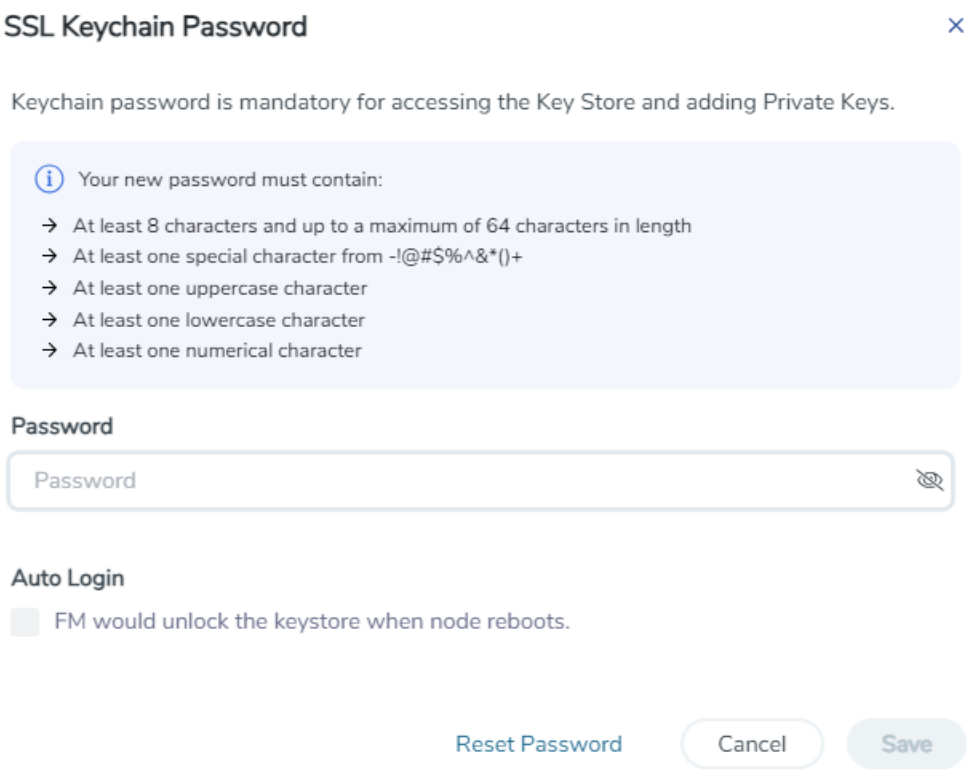


Figure 121*SSL Keychain Password Setup Page*

3. In the **Password** field, enter a valid password. Ensure that the password meets the following specifications:
   - Password must be at least eight (8) characters in length.
   - Password must contain at least one:
     - uppercase character
     - lowercase character
     - numerical character
     - special character
4. Click on the **Auto Login** check box to allow GigaVUE-FM,GigaVUE HC Series to unlock the keystore when the node reboots.

5. Click **Save**.

Upload TLS/SSL Private Keys

To upload an TLS/SSL private key:

1. From the device view, go to **GigaSMART >Passive SSL > Key Store** to open the Key Store page.
2. Click **Add**. The SSL Key page appears.
3. In the **Alias** field, enter a name for the SSL key.
4. Select the **Key Upload Type** from the options **PEM, PKCS12, Entrust nShield**, **Luna-HSM**.
   a. Passphrase- SSH passphrases allows you to protect your private key from being used without the passphrase. Enter the passphrase created with the private key.
   b. Private key- Enter the Private Key using any of the following options:
      - Copy and Paste
      - Install from URL
      - Install from Local Directory
   c. Certificate- Enter the Certificate using any of the following options:
      - Copy and Paste
      - Install from URL
      - Install from Local Directory
   d. Key Label - The Key label that is imported from the HSM server. This field is applicable when you configure Thales-Luna HSM.

   NOTE: Ensure that Luna key labels match the correct certificate and decryption keys when using the default service in Passive SSL.

   e. Path and Password- Configure the file path and password if you select 'Install form URL', or else choose the file if you select 'Install form Local directory'. The supported protocols are **HTTP, HTTPS, FTP, TFTP, SCP, and SFTP**.
5. Click **Save**.

Configure TLS/SSL Service

After you have uploaded a private key, you can add a service. A service maps to a physical server, such as an HTTP server. One server can run multiple services. A service is a combination of an IP address and a server port number. Also, the key and the service must be tied together.

To create an TLS/SSL service:

1. From the device view, go to **GigaSMART > Passive SSL > SSL Services.**
2. Click **New**. The SSL Service page appears.

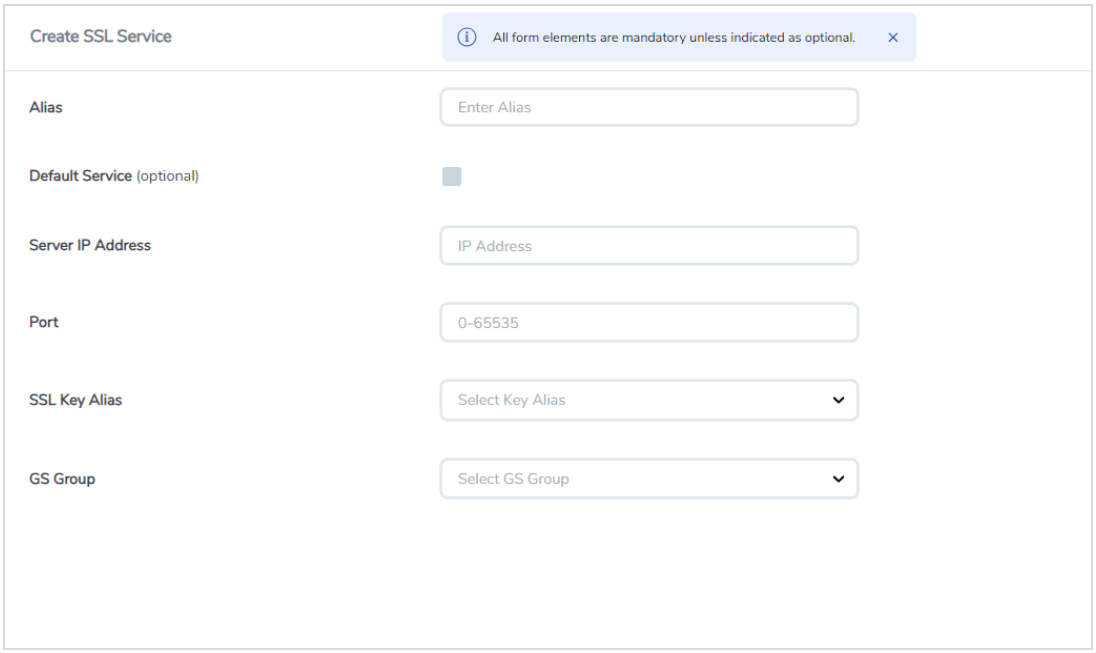


Figure 122*SSL Service*

3. In the **Alias** field, enter a name for the SSL service.
4. Map the TLS/SSL service to a server IP address and a server port using one of the following methods:
   - Select the **Enabled** check box next to the **Default Service** field to dynamically map the server IP address and server port.

   NOTE: If you select the **Enabled** check box, the **Server IP Address** and **Server Port** fields are disabled.

   - In the **Server IP Address** and **Server Port** fields, enter an IP address and port to which you want to map the TLS/SSL service.
5. From the **SSL Key Alias** drop-down list, select the name of the SSL Key previously uploaded.

6. From the **GS Group** drop-down list, select the GigaSMART group with TLS/SSL decryption enabled to associate with this TLS/SSL service.
7. Click **Apply**.

Configure Maps

1. From the device view, go to **Traffic> Maps > Maps.**
2. Click **New**.
3. Configure the map.

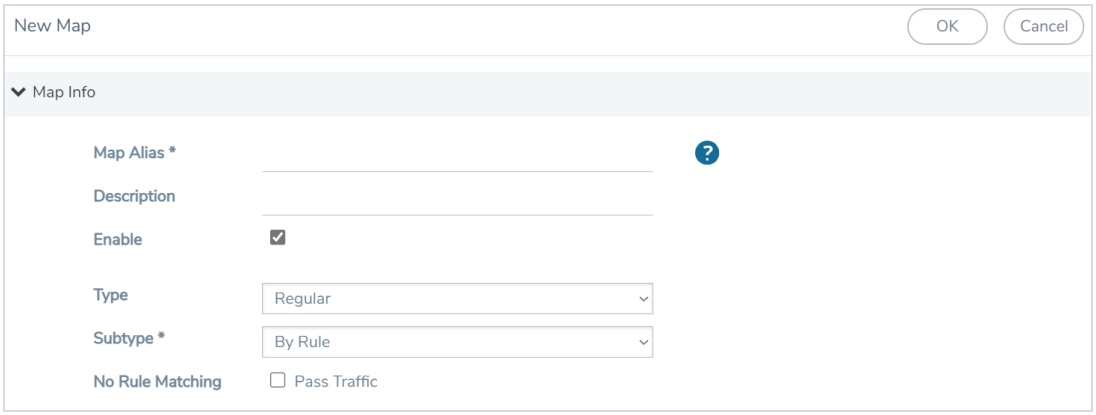


Figure 123*Create New Map*

- Type **map11** in the Alias field.
- Select Regular for **Type**.
- Select **ByRule** for **Subtype**.
- Select the network port for the Source.
- Select **Tool port/Hybrid port** for Destination.

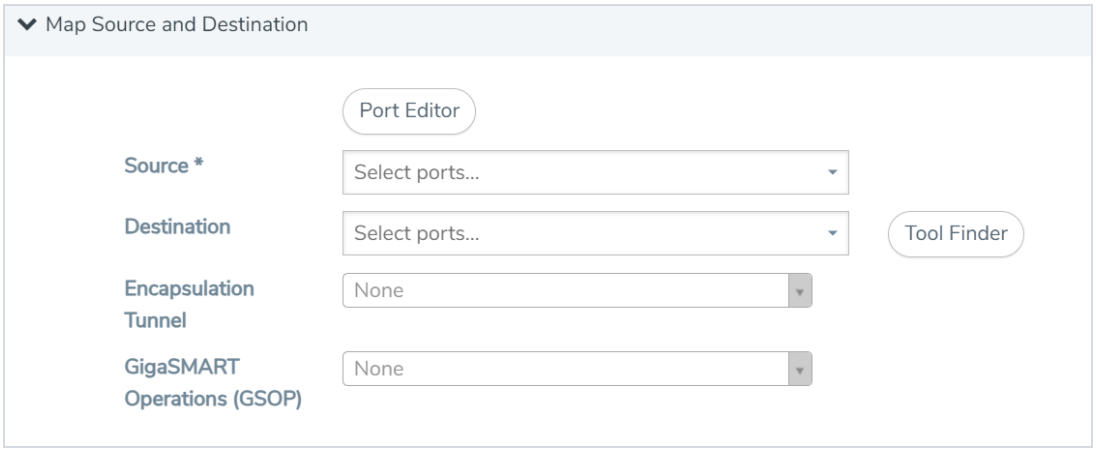


Figure 124*Configure Map Details*

4. Add **Configuration & Rules**.

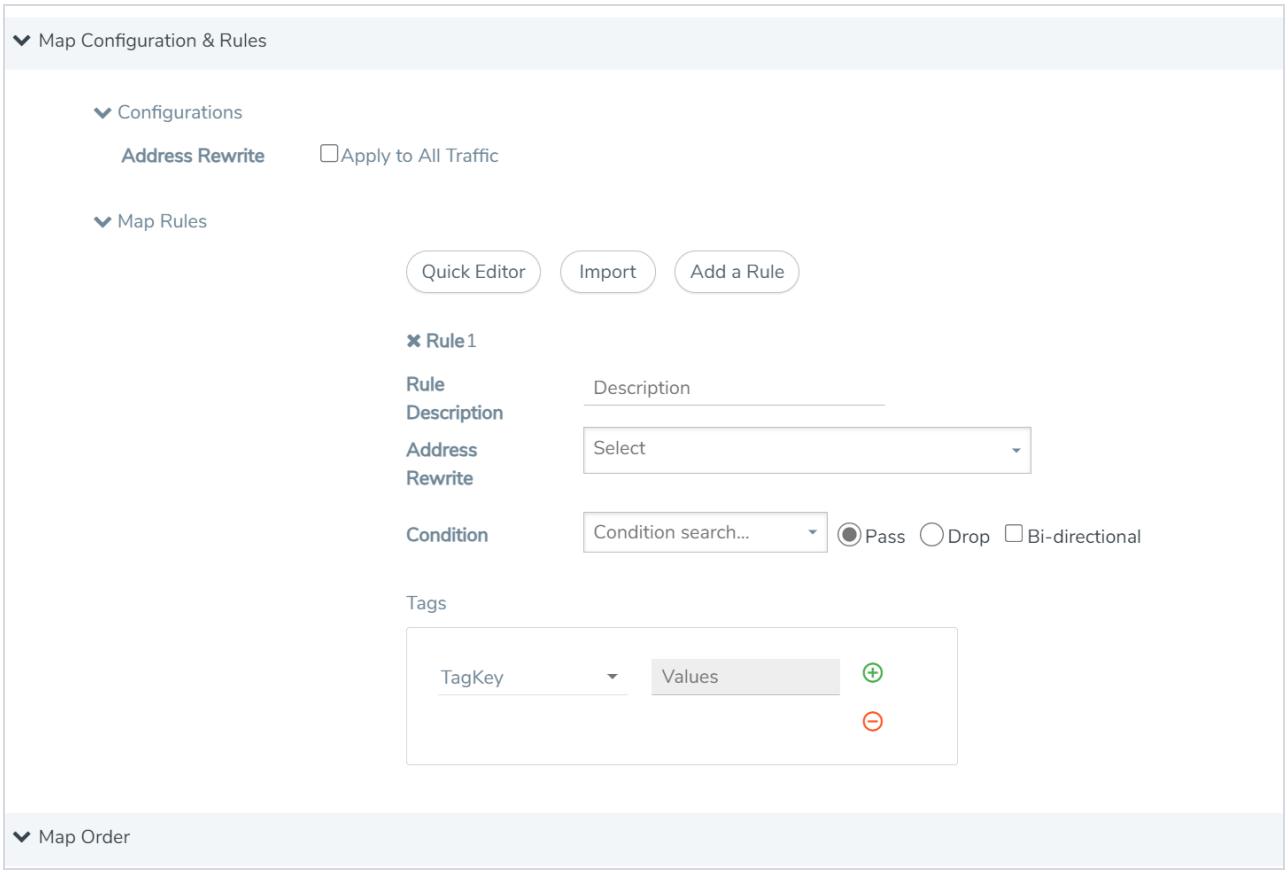


Figure 125 *Figure 20-123: Map Details - Create Rule*

a. Under Configurations, select the **Apply to All Traffic** check box to rewrite the address.

b. Select the address from the drop-down list.

c. Under Map Rules, Click **Add a Rule**.

d. Select **Pass** as condition.

e. Select **IPv4 Version** and set **Version to v4.**

5. Click **OK**.

To view the configured Thales-Luna HSM Group status, select the **Thales Luna profile> Action> Diagnostic.** You can view the statistics of the below components:

1. Ping Result
2. Verify
3. High Availability
4. Luna Key Label

PCAPng Application

The PCAPng application is a GigaSMART parser application that reads the various blocks in the received PCAPng files and validates the blocks to be sent to the destination application or to the tools. The PCAPng file contains the following blocks:

- Mandatory Blocks
  - Section Header Block (SHB)
- Optional Blocks
  - Interface Description Block (IDB)
  - Enhanced Packet Block (EPB)
  - Simple Packet Block
  - Name Resolution Block
  - Interface Statistics Block

The actual packets are present in the Enhanced Packet Block. The block data is parsed to find the start and end offset of the valid packets and the packet is sent out to the next application.

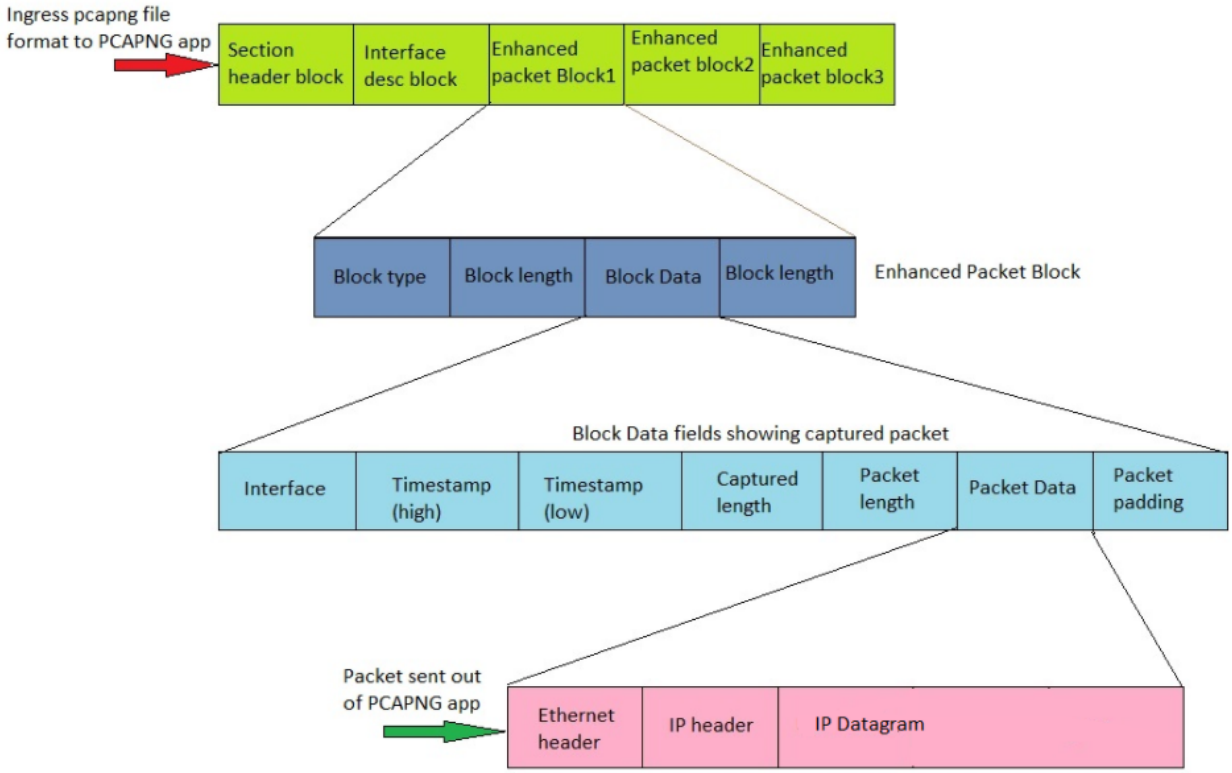
NOTE: Only one EPB in a PCAPng file is supported.

The PCAPng application processes the data depending on the packet type that contains a combination of the blocks mentioned above:

Block Combination	Process
SHB+IDB+EPB+data	Packets are parsed, validated, and the data packet is sent out.
SHB+IDB	Packets are dropped.
EPB+Data	Packets are parsed, validated, and the data packet is sent out.

The PCAPng application validates if the incoming data matches any of the above three formats in the same order, and processes the packets accordingly.

The following figure shows a sample PCAPng file format that contains one section header block:



Tunneling Operations

GigaSMART tunneling operations:

GigaSMART IP Encapsulation/Decapsulation (GigaSMART Tunnel)

Required License for IP Decapsulation: Base (GigaVUE-HC3), Tunneling (GigaVUE-HC1)Required License for IP Encapsulation: Advanced Tunneling (GigaVUE-HC3), Tunneling (GigaVUE-HC1)

Use GigaSMART encapsulation and decapsulation operations to send traffic arriving on one GigaSMART-enabled node over the Internet to a second GigaSMART-enabled node. There, the traffic is decapsulated and made available to local tool ports.

This feature is useful when instrumenting remote data centers – you can tunnel selected portions of the traffic from the remote GigaSMART-enabled node to tools in a central location. Traffic is encapsulated at the sending end of the tunnel and decapsulated at the receiving end.

GigaSMART drops packets when it receives packets with multiple encapsulation headers that are added due to unpredictable looping in the network.

IP fragmentation and reassembly are supported. Refer to IP Fragmentation and Reassembly on L2GRE and GMIP Tunnels.

The source of the GigaSMART tunnel can be any of the following:

- **GigaSMART-Enabled GigaVUE-FM,GigaVUE HC Series H Series Node**
  - Standalone GigaVUE-HC3 node with SMT-HC3-C05 modules installed.
  - Standalone GigaVUE-HC1 nodes.
  - Any GigaVUE-FM,GigaVUE HC Series HC Series node operating in a cluster with the previous node types.
- **GigaVUE V Series node or a GigaVUE-VM**

NOTE: You can also create GigaSMART operations that allow a GigaVUE-FM,GigaVUE HC Series H Series node to act as the receiving end of an ERSPAN tunnel for data mirrored over the Internet from Cisco equipment. However, this feature requires the Advanced Tunneling license; refer to GigaSMART ERSPAN Tunnel Decapsulation.

Display GMIP Tunnel Decapsulation Statistics

To display tunnel decapsulation statistics, select **GigaSMART > GigaSMART Operations > Statistics** and click a on the GS Operation in the table to open the Quick View for GS Operations Statistics.

Refer to Tunnel Decapsulation Statistics Definitions and GigaSMART Operations Statistics Definitions for descriptions of these statistics.

Display GMIP Tunnel Encapsulation Statistics

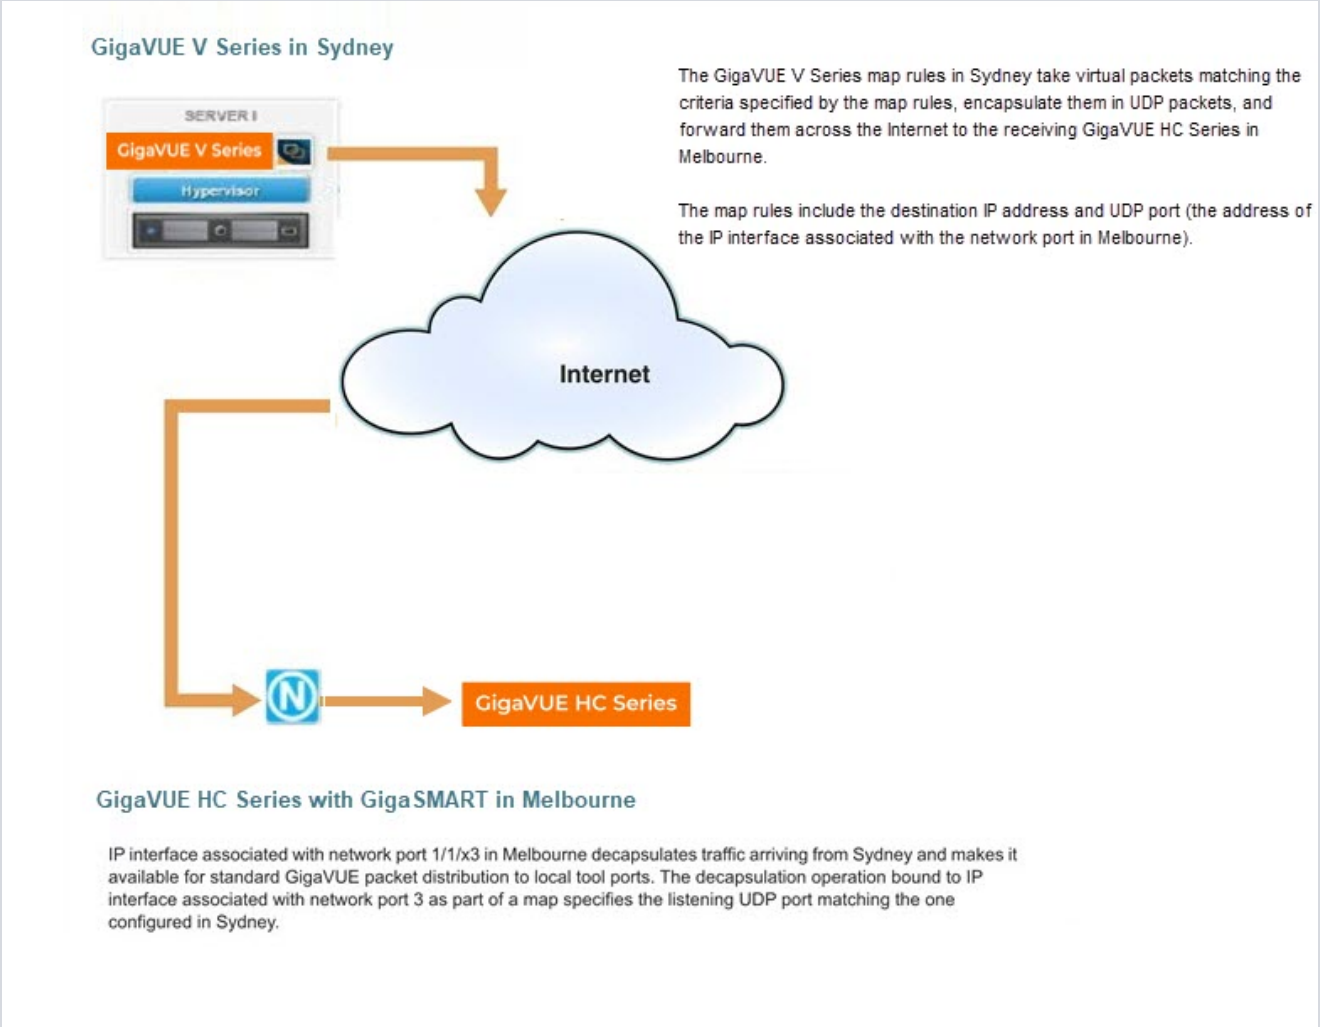
To display tunnel encapsulation statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics** and click on the GS Operation in the table to open the Quick View for GS Operations Statistics.

Refer to Tunnel Encapsulation Statistics Definitions and GigaSMART Operations Statistics Definitions for descriptions of these statistics.

Example: GigaSMART Encapsulation/Decapsulation (GigaVUE V Series)

The following figures demonstrate how to create a sample IP tunnel between a sending GigaVUE V Series node in Sydney and a receiving GigaVUE-FM,GigaVUE HC Series HC Series in Melbourne. First, the overall tunnel is summarized, followed by configuration

descriptions for the sending and receiving ends.



Configure Sending End of Tunnel: GigaVUE V Series map rules in Sydney

A GigaVUE V Series node in this location is configured with map rules that will send data over the Internet to the IP interface associated with a network port on a GigaVUE-FM,GigaVUE HC Series HC Series.

Map rules are created in the GigaVUE-FM user interface – Step 2 in the Create Map wizard includes **Tunnel Traffic To** settings that specify where matching traffic should be sent:

Create "Tunnel Traffic To" Option	Setting
UDP IP	This is the destination IP address for the IP interface associated with network port on the GigaVUE-FM,GigaVUE HC Series H Series in Melbourne. We will set it to 10.150.68.222
UDP Source Port	This is the UDP source port from which tunneled packets will be sent. We will set this to 5000.
UDP Destination Port	This is the listening port on the receiving GigaVUE-FM,GigaVUE HC Series H Series IP interface associated with network port. We will set this to 10000.

Configure Receiving End of Tunnel: GigaVUE-FM,GigaVUE HC Series HC Series with GigaSMART in Melbourne

Now we need to configure the receiving end of the tunnel with an IP interface associated with network port. The GigaVUE-FM,GigaVUE HC Series H Series in this location will have an IP interface associated with network port configured on network port 1/1/3 with an IP address of 10.150.68.222 and a GigaSMART decapsulation operation that listens on UDP port 10000.

The following table summarizes the steps necessary to configure the receiving end of the tunnel using the UI:

Task	UI Steps
Start by designating port 1/1/x3 as an IP interface with network port, configuring its IP profile, and assigning its GigaSMART operations to a GigaSMART group. This command sets the IP address, subnet mask, default gateway, and MTU for the IP interface associated with a tool port on port 1/1/x3.	**1.** Select Ports > IP Interfaces. **2.** Click New. **3.** Configure the IP Interface: ▪ Alias: 1_1_x3 ▪ Port: 1/1/x3 ▪ IP Address: **10.150.68.222** ▪ IP Mask: 255.255.255.255 ▪ Gateway: **10.150.68.1** ▪ MTU: 9400 ▪ GigaSMART Group: GS2 **4.** Save
Now, create an IP decapsulation GigaSMART operation (**gmipdecap**) that will decapsulate traffic received on UDP port 10000. Recall that we configured the sending end of the tunnel to send	**1.** From the device view, select GigaSMART > GigaSMART Operations > GigaSMART Operation. **2.** Click New. **3.** Configure the GigaSMART Operation:

Task	UI Steps
to that UDP port. The operation has the alias **gv_ipdecap**. Note that this operation uses the same GigaSMART group (GS2) as the IP interface associated with network port we set up in the first step.	▪ Alias: **gv_ipdecap** ▪ GigaSMART Groups: GS2 ▪ GigaSMART Operations (GSOP): Tunnel Decapsulation **4.** Configure the Tunnel Encapsulation: ▪ GMIP ▪ GMIP Port: 10000 **5.** Save.
Once we have our IP decapsulation operation, we can include it as part of a map. ▪ Open the map configuration page to create a map named **decapper**. ▪ The **Source** field specifies the ingress ports for this map. ▪ The **GSOP** field applies the gv_ipdecap GigaSMART operation to all packets matching the rules in the map, decapsulating them from the tunnel. ▪ The **Destination** field specifies where matching packets will be sent (tool port 1/1/x11). ▪ The rule with **Pass** selected specifies that packets arriving on this port with an IP Source address of 10.10.10.10 /32 will be processed by the gv_ipdecap GSOP and sent to tool port 1/1/x11.	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. ▪ Alias: decapper ▪ Type: Regular ▪ Subtype: By Rule ▪ Source: 1/1/x3 ▪ Destination: 1/1/x11 ▪ GigaSMART Operation (GSOP): gv_ipdecap (GS2) **4.** Click Add Rule. ▪ Select Pass. ▪ Select IPv4 Source for Rule 1. ▪ Set the IPv4 Address to 10.10.10.10 ▪ Set the Net Mask to 255.255.255.255 **5.** Click Save.

Configure IP Interfaces for GigaSMART Tunnels

When you configure IP interfaces for GigaSMART tunnels:

- You can associate the IP interfaces with multiple GigaSMART groups that are created either in the same node or in another node that resides in the same cluster.
- You can associate multiple GigaSMART engines to a GigaSMART group.
- You can also associate NetFlow exporters to the IP interface.

About IP Interface Centralization

A tunnel that originates from a node in a cluster can terminate on a remote port in another cluster. Also, a tunnel can have multiple termination points. You can associate the IP interface with multiple GigaSMART groups that are created either in the same node or in another node that resides in the same cluster. Moreover, you can associate multiple GigaSMART engines to a GigaSMART group.

The following figure illustrates the tunnel centralization feature.

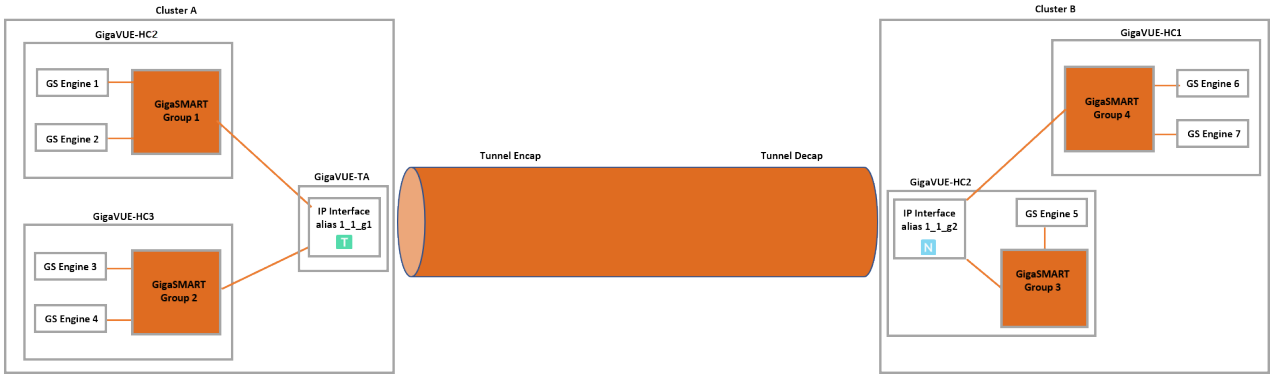


Figure 126 *Tunnel Centralization*

In this example, a GigaVUE-HC3, and GigaVUE-FM,GigaVUE HC Series-TA nodes reside in cluster A. In GigaSMART engines 1 and 2 are associated with the GigaSMART group 1. Similarly, in GigaVUE-HC3, GigaSMART engines 3 and 4 are associated with GigaSMART group 2. An IP interface with alias 1_1_g1 is configured with a tool port in GigaVUE-FM,GigaVUE HC Series-TA. Both the GigaSMART groups, 1 and 2 are associated with the IP interface. The IP interface 1_1_g1 is the originating point for the tunnel, where encapsulation happens.

Now, let us look at the termination point of the tunnel. The GigaVUE-HC1 a nodes reside in cluster B. In GigaVUE-HC1-Plus, the GigaSMART engine 5 is associated with GigaSMART group 3 and in GigaVUE-HC1, the GigaSMART engines 6 and 7 are associated with GigaSMART group 3. An IP interface with alias 1_1_g2 is configured with a network port in GigaVUE-HC1-Plus. Both the GigaSMART groups, 3 and 4 are associated with the IP interface. The IP interface 1_1_g2 is the termination point for the tunnel, where decapsulation happens. Thus the tunnel terminates on a remote port in another cluster.

> **NOTE:** Do not attempt to configure with IP /Tunnel Interface and GigaVUE-FM Management Interface in the same subnet. As per routing functionality this is not a valid configuration.

For configuring IP Interface for GigaSMART tunnels, refer to Configure IP Interface.

Configure GigaSMART IP Encapsulation/Decapsulation

This section contains the following topics:

- Configure Both Ends of GigaSMART Tunnel
- Configure Sending End of Tunnel: GigaVUE-HC1 in Reno
- Configure Receiving End of Tunnel: GigaVUE-HC3 with GigaSMART in San Francisco

Configure Both Ends of GigaSMART Tunnel

Creating a GigaSMART tunnel requires configuration on both the sending and receiving ends:

Sending End of Tunnel	Receiving End of Tunnel
The sending end of a GigaSMART tunnel can be either a GigaVUE-VM deployment or a GigaSMART-enabled GigaVUE-FM,GigaVUE HC Series H Series node. **Sending Data from a GigaSMART-Enabled GigaVUE-FM,GigaVUE HC Series H Series Node** o Configure an IP interface with an IP address, subnet mask, default gateway, MTU setting and assign it to a GigaSMART group. o Create a GigaSMART operation with a tunnel-encap component. The encapsulation settings include the IP address and listening UDP port of the P interface that is associated with a network port on the destination GigaVUE-FM,GigaVUE HC Series H Series. o Bind the GigaSMART operation to one or more network ports as part of a map. The network ports must be mapped to the IP interface associated with a tool port. **Sending Data from GigaVUE-VM/GigaVUE-FM** When you provision a vMap for a GigaVUE-VM node in GigaVUE-FM, in addition to selecting the virtual traffic to be forwarded, you also specify the destination to which traffic should be tunneled with the following settings: o **UDP IP** – The IP address of the P interface that is associated with a network port on the receiving end of the tunnel. o **UDP Source Port** – The source port from which traffic will be sent to the receiving end of the GigaSMART tunnel. o **UDP Destination Port** – The listening UDP port at the destination end of the GigaSMART tunnel. **Sending Data from GigaVUE-FM,GigaVUE HC Series-2404/GigaSMART-6X** o Configure an IP interface with an IP address, subnet mask, default gateway, and MTU setting. Associate the IP interface with a tool port. o Create a GigaSMART operation with an encapsulation component. The encapsulation settings include the IP address and listening UDP port of the IP interface thst is associated with network	o Configure an IP interface with an IP address, subnet mask, and default gateway. The IP address must match the destination IP address specified at the sending end of the tunnel. o Create a GigaSMART operation with a decapsulation component. The decapsulation settings include the same listening UDP port you specified as the destination port at the sending end of the tunnel. o Bind the GigaSMART operation to the IP interface that is associated with a network port as part of a map that distributes arriving traffic to local tool ports for analysis with local tools.

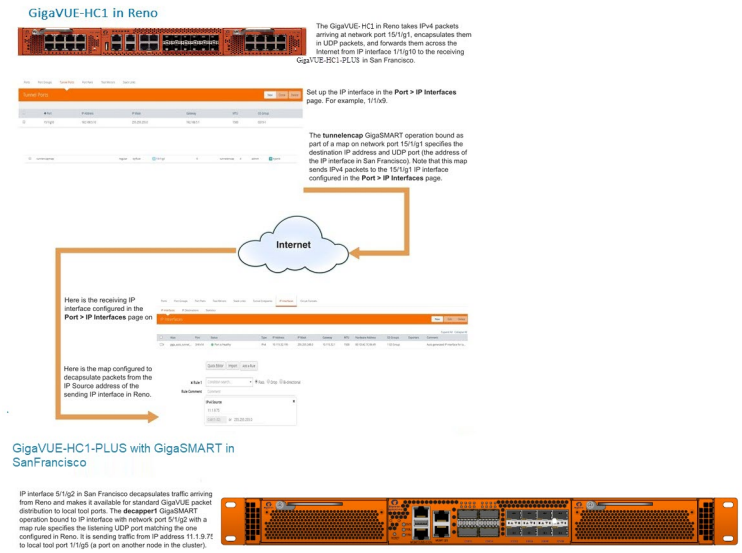
Sending End of Tunnel	Receiving End of Tunnel
port on the destination device. o Bind the GigaSMART operation to one or more network ports as part of a map rule with at least one regular map rule criterion. The network ports must be mapped to the IP interface associated with a tool port.	

Keep in mind the following when configuring GigaSMART operations with encapsulation/decapsulation components:

Feature	Description
Viewing Statistics	Use the **show tunneled-port** commands to see statistics related to ongoing tunnel operations. Refer to View GigaSMART Statistics for more information.
Packet Order	Packer sequence is not preserved if the packets are reordered while traversing the Internet. The receiving GigaSMART delivers them in the same order received.
GMIP Header	The GMIP header is 46 bytes consisting of 14 Ethernet + 20 IP + 8 UDP + 4 tunnel version.
Tunnel Decap Type GMIP portdst	Use the GigaSMART Operations page to specify the UDP port on which the P interface that is associated with a network port on the receiving GigaVUE HC Series is listening. Use this option when decapsulating traffic from a either GigaSMART-enabled node or a GigaVUE-VM deployment. The setting must match the configuration of the portdst configured on the sending end of the tunnel.
GigaSMART Engine Ports	GigaSMART operations with a tunnel component can be assigned to GigaSMART groups consisting of multiple GigaSMART engine ports. Refer to Groups of GigaSMART Engine Ports for more information.

Example: GigaSMART Encapsulation/Decapsulation (GigaVUE-HC1 Node)

The following figures demonstrate how to create a sample IP tunnel between a sending GigaVUE-HC1 in Reno and a receiving GigaVUE-FM,GigaVUE HC Series H Series cluster in San Francisco. First, the overall tunnel is summarized, followed by configuration descriptions for the sending and receiving ends.



Configure Sending End of Tunnel: GigaVUE-HC1 in Reno

The GigaVUE-HC1 in this location has an IP interface configured on tool port 1/1/g1 with an IP address of 11.1.9.75. Maps to this port that use a tunnel encapsulation GigaSMART operation can send data over the Internet. The following table summarizes the commands necessary to configure the sending end of the tunnel in the GigaVUE-FM:

Task	UI Steps
Start by designating port 1/1/g1 as a tool port.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** In the Quick View Editor find port 1/1/g1. **4.** Set Type to Tool. **5.** Select Enable **6.** Click OK. **7.** Close the Quick Port Editor.
Use the **IP Interfaces** page to set up the network parameters for 1/1/g1. This page sets the IP address, subnet mask, default gateway, and MTU for the IP interface associated with a tool port on port 1/1/g1. Notice that the GigaSMART group in this example has the alias **gsport1**.	**1.** Select Ports > IP Interfaces. **2.** Click New. **3.** Configure the IP interface: ▪ Alias: 1_1_g1 ▪ Port: 1/1/g1 ▪ IP Address: **11.1.9.75** ▪ IP Mask: 255.255.255.0 ▪ Gateway: 11.1.9.1 ▪ MTU: 9400 ▪ GigaSMART Group: gsport1 **4.** Click OK.

Task	UI Steps
Now, create a tunnel encapsulation GigaSMART operation (**tunnelencap**) that will send traffic to IP address 21.2.9.75 on destination UDP port 10000 from source port 5000. The operation has the alias **tunnelenc**.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **2.** Click New. **3.** Configure the GigaSMART Operation: ▪ Alias: tunnelenc ▪ GigaSMART Groups: gsport1 ▪ GigaSMART Operations (GSOP): Tunnel Encapsulation **4.** Configure Tunnel Encapsulation: ▪ GMIP ▪ Destination: IPv4 ▪ Port Source: 5000 ▪ Port Destination: 10000 ▪ Destination IP: 21.2.9.75 ▪ DSCP: 0 ▪ Precision: 1 ▪ TTL: 64 **5.** Click Save.
Once you have the tunnel encapsulation operation, you can include it as part of a map rule. This map rule matches IPv4 packets and sends them to 21.2.9.75:10000 (the socket specified by the GigaSMART operation named **tunnelencap** that you created in the previous step).	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. ▪ Alias: tunnelencap ▪ Type: Regular ▪ Subtype: By Rule ▪ Source: 1/1/x3 ▪ Destination; 1/1/x1 ▪ GigaSMART Operations (GSOP): tunnelencap (gsport1) **4.** Click Add Rule. **5.** Select Pass. **6.** Select IP Version for Rule 1. **7.** Select v4 Version. **8.** Save.

Configure Receiving End of Tunnel: GigaVUE-HC3 with GigaSMART in San Francisco

Now we need to configure the receiving end of the tunnel with an IP interface associated with network port. The GigaVUE-HC3 in this location will have an IP interface associated with network port configured on network port 5/1/x2 with an IP address of 21.2.9.75 and a

GigaSMART decapsulation operation that listens on UDP port 10000.

The following table summarizes the steps necessary to configure the receiving end of the tunnel using the UI:

Task	UI Steps
Start by designating port 5/1/x2 as a network port.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** In the Quick View Editor find port 5/1/x2. **4.** Set Type to Network. **5.** Select Enable **6.** OK. **7.** Close the Quick Port Editor.
Use the **IP Interfaces** page to set up the network parameters for 5/1/x2. This command sets the IP address, subnet mask, default gateway, and MTU for the IP interface associated with network port on port 5/1/x2. Note that this port uses the same IP address to which the GSOP in Reno is configured to send data (21.2.9.75).	**1.** Select Ports > IP Interfaces. **2.** Click New. **3.** Configure the IP Interface: ▪ Alias: 1_1_x2 ▪ Port: 1/1/x2 ▪ IP Address: 21.2.9.75 ▪ IP Mask: 255.255.255.0 ▪ Gateway: 21.2.9.1 ▪ MTU: 9400 ▪ GigaSMART Group: gsport5 **4.** Save.

Task	UI Steps
Now, create a tunnel decapsulation GigaSMART operation (**tunnel-decap**) that will decapsulate traffic received on UDP port 10000. Recall that we configured the sending end of the tunnel to send to that UDP port. The operation has the alias **hd-decap1**.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP). **2.** Click New. **3.** Configure the GigaSMART Operation: ▪ Alias: hd-decap1 ▪ GigaSMART Groups: gsport5 ▪ GigaSMART Operations (GSOP): Tunnel Decapsulation **4.** Configure the Tunnel Decapsulation. ▪ GMIP ▪ GMIP Port: 10000 **5.** Save.
Once you have your tunnel decapsulation operation, you can include it as part of a map rule. This map decapsulates all traffic arriving at 5/1/x2 from IP address 21.2.9.25 (the start of the tunnel) and sends it to port 1/1/x5. This is a tool port on the chassis with box ID 1 in this cluster.	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. ▪ Alias: decapper ▪ Type: Regular ▪ Subtype: By Rule ▪ Source: 5/1/x2 ▪ Destination: 1/1/x5 ▪ GigaSMART Operations (GSOP): hd-decap1 (gsport5) **4.** Click Add Rule. **5.** Select Pass. **6.** Select IPv4 Source for Rule 1. **7.** Set the IPv4 Address to 11.1.9.75 **8.** Set the Net Mask to 255.255.255.0 **9.** Save.

GigaSMART IP Encapsulation (GigaSMART Tunnel)

Required License for IP Encapsulation: Advanced Tunneling (GigaVUE-HC3), Tunneling (GigaVUE-HC1)

GigaSMART-enabled nodes with the Advanced Tunneling license installed can encapsulate traffic and send it through a GigaSMART tunnel to a destination GigaSMART-enabled node.

1. Configure an IP interface with an IP address, subnet mask, default gateway, and MTU setting and assign it to a GigaSMART group.

2. Create a GigaSMART operation with a **Tunnel Encapsulation** component. The encapsulation settings include the IP address and listening UDP port of the IP interface associated with network port on the destination GigaVUE-FM,GigaVUE HC Series H Series.
3. Bind the GigaSMART operation to one or more network ports as part of a map. The network ports must be mapped to the IP interface associated with a tool port.

Refer to the sections beginning with Configure Both Ends of GigaSMART Tunnel for examples of the end-to-end configuration of a GigaSMART tunnel.

GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation

Required License for L2GRE Decapsulation: Base (GigaVUE-HC3), Tunneling (GigaVUE-HC1)Required License for L2GRE Encapsulation: Advanced Tunneling (GigaVUE-HC3), Tunneling (GigaVUE-HC1)

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

GigaVUE-HC1 Gen 3 module supports IPv6 for L2GRE tunneling.

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

Use GigaSMART Layer 2 (L2) Generic Routing Encapsulation (GRE) tunnel encapsulation to send traffic from one GigaSMART node over the Internet to a second GigaSMART node using L2GRE encapsulation. Use GigaSMART L2GRE tunnel decapsulation at the second GigaSMART node to decapsulate the traffic before sending it to local tool ports.

GigaSMART Layer 2 GRE tunnel encapsulation/decapsulation provides the following:

- L2GRE tunnel initiation and encapsulation on the tool port at the sending end of the tunnel (for example, at a remote site)
- L2GRE tunnel termination and decapsulation on the network port at the receiving end of the tunnel (for example, at a main office site)

The GigaSMART at the remote site encapsulates the filtered packets, adds an encapsulation header, and routes it to the main office site. The encapsulation protocol is GRE and the delivery protocol is IP or IPv6, so the encapsulation header consists of Ethernet + IP + GRE or Ethernet + IPv6 + GRE headers.

The parameters of the encapsulated header are user-configurable, such as the IPv4 address of the IP interface on the destination GigaSMART node and the GRE key that identifies the source of the tunnel.

At the remote end, packets are decapsulated, the L2GRE header is stripped off, and packets are sent to the specified tool port.

IP fragmentation and reassembly are supported. Refer to IP Fragmentation and Reassembly on L2GRE and GMIP Tunnels.

Figure 127L2GRE Tunnel Encapsulation/Decapsulation shows the remote site encapsulating the filtered traffic and routing it to the main office from the remote end.

The encapsulated packet is sent out of the tool port, which is connected to the public network (the Internet). This packet is routed in the public network to reach the main office site. It ingresses at the routed network port of the GigaVUE-FM,GigaVUE HC Series node at the main office.

The ingress encapsulated packet is then sent to the GigaSMART at the main office, where the packet is decapsulated and sent to the tool port. The received packet's destination IP is checked against the source IP/IPv6 configured for the network port. If they match, decapsulation is applied. The Ethernet + IP + GRE or Ethernet + IPv6 + GRE header is stripped and the remaining packet is sent to the tool port.

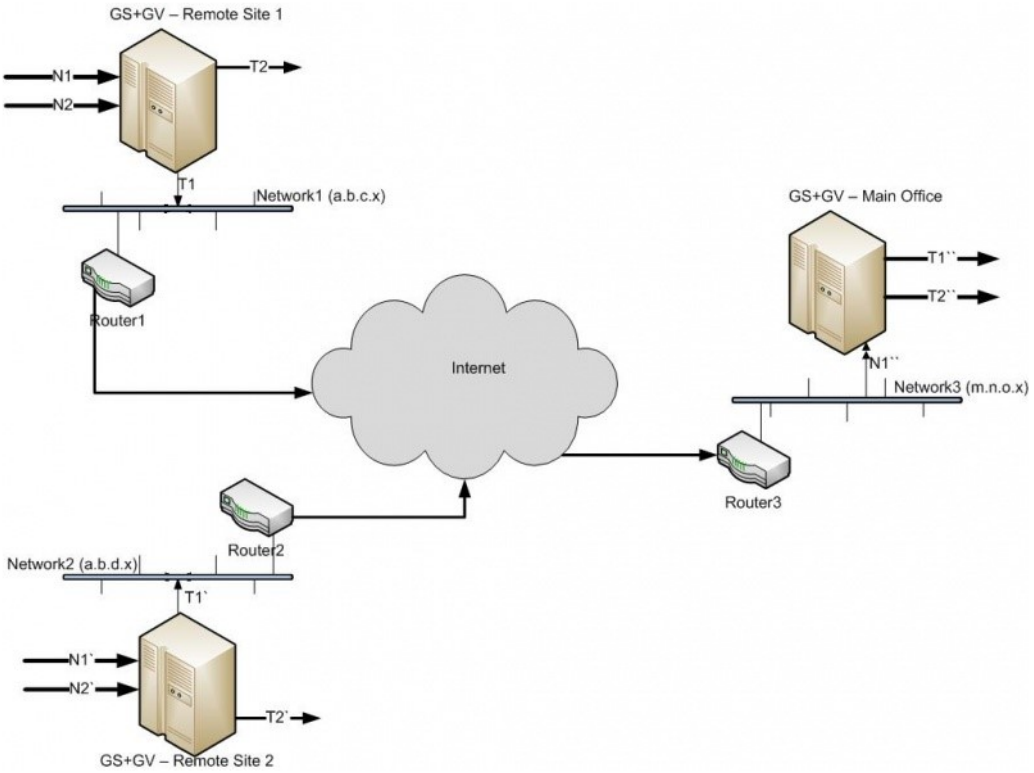


Figure 127 *L2GRE Tunnel Encapsulation/Decapsulation*

For L2GRE tunnel encapsulation/decapsulation configuration examples, refer to GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation and GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation.

For statistics for encapsulated packets, refer to Display L2GRE Tunnel Encapsulation Statistics. For statistics for decapsulated packets, refer to Display L2GRE Tunnel Decapsulation Statistics.

Layer 2 GRE Header Length

The L2GRE header length is as follows:

Header	Length in Bytes
With Key	42 bytes consisting of 14 Ethernet + 20 IP + 4 GRE + 4 GRE Key.
Without Key	38 bytes consisting of 14 Ethernet + 20 IP + 4 GRE.

Load Balancing to Multiple Destinations

Starting in software version 5.1, L2GRE tunnel encapsulation supports Load Balancing. Traffic from an IP Interface can be sent to multiple destinations Defined by IP address. The traffic is distributed using stateful Load Balancing or stateless hashing.

For information on stateful and stateless Load Balancing, refer to GigaSMART Load Balancing.

For examples of Load Balancing on L2GRE encapsulation, refer to GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation and GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation.

L2GRE IPv4 Encapsulation/Decapsulation

GigaSMART L2GRE IPv4 lets you route filtered traffic to the remote

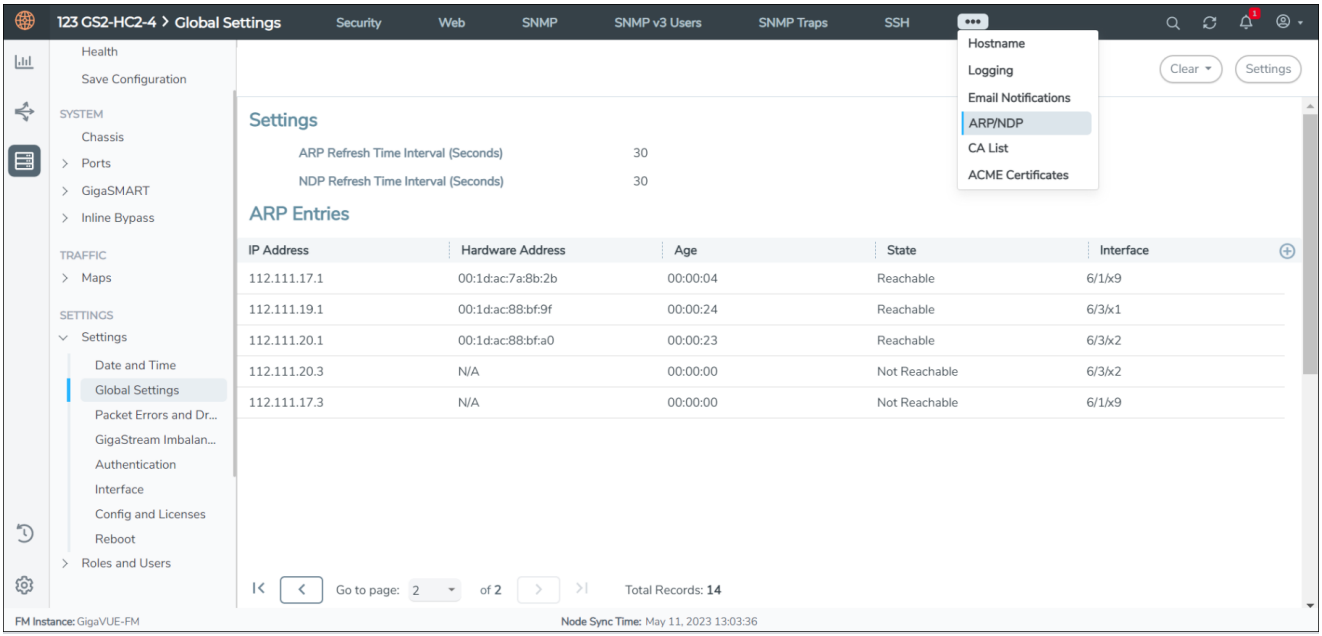
end using IPv4-based L2GRE tunneling. At the receiving end, filtered traffic is sent to GigaSMART, which adds an L2GRE header and a IPv4 header to make it routable.

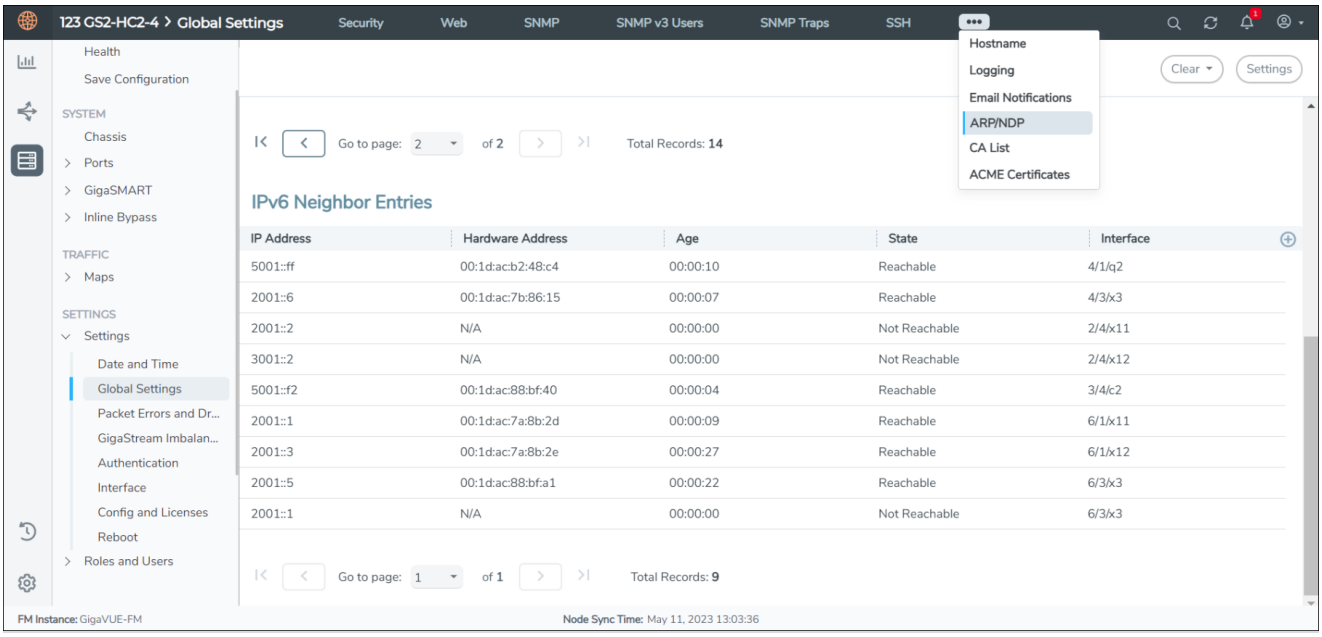
The remote end decapsulates the packet and sends it to the tool port.

GigaVUE-FM,GigaVUE HC Series nodes act as L2GRE encapsulation and decapsulation devices.

The IPv6 protocol is used to deliver all packets received in the encap tunnel to the termination node using the configured source and destination IPv4 address. The tunnel termination (decap) node strips the IPv4 + GRE header and sends the payload to the tool port.

The arp protocol is used by the tool port on the encapsulation node for arp request and arp response messages to resolve the gateway MAC address as well as destination tool MAC address in local network and respond to arp request messages received from the gateway in the tunnel decapsulation/termination node. arp echo request/reply messages are also sent and received.





For a configuration example, refer to GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation

L2GRE IPv6 Encapsulation/Decapsulation

GigaSMART L2GRE IPv6 lets you route filtered traffic to the remote

end using IPv6-based L2GRE tunneling. At the receiving end, filtered traffic is sent to GigaSMART, which adds an L2GRE header and a IPv6 header to make it routable.

The remote end decapsulates the packet and sends it to the tool port.

GigaVUE-FM,GigaVUE HC Series nodes act as L2GRE encapsulation and decapsulation devices.

The IPv6 protocol is used to deliver all packets received in the encap tunnel to the termination node using the configured source and destination IPv6 address. The tunnel termination (decapsulation) node strips the IPv6 + GRE header and sends the payload to the tool port.

The ICMPv6 protocol is used by the tool port on the encapsulation node for Neighbor Solicitation (NS) and Neighbor Advertisement (NA) messages to resolve the gateway MAC address as well as destination tool MAC address in local network and respond to NS messages received from the gateway in the tunnel decapsulation/termination node. ICMPv6 echo request/reply messages are also sent and received.

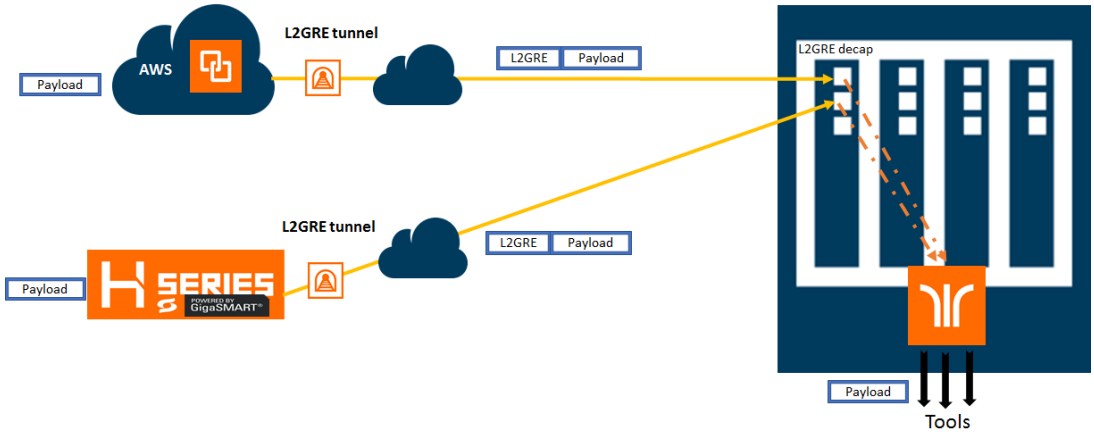
For a configuration example, refer to GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation

L2GRE Tunnel Termination

L2GRE tunnel termination is supported on physical devices, and the decapsulation happens through the GigaSMART engine. Tunneled traffic coming in the chassis is sent to the GigaSMART engine, which is sent to the tools using a hybrid port. The maps created are then applied to this decapsulated traffic.

Starting with version 5.4, tunnel termination is supported for VXLAN and L2GRE tunnel in the front panel ports of the switch. This feature provides line rate tunneling on all faceplate ports and also allows Flow Mapping® to be applied for the incoming tunneled traffic on the same ports.

The following diagram illustrates how the traffic from two sources—a GigaVUE V Series appliance running on an AWS platform and a GigaVUE HC Series device at a remote site traverses through the L2GRE tunnel and reaches the GigaVUE HC Series node in the main office site. In each case, traffic is tapped at the remote source and is then tunneled through L2GRE encapsulation across the cloud before it reaches the GigaVUE HC Series device at the main office site, which is connected to the actual tools. The L2GRE tunnel termination is executed on an ingress circuit port (IP interface) on the destination GigaVUE HC Series device. After tunnel termination, the packet is presented to the Flow Mapping® module to filter based on map rule parameters.



Configure L2GRE Tunnel Encapsulation and Decapsulation

Refer to the following configuration examples:

- GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation
- GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation
- GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation
- GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation
- GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation

Display L2GRE IPV4/IPv6 Tunnel Statistics

To view IP Interfaces statistics, **select Ports > IP Interfaces > Statistics** to open the IP Interfaces Statistics page.

The IP destinations pane displays the gateway or local tool status as **UP** if neighbor discovery is completed with gateway or local tool or **down** if neighbor discovery failed. Neighbor discovery is made only on the encapsulation node. On the decapsulation node, the gateway or local tool status will be **Not Applicable**.

Display L2GRE Tunnel Encapsulation Statistics

To display Layer 2 GRE tunnel encapsulation statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics**. The statistics for tunnel encapsulation will be in the row labeled Tunnel Encap in the GS Operations column.

Refer to Tunnel Encapsulation Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

Display L2GRE Tunnel Decapsulation Statistics

To display Layer 2 GRE tunnel decapsulation statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics** and click on the GS Operation in table to open the Quick View for GS Operation Statistics.

Refer to Tunnel Decapsulation Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

Configure GigaSMART Operation for Layer 2 GRE

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

To configure the L2GRE encapsulation/decapsulation types and options, use the GigaSMART Operations (GSOP) page:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.
2. Click **New**.
3. On the GigaSMART Operations page, do the following:
   a. Type an alias in the **Alias** field.
   b. In the **GigaSMART Groups** field, select the gsgroup for this operation.
   c. In the **GigaSMART Operations (GSOP)** field, select either **Tunnel Decapsulation** or **Tunnel Encapsulation** from the drop-down list, depending on whether you want decapsulation or encapsulation.
   d. Select **L2GRE**, and then enter options in the fields that display.
4. Click **Save**.

Example 1 – GigaSMART L2GRE Tunnel Encapsulation

In this example, an IP interface is configured on the tool port. A GigaSMART operation for tunnel encapsulation is configured to encapsulate the filtered packets. A map is configured that uses the L2GRE tunnel encapsulation GigaSMART operation, which sends packets from the remote site over the Internet to the main office using the IP interface associated with a tool port. Staring with software version 5.4 GigaSMART L2GRE Tunnel Encapsulation provides support for IPv6 with load-balancing.

Task	Description	UI Steps
1.	Configure a tool type of port and a network type of port.	**a.** Select Ports > All Ports. **b.** Click Quick Port Editor. **c.** Use Quick search to find the ports to configure. **d.** Set the type (Network or Tool) for each port and select Enable.
2.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**a.** From the device view, select GigaSMART > GigaSMART Groups. **b.** Click New. **c.** Type an alias in the Alias field and enter an engine port in the Port List field. **d.** Save.
3.	Configure the IP interface with an IP address, subnet mask, default gateway, and MTU setting. Assign it to the GigaSMART group.	**a.** Select Ports > Ports > IP Interfaces. **b.** Click New. **c.** On the IP Interfaces page, in the Alias and Description fields, enter a name and description for the IP interface. **d.** Click the Ports field and select the network or tool port from the drop-down list. **e.** Select Type: IPv4 or IPv6 **f.** Enter the IP address, subnet mask, gateway, and MTU settings in the respective fields. **g.** Click on the GigaSMART Group field to select the GigaSMART group. **h.** Save.
4.	Configure the GigaSMART operation for tunnel encapsulation and assign it to the GigaSMART group. The tunnel encapsulation settings include the IP address (IPv4) of the IP interface on the destination GigaSMART node and the GRE key that identifies the source of the tunnel.	**a.** From the device view, select GigaSMART > GigaSMART Operations > GigaSMART Operation. **b.** Click New. **c.** Type an alias in the Alias field. **d.** From the GigaSMART Groups drop-down list, select the GigaSMART Group that you created in the second task. **e.** From the GigaSMART Operations (GSOP) drop-down list select Tunnel Encapsulation. **f.** Select L2GRE for the encapsulation type. **g.** Enter the IP address of the IP interface in the Destination IP field. **h.** IPv4, iPv6 or Port Group **i.** Enter the key parameter in the Key field. **j.** Save.

Task	Description	UI Steps
5.	Create a map using the tunnel encapsulation GigaSMART operation, with packets coming from the network port and being sent to the Internet through the tool port.	**a.** Select Maps > Maps. **b.** Click New. **c.** Type an alias in the Map Alias field that will help you identify this map. **d.** Select Regular for Type and By Rule for Subtype. **e.** Specify the network and tool ports that you configured in task one in the Source and Destination fields, respectively. **f.** From the GigaSMART Operations (GSOP) drop-down list, select the GigaSMART operation configured in task 4. **g.** Click Add a Rule under Map Rules and create the following rule: Select IP Version from the drop-down list and select v4 for Version, and then select Pass. **h.** Click Save.

Example 2 – GigaSMART L2GRE Tunnel Encap Stateful LB

Example 2 configures stateful load balancing of tunnel traffic to tunnel endpoints based on a metric. Each tunnel endpoint is assigned a weight.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

To configure this example:

1. Go to **Ports > Ports > All Ports**. Make sure you have one tool type of port and one network type of port enabled. Also make sure you have a GigaSMART port (eport).
2. Go to **GigaSMART > GigaSMART Groups**.
3. Click **New**, and then configure an Alias for the GigaSMART group and associate it with a GigaSMART engine port.
4. Click **OK**.
5. Go to **Ports > Ports > IP Interfaces**.
6. Click **New**, and then in the Alias and Description fields, enter the alias and description of the IP interface.

7. Select a port and configure it with an IP version Type, IP Address, IP Mask, Gateway, and MTU. Assign the IP interface to the GigaSMART group.
8. Click **OK**.
9. Go to **Ports > Tunnel Endpoints**.
10. Click **New**, then configure one or more tunnel endpoint IDs and their IP Addresses. The Alias is optional.
11. Click **OK**.
12. Go to **Ports > Port Groups**.
13. Click **New**, then type an alias for the port group, select GigaSMART Load Balancing, select the previously configured tunnel endpoints. Optionally, you can specify weights for each tunnel endpoint in the port group.
14. Click **OK**.
15. Go to **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.
16. Click **New**, then select the same GigaSMART Group and Tunnel Encapsulation for the GSOP. Under Load Balancing, select Stateful, the type Tunnel, and the metric, such as Round Robin. Under Tunnel Encapsulation, select L2GRE, Destination Port Group, select the Port Group, Session Field, and Session Position. Refer to Figure 128New GigaSMART Operation for Stateful Load Balancing.

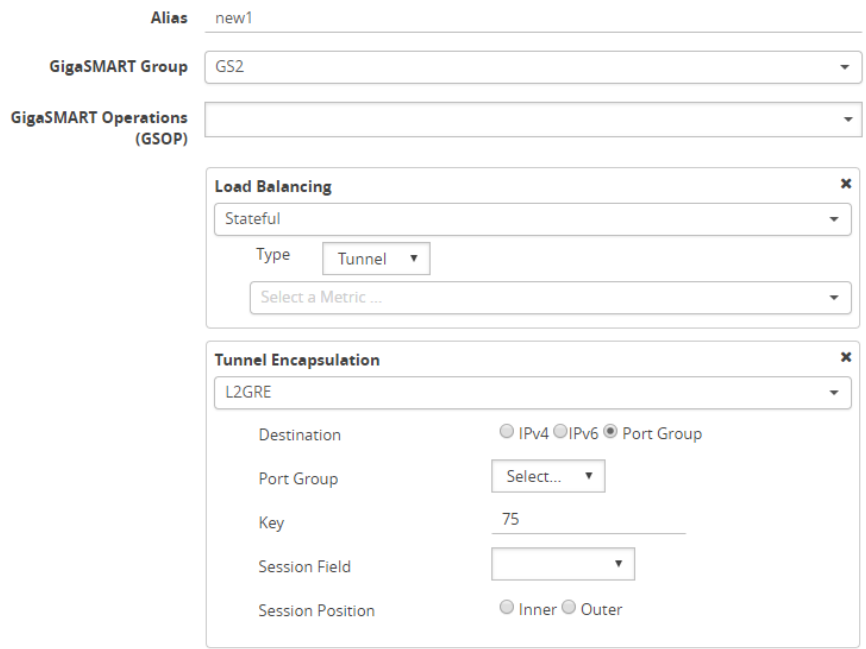


Figure 128 *New GigaSMART Operation for Stateful Load Balancing*

17. Click **OK**.

18. Go to **Maps > Maps**.
19. Click **New**, then type an alias for the map, select type Regular and subtype ByRule. Under Map Source and Destination, select a network port as the Source and a tool port as the Destination, then select the GigaSMART operation. Under Map Rules, configure a map rule. Refer to Figure 129New Map Configuration.

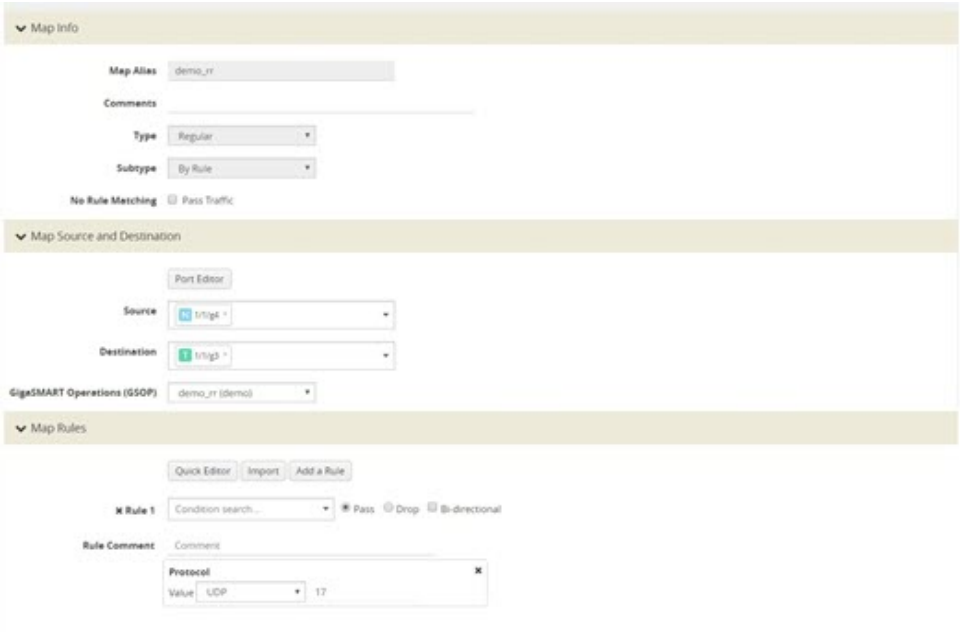


Figure 129 *New Map Configuration*

20. Click **OK**.

Example 3 – GigaSMART L2GRE Tunnel Encap Stateless LB

Example 3 configures stateless load balancing of tunnel traffic to tunnel endpoints based on a hash value.

Example 3 has the same configuration steps as Example 2 except for the GigaSMART operation (gsop) in Step 16. Under Load Balancing, select Stateless and the metric, such as Five Tuple. Under Tunnel Encapsulation, select L2GRE, Destination Port Group, and select the Port Group. Refer to Figure 130New GigaSMART Operation for Stateless.

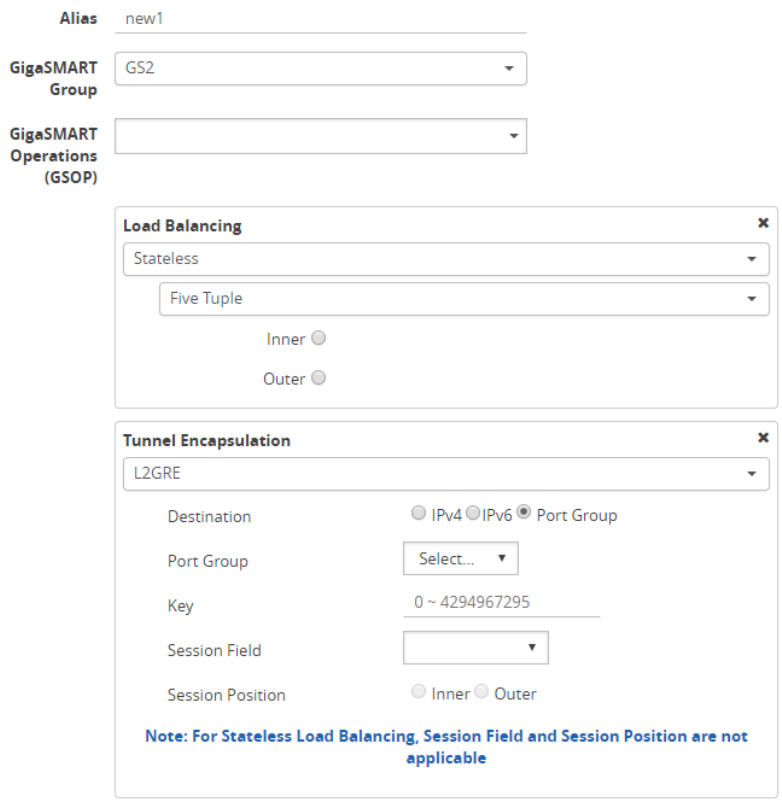


Figure 130 *New GigaSMART Operation for Stateless*

Example 4 – GigaSMART L2GRE Tunnel Decapsulation

In this example, an IP interface is configured on the network port. A GigaSMART operation for tunnel decapsulation is configured to decapsulate the filtered packets. A map is configured that uses the L2GRE tunnel decapsulation GigaSMART operation, which receives packets from the remote site over the Internet to the main office using the IP interface associated with a tool port and then forwards packets over the tool port. Staring with software version 5.4 GigaSMART L2GRE Tunnel Decapsulation provides support for IPv6 with load-balancing.

Task	Description	UI Steps
1.	Configure a network type of port and a tool type of port.	**a.** Select Ports > Ports > All Ports. **b.** Click Quick Port Editor. **c.** Use Quick search to find the ports to configure. **d.** Set the type for each port and select Enable.

Task	Description	UI Steps
2.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**a.** From the device view, select GigaSMART > GigaSMART Groups. **b.** Click New. **c.** Type an alias in the Alias field and enter an engine port in the Port List field. **d.** Click Save.
3.	Configure the IP interface with an IP address, subnet mask, default gateway, and MTU setting. Assign it to the GigaSMART group. The IP address must match the destination IP address specified at the sending end of the tunnel.	**a.** Select Ports > IP Interfaces. **b.** Click New. **c.** On the IP Interfaces page, in the Alias and Description fields, enter the name and description for the IP interface. **d.** Click the Ports field and select the port from the drop-down list. **e.** Enter the IP address, subnet mask, gateway, and MTU settings in the respective fields. **f.** Click on the GigaSMART Group field to select the GigaSMART group. **g.** Click Save.
4.	Configure the GigaSMART operation for tunnel decapsulation and assign it to the GigaSMART group. The tunnel decapsulation settings include the GRE key that identifies the source of the tunnel.	**a.** From the device view, select GigaSMART > GigaSMART Operations > GigaSMART Operation. **b.** Click New. **c.** Type an alias in the Alias field. **d.** From the GigaSMART Groups drop-down list, select the GigaSMART Group that you created in the second task. **e.** From the GigaSMART Operations (GSOP) drop-down list select Tunnel Decapsulation. **f.** Select L2GRE for the decapsulation type. **g.** Enter the GRE key in the Key field. **h.** Click Save.
5.	Create a map using the tunnel decapsulation GigaSMART operation, with packets coming from the Internet through the network port and being sent to the local tool port.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type an alias in the Map Alias field that will help you identify this map. **d.** Select Regular and By Rule for the map type and subtype. **e.** Specify the network and tool ports that you configured in task one in the Source and Destination fields, respectively. **f.** From the GSOP drop-down list, select the

Task	Description	UI Steps
		GigaSMART operation configured in task 4. **g.** Click Add a rule under Map Rules and create the following rule: Select IP Version from the drop-down list and select v4 for Version, and then select Pass. **h.** Click Save.

Example 5 – GigaSMART L2GRE IPv6 Tunnel Encamp/Decap with Load-Balancing

In this example, the encapsulation and decapsulation nodes are configured with IP interfaces using IPv6 addresses and load-balancing. IPv6 tunnel load-balancing feature supports the distribution of traffic across multiple IPv6 tunnel destination through the same tool port. Two types of load-balancing is supported, stateful and stateless.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

Step 1: Configure a tool type of port and a network type of port.

Create a map using the tunnel encapsulation GigaSMART operation, with packets coming from the network port and being sent to the Internet through the tool port.

1. Select **Ports > Ports > All Ports**.
2. Click **Quick Port Editor**.
3. Use **Quick search** to find the ports to configure.
4. Set the **type for each port** and select **Enable**.
   a. type: tool - port 1/3/x7
   b. type: network - 1/3/x8

Step 2: Configure a GigaSMART group and associate it with a GigaSMART engine port.

1. From the device view, select **GigaSMART > GigaSMART Groups.**
2. Click **New**.

3. Type an **alias** in the Alias field and enter an **engine port** in the Port List field.
4. Click **OK**.

Step 3: Configure the IP Interface

1. Select **Ports > IP Interfaces**.
2. Click **New**.
3. In the **Alias** and **Description** fields, enter the name and description for the IP interface.
4. Click the **Ports field** and select the port from the drop-down list.
5. Select the Port address: **IPv6**
6. Enter the **IP address, subnet mask, gateway, and MTU** settings in the respective fields.
7. From the **GS Group** drop-down list, select the required GigaSMART group.
8. Click **OK.**

Step 4: Configure the GigaSMART operation for tunnel encapsulation and load balancing and assign it to the GigaSMART group

1. From the device view, select **GigaSMART > GigaSMART Operations > GigaSMART Operation.**
2. Click **New**.

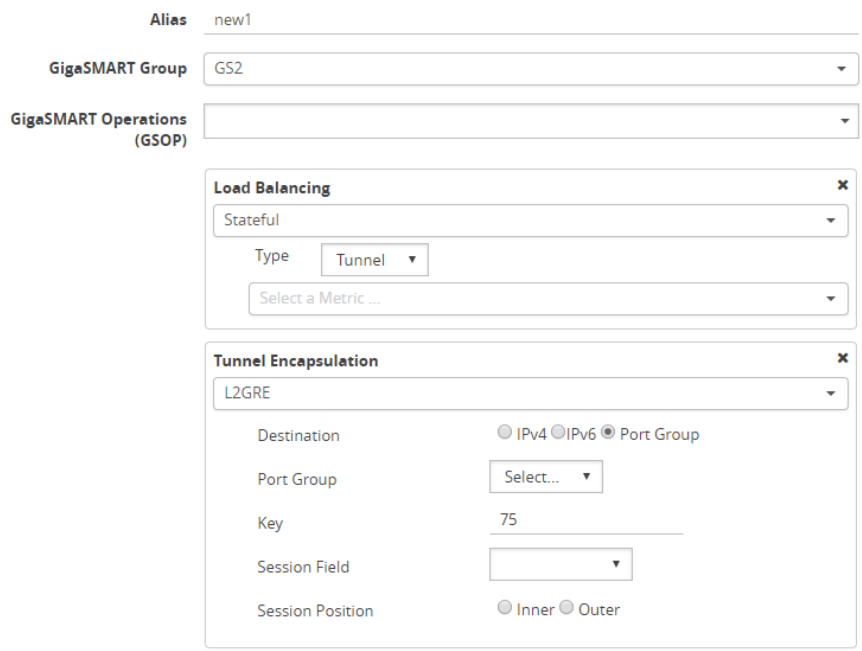


Figure 131 *GigaSMART Operation (GSOP)*

3. In the **Alias** field, enter a name for the GigaSMART operation.
4. From the **GigaSMART Group** drop-down list, select the **GigaSMARTGroup** that you created in the step 2.
5. From the **GigaSMART Operations (GSOP)** drop-down list, select **Tunnel Encapsulation**.
6. Select **L2GRE** for the encapsulation type.
7. Enter the **GRE key 123214** in the Key field.
8. Click **OK**.

Step 5: Create a map using the tunnel encapsulation GigaSMART operation

1. Select **Maps > Maps > Maps.**
2. Click **New**.

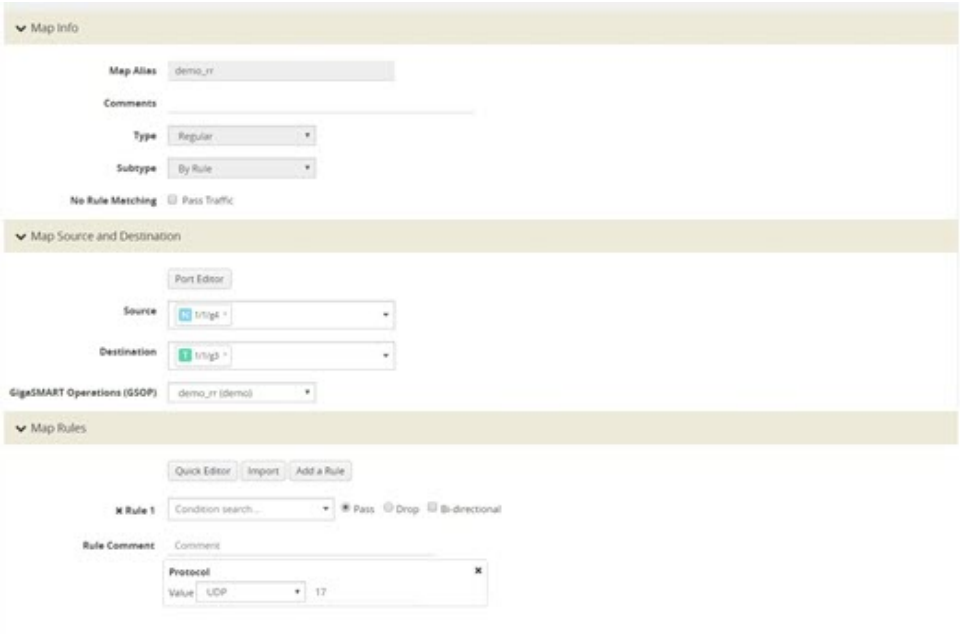


Figure 132 *Map Configuration*

3. Type an **alias** in the Map Alias field that will help you identify this map.
4. Select **Regular** and **By Rule** for the map type and subtype.
5. Specify the **network and tool ports** that you configured in step one in the Source and Destination fields, respectively.
6. From the GSOP drop-down list, select the **GigaSMART operation** configured in step 4.
7. Click **Add** a rule under Map Rules.
8. Select **IP Version** from the drop-down list and select **v4** for Version.
9. Select **Pass**.
10. Click **Add a rule** under Map Rules and create the following rule:
11. Select **IP Version** from the drop-down list and select **v6** for Version.
12. Select **Pass**.
    a. Source: **From - 1/4x24**
    b. Destination: **To: 1/4x7**
13. Click **OK**.

On the decapsulation node, configure the receiving end of the tunnel

Step 6: Configure a tool type of port and a network type of port.

1. Select **Ports > Ports > All Ports**.
2. Click **Quick Port Editor**.
3. Use **Quick search** to find the ports to configure.
4. Set the **type for each port** and select **Enable**.
   a. type: tool - port 1/4/x7
   b. type: network - 1/4/x24

Step 7: Configure a GigaSMART group and associate it with a GigaSMART engine port.

1. From the device view, select **GigaSMART > GigaSMART Groups.**
2. Click **New**.
3. Type an **alias** in the Alias field and enter an **engine port** in the Port List field.
   a. engine port: **1/3/e1**
4. Click **OK**.

Step 8: Configure the IP Interface with an IPv6 address, prefix length, default gateway, and MTU setting. Assign it to the GigaSMART group.

1. Select **Ports > IP Interfaces**.
2. Click **New**.
3. In the **Alias** and **Description** fields, enter the name and description for the IP interface.
4. Click the **Ports field** and select the port from the drop-down list.
5. Enter the **IP address**, **subnet mask, gateway, and MTU settings** in the respective fields.
6. From the **GS Group** drop-down list, select the required GigaSMART group.
7. Click **OK**.

Step 9: Configure the GigaSMART operation for tunnel decapsulation and assign it to the GigaSMART group.

1. From the device view, select **GigaSMART > GigaSMART Operations > GigaSMART Operation.**
2. Click **New**.
3. In the **Alias** field, enter a name for the GigaSMART operation.
4. From the **GigaSMART Groups** drop-down list, select the GigaSMART Group.
5. From the **GigaSMART Operations (GSOP)** drop-down list, select **Tunnel Decapsulation**.
6. Select **L2GRE** for the decapsulation type.
7. Enter the **GRE** key in the Key field.
8. Click **OK**.

Step 10: Create a map using the tunnel decapsulation GigaSMART operation.

1. Select **Maps > Maps > Maps**.
2. Click **New**.
3. Type an **alias** in the Map Alias field that will help you identify this map.
4. Select **Regular and By Rule** for the map type and subtype.
5. Specify the **network and tool ports** that you configured in task one in the Source and Destination fields, respectively.
6. From the **GSOP** drop-down list, select the GigaSMART operation.
7. Click **Add a rule** under Map Rules and create the following rule:
8. Select **IP Version** from the drop-down list and select **v4** for Version.
9. Select **Pass**.
10. Click **Add a rule** under Map Rules and create the following rule:
11. Select **IP Version** from the drop-down list and select **v6** for Version.
12. Select **Pass**.
    a. Source: **From - 1/4x24**
    b. Destination: **To: 1/4x7**
13. Click **OK**.

Orchestrated Workflow Configuration of GigaSMART L2GRE Tunnels

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

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

Field	Description
Tunnel Name	The name of the tunnel. **Note**: Alias must not have spaces or any of these characters: *"?;:,/%@.
Tunnel Description	The description of the tunnel endpoint.
Traffic Direction	Select the traffic direction of the tunnel. Choose DECAP for decapsulation or choose ENCAP for encapsulation. • If you choose Decap, select the IP Interface from the list. You can create one using the Create an IP interface option if you do not have an IP interface. Refer to IP Interfaces to learn more about creating an IP Interface. • If you choose Encap enter the Source Port. You can select from the drop-down list. To edit the Port type or to enable the port for Admin privileges use the Port Editor window.
Nodes	Select the Nodes on which you intend to create the tunnel.
GigaSMART Group	Select the GigaSMART groups created. If you do not have any groups, you can configure them using the **Create GS Groups** option. Refer Create GigaSMART Operations – A Summary to know more about configure groups.
Key	Select the key to be used for Decapsulation tunnel. You can choose a value from 0-4297964295. For Encapsulation Tunnel the key range is from 1-4297964295.
Rules	Configure the map rules that needs to be adhered to while decapsulating or encapsulating the traffic. Click on **Rules Editor** to configure the rules. **NOTE:** Configuring rules are mandatory for Encapsulation and Decapsulation GigaSMART tunnels.
Encapsulation IP interface	Enter the interface IP address of the node (Destination IP) for an encapsulation tunnel.
Time to Live	For encapsulation tunnel, enter the value of the time interval till which the session needs to be available. The value ranges from 1 to 255.
DSCP	Enter the Differentiated Services Code Point (DSCP) value for an encapsulation.
Precedence	Add a value between 0-7 for an encapsulation GigaSMART L2GRE tunnel. This field is applicable only for Gen 2 cards.
Remote Tunnel IP	Enter the Remote Tunnel IP values for and encapsulation tunnel.

Field	Description
	NOTE: If the Tunnel is load balanced ,Remote Tunnel Port group must be configured.
Destination Port	Select the destination port for the decapsulation tunnel. You can edit the port details from the Port editor screen. **NOTE:** If the Tunnel is load balanced ,Destination Port group must be configured.
Additional GigaSMART operations	Enable this option to add GigaSMART operations. When enabled you will be provided with GigaSMART Operations section which would allow you to choose among the following: • Add Header • Add Trailer • De-duplication • Load Balancing • Masking • Remove Trailer • Slicing

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

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

GigaSMART VXLAN Tunnel Encapsulation

Required License for VXLAN Encapsulation: Base (GigaVUE-HC1 Gen 3), Tunneling (GigaVUE-HC3 Gen 3)

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

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

Starting in software version 5.11, support for VXLAN tunnel origination is added to GigaSMART. After receiving packets on the network ingress ports, they are sent to the GSOP chain. When a packet egresses out of the GS card, it is encapsulated with a VXLAN header, along with all outer headers (UDP, IP, and MAC), and then sent out to the appropriate multicast ID.

Use GigaSMART VxLAN tunnel encapsulation to send traffic from one GigaSMART node over the Internet to a second GigaSMART node using VxLAN encapsulation.

GigaSMART VxLAN tunnel encapsulation provides the VxLAN tunnel initiation and encapsulation on the tool port at the sending end of the tunnel (for example, at a remote site)

The GigaSMART at the remote site encapsulates the filtered packets, adds an encapsulation header, and routes it to the main office site. The encapsulation protocol is VxLAN and the delivery protocol is IP or IPv6, so the encapsulation header consists of Ethernet + IP + VxLAN or Ethernet + IPv6 + VxLAN headers.

The encapsulated packet is sent out of the tool port, which is connected to the public network (the Internet). This packet is routed in the public network to reach the main office site. It ingresses at the routed network port of the GigaVUE-FM,GigaVUE HC Series node at the main office.

For statistics for encapsulated packets, refer to Display VXLAN Tunnel Encapsulation Statistics.

VXLAN Header Length

The VXLAN header length is as follows:

Header	Length in Bytes
With Key	50 bytes consisting of 14 Ethernet + 20 IP + 8 UDP + 8 VxLAN Key.

VXLAN IPv4 Encapsulation/Decapsulation

GigaSMART VXLAN IPv4 lets you route filtered traffic to the remote

end using IPv4-based VXLAN tunneling. At the receiving end, filtered traffic is sent to GigaSMART, which adds an VXLAN header and a IPv4 header to make it routable.

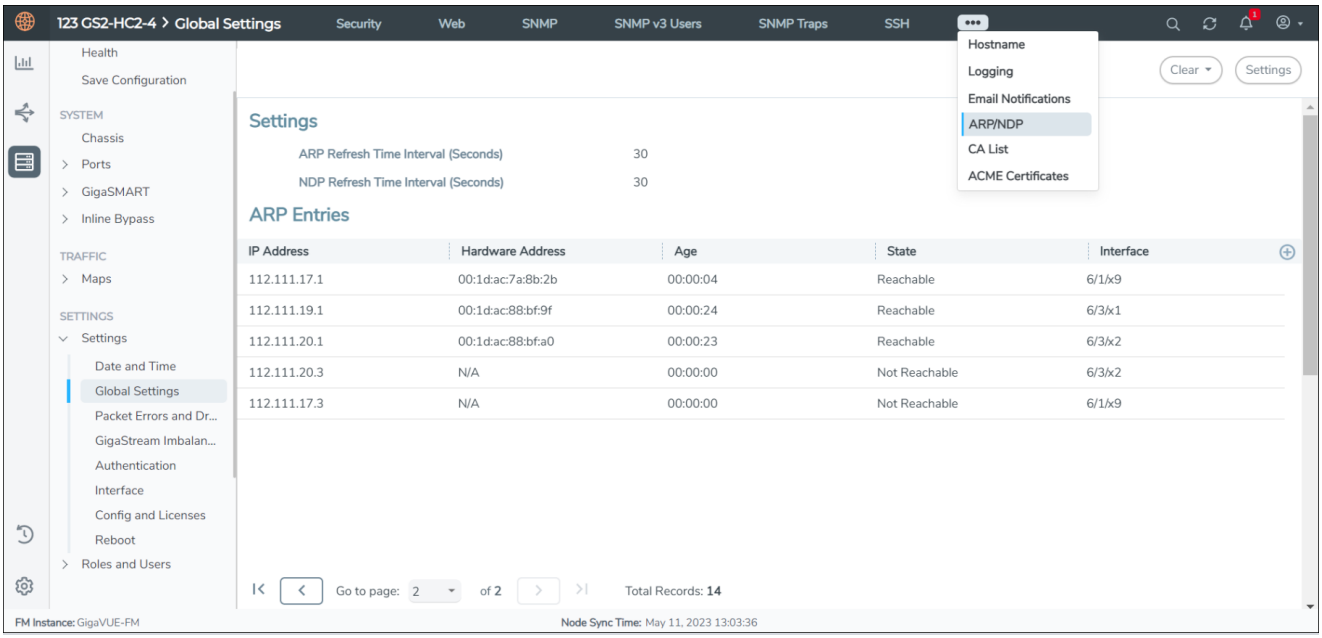
The remote end decapsulates the packet and sends it to the tool port.

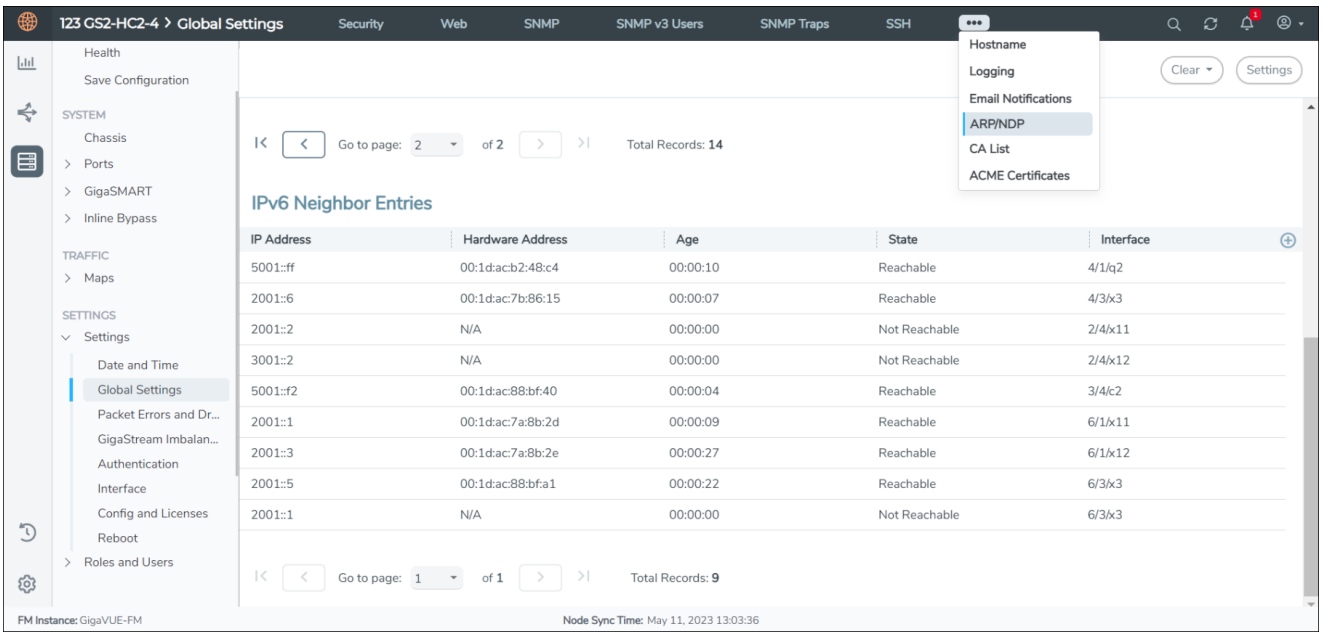
GigaVUE-FM,GigaVUE HC Series nodes act as VXLAN encapsulation and decapsulation devices.

The IPv4 protocol is used to deliver all packets received in the encap tunnel to the termination node using the configured source and destination IPv4 address. The tunnel termination (decap) node strips the IPv4 + VXLAN header and sends the payload to the tool port.

The arp protocol is used by the tool port on the encapsulation node for arp request and arp response messages to resolve the gateway MAC address as well as destination tool

MAC address in local network and respond to arp request messages received from the gateway in the tunnel decapsulation/termination node. arp echo request/reply messages are also sent and received.





VXLAN IPv6 Encapsulation/Decapsulation

GigaSMART VxLAN IPv6 lets you route filtered traffic to the remote

end using IPv6-based VxLAN tunneling. At the receiving end, filtered traffic is sent to GigaSMART, which adds an VXLAN header and a IPv6 header to make it routable.

The remote end decapsulates the packet and sends it to the tool port.

GigaVUE-FM,GigaVUE HC Series nodes act as VXLAN encapsulation and decapsulation devices.

The IPv6 protocol is used to deliver all packets received in the encap tunnel to the termination node using the configured source and destination IPv6 address. The tunnel termination (decap) node strips the IPv6 + VXLAN header and sends the payload to the tool port.

The ICMPv6 protocol is used by the tool port on the encapsulation node for Neighbor Solicitation (NS) and Neighbor Advertisement (NA) messages to resolve the gateway MAC address as well as destination tool MAC address in local network and respond to NS messages received from the gateway in the tunnel decapsulation/termination node. ICMPv6 echo request/reply messages are also sent and received.

Display VXLAN IPV4/IPv6 Tunnel Statistics

To view IP Interfaces statistics, **select Ports > IP Interfaces > Statistics** to open the IP Interfaces Statistics page.

The IP destinations pane displays the gateway or local tool status as **UP** if neighbor discovery is completed with gateway or local tool or **down** if neighbor discovery failed. Neighbor discovery is made only on the encapsulation node. On the decapsulation node, the gateway or local tool status will be **Not Applicable**.

Display VXLAN Tunnel Encapsulation Statistics

To display Layer 2 VXLAN tunnel encapsulation statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics**. The statistics for tunnel encapsulation will be in the row labeled Tunnel Encap in the GS Operations column.

Refer to Tunnel Encapsulation Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

Orchestrated Workflow Configuration of GigaSMART VXLAN Tunnels

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

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

Field	Description
Tunnel Name	The name of the tunnel. **NOTE:** Alias must not have spaces or any of these characters: *"?;,/%@.
Tunnel Description	The description of the tunnel endpoint.
Traffic Direction	Select the traffic direction of the tunnel. Choose DECAP for decapsulation or choose ENCAP for encapsulation. • If you choose Decap, select the IP Interface from the list. You can create one using the Create an IP interface option if you do not have an IP interface. Refer to IP Interfaces to learn more about creating an IP Interface. • If you choose Encap enter the Source Port. You can select from the drop-down list. To edit the Port type or to enable the port for Admin privileges use the Port Editor window. **NOTE:** For Encapsulation Tunnel, only Gen 3 nodes are supported.
Nodes	Select the Nodes on which you intend to create the tunnel.
GigaSMART Group	Select the GigaSMART groups created. If you do not have any groups, you can configure them using the **Create GS Groups** option. Refer Create GigaSMART Operations – A Summary to know more about configure groups.
Time to Live	or encapsulation tunnel, enter the value of the time interval till which the session needs to be available. The value ranges from 1 to 255
DSCP	Enter the Differentiated Services Code Point (DSCP) value for an encapsulation.
Rules	Configure the map rules that needs to be adhered to while decapsulating or encapsulating the traffic. Click on **Rules Editor** to configure the rules. **NOTE:** Configuring rules are mandatory for Encapsulation and Decapsulation GigaSMART tunnels.
Encapsulation IP interface	Enter the interface IP address of the node (Destination IP) for an encapsulation tunnel.
Remote Tunnel IP, L4 Port & Source Port	Enter the IP address, L4 Port and Source Port values for an encapsulation tunnel. **NOTE:** For Encapsulation Tunnels, the destination and L4 port can be configured only on three ports 4789, 8472 and 48879. **NOTE:** For Decapsulation Tunnel, the Destination and Source port range is from 1 - 65535.

Field	Description
Key	Select the key to be used for Decapsulation tunnel. You can choose a value from 0-4297964295. For Encapsulation Tunnel the key range is from 1-4297964295.
Destination Port	Select the destination port for the decapsulation tunnel. You can edit the port details from the Port editor screen.
Additional GigaSMART operations	Enable this option to add GigaSMART operations. When enabled you will be provided with GigaSMART Operations section which would allow you to choose among the following: • Add Header • Add Trailer • De-duplication • Load Balancing • Masking • Remove Trailer • Slicing

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

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

IP Fragmentation and Reassembly on L2GRE and GMIP Tunnels

Required Licenses:
IP Fragmentation on L2GRE and GMIP Tunnels-Advanced Tunneling license on GigaVUE-HC3. Referred to as “Tunneling license” on GigaVUE-HC1. Reassembly on L2GRE and GMIP Tunnels -Base (GigaVUE-HC3, GigaVUE-HC1-Plus), Tunneling (GigaVUE-HC1).

Starting in software version 4.6, L2GRE and GMIP tunnels support IP fragmentation and reassembly of packets. IP fragmentation occurs with encapsulation. Fragmented packets are sent on the tool port at the sending end of the tunnel (for example, at a remote site). IP reassembly occurs with decapsulation. Fragmented packets reaching the network port at the receiving end of the tunnel (for example, at a main office site), are decapsulated and reassembled before being sent to the destination.

IP Fragmentation on Encapsulation

The tool port at the remote site is configured with a tunnel MTU. If a packet exceeds the tunnel MTU, the packet will be fragmented, and the fragmented packets will be sent out the IP interface.

> **NOTE:** The first fragmented packet contains the tunnel header (Eth+IP+GRE). The rest of the fragments have the Ethernet and IP headers.

The packet size plus the tunnel header size is calculated and checked against the tunnel MTU. For example, if the tunnel MTU is 1518 and the packet is 1526, the packet exceeds the tunnel MTU. If the tunnel MTU is 1518 and the packet is 1518, the packet will also exceed the tunnel MTU due to the addition of the tunnel header.

IPv6 tunnel supports Path MTU. When a GigaSMART (GSOP) is associated to a map, Path MTU discovery message is sent to the tunnel destination and the Path MTU learnt is used for IPv6 fragmentation.

IP Reassembly on Decapsulation

The network port at the main office site receives the fragmented packets sent from the remote site. The tunnel header is removed from all fragmented packets, and they are buffered in memory. After all the fragmented packets are available, they are reassembled. The reassembled packet is then sent to the tool.

Notes and Considerations

Take into account the following notes and considerations:

Feature	Description
IPv4 and IPv6 Support	IPv4 and IPv6 packets are supported. **Note:** To avoid the overhead and improve the performance, existing GMIP IPv4 tunneling does not calculate the UDP Checksum on the Encapsulated Packets. The same will be adopted for IPv6 and IPv6 RFCs, where the checksum value will be set to 0. UDP checksum of the out header is not mandatory in the IPv6 tunneling.
Always Enabled	IP fragmentation and reassembly are always enabled. No configuration is required.
Tunnel MTU	The tunnel MTU is configured using the MTU field on the **IP Interfaces** configuration page. (Select **Ports > IP Interfaces**, and then click **New** to open the page.)

Feature	Description
	The MTU is fixed at 9400 for all network/tool ports on the following platforms: o GigaVUE-HC1 o GigaVUE-HC3 o GigaVUE-FM,GigaVUE HC Series-TA100
Encapsulation Statistics	The encapsulation statistics count the number of fragmented packets. Refer to Display GMIP Tunnel Encapsulation Statistics and Display L2GRE Tunnel Encapsulation Statistics. For definitions, refer to Tunnel Encapsulation Statistics Definitions.
Decapsulation Statistics	The decapsulation statistics count the number of reassembled packets. Refer to Display GMIP Tunnel Decapsulation Statistics and Display L2GRE Tunnel Decapsulation Statistics. For definitions, refer to Tunnel Decapsulation Statistics Definitions.
GigaSMART Engine Ports	GigaSMART operations with a tunnel component can be assigned to GigaSMART groups consisting of multiple GigaSMART engine ports. Refer to Groups of GigaSMART Engine Ports for details.

GigaSMART ERSPAN Tunnel Decapsulation

Required License for ERSPAN Decapsulation: Advanced Tunneling (GigaVUE-HC1 Gen3, GigaVUE-HC3, GigaVUE-HC1-Plus, and GigaVUE-HCT Gen 3), Tunneling (GigaVUE-HC1 and GigaVUE-HCT Gen 3).

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

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

Some Cisco equipment provides the ability to mirror monitored traffic to a remote destination through an ERSPAN tunnel. Using ERSPAN tunnel decapsulation, GigaSMART can act as the receiving end of an ERSPAN tunnel, decapsulating mirrored traffic sent over the Internet from a Cisco switch or router.

Both ERSPAN Type II and Type III header decapsulation are supported. For ERSPAN Type III details, refer to ERSPAN Type III.

You can configure a GigaSMART-enabled node to act as the receiving end of an ERSPAN tunnel by configuring a GigaSMART **Tunnel Decapsulation** operation with type set to **ERSPAN** and a **Flow ID** matching the sending end of the tunnel.

The high-level steps are as follows:

1. Configure an IP interface associated with network port and assign an IP address, subnet mask, and default gateway to the IP interface. The IP address must match the destination IP address specified at the sending end of the tunnel.
2. Create a GigaSMART operation with an ERSPAN tunnel decapsulation component. The decapsulation settings include the same flow ID specified at the sending end of the tunnel. The flow ID is a value from 0 to 1023. Use this options when decapsulating traffic received over a Cisco-standard ERSPAN tunnel. A flow ID of 0 decapsulates all ERSPAN tunnel traffic regardless of flow ID.
3. For ERSPAN Type III, a trailer timestamp may be specified.
4. Bind the GigaSMART operation to the IP interface associated with network port as part of a map that distributes arriving traffic to local tool ports for analysis with local tools.

For example configurations, refer to ERSPAN Tunnel Header Removal and GigaSMART ERSPAN Tunnel Decapsulation.

For an example of APF and ERSPAN tunneling, refer to ERSPAN Tunneling.

ERSPAN Type III

ERSPAN Type III is similar to ERSPAN Type II but has a hardware timestamp in the packet. The hardware timestamp needs to be translated into a usable timestamp.

The UTC timestamp can be calculated, based on the reference hardware timestamp and the reference UTC timestamp carried in marker packets that are periodically sent over UDP. The calculated UTC timestamp can then be appended to the packets as a trailer.

Marker packets have a fixed length and are identified by a signature of 0xA5A5A5A5. If the marker packet session ID matches the ERSPAN session ID, the UTC timestamp can be extracted from the marker packet. An ERSPAN session is defined by a map that uses an ERSPAN GigaSMART operation (gsop).

There are three timestamp formats: **None**, **GigaSMART**, and **X12-TS** (for PRT-H00-X12TS). The timestamp options are set from the GigaSMART Group page, which is accessed by selecting **GigaSMART > GigaSMART Groups > GigaSMART Groups**, and then clicking **New** or editing an existing GigaSMART Group. Figure 133ERSPAN Type III Timestamp Formats on GigaSMART Groups Page shows the timestamp format options. If the timestamp format is **Disabled**, ERSPAN Type III packets are parsed and the ERSPAN

header is removed by GigaSMART. The inner packets are forwarded to a tool port. If the timestamp format is **GigaSMART** or **X12-TS**, a trailer containing the recovered timestamp is added to the inner packets before they are forwarded to a tool port.

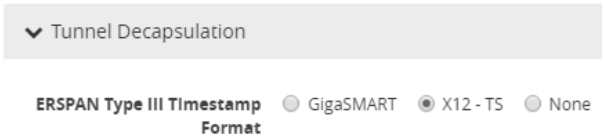


Figure 133 *ERSPAN Type III Timestamp Formats on GigaSMART Groups Page*

The GigaSMART timestamp is added to the GigaSMART trailer. For the format of the GigaSMART trailer, refer to GigaSMART Trailer Reference. The x12-ts timestamp is added to the PRT-H00-X12TS trailer. For the format of the PRT-H00X12TS trailer, refer to the *GigaVUE-OS CLI Reference Guide*.

Only 10 ERSPAN sessions are supported per GigaSMART Group (gsgroup) when the timestamp format is configured to **GigaSMART** or **X12-TS**.

In summary for ERSPAN Type III encapsulation, **GigaSMART** does the following:

- strips encapsulating Ethernet + outer IP + GRE + ERSPAN Type III headers from incoming packets
- uses the timestamp field in ERSPAN packets and calculates the UTC timestamp, based on the timestamp in marker packets
- appends the UTC timestamp to the GigaSMART trailer or the PRT-H00-X12TS trailer if either **GigaSMART** format or PRT-HD00-X12TS (**X12-TS**) format is configured
- forwards packets to tool ports

ERSPAN Granularity

ERSPAN granularity is a setting that can be configured on the Cisco switch for the level of detail of the hardware timestamp in marker packets.

A marker packet will be considered overdue if it does not arrive by the following times:

- 00: Granularity—overdue after 119 hours
- 01: Granularity—overdue after 430 seconds (7 minutes)
- 10: 1588 PTP—overdue after 4.3 seconds

ERSPAN statistics include a count of overdue packets. Refer to Display ERSPAN Statistics for how to display the output and to ERSPAN Statistics Definitions for descriptions of these statistics.

PRT-H00-X12TS Unique ID

For the PRT-H00-X12TS format, you can obtain a unique ID identifying the port on which packets arrive. Use the following CLI command to display the mapping of ports to unique IDs:

```
(config) # show apps netflow port-id
============================
Port        Netflow port-id
----------------------------
1/1/x1      1
1/1/x2      2
1/1/x3      3
1/1/x4      4
1/1/x5      5
1/1/x6      6
1/1/x7      7
1/1/x8      8
1/1/x9      9
1/1/x10     10
1/1/x11     11
1/1/x12     12
----------------------------
```

Configure GigaSMART Operations for ERSPAN

Use the GigaSMART Operation (GSOP) page to configure the ERSPAN decapsulation types and options. For example, you can specify an ERSPAN flow ID, from 0 to 1023. Use this option when decapsulating traffic received over a Cisco-standard ERSPAN tunnel. Both ERSPAN Type II and Type III header decapsulation are supported.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

ERSPAN Tunnel Header Removal

To configure a tunnel to capture the ERSPAN packets, remove the ERSPAN header, and then forward the packets to a tool port, set the ERSPAN Decapsulation Flow ID to zero when creating the GigaSMART operation as shown in Figure 134Decapsulation Flow ID Set to Zero..

NOTE: A flow ID of zero is a wildcard value that matches all flow IDs.

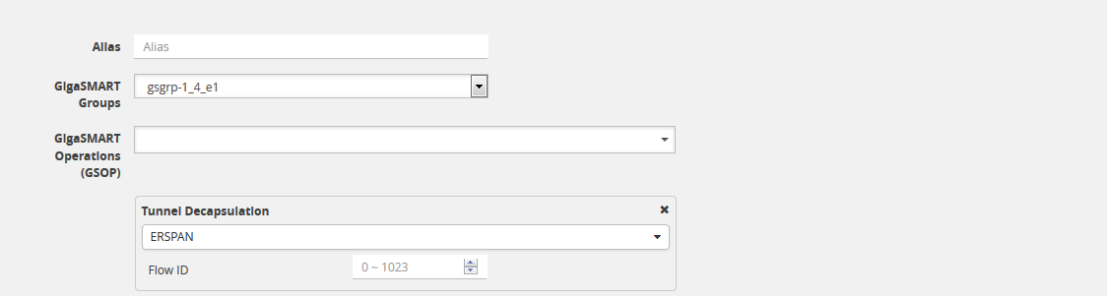


Figure 134*Decapsulation Flow ID Set to Zero.*

In the following example, a tunnel is configured to capture ERSPAN packets, then the ERSPAN header is removed and the packets are forwarded to a tool port.

Task	Description	UI Steps
1.	Configure a tool type of port.	**a.** Select Ports > All Ports. **b.** Click Quick Port Editor. **c.** Use Quick search to find the ports to configure. For example,1/1/g1. **d.** Set the type to Tool and select Enable. **e.** Click OK.
2.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type an alias in the Alias field (for example, gsgrp1) and enter an engine port in the Port List field (for example 1/3/e1). **d.** Click Save.
3.	Configure the IP interface.	**a.** Select Ports > IP Interfaces. **b.** Click New. **c.** On the IP Interfaces page, in the Alias and Description fields, enter the name and description for the IP interface. **d.** Click the Ports field and select the port from the drop-down list. **e.** Enter the IP address, subnet mask, gateway, and MTU settings in the respective fields. For example, port 1/1/g2, IP address 10.10.10.10, mask 255.255.225.0, gateway 0.10.10.1, and MTU 1500. **f.** Click on the GigaSMART Group field to select the GigaSMART group. **g.** Click Save.

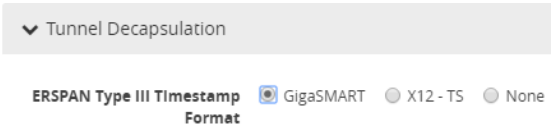
Task	Description	UI Steps
4.	Configure the GigaSMART operation and assign it to the GigaSMART group. **Note**: A flow ID of zero is a wildcard value that matches all flow IDs.	**a.** From the device view, select GigaSMART > GigaSMART Operations > GigaSMART Operation. **b.** Click New. **c.** Type an alias in the Alias field. **d.** From the GigaSMART Groups drop-down list, select the GigaSMART Group that you created in the second task. **e.** From the GigaSMART Operations (GSOP) drop-down list, select Tunnel Decapsulation. **f.** Select ERSPAN for the decapsulation type. **g.** Enter a value of 0 in the Flow ID field. The configuration should look like the example shown in Figure 134Decapsulation Flow ID Set to Zero.. **h.** Click Save.
5.	Create a map.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type an alias in the Map Alias field that will help you identify this map. **d.** Select Regular and By Rule for the map type and subtype. **e.** Specify the network and tool ports in the Source and Destination fields, respectively. **f.** From the GSOP drop-down list, select the GigaSMART operation configured in task 4. **g.** Click Add a rule under Map Rules and create the following rule: Select IPv4 Protocol from the drop-down list and select GRE for Value, and then select Pass. **h.** Click Save.

ERSPAN Type III Tunnel Header Removal

In this example, a tunnel is configured to capture ERSPAN packets. ERSPAN Type III packets are parsed, the ERSPAN header is removed, and the timestamp is calculated. A timestamp trailer is added before the packets are forwarded to a tool port.

NOTE: A flow ID of zero is a wildcard value that matches all flow IDs.

Task	Description	UI Steps
1.	Configure a port of type tool.	**a.** Select Ports > Ports > All Ports. **b.** Click Quick Port Editor. **c.** In the Quick View Editor, find the port to configure. **d.** Set Type to Tool. **e.** Select Enable **f.** Click OK. **g.** Close the Quick Port Editor.
2.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Enter a name in the Alias field **d.** Select the engine port in the Port List field. **e.** Click Save.
3.	Configure the IP interface.	**a.** Select Ports > IP Interfaces. **b.** Click New. **c.** On the IP Interfaces page, in the Alias and Description fields, enter the name and description of the IP interface. **d.** Click the Ports field and select the tool port from the drop-down list. **e.** Enter the IP address, subnet mask, gateway, and MTU settings in the respective fields. For example, port 1/1/g2, IP address 10.10.10.10, mask 255.255.225.0, gateway 0.10.10.1, and MTU 1500. **f.** Click on the GigaSMART Group field to select the GigaSMART group. **g.** Click Save.
4.	Configure the GigaSMART operation and assign it to the GigaSMART group. **Note**: A flow ID of zero is a wildcard value that matches all flow IDs.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Type an alias in the Alias field. **d.** From the GigaSMART Groups drop-down list, select the GigaSMART Group that you created in the second task. **e.** From the GigaSMART Operations (GSOP) drop-down list, select Tunnel Decapsulation. **f.** Select ERSPAN for the decapsulation type. **g.** Enter a value of 0 in the Flow ID field. The

Task	Description	UI Steps
		configuration should look like the example shown in Figure 134Decapsulation Flow ID Set to Zero.. **h.** Click Save.
5.	Configure a timestamp trailer format.	**a.** From the device view, select GigaSMART > GigaSMART Group. **b.** Select the GigaSMART Group created in Task 2. **c.** Under GigaSMART Parameters, go to Tunnel Decapsulation. **d.** For ERSPAN Type III Timestamp Format, select GigaSMART  **e.** Click Save.
6.	Create a map. The map contains a rule to allow marker packets (UDP) to be processed.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type an alias in the Map Alias field that will help you identify this map. **d.** Select Regular and By Rule for the map type and subtype. **e.** Specify a network ports in the Source fields. **f.** Select the tool port configured in Task 1 in the Destination field **g.** From the GSOP drop-down list, select the GigaSMART operation configured in task 4. **h.** Click Add a Rule and create the first rule. Select Pass, then select IPv4 Protocol, and then select GRE fro Value. **i.** Click Add a Rule and create the second rule. Select Pass, then select IPv4 Protocol, and then select UDP for Value. **j.** Click Save.

Display ERSPAN Statistics

To display ERSPAN statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics**. The ERSPAN statistics will be in the row labeled Tunnel Decap in the GS Operations column.

Refer to ERSPAN Statistics Definitions for descriptions of these statistics as well as to GigaSMART Operations Statistics Definitions.

Orchestrated Workflow Configuration of GigaSMART ERSPAN Tunnels

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

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

Field	Description
Tunnel Name	The name of the tunnel. **NOTE:** Alias must not have spaces or any of these characters: *"?;:,/%@.
Tunnel Description	The description of the tunnel endpoint.
Traffic Direction	Select the traffic direction of the tunnel. Choose DECAP for decapsulation . **NOTE:** GigaSMART ERSPAN tunneling does not support Encapsulation.
Nodes	Select the Nodes on which you intend to create the tunnel.
IP Interface	Select the IP Interface from the list. You can create one using the Create an IP interface option if you do not have an IP interface. Refer to IP Interfaces to learn more about creating an IP Interface.
GigaSMART Group	Select the GigaSMART groups created. If you do not have any groups, you can configure them using the **Create GS Groups** option. Refer Create GigaSMART Operations – A Summary to know more about configure groups.
Key	Select the key to be used for decapsulation. You can choose a value from 0-1023.
Rules	Configure the map rules that needs to be adhered to while decapsulating or encapsulating the traffic. Click on **Rules Editor** to configure the rules.
Additional GigaSMART operations	Enable this option to add GigaSMART operations. When enabled you will be provided with GigaSMART Operations section which would allow you to choose among the following: • Add Header • Add Trailer • De-duplication

Field	Description
	• Load Balancing • Masking • Remove Trailer • Slicing
Destination Node	Select the destination node for the decapsulation tunnel.
Destination Port	Select the destination port for the decapsulation tunnel. You can edit the port details from the Port editor screen.

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

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

GigaSMART VXLAN Tunnel Decapsulation

Required License for VXLAN Decapsulation: Base (GigaVUE-HC3), Tunneling (GigaVUE-HC1)

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

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

Starting in software version 5.3, support for VXLAN tunnel termination is added to GigaSMART. VXLAN encapsulated packets originating from any device, such as the Gigamon cloud or a customer-specific device, will be received on a network port, then will be terminated at GigaSMART. The VXLAN payload (the inner packet) will be sent to tools. The reassembly of fragmented IP packets is also supported.

This section only includes VXLAN tunnel termination. It does not include VXLAN origination. To terminate a custom tunnel header that is not known to GigaSMART, use custom tunnel termination. Refer to GigaSMART Custom Tunnel Decapsulation

You can configure a GigaSMART-enabled node to act as the receiving end of a VXLAN tunnel by configuring a GigaSMART **tunnel-decap** operation with **type** set to **vxlan**. The high-level steps are as follows:

1. Configure an IP interface associated with network port and assign an IP address, subnet mask, and default gateway to the IP interface. The gateway forwards the encapsulated packet to the network port.
2. Create a GigaSMART operation with a **vxlan** decapsulation component.

3. Bind the GigaSMART operation to the IP interface associated with network port as part of a map.

At GigaSMART, VXLAN encap packets are received on the network port. After validation of the source port, destination port, and VXLAN Network Identifier (VNI) of the packet, the VXLAN tunnel header will be removed and the inner payload will be sent to a subsequent GSOP or to the tools. The VNI in the VXLAN header is validated against the user VNI provided. If it does not match, the packet will be dropped and counted as an error.

A VXLAN packet is identified using the **portdst** parameter. The destination port can be 4789, or any user-configured port number from 1 to 65535.

For an example configuration, refer to VXLAN Tunnel Termination Example

NOTE: GigaSMART operations with a tunnel component can be assigned to GigaSMART groups consisting of multiple GigaSMART engine ports.

VXLAN Tunnel Termination Example

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

To configure a VXLAN tunnel termination:

Step 1: Configure a Tool Port

1. Select **Ports > Ports > All Ports**.
2. Click **Quick Port Editor**.
3. Configure an available port as follows:
   - Select Tool in the **Type** field.
   - Enter an alias in the **Alias** field. For example, 1/4/x2.
   - Check the **Enable** check box.
4. Click **OK** to save the port.
5. Close the Quick Port Editor.

Step 2: Configure a GigaSMART group and associate it with a GigaSMART engine port.

1. From the device view, select **GigaSMART >GigaSMART Groups > GigaSMART Groups**.
2. Click **New**.
3. Type an aliasin the **Alias** field. For example, gsg.
4. Select an engine port in the **Port List** field. For example, 1/3/e2. (All engine ports have an 'e'.)
5. Click **OK** to save the GigaSMART group.

Step 3: Configure the IP Interface

1. Select **Ports > IP Interfaces**.
2. Click **New**.
3. In the **Alias** and **Description** fields, enter the name and description for the IP interface.
4. From the **Port** field, select any available network port. In this example, port 1/2/x1.
5. Complete the fields to configure the IP Interface:
   - Enter an **IP Address**. For example, 10.115.9.5.
   - Enter a **Mask**. For example, 255.255.255.0.
   - Enter a **Gateway**. For example, 10.115.9.1.
   - Enter the maximum transmission unit (MTU) for this port in the **MTU** field. For example, 1500.
   - Select the **GigaSMARTGroup** you created in step 2 of this process (gsg).
6. Click **OK** to save the IP interface configuration.

Step 4: Configure GigaSMART operation and assign to the GigaSMART Group

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.
2. Click **New**.
3. Type an alias in the **Alias** field to identify this operation. For example, vxlan2.
4. For **GigaSMART Groups**, select the group created in step 2 of this process (gsg).
5. For **GigaSMART Operations (GSOP)**, select Tunnel Decapsulation.

A Tunnel Decapsulation form appears. Complete the fields as follows:
- Select “VXLAN” as the **Type**.
- Enter a **Source Port**.
- Enter a **Destination Port**.
- Enter a **VNI**.

6. Click **OK** to save the GSOP.

Step 5: Create a Map

1. Select **Maps > Maps > Maps**.
2. Click **New**.
3. Configure the map.
   - Type an alias the **Alias** field.
   - For **Type**, leave the default (Regular).
   - For **Subtype**, leave the default (By Rule).
   - For **Source**, select the IP interface you configured in step 3 (1/2/x1).
   - For **Destination**, select the tool port you configured in step 1 (1/4/x2).
   - Select the GigaSMART Operations (GSOP) associated with VXLAN decapsulation.
4. Under Map Rules, click **Add a Rule**.
   - Select **IP Version** from the drop list and set IP Version to **v4** when prompted.
   - Select **Pass** radio button for rule type.
   - Click **OK**.

Display VXLAN Tunnel GSOP

To display the VXLAN Tunnel GigaSMART operation:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP)**.
2. Select the VXLAN tunnel GSOP that you created in step 4 (vxlan2). The GSOP quick view appears.

Display VXLAN Tunnel Statistics

To display VXLAN tunnel statistics:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP)> Statistics**.

   The Statistics page displays all GSOP statistics in a table format.

2. In the table view, click the VXLAN tunnel GSOP alias that you created in step 4 (vxlan2) to display the Statistics quick view.

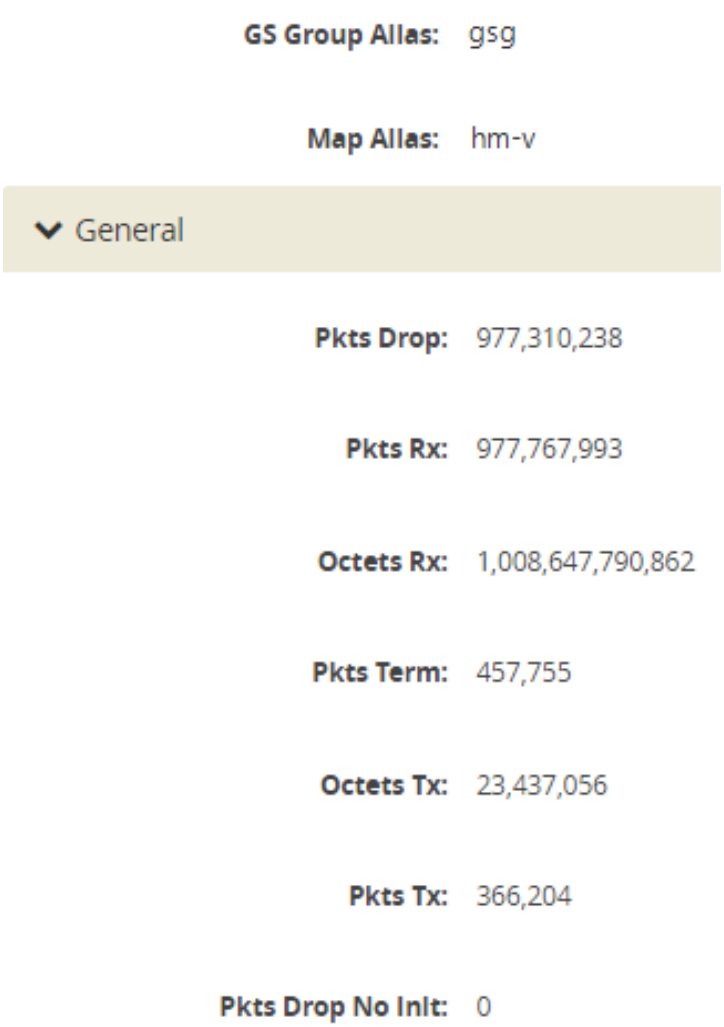


Figure 135 *GSOP VXLAN Statistics Quick View*

GigaSMART Custom Tunnel Decapsulation

Required Licenses for Custom Decapsulation: Base (GigaVUE-HC3), Tunneling (GigaVUE-HC1)and Header Stripping

Starting in software version 5.3, support for custom tunnel termination is added to GigaSMART. Use custom tunnel termination to terminate a custom tunnel header that is received at the IP interface that is associated with a network port, but is not known to GigaSMART. The destination IP and MAC addresses must match the IP and MAC addresses of the network tunnel.

The packets that are successfully received at GigaSMART on a custom tunnel can be stripped, after some validations are performed, or can be sent to tools. The existing generic Header Stripping operation can be leveraged to remove the tunnel header if required. The reassembly of fragmented IP packets is also supported.

You can configure a GigaSMART-enabled node to act as the receiving end of a tunnel by configuring a GigaSMART **tunnel-decap** operation with **type** set to **custom** and Layer 4 (L4) source and destination ports. The high-level steps are as follows:

1. Configure an IP interface associated with network port and assign an IP address, subnet mask, and default gateway to the IP interface. The gateway forwards the encapsulated packet to the IP interface that is associated with a network port.
2. Create a GigaSMART operation (GSOP) with a **custom** decapsulation component.
3. (Optional) Create a chain of GigaSMART operations containing the custom tunnel decap GSOP and a generic Header Stripping GSOP.
4. Bind the GigaSMART operation to the P interface that is associated with a network port as part of a map.

The encapsulated packet will go from the P interface that is associated with a network port to GigaSMART, where basic validation against configured values will be performed. The packet will then be sent to the chained GSOPs, where the encapsulated header will be stripped off (if configured to strip) and sent to the tools.

NOTE: The generic Header Stripping operation is performed on the inner payload of the tunneled packet. Use a hybrid port for the Header Stripping GSOP.

For an example configuration, refer to GigaSMART Custom Tunnel Decapsulation

NOTE: GigaSMART operations with a tunnel component can be assigned to GigaSMART groups consisting of multiple GigaSMART engine ports.

Custom Tunnel Decapsulation Configuration Example

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to Access GigaSMART from GigaVUE-FM for details.

To configure a custom tunnel termination, do the following:

Configure a GigaSMART group/associate group with GigaSMART Engine Port

1. From the device view, select **GigaSMART > GigaSMART Groups**.
2. Click **New**.
3. Type an **alias** in the Alias field. For example, gsg.
4. Select **Ports > Ports.**
5. Select an engine port. For example,**1/1/g13**.
6. Click **Edit.**
7. Select **Network** for Type and enable **Admin**.
8. Click **OK**.

Configure the IP Interface

1. Select **Ports > IP Interfaces**.
2. Click **New**.
3. In the **Alias** and **Description** fields, enter the name and description for the IP interface.
4. From the **Port** drop-down list, select port 1/1/g3
5. Enter **10.115.9.5** in the IP Address field.
6. Enter **255.255.255.255** in the Mask field
7. Enter **10.115.9.1** in the Gateway field.
8. Enter **1500** in the MTU field.
9. Select the **GigaSMARTGroup** you created in the first step of this process. For example: **gsg**.
10. Click **Save**.

Configure GigaSMART operation and assign to GigaSMART Group

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.
2. Click **New**.
3. Type an **alias** in the Alias field to identify this operation. For example, custom2
4. For **GigaSMART Groups**, select **gsg**.

5. For **GigaSMART Operations (GSOP),** select **TunnelDecapsulation**.
   a. Select **Custom**
   b. Type **Source Port**
   c. Type **Destination Port**
6. Click **Save**

Create a Map

1. Select **Maps > Maps > Maps.**
2. Click **New**.
3. Configure the map.
   - Type map11 in the Alias field.
   - Select T **or tool** for **Type**.
   - Select **By Rule** for Subtype.
   - Select the from **network port 1/1/g1** for the Source.
   - Select the virtual port vp1 for the Destination.
4. Add a Rule.
   a. Click **Add a Rule**.
   b. Select **Pass**.
   c. Select **IPv4 Version** and set Version to **v4**.
5. Click **Save**.

Display Custom Tunnel GSOP

To display the custom tunnel GigaSMART operation, do the following:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP)**.
2. Select the custom tunnel from the list. The GSOP displays.

Display Custom Tunnel Statistics

To display custom tunnel statistics, do the following:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP)> Statistics**.

2. Select **Statistics**. The Statistics page displays basic tunnel termination details in a table format.
3. Click the **GSOP alias** to display all the statistics available for this GS Operation.

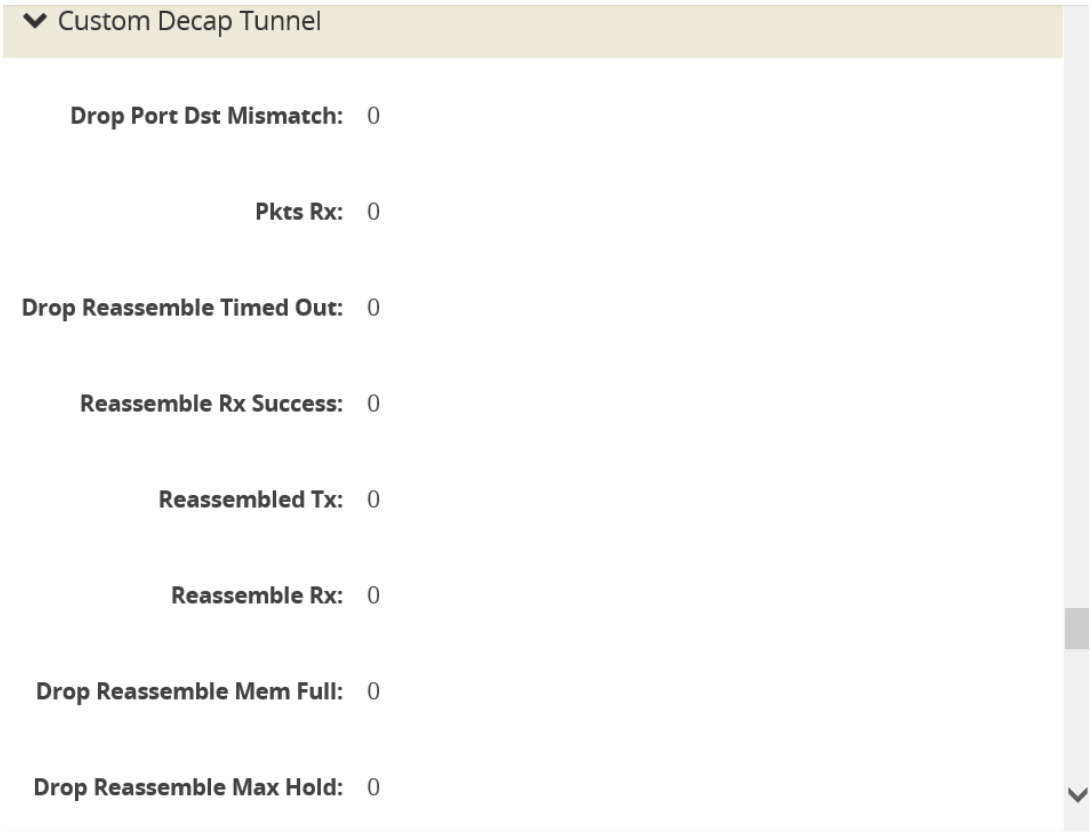


Figure 136

4. **Click** the **X** to close the **GS Operation Statistics** pane.

GigaSMART TCP tunnel

Required License- Advanced Tunneling license on GigaVUE-HC3. Referred to as “Tunneling license” on GigaVUE-HC1.

TCP tunnel feature routes the mirrored traffic from GigaVUE-VM to remote GigaVUE HC Series node reliably and without any reorder issues. TCP tunnel encapsulation is supported in the GigaVUE-VM node and the TCP tunnel decapsulation is supported in the GigaVUE HC Series node. Tunnel decapsulation can terminate more than one TCP connection initiated by the GigaVUE-VM node.

Configuration

The following are the steps to configure TCP tunnel between GigaVUE-VM and GigaVUE-FM,GigaVUE HC Series HC Series:

Steps	Refer to..
Configure GigaVUE-VM For Encapsulation	Configure TCP Tunnel in GigaVUE-VM User's Guide
Configure GigaVUE-FM,GigaVUE HC Series H-Series for Decapsulation	Configure GigaVUE-FM,GigaVUE HC Series HC Series for Decapsulation
Configure vMap for VMware	Configure vMap for VMware in GigaVUE-VM User's Guide

Supported Devices

TCP tunnel Decapsulation is supported only in GigaVUE-HC1, and GigaVUE-HC3 platforms.

Secure Tunnels

Secure Tunnels allow the traffic to be transported securely between a GigaVUE V Series to a GigaVUE HC Series or a GigaVUE HC Series to a GigaVUE HC Series device. The encapsulated traffic is sent to the GigaVUE HC Series device via secure tunnels in PCAPng format over a TLS connection.

Refer to the PCAPng Application section for more.

Secure Tunnels in GigaVUE HC Series has three primary use cases:

Secure Tunnels from GigaVUE V Series to GigaVUE HC Series

In this scenario, the encrypted traffic originating from a GigaVUE V Series node is transmitted to a remote GigaVUE HC Series device. This device then strips off the encryption, processes the traffic by utilizing GigaSMART applications (optional) and then forwards it to the tool. The GigaVUE HC Series can have a single map, multiple map rules, or multiple tools, as shown below.

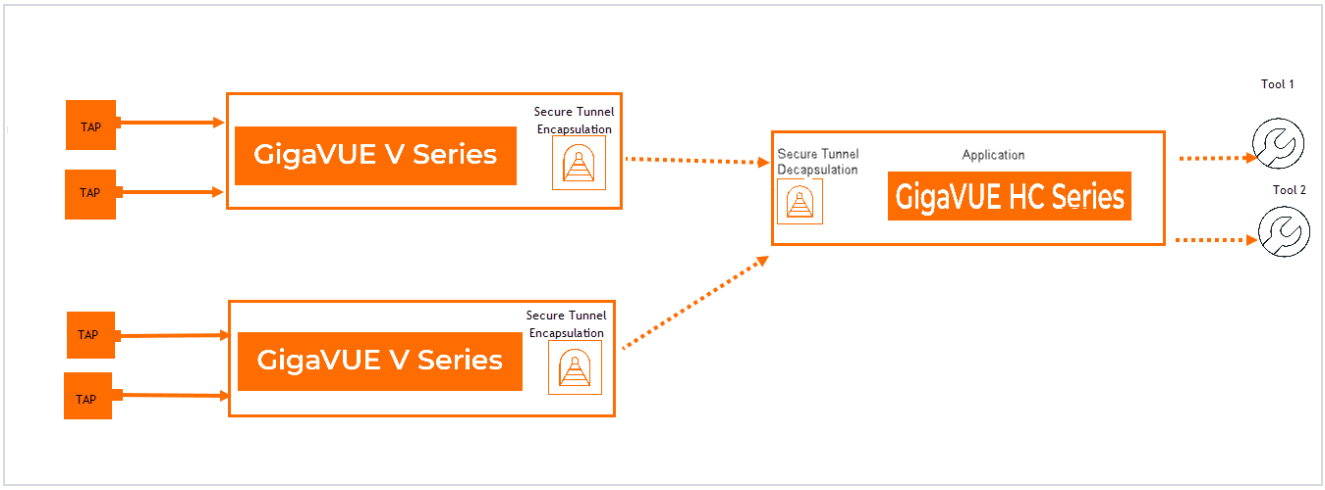


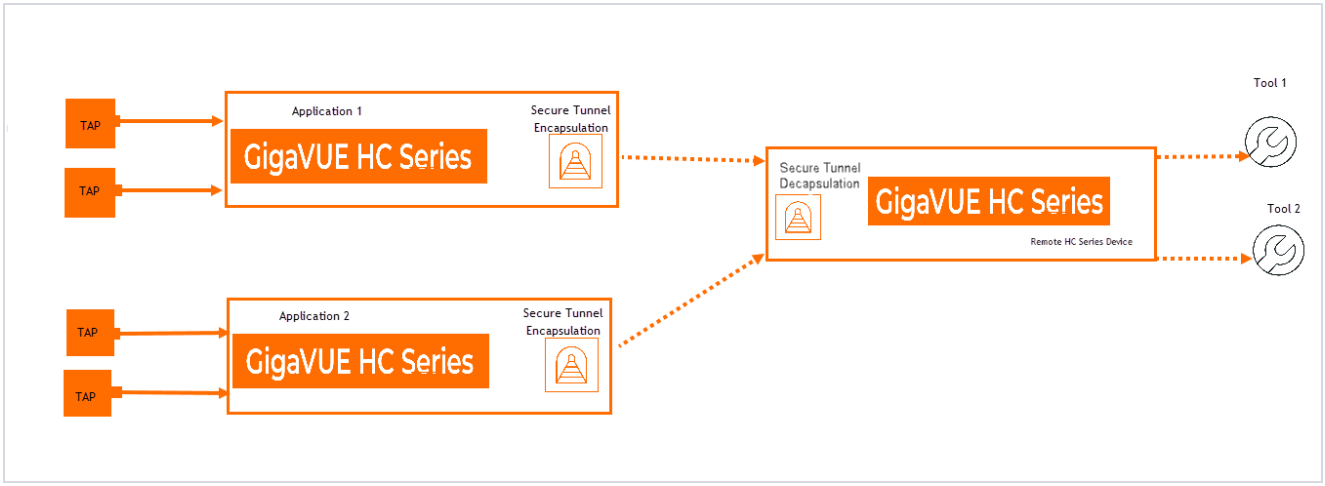
Figure 137 *GigaVUE HC Series with a single IP Interface and multiple Map rules*

NOTE: Utilizing GigaSMART applications on the traffic is optional.

Secure Tunnels between the GigaVUE HC Series and a remote GigaVUE HC Series

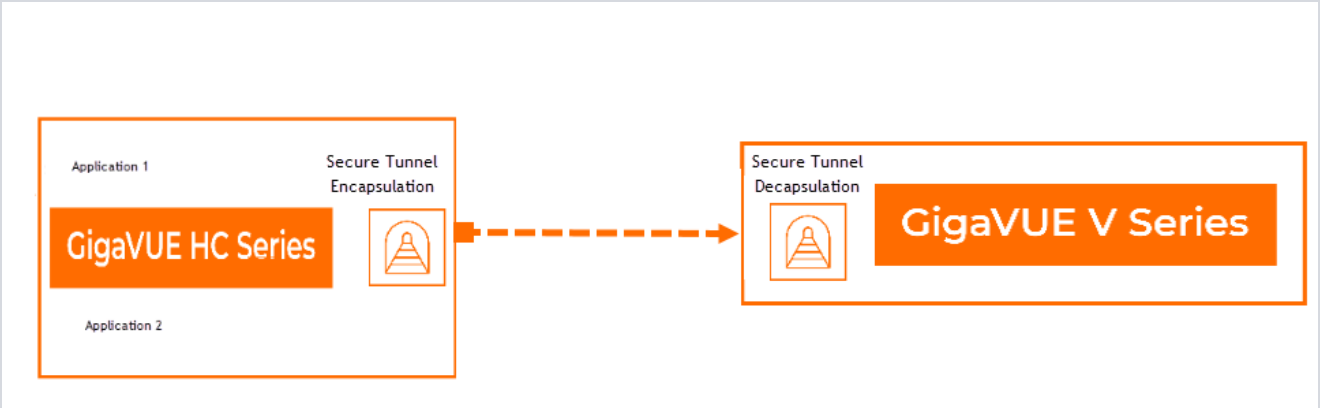
In this scenario, encrypted traffic is sent from the GigaVUE HC Series devices to a remote GigaVUE HC Series via secure tunnels, where it is decrypted, and forwarded to the tool.

In locally available GigaVUE HC Series devices you can utilize physical connectivity to securely send the traffic.



Secure Tunnels between GigaVUE HC Series to a GigaVUE V Series.

In this scenario , the encapsulated traffic from the GigaVUE HC Series is transported to GigaVUE V Series node. The GigaVUE V Series node decapsulates and processes the packet per the configuration.



Supported Platforms

Secure Tunnels is supported in the following platforms:

- GigaVUE-HC3 Gen3 (SMT-HC30C08)
- GigaVUE-HC1 Gen3 (SMT-HC1-S)
- GigaVUE-HCT (SMT-HC1-S)
- GigaVUE-HC1-Plus (onboard Gen3 GigaSMART engine)

Configure Secure Tunnels

To configure Secure Tunnels, follow the below steps:

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

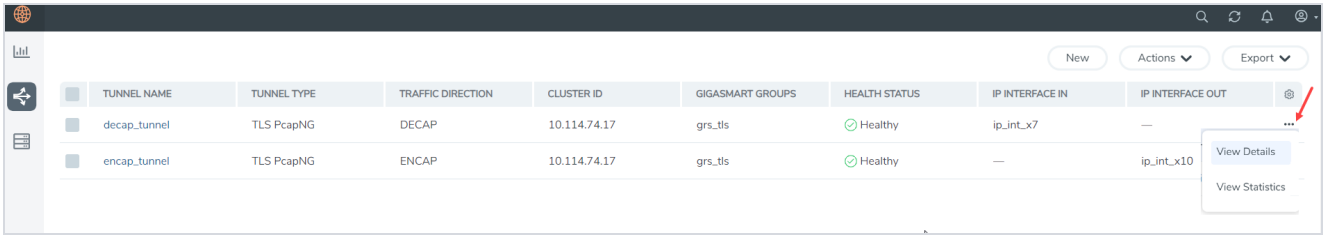
Field	Description
Tunnel Name	The name of the tunnel. **NOTE:** Alias must not have spaces or any of these characters: *'"?;:,/%@.
Tunnel Description	The description of the tunnel endpoint.
Tunnel Type	The tunnel type TLS-PCAPNG is selected by default.
Traffic Direction	Select the traffic direction of the tunnel. Choose DECAP for decapsulation or choose ENCAP for encapsulation. • If you choose Decap, select the IP Interface from the list. You can create one using the Create an IP interface option if you do not have an IP interface. Refer to IP Interfaces to learn more about creating an IP Interface. • If you choose Encap enter the Source Port. You can select from the drop-down list. To edit the Port type or to enable the port for Admin privileges use the Port Editor window.

Field	Description
	NOTE: In the scenario where secure tunnels needs to be established between GigaVUE V Series to a GigaVUE HC Series , you can utilize the **Configure Physical Tunnel** option provided at the GigaVUE V Series Secure Tunnel page. This allows you to conveniently configure secure tunnels at your physical device . Refer to Create Ingress and Egress Tunnels section in the respective cloud guides.
Nodes	Select the Nodes on which you intend to create the tunnel.
GigaSMART Group	Select the GigaSMART groups created. If you do not have any groups, you can configure them using the **Create GS Groups** option. Refer Create GigaSMART Operations – A Summary to know more about configure groups. **NOTE:** Secure tunnels can be configured only for Gen 3 nodes.
TCP Profile	Configure the TCP profile parameters such as: • TCP Keep Alive Timer - Enter the time duration between keep alive messages. The value ranges from 30-7200. The default value is 60. • Selective Acknowledgment- Choose Enable to turn on the TCP selective acknowledgments. • SYN Retries- Enter the value for number of times the SYN has to be tried. The value ranges from 1 to 6. The default value is 3.
SSL Profile	Enter the following SSL Profile parameters: • MTLS- MTLS is disabled by default. • Cipher - Only SHA 256 is supported. • TLS Version- Only TLS Version1.3 is supported. • Key Alias- Select the key to be used for decapsulation. You can choose a key from Node or GigaVUE-FM. o Add Keychain Password o Add Key Alias • Trust Store - Select the Trust Store key or click Append to add a new key for encapsulation.
Listener Port	For a decapsulation , specify the Listener port range. The range value is between 1- 30000. **NOTE:** If you update the Listener Ports in an existing secure tunnel destination, the tunnel will be re-established by deleting and recreating the maps.
Destination Node	Select the destination node for the decapsulation tunnel.
Destination Port	Select the destination port for the decapsulation tunnel. You can edit the port details from the Port editor screen.
DSCP	Enter the Differentiated Services Code Point (DSCP) value for an encapsulation.
Time to Live	For encapsulation tunnel, enter the value of the time interval till which the session needs to be available. The value ranges from 1 to 255. The default value is 64.

Field	Description
	NOTE: This is ignored for IPv6 based tunnel.
Encapsulation IP interface	Enter the interface IP address of the node (Destination IP) for an encapsulation tunnel.
Remote Tunnel IP, L4 Port & Source Port	Enter the IP address , L4 Port and Source Port values for an encapsulation tunnel. The Source Port range is from 30001 - 65535. **NOTE:** For decapsulation tunnel, It is recommended to match your L4 destination port rule to your listener port range. **NOTE:** The 50 ports from the configured exporter Source Port will be reserved for multi-engine GigaSMART group.
Rules	Configure the map rules that needs to be adhered to while decapsulating or encapsulating the traffic. Click on **Rules Editor** to configure the rules. For detailed information about map rules, refer to Map Rules. You can configure Inner Header qualifiers and MPLS Header qualifiers for GIgaVUE-TA400 device. Refer to Inner Header and MPLS Header Filtering.
Additional GigaSMART operations	Enable this option to add GigaSMART operations. When enabled you will be provided with GigaSMART Operations section which would allow you to choose among the following: • Add Header • De-duplication • Load Balancing • Masking • Slicing

- Click on **Save.**

The configured Secure tunnels will be displayed. To get a quick view of the secure tunnel configuration go to ••• **>View Details** or click on the tunnel name. To know if your Secure tunnel is healthy refer the **Health Status** column.



NOTE: The Tunnels page accessible here at **Traffic>Physical> Orchestrated Flows> Tunnels** is solely for configuring Secure tunnels in GigaVUE HC Series and is an Early Access page. To learn more about configuring other Tunneling operations, refer to the Tunneling Operations section.

Edit Secure Tunnel

To edit a Secure Tunnel follow the below steps:

- Select the secure tunnel from the tunnels page.
- Click on **Actions > Edit** .
- Edit your tunnel parameters.
- Click **Save**.

Delete Secure Tunnel

Select the secure tunnel from the tunnels page.

- Click on **Actions> Delete**.
- Delete your tunnel.

Export Secure Tunnels

To export the configured Secure Tunnels, click **Export**. The secure tunnels is downloaded in either CSV or XLSX format.

View Tunnel Statistics

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

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

Limitations

The following limitations apply when configuring Secure Tunnel in GigaVUE HC Series.

- Secure-tunnel is not supported with the following applications:
  - GTP Correlation
  - Passive SSL
  - Inline SSL
  - Internet Content Adaptation Protocol (ICAP)
- GigaSMART operations with Secure Tunnel configuration is supported only with a regular map.
- The Tunnel Decapsulation listener port cannot be same as the Tunnel Encapsulation Source Port.

Additional Sources of Information

This appendix provides additional sources of information. Refer to the following sections for details:

- Documentation
- Documentation Feedback
- Contact Technical Support
- Contact Sales
- The VÜE Community

Documentation

©This table lists all the guides provided for GigaVUE-FM,GigaVUE HC Series 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-FM,GigaVUE HC Series 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-FM,GigaVUE HC Series devices; reference information and specifications for the respective GigaVUE-FM,GigaVUE HC Series 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-TA200 Hardware Installation Guide

GigaVUE-FM,GigaVUE HC Series 6.14 Hardware and Software Guides
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,GigaVUE HC Series 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 HC Series
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 Cloud Suite Deployment Guide - Third Party Orchestration

GigaVUE-FM,GigaVUE HC Series 6.14 Hardware and Software Guides
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,GigaVUE HC Series 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 HC Series, 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,GigaVUE HC Series 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,GigaVUE HC Series" and Release: "5.6," enter "pdf" in the search box, and then click **GO** to view all PDF documentation for GigaVUE-FM,GigaVUE HC Series 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)