GigaVUE-FM User's Guide

Version 5.6.00

Document Version: 3.0 (*Change Notes*)

DOCUMENT REVISION – 5/31/19

Change Notes

When a document is updated, the document revision number on the cover page will indicate a new revision number, the Document Revision date is updated on the title page, and this table will describe what changed.

Rev	Date	Change
rev 1	03/29/2019	Original release of document with the 5.6.00 release.
rev 2	04/13/2019	As part of effort to streamline the documentation set, the GigaVUE-OS CLI User's Guide has been transformed into the GigaVUE-OS CLI Reference Guide. Detailed feature descriptions that were previously provided in the GigaVUE-OS CLI User's Guide are now incorporated into the GigaVUE-FM User's Guide. Updated the following sections: • *About Cluster* on page 193 • *Cluster Node Limit* on page 194 • *Separate Paths for Cluster Control and Stack Traffic* on page 197 • *About Cluster Roles* on page 199 • *Creating Clusters: A Roadmap* on page 204 • *Cluster Rules and Recommendations* on page 205 • *GigaVUE-TA Series and GigaVUE-HC3 Clustering Recommendations* on page 206 • *Designated Port Types* on page 487 • *Backwards Compatibility* on page 495 • *GigaSMART Perpetual Licenses* on page 743 • *Licensing GigaSMART Applications* on page 751 • *Limitations of Gigamon Discovery* on page 302
rev 3	05/30/2019	• Fixed cross-references and streamlined instructions throughout the guide. • Updated the GTP correlation diagram under *Configure GTP Correlation Examples* on page 897. • Consolidated the Inline SSL Decryption section into a chapter in the Traffic section: *Work With Inline SSL Decryption* on page 645.

Contents

1 About this Guide

This guide describes how to install, deploy, and operate the GigaVUE® Fabric Manager (GigaVUE-FM and GigaVUE-VM) fabric management system from Gigamon® Inc.

Use the chapters in *Part 2: Installation and Upgrade* on page 63 of this document to get GigaVUE-FM installed and running in your network environment. Then, turn to later chapters for information on using product features.

Part 1: Getting Started

This section provides an introduction to GigaVUE-FM, describing the new features in the current release, and the licensing options for GigaVUE-FM. The topics covered are:

2 About GigaVUE-FM

This section introduces the GigaVUE-FM, fabric management, and software, describing their features and functions and summarizing the relationships between the products. This chapter includes the following major sections:

GigaVUE-FM Overview

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

- **Physical GigaVUE Nodes** – GigaVUE-FM manages GigaVUE G Series, TA Series, and H Series nodes, allowing for a unified workspace, while also providing an easy-to-use wizard-based approach for configuring Flow Mapping® and GigaSMART® traffic policies. For a list of GigaVUE G Series, TA Series, and H Series nodes supported for management in this release, refer to *Supported GigaVUE G Series, TA Series, and H Series Nodes*.
- **Virtual GigaVUE Nodes** – GigaVUE-FM also extends the GigaVUE feature set into the virtual space by allowing for the discovery, configuration, and management of the new GigaVUE-VM virtual traffic visibility node. GigaVUE-VM provides powerful Flow Mapping technology for the traffic flowing between virtual machines, allowing you to distribute cloud-based traffic to physical tool ports in the visibility fabric.

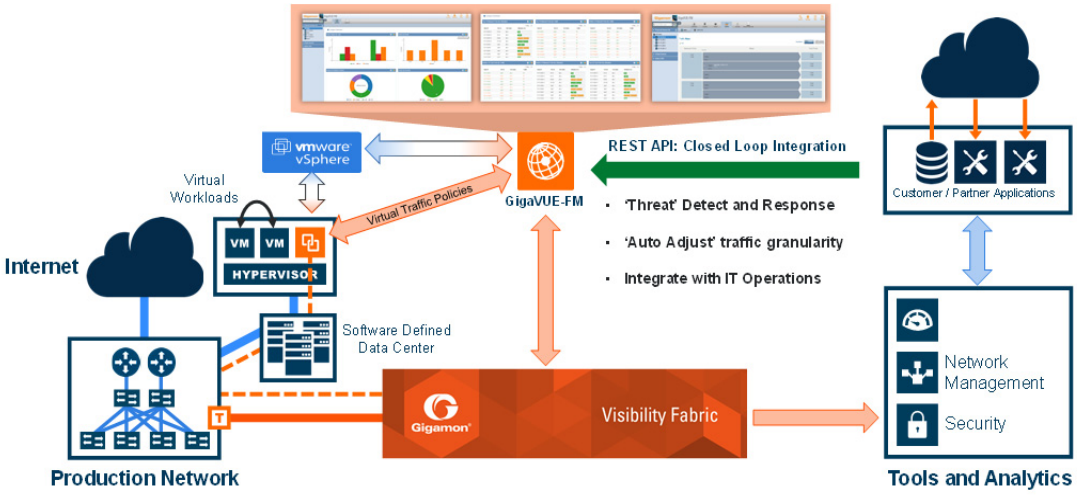

Figure 2-1: GigaVUE-FM Overview

GigaVUE-FM Features and Benefits

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

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

Table 2-1: Features and Benefits of GigaVUE-FM

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

Supported GigaVUE G Series, TA Series, and H Series Nodes

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

Device Management

During the initial start up of GigaVUE-FM, it performs a *Discover* on GigaVUE G Series, TA Series, and H Series nodes listed in Table 2-2.

Table 2-2: GigaVUE-FM Managed GigaVUE Nodes and Software Versions

GigaVUE Series	Node	Supported Releases for GigaVUE-FM Management
GigaVUE G Series Nodes	GigaVUE-2404	**v8.6.10 v8.6.11**
	GigaVUE-420	
	GigaVUE-212	
	G-SECURE-0216	N/A
GigaVUE H Series Nodes	GigaVUE-HD4 or GigaVUE-HD8 with **HD CCv2 Control Card**	**v4.7.xx, v4.8.xx, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx**
	GigaVUE-HC1	**v4.8.xx, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**
	GigaVUE-HC2	**v4.7.xx, v4.8.xx, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**
	GigaVUE-HC2 with **HC CCv2 Control Card**	**v4.7.xx, v4.8.xx, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**
	GigaVUE-HC3	**v5.0.00, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**
	GigaVUE-HB1	**v4.7.xx, v4.8.xx, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx**
GigaVUE TA Series Nodes	GigaVUE-TA1	**v4.7.xx, v4.8.xx, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx**
	GigaVUE-TA10	**v4.7.xx, v4.8.xx, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**
	GigaVUE-TA40	**v4.7.xx, v4.8.xx, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**
	GigaVUE-TA100	**v4.7.xx, v4.8.xx, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**
	GigaVUE-TA100-CXP	**v4.8.01, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**
	GigaVUE-TA200	**v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**
	Certified Traffic Aggregation White Box with GigaVUE-OS	**v4.7.xx, v4.8.xx, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**

NOTE: Although GigaVUE-FM may recognize earlier versions of GigaVUE G Series, GigaVUE-TA1 and H Series nodes, the versions listed in Table 2-2 are the officially supported versions. Earlier versions are not managed by GigaVUE-FM.

Configuration Management

GigaVUE-FM allows you to perform configuration tasks on GigaVUE H Series and TA Series nodes only. It supports the versions listed in Table 2-3.

Table 2-3: GigaVUE Nodes and Software Versions Configurable by GigaVUE-FM

GigaVUE Series	Node	Supported Releases for GigaVUE-FM Management
GigaVUE H Series Nodes	GigaVUE-HD4 or GigaVUE-HD8 with **HD CCv2 Control Card**	**v4.7.xx, v4.8.xx, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx**
	GigaVUE-HC1	**v4.8.xx, v5.0.xx. v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**
	GigaVUE-HC3	**v5.0.00, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**
	GigaVUE-HC2	**v4.7.xx, v4.8.xx, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**
	GigaVUE-HC2 with **HC CCv2 Control Card**	**v4.7.xx, v4.8.xx, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**
	GigaVUE-HB1	**v4.7.xx, v4.8.xx, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx**
GigaVUE TA Series Nodes	GigaVUE-TA1	**v4.7.xx, v4.8.xx, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx**
	GigaVUE-TA10	**v4.7.xx, v4.8.xx, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**
	GigaVUE-TA40	**v4.7.xx, 4.8.xx, v5.0.xx, v5, 5.6.xx.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**
	GigaVUE-TA100	**v4.8.xx, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**
	GigaVUE-TA100-CXP	**v4.8.01, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**
	GigaVUE-TA200	**v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**
	Certified Traffic Aggregation White Box with GigaVUE-OS	**v4.7.xx, v4.8.xx, v5.0.xx, v5.1.xx, v5.2.xx, v5.3.xx, v5.4.xx, v5.5.xx, 5.6.xx**

How to enable Web Server for Node Management

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

3 About GigaVUE-FM Licenses

This section describes how to obtain and apply licenses for GigaVUE-FM. It consists of the following main sections:

- *GigaVUE-FM License* on page 31 describes the licenses available and how to obtain and apply them.
- *About GigaVUE-FM License Types* on page 32 lists the available licenses and features available with each license type.
- *Apply Licenses* on page 33 describes the process to apply the licenses.
- *Upgrade and Downgrade License Packages* on page 35 covers the best practices when upgrading or downgrading license packages.

NOTE: For information about GigaVUE-VM licensing, refer to the *GigaVUE-VM User's Guide*.

GigaVUE-FM License

GigaVUE-FM is provisioned by default with a Base License that lets you add one physical node and one virtual node. To manage additional physical or virtual nodes, you must obtain and apply licenses, as described in this section.

NOTE: To run only GigaVUE-VM, there is no requirement to purchase additional licenses for GigaVUE-FM. For information about GigaVUE-VM licensing, refer to the *GigaVUE-VM User's Guide*.

Obtain New License

Contact your Sales representative to obtain a new license for GigaVUE-FM Nodes (see *Contacting Sales* for the contact information).

Retrieve Lost License

If you lost an existing license, contact Gigamon Technical Support for assistance. For the contact information, refer to *Contacting Technical Support*.

About GigaVUE-FM License Types

GigaVUE-FM are available in multiple tiered options along with optional Add-On Features which are also available as a special license (add-on are included with the Prime Package as free-of-charge). All GigaVUE-FM are available with base option and with base feature of 1 free physical node and 1 free virtual node and 10 virtual tap points for OpenStack, AWS and Azure. No licenses are required to activate this option.

NOTE: For information about GigaVUE-VM licensing, refer to the *GigaVUE-VM User's Guide*.

Additional GigaVUE-FM licenses are available for purchase. The following tables summarizes the available packages and support features with each package.

Table 3-1: GigaVUE-FM Evaluation License Packages

License Types	Physical Nodes	Virtual Nodes	OpenStack/ AWS/Azure	Features available	Notes
GigaVUE-FM Evaluation	Up to 200	1 (included as Base)	10 Virtual TAP Points	All features available with Prime for the evaluation period.	License automatically expires after 45 days.

NOTE: Evaluation licenses are not recommended for deployment in production environment. At the end of the evaluation period, if the license is not upgraded to a fully licensed version, the features are disabled automatically. For an evaluation license, contact your Gigamon representative.

GigaVUE-FM License Packages

The following table summarizes the GigaVUE-FM License packages.

Table 3-2: GigaVUE-FM License Packages

Features	Base (Free-of-charge)	5-Pack	10-Pack	Prime
Physical Node Count	1	Up to 5	Up to 10	Up to 200
Rest API	Yes	Yes	Yes	Yes
Audit, Events Logs	Yes	Yes	Yes	Yes
Firmware Upgrade	Yes	Yes	Yes	Yes
Configuration Backup	Yes	Yes	Yes	Yes

Table 3-2: GigaVUE-FM License Packages

Features	Base (Free-of-charge)	5-Pack	10-Pack	Prime
Dashboard	Only the following Static Widgets are displayed: • Top 10 Network Ports by Traffic • Top 10 Tool Ports by Traffic • Top 10 Maps by Traffic • Audit Logs By Result • Events By Severity • Ports Link Status Summary • Unhealthy Maps Status Summary • Unhealthy Flows	Customizable	Customizable	Customizable
Reports	No	Yes	Yes	Yes
Trending Data	1 Day	1 Week	1 Month	1 Month
Add-On Features: **FabricVUE® Traffic Analyzer**	No	Available for Purchase	Available for Purchase	Included

There are also upgrade packages for GigaVUE-FM available for customers that have already purchased GigaVUE-FM. The packages allow users to upgrade from a 5-pack option to a 10-pack or a Prime package. There is also an option to upgrade from a 10-pack to a Prime package. To find out more about the upgrade purchase, contact your Gigamon Sales representative.

For upgrade option, your GigaVUE-FM information should match what is in the record for MAC address and customer information.

Apply Licenses

Use the following procedure to license your products on the **License** page.

To obtain and apply the GigaVUE-FM license:

1. Locate the email sent to you by Gigamon containing the licensing information for your installation. This email contains one or more **GIK** (**G**igamon **I**nstallation **K**eys) values. You will use these GIKs to generate License Keys on the Gigamon Licensing Website.

2. Locate the MAC address of the virtual network adapter associated with the GigaVUE-FM installation. To locate the address, click **Administration** on the top navigation link. Refer to Figure 3-1 on page 34.

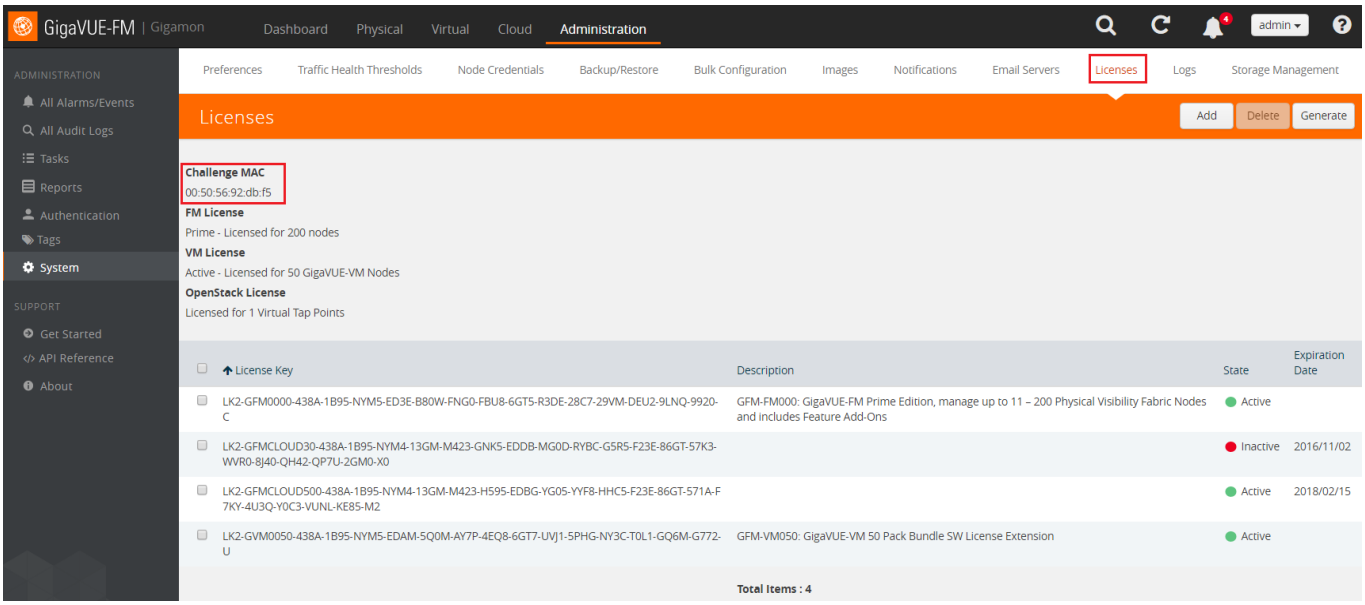

Figure 3-1: Locate the MAC Address

3. In the left navigation pane, select **System**.
4. Click **Licenses**. The Licenses page is displayed as shown in Figure 3-1 on page 34.

   Refer to the Challenge MAC displayed in the Licenses page. The license is only valid with the corresponding MAC address. If GigaVUE-FM is deleted or re-installed, contact Gigamon Support.

5. GigaVUE-FM licenses can be generated by clicking on **Generate** from the Licenses tab or by using the Gigamon Licensing Portal from the following location:

   https://licensing.gigamon.com

6. Enter the MAC address and GIKs of the purchased licenses in the portal. Multiple GIKs can be entered by clicking the **+** button. Once all the information is entered and submitted, the license key(s) are displayed on the screen.
7. Write down the license key or keys.
8. Login to GigaVUE-FM as an administrator and navigate to the **Administration > System > License** page. Refer to Figure 3-1 on page 34.

9. Click the **+** (Add) button and enter the license key or keys provided by Gigamon into the dialog box that appears, and then click **Save** to apply the license.

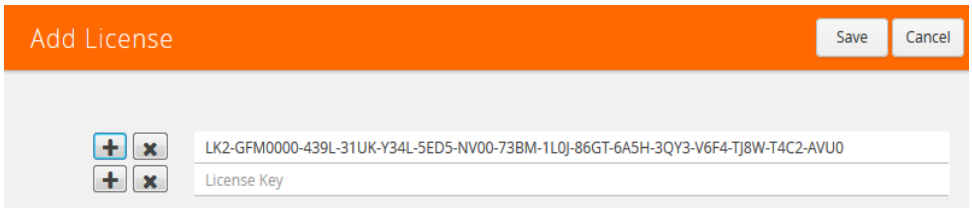

Figure 3-2: Add License Page

Upgrade and Downgrade License Packages

- Upgrading of license packages is available at all times.
- To purchase a new license, please contact Gigamon Sales Representative.
- All licenses are perpetual therefore they carry in to any software upgrades without re-applying the licenses, except evaluation licenses are set for expiration. Software upgrades can be managed during valid evaluation period.
- Licenses can be upgraded from Base to either a 45 day evaluation or to paid version.
- If an evaluation license is upgraded to Express or Advanced version, the Add-on features are automatically disabled. To retain the Add-On features, please purchase the license for Add-On features such as FabricVUE Traffic Analyzer, or upgrade to the Prime Package.
- Purchased licenses cannot be downgraded.
- Licenses can be deleted and re-entered as long as the MAC address tied to the license is still valid.

In case of expiration of the evaluation license, GigaVUE-FM will revert back to supporting only 1 physical and 1 virtual node.

The list below shows the node priorities in case a License is invalid and nodes are deactivated. In such a case, the nodes will be visible but deactivated. They can be re-activated if the license is reinstalled. This is especially important if the evaluation license expires and you need extra time to enter a valid license.

(1) If a cluster exists:

- Master node

  In case of multiple clusters, the cluster with the top level priorities as shown in standalone will take over. For example, a cluster with master as GigaVUE-2404 will have preference over a cluster with master as GigaVUE-HC2.
- Standalone node *(Based on the node priority levels as shown below)*
- Standby Master, in case the master is removed
- Stack/Cluster member
- Unreachable
- Unknown

(2) If there is no cluster *(G Series nodes have top preference)*

- GigaVUE-212
- GigaVUE-420
- GigaVUE-2404
- GigaVUE-0216
- GigaVUE-TA100-CXP
- GigaVUE-HC3
- GigaVUE-HC2
- GigaVUE-HC1
- GigaVUE-TA10
- GigaVUE-TA40
- GigaVUE-TA100
- GigaVUE-TA10A
- White box with GigaVUE-OS

In cases where nodes have connectivity issues as listed below, the next level nodes as shown in Scenarios (1) and (2) will take effect:

- Nodes should not have any connectivity issue
- Nodes can be reached but node response has parsing problem
- Nodes can be reached but authentication is invalid
- Node cannot be reached
- Nodes cannot be added

4 Get Started with GigaVUE-FM

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

It includes the following major sections:

Log In to GigaVUE-FM

The GigaVUE-FM login page provides information about the security policy login banner beside the username and password fields. The login banner is customizable. For more information about configuring a custom banner, refer to *Configure a Custom Banner* on page 44.

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

Log Out of GigaVUE-FM

To logout of GigaVUE-FM, click the **Admin** drop-down list on the top right of GigaVUE-FM and select **Logout**.

GigaVUE-FM Navigation

Starting in software version 5.1, the navigation in GigaVUE-FM has been changed to improve the usability and customer experience.

Dashboard

When you first login to GigaVUE-FM, the Dashboard - Physical & Virtual page is displayed by default as shown in Figure 4-1 on page 39. You can navigate to Health Monitor or Traffic Analyzer Dashboards from the left navigation pane.

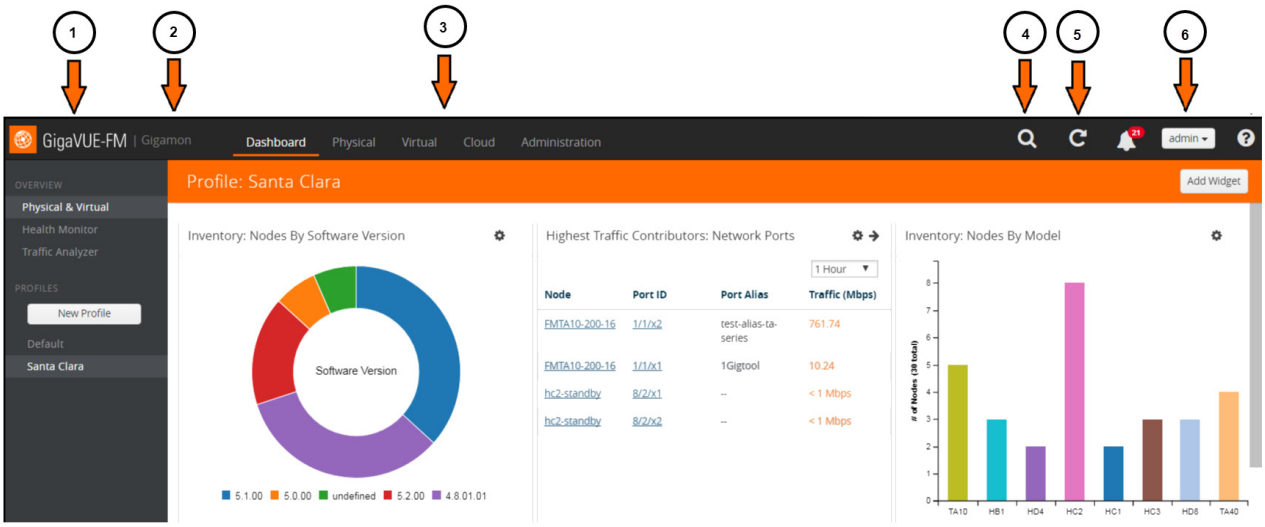

Figure 4-1: GigaVUE-FM Home Page.

1 - Click on GigaVUE-FM to return to the Dashboard
2 - Customized GigaVUE-FM instance title
3- Top navigation links
4 - Search
5 - Refresh
6 - Click to upgrade, change the password and logout

Table 4-1 provides descriptions of the top navigation links and the left navigation pane in GigaVUE-FM.

Table 4-1: Summary of Dashboard Navigation

Top Navigation Links	Left Navigation Pane
Dashboard	
	Physical & Virtual —Displays the informational widgets for the physical and Virtual nodes.
	Health Monitor—Displays information about GigaVUE-FM such as CPU usage, amount of storage used and available.
	Traffic Analyzer—Displays the FabricVUE Traffic Analyzer, which provides traffic flow patterns based on FabricVUE Traffic Analyzer. This dashboard displays information only when NetFlow is enabled on the port or port under observation.
	Profiles—Allows you to create and view the profiles. A profile allows you to create a customized dashboard to monitor the physical and virtual nodes.

Physical

Click Physical to view the physical nodes and clusters managed by GigaVUE-FM. Under Physical, you can add and delete physical nodes, create and manage clusters, perform image upgrades, take backups, and restore and reboot images.

Table 4-2 provides descriptions of the Physical page in GigaVUE-FM:

Table 4-2: Summary of Physical

Top Navigation Links	Left Navigation Pane
Physical	
	Sites—Lets you choose all sites or a site value. Sites are associated to only clusters. The All Sites drop-down list lets you choose a site value and view the related topology, alarms and events, and audit logs of the clusters associated to the site.
	Physical Nodes—Provides information about the nodes and clusters managed by GigaVUE-FM.
	Topology—Provides graphical view of the nodes and clusters currently managed by GigaVUE-FM. It also displays the nodes and clusters that are discoverable. For example, the neighbor nodes can be auto-discovered using the Link Layer Discovery Protocol (LLDP), Cisco Discovery Protocol (CDP) and Gigamon Discovery Protocol (GDP).
	Alarms/Events—Displays all notifiable alarms and events pertaining to physical nodes. For example, you can view user authentication failure, high CPU utilization, node discovery, device configuration state changes, and many others. You can filter and manage the Alarm/Event logs from this page.
	Audit Logs—Displays the tasks pertaining to physical nodes managed by GigaVUE-FM. You can filter and manage the audit logs from this page.

Virtual

Click **Virtual** to view the virtual nodes managed by GigaVUE-FM. Under Virtual, you can configure GigaVUE-VMs in different virtual environments such as VMware ESXi and VMware NSX.

Table 4-3 provides descriptions of the Virtual page in GigaVUE-FM:

Table 4-3: Summary of Virtual

Top Navigation Links	Left Navigation Pane
Virtual	
	VMware vCenter
	Virtual Nodes—Displays information about all the deployed GigaVUE-VMs and their status.
	Virtual Maps—Displays information about configuring the virtual maps on the virtual nodes for VMware vCenter.
	Management—Allows you to connect to virtual center, deploy GigaVUE-VMs to multiple ESXi hosts, and configure tunnel endpoints.
	VMware vCenter/NSX-V
	Virtual Nodes—Displays information about all the deployed GigaVUE-VMs in the NSX environment and their statuses.
	Virtual Maps—Displays information about configuring the virtual maps on the virtual nodes for VMware NSX.
	Management—Allows you to connect to virtual center, deploy GigaVUE-VMs to multiple ESXi hosts, and configure tunnel endpoints.
	All VIRTUAL
	Alarms/Events—Displays all notifiable alarms and events pertaining to physical nodes. For example, you can view user authentication failure, high CPU utilization, node discovery, device configuration state changes, and many others. You can filter and manage the Alarm/Event logs from this page.
	Audit Logs—Displays the tasks pertaining to virtual nodes managed by GigaVUE-FM. You can filter and manage the audit logs from this page.

Cloud

Click **Cloud** to configure the components of the Gigamon Visibility Platform for Amazon Web Services (AWS), Azure and OpenStack.

Table 4-4 provides descriptions of the Virtual page in GigaVUE-FM:

Table 4-4: Summary of Cloud

Top Navigation Links	Left Navigation Pane
Cloud	
	AWS **NOTE:** For documentation, refer to *GigaSECURE® Cloud for AWS Configuration Guide*, available in the Customer Portal.
	Monitoring Session—Displays information about configuring flow maps to filter and send traffic from GigaVUE V Series nodes to the monitoring tools.
	Topology—Displays the graphical view of VPCs, subnets and instances for the selected monitoring session.
	Visibility Fabric—Displays the status of G-vTAP Controllers, V Series Controllers, V Series Nodes, and Tunnel Endpoints that are currently configured in the monitoring session.
	Configuration—Allows you to configure G-vTAP Controllers, V Series Controllers, V Series Nodes, and Tunnel Endpoints.
	Azure **NOTE:** For documentation, refer to *GigaSECURE® Cloud for Azure Configuration Guide*, available in the Customer Portal.
	Monitoring Session—Displays information about configuring flow maps to filter and send traffic from GigaVUE V Series nodes to the monitoring tools.
	Topology—Displays the graphical view of VNets, subnets and instances for the selected monitoring session.
	Visibility Fabric—Displays the status of G-vTAP Controllers, V Series Controllers, V Series Nodes, and Tunnel Endpoints that are currently configured in the monitoring session.
	Configuration—Allows you to configure G-vTAP Controllers, V Series Controllers, V Series Nodes, and Tunnel Endpoints.
	OpenStack **NOTE:** For documentation, refer to *GigaSECURE® Cloud for OpenStack Configuration Guide*, available in the Customer Portal.
	Monitoring Session—Displays information about configuring the OpenStack monitoring sessions.
	Topology—Displays the graphical view of virtual nodes and clusters currently managed by GigaVUE-FM.
	Visibility Fabric—Displays information about the G-vTAP Controllers and GigaVUE-VMs that are currently configured.

Table 4-4: Summary of Cloud

Top Navigation Links	Left Navigation Pane
	Configuration—Displays information about configuring G-vTAP Controllers, V Series Controllers, V Series Nodes, and Tunnel Library.
	CLOUD
	Alarms/Events—Displays all notifiable alarms and events pertaining to AWS and OpenStack. For example, you can view AWS connection status changed, AWS license expiry, and others.
	Audit Logs—Displays the tasks pertaining to AWS and OpenStack. You can filter and manage the audit logs from this page.

Administration

Click **Administration** on the top navigation link.

Table 4-4 provides descriptions of the Administration page in GigaVUE-FM:

Table 4-5: Summary of Administration

Top Navigation Links	Left Navigation Pane
Administration	
	All Alarms/Events—Displays the alarms and events for Physical, Virtual, and Cloud.
	All Audit Logs—Displays the audit logs for Physical, Virtual, and Cloud. You can easily filter and manage the audit logs.
	Tasks—Displays the admin and scheduled tasks. These pages are used to add and manage scheduled events such as image upgrades.
	Reports—Allows you to generate visibility fabric performance reports, visibility fabric node details report, GigaVUE-VM report, visibility fabric inventory report, and so on using various templates. The reports can be downloaded in HTML or PDF format.
	Authentication—Allows you to configure authentication and authorization settings for GigaVUE-FM.
	Tags—Allows you to create a site or a user-defined tag that can be associated to clusters, ports, port groups, and GigaSMART groups.
	System—Allows you to configure licenses, set up preferences, node credentials, backup and image file storage location, notifications, email servers, storage management, and so on.

Configure a Custom Banner

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

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

To configure the custom banner:

1. Click **Administration** on the top navigation link and select **Systems** > **Preferences**.
2. Click **Edit** and the Edit Preferences page appears.
3. Enter the custom banner message in the **Login Banner** text box.
4. Click **Save**.

Quick Views

A quick view provides easy access to Visibility Fabric details such as nodes, ports, and traffic policies. In GigaVUE-FM, you can click on items such as port ID, map alias, port error counts, and so on, and get detailed information about the selected item. Figure 4-2 shows the details displayed on a Port quick view.

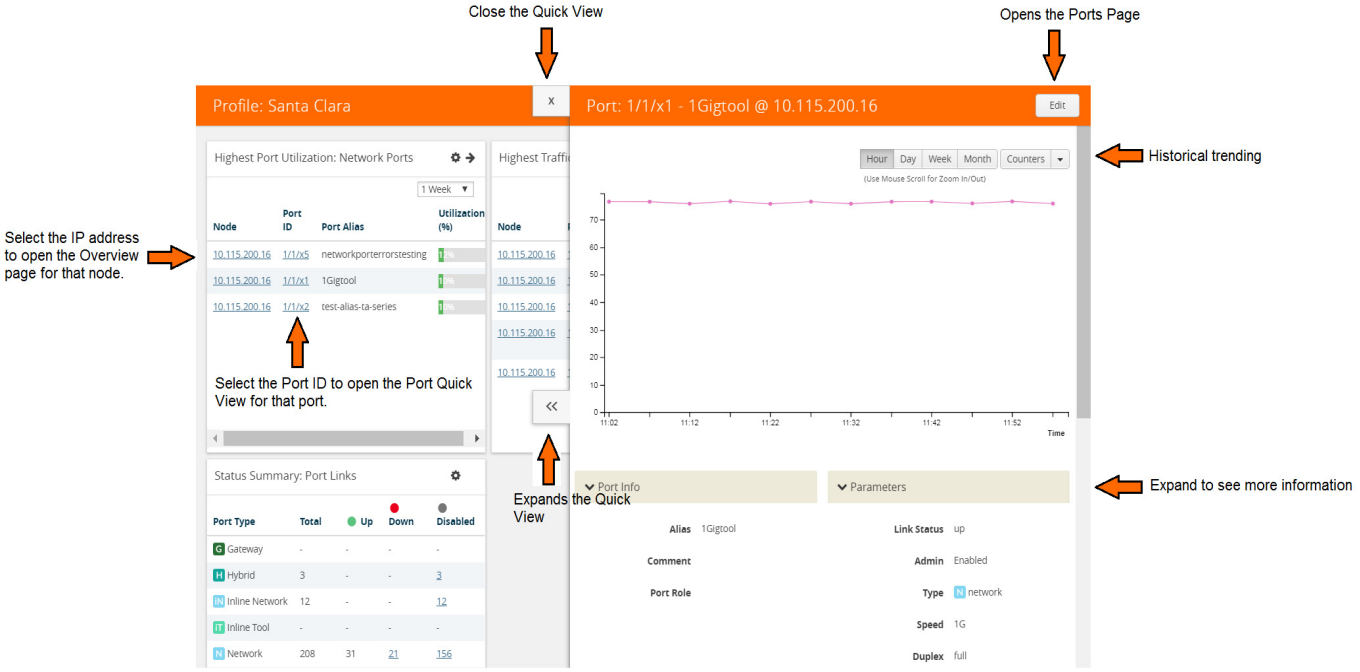

Figure 4-2: Quick View Options with GigaVUE-FM

Return to the Dashboard

At any time, to return back to the Dashboard, click the GigaVUE-FM text on the top left of GigaVUE-FM. Refer to Figure 4-1 on page 39. By default, the Physical & Virtual Dashboard page is displayed.

Table View Customization

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

To customize the columns:

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

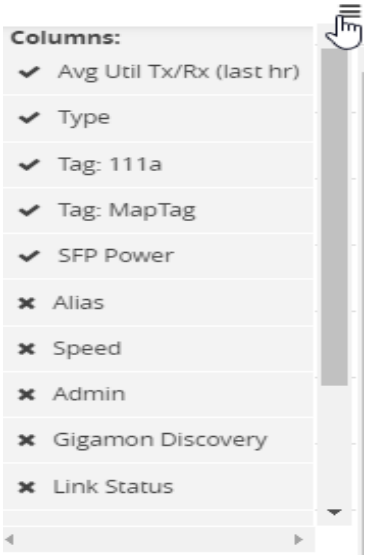

Figure 4-3: Table menu to configure columns

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

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

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

GigaVUE-FM allows you to download the tables in a CSV file format using the **Download all data as csv** option.

Notifications Panel

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

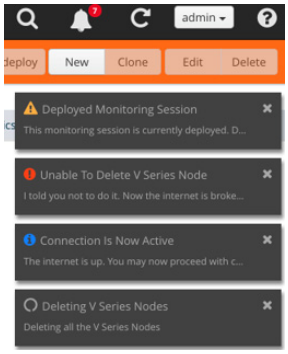

Figure 4-4: Pop-up messages

Long-term Notifications

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

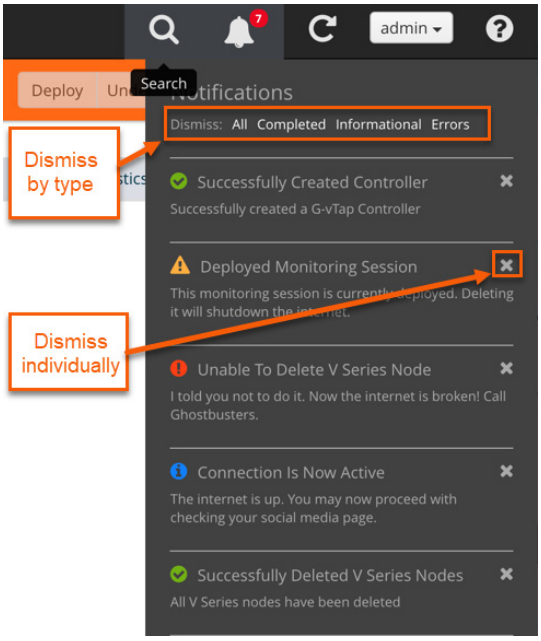

Figure 4-5: Long-term messages

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

Notification Type Icons

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

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

How to Add the GigaVUE-FM Instance Name

The default name displayed on a GigaVUE-FM instance tab is GigaVUE-FM. Refer to Figure 4-6 on page 49.

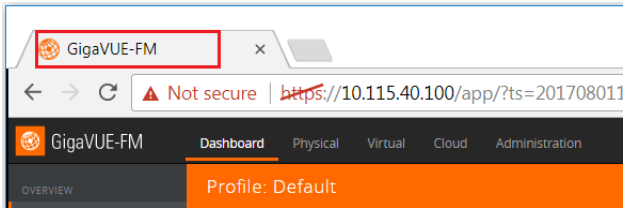

Figure 4-6: Default Tab Name

When you have multiple GigaVUE-FM Instances running in your system, it becomes difficult to differentiate the instances and switch between tabs. Starting in software version 5.1, GigaVUE-FM provides the ability to customize the GigaVUE-FM instance name.

To customize the GigaVUE-FM instance name:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, go to **System** > **Preferences**.
3. Click **Edit**.
4. In the **FM Instance Name** box, enter a name for the GigaVUE-FM instance.
5. Click **Save**.

   The customized GigaVUE-FM instance name is displayed instantly in the tab. The Instance name is also displayed beside the GigaVUE-FM logo. Refer to Figure 4-7 on page 49.

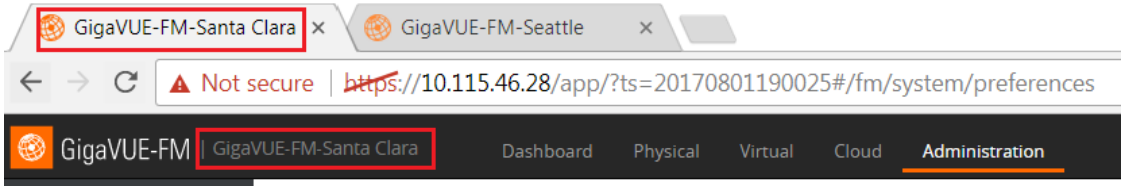

Figure 4-7: GigaVUE-FM Instance Name

How to Search in GigaVUE-FM

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

GigaVUE-FM has an Elastic Search feature that allows you to search for information in GigaVUE-FM as well all the devices and device configurations managed by GigaVUE-FM. Essentially, if it is part of the GigaVUE-FM UI or managed by GigaVUE-FM, you can search for items based on keywords because almost all items are indexed for elastic search. The search feature allows you to search for items across the following:

- Maps
- Roles and Users
- GigaSMART
  - GigaSMART Operations
  - GigaSMART Groups
  - Virtual Ports
  - NetFlow/IPFIX Generation
  - SSL Decryption
  - GTP Whitelists
  - Application Session Filtering (ASF)
- Ports
  - Port Groups
  - Port Pairs
  - Tool Mirrors
  - Stack Links
  - Tunnel Endpoints
  - IP Interfaces
  - Circuit Tunnels
  - GigaStreams
- Chassis and port inventory
- Node Clusters
- VMs
- IP, DNS, and MAC address

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

Categories are another important component of Elastic Search. When you provide a keyword, the search system displays the matching category or categories related to the

keyword search. This provides an automatic filtering of the keyword, helping to narrow your search. Some of the general categories are:

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

You can refine the search categories by using the Filter feature (refer to *Filtering Search Results* on page 61) to further narrow the search results. For example, if the search keyword falls into the GigaSMART category, you can narrow the search further to NetFlow Records.

Performing a Search

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

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

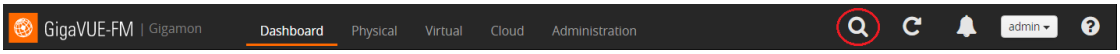

Figure 4-8: GigaVUE-FM Top Menu

2. Type a keyword in the text field.

   As you enter the keyword, the system displays the categories in which the keyword appears and the total matches in that category along with the specific instances. For example, as shown in Figure 4-9, you start to enter an IP address. The search shows that 10.115 occurs in the Cluster category 227 times, Ports 252 times, Filter Templates 661 times, and Action 3390 times.

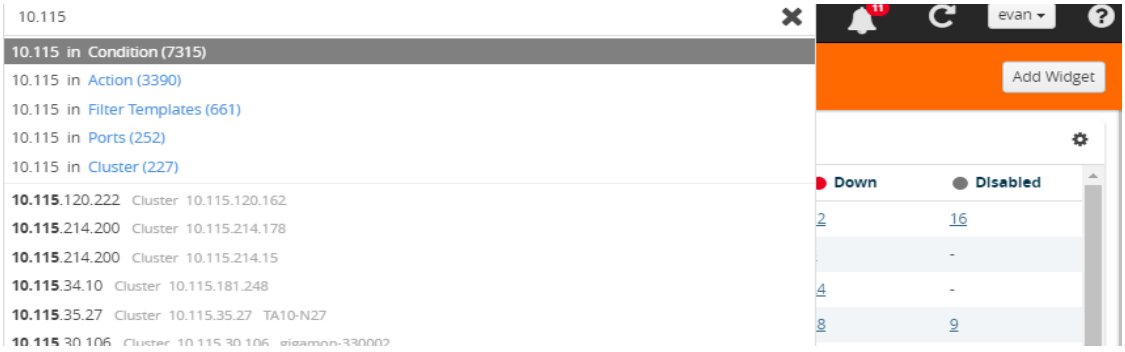

Figure 4-9: GigaVUE-FM Search

3. Once you are done entering the keyword, you can scroll through the list by using the up and down keys on the keyboard. Select an item by pressing the Enter key.
   - Selecting a category, displays all the results in that category, using the category as a filter. You can further refine the results by clicking the **Filter** button. For details on filtering, refer to *Filtering Search Results* on page 61.
   - Selecting a specific result opens the page for that item. For example, as shown in the following figure, if the item is a port, the Port page on the node opens.

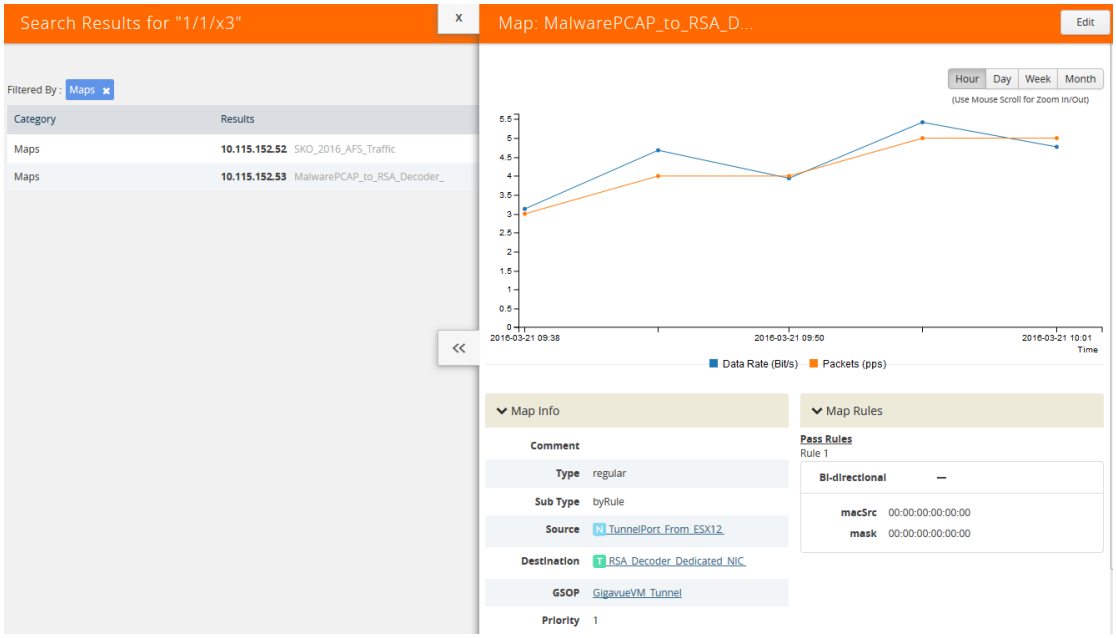

Figure 4-10: Search Results

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

Search Examples

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

- *Searching Maps* on page 52
- *Searching for Roles and Users* on page 56
- *Searching Ports* on page 59

Searching Maps

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

This section provides several examples of searching Maps:

- *Example 1: Searching for a Map by Alias* on page 53
- *Example 2: Searching for a Map with an IP Address* on page 55
- *Example 3: Searching for a Map with a MAC Address* on page 55
- *Example 4: Searching for Maps with Down Ports* on page 56

Example 1: Searching for a Map by Alias

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

1. Click the Search icon in the GigaVUE-FM header to open the Keyword field and start typing the map alias. In this example, APS_Map_L1.

   As you type the search displays the categories and items that match the keyword. Figure 4-11 shows the search results and you can see that the map with Alias APS_Map_L1 is on node 10.115.152.55 without typing the entire string.

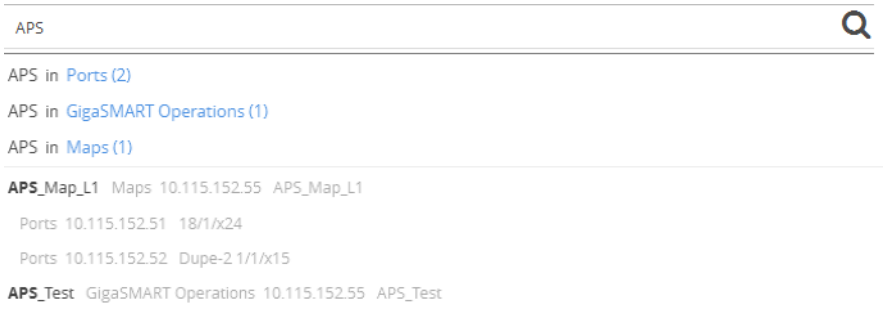

Figure 4-11: Searching by Map Alias

2. In the search results, click on APS_Map_L1.

   GigaVUE-FM opens the Map page as shown in Figure 4-12 on page 54.

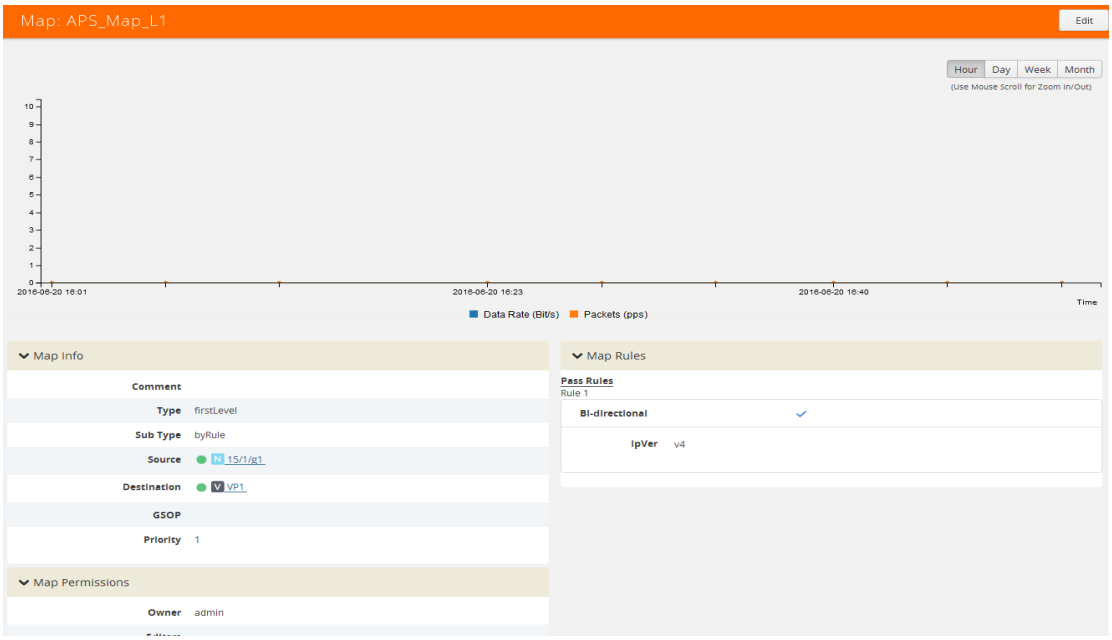

Figure 4-12: Map Page

Example 2: Searching for a Map with an IP Address

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

1. Click the Search icon in the GigaVUE-FM header to open the Keyword field and start typing an IP address. In this example, the IP address is 10.10.8.240.

   As you type the search displays the categories and items that match the keyword. Figure 4-13 shows the search results and you can see that the two maps partially map the address that you are typing.

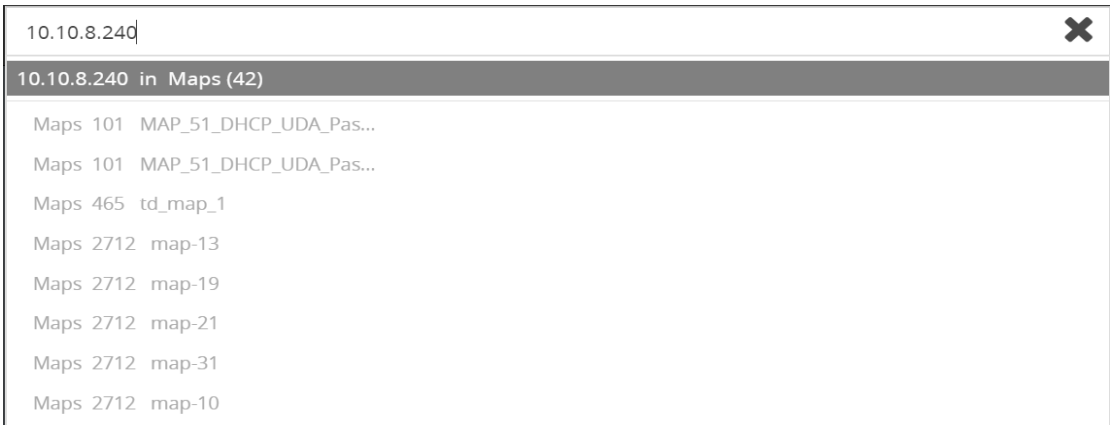

Figure 4-13: Search for IP Address

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

   The Filter page opens. Figure 4-14 on page 55 shows the results with maps the that contain the IP address. In this scenario, the item of interest is the cluster MAP_51_DHCP_UDA_Pass_Drop_slot2, so you click on the cluster name.

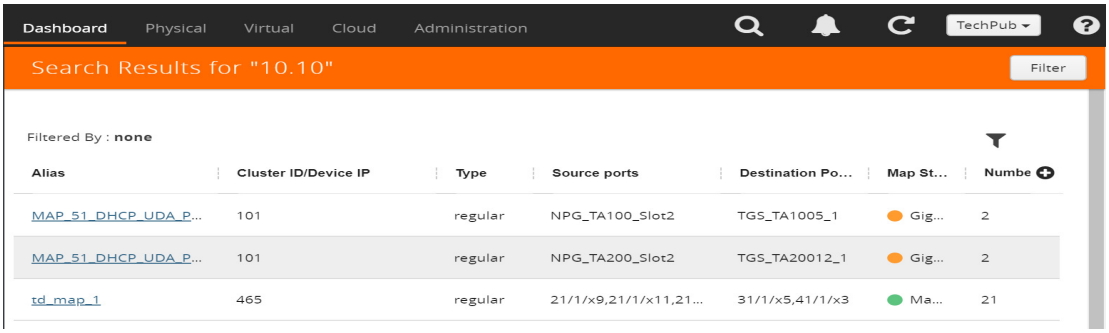

Figure 4-14: Search Result Filter Page

   A Map quick view opens, showing the information for the map and that Rule 1 contains a IP4 Destination rule with the IP address 10.10.8.240.

Example 3: Searching for a Map with a MAC Address

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

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

After entering the MAC address in the Keyword field, only one map is found as shown in Figure 4-15.

Figure 4-15: Search for MAC Address

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

Example 4: Searching for Maps with Down Ports

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

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

   As shown in Figure 4-16, the keyword occurs in several categories.

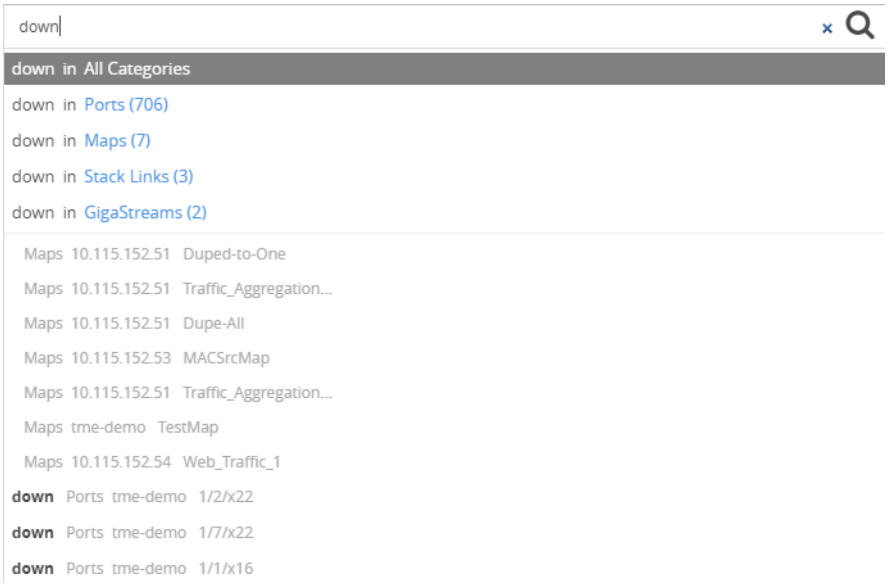

Figure 4-16: Search Results for down Keyword

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

Searching for Roles and Users

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

- *Example 1: Searching for Monitor Role* on page 56
- *Example 2: Searching for a User* on page 58

Example 1: Searching for Monitor Role

In this example, you are looking for the Monitor Role is applied.

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

   As shown in Figure 4-17, the keyword occurs in several categories.

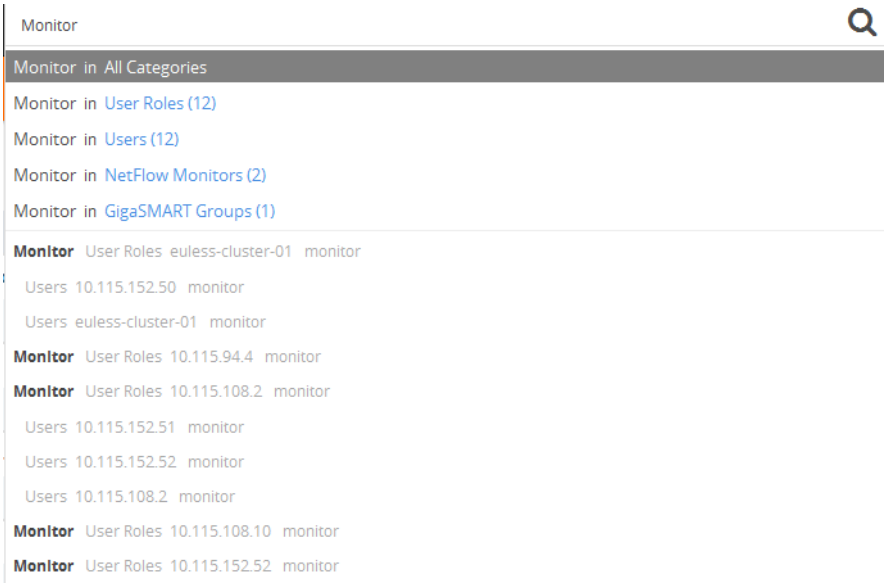

Figure 4-17: Categories for Monitor

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

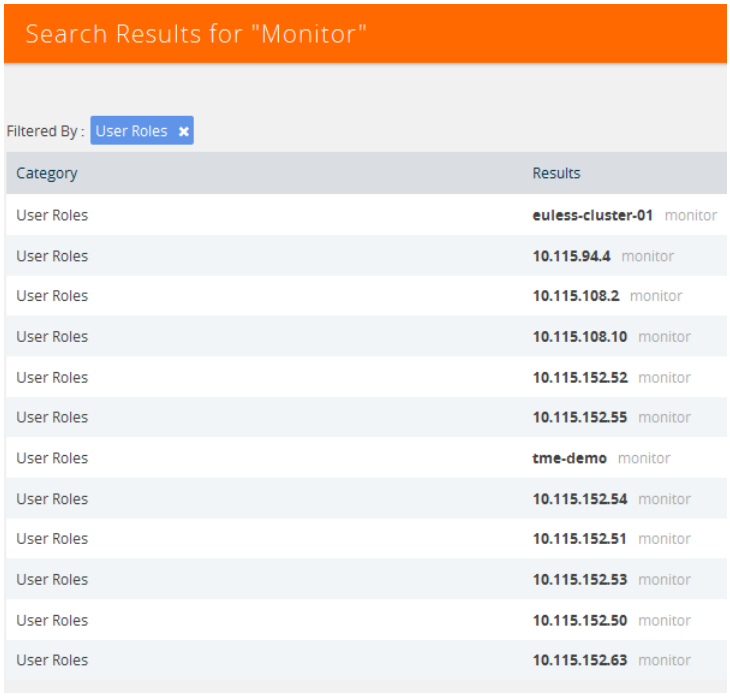

Figure 4-18: Search Results for Monitor

3. From the search results, drill down further by selecting one of the results. For example, the cluster tme-demo.

   GigaVUE-FM takes you to the Roles page for the selected cluster as shown in Figure 4-19 and indicates in the Keywords field that further searches are restricted to the cluster.

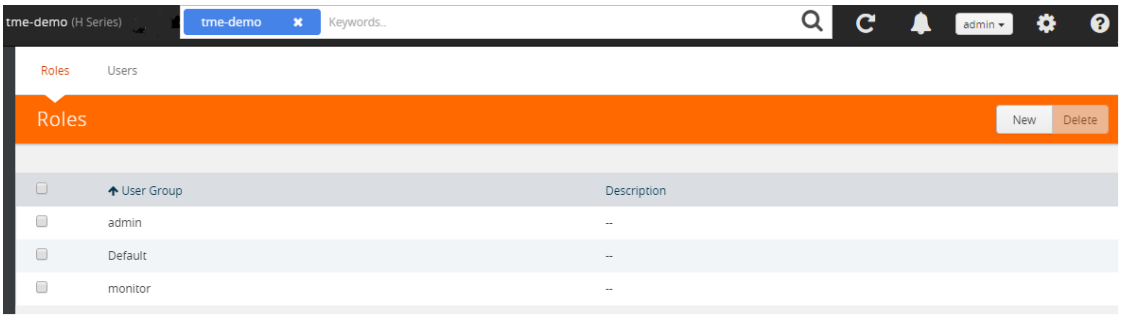

Figure 4-19: Role From Search Result

Example 2: Searching for a User

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

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

   In this case, the search discovers only one item as shown in the following figure.

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

Searching Ports

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

This section provides few examples related to searching for ports:

- *Example 1: Searching for Down Ports* on page 59
- *Example 2: Searching for Port Details of Devices Managed by GigaVUE-FM* on page 60

Example 1: Searching for Down Ports

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

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

   As shown in Figure 4-20, the 1/1/x5 and down occur 60 times in the Port category.

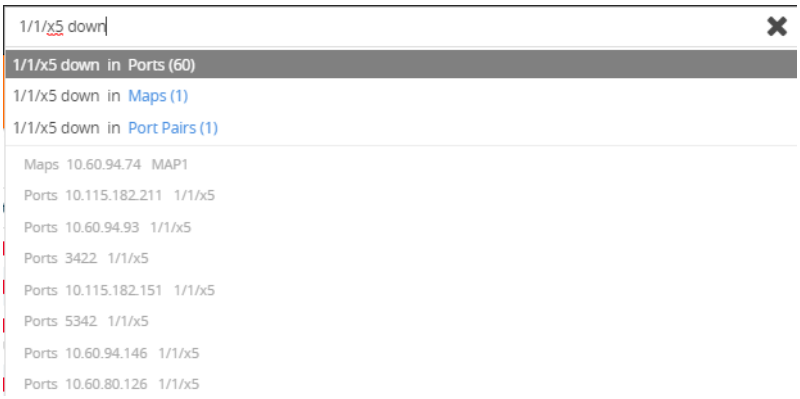

Figure 4-20: Categories Returned for Port ID and Down

2. Click on the Ports category to view the results.
3. Click on an item search results to see the port information. A Port quick view displays for the selected port as shown in Figure 4-21.

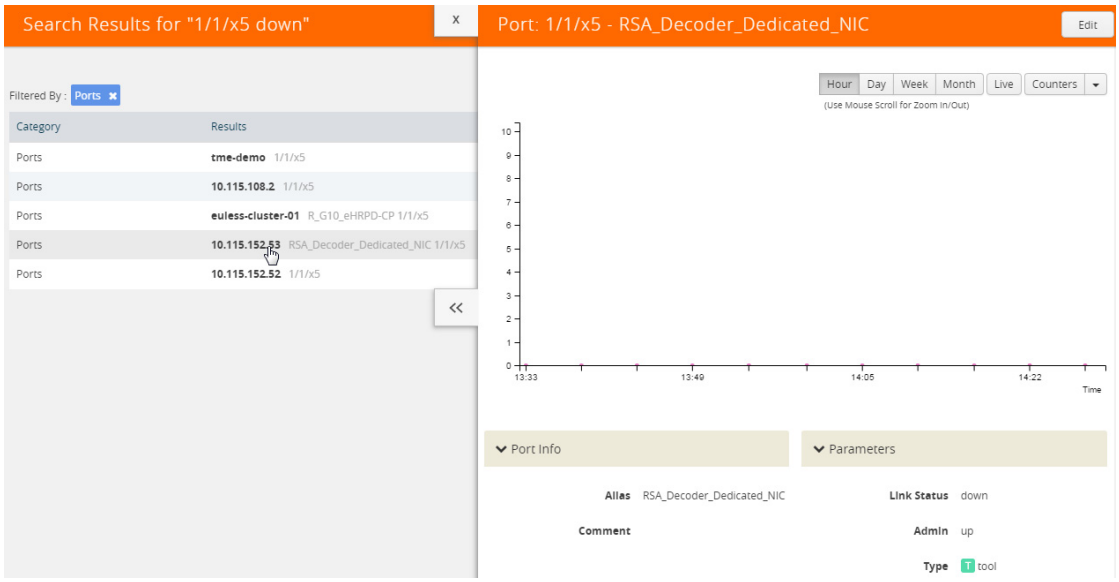

Figure 4-21: Port Quick View for Search Results

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

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

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

   As shown in Figure 4-22, the search returns results in several categories.

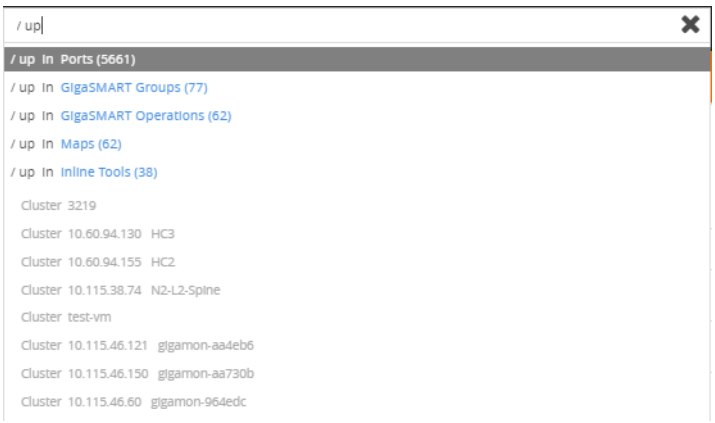

Figure 4-22: Categories Returned for Port Details of All Devices and up

2. Click on the Ports category to view the results.

3. Click on an item search results to see the port information.

Filtering Search Results

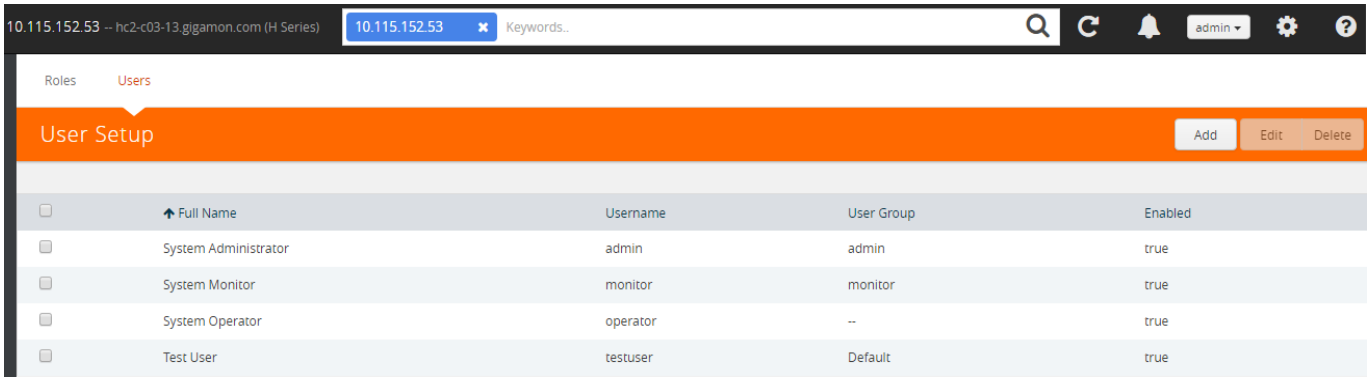

When selecting a search result for a category, GigaVUE-FM opens a page that lists the results for that category. Refer to Figure 4-23 for an example.

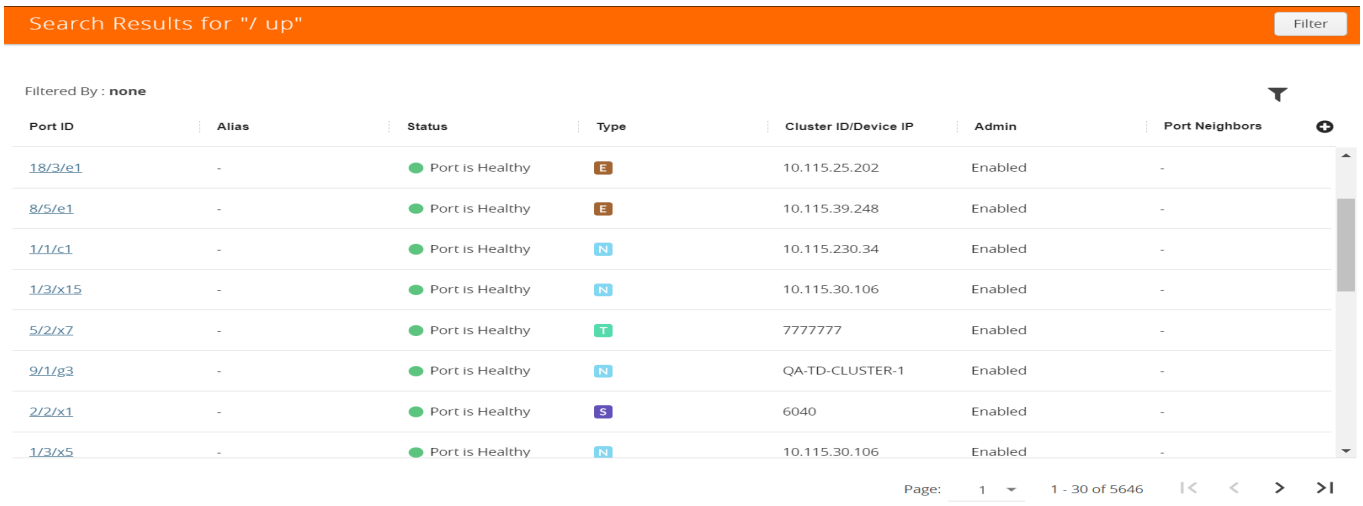

Figure 4-23: Filter Results by Category

To change the filter, do the following:

1. Click **Filter**.

   The Filter quick view displays.

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

Part 2: Installation and Upgrade

This section provides information about installing GigaVUE-FM on ESX, MS HyperV, and KVM. Upgrade information is also provided. The following topics are covered:

5 Install GigaVUE-FM on VMware ESXi

This section describes how to install GigaVUE-FM on VMware hypervisor, ESXi. It consists of the following main sections:

- *Before You Install* on page 65 describes the minimum hardware and computing requirements.
- *Install New GigaVUE-FM on VMware ESXi* on page 67 describes the steps to install and deploy GigaVUE-FM on VMware ESXi hypervisor.
- *Configure SSH Settings* on page 86 describes the CLI for setting SSH.
- *HTTP/HTTPS Ports* on page 87 describes the CLI for setting the Web access.

Before You Install

This section describes the hardware and virtual computing requirements for GigaVUE-FM. Ensure that the GigaVUE-FM time is set correctly to ensure accuracy of the trending data that is captured.

Prerequisites for GigaVUE-FM

Before Installing GigaVUE-FM, ensure that VMware vSphere Standard, Enterprise, or Enterprise Plus is installed on hardware that meets minimum requirements. The following VMware vSphere versions are supported. Note the minimum version requirements under *Hardware Requirements* on page 65.

VMware ESXi and NSX-V Hardware Requirements

The following table describes the hardware requirements on which VMware ESXi runs GigaVUE-FM.

Table 5-1: Hardware Requirements for VMware Hypervisor

Hardware Requirements	
VMware Hypervisor	vSphere ESXi: v5.5 and above

Table 5-1: Hardware Requirements for VMware Hypervisor

Hardware Requirements	
CPU	One or more 64-bit x86 CPUs with virtualization assist (Intel-VT or AMD-V) enabled. **Note**: To run GigaVUE-FM, hardware support for virtualization must be enabled on the VMware ESXi host. Make sure that the BIOS option for virtualization support is not disabled. For more information, see your BIOS documentation
RAM	At least 8GB
Disk Space	At least 40GB shared (FC, iSCSI, NFS, or FCoE) or locally attached storage (PATA, SATA, SCSI)
Network	At least one 1Gb NIC

The following table lists the virtual computing resources for GigaVUE-FM.

Table 5-2: Minimum Virtual Computing Requirements for VMware Hypervisor

Minimum Virtual Computing Requirements	
Memory	Minimum 8GB memory
Virtual CPU (vCPU)	2 vCPU
Virtual Storage for Guest	40GB using PVSCSI (VMware Paravirtual SCSI) 80GB (or more) using FabricVUE Traffic Analyzer
Virtual Network Interfaces	One vNIC using VMXNET3 (VMware 3rd Generation Paravirtual NIC)

Supported Browsers

GigaVUE-FM has been tested on the following browsers:

Browser	Version
Mozilla Firefox™	• Version 49.00
Windows® Internet Explorer®	• Version 11 and higher
Apple® Safari®	• Version 9.1 and Higher
Google®Chrome®	• Version 54 and higher
Microsoft Edge	• Version 38

Notes:

- Only the browsers that support TLS v1.2 can access GigaVUE-FM.
- DNS prefetch is a known limitation of Internet Explorer 11. If GigaVUE-FM is configured with DNS and you are using Internet Explorer 11, every new screen can be slowed significantly. If a direct IP address is used instead of a DNS name, the UI response is similar to other browsers. It is recommended that you use the GigaVUE-FM IP when using Internet Explorer 11 or use either a FireFox or Chrome browser instead.

- IE11 Compatibility view mode is not supported.

Install New GigaVUE-FM on VMware ESXi

The GigaVUE-FM software package is distributed as an OVA file. The following sections describe how to deploy a fresh installation of GigaVUE-FM on an ESXi host and perform its initial configuration:

- *Deploy GigaVUE-FM from an OVA File* on page 67
- *Initial GigaVUE-FM Configuration* on page 76

Upgrades and OVA Files

You can also use the GigaVUE-FM OVA file to upgrade an existing deployment. However, settings and data are not retained when updating from an OVA file. Upgrade using the provided image file to retain settings and data across an upgrade. For details, refer to *Upgrade an Existing GigaVUE-FM Deployment* on page 113.

Deploy GigaVUE-FM from an OVA File

Use the vSphere Client to install the GigaVUE-FM OVA file. Starting from software version 5.3, you cannot deploy GigaVUE-FM directly from the ESXi host. You must login to the VCenter on the vSphere client to deploy a GigaVUE-FM instance.

NOTE: The OVA file must be stored in a location that is accessible to the vSphere Client. This location cannot be a datastore accessible to the ESXi host which will be the target of the deployment.

The following steps are shown using the ESXi version 5.5 Update 2b. ESXi version 6.x will use the same steps to deploy GigaVUE-FM; however, the screens may look different from the ESXi version 5.5.

NOTE: Starting in software version 5.4.01, you cannot deploy the GigaVUE-FM OVA file on older versions of VMware ESXi. If you have an older version of ESXi, then you must upgrade your VMware ESXi to at least version 5.5. Otherwise, GigaVUE-FM OVA deployment will fail.

If you apply the current GigaVUE-FM release as an image upgrade, then it will not change the VM Virtual Hardware version of the existing GigaVUE-FM virtual machine. To change the Virtual Hardware version of an existing GigaVUE-FM installation, you must shut-down the VM instance and use the 'Upgrade Virtual Hardware' dialog in the vSphere client.

To deploy a GigaVUE-FM instance:

1. Log in to vCenter on the vSphere Client. The main page of the vSphere Client opens as shown in Figure 5-1.

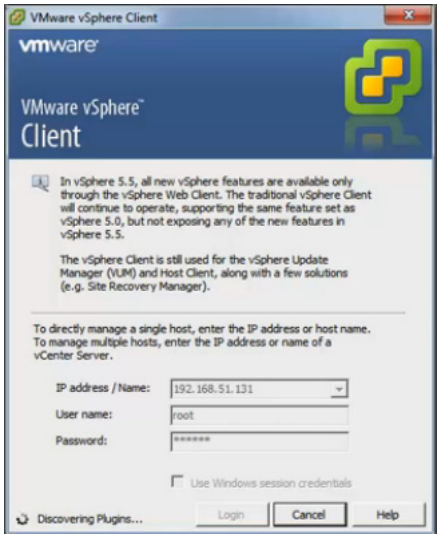

Figure 5-1: vSphere Client 5.5

2. Select the entry for the ESXi Host or Data Center on which you would like to install the GigaVUE-FM instance in the inventory panel.
3. From the vSphere Client, click the **File** menu and select **Deploy OVF Template** as shown in Figure 5-2.

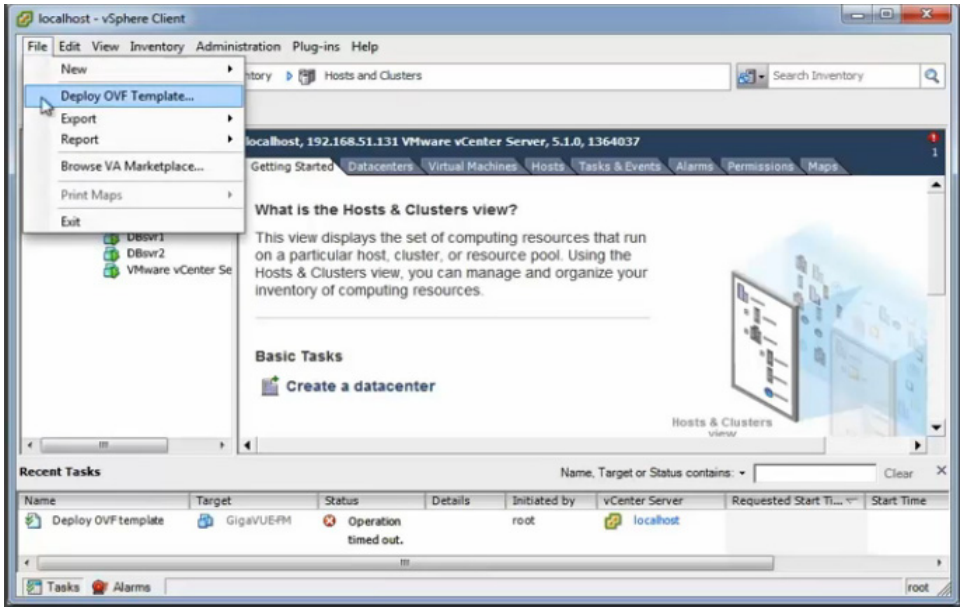

Figure 5-2: vSphere Client: OVF Template

4. When the **Source** page of the **Deploy OVF Template** wizard opens, do the following to open the OVA file:

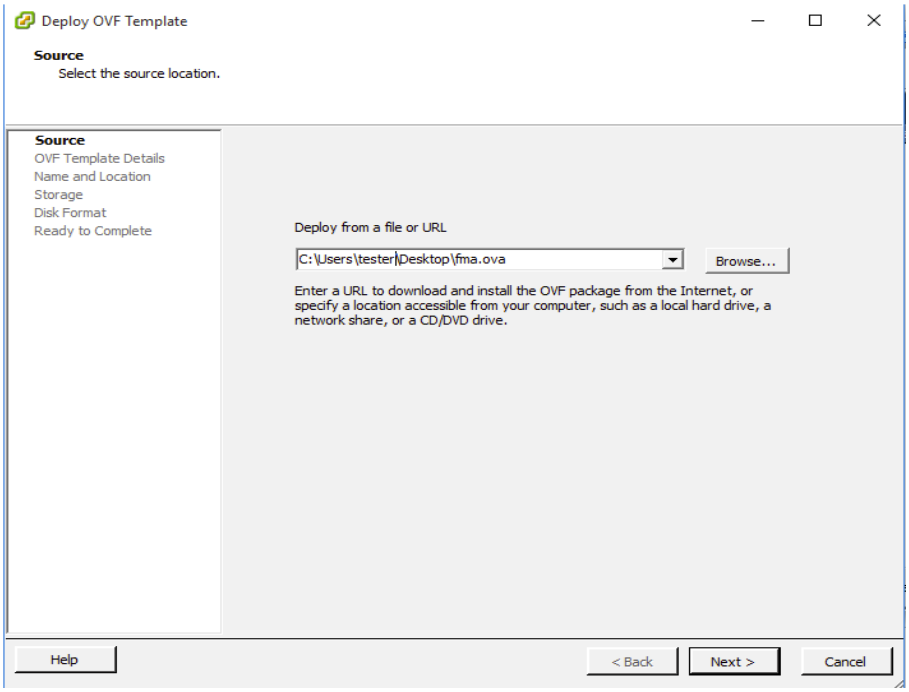

Figure 5-3: vSphere Client: Find the Location of the OVA File

a. Click **Browse** to navigate to the OVA file available on your local machine and its accessible network shares or to an HTTP URL.

b. Select the GigaVUE-FM OVA file and click **Open**.

The **Open** dialog closes, returning you to the **Source** page with the OVA file displayed in the field on the page.

c. Click the **Next**.

The **OVF Template Details** page opens, showing the details of the OVA file. Figure 5-4 shows an example of the details page.

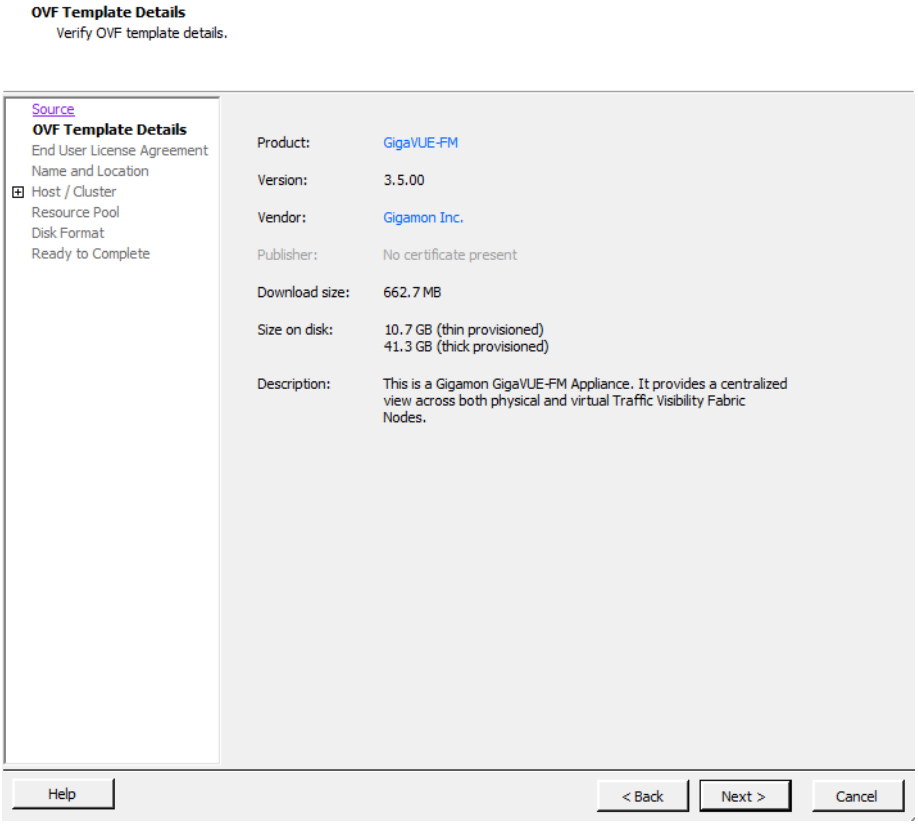

Figure 5-4: vSphere Client: OVF Template Details

5. Review the EULA for the OVA file, and then click **Accept**, and then **Next**.
6. Select the name of the GigaVUE-FM instance and the host to which to deploy it.
   a. When the **Name and Location** page opens, enter a name for this GigaVUE-FM instance, select the location to deploy it to, and then click **Next**.

      Figure 5-5 shows an example of the **Name and Location** page, where the specified name is FM and the specified location is DC-2/discovered virtual machine.

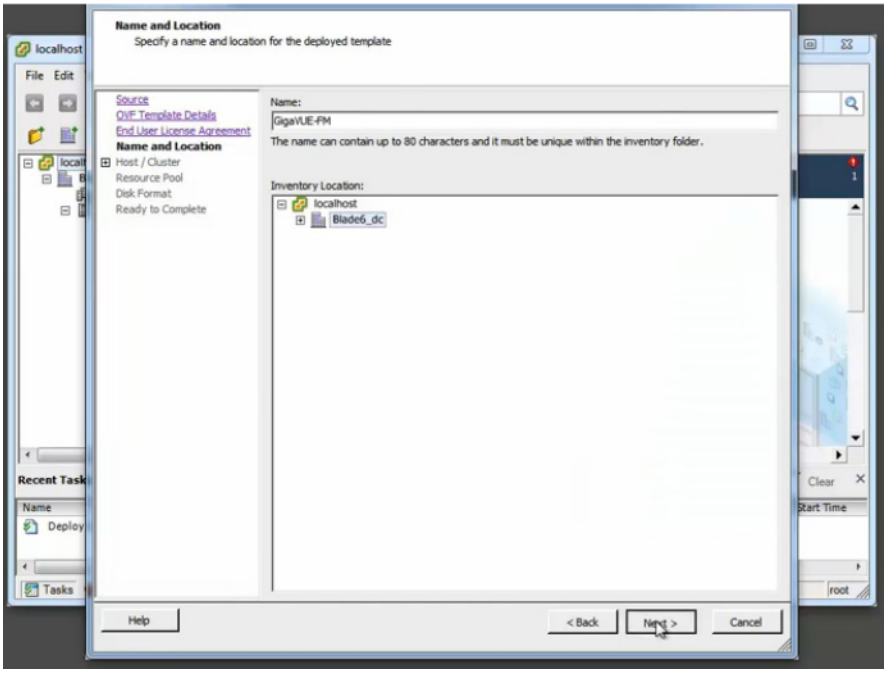

Figure 5-5: vSphere Client: Enter Location for OVA File on ESXi Server

b. Select the host to which you wish to deploy this GigaVUE-FM instance.

NOTE: If you selected a Data Center rather than an ESXi host in Step 2, you are prompted to select a host now.

c. Click **Next**.

The OVF Wizard performs a validation to ensure that the selected host has all the resources required for this GigaVUE-FM deployment. and presents the **Storage** page.

7. Select the storage location for the virtual machine files by doing the following:

a. After the **Storage** page opens, choose the datastore where the virtual machine's files will be stored.

b. Click **Next**.

Figure 5-6 shows an example of the Storage page with the datastore selected.

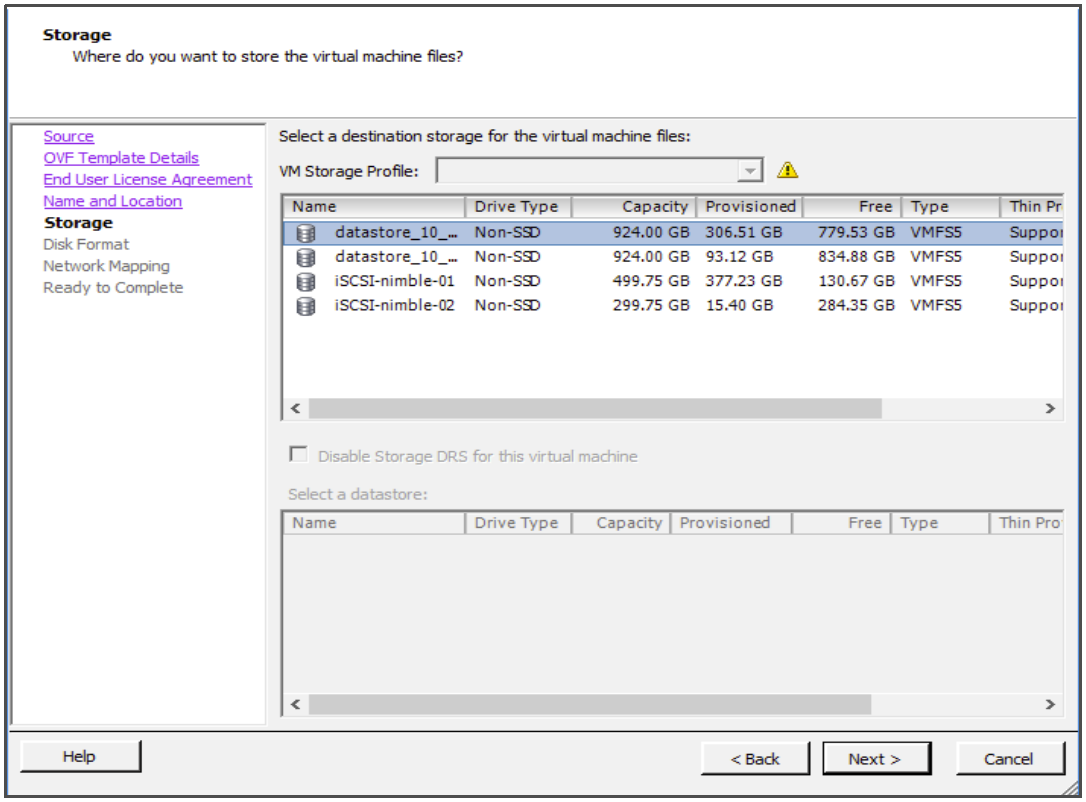

Figure 5-6: vSphere Client: Storage Information Verification

8. Set the disk format by doing the following:

   a. After **Disk Format** page opens, select **Thick Provisioning** as the format for the virtual disks and provisioning.

   NOTE: You *must deploy* FM using **Thick Provisioning**. Any other choice results in FM not working correctly.

   b. Click **Next**.

   Figure 5-7shows an example of the Disk Format page with Thick Provisioning selected.

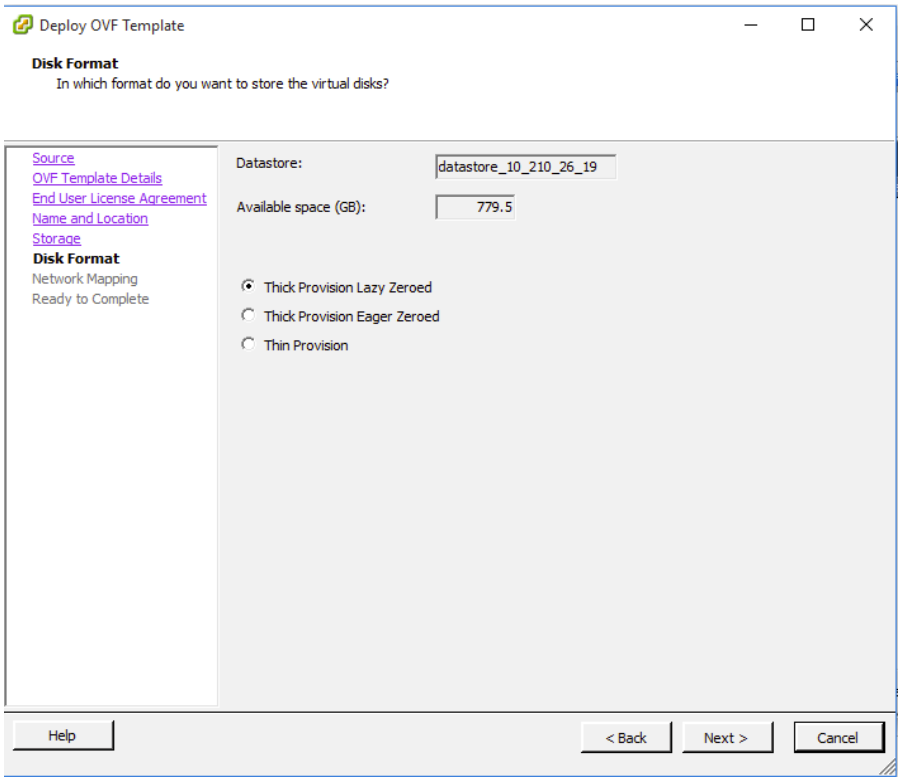

Figure 5-7: vSphere Client: Select Datastore

9. After the **Network Mapping** page opens, set network mapping by doing either of the following, depending on how your are deploying, and then click **Next**:

   - ***If you are not deploying on a standalone ESXi host***, the **Network Mapping** displays under **Source Networks**. Use the drop-down lists to assign the correct **Destination Network** to the source network.

   - ***If you are deploying GigaVUE-FM on a standalone ESXi host***, the network mapping is set automatically by assigning the destination network to the VM Network. In case of multiple port groups, you need to manually assign the destination network to the VM Network.

     Figure 5-8 shows an example of the Network Mapping page.

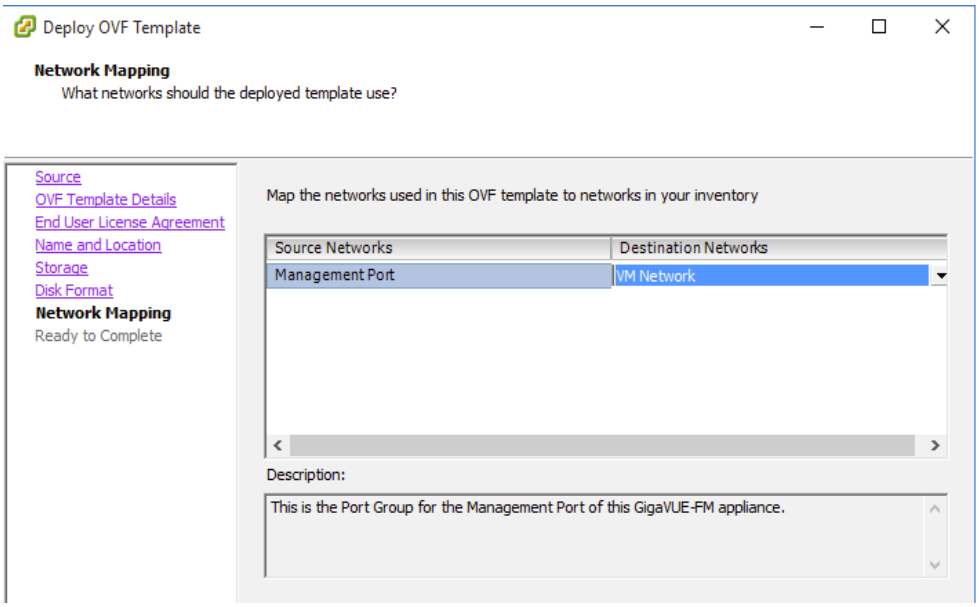

Figure 5-8: vSphere Client: Network Mapping Menu

10. In the **Properties** page, enter the hostname of the GigaVUE-FM instance and set the admin password. Configure the IP networking information and click **Next**.

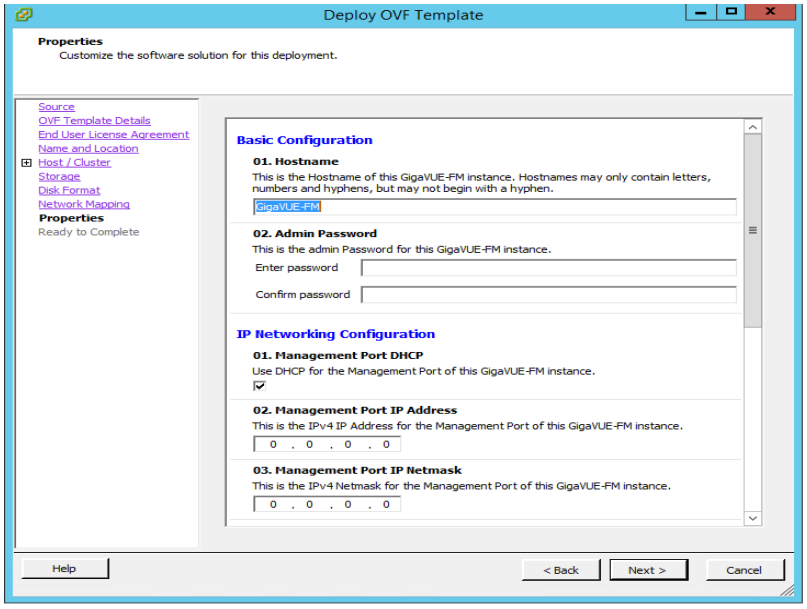

Figure 5-9: vSphere Client: Properties Page

11. After the **Ready to Complete** page opens, do the following:

    a. Verify that all of the settings are correct.

    b. (Optional) Select **Power on after deployment**.

    NOTE: Do not select **Power on after deployment** if you want to change the default configuration of GigaVUE-FM. The configuration changes could be as follows—adding vCPUs, increasing the memory size, or adding another Network Interface

Card. For more information on these configurations, refer to *Perform Initial Configuration* on page 77.

c. Click **Finish**.

Figure 5-10 shows an example of the Ready to Complete page.

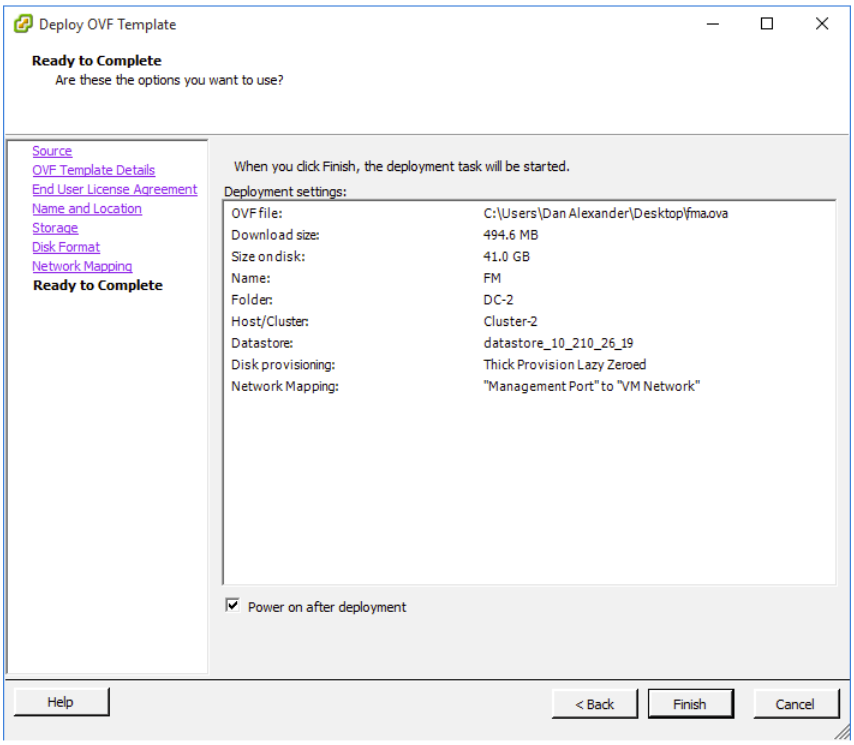

Figure 5-10: vSphere Client: Final Verification before Installing

After clicking Finish, a dialog opens (refer to Figure 5-11), showing the progress of the deployment operation. When the operation completes, you have successfully deployed a GigaVUE-FM instance.

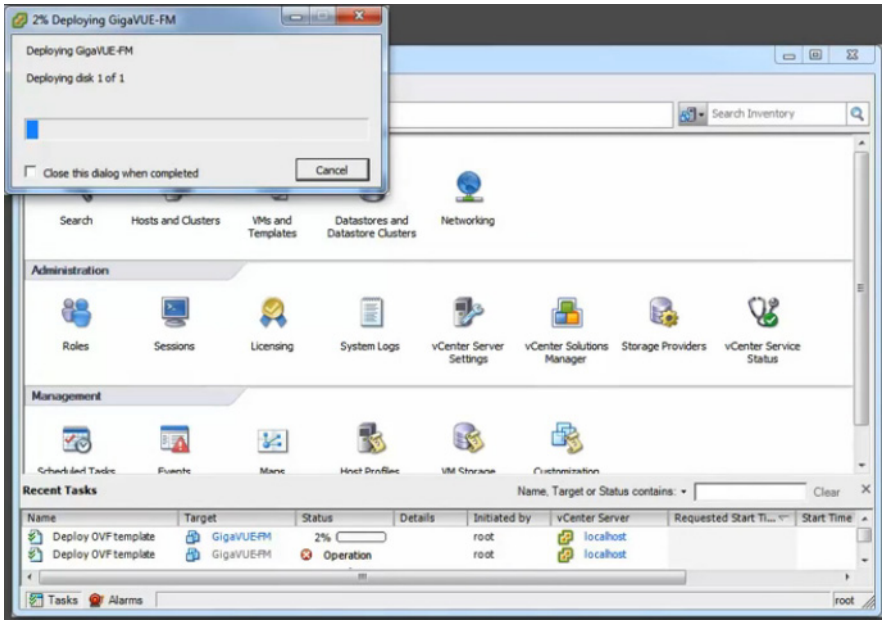

Figure 5-11: vSphere Client: Dialog Showing the Installation Progress

Important: Clear the browser cache before logging in to GigaVUE-FM!

Redeploy GigaVUE-FM Instance (with VMs Already Deployed)

If there is a need to remove an existing instance of GigaVUE-FM and reinstall it, delete all the virtual centers configured in the Virtual > Management > Virtual Center tab prior to deleting the GigaVUE-FM. To re-install GigaVUE-FM, refer to the *Install New GigaVUE-FM on VMware ESXi* or *Install GigaVUE-FM for Microsoft Hyper-V*.

By deleting the virtual centers, you will lose all the GigaVUE-VM nodes and vMaps configured on those virtual centers and they will need to be recreated.

Initial GigaVUE-FM Configuration

After you have deployed a new GigaVUE-FM instance, you need to perform an initial configuration before you can start using GigaVUE-FM. This procedure only needs to be performed once for each GigaVUE-FM instance deployed.

NOTE: Use Care when Shutting Down or Rebooting a GigaVUE-FM. **Never** directly Power-Off the virtual machine. In VMware ESXi environment when using vSphere client, ALWAYS use **Shut Down Guest OS** or **Restart Guest** functions from VMware. Access is available from either the FILE menu or from the appropriate buttons on the GigaVUE-FM console. Failure to follow these steps may lead to database corruption issues.

How to Use Fault Tolerance for GigaVUE-FM Deployments (VMware ESXi only)

Gigamon recommends that you enable the VMware Fault Tolerance feature for the GigaVUE-FM virtual machine, providing redundancy in the case of a failure. Enabling the VMware Fault Tolerance feature provides a “hot” GigaVUE-FM virtual machine instance on another ESXi host in the cluster. If the ESXi host with the primary GigaVUE-FM instance goes down, you can take advantage of the Fault Tolerance feature to continue GigaVUE-FM operations.

When in Fault Tolerance mode, the MAC address and the UUID for both the primary GigaVUE-FM and the “hot” GigaVUE-FM virtual machine instance remains the same, therefore there is no need to update the existing licenses for GigaVUE-FM or GigaVUE-VM that are installed on the primary.

Both instances of the GigaVUE-FM implement VMware vLockstep technology to keep in virtual lockstep with each other. Any events are executed on the primary and then transmitted over a Gigabit Ethernet network to the other instance. Both instances access a common disk and appear as a single instance since they share the MAC address and UUID.

NOTE: Refer to the VMware Fault Tolerance documentation for deployment requirements and instructions.

Depending on the host configurations, there may be a need to shutdown GigaVUE-FM (primary) to enable the Fault Tolerance mode.

Perform Initial Configuration

Before powering on GigaVUE-FM, you can optionally perform the following:

- *Add Additional vNIC* on page 77
- *Increase Memory* on page 79
- *Add vCPUs* on page 80

Add Additional vNIC

Gigamon allows you to configure GigaVUE-FM with two network interfaces—eth0 and eth1. The network interface eth0 can be configured to connect to a network used to manage Gigamon devices. The other network interface eth1 can be configured to connect to a network hosting different servers like SMTP server, Archive server, and so on.

To add an additional vNIC:

1. Right-click the GigaVUE-FM instance and select **Edit Settings...** Refer to Figure 5-12.

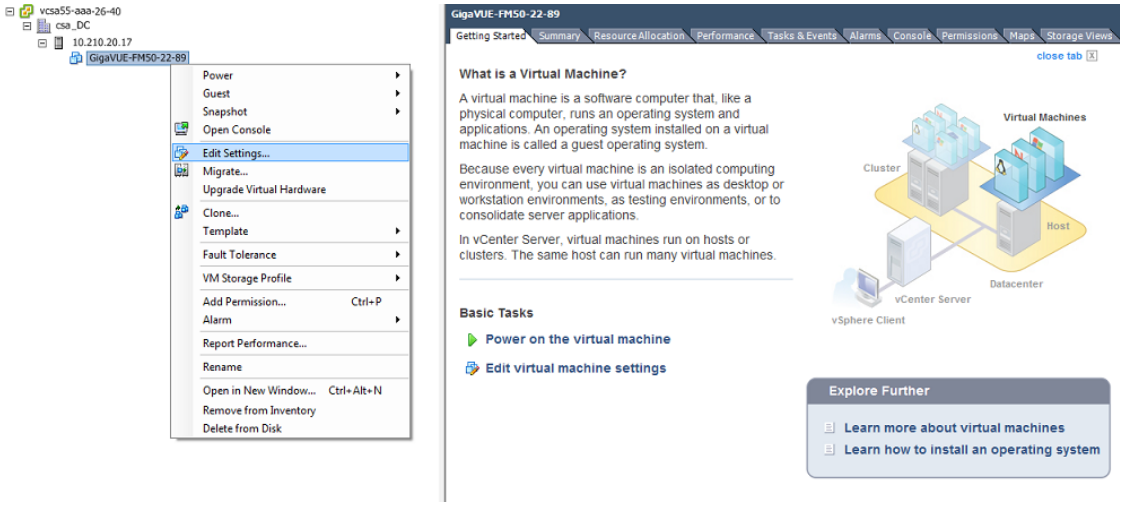

Figure 5-12: vSphere Client: Edit Settings

2. In the Hardware tab, click **Add**.
3. In the Add Hardware dialog box, select **Ethernet Adapter** and then click **Next**. Refer to Figure 5-13.

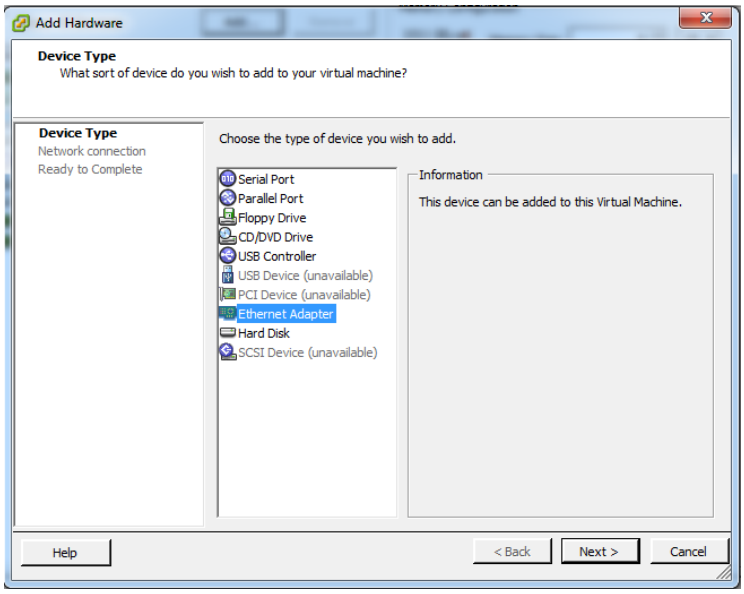

Figure 5-13: vSphere Client: Selecting Ethernet Adapter

4. In the **Adapter Type** drop-down list, select an appropriate adapter type. Refer to Figure 5-14.

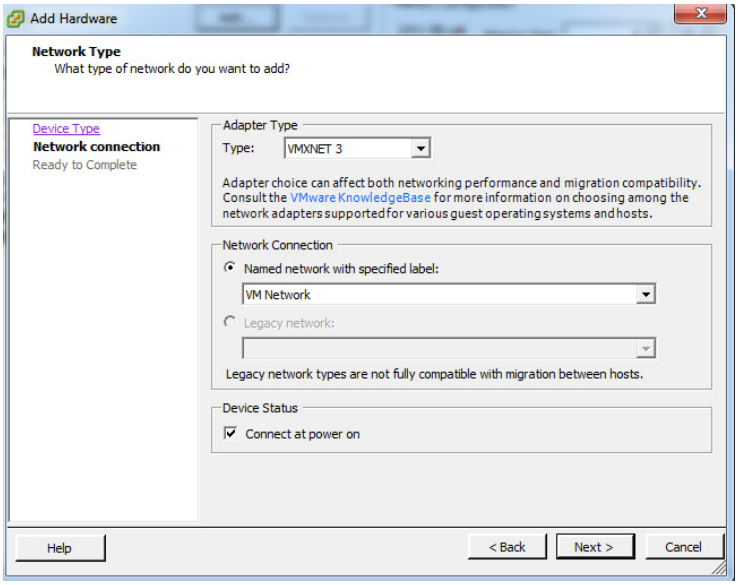

Figure 5-14: vSphere Client: Selecting Device Type

5. In the **Named network with specified label** drop-down list, select an appropriate network and click **Next**.

   The Network adapter 2 is added to GigaVUE-FM.

Increase Memory

Based on the requirement, you can increase the memory of the GigaVUE-FM instance.

To increase the memory:

1. Right-click the GigaVUE-FM instance and select **Edit Settings...** Refer to Figure 5-15.

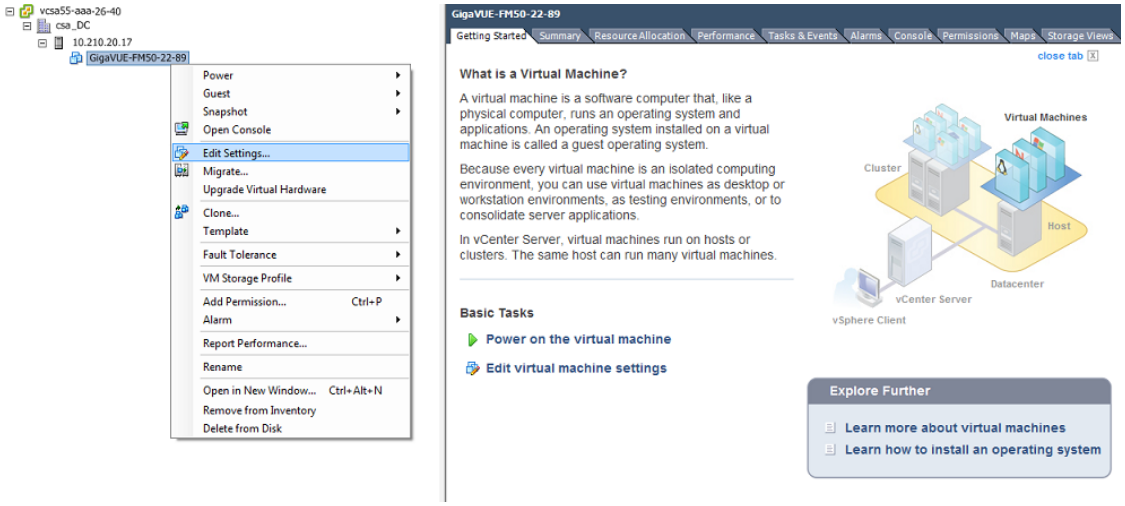

Figure 5-15: vSphere Client: Edit Settings

2. In the Hardware tab, select **Memory**.
3. In Memory Configuration, increase the size of the memory as per your requirement. Refer to the recommended size for your guest OS in the dialog box. Refer to Figure 5-16.

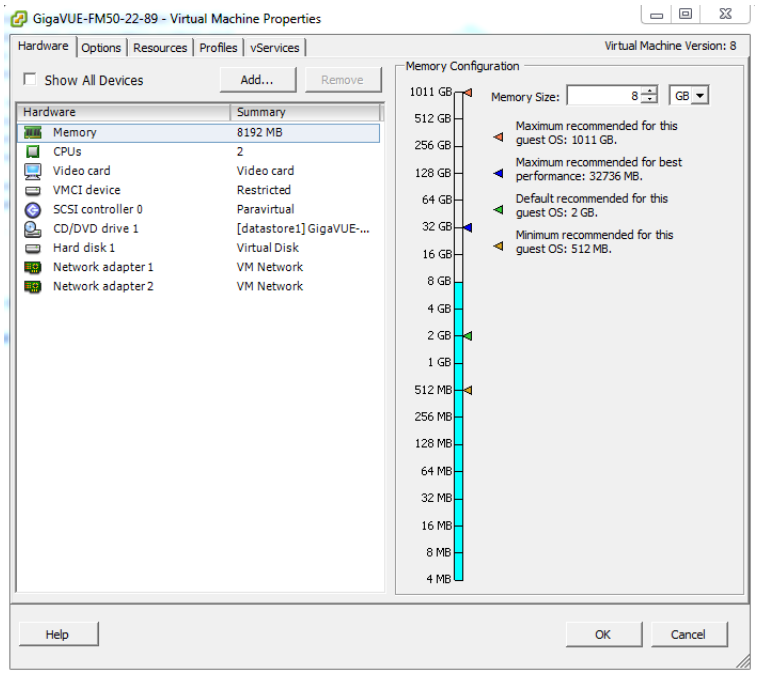

Figure 5-16: vSphere Client: Increasing Memory

4. Click **OK**.

Add vCPUs

Based on the requirement, you can add additional vCPUs to the GigaVUE-FM instance.

1. Right-click the GigaVUE-FM instance and select **Edit Settings...**

2. In the Hardware tab, select **CPUs**. Refer to Figure 5-17.

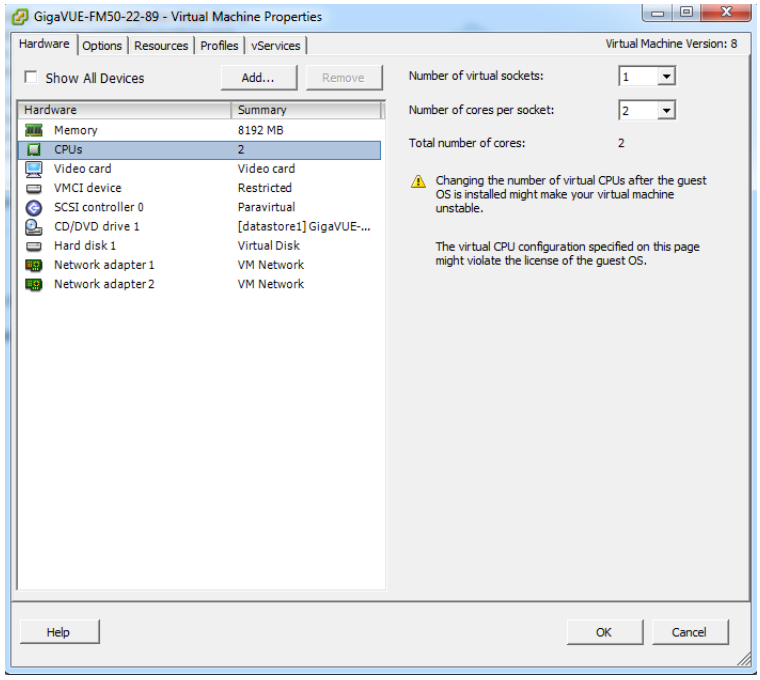

Figure 5-17: vSphere Client: Adding vCPUs

3. In the **Number of virtual sockets** field, enter the appropriate number of sockets.
4. In the **Number of cores per socket** field, enter the appropriate number of cores per socket.
5. Click **OK**.

Following are the steps to perform the initial configuration of GigaVUE-FM after installing on VMware ESXi:

1. Log in to vCenter in the vSphere Client.
2. Ensure that the UTC time for GigaVUE-FM is configured correctly. Refer to the vSphere documentation for instructions on how to set the time.
3. If you checked the **Power on after deployment** box at the end of the GigaVUE-FM deployment in the previous procedure, then the GigaVUE-FM instance starts automatically in vSphere Client.

   If you did not check the box, you can power GigaVUE-FM on now by right-clicking the GigaVUE-FM instance (refer to Figure 5-18) in the vSphere Client by selecting **Power**, and then **Power On**.

   A GigaVUE-FM console displays a login prompt.

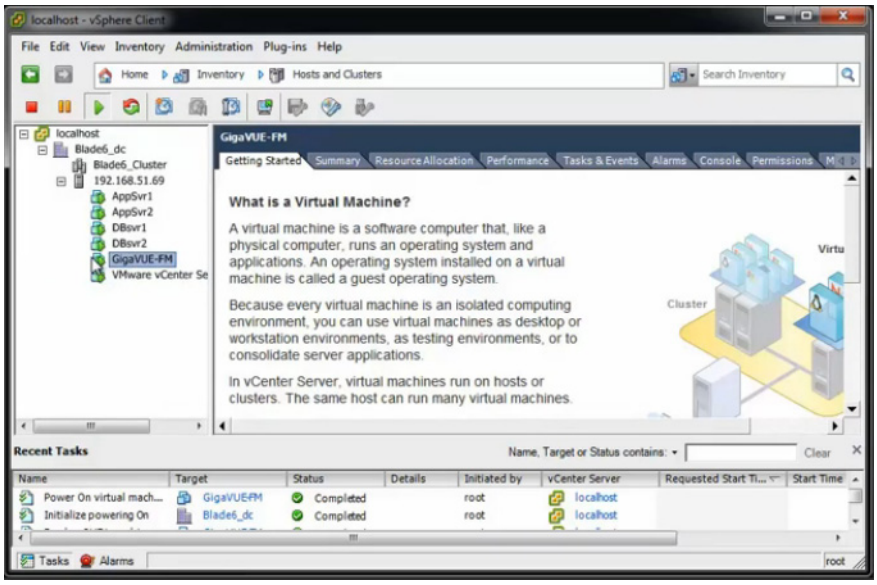

Figure 5-18: GigaVUE-FM in vSphere Client

4. Log in as **admin** with password **admin123A!**

   For a new installation of GigaVUE-FM, a password is required.

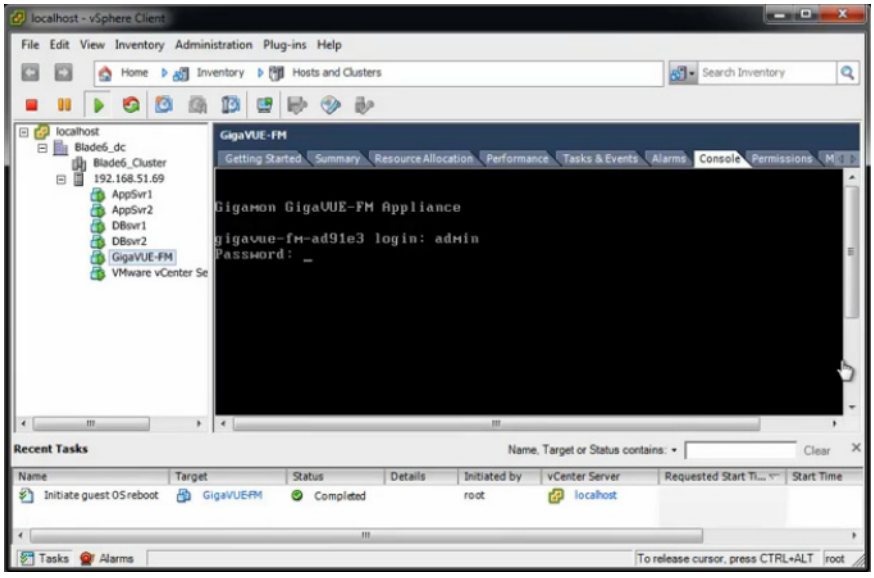

Figure 5-19: Jump Start Configuration Starts Automatically

5. The jump start configuration for GigaVUE-FM starts automatically.

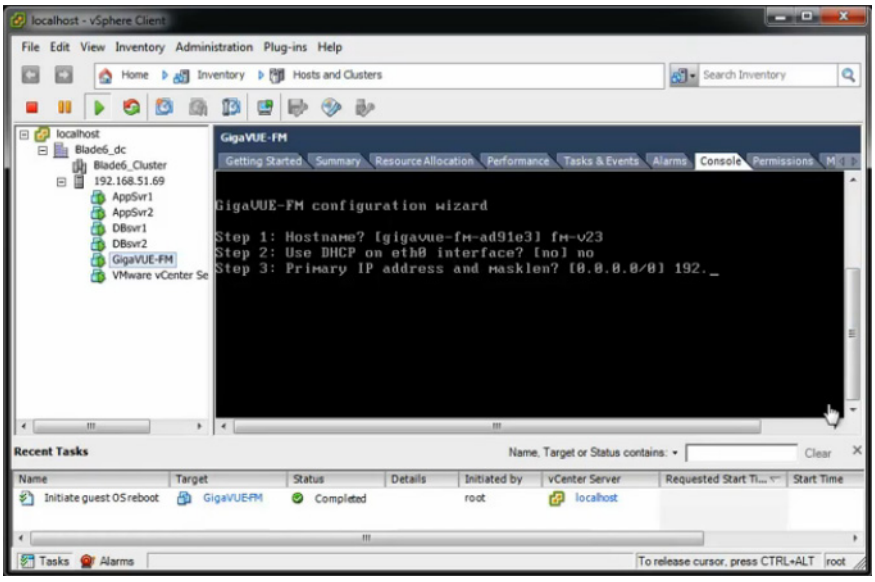

Figure 5-20: Jumpstart Wizard for GigaVUE-FM

6. Provide a unique hostname for GigaVUE-FM. Note that the hostname may contain letters, numbers, periods (**.**), and hyphens (**-**), but may not begin with a hyphen. No other special characters are permitted. The hostname will display as part of the command line prompt after configuration jump-start completes.
7. To enable DHCP on eth0 interface, type yes and press enter.
8. Enter the primary IP address and the mask length, and then press enter.
9. Enter the default gateway and press enter.
10. Enter the primary DNS server address and press enter. Refer to Figure 5-21.

```
Step 1: Hostname? [FM50-22-89]
Step 2: Use DHCP on eth0 interface? [no]
Step 3: Primary IP address and masklen? [10.210.22.89/21]
Step 4: Default gateway? [10.210.16.1]
Step 5: Primary DNS server? [10.10.1.20]
Step 6: Domain name? [gigamon.com]
Step 7: Use DHCP on eth1 interface? [no]
Step 8: eth1 IP address and masklen? [10.210.22.90/21]
Step 9: Admin password (Enter to leave unchanged)?
Step 10: Additional Domain Name Server IP addresses? [10.10.1.21]
Step 11: Additional DNS Domains? [fmqa.com]
Step 12: Enable NTP? [yes]
Step 13: NTP Server IP address?
Step 14: NTP Server version?
```

Figure 5-21: Jump-Start Wizard for Network Interface Configuration

11. (optional) To enable DHCP on eth1 interface, type yes and press enter. Follow steps 7 to 10 to enable DHCP on eth1 interface.
12. Provide an appropriate password for your environment. (Type a password and press **Enter**, or just press **Enter** to leave the password unchanged.)

    NOTE: GigaVUE-FM requires a password.
13. For configuration options:

a. **Additional Domain Name Server IP Addresses?** - the address of any additional name servers required. These must be provided as a set of IP addresses with spaces as shown in the Figure 5-22.

b. **Additional DNS Domains?** - Multiple DNS domains can be defined in the jump start configuration with spaces in between as shown in Figure 5-22.

c. **Enable NTP? [no]** - the default is set to “yes”, the following options are available:

   NTP Server IP Address? - enter the NTP server address

   NTP Server Version? - enter the NTP version number of the NTP server

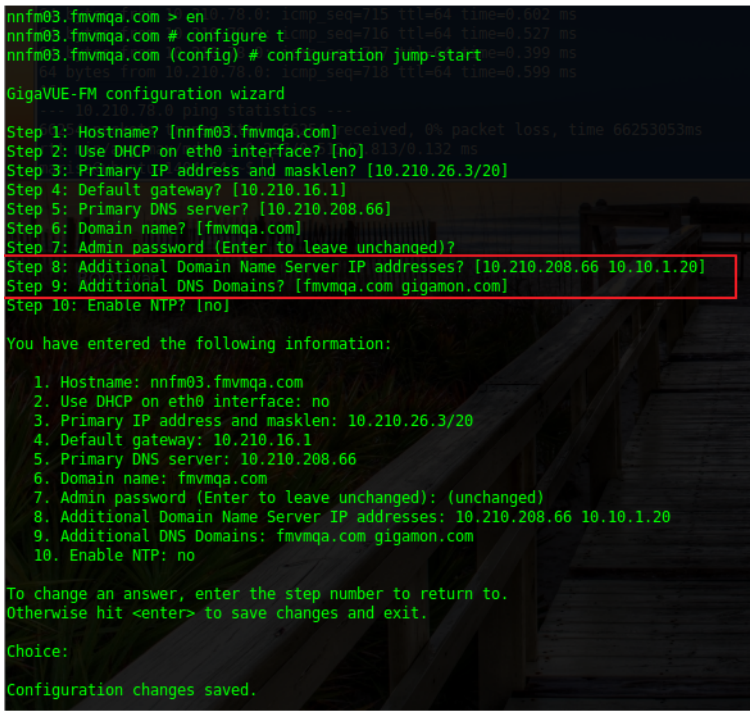

Figure 5-22: Jumpstart Wizard for Additional Domain Name Server IP Addresses

14. The console displays the summary of the chosen selections with instructions on how to make changes, as needed.

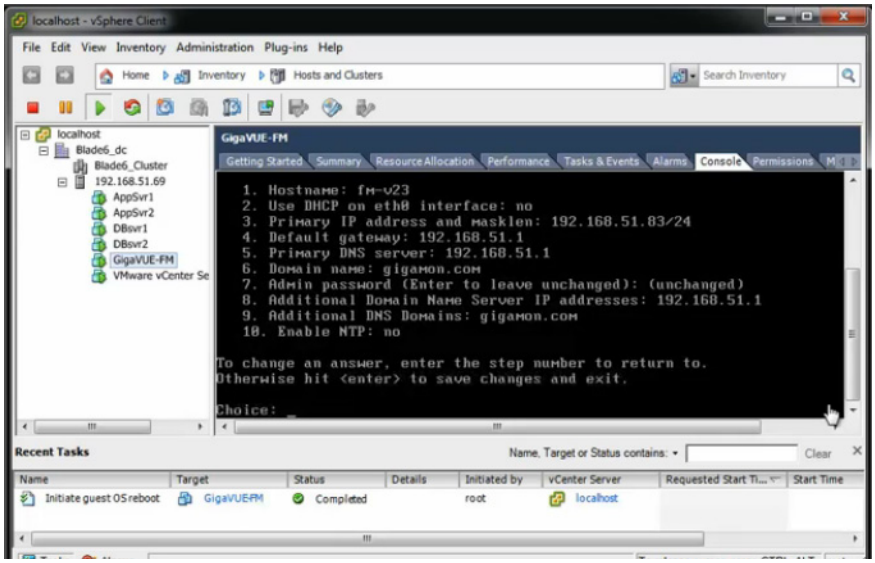

Figure 5-23: Summary of the Selections from Jumpstart Menu

15. Press **Enter** to save your choices and exit the wizard.

The initial configuration is saved and GigaVUE-FM is up and running. GigaVUE-FM is now accessible using a web browser, using IP address specified in the jumpstart steps. Also the first time GigaVUE-FM starts, a EULA is presented. Accept the EULA to continue and see a dashboard similar to the one shown in Figure 5-24.

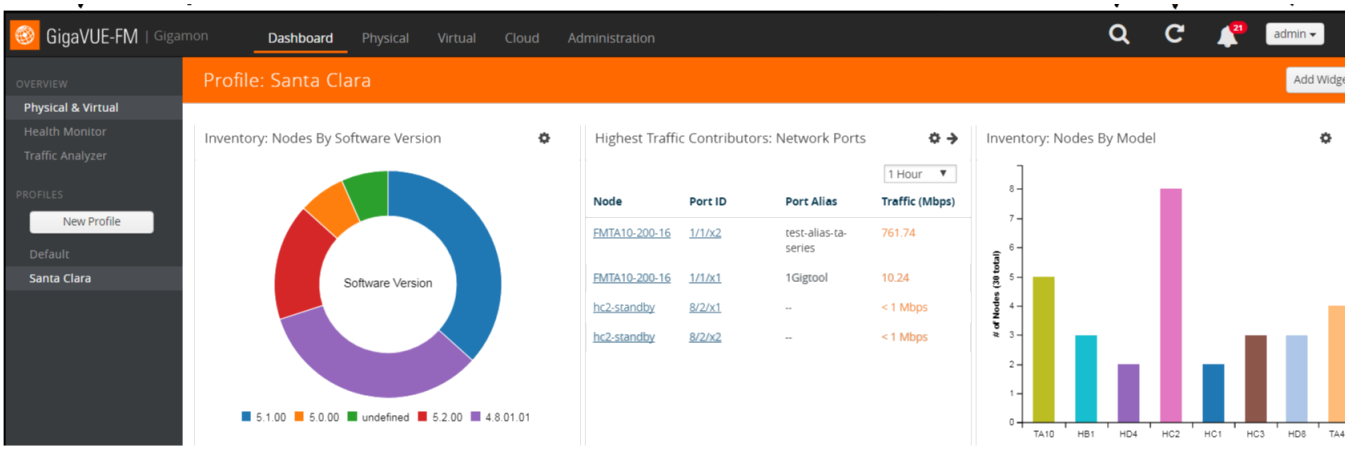

Figure 5-24: Instance of GigaVUE-FM from Web Client

Configure SSH Settings

SSH access is enabled by default on new GigaVUE-FM and GigaVUE-VM deployments. By default, the SSH server runs on port 22.

GigaVUE-FM can be configured to use a custom port for its SSH server with the **ssh server ports <port number>** command followed by a **write memory** command to save the configuration. For example, the following commands change the SSH port number to 2222.

```
(config) # ssh server ports 2222
(config) # write memory
(config) #
```

After making the settings shown above in the GigaVUE-FM CLI, you can connect an SSH session to GigaVUE-FM using the new port number from a web client.

HTTP/HTTPS Ports

GigaVUE-FM uses the following standard HTTP configuration shown in Figure 5-25

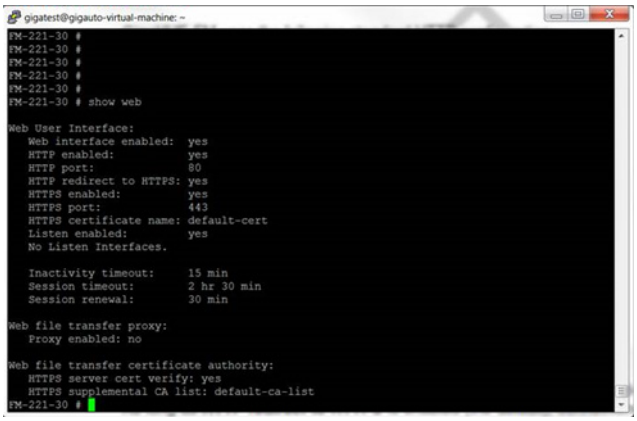

Figure 5-25: GigaVUE-FM CLI Screen to Configure Web Client

HTTPS port can be changed for GigaVUE-FM but the HTTP port is hard-coded to 80. As long as **HTTP redirect to HTTPS** is enabled (the default), connections to the fixed HTTP port of 80 will redirect to whatever the configured HTTPS port is.

Install Third-Party Certificate

Use the following procedure to install a third-party certificate on GigaVUE-FM:

1. Generate a certificate and a private key file in pem format. Use the following command on Linux or a Linux app (such as Cygwin) for generating the files:

   openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout privatekey.pem -out certificate.pem

2. Copy the contents of the certificate file. You will use the contents in Step 5.

NOTE: When you copy the contents, do not delete the EOL characters at the start of each line.

3. Log in to the GigaVUE-FM CLI.
4. Enable configuration mode by entering the following commands:

   > en
   # conf t

5. Use the **crypto certificate** command to add a new certificate. In the following example, the certificate is named “my-cert” an and the contents of the public certificate pem file is copied inside the quotes.

   (config) # crypto certificate name my-cert public-cert pem “<contents-of-public-certificate-pem>”

6. Copy the contents of the private key file. You will use the contents in Step 7.
7. Use the following command to add the private key, which is “my-cert” in this example. The contents of the private key pem file is copied between the quotes.

   (config) # crypto certificate name my-cert private-key pem “<contents-of-private-key-pem-file>”

   The private key file and certificate are installed and ready to use.

8. Set the certificate (“my-cert” in this example) to be the default self-signed certificate by using the following command:

(config) # crypto certificate default-cert name my-cert

The system will now start using the newly installed certificate.

Install Third-Party Certificate on GigaVUE-FM in AWS

Use the following procedure to install a third-party certificate on GigaVUE-FM that is hosted in AWS:

1. Log in to GigaVUE-FM.
2. Execute the following steps from the shell prompt as a root user (**sudo**):

- Replace SSLCertificateFile: **/etc/pki/tls/certs/localhost.crt**
- Replace SSLCertificateKeyFile: /etc/pki/tls/private/localhost.key
- Provide access to certificate and key files: chmod 777
- Restart apache as root systemctl restart hhtps.

6 Install GigaVUE-FM on MS Hyper-V

This section describes how to install and configure GigaVUE-FM in a Microsoft Hyper-V environment. It consists of the following main sections:

- *System Requirements* on page 89 describes the hardware requirements.
- *Install GigaVUE-FM for Microsoft Hyper-V* on page 91 describes the steps to install and deploy GigaVUE-FM.
- *Initial GigaVUE-FM Configuration* on page 99 describes the steps to start GigaVUE-FM instance and configure it.
- *Configure SSH Settings* on page 100 describes the SSH settings
- *HTTP/HTTPS Ports* on page 101 describers how to setup the HTTP client

System Requirements

This section describes the hardware and virtual computing requirements for GigaVUE-FM. Before installing GigaVUE-FM, ensure that a supported version of Windows Server is installed on hardware that meets minimum requirements (see *Windows Server Hardware Requirements* on page 89 for hardware requirements). Ensure that the GigaVUE-FM time is set correctly to ensure accuracy of the trending data that is captured.

The Hyper-V implementations provided by the following Windows Server versions have been tested and found to operate acceptably with GigaVUE-FM:

- Windows Server 2012 R2 and later

Windows Server Hardware Requirements

The following table describes the minimum requirements for the hardware on which Microsoft Hyper-V runs GigaVUE-FM.

Minimum Hardware Requirements	
Hypervisor	Microsoft Hyper-V

Minimum Hardware Requirements	
CPU	One or more 64-bit x86 CPUs with virtualization assist (Intel-VT or AMD-V) enabled. **Note**: To run GigaVUE-FM, hardware support for virtualization must be enabled. Make sure that the BIOS option for virtualization support is not disabled. For more information, see your BIOS documentation
RAM	At least 8GB
Disk Space	At least 41GB shared (FC, iSCSI, NFS, or FCoE) or locally attached storage (PATA, SATA, SCSI)
Network	At least one 1Gb NIC

NOTE: Refer to the Microsoft documentation for information on enabling Hyper-V.

The following table lists the virtual computing resources that the Windows Server must provide for each GigaVUE-FM instance.

Minimum Virtual Computing Requirements	
Memory	Minimum 8GB memory
Virtual CPU	2 vCPU
Virtual Storage for Guest	41GB using Virtual IDE (the Hyper-V default)
Virtual Network Interfaces	1 vNIC using Hyper-V Virtualized NIC (the Hyper-V default)

Supported Browsers

GigaVUE-FM v3.5 has been tested on the following browsers:

Browser	Version
Mozilla Firefox™	• Version 47
Windows® Internet Explorer®	• Version 11
Apple® Safari®	• Version 9.1
Google®Chrome®	• Version 52
Microsoft Edge	• Version 38

Notes:

- Only the browsers that support TLS v1.2 can access GigaVUE-FM.
- DNS prefetch is a known limitation of Internet Explorer 11. If GigaVUE-FM is configured with DNS and you are using Internet Explorer 11, every new screen can be slowed significantly. If a direct IP address is used instead of a DNS name, the UI response is similar to other browsers. It is recommended that you use the GigaVUE-FM IP when using Internet Explorer 11 or use either a FireFox or Chrome browser instead.
- IE11 Compatibility view mode is not supported.

Install GigaVUE-FM for Microsoft Hyper-V

The GigaVUE-FM software package for Microsoft Hyper-V environments is distributed as an **ISO image** file. The following sections describes how to deploy a fresh installation of GigaVUE-FM on a Hyper-V host and perform its initial configuration:

- *Install GigaVUE-FM from an ISO Image File*
- *Initial GigaVUE-FM Configuration* on page 99

Install GigaVUE-FM from an ISO Image File

Use the Hyper-V Manager to install the GigaVUE-FM ISO image file.

NOTE: The ISO image file must be stored in a location that is accessible to the Hyper-V Manager.

To create the Virtual Machine for GigaVUE-FM in Microsoft Hyper-V:

1. Open Hyper-V Manager by clicking **Start > Administrative Tools > Hyper-V Manager**.
2. From the Actions pane, click **New > Virtual Machine**. Refer to Figure 6-1.

   The **New Virtual Machine Wizard** opens.

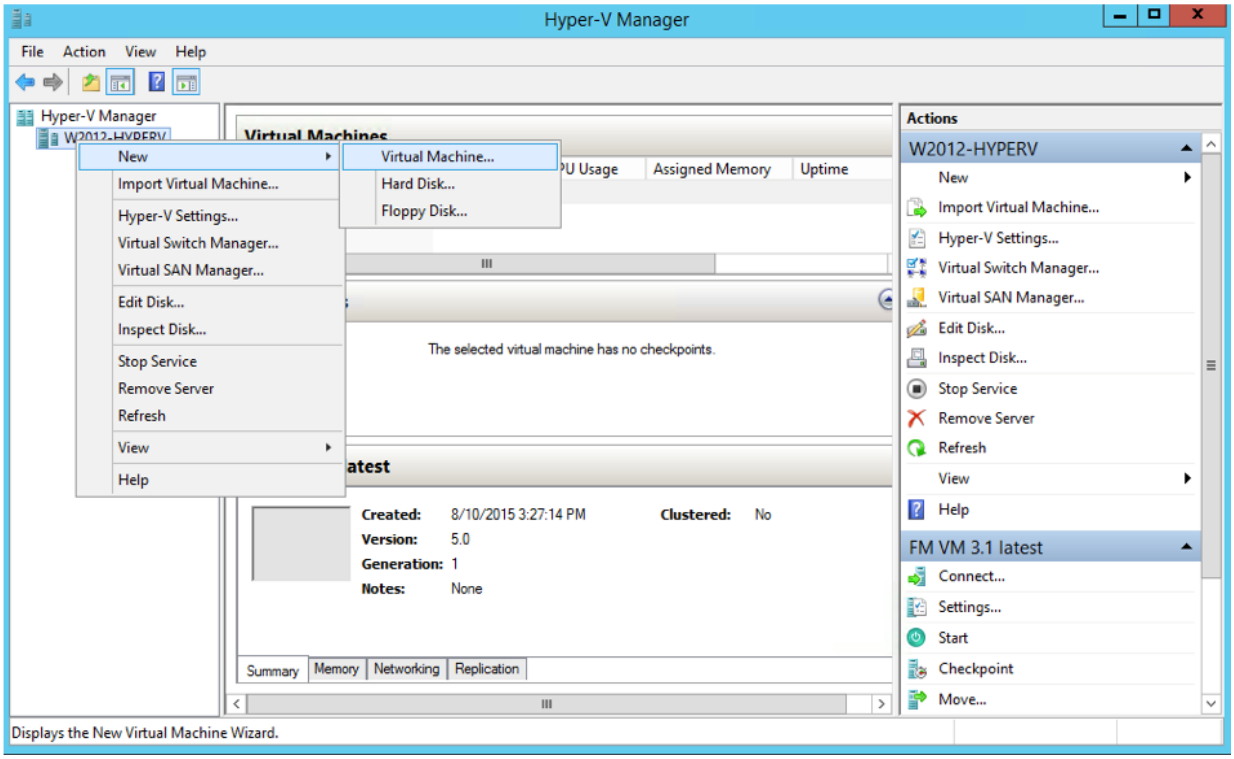

Figure 6-1: Opening the Virtual Machine Wizard

3. Read the notes on the Before You Begin screen (refer to Figure 6-2), and then click **Next** to continue.

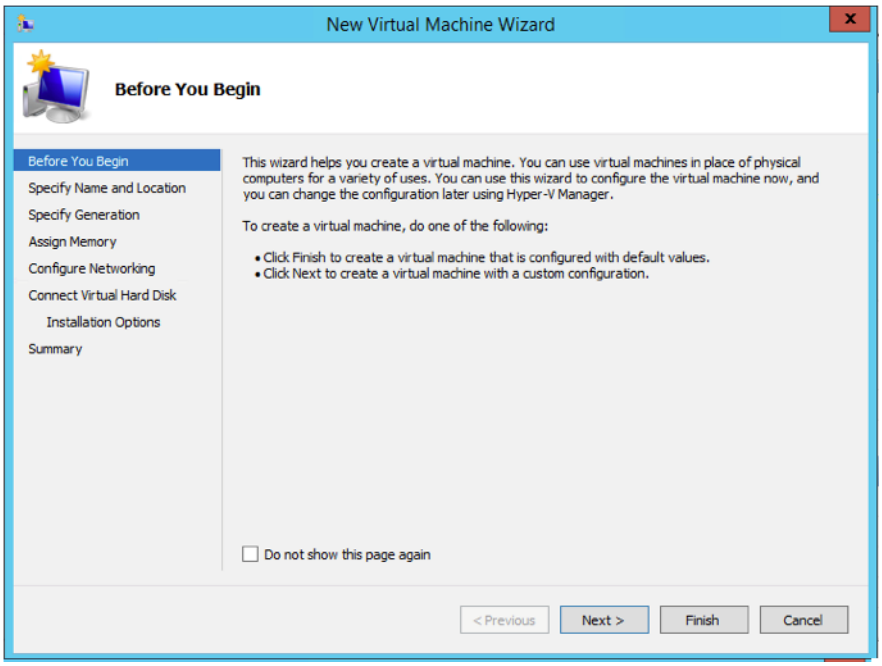

Figure 6-2: Before You Begin Screen

4. After the **Specify Name and Location** page of the **New Virtual Machine Wizard** opens, which is shown in Figure 6-3, do the following:

   a. Supply a descriptive name for the GigaVUE-FM virtual machine in the **Name** field.

      By default, the virtual machine will be stored in the default configuration folder shown in the **Location** text box. You can change this default location by checking the **Store the virtual machine in a different location** checkbox and providing a custom path.

b. Select **Next** to continue.

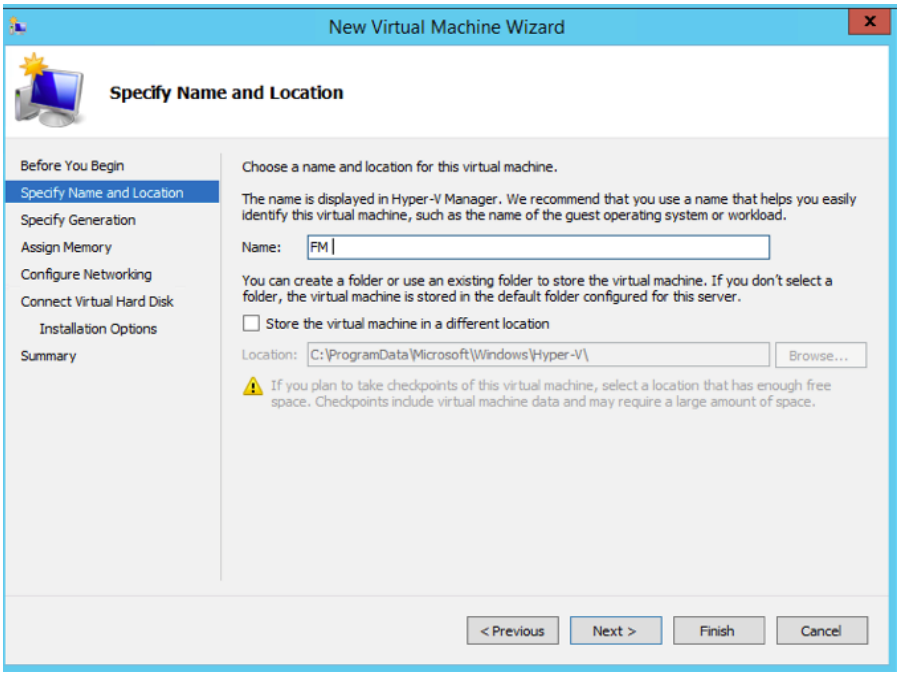

Figure 6-3: Name and Location Page

5. On the **Specify Generation,** select **Generation 1** as shown in Figure 6-4

   It is important to select Generation 1 and not Generation 2. Selecting Generation 2 may lead to failure of the GigaVUE-FM installation process because the CD Drive is presented as an SCSI device and not IDE.

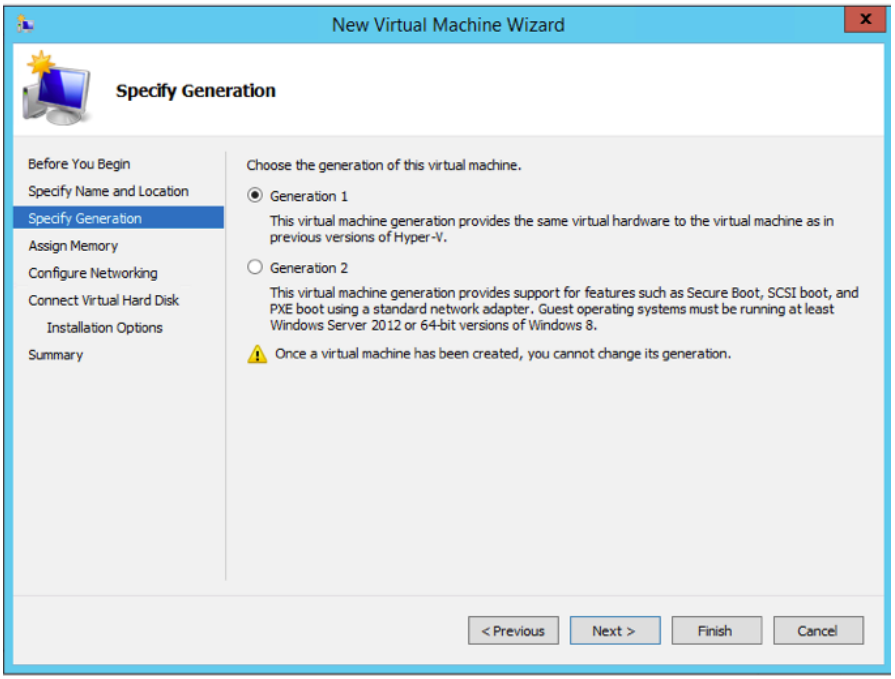

Figure 6-4: Specify Generation Page

6. Click **Next** to continue.

   The **Assign Memory** page of the **New Virtual Machine Wizard** opens, which is shown in Figure 6-5.

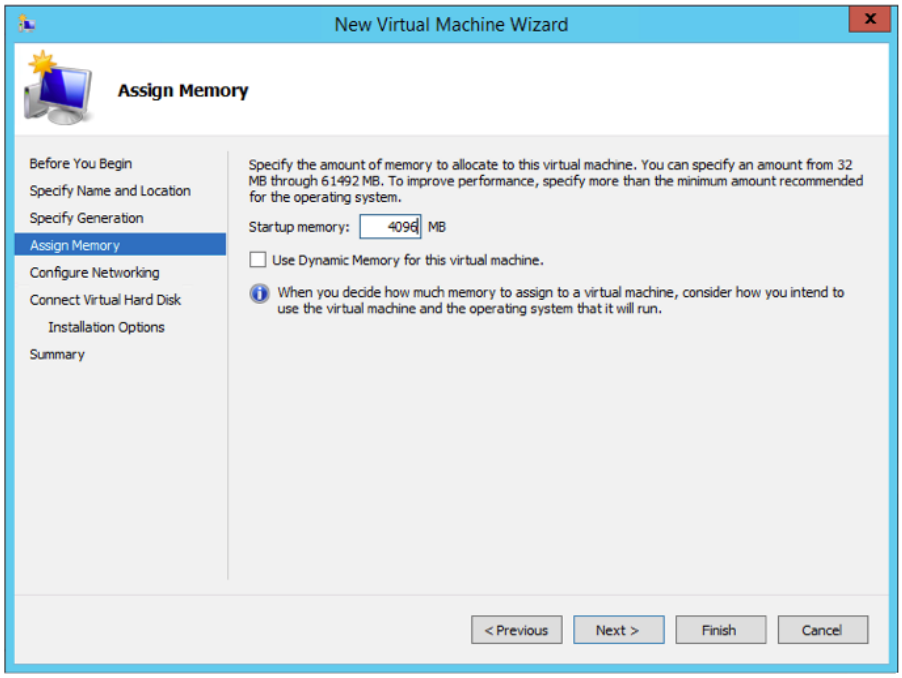

Figure 6-5: Assign Memory Page

7. Change the **Memory** assigned to this virtual machine to **4096** MB, and then click **Next** to continue.

   The **Configure Networking** page of the **New Virtual Machine Wizard** opens, which is shown in Figure 6-6.

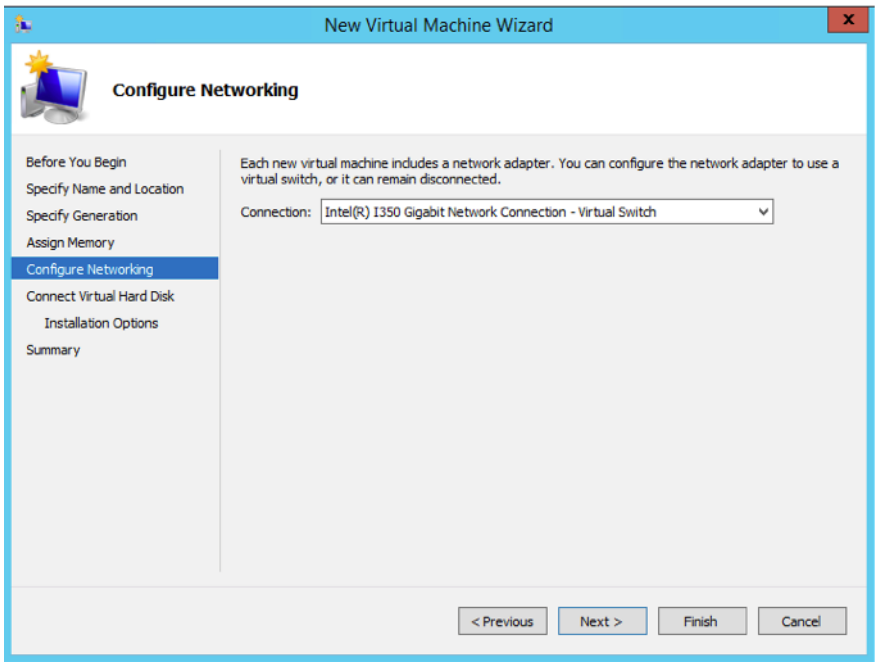

Figure 6-6: Configure Networking Page

8. Choose the virtual network to which GigaVUE-FM will connect from the drop-down list, and then click **Next** to continue.

   The **Connect Virtual Hard Disk** page of the **New Virtual Machine Wizard** opens, which is shown in Figure 6-7.

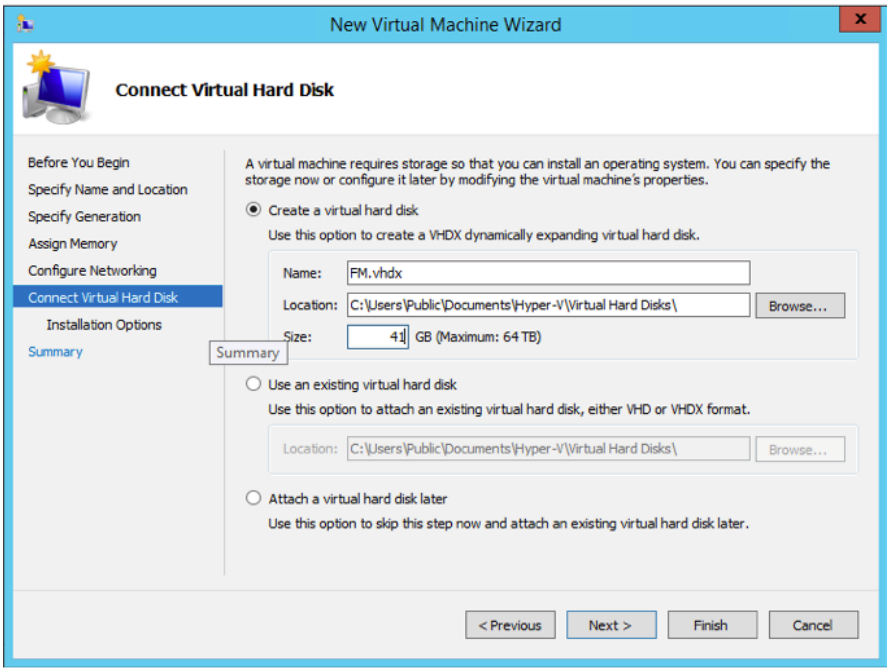

Figure 6-7: Connect Virtual Hard Disk Page

9. Select the **Create a virtual hard disk** option and set the **Size** to **41 GB**.

   You can accept the default **Name** and **Location** or customize them according to your needs. When you have finished, click **Next** to continue.

   The **Installation Options** page of the **New Virtual Machine Wizard** opens., which is shown in Figure 6-8.

10. Use this dialog box to select the ISO image file for GigaVUE-FM. As shown in the figure Figure 6-8, set the following options:

    a. Select the option **Install an operating system from a boot CD/DVD-ROM**.

    b. Set the **Media** option to **Image file (.iso)**.

    c. Use the **Browse** button to navigate to the GigaVUE-FM ISO image file.

d. Click **Next** to continue.

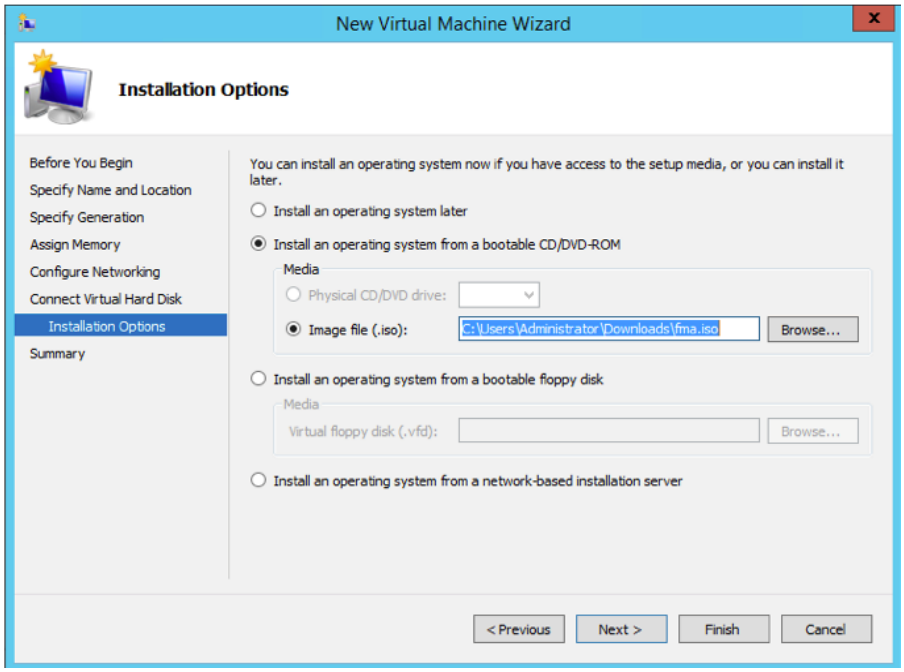

Figure 6-8: Install Options Page

After clicking **Next**, the summary page of the **New Virtual Machine Wizard** opens, showing the settings that you configured for the GigaVUE-FM virtual machine. An example is shown in Figure 6-9.

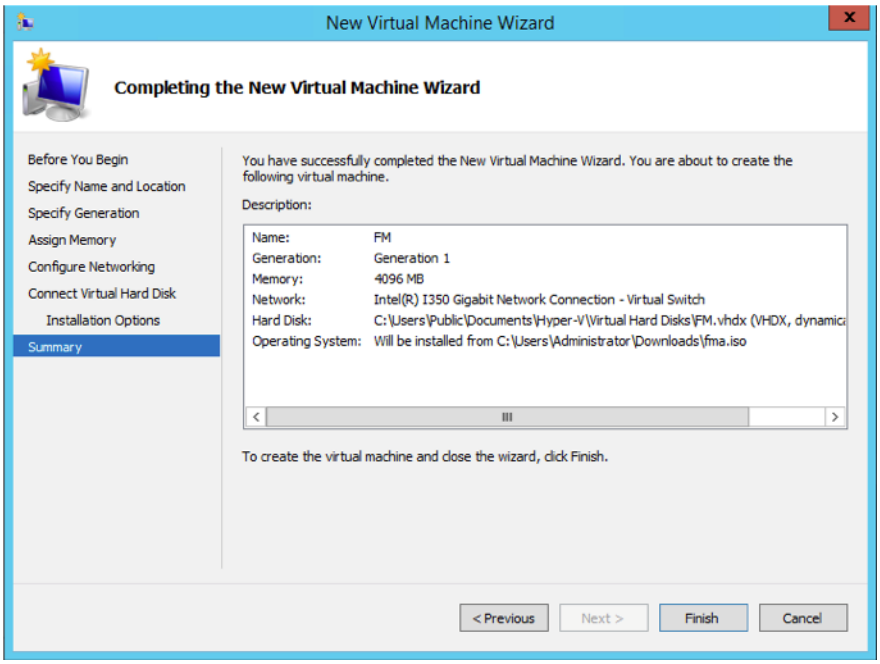

Figure 6-9: Summary Page

11. Click **Finish** to create the GigaVUE-FM virtual machine as configured. Alternatively, you can use the **Previous** button to go back and change settings.

    The New Virtual Machine Wizard only creates the GigaVUE-FM virtual machine, returning you to the Hyper-V Manager when creation is complete. The new GigaVUE-FM virtual machine is listed in the **Action** pane by the name you supplied during installation. The rest of this procedure will take you through the steps of actually installing GigaVUE-FM.

Connect and Power On the GigaVUE-FM Virtual Machine

The next step is to connect to the GigaVUE-FM virtual machine from within Hyper-V Manager and start it. This begins the actual installation of the GigaVUE-FM Virtual Appliance from the connected ISO image file. Once GigaVUE-FM finishes installing from the ISO image file, you will then disconnect the ISO image file and restart the virtual machine.

1. In Hyper-V Manager, locate the **Virtual Machines** entry in the results pane, right-click the GigaVUE-FM virtual machine, and click the **Connect** option.
2. The Virtual Machine Connection tool opens for the GigaVUE-FM virtual machine.

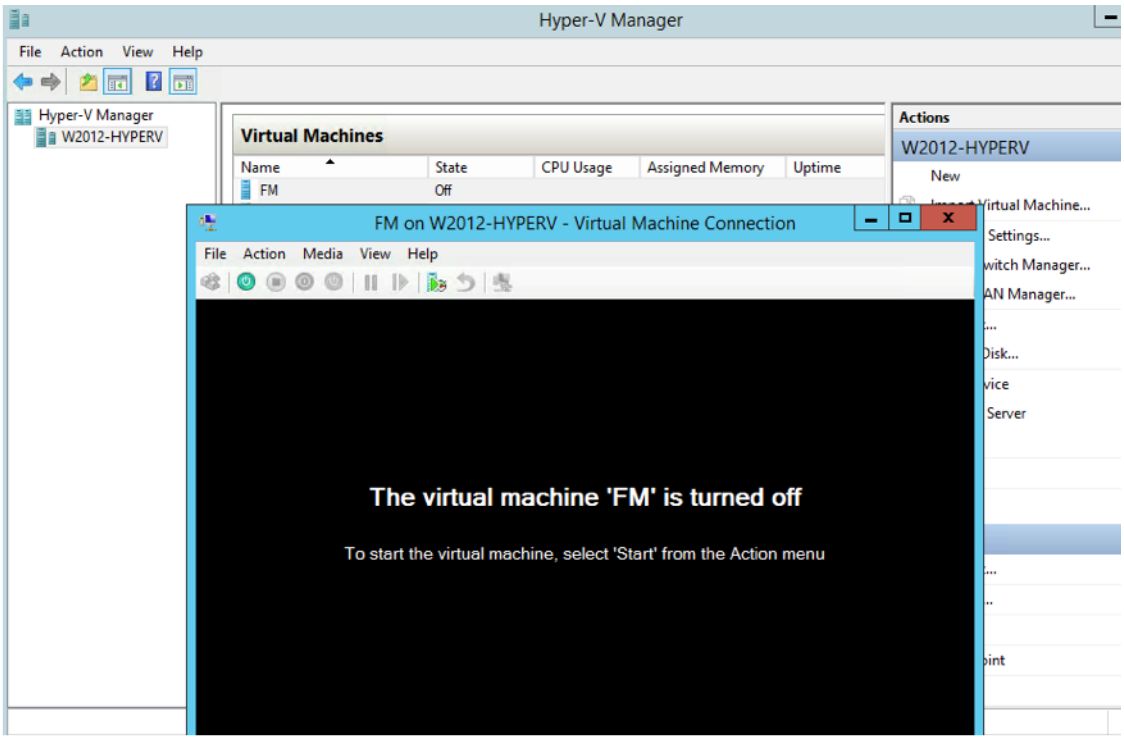

3. Select **Action > Start** to start the GigaVUE-FM virtual machine from the Virtual Machine Connection tool.

4. The GigaVUE-FM virtual machine powers on. You can monitor the progress of the system start in the Virtual Machine Connection tool.

   The system power-on can take several minutes as GigaVUE-FM is installed from the ISO image file. Disconnect the media before Powering On as shown below:

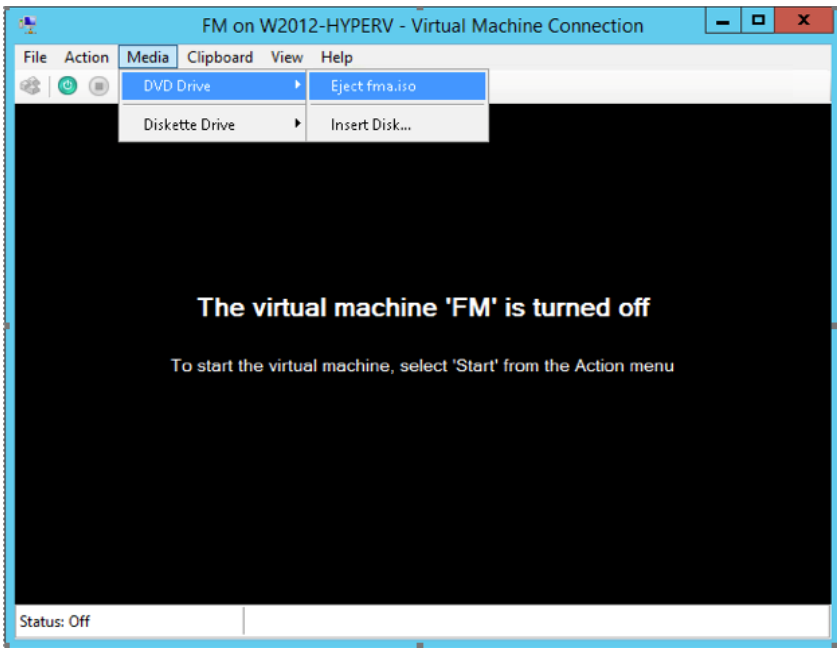

Disconnect the ISO Image File

It is important to disconnect the ISO image file before you power on GigaVUE-FM again so you don't have to go through the image install process again. Disconnect the ISO image file as follows:

1. In Hyper-V Manager, locate the **Virtual Machines** entry in the results pane and select the GigaVUE-FM virtual machine.
2. In the **Actions** pane, click the **Settings** entry under the GigaVUE-FM virtual machine name.

   A Settings dialog box for the GigaVUE-FM virtual machine appears.
3. Select the **DVD Drive** entry in the panel on the left of the Settings dialog box and change its setting from **Image file** to **None**, as shown in the figure below.
4. Click the **OK** button to apply the changes.

This concludes the installation procedure for GigaVUE-FM on Hyper-V. The next step is to power on the virtual machine and perform its initial configuration, as described in *Initial GigaVUE-FM Configuration* on page 99.

IMPORTANT: Clear the browser cache before logging in to GigaVUE-FM!

Initial GigaVUE-FM Configuration

After you have deployed a new GigaVUE-FM instance, you need to perform an initial configuration before you can start using GigaVUE-FM. This procedure only needs to be performed once for each GigaVUE-FM instance deployed.

NOTE: Use Care When Shutting Down or Rebooting a GigaVUE-FM. **Never** directly Power-Off the virtual machine. For Microsoft Hyper-V environment, you cannot use any of the reset, or turn-off hooks. Using either of these may lead to corruption that will prevent proper GigaVUE-FM operation.

The best ways to **shutdown** a GigaVUE-FM on Hyper-V is to use either Shutdown or Ctrl+Alt+Del from the **Action** button on the virtual console.

To perform the initial configuration:

1. Open Hyper-V Manager by clicking **Start > Administrative Tools > Hyper-V Manager**.
2. Make sure you have already disconnected the ISO image file used to install GigaVUE-FM. Refer to *Disconnect the ISO Image File* on page 98 for details.
3. Locate the **Virtual Machines** entry in the results pane, right-click the GigaVUE-FM virtual machine, and click the **Connect** option.

   The Virtual Machine Connection tool opens for the GigaVUE-FM virtual machine.
4. Select **Action > Start** to start the GigaVUE-FM virtual machine from the Virtual Machine Connection tool.

   The GigaVUE-FM virtual machine powers on and displays a login prompt.
5. Log in as **admin** with password **admin123A!**

   The configuration wizard starts automatically, as shown in the following figure.

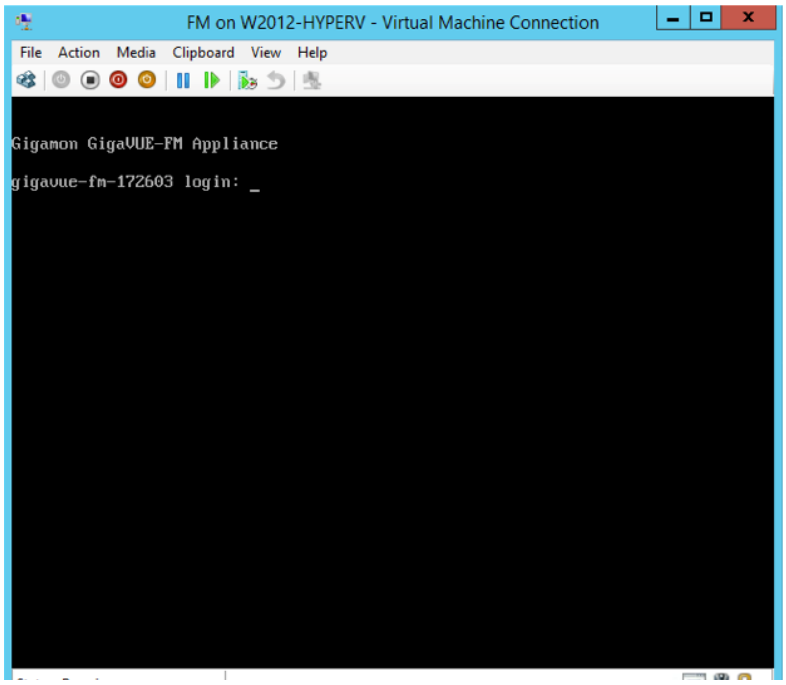

6. At this point, the wizard presented a series of steps for you to provide the initial configuration for GigaVUE-FM. These are the steps

   a. Provide a unique hostname for GigaVUE-FM. Note that the hostname may contain letters, numbers, periods (**.**), and hyphens (**-**), but may not begin with a hyphen. No other special characters are permitted.

   b. Decide whether to use DHCP for the management interface.

      If you choose **no**, you will be prompted to provide the following:

      - **IPv4** address and **masklen**

      - Default gateway

      - Primary DNS server

      - Domain name

      If you choose yes, skip to Step c

   c. If you choose Yes for Step b, follow these instructions. The same options are repeated if DHCP is selected as No, but only one DNS IP address and domain server can be listed.

      For configuration options:

      •**Additional Domain Name Server IP Addresses?** - the address of any additional name servers required must be provided as a set of IP addresses with spaces as shown in the following figure.

      •**Additional DNS Domains**? - Multiple DNS domains can be defined in the jump start configuration with spaces in between as shown in the following figure.

      •**Enable NTP**? [yes] - the default is set to “yes”. The following options are available:

         •**NTP Server IP Address?** - enter the NTP server address

         •**NTP Server Version**? - enter the NTP version number of the NTP server

7. Provide an appropriate password for your environment. (Type a password and press **Enter**, or just press **Enter** to leave the password unchanged.)

   NOTE: Blank passwords are not permitted.

   The console displays your selections with instructions on how to make changes, if necessary.

8. Press **Enter** to save your choices and exit the wizard.

9. Your initial configuration is saved and GigaVUE-FM is up and running. You should now be at a standard mode command prompt.

   You can now access GigaVUE-FM by opening a browser and entering its IP address (the IP address you specified).

Configure SSH Settings

SSH access is enabled by default on new GigaVUE-FM deployments. You can enable SSH from the CLI using the **ssh server enable** command. By default the SSH server runs on port 22.

You can configure GigaVUE-FM to use a custom port for its SSH server with the **ssh server ports <port number>** command followed by a **write memory** command to save the configuration. For example, the following CLI commands change the SSH port number to 2222.

```
(config) # ssh server ports 2222
(config) # write memory
(config) #
```

After making the settings shown above with the GigaVUE-FM CLI command, you can connect an SSH session to GigaVUE-FM using the new port number.

HTTP/HTTPS Ports

GigaVUE-FM uses the standard HTTP configuration shown in the following figure:

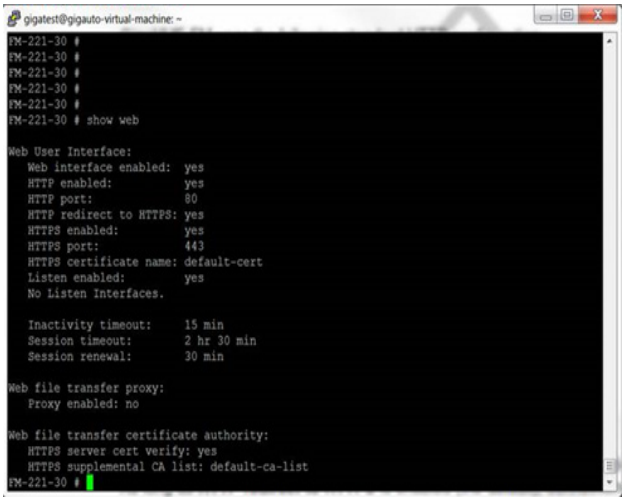

In this release, you can change the HTTPS port for GigaVUE-FM but the HTTP port is hard-coded to 80. As long as **HTTP redirect to HTTPS** is enabled (the default), connections to the fixed HTTP port of 80 will redirect to whatever the configured HTTPS port is.

Make Sure the Web Server is Enabled on Nodes to be Managed

GigaVUE-FM can only discover and manage nodes with their web servers enabled and operating on the default HTTP port of 80. Both G and H Series nodes have their web servers enabled by default. However, if you disabled a node's web server or changed its HTTP port, you will need to restore the settings before GigaVUE-FM can manage it.

GigaVUE-FM can manage nodes operating on custom HTTPS ports. Incoming HTTP connections redirect to the custom HTTPS port.

The **show web_server** and **show web** output listed below summarizes the necessary HTTP settings for G and H Series GigaVUE nodes managed by GigaVUE-FM. The items shown in red are required settings.

G Series

```
G Series>show web_server
Admin      : 1 (Must be enabled)
Operation  : 1
HTTP port  : 80 (Must remain at its default setting of 80)
HTTPS port : 8000 (Can be set to any custom value; HTTP redirects here)
Timeout    : 5 (minutes)
HTTPS Cert : Default Certificate
```

H Series

```
H Series (config) # show web
Web-based management console enabled: yes
HTTP enabled:           yes (Must be enabled)
HTTP port:              80 (Must remain at its default setting of 80)
HTTP redirect to HTTPS: yes (Must remain enabled)
HTTPS enabled:          yes (Must remain enabled)
HTTPS port:             443 (Can be set to any custom value; HTTP                    redirects here)
Listen enabled:         yes
No Listen Interfaces.
Inactivity timeout:     15 min
Session timeout:        2 hr 30 min
Session renewal:        30 min
Web proxy enabled: no
```

7 Install GigaVUE-FM on KVM

This section describes how to install and configure GigaVUE-FM in a KVM environment. It consists of the following main sections:

- *System Requirements* on page 103 describes the hardware requirements.
- *Install GigaVUE-FM for KVM* on page 105 describes the steps to install and deploy GigaVUE-FM.
- *Initial GigaVUE-FM Configuration* on page 109 describes the steps to start GigaVUE-FM instance and configure it.
- *Configure SSH Settings* on page 110 describes the SSH settings.
- *HTTP/HTTPS Ports* on page 111 describers how to setup the HTTP client.

Limitations

You can install GigaVUE-FM in a KVM environment, but you cannot access GigaVUE-FM through CLI in a KVM environment using SSH. In KVM, you can only access the GigaVUE-FM CLI using the VNC console.

System Requirements

This section describes the hardware and virtual computing requirements for GigaVUE-FM. Ensure that the GigaVUE-FM time is set correctly to ensure accuracy of the trending data that is captured.

Linux Server Hardware Requirements

The following table describes the minimum requirements for the hardware on which KVM runs GigaVUE-FM.

Minimum Hardware Requirements	
Hypervisor	KVM **Supported** (tested on previous versions of GigaVUE-FM) • v2.0.0

Minimum Hardware Requirements	
CPU	One or more 64-bit x86 CPUs with virtualization assist (Intel-VT or AMD-V) enabled. **Note**: To run GigaVUE-FM, hardware support for virtualization must be enabled. Make sure that the BIOS option for virtualization support is not disabled. For more information, see your BIOS documentation.
RAM	At least 8GB.
Disk Space	At least 40GB shared (FC, iSCSI, NFS, or FCoE) or locally attached storage (PATA, SATA, SCSI).
Network	At least one 1Gb NIC.

The following table lists the virtual computing resources that the Linux Server must provide for each GigaVUE-FM instance.

Minimum Virtual Computing Requirements	
Memory	Minimum 8GB memory
Virtual CPU	2 vCPU
Virtual Storage for Guest	41GB
Virtual Network Interfaces	One vNIC

Supported Browsers

GigaVUE-FM has been tested on the following browsers:

Browser	Version
Mozilla Firefox™	• Version 47
Windows® Internet Explorer®	• Version 11
Apple® Safari®	• Version 9.1
Google®Chrome®	• Version 52
Microsoft Edge	• Version 38

Notes:

- Only the browsers that support TLS v1.2 can access GigaVUE-FM.
- DNS prefetch is a known limitation of Internet Explorer 11. If GigaVUE-FM is configured with DNS and you are using Internet Explorer 11, every new screen can be slowed significantly. If a direct IP address is used instead of a DNS name, the UI response is similar to other browsers. It is recommended that you use the GigaVUE-FM IP when using Internet Explorer 11 or use either a FireFox or Chrome browser instead.
- IE11 Compatibility view mode is not supported.

Install GigaVUE-FM for KVM

The GigaVUE-FM software package for KVM environments is distributed as an **ISO image** file. The following sections describes how to deploy a fresh installation of GigaVUE-FM on a KVM host and perform its initial configuration:

- *Install GigaVUE-FM from an ISO Image File*
- *Initial GigaVUE-FM Configuration* on page 109

These steps are only valid for new installations of GigaVUE-FM.

Install GigaVUE-FM from an ISO Image File

Use the Virtual Machine Manager to install the GigaVUE-FM ISO image file.

NOTE: The ISO image file must be stored in a location that is accessible to the Manager.

To create the Virtual Machine for GigaVUE-FM in KVM:

NOTE: These instructions use “Virtual Machine Manager” to create and manage the virtual machines (VMs).

1. Open the Virtual Machine Manager by using **virt-manager** from the command line. Select **Create a new virtual machine**. The **New Virtual Machine Wizard** opens.

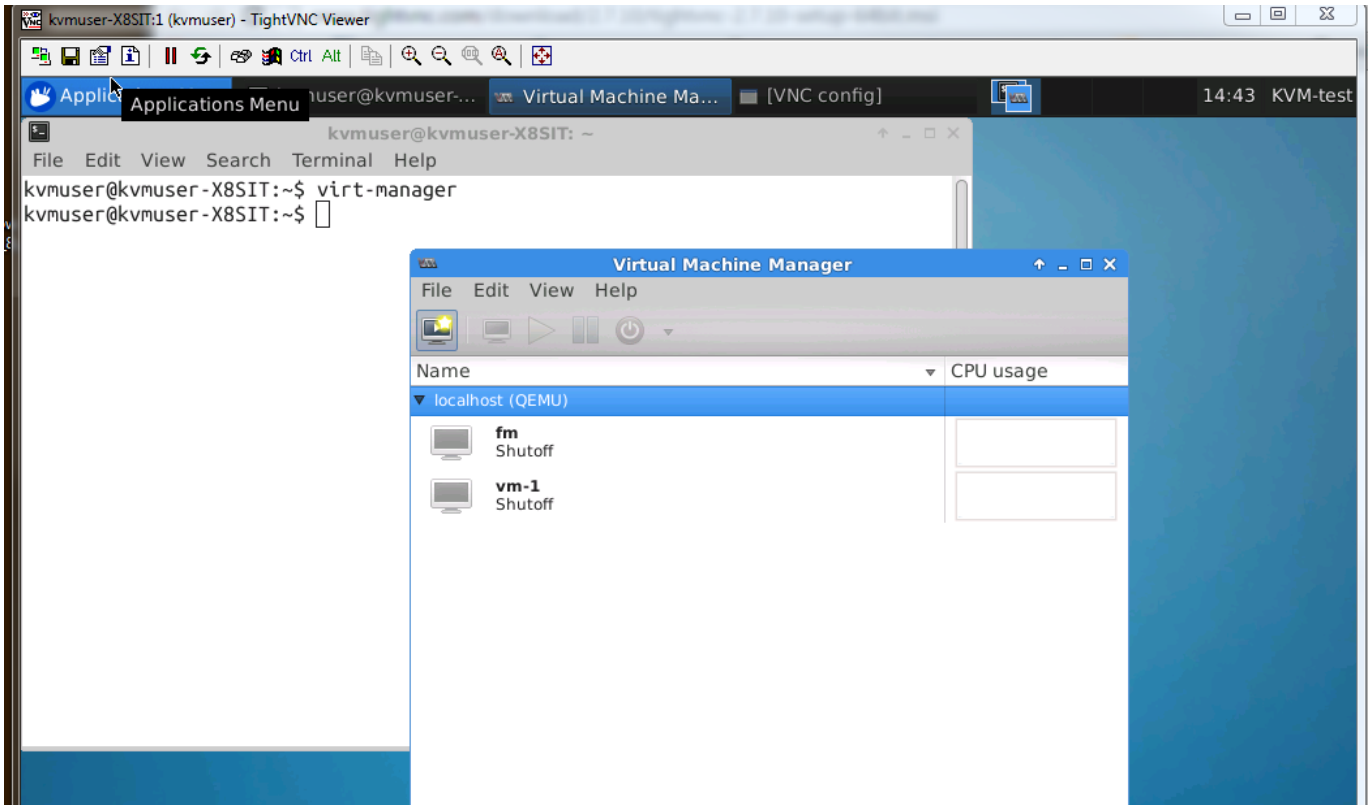

2. The **Specify Name and Location**.

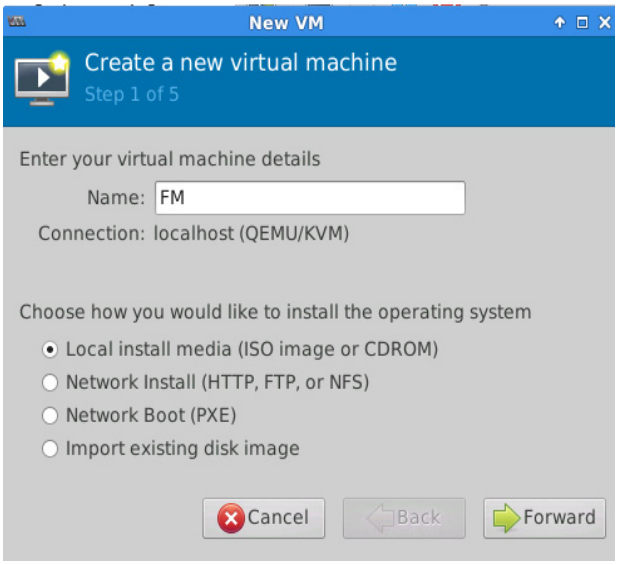

It is recommended to supply a descriptive name for the GigaVUE-FM virtual machine in the **Name** field.

Click **Forward** to continue.

3. Enter the Location from where to upload the GigaVUE-FM iso image and choose the OS type and Version.

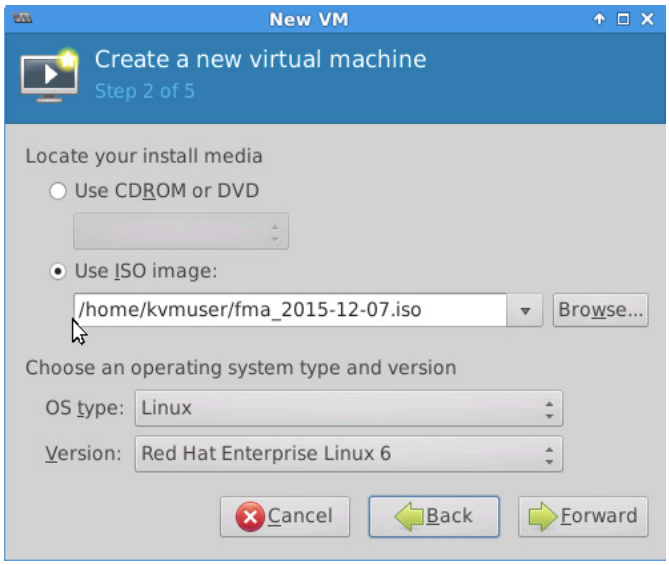

Click **Forward** to continue.

4. Set the **Memory** and **CPU** Settings.

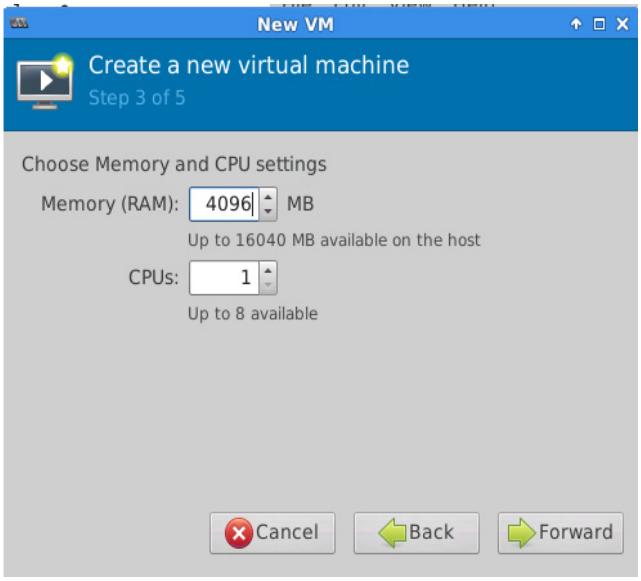

Click **Forward** to continue.

5. **Enable the storage for this virtual machine** option and set the **Size** to **41 GB**. Select **Allocate entire disk now.** When you have finished, click **Forward** to continue.

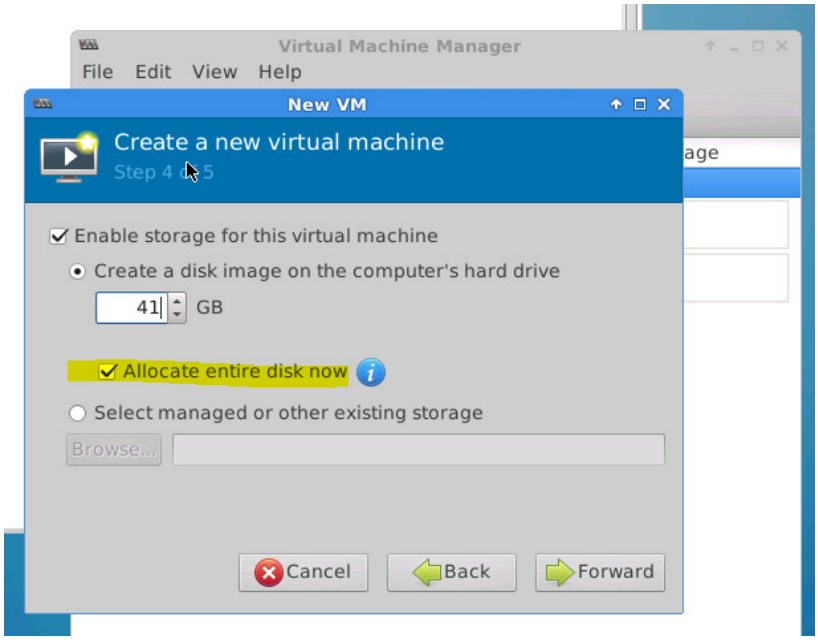

6. On the next screen, select Customize configuration before install.
   - Use the Advanced options to set the VM connection to the network adapter.
   - Connect the VM to the network that you have configured on your hypervisor that ensures the network connectivity to your managed VMs.
   - Ensure that the Virt Type is set to KVM

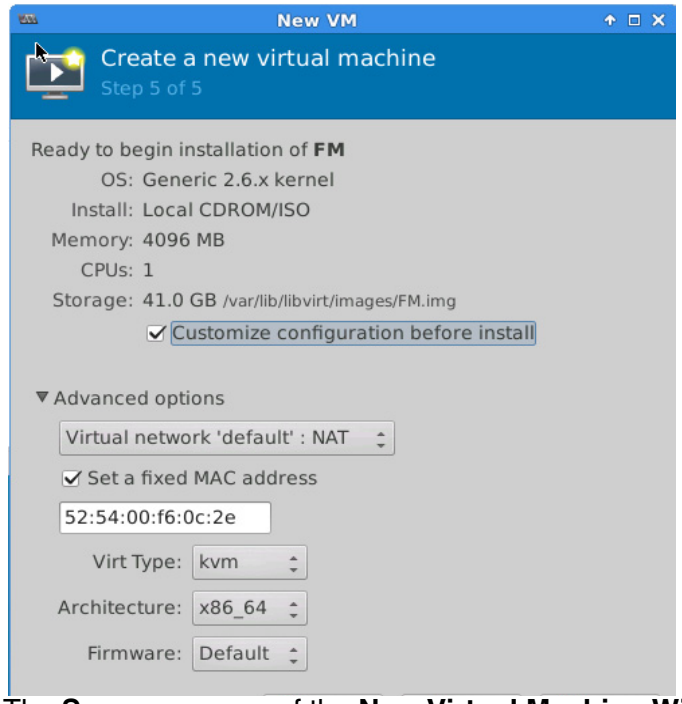

7. The **Summary** page of the **New Virtual Machine Wizard** opens, showing the settings you have configured for the GigaVUE-FM virtual machine.

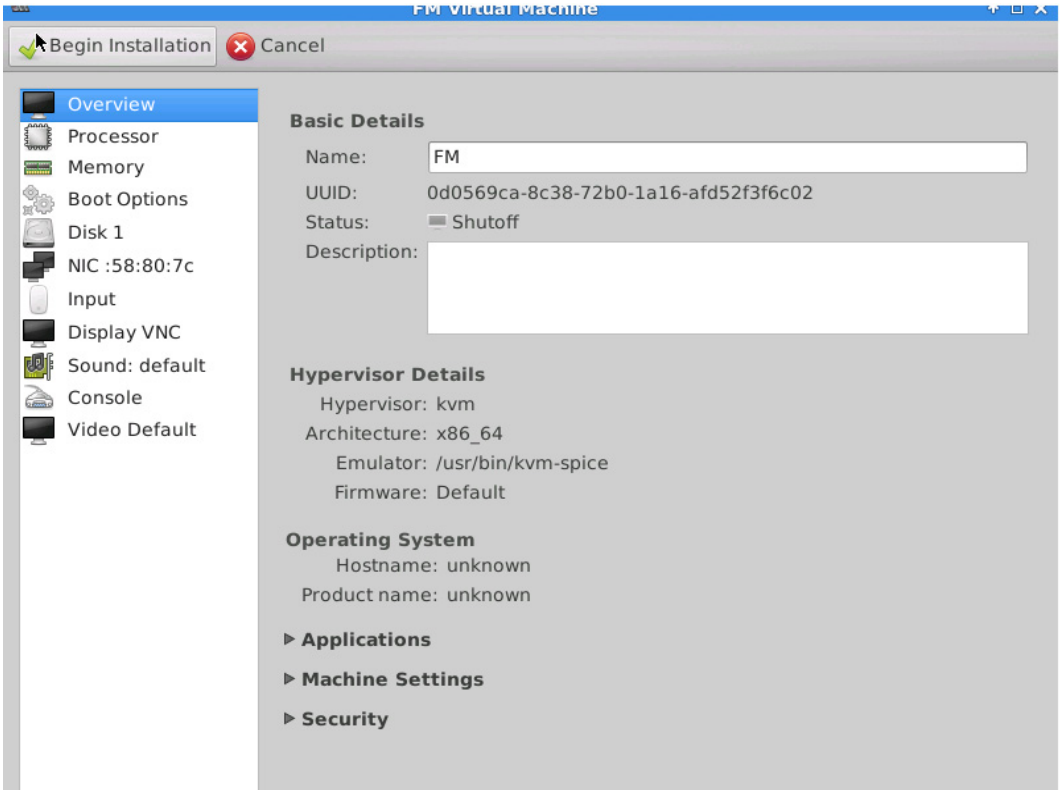

8. Connect and power on the GigaVUE-FM Virtual Machine.

Initial GigaVUE-FM Configuration

After you have deployed a new GigaVUE-FM instance, you need to perform an initial configuration before you can start using GigaVUE-FM. This procedure only needs to be performed once for each GigaVUE-FM instance deployed.

To perform the initial configuration:

1. Open Virtual Machine Manager.
2. Locate the **Virtual Machines** entry in the results pane, right-click the GigaVUE-FM virtual machine, and click the **Open** option.

   The Virtual Machine Connection tool opens for the GigaVUE-FM virtual machine.
3. Open the GigaVUE-FM virtual machine from the Virtual Machine tool.

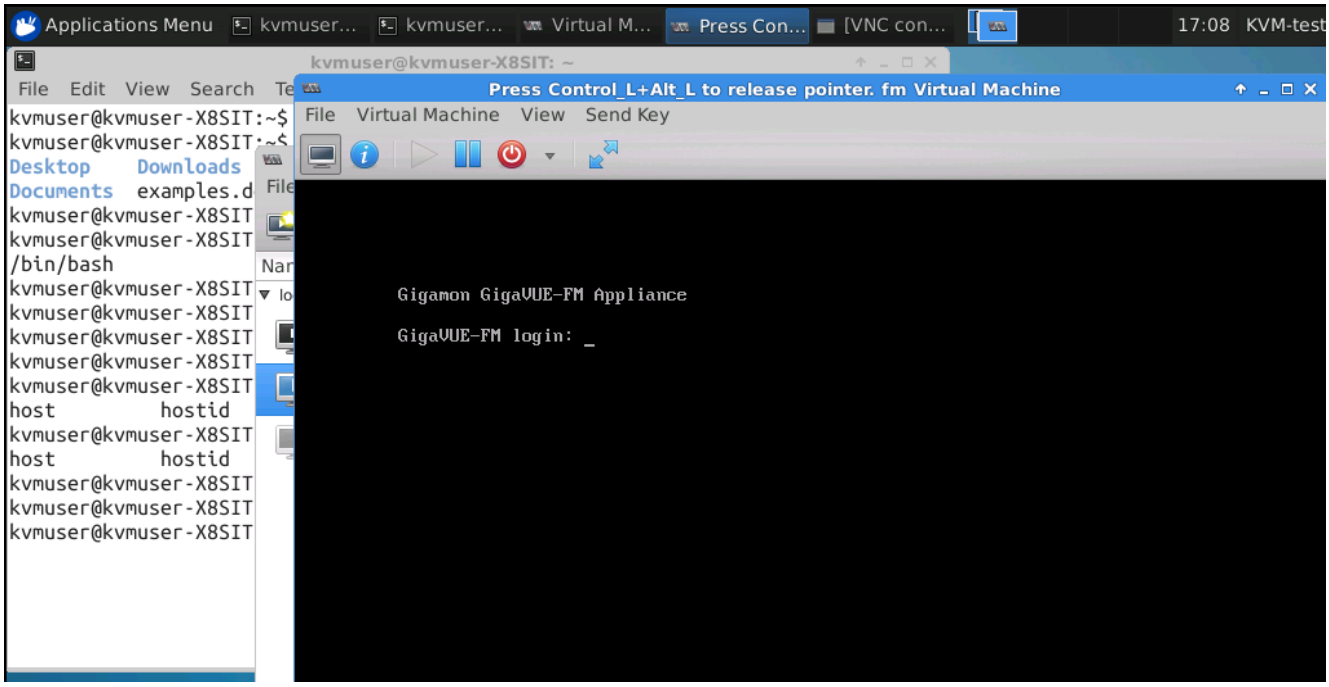

4. The GigaVUE-FM virtual machine powers on and displays a login prompt.
5. Log in as **admin** with password **admin123A!**

   The configuration wizard starts automatically, as shown in the figure below.
6. At this point, you will be presented with a series of prompts for you to provide the initial configuration for GigaVUE-FM.
   a. Provide a unique hostname for GigaVUE-FM. Note that the hostname may contain letters, numbers, periods (**.**), and hyphens (**-**), but may not begin with a hyphen. No other special characters are permitted.

b. Decide whether to use DHCP for the management interface.

   If you choose **no**, you will be prompted to provide the following:
   - IPv4 address and **masklen**
   - Default gateway
   - Primary DNS server
   - Domain name

   If you choose yes, skip to Step c

c. If you choose Yes for Step b, follow these instructions. The same options are repeated if DHCP is selected as No, but only one DNS IP address and domain server can be listed.

   For configuration options:

   - **Additional Domain Name Server IP Addresses**? - the address of any additional name servers required must be provided as a set of IP addresses with spaces as shown in the following figure.
   - **Additional DNS Domains**? - Multiple DNS domains can be defined in the jump start configuration with spaces in between as shown in the following figure.
   - **Enable NTP**? [yes] - the default is set to “yes”. The following options are available:

     NTP Server IP Address? - enter the NTP server address

     NTP Server Version? - enter the NTP version number of the NTP server

7. Provide an appropriate password for your environment. (Type a password and press **Enter**, or just press **Enter** to leave the password unchanged.)

   NOTE: Blank passwords are not permitted.

   The console displays your selections with instructions on how to make changes, if necessary.

8. Press **Enter** to save your choices and exit the wizard.
9. Your initial configuration is saved and GigaVUE-FM is up and running. You should now be at a standard mode command prompt.

   You can now access GigaVUE-FM by opening a browser and entering its IP address (the IP address you specified).

Configure SSH Settings

SSH access is enabled by default on new GigaVUE-FM deployments. You can enable SSH from the CLI using the **ssh server enable** command. By default the SSH server runs on port 22.

You can configure GigaVUE-FM to use a custom port for its SSH server with the **ssh server ports <port number>** command followed by a **write memory** command to save the configuration. For example, the following commands change the SSH port number to 2222.

```
(config) # ssh server ports 2222
(config) # write memory
(config) #
```

After making the settings shown above in the GigaVUE-FM CLI, you can connect an SSH session to GigaVUE-FM using the new port number.

HTTP/HTTPS Ports

GigaVUE-FM uses the following standard HTTP configuration:

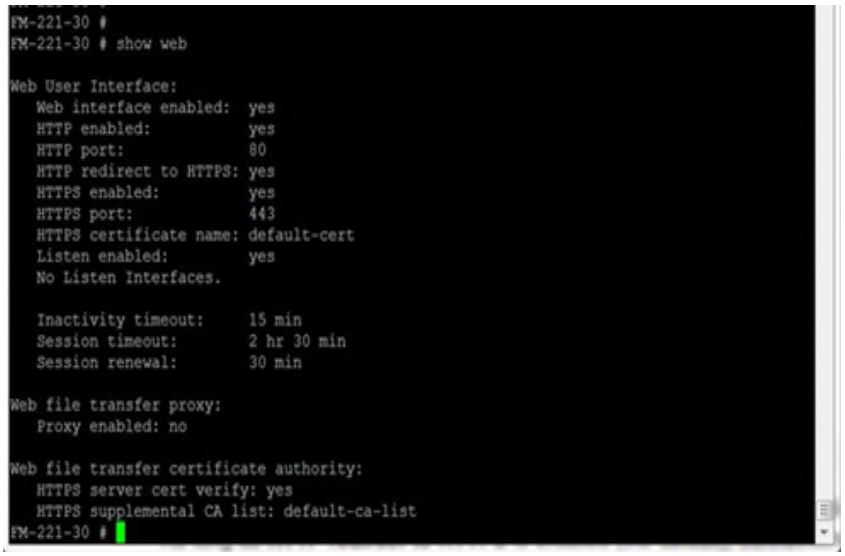

In this release, you can change the HTTPS port for GigaVUE-FM but the HTTP port is hard-coded to 80. As long as **HTTP redirect to HTTPS** is enabled (the default), connections to the fixed HTTP port of 80 will redirect to whatever the configured HTTPS port is.

Make Sure the Web Server is Enabled on Nodes to be Managed

GigaVUE-FM can only discover and manage nodes with their web servers enabled and operating on the default HTTP port of 80. Both G and H Series nodes have their web servers enabled by default. However, if you disabled a node's web server or changed its HTTP port, you will need to restore the settings before GigaVUE-FM can manage it.

GigaVUE-FM can manage nodes operating on custom HTTPS ports. Incoming HTTP connections redirect to the custom HTTPS port.

The **show web_server** and **show web** output listed below summarizes the necessary HTTP settings for GigaVUE nodes managed by GigaVUE-FM. The items shown in red are required settings.

H Series

H Series (config) # show web
Web-based management console enabled: yes
HTTP enabled: yes (Must be enabled)
HTTP port: 80 (Must remain at its default setting of 80)
HTTP redirect to HTTPS: yes (Must remain enabled)
HTTPS enabled: yes (Must remain enabled)
HTTPS port: 443 (Can be set to any custom value; HTTP redirects here)
Listen enabled: yes
No Listen Interfaces.
Inactivity timeout: 15 min

```
Session timeout:       2 hr 30 min
Session renewal:     30 min
Web proxy enabled: no
```

8 Upgrade GigaVUE-FM

This section describes how to upgrade GigaVUE-FM to the latest revision in either a VMware ESXi host or in Microsoft HyperV environment. Starting with release 3.1, Gigamon supports KVM environments. Previous versions of GigaVUE-FM have not been tested in the KVM environment.

NOTE: To upgrade software on H Series or TA Series nodes, refer to *Upgrade Software on a GigaVUE Node or a Cluster from GigaVUE-FM* on page 228

The topic covered in this sections cover:

- *Upgrade an Existing GigaVUE-FM Deployment* on page 113, which describes the overall upgrade path from an existing GigaVUE-FM deployment.
- *How to Use the "Snapshot" Feature* on page 120, which describes how to upgrade GigaVUE-FM 3.1 and above to the current version of GigaVUE-FM.

Upgrade an Existing GigaVUE-FM Deployment

Before starting an upgrade to GigaVUE-FM version, be sure to get the latest image, upgrade information, and release notes from the customer portal. Be sure to review the release notes for the latest release prior to upgrading your instance of GigaVUE-FM.

Once the GigaVUE-FM image is obtained, download it to a server within your environment from which the current instance of GigaVUE-FM can upload it. It is important to save your current running configuration using the facilities provided by the hypervisor before upgrading.

Ensure that the GigaVUE-FM time is set correctly to ensure accuracy of the trending data that is captured.

You can upgrade the existing deployment of GigaVUE-FM either from the CLI or from the GigaVUE-FM UI.

NOTE: When upgrading GigaVUE-FM, you must also upgrade GigaVUE-VM. For the steps to upgrade GigaVUE-VM, refer to the *"Bulk Upgrading GigaVUE-VM Nodes"* section in the *GigaVUE-VM User's Guide*.

Upgrade from CLI

There are five steps on how to upgrade an existing GigaVUE-FM deployment to the current release.

1. Verify that less than two images are present on the GigaVUE-FM server.
2. Download the new image into GigaVUE-FM using either HTTP, HTTPS, FTP, TFTP, SCP, or SFTP.
3. Install the new image.
4. Change boot partition.
5. Reboot GigaVUE-FM.

Notes:

- It is important to log in with the **admin** account/username when upgrading the image on the existing GigaVUE-FM.
- GigaVUE-FM 3.2 and higher versions compute node health status differently than previous versions. After the upgrade completes, rediscover the nodes to recompute node health status.
- Prior to GigaVUE-FM 3.2, backup files for physical nodes were in a binary format. Starting with GigaVUE-FM 3.2, backup and restore files use a text based format and binary backup or restore on physical nodes is not supported. When upgrading from a version lower than version 3.2, backup your configuration prior to upgrading to the current version of GigaVUE-FM if you desire, but the files will be in a binary format. Existing binary backups are not visible to GigaVUE-FM. For binary backups, you must back up the node using the CLI commands rather than GigaVUE-FM. For more information about the CLI commands, refer to the *GigaVUE-OS CLI User's Guide*.
- When using the Firefox or IE browser, clear the cache before upgrading to prevent issues with the browser.

Step 1: Verify that only two images are present on GigaVUE-FM server

NOTE: It is important that you log in with the **admin** account/username when upgrading the image on the existing GigaVUE-FM.

1. To begin an upgrade, open a SSH session or console session within the vSphere Client, and log int o GigaVUE-FM and change to configure mode, by entering the following on the command line:
   a. Type **en <Enter>** to switch to Enable mode.

      The system prompt changes from ***[hostname]*** **>** to ***[hostname]*** **#**.

   b. Type **config t <Enter>** to switch to Configure mode.

      The system prompt changes from ***[hostname]*** **#** to ***[hostname]*** **(config) #**

   Figure 8-1shows an example of the login console.
2. Check the number of images currently available for installation with the following command from the GigaVUE-FM CLI:

   (config) # show images

Important: If there are more than two images listed in the **Images available to be installed** section of the **show images** output, Gigamon recommends that you use the **image delete** command to remove existing images until the system has only a single image. Both GigaVUE-FM and GigaVUE-VM will display a warning if you attempt to fetch a third image.

3. To delete an existing image from the server use the following command:

   (config) # image delete fma3300.img

4. Go to *Step 2: Fetch the latest release of GigaVUE-FM* on page 115.

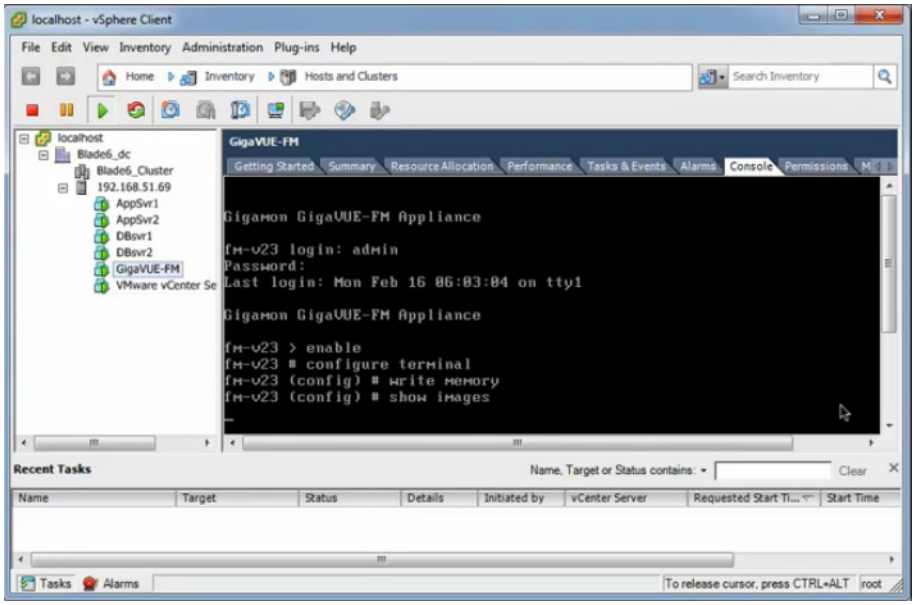

Figure 8-1: Console Login to vSphere Client

Step 2: Fetch the latest release of GigaVUE-FM

Gigamon provides an FTP site where the new release image file resides. To fetch the latest release, do the following:

1. Locate the image file for the new release. Image files are named using a fma***xxxx***.img format. The ***xxxx*** indicates the version and build number (for example, **fma3500.img** for the v3.5 release).
2. Copy the image to your file server.
3. Use the **image fetch** command to retrieve the software image from your file server. The CLI shows you the progress of the image fetch with a series of hash marks, returning you to the system prompt when complete.

   Note: The **image fetch** command supports the use of HTTP, HTTPS, FTP, TFTP, SCP, or SFTP for the transfer of images.

   a. The following command uses SCP to retrieve the **fma3500** image from the image server with the IP address of 10.115.0.100 using login and password.

      (config) # image fetch scp://user:password@10.115.0.100/fma3500.img

   b. The following command uses FTP to retrieve the same image using login and password as well.

      (config) # image fetch ftp://user:password@10.115.0.100/fma3500.img

c. The following command uses TFTP to retrieve the same image without using a password but using a DNS server instead of an IP address for the download server.

(config) # image fetch tftp://myserver.gigamon.com/tftpboot/fma3500.img

Ensure that you specify the base directory when using TFTP

4. Go to *Step 3: Install the latest release of the GigaVUE-FM* on page 116.

Step 3: Install the latest release of the GigaVUE-FM

Use the **image install** command to install the downloaded image file. When running the following command, the process will first verify that the filename used for the image is suitable for installation prior to installing the image. For example, to install the image downloaded in the previous step:

(config) # image install fma3400.img

Step 4: Change the boot partition

Set the image you just installed to boot next with the following command. This ensures that at the next boot the latest image will be picked up.

(config) # image boot next

Step 5: Reboot

The following command shuts down the current instance of GigaVUE-FM and reloads. If Step 4 is performed, upon reboot, the new image is used.

(config) # reload

Step 6: Upgrade GigaVUE-VM

After upgrading GigaVUE-FM, you must also upgrade any deployed GigaVUE-VMs. Otherwise, maps may not work and the GigaVUE-VMs will be unreachable. For information about upgrading GigaVUE-VM, refer to the *"Bulk Upgrading GigaVUE-VM Nodes"* section in the *GigaVUE-VM User's Guide*.

CLI Summary of the Upgrade Path

The following summarizes the CLI commands used to upgrade an image after logging in from the console:

```
> en
# config t
(config) # write memory
(config) # show images
(config) # image delete fma3400.img
(config) # image fetch tftp://192.158.51.41/fma3500.img
(config) # image install fma3500.img
(config) # image boot next
(config) # reload
```

Upgrade from GigaVUE-FM UI

This section describes the steps to upgrade GigaVUE-FM from the UI. You can upgrade by using an image that is located on an external image server, or you can use GigaVUE-FM as the image server.

NOTE:

- When using the GigaVUE-FM UI to upgrade GigaVUE-FM, you can only upgrade to the currently available version or to the next version. You cannot downgrade.
- When using the Firefox or IE browser, clear the cache before upgrading to prevent issues with the browser

Upgrade from an External Image Server

This section provides the steps for upgrading the GigaVUE-FM from an image stored on an external server. The image can be transfered from the server to the GigaVUE-FM using either SCP or TFTP file protocols.

To upgrade with an image stored on an external image server, do the following:

1. Upload the image to the external image server to make it available to GigaVUE-FM.

   To obtain software images, register on the customer portal and download the software. To reach the customer portal, go to https://gigamoncp.force.com/gigamoncp/.

2. Add the image server to GigaVUE-FM. This stores the credentials, image file name, and IP address of the server on GigaVUE-FM.

   To add the image server:

   a. In GigaVUE-FM, click **Administration** on the top navigation link of the window.

   b. Select **System** on the left navigation panel and go to **Images > External Servers**. The External Servers page displays as shown in Figure 8-2.

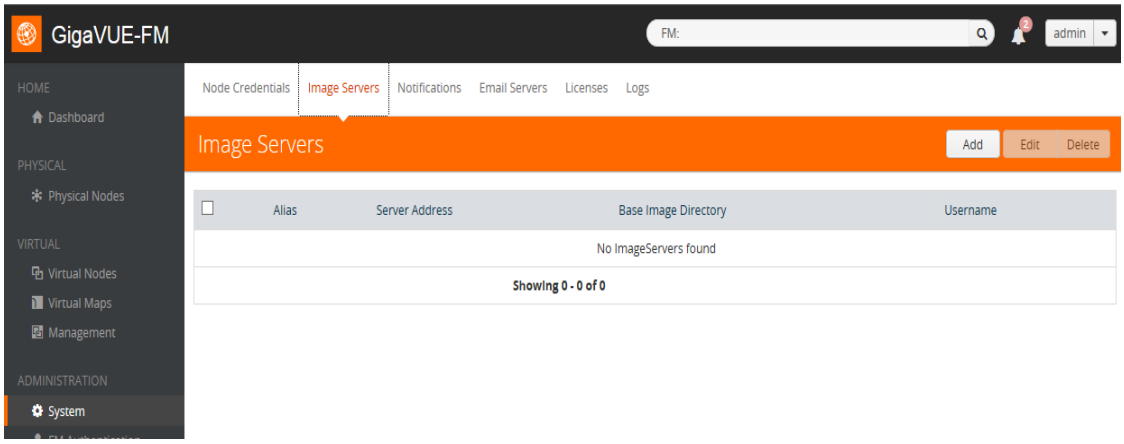

Figure 8-2: Adding Image Servers

   c. Click **Add**. The Add External Server page displays as shown in Figure 8-3

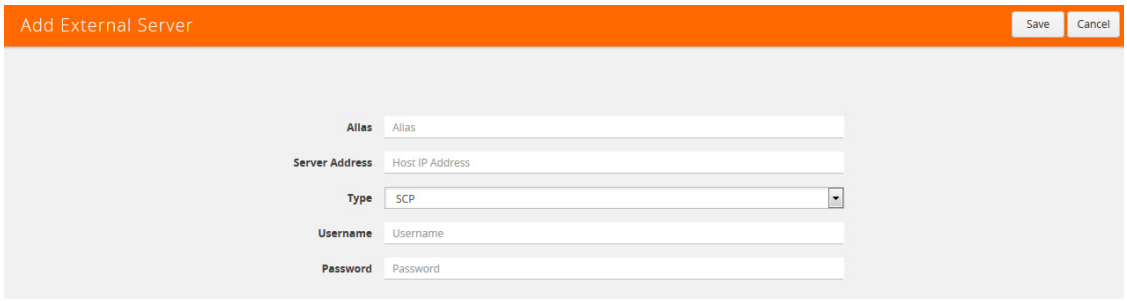

Figure 8-3: Add External Server

d. On the Images Server page, specify the following:

- An alias to help identify the image server.
- The host IP address of the server.
- The protocol to use for the download: SCP or TFTP.
- The user name and password if you selected SCP. They are not needed for TFTP.

e. Click **Save**.

The External Server page displays the newly added external server.

3. From the **Admin** drop-down list in the top right corner of the window, select **Upgrade** to open the FM Image Upgrade page shown in Figure 8-4.

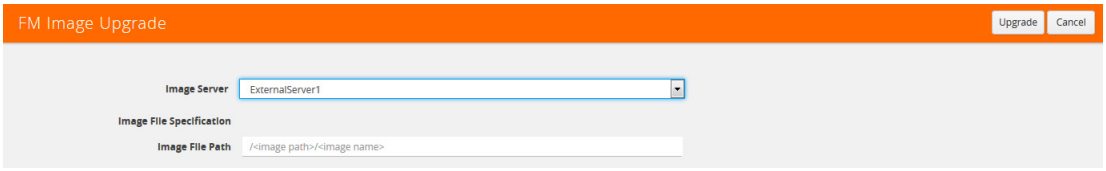

Figure 8-4: FM Image Upgrade Page

To monitor the progress and status of the upgrade, click **Administration** and go to **All Alarms/Events**. Also, email notifications are sent if Email Notifications have been configured. For more information about Email Notifications, refer to *Notifications* on page 1320.

4. On the FM Image Upgrade page, click on the Image Server field and select the server added in Step 2.

5. In the Image File Path, enter the image path and filename on the external file server.

6. Upgrade any deployed GigaVUE-VMs.

After upgrading GigaVUE-FM, you must also upgrade any deployed GigaVUE-VMs. Otherwise, maps may not work and the GigaVUE-VMs will be unreachable. For information about upgrading GigaVUE-VM, refer to the *"Bulk Upgrading GigaVUE-VM Nodes"* section in the *GigaVUE-VM User's Guide*.

NOTE: If you are using FireFox or Internet Explorer, you must refresh the browser to ensure that the cached information is not displayed after upgrading to the latest version of GigaVUE-FM.

Upgrade with GigaVUE-FM as the Image Server

This section provides the steps for upgrading GigaVUE-FM when GigaVUE-FM is used as the file server instead of an external server.

To upgrade a GigaVUE-FM using internal image files, do the following:

1. Download the images from the Gigamon website and place them where they can be available for uploading to GigaVUE-FM.

   To obtain software images, register on the customer portal and download the software. To reach the customer portal, go to https://gigamoncp.forc.com/gigamoncp/.

2. Upload the images file to GigaVUE-FM.

   a. Click the **Administration** on the top navigation link. In the left navigation panel, go to **System > Images > Internal Image Files**.

   b. On the Internal Image File page, click **Upload**.

   c. Click **Browse** to locate the image file.

   d. Click **OK** to upload the file. The page displays the progress of the upload.

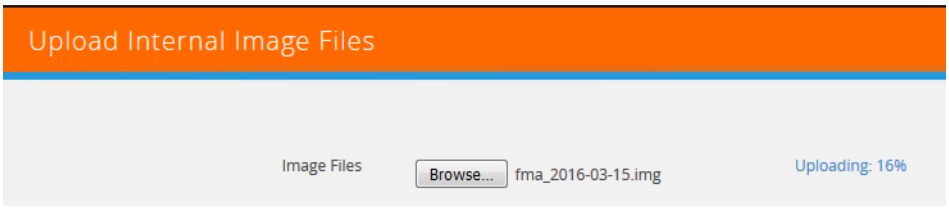

   After the upload completes, you can see the GigaVUE-FM image to use for the upgrade on the Internal Images Files page.

3. Click the **Admin** drop-down list on the top right of the window and select **Upgrade** as shown in Figure 8-5.

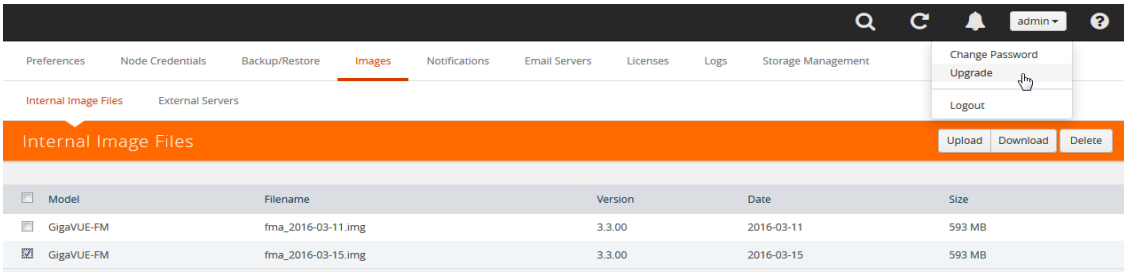

Figure 8-5: Selecting Upgrade

4. On the FM Image Upgrade page, click in the **Image Server** field and select **Internal Image Server**.

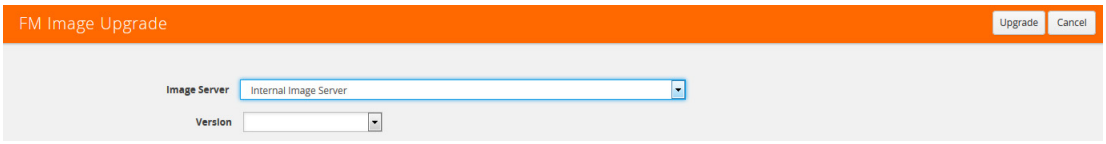

5. From the **Version** drop-down list, select the version to which you are upgrading.

   NOTE: You can only upgrade to another instance of the current version or the immediate next version. Downgrading to a lower version is not supported through the UI. To downgrade to an lower version, use the CLI.

6. Click **Upgrade**.

   To monitor the progress and status of the upgrade, click **Administration** on the top navigation link and select **All Alarms/Events** on the left navigation panel. Also, email notifications are sent if email notifications have been configured.

7. Upgrade any deployed GigaVUE-VMs.

   After upgrading GigaVUE-FM, you must also upgrade any deployed GigaVUE-VMs. Otherwise, maps may not work and the GigaVUE-VMs will be unreachable. For information about upgrading GigaVUE-FM, refer to the *"Bulk Upgrading GigaVUE-VM Nodes"* section in the *GigaVUE-VM User's Guide*.

NOTE: If you are using FireFox or Internet Explorer, clear the cache before upgrading to prevent issues with the browser.

How to Use the "Snapshot" Feature

This procedure is only valid for upgrading from GigaVUE-FM v3.0 and above. For upgrades from pre-3.0 releases, review the GigaVUE-FM v3.0 User's Guide and upgrade to release GigaVUE-FM v3.1. Then follow the steps below to upgrade to the current release version of GigaVUE-FM.

NOTE: You cannot directly upgrade from a pre-3.3 releases to the current release. You can only upgrade from GigaVUE-FM v3.3 or v3.4 release.

1. Prior to upgrading, ensure that the available **memory size is at least 8GB** prior to upgrading to the new GigaVUE-FM release. If the available memory size is less than 8GB, it will cause out of memory issues. Also, at least 2 vCPU are required.
2. When upgrading from v3.1, it's a good idea to use the vSphere client's **Snapshot** feature to record the current state of the GigaVUE-FM virtual machine. Steps to use **Snapshot** feature are as follows:
   a. Log into the vSphere client and navigate to the Datacenter or Cluster level where the GigaVUE-FM installation is located.
   b. Right-click the GigaVUE-FM entry in the vSphere client and select the **Take Snapshot** option.

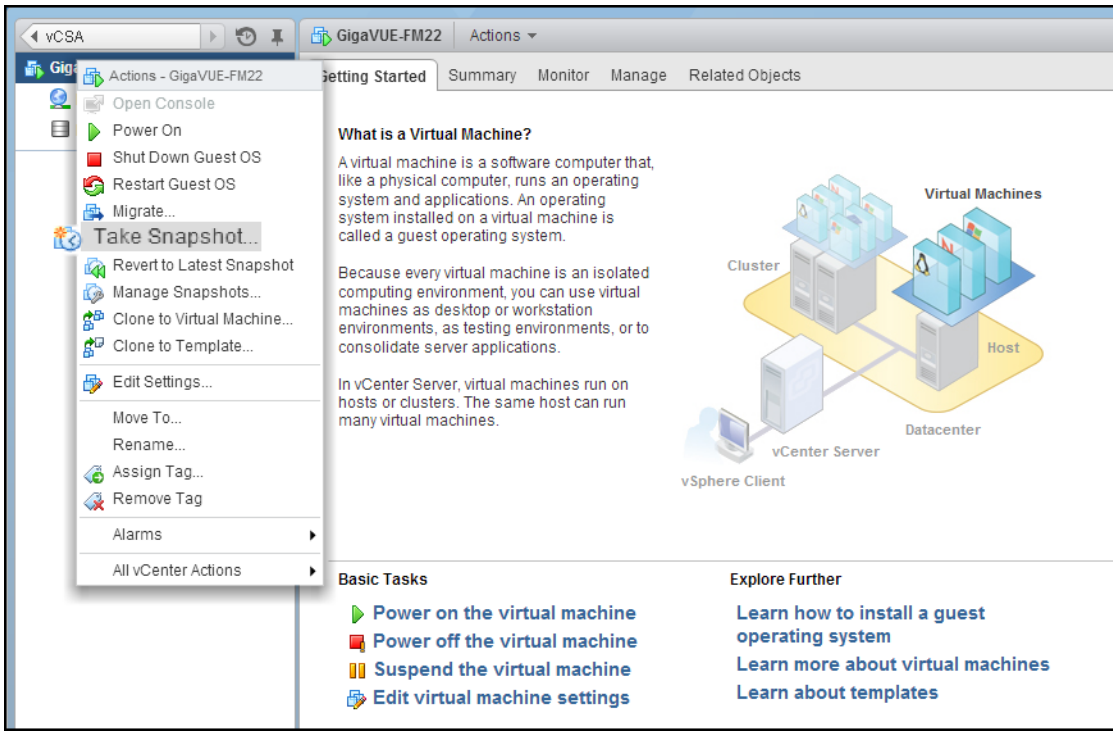

Figure 8-6: "Take Snapshot" Command to Preserve Current Settings Prior to Upgrade

c. Follow the system prompts to record a snapshot of GigaVUE-FM's current state.

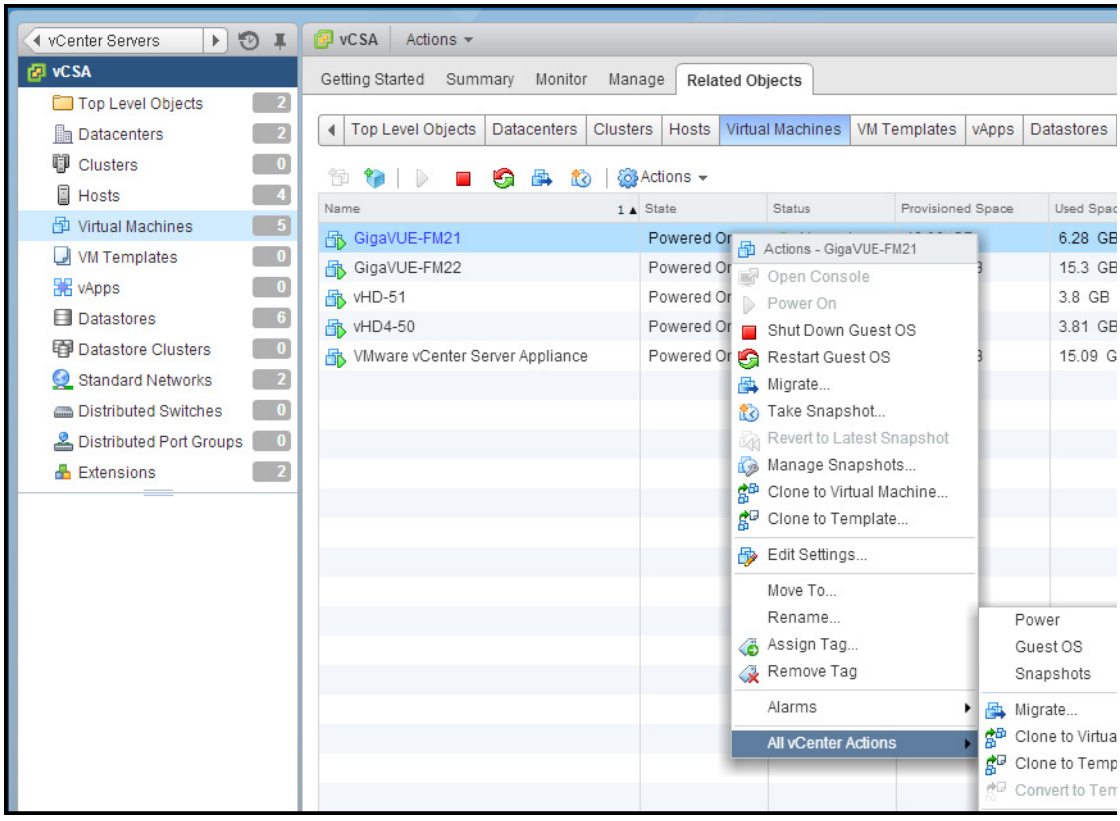

Figure 8-7: Power Off Command

Part 3: Dashboard

This section describes the Dashboards available in GigaVUE-FM that display information about the physical and virtual nodes manged by GigaVUE-FM as well as information about GigaVUE-FM itself.

The section provides information about the following topics:

- *Physical Dashboard* on page 125

  Describes the dashboards that display physical nodes manged by GigaVUE-FM and GigaVUE-FM itself.

- *Virtual Dashboard* on page 173

  Describes the dashboards that display virtual nodes manged by GigaVUE-FM.

- *Health Monitor Dashboard* on page 163

  Describes the health monitor dashboards that provides health information about GigaVUE-FM.

- *FabricVUE Traffic Analyzer* on page 169

  Describes the Traffic Analyzer, which provides a larger insight into traffic flows through the GigaVUE nodes, allowing visibility into the type of traffic that may not be flowing to the tool.

9 Physical Dashboard

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

This chapter covers the following topics:

Overview of the Physical Dashboard

The Physical Dashboard is a central location to monitor all the physical nodes and clusters that are managed by GigaVUE-FM. The widgets in the dashboard provides a quick visual overview of inventory and events, GigaSMART traffic, highest and lowest traffic by maps and ports, traffic comparison by tags, most and least utilized traffic, and health status. Refer to Figure 9-1 on page 125.

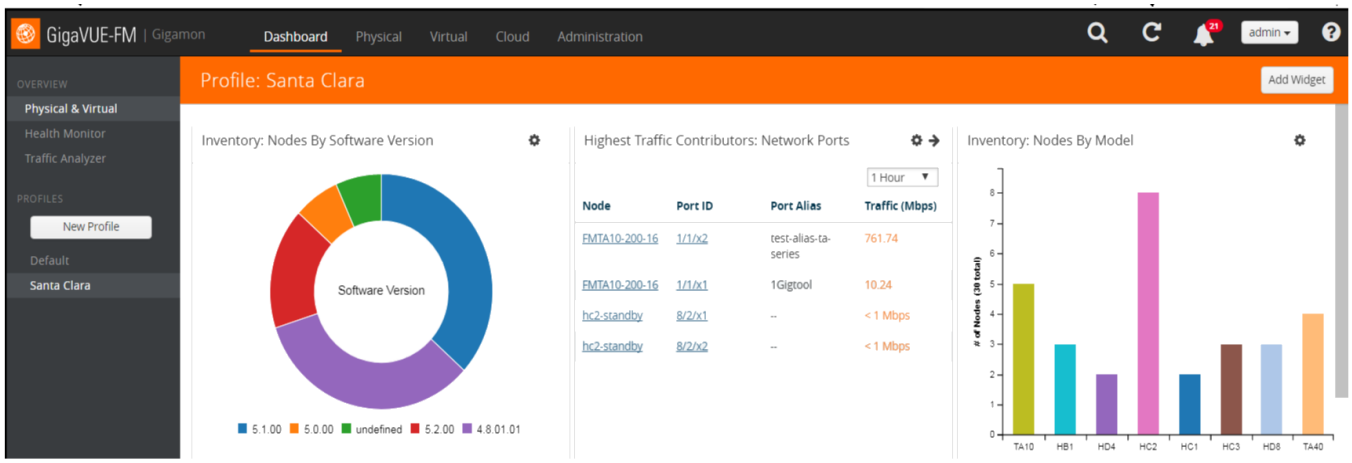

Figure 9-1: GigaVUE-FM Overview

Physical Dashboard Profiles

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

To create a new profile:

1. Click **Dashboard** on the top navigation link.
2. On the Physical & Virtual dashboard page, click **New Profile**. The Enter profile name box is displayed. Refer to Figure 9-2 on page 126.

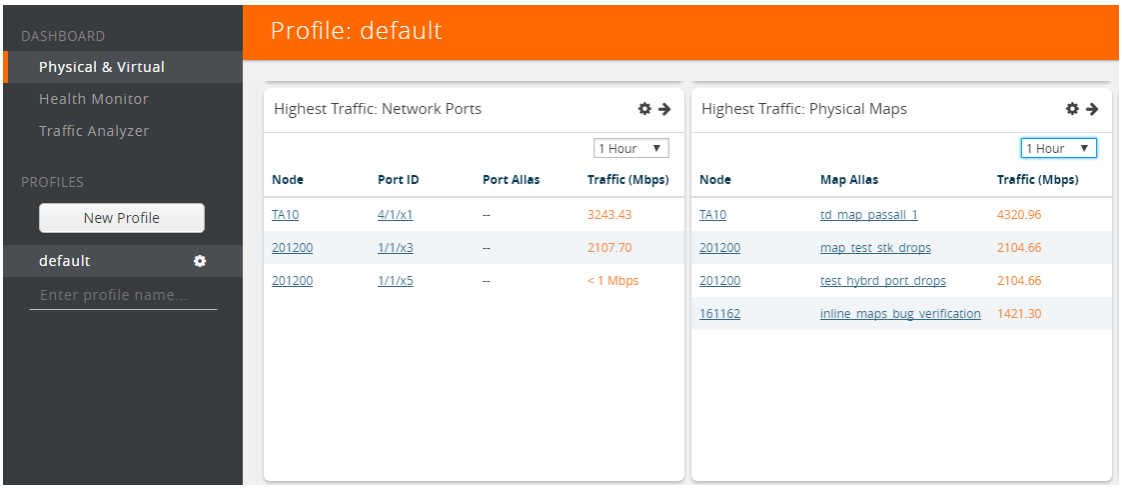

Figure 9-2: Create New Profile

3. In the **Enter profile name**... box, enter the name of the new profile and click **Enter**.

   The new profile name is displayed under Profiles in the left navigation. The new profile page is displayed as shown in Figure 9-3 on page 126.

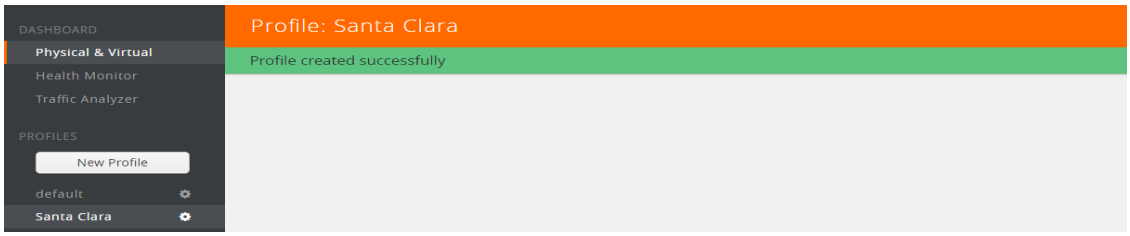

Figure 9-3: Create Profile for Physical Dashboard Selections

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

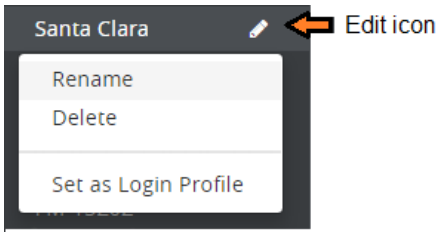

Figure 9-4: Profile Settings (add the setting icon label)

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

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

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

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

- Widget and trending data is available based on the GigaVUE-FM license purchased. For the base package, the data is not stored for more than 1 day. The prime package users can select any option including 1 month.
- Individual widgets can be resized and saved as part of the profile. Each widget can expand in both horizontal and vertical planes. The other widgets self-adjust when the widgets are manipulated.
- The widgets can also be dragged and dropped to different section of the page. Refer to Figure 9-5 on page 128.
- The data points can be viewed when the mouse is hovered over the graph as shown in Figure 9-5 on page 128.
- The widgets such as Unhealthy Maps opens a quick view when clicked on the cluster info or the map alias. The quick view shows more details relating to that specific map.
- The trending information can also be changed for each widget on the same dashboard.
- The port and map health status changes on the device reflect instantly on the screen.
- The color-coded legends are available at the bottom of each widget.

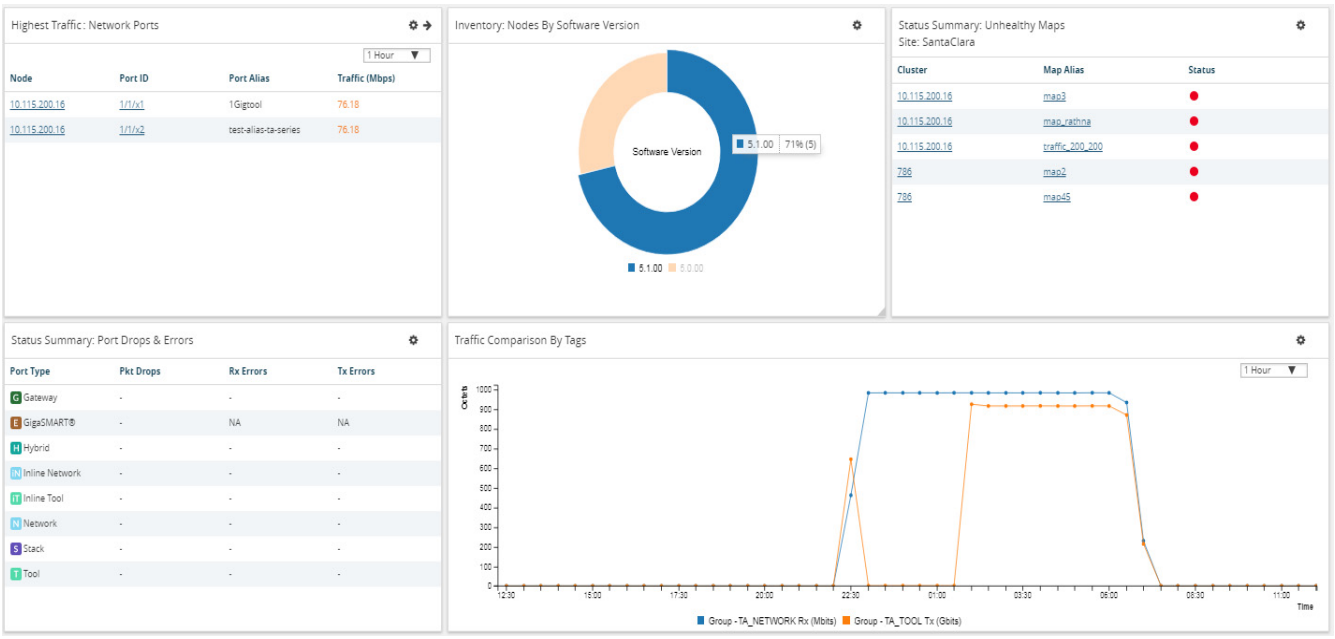

Figure 9-5: Resized and Repositioned Widgets

Physical Dashboard Quick Views

When reviewing the widgets available on the Physical & Virtual dashboard, clicking on the options in the widgets takes you to the details page relating to the information for

that node. For example, on the Nodes by Model or Software Version widget, you click on the node and it takes you to the Physical Nodes page.

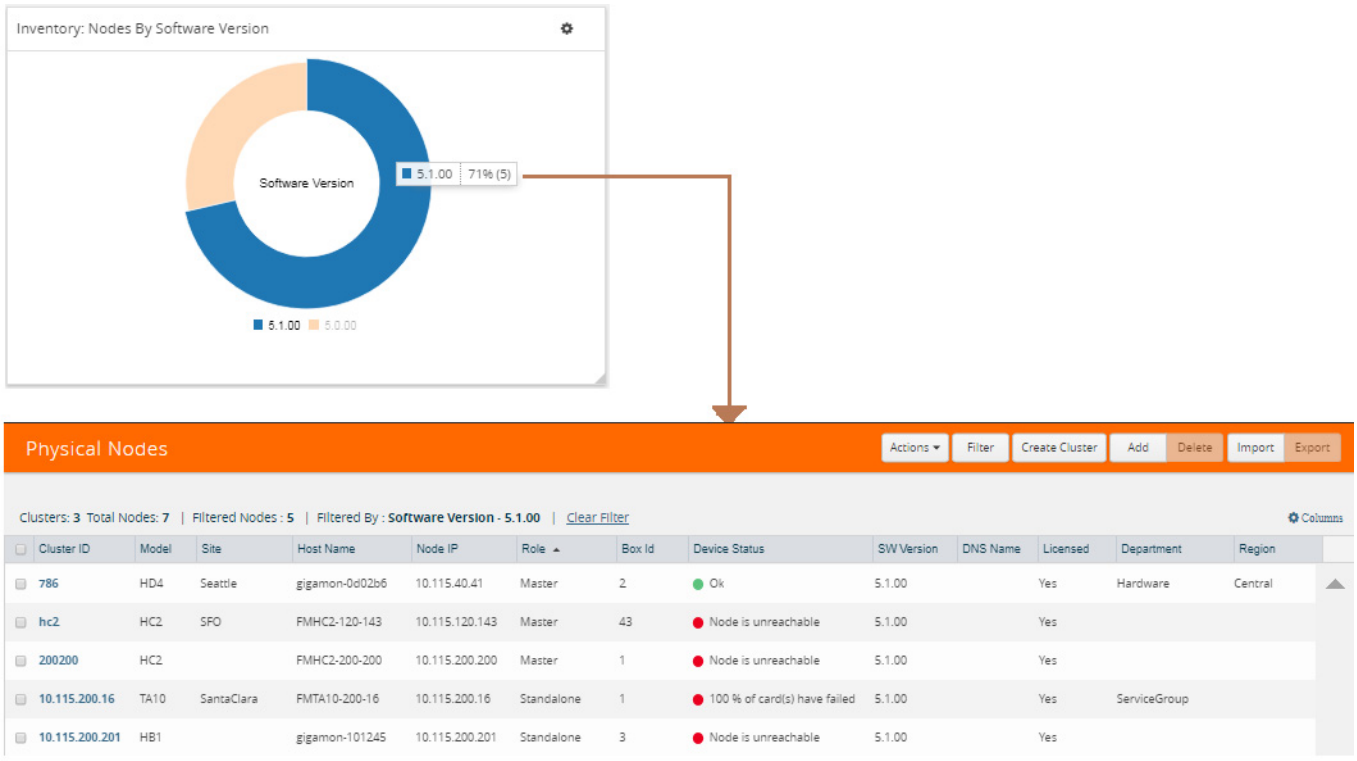

Figure 9-6: Nodes By Software Version Widget

For more information about Physical Nodes, refer to *Manage GigaVUE Nodes and Clusters* on page 187.

Physical Dashboard Widgets

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

- *Highest Traffic* on page 130
- *Lowest Traffic* on page 134
- *Traffic Comparison By Tags* on page 135
- *Most Utilized Traffic* on page 139
- *Least Utilized Traffic* on page 142
- *Inventory* on page 143
- *Status Summary* on page 146

The default profile displays the following widgets:

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

- Audit Logs
- Events

You can customize the widgets by modifying the physical dashboard profiles. Refer to *Physical Dashboard Profiles* on page 126 for more information.

Highest Traffic

The Highest Traffic widget can be created for the following:

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

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

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

The highest traffic can be displayed as either a table or a graph. By default, a table is displayed. You can click the arrow to change the display to a graph as shown in Figure 9-7 on page 131.

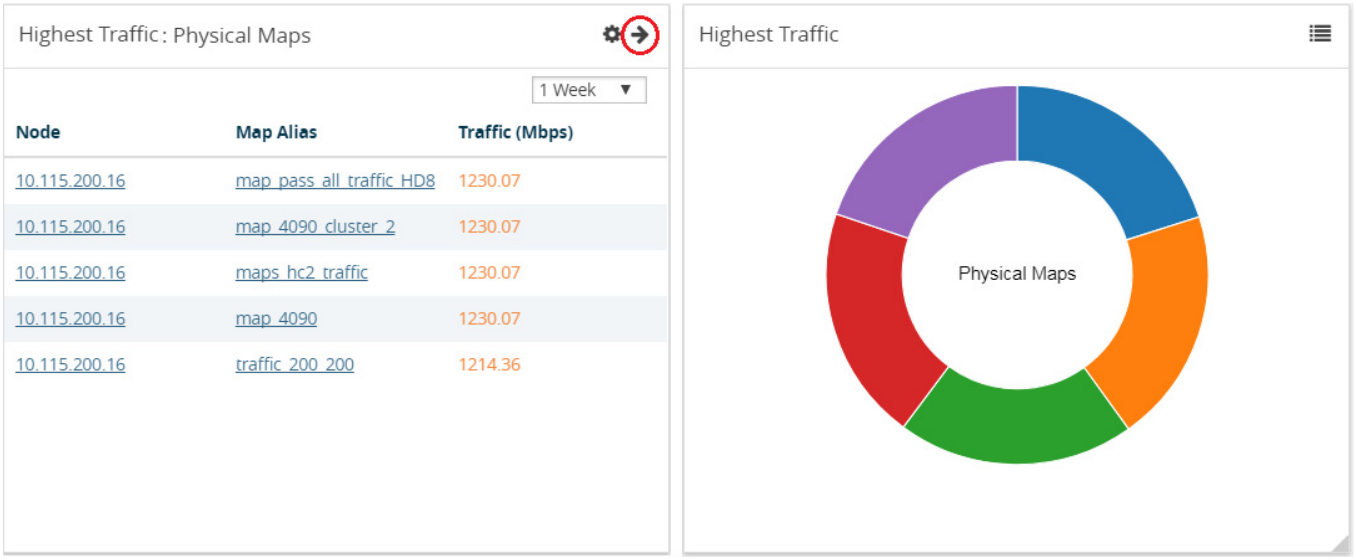

Figure 9-7: Highest Traffic: Physical Maps Example

In the graph view, each ring represents a map or a port. You can hover your mouse over the graph to view the percentage of traffic handled by the map or the port.

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

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

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

The highest traffic for GigaSMART operations or GigaSMART group can be displayed as shown in Figure 9-8 on page 132.

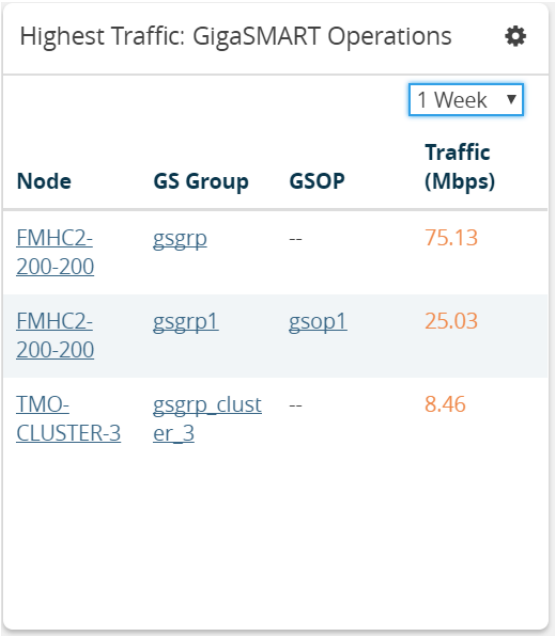

Figure 9-8: Highest Traffic GigaSMART

To configure the Highest Traffic widget:

1. Click **Dashboard** on the top navigation link.
2. On the Physical & Virtual dashboard page, select the profile in which you want to add the widget.

3. Click **Add New Widget**. The Add New Widget window is displayed. Refer to Figure 9-9 on page 133.

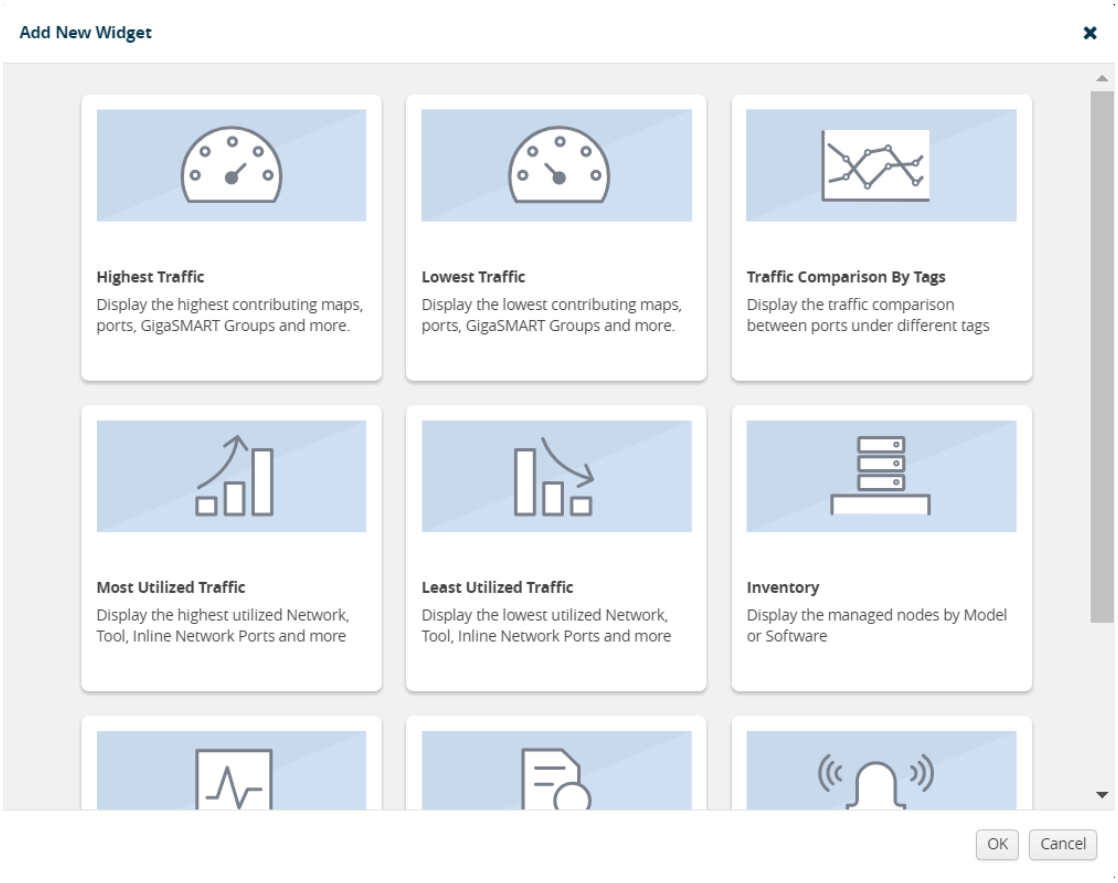

Figure 9-9: Add New Widget

4. In the Add New Widget window, select **Highest Traffic** and click **OK**. The Highest Traffic configuration window is displayed. Refer to Figure 9-10 on page 134.

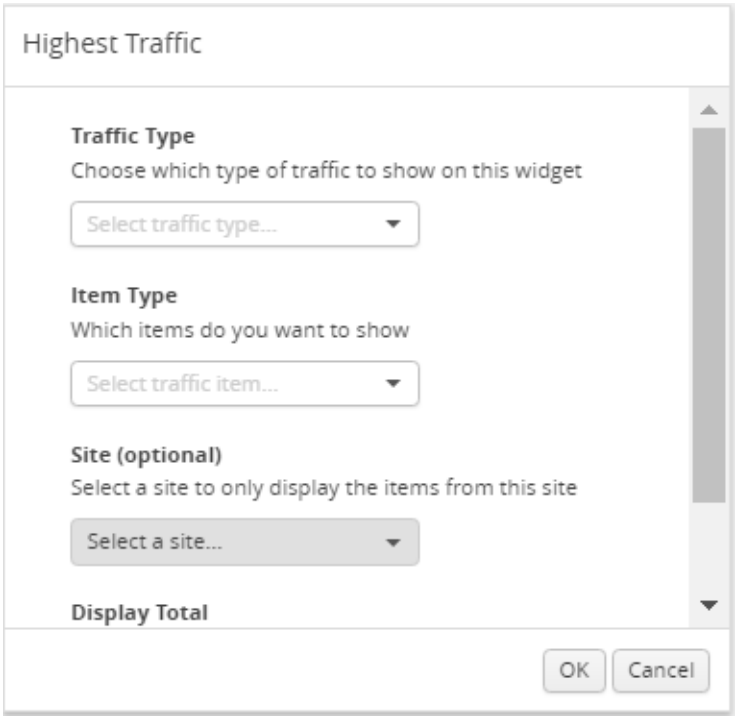

Figure 9-10: Highest Traffic Configuration

5. From the **Traffic Type** drop-down list, select one of the following traffic types:
   - Physical—Allows you to view the physical maps and ports that contribute to the highest traffic distribution.
   - GigaSMART—Allows you to view the virtual ports, GigaSMART groups, and GigaSMART operations that contribute to the highest traffic distribution.
6. From the **Item Type** drop-down list, select the item you want to view. The options displayed are based on the traffic type you selected in step 5.
7. (Optional) From the Site drop-down list, select a site to view only the ports associated to the selected site.
8. From the **Display Total** drop-down list, select the number of items to be displayed. By default, the number of items selected for display is 5.
9. Click **OK**.

Lowest Traffic

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

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

The physical maps and ports can be displayed as either a table or a graph. By default, a table is displayed. You can click the arrow to change the display to a graph as shown in Figure 9-11 on page 135.

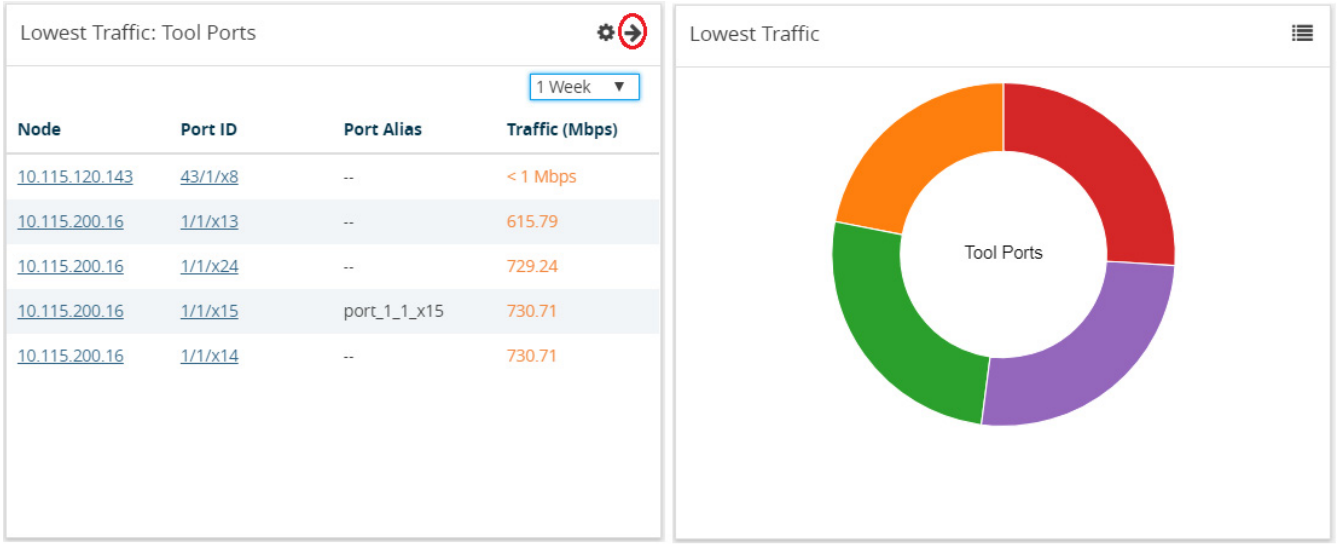

Figure 9-11: Lowest Traffic

In the graph view, each ring represents a map or a port. You can hover your mouse over the graph to view the percentage of traffic flowing through the map or the port.

The Lowest Traffic widget is configured exactly the same way as the Highest Traffic widget. To configure the Lowest Traffic widget, refer to the configuration steps provided in *Highest Traffic* on page 130. In **step 4**, select **Lowest Traffic** and click **OK**. The Lowest Traffic configuration window is displayed.

Traffic Comparison By Tags

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

In this example, there is traffic flowing from GigaVUE-TA10 to GigaVUE-HC2. You can group the tool ports in GigaVUE-TA10 and create a tag as TA_TOOL. Then, you can

group the network ports in GigaVUE-HC2 and create a tag as HC2-NETWORK. Refer to Figure 9-12 on page 136.

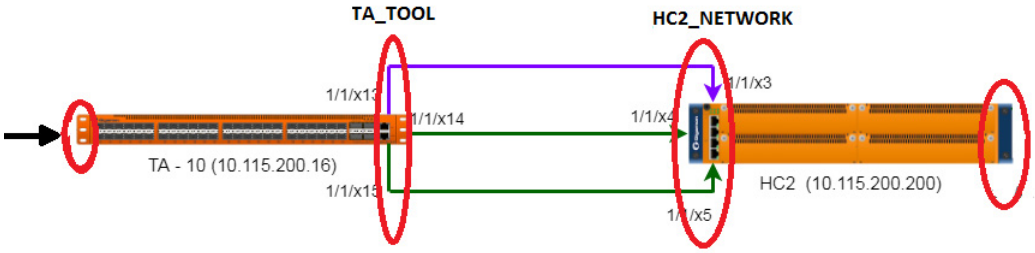

Figure 9-12: Example for Traffic Comparison By Tags Widget

Using the Traffic Comparison By Tags widget, you can compare the egress traffic passing through the ports associated with TA_TOOL with the ingress traffic passing through the ports associated with HC2-NETWORK, and quickly analyze if there is any packet loss associated. Refer to Figure 9-13 on page 136

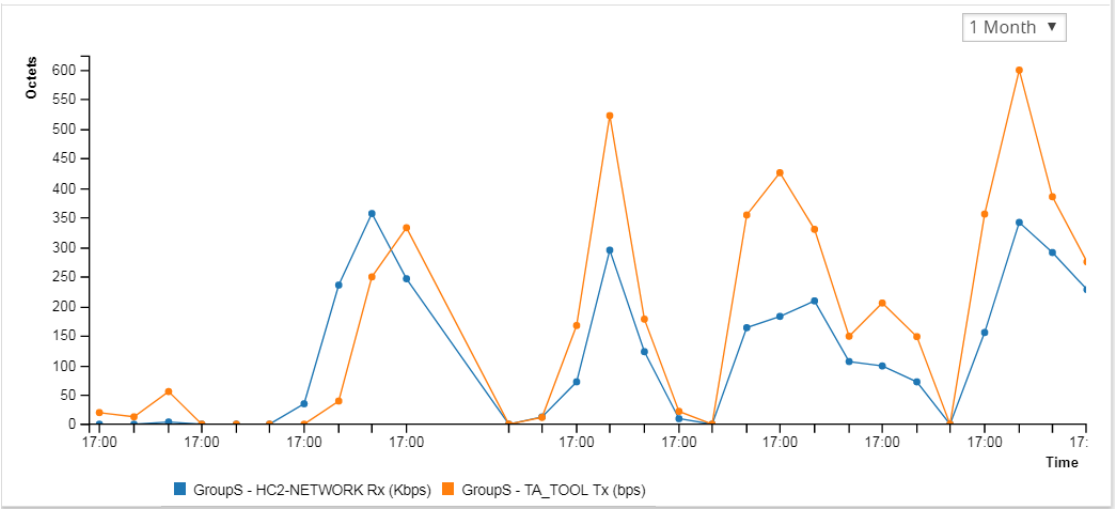

Figure 9-13: Traffic Comparison By Tags

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

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

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

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

The aggregated traffic comparison is displayed as a graph. You can choose to display the data over a day, an hour, a week, or a month. However, when you select a week or a month, the time period is not persisted. The data is defaulted to 1 day when you navigate away from the Physical Dashboards page and then return to the page. Click the notification icon at the top of the window and view the alarms and notifications displayed (refer to Figure 9-14 on page 137):

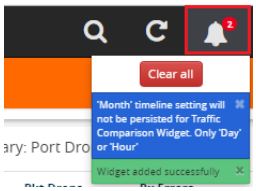

Figure 9-14: Notifications

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

To configure the Traffic Comparison By Tags widget:

1. Click **Dashboard** on the top navigation link.
2. On the Physical & Virtual dashboard page, select the profile in which you want to add the widget.

3. Click **Add New Widget**. The Add New Widget window is displayed. Refer to Figure 9-9 on page 133.

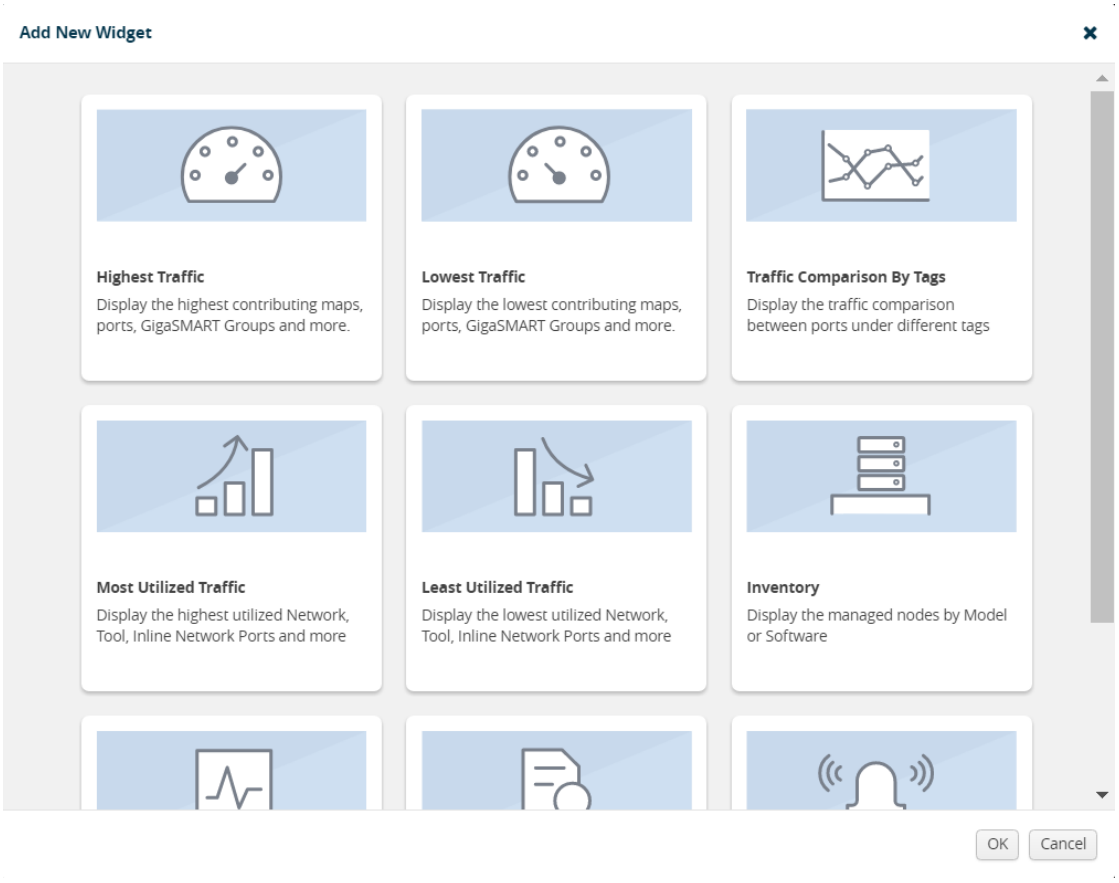

Figure 9-15: Add New Widget

4. In the Add New Widget window, select **Traffic Comparison By Tags** widget and click **OK**. The Traffic Comparison by Tags configuration window is displayed. Refer to Figure 9-16 on page 138.

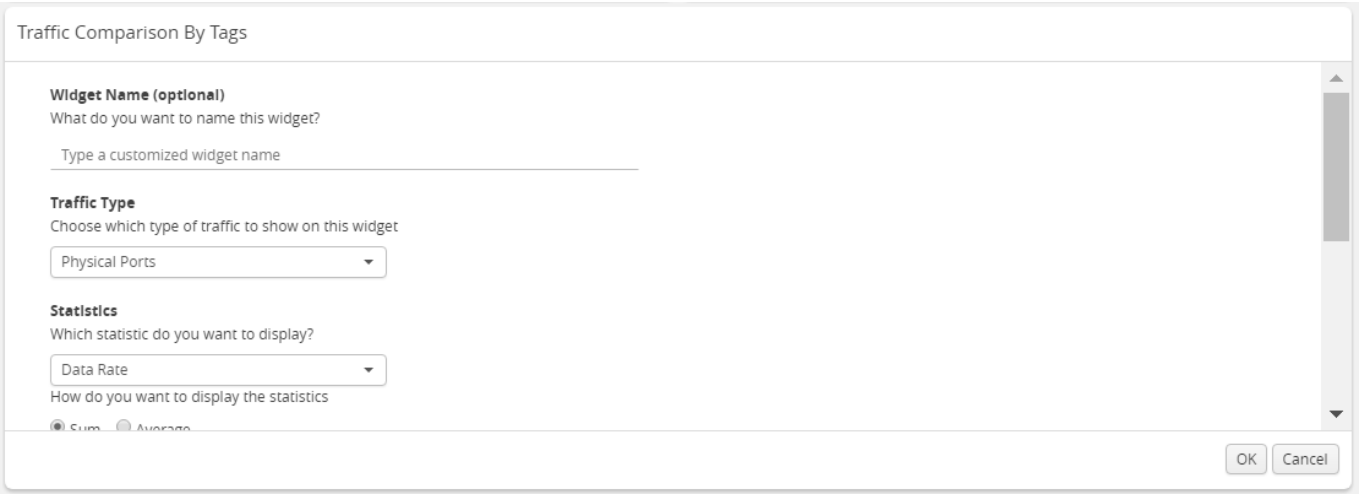

Figure 9-16: Traffic Comparison By Tags Configuration

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

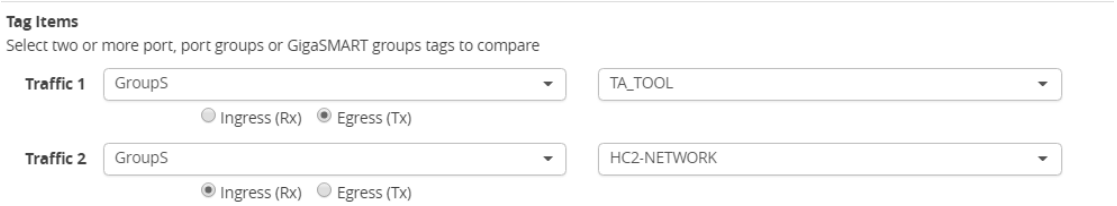

   c. Repeat step a and step b to select the next traffic for comparison.
10. Click **OK**.

Most Utilized Traffic

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

The Most Utilized Traffic widget lists the ports with the node ID, port number, port alias, and the utilization percentage. The most utilized traffic percentage is displayed as

either a table or a bar graph. By default, a table is displayed. You can click the arrow to change the display to a graph as shown in Figure 9-17 on page 140.

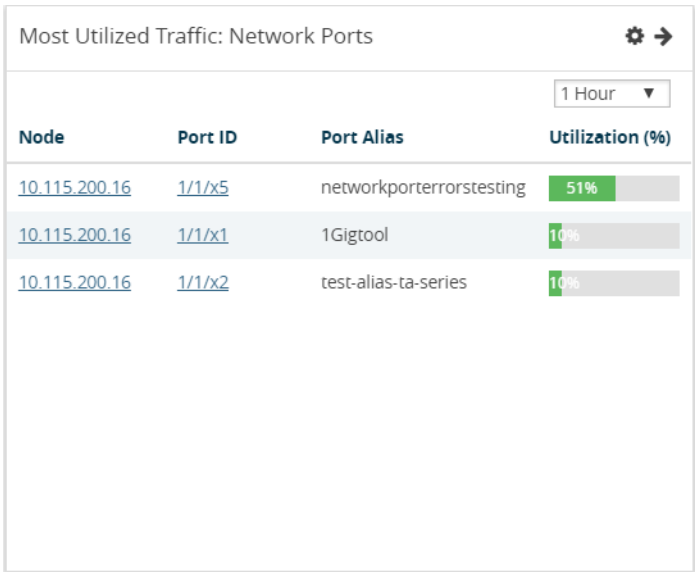

Figure 9-17: Most Utilized Traffic Widget

To configure the Most Utilized Traffic widget:

1. Click **Dashboard** on the top navigation link.
2. On the Physical & Virtual dashboard page, select the profile in which you want to add the widget.

3. Click **Add New Widget**. The Add New Widget window is displayed. Refer to Figure 9-18 on page 141.

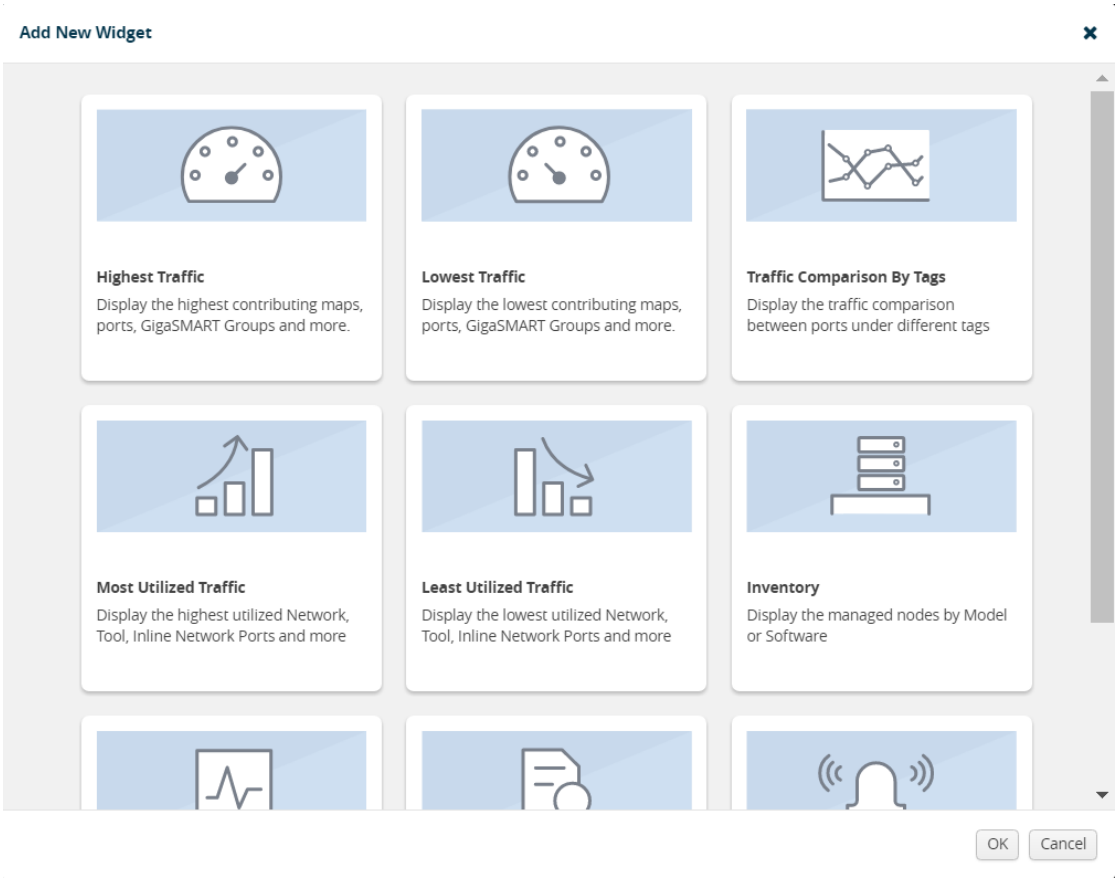

Figure 9-18: Add New Widget

4. In the Add New Widget window, select **Most Utilized Traffic** and click **OK**. The Most Utilized Traffic configuration window is displayed. Refer to Figure 9-19 on page 141.

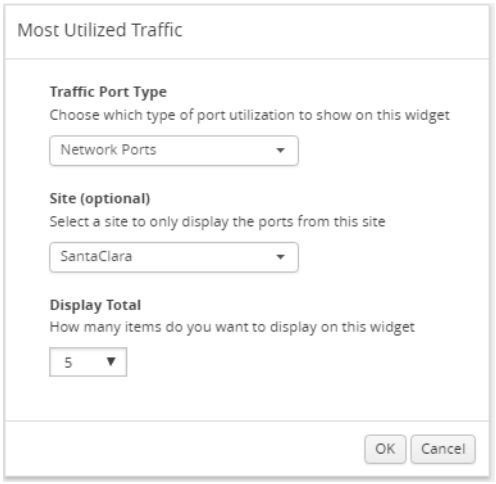

Figure 9-19: Most Utilized Traffic Configuration

5. From the **Traffic Port Type** drop-down list, select one of the following port types:

- Network Ports
- Tool Ports
- Stack Ports
- Hybrid Ports
- Inline Network Ports
- Inline Tool Ports

6. (Optional) From the **Site** drop-down list, select a site to display the utilization percentage of the ports associated to the site.
7. From the **Display Total** drop-down list, select the number of items to be displayed. By default, the number of items selected for display is 5.
8. Click **OK**.

Least Utilized Traffic

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

The Least Utilized Traffic widget lists the ports with the node ID, port number, port alias, and the utilization percentage. The utilization percentage is displayed as either a table or a bar graph. By default, a table is displayed. You can click the arrow to change the display to a graph as shown in Figure 9-20 on page 142.

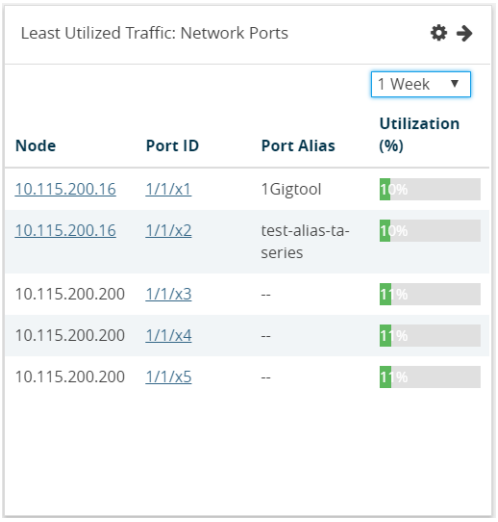

Figure 9-20: Least Utilized Traffic

The Least Utilized Traffic widget is configured exactly the same way as the Most Utilized Traffic widget. To configure the Least Utilized Traffic widget, refer to the configuration steps provided in *Most Utilized Traffic* on page 139. In **step 4**, select **Least Utilized Traffic** and click **OK**.

Inventory

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

Nodes by Model

The Nodes by Model widget displays the number of nodes managed by the current instance of GigaVUE-FM as a bar graph. Each bar in the graph indicates the number of each device model managed. Hovering the mouse over a bar in the graph displays the model name and the total number. Figure 9-21 on page 143 shows a Node by Model widget displaying six different nodes managed by GigaVUE-FM. Hover the mouse over the bar to view the number of devices in each node.

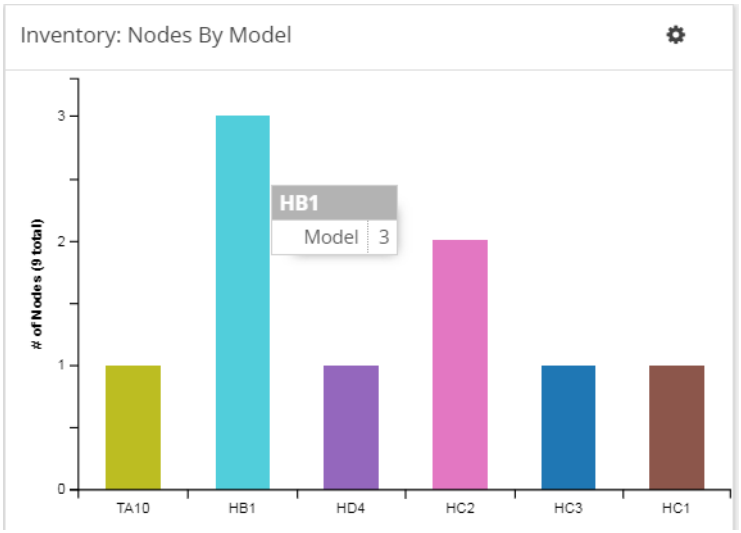

Figure 9-21: Nodes by Model

Nodes by Software Version

The Nodes by Software Version widget presents a graph that helps you to quickly view the software versions of the nodes that GigaVUE-FM is managing and the total percentage of each version. Each software version is assigned a color in the graph, which is specified by the legend. Hovering the mouse over an area in the graph displays the total number of software version used as a percentage. In Figure 9-22 on page 144, the Nodes by Software Version widget shows that there are 7 instances of version 5.1 and 2 instances of version 5.0, which is 22 percent of the total versions installed.

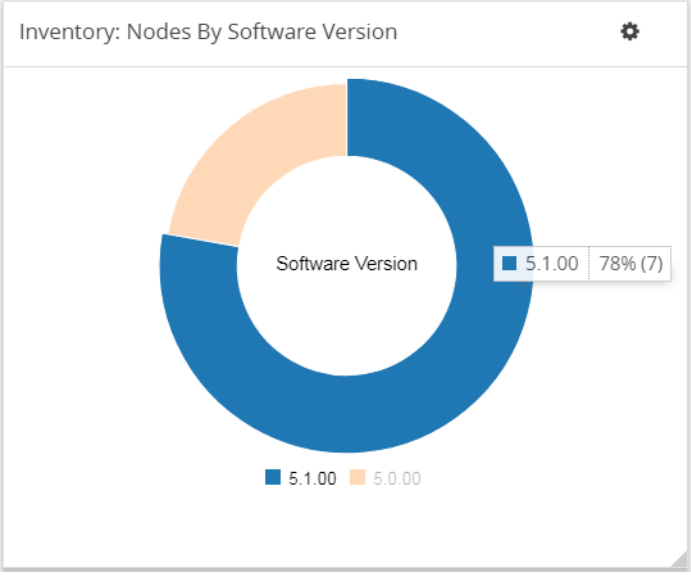

Figure 9-22: Nodes by Software Version

To configure the Inventory widget:

1. Click **Dashboard** on the top navigation link.
2. On the Physical & Virtual dashboard page, select the profile in which you want to add the widget.

3. Click **Add New Widget**. The Add New Widget window is displayed. Refer to Figure 9-23 on page 145.

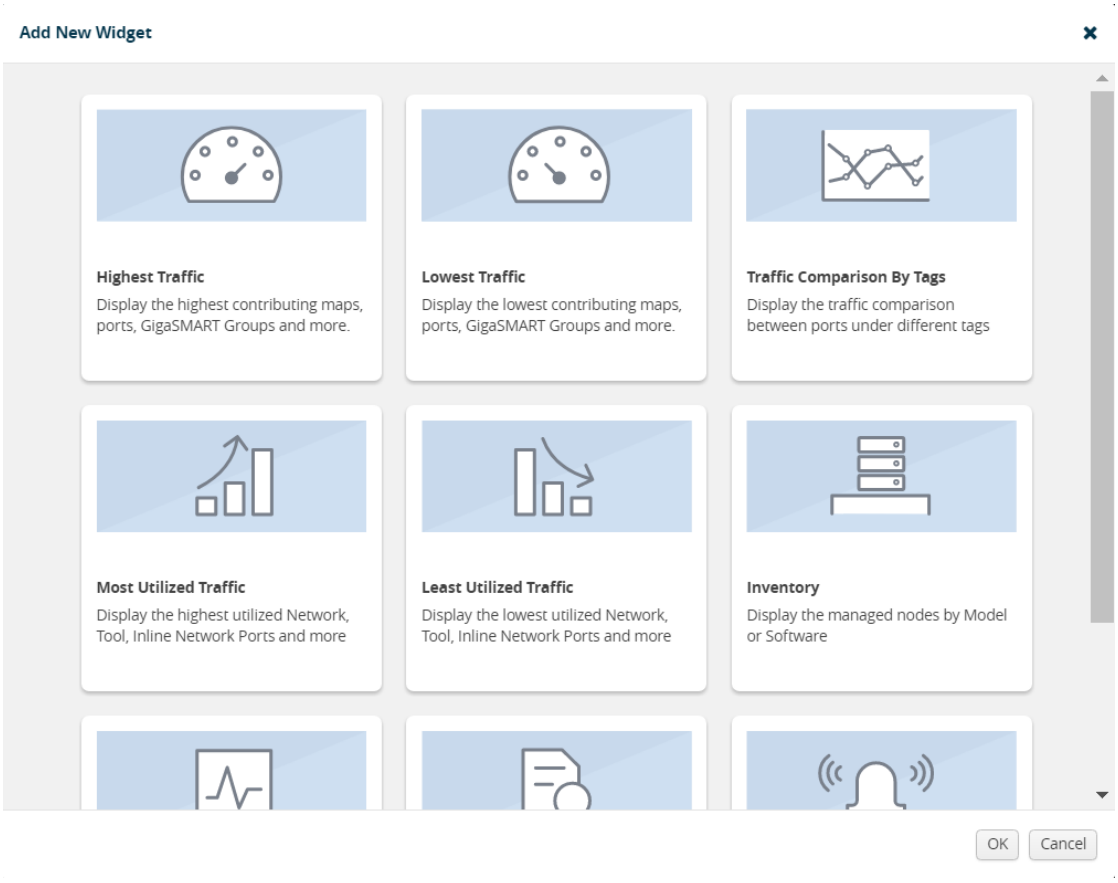

Figure 9-23: Add New Widget

4. In the Add New Widget window, select **Inventory** and click **OK**. The Inventory configuration window is displayed. Refer to Figure 9-19 on page 141.

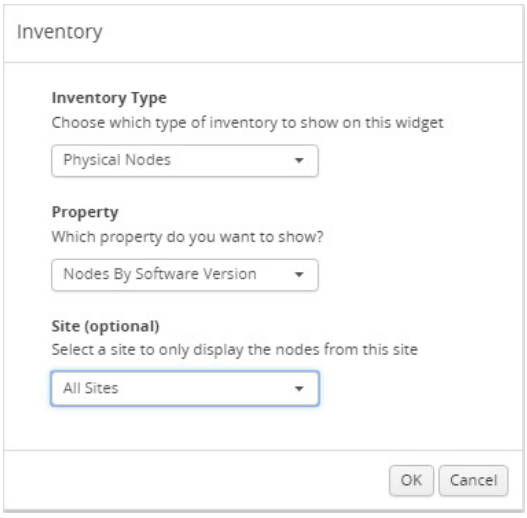

Figure 9-24: Inventory Configuration

5. From the **Inventory Type** drop-down list, select the Physical Nodes.

6. From the **Property** drop-down list, select one of the following:
   - Nodes by Model
   - Nodes by Software Version
7. (Optional) From the **Site** drop-down list, select a site to display the nodes associated to the site.
8. Click **OK**.

Status Summary

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

- *Nodes' Status Summary* on page 146
- *Port Link Status Summary* on page 147
- *Unhealthy Maps* on page 149
- *Unhealthy Flows* on page 149
- *Port Drops and Errors* on page 151

Nodes' Status Summary

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

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

For information about how the device status is computed, refer to *Node Health Status* on page 1361.

Hovering the mouse over an area in the graph displays the percentage of nodes in that status. In Figure 9-25 on page 147, the widget shows that there are 5 nodes in Normal status, 4 nodes in Warning status, and 5 nodes in Critical status. There are no nodes in Error status.

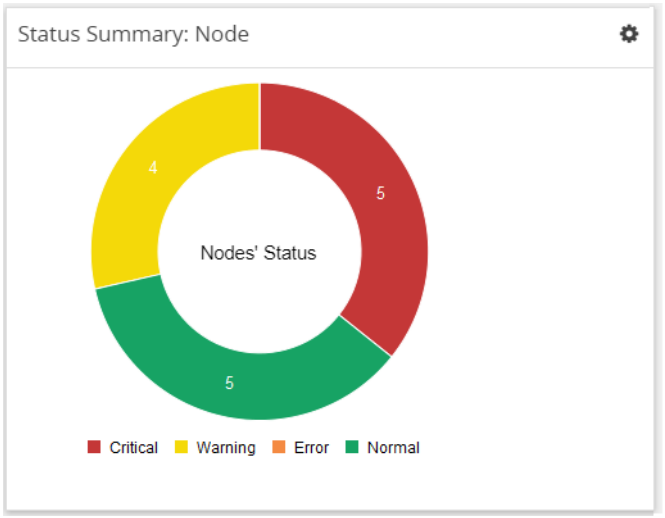

Figure 9-25: Nodes' Status Widget

Port Link Status Summary

The Port Link status summary widget allows you to view the current link status of all the ports available in the physical nodes currently managed by GigaVUE-FM. Optionally, the Port Link status can be displayed for a specified site. When a particular site is selected in the configuration window, the port link status of all the ports available in the nodes associated to the site are displayed in the Port Link Status Summary dashboard.

Refer to Figure 9-26 on page 147 for Status Summary: Port Links dashboard for all sites.

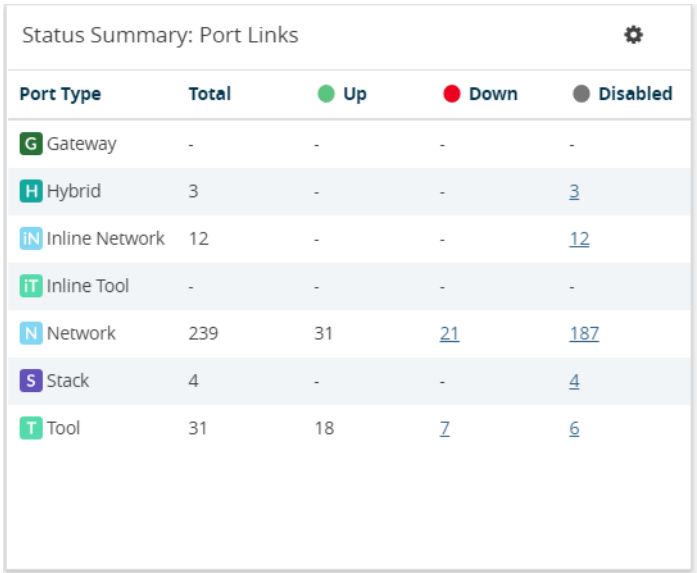

Port Type	Total	Up	Down	Disabled
G Gateway	-	-	-	-
H Hybrid	3	-	-	3
iN Inline Network	12	-	-	12
iT Inline Tool	-	-	-	-
N Network	239	31	21	187
S Stack	4	-	-	4
T Tool	31	18	7	6

Figure 9-26: Status Summary: Port Links Widget

Refer to Figure 9-27 on page 148 for Status Summary: Port Links dashboard for Santa Clara site.

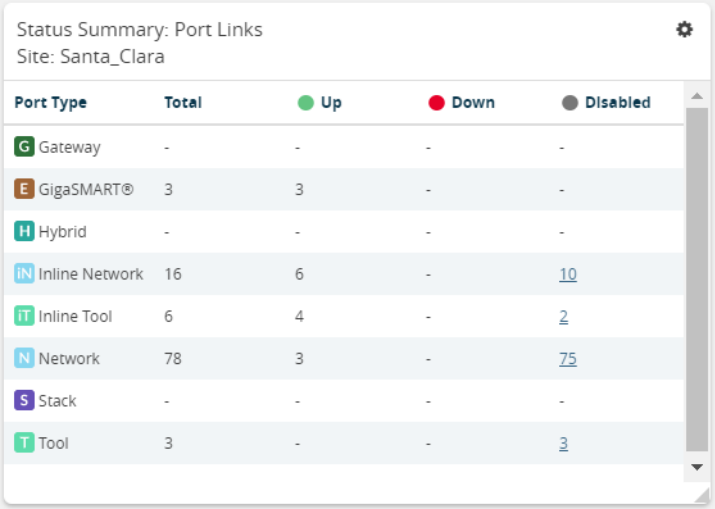

Figure 9-27: Status Summary: Port Links Widget For Santa Clara Site

The Status Summary: Port Links widget lists the following:

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

Click the numbers in the down or disabled column. A quick view provides detailed information with the cluster ID, device host name, port ID, and port alias of all the ports in the down or disabled state. If the port status is down, the quick view also provides information about the time since when the port has been in down state. For example, in Figure 9-28 on page 148, the stack port 1/1/q1 has been down since 10 days. The down time is displayed in minutes, hours, days, or months.

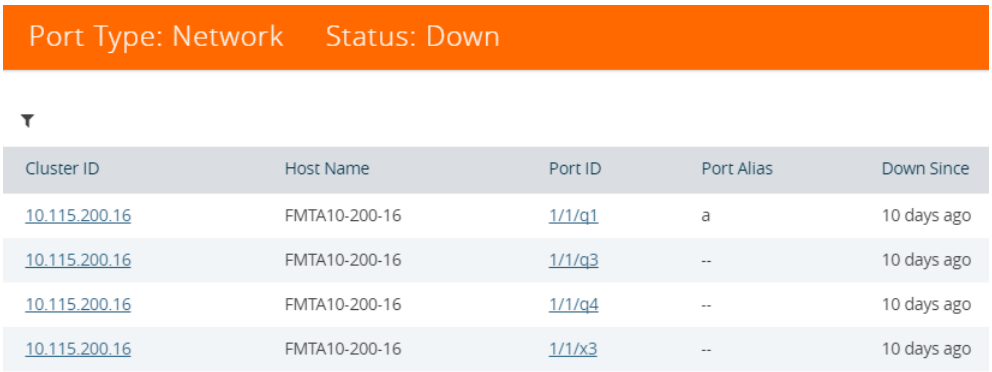

Figure 9-28: Port Type Quick View

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

NOTE: Prior to software GigaVUE-OS version 4.7, GigaVUE TA Series Traffic Aggregator nodes did not support tool ports. Starting in version 4.7, all gateway ports on GigaVUE TA Series nodes are tool ports. If GigaVUE-FM 3.5 or later is managing nodes running a software version earlier than 4.7, the Ports by Link Status widget may display the number of Gateway nodes.

Unhealthy Maps

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

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

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

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

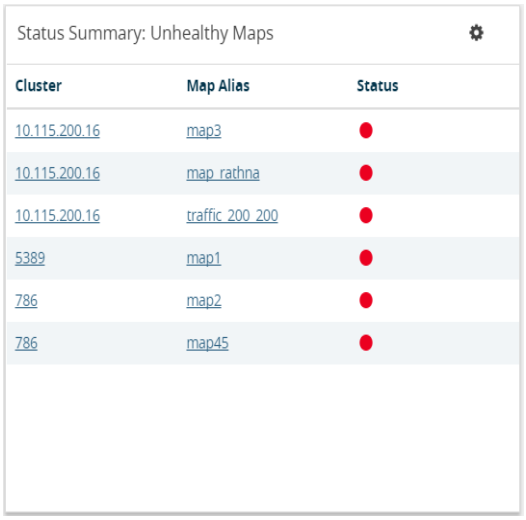

Figure 9-29: Top 10 Unhealthy Maps

Unhealthy Flows

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

A priority map group consists of one or more maps configured with the same source ports. The health of a priority map group is determined by the aggregated health of the

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

The Unhealthy Flows widget shows the names of the flows that are in unhealthy state and the names of the maps that are unhealthy in the flow. Refer to Figure 9-30 on page 150. Click the Flow Name to open the flow view page.

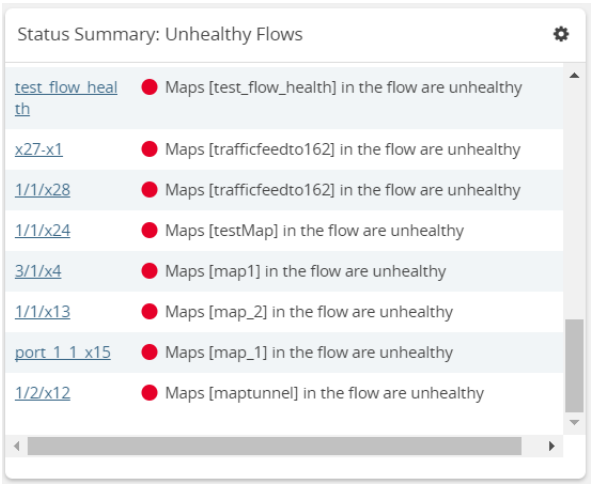

Figure 9-30: Top 10 Unhealthy Flows

For more information about Flows, refer to *Flows* on page 347.

Port Drops and Errors

The Port Drops and Errors status summary widget helps in identifying the ports with packet drops or packet errors in the network. When a particular site is selected in the configuration window, the status summary widget lists the port types associated to the site and the number of ports having packet drops, transmitting errors, or receiving errors in the site. You can view the number of ports with packet drops or packet errors occurring on a daily or an hourly basis. Refer to Figure 9-31 on page 151.

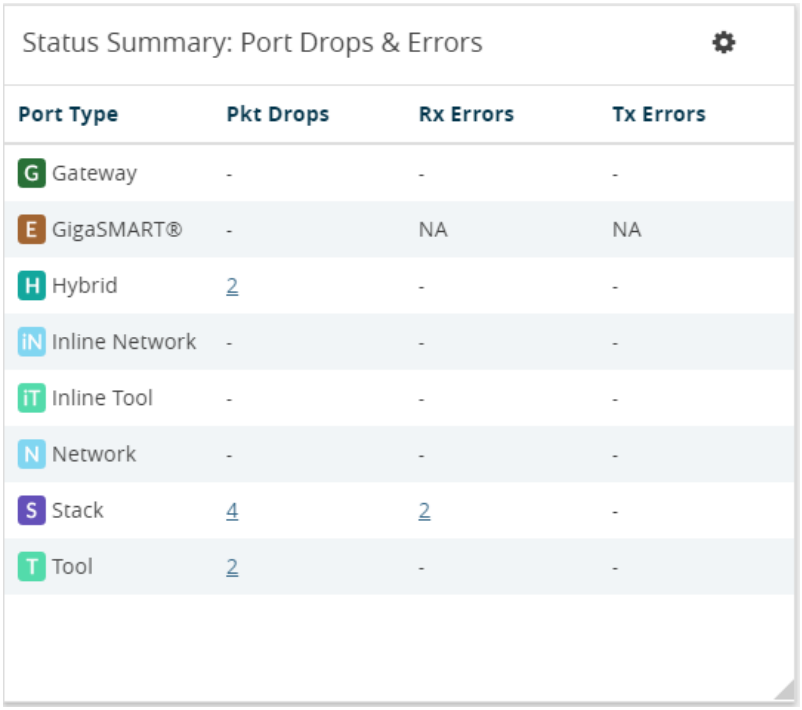

Figure 9-31: Port Drops and Errors

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

To view detailed information about the port drops and errors, click the number in the Pkt Drops, Rx Errors, or Tx Errors column. A quick view displays the cluster ID, host name, port ID, port alias, and the number of packet drops or errors for the list of unhealthy ports in the port type. Refer to Figure 9-32 on page 152. If there are too

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

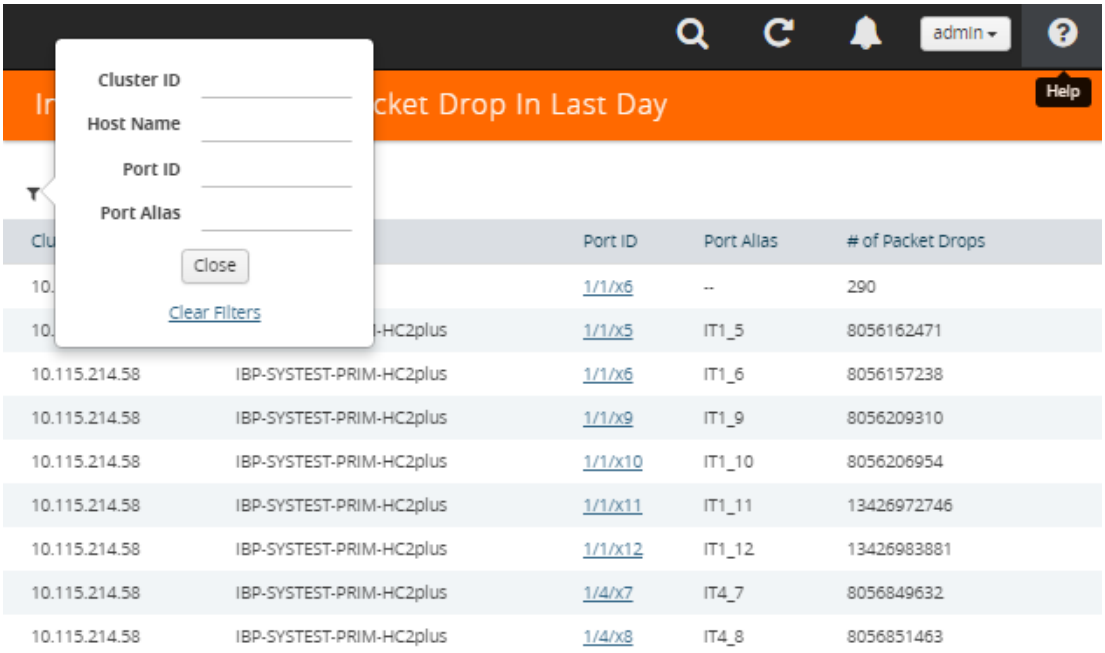

Figure 9-32: Unhealthy Ports Quick View

In the Unhealthy Ports quick view, click the port ID for a detailed view of the type of packet errors or packet drops occurring on a daily or an hourly basis. Refer to Figure 9-33 on page 153. You can also view the related maps, transceiver type, speed,

and other detailed information about the port, which helps to investigate the reason for the packet drops or packet errors. To return to the Ports quick view, click **Back**.

Figure 9-33: Ports Quick View

To configure the Inventory widget:

1. Click **Dashboard** on the top navigation link.
2. On the Physical & Virtual dashboard page, select the profile in which you want to add the widget.
3. Click **Add New Widget**. The Add New Widget window is displayed. Refer to Figure 9-9 on page 133.

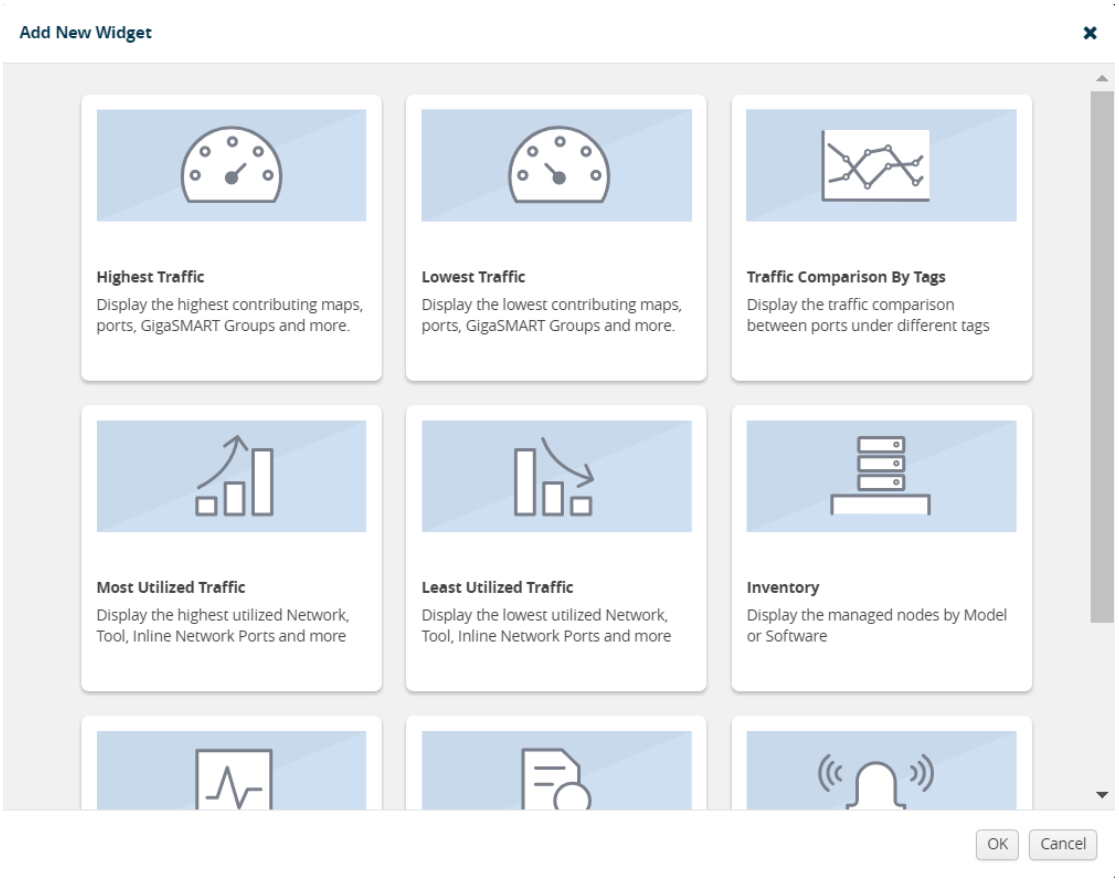

Figure 9-34: Add New Widget

4. In the Add New Widget window, select **Status Summary** and click **OK**. The Status Summary configuration window is displayed. Refer to Figure 9-19 on page 141.

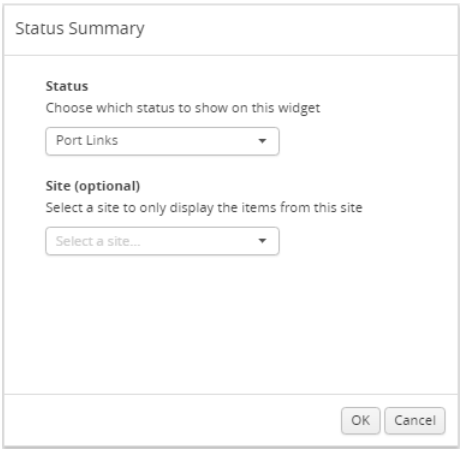

Figure 9-35: Status Summary Configuration

5. From the **Status** drop-down list, select one of the following:
   - Node—For information, refer to *Nodes' Status Summary* on page 146.
   - Port Links—For information, refer to *Port Link Status Summary* on page 147.
   - Unhealthy Maps—For information, refer to *Unhealthy Maps* on page 149.
   - Port Drops & Errors—For information, refer to *Port Drops and Errors* on page 151.
6. (Optional) From the **Site** drop-down list, select a particular site.
7. Click **OK**.

Audit Logs

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

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

- Success (green)
- Failure (red)

Hovering the mouse over an area in the graph displays the percentage of audit logs in that result. You can also specify the audit log statuses that have occurred over the past hour, day, week, or month. Figure 9-36 shows the audit log results over each of the time periods.

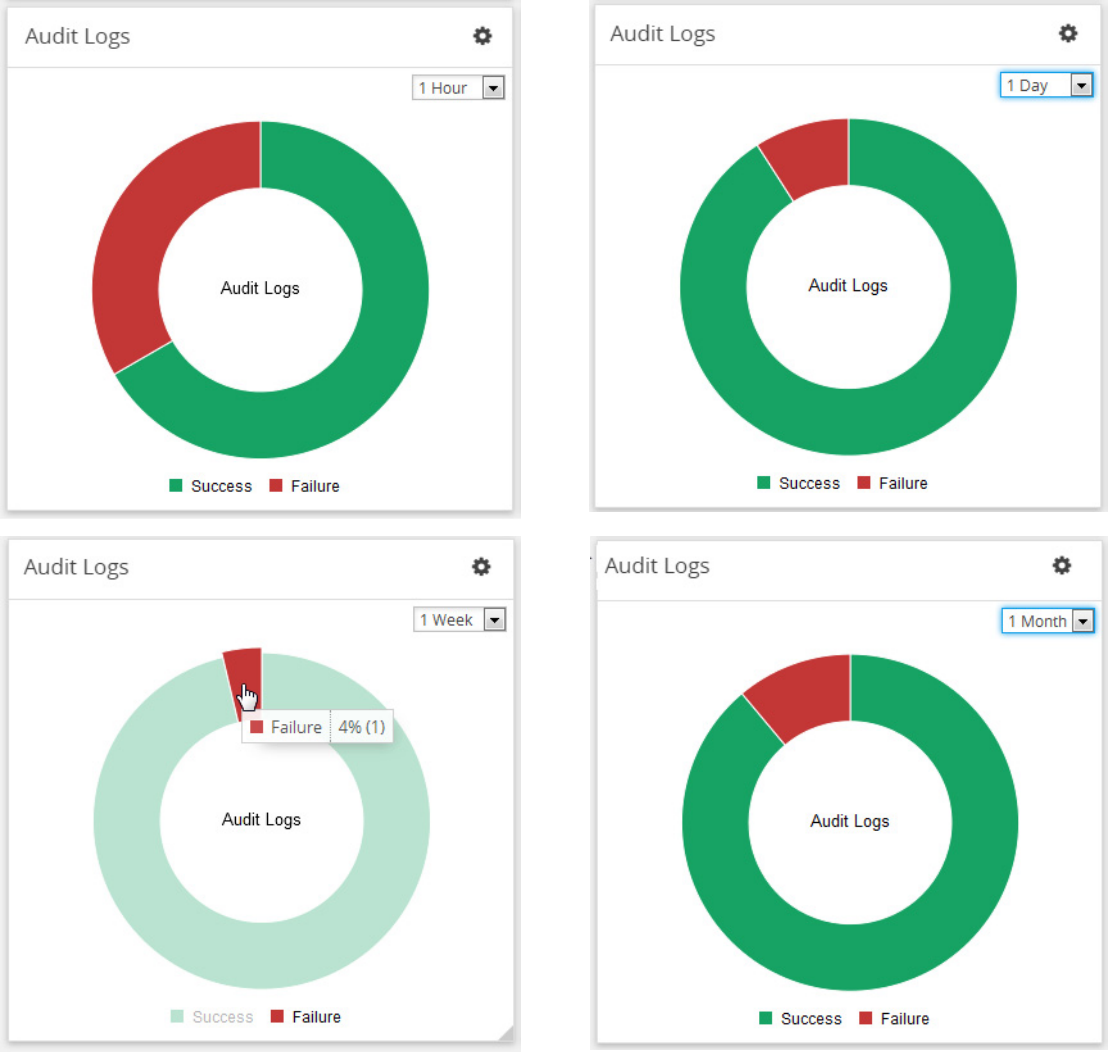

Figure 9-36: Audit Logs by Result

In this example, the Audit Logs widget shows that there is a log with the failure status that has occurred in the last hour. When you go to the audit logs page, you can see the entries shown in Figure 9-37 on page 156, which matches with the information displayed in Audit Logs widget: two successes and one failure due to an incorrect log in.

Audit Logs

Filter | Manage

Time	User	Operation Type	Source	Status	Description
2016-06-14 12:29:38	admin	login fmUser admin	FM	SUCCESS	
2016-06-14 12:29:29	admin	login fmUser admin	FM	FAILURE	Login failed
2016-06-14 12:29:12	admin	logout fmUser admin	FM	SUCCESS	

Figure 9-37: Audit Log Entries

To configure the Audit Logs widget:

1. Click **Dashboard** on the top navigation link.
2. On the Physical & Virtual dashboard page, select the profile in which you want to add the widget.
3. Click **Add New Widget**. The Add New Widget window is displayed. Refer to Figure 9-9 on page 133.

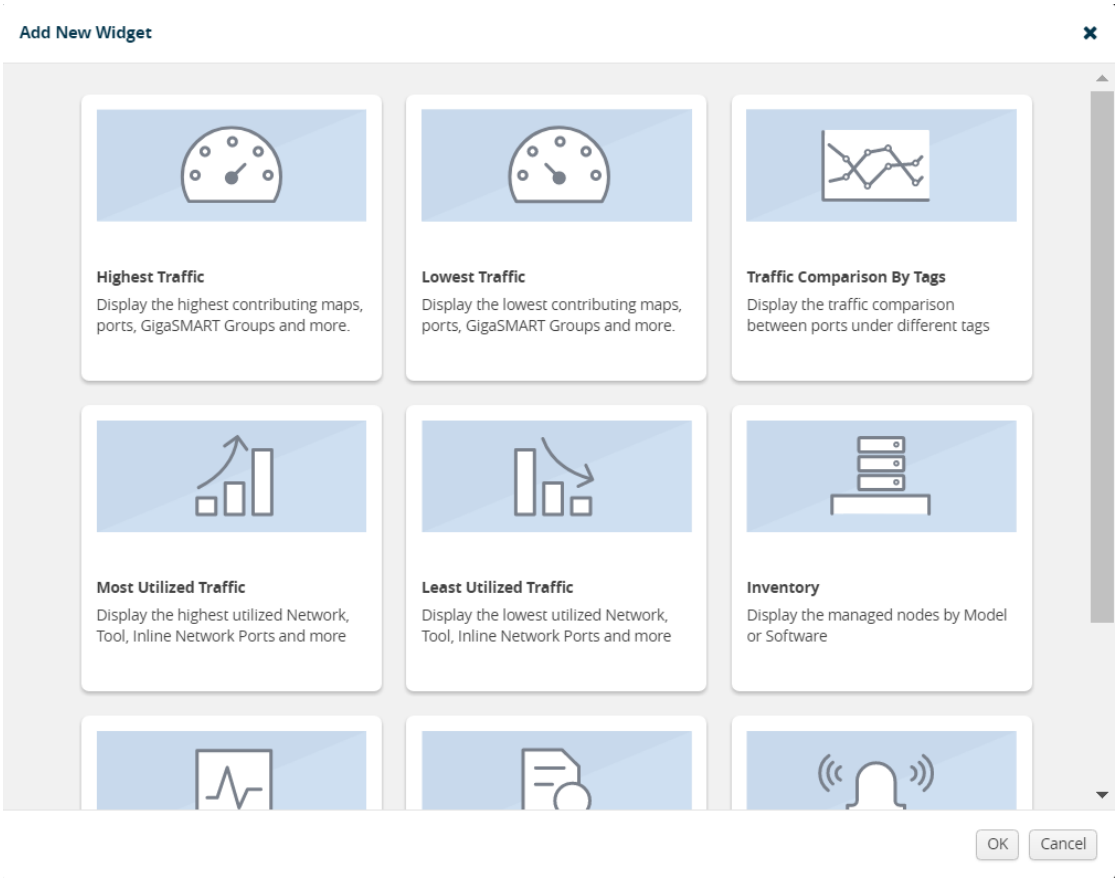

Figure 9-38: Add New Widget

4. In the Add New Widget window, select **Audit Logs** and click **OK**. The Audit Logs configuration window is displayed. Refer to Figure 9-19 on page 141.

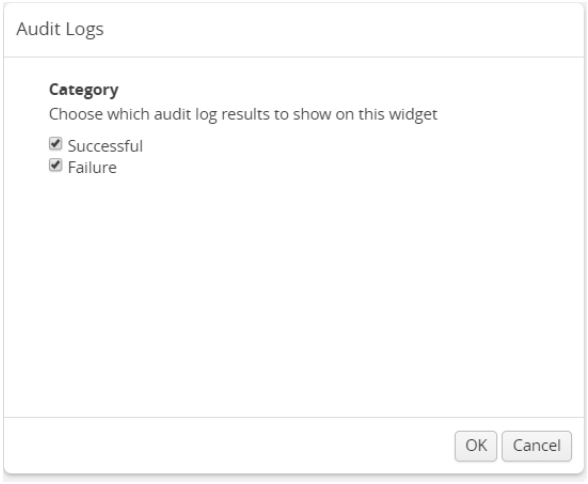

Figure 9-39: Status Summary Configuration

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

Events

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

The possible severity levels are:

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

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

Figure 9-40 shows the number of events that have occurred in the past week in each severity level for all sites. If you want more detail about the events, select **Alarms/ Events** in the Physical page.

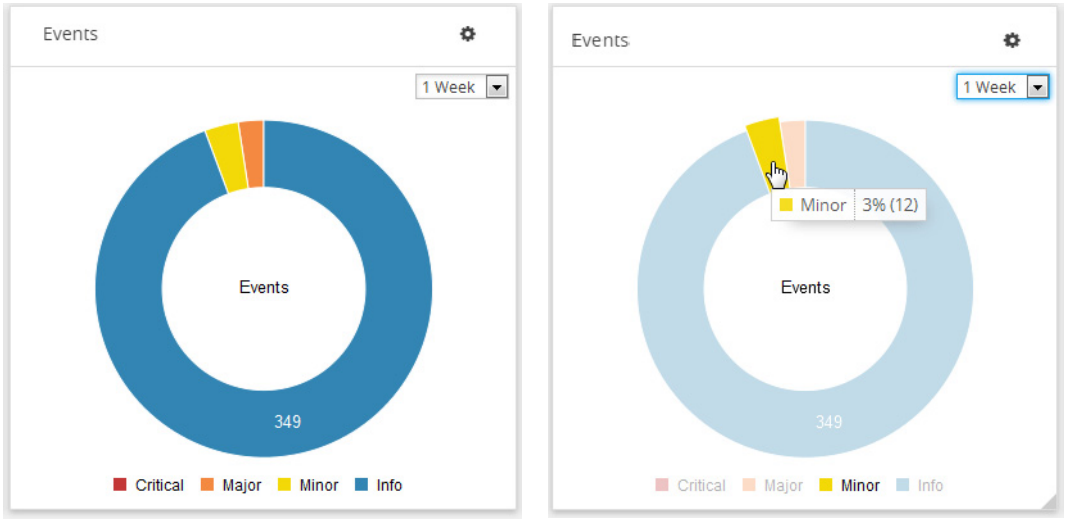

Figure 9-40: Events Widget

Figure 9-40 shows the number of events that have occurred in the past week in each severity level for Santa Clara site.

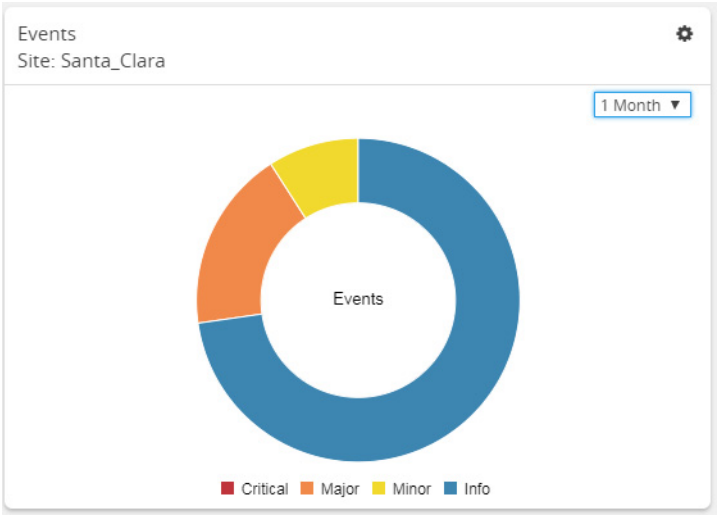

Figure 9-41: Events Widget for Santa Clara

To configure the Events widget:

1. Click **Dashboard** on the top navigation link.
2. On the Physical & Virtual dashboard page, select the profile in which you want to add the widget.
3. Click **Add New Widget**. The Add New Widget window is displayed. Refer to Figure 9-9 on page 133.

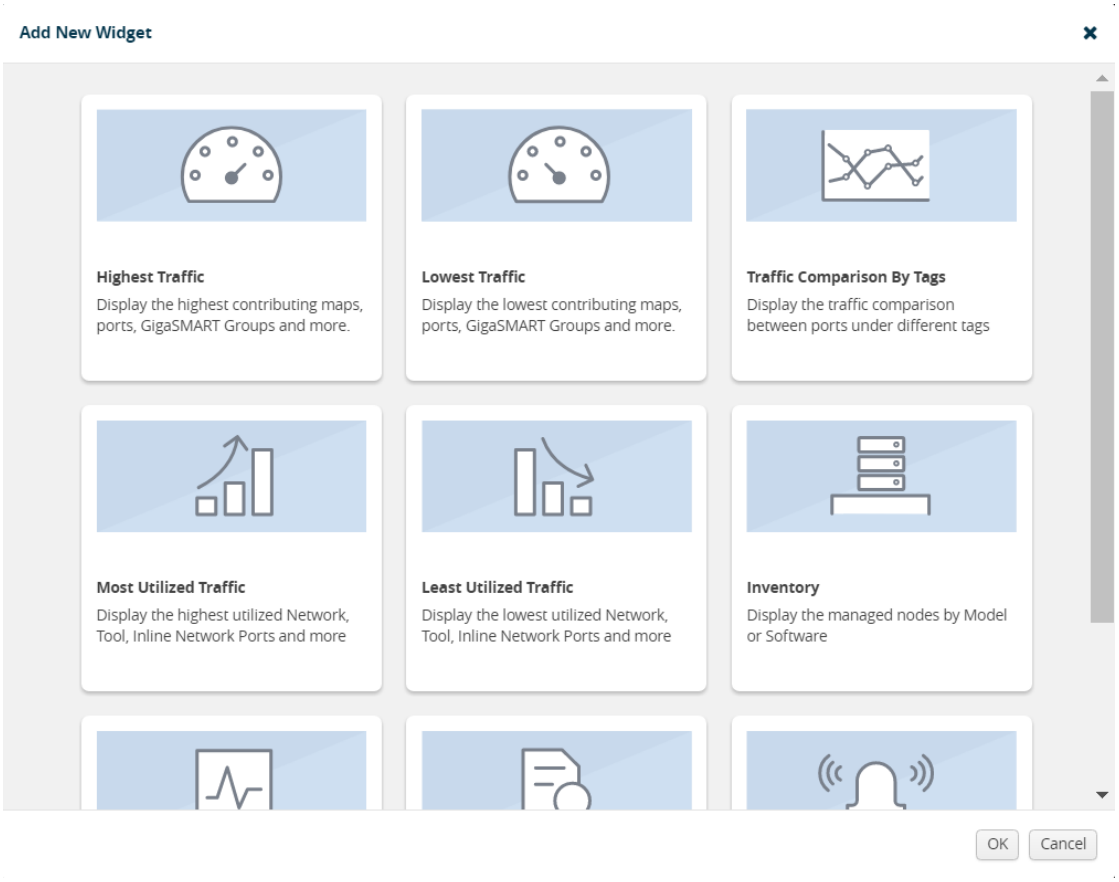

Figure 9-42: Add New Widget

4. In the Add New Widget window, select **Events** and click **OK**. The Events configuration window is displayed. Refer to Figure 9-19 on page 141.

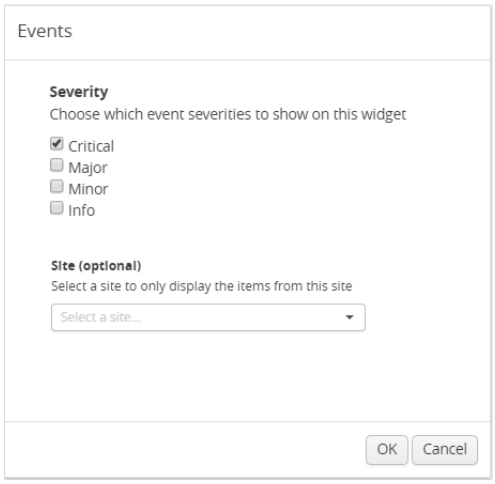

Figure 9-43: Status Summary Configuration

5. Choose the event you want to view in the widget:
   - Critical
   - Major
   - Minor
   - Info
6. (Optional) From the Site drop-down list, select a particular site.
7. Click **OK**.

10 Health Monitor Dashboard

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

This chapter covers the following topics:

Overview of the Health Monitor Dashboard

GigaVUE-FM is the central management appliance for the visibility fabric. Therefore, knowing its current health is important in order to maximize the availability of the appliance. Staring in GigaVUE-FM 3.3, a Health Monitor is available that provides health information about GigaVUE-FM. Refer to Figure 10-1 on page 164. The Health Monitor makes it possible helps do the following:

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

The Health Monitor provides the following monitors:

- CPU utilization
- Memory utilization
- Disk utilization

- Service status

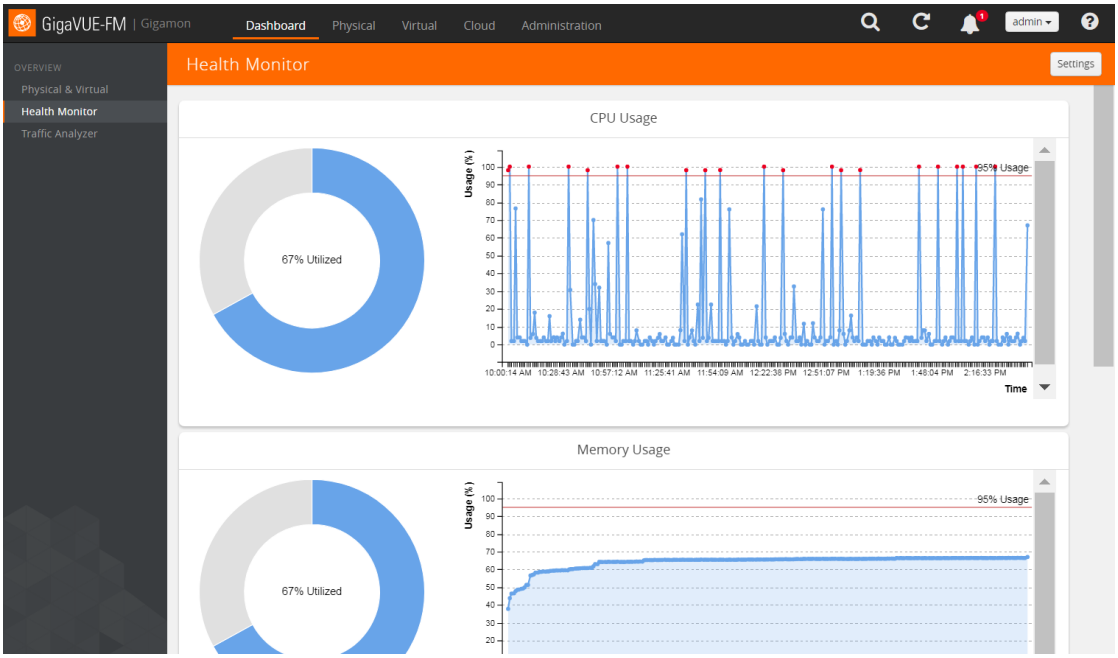

Figure 10-1: Health Monitor Dashboards

CPU Utilization

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

The CPU Utilization Monitor display utilization as a donut and time charts. The donut chart show percentage of utilized and available CPU. The time chart shows utilization at specific intervals. By clicking on a point in the time chart, you can see the utilization at a specific point in time. In Figure 10-2, the CPU utilization at 10:49:55 AM is 1.0 percent.

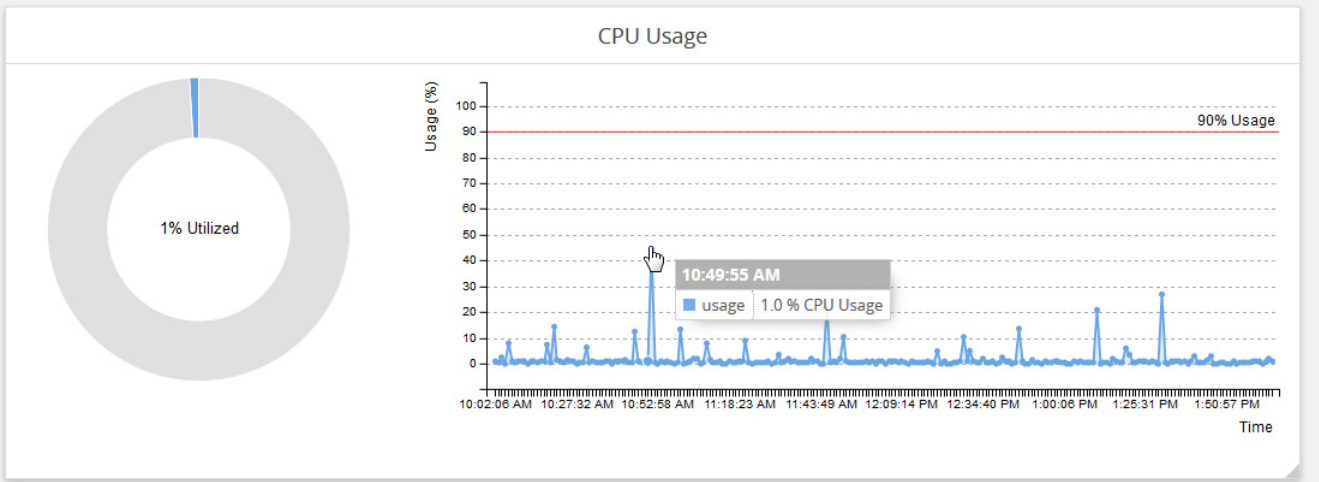

Figure 10-2: CPU Utilization Monitor

Memory Utilization

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

The Memory Utilization Monitor display utilization as a dough nut and time charts. The dough nut chart show the percentage of utilized and available memory. The time chart shows utilization as specific intervals. By clicking on a point in the time chart, you can see the utilization at a specific point in time. In Figure 10-3, the memory utilization at 11:48:49 AM is 17percent.

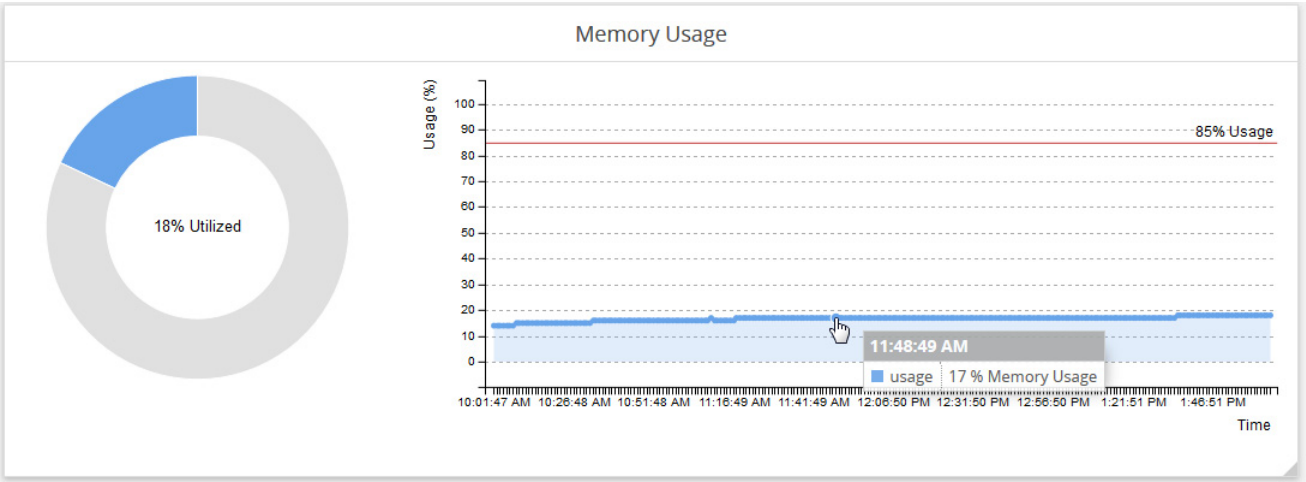

Figure 10-3: Memory Utilization Monitor

Disk Utilization

The Disk Utilization Monitor displays disk usage levers over time for individual partitions, providing information about peak disk usage for FM logs and FM data. This provides information that can help prevent outages due to disk out-of-space issues. For example, a partition using more than 75 percent of allocated space is generally a warning sign and utilization of over 90 percent of allocated space could justify sending an alert.

The Disk Utilization Monitor display utilization as two bar graphs, one for FM logs and one for FM data, and a time chart. The bar charts show the percentage of disk utilization in the partitions for FM logs and FM data as well as total number of Gigabytes available versus used. The time chart shows utilization as specific intervals for both partition. By clicking on a point in the time chart, you can see the utilization at a specific point in time. In Figure 10-3, the disk usage at 11:16:01 AM for FM logs is 18.5 percent and the disk usage for FM data is 12%.

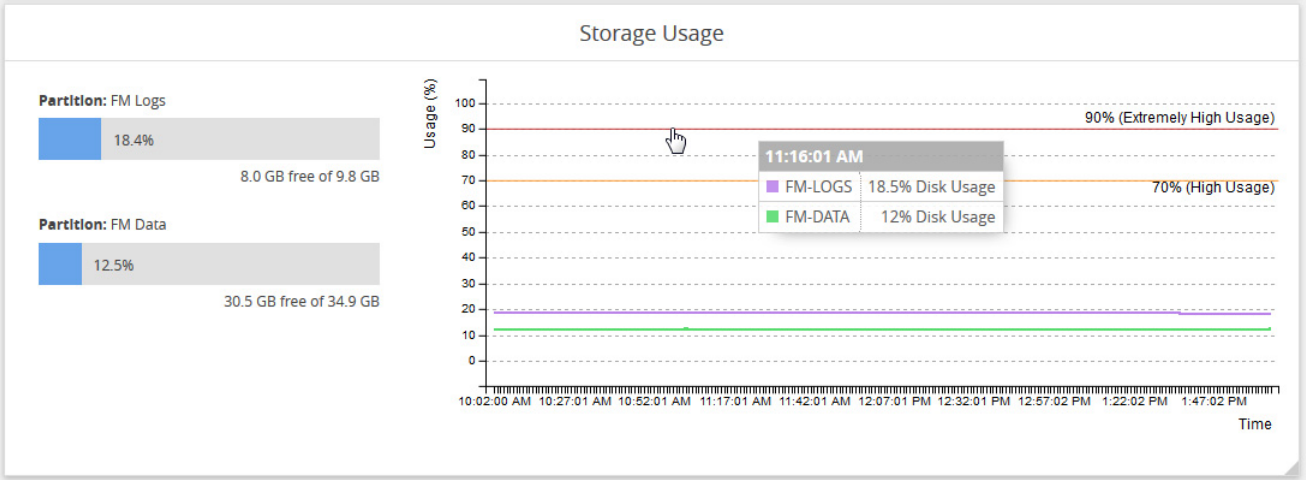

Figure 10-4: Disk Utilization Monitor

Services

The Service Monitor provides the current status of GigaVUE-FM services, indicating whether they are Up or Down. The services monitored are:

- NetFlow records storage and collection
- Time-series data storage and access
- Fabric configuration storage
- Search engine

In Figure 10-5, all of the services are currently up.

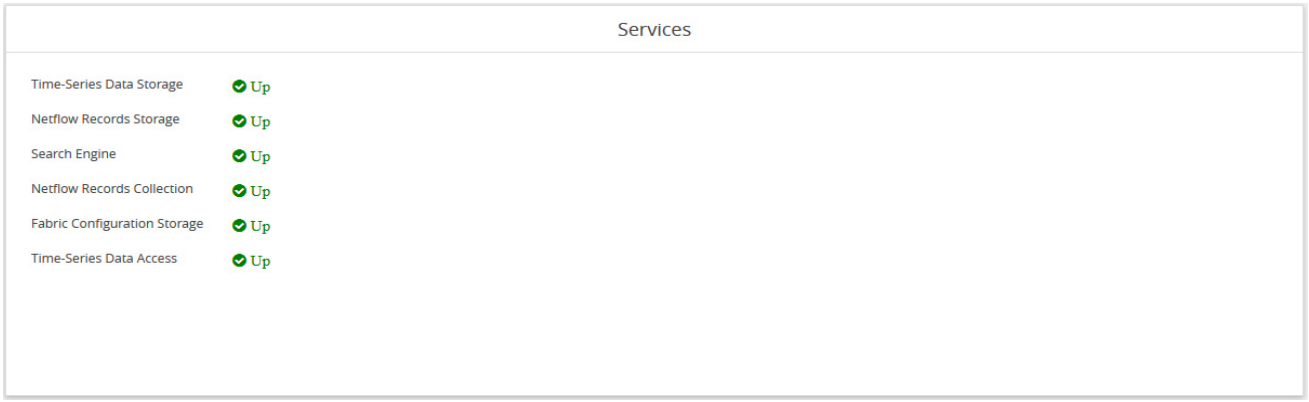

Figure 10-5: Services Monitor

How to Set Health Monitor Alarm Thresholds and Notifications

For CPU, memory, and disk utilization monitoring you can define thresholds that cause an alert if the threshold is exceeded. You can also set up email notifications.

Set Alarm Thresholds

To set thresholds for the Health Monitor, do the following:

1. Click **Dashboard** on the top navigation link.
2. On the Health Monitor Dashboard page, click **Settings**.
3. On the Health Monitor Thresholds Settings page, select the percentages for CPU, Memory, and Disk utilization that will be the threshold for an alarm.
4. Click **Save**.

After you have set the utilization thresholds, the threshold is displayed as a red line on the time chart of the monitors. For example, in Figure 10-6, the threshold for CPU utilization is set to 85 percent and the red line is displayed at 85% on the chart. Also, when a specified threshold has been exceed, the Alarm/Events page shows an alarm. For more information about Alarm and Events, refer to *All Alarms/Events* on page 1269.

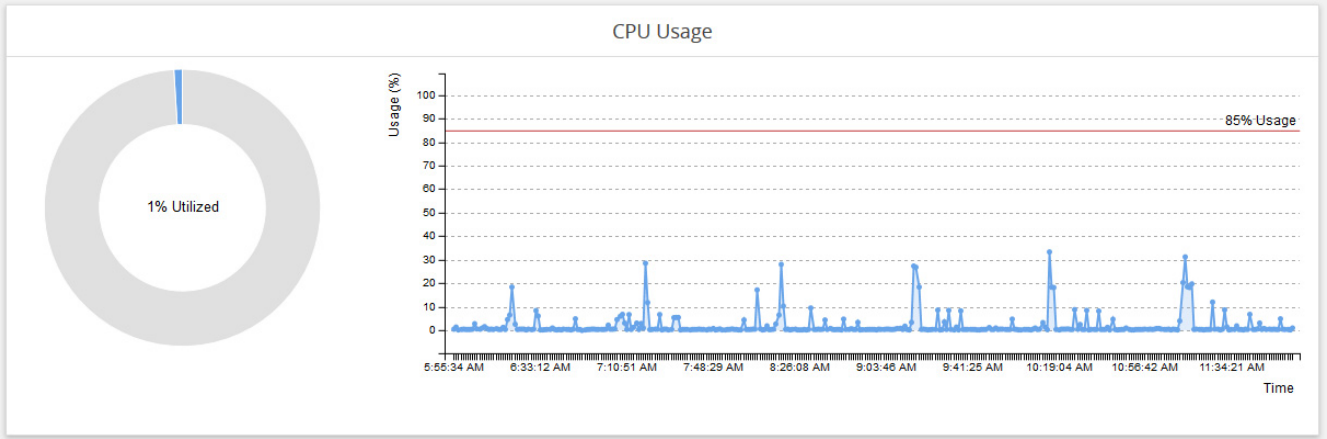

Figure 10-6: CPU Utilization with Threshold Set

Set Notifications for Health Monitor Thresholds

You can set GigaVUE-FM to send email notifications when ever a utilization threshold is reached by one of the health monitors. To set the threshold notification, do the following:

1. Click **Administration** on the top navigation link.
2. Select **System > Notifications**.
3. Select one or more health monitor notifications.
4. Click **Configure**.
5. Enter the email address for the recipient or recipients for the notifications.
6. Click **Save**.

For more detailed information about Notifications, refer to *Notifications* on page 1320.

11 FabricVUE Traffic Analyzer

This chapter describes the FabricVUE Traffic Analyzer Dashboard that is available to users with the prime package license. If you have other packages, you can purchase the license for FabricVUE Traffic Analyzer. Also, to use Traffic Analyzer, the recommended minimum configuration for GigaVUE-FM is 2 vCPU, 8GB RAM, and an 80GB HDD.

NOTE: To track traffic patterns using the FabricVUE Traffic Analyzer, you will need to purchase and enable the NetFlow license, which is offered as part of the GigaSMART licenses. Ensure that the NetFlow license is installed on the node that you want to view.

This chapter covers the following topics:

Overview of FabricVUE Traffic Analyzer

Traffic Analyzer gives you a larger insight into traffic flows through the GigaVUE nodes, allowing you to gain visibility into the type of traffic that may not be flowing to the tool. This allows you to see the overall traffic passing through your network and how to modify flow mapping to accommodate any changes in the traffic patterns.

Traffic Analyzer Widgets

Traffic Analyzer displays the following top N values, where N is the value (5, 10, 15, 20, 50, or 100) specified in the configuration settings for the Traffic Analyzer:

- Top N conversations, applications, protocols and endpoints.
- Top N Gigamon Interfaces
- Top N Endpoints
- Top N External interfaces
- Top N Visibility Fabric Interfaces
- Top N URLs
- Top N endpoints with HTTP Errors (401, 402, 403)

Notes:

- The Top N external interfaces are discovered by enabling CDP/LLDP on the Gigamon ports.
- The Top N URLs and Top N Endpoints with Error Codes are only supported with IPFIX and requires Gigamon private extensions to be enabled in the Netflow Record configuration.

The traffic analyzer displays the top 5 to top 20 statistics from the aggregated NetFlow data. The analyzer can also be used to see the historical trending data from 1 hr to 1 month. The remaining data is not available for viewing.

Each instance of Traffic Analyzer shows data from the aggregated database of all physical nodes that are configured for NetFlow and tunneled to GigaVUE-FM.

The data displayed in the Top N widgets of the traffic analyzer can be selected and viewed as a time-series trend in the quick view.

Figure 11-1 on page 171 shows the traffic between two end points over a 1 week period while the Figure 11-2 on page 172 shows the TCP data from the Top Protocols traffic trending over 1 month.

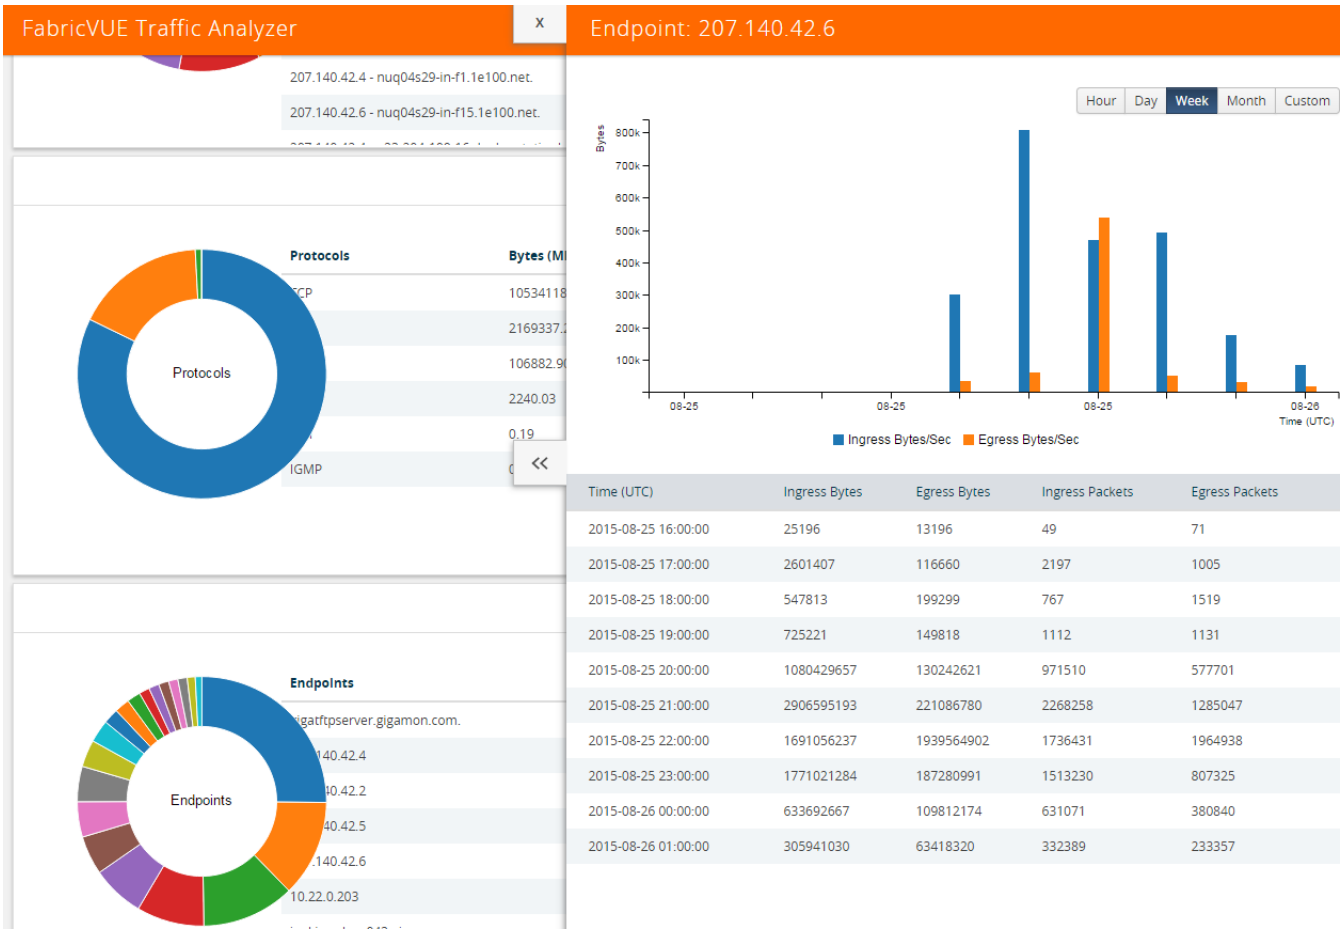

Time (UTC)	Ingress Bytes	Egress Bytes	Ingress Packets	Egress Packets
2015-08-25 16:00:00	25196	13196	49	71
2015-08-25 17:00:00	2601407	116660	2197	1005
2015-08-25 18:00:00	547813	199299	767	1519
2015-08-25 19:00:00	725221	149818	1112	1131
2015-08-25 20:00:00	1080429657	130242621	971510	577701
2015-08-25 21:00:00	2906595193	221086780	2268258	1285047
2015-08-25 22:00:00	1691056237	1939564902	1736431	1964938
2015-08-25 23:00:00	1771021284	187280991	1513230	807325
2015-08-26 00:00:00	633692667	109812174	631071	380840
2015-08-26 01:00:00	305941030	63418320	332389	233357

Figure 11-1: FabricVUE Traffic Analyzer Quick View Window For Endpoints

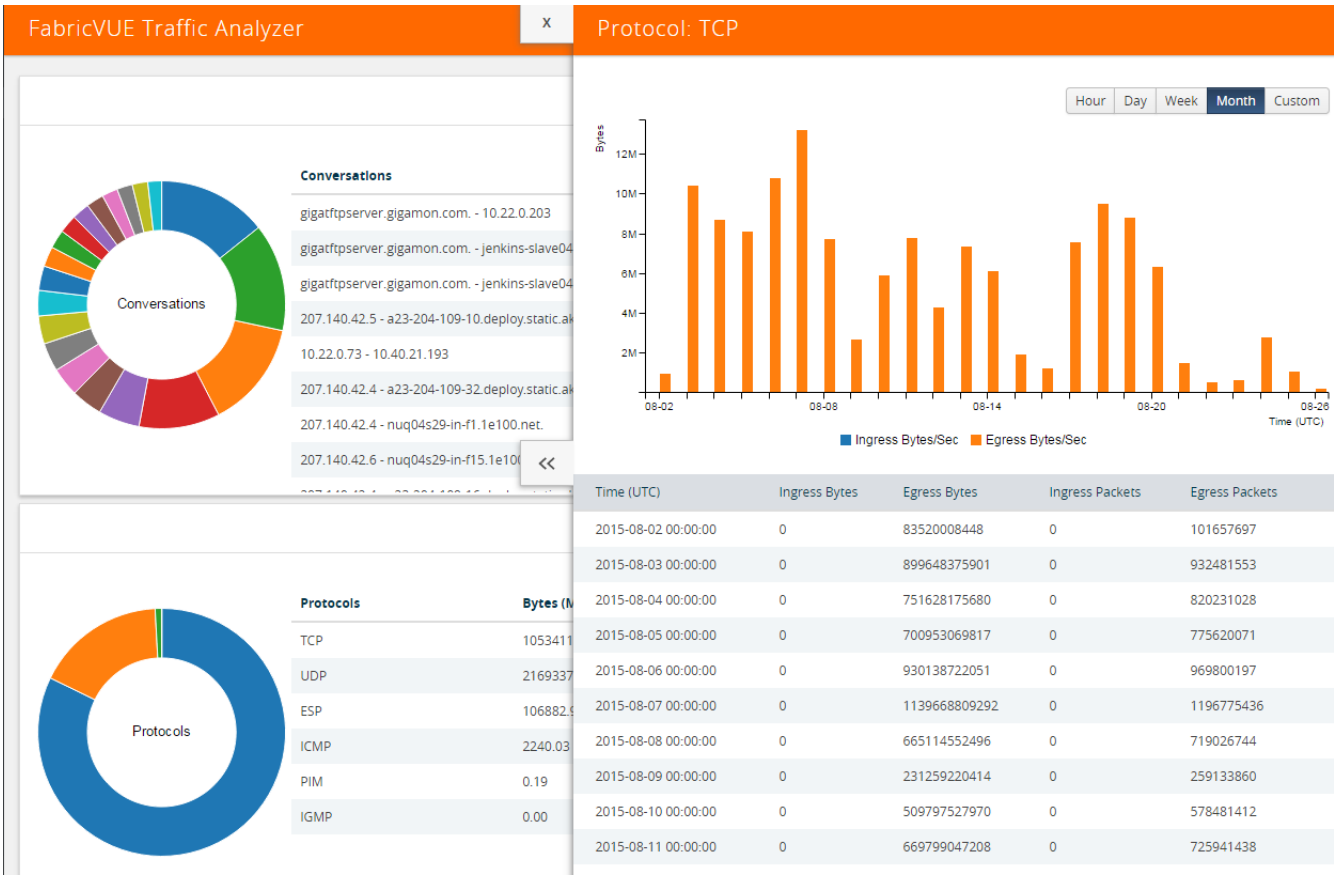

Figure 11-2: FabricVUE Traffic Analyzer Quick View Window for Endpoints

The Top N Conversations and Top N URL widgets include a search feature. The Top N Conversation allows you to search for conversations between two IP addresses. The Top N URL allows you to search for a specific URL. Figure 11-3 show an example of traffic from the Top 10 URLs.

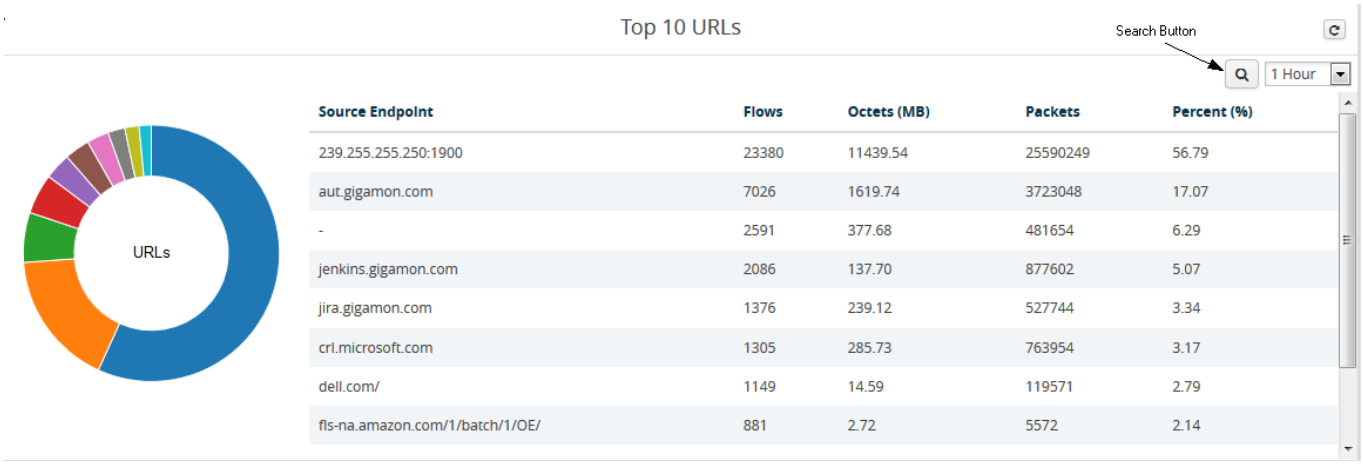

Figure 11-3: Top 10 URLs

Click the Search Button on the widget to open the Search URLs page. To search for a specific source endpoint, enter a search term in the URL Contains field, and then click

the **Search** button. The URLs that contain the search term will be displayed. To return to the Traffic Analyzer, click **Close**.

On returning to the Top N URLs widget, the widget displays the information based on the specified search.

Create Data for FabricVUE Traffic Analyzer

Before setting up the Traffic Analyzer, you will need to set up NetFlow, which requires a GigaSMART NetFlow license, and create an IP interface with GigaVUE-FM as the endpoint. Table 11-1 summarizes the NetFlow records needed for the FabricVUE Traffic Analyzer.

Table 11-1: NetFlow Records for FabricVUE Traffic Analyzer Widgets

FabricVUE Traffic Analyzer Widget	Netflow Records Needed
Top 10 Conversations	IP Source, IP Destination, Packets, Protocols
Top 10 Applications	Protocols, Source Port, Destination Port, Bytes, Packets
Top 10 Protocols	Protocols
Top 10 Endpoints	IPv4 Source Address, Bytes, Protocols
Top 10 External Interfaces (CDP/LLDP)	Interface Name **Note:** For this widget, CDP/LLDP needs to be enabled on the port that is receiving traffic on the node and sending CDP/LLDP.
Top 10 Visibility Fabric Interfaces	IPv4 Address, Bytes, Packets, Interface Name
Top 10 URLs	URL, IPv4 Address, Bytes, Packets
Top 10 Endpoints with Error Code 401	Bytes, Packets, IPv4 Address, HTTP Response Code
Top 10 Endpoints with Error Code 402	Bytes, Packets, IPv4 Address, HTTP Response Code
Top 10 Endpoints with Error Code 403	Bytes, Packets, IPv4 Address, HTTP Response Code

NOTE: You need an active NetFlow GigaSMART license enabled on the physical node to use the FabricVUE Traffic Analyzer.

NOTE: Bytes, Packets, Protocol, Source IP, Destination IP, Source Port, and Destination Port should always be added. These are required in most cases.

To set up the NetFlow for the Traffic Analyzer, use the following steps or you can also use Workflows to set up NetFlow. For more information about Workflows, refer to *Workflows* on page 240.

1. Click **Physical** on the top navigation link.
2. On the Physical Nodes page, click the IP address of the node (which has NetFlow licensed) that you want to review for traffic patterns using FabricVUE Traffic Analyzer. This Overview page of the node is displayed.
3. Configure a GigaSMART Group. If you want to use an existing GigaSMART Group, you can go to *Step 4*.

   The GigaSMART Group needs to be setup before creating an IP interface.

   a. Select **GigaSMART > GigaSMART Groups > GigaSMART Groups** and click **New** to create a new GigaSMART Group.

   b. Enter an **Alias** and select an engine port in the **Port List** field.

   c. Click **Save**.
4. Crete an IP interface with a tool port.

a. Select **Ports > All Ports**.

b. Use the Quick Port Editor to configure a port as a tool port.

c. Select **IP Interfaces > IP Interfaces** and click **New** to setup the tool port configured in *Step b* as the IP interface.

5. Configure NetFlow.

a. Select **GigaSMART** > **NetFlow / IPFiX Generation**, and then do the following:

b. Configure the **NetFlow Exporter**.

When you configure the NetFlow Exporter, make sure to do the following:

•Set the Destination IP to the address of the GigaVUE-FM IP that will receive the data.

•Set the Transport Protocol to UDP.

•Set the Destination Port to 2055.

c. Configure the **NetFlow Record**.

The minimum **Key Field (Match)** to select are as follows:

•Select IPv4 and enable **Protocol**, **Source Address**, and **Destination Address**.

•Select **Transport** and enable **Source Port** and **Destination Port**.

The minimum **Non-Key Fields (Collect)** to select are as follows:

•Select **Counter** and enable **Bytes** and **Packets**.

•Select **IPv4** and enable **Protocol**, **Source Address**, and **Destination Address**.

•Select **Transport** and enable **Source Port** and **Destination Port**.

d. Configure the **NetFlow Monitor**.

e. Add the NetFlow monitor to the GigaSMART Group configured in *Step 3*.

6. Configure a GigaSMART Operation with NetFlow using the GigaSMART Group from Step 3.

a. Select **GigaSMART Operations > GigaSMART Operations**, and then click **New**.

b. For **GigaSMART Groups**, select the GigaSMART Group from *Step 3*.

c. For the **GigaSMART Operations**, select and enable **NetFlow**.

7. Select **Maps > Maps > Maps**.

Create a By Rule map from a single or multiple network ports to the (GigaVUE-FM) IP interface.

8. Return to GigaVUE-FM and select **Traffic Analyzer** in the Navigation pane.

You can view the traffic patterns.

Part 4: Physical

This section describe management information and activities for the physical nodes that you can perform with GigaVUE-FM. The following topics are covered:

12 GigaVUE Nodes and Clusters

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

GigaVUE H Series and TA Series Overview

The GigaVUE H Series and TA Series delivers performance and intelligence in each of its Traffic Visibility Fabric nodes, with port density and speeds that scale to your needs, from 1Gb to 100Gb. With an intuitive web-based interface (H-VUE or GigaVUE-FM) and a powerful CLI based GigaVUE-OS, the Visibility Fabric is able to replicate, filter, and selectively forward network traffic to monitoring, management, and security tools. The GigaVUE-FM provides a web-based centralized interface to configure and manage all of your GigaVUE nodes through a single pane of glass.

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

- GigaVUE-HD8
- GigaVUE-HD4
- GigaVUE-HC1
- GigaVUE-HC2
- GigaVUE-HC3
- GigaVUE-HB1
- GigaVUE-TA1
- GigaVUE-TA10A
- GigaVUE-TA40
- GigaVUE-TA100
- GigaVUE-TA100-CXP
- GigaVUE-TA200
- Certified Traffic Aggregation White Boxes

Table 12-1 provides overviews of the GigaVUE H_Series and TA_Series nodes.

Table 12-1: GigaVUE H Series and GigaVUE TA Series Nodes

GigaVUE-HC1	• **1RU Footprint** • **Built-in GigaSMART functionality** • **Standard GigaVUE-OS CLI and GigaVUE-FM GUI** • **Cluster with GigaVUE H Series and GigaVUE TA Series Nodes**	
GigaVUE-HC2	• **2RU Footprint** • **Four front-facing bays for port, TAP, BPS, and GigaSMART front modules** • **One rear bay for a GigaSMART rear module** • **Standard GigaVUE-OS CLI and GigaVUE-FM GUI** • **Cluster with GigaVUE H Series and GigaVUE TA Series Nodes**	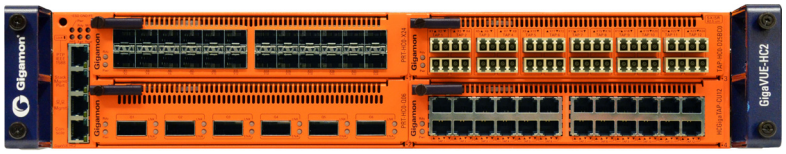
GigaVUE-HC3	• **3RU Footprint** • **Four Module Slots (Bays)** • **Internal Control Card** • **Extension Board** • **Dedicated Cluster Management Port** • **Standard GigaVUE-OS CLI and GigaVUE-FM GUI** • **Supports all GigaVUE-HC3 Modules** • **Cluster with GigaVUE H Series and GigaVUE TA Series Nodes**	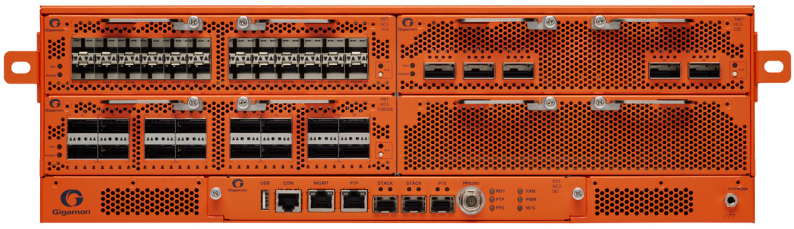

Table 12-1: GigaVUE H Series and GigaVUE TA Series Nodes

GigaVUE-HD4	• **5RU Footprint** • **Four Line Card Slots** • **Single Control Card** • **Dedicated Cluster Management Port** • **Supports all GigaVUE HD Series Line Cards** • **Standard GigaVUE-OS CLI and GigaVUE-FM GUI** • **Cluster with GigaVUE H Series and GigaVUE TA Series Nodes**	
GigaVUE-HD8	• **14RU Footprint** • **Eight Line Card Slots** • **Dual Control Cards** • **Dedicated Cluster Management Port** • **Supports all GigaVUE HD Series Line Cards** • **Standard GigaVUE-OS CLI and GigaVUE-FM GUI** • **Cluster with GigaVUE H Series and GigaVUE TA Series Nodes**	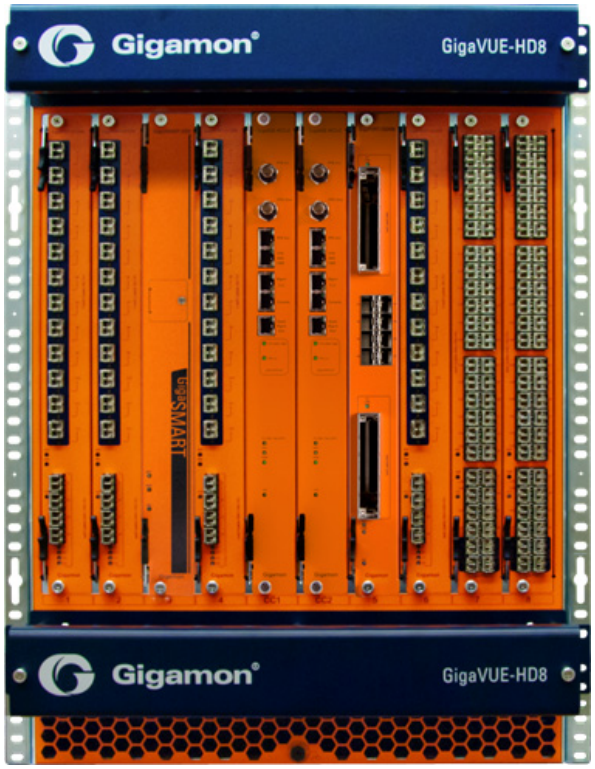

Table 12-1: GigaVUE H Series and GigaVUE TA Series Nodes

GigaVUE-TA1	• 1RU Footprint • Flexible 10Gb/40Gb Modes for 40Gb Ports • Standard GigaVUE-OS CLI and GigaVUE-FM GUI • Cluster with GigaVUE H Series and GigaVUE TA Series Nodes	
GigaVUE-TA10	• 1RU Footprint • Flexible 10Gb/40Gb Modes for 40Gb Ports • Standard GigaVUE-OS CLI and GigaVUE-FM GUI • Cluster with GigaVUE H Series and GigaVUE TA Series Nodes	
GigaVUE-TA40	• 1RU Footprint • Flexible 10Gb/40Gb Modes for 40Gb Ports • Standard GigaVUE-OS CLI and GigaVUE-FM GUI • Cluster with GigaVUE H Series and GigaVUE TA Series Nodes	
GigaVUE-TA100	• 1RU Footprint • 32 x 100Gb/40Gb Ports • Standard GigaVUE-OS CLI and GigaVUE-FM GUI • Cluster with GigaVUE H Series and GigaVUE TA Series Nodes	

Table 12-1: GigaVUE H Series and GigaVUE TA Series Nodes

GigaVUE-TA100 CXP	• **1RU Footprint** • **20 100Gb CXP Ports, 8 100Gb QSFP28 Ports** • **Standard GigaVUE-OS CLI and GigaVUE-FM GUI**	
GigaVUE-TA200	• **2RU Footprint** • **64 100Gb/40Gb ports** • **Optional patch or breakout panel support** • **Cluster with H Series nodes**	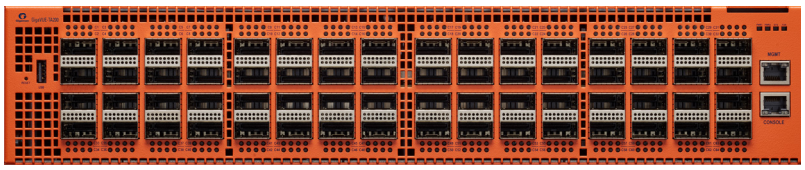
Certified Traffic Aggregation White Box	• **1RU Footprint** • **10Gb/40Gb Ports** • **Standard GigaVUE-OS CLI and GigaVUE-FM GUI** • **Cluster with GigaVUE H Series and GigaVUE TA Series Nodes**	

Notes on TA Series Nodes

- A twenty-four (24) port GigaVUE-TA10 version, called the GigaVUE-TA10A is available with only the first 24 1Gb/10Gb ports enabled. A license is needed to expand a GigaVUE-TA10A to include all 48 1Gb/10Gb ports as well as the four (4) 40Gb ports.
- On the GigaVUE-TA100, only the first 16 out of 32 100Gb ports are enabled. Two port licenses are available to enable an additional 8 or 16 ports to expand from 16 ports to 24 ports or from 16 ports to 24 ports and then to 32 ports.
- The ports on the GigaVUE-TA100 can be used as network, tool, or hybrid ports.
- For more information about the TA Series nodes, refer to the *GigaVUE TA Series Hardware Installation Guide*.

For adding a physical node to GigaVUE-FM, refer to *Configure Physical Nodes* on page 188.

About Cluster

Starting in software version 5.1, you can use GigaVUE-FM to create a cluster. Cluster is created from standalone nodes that are currently managed by GigaVUE-FM. The type of cluster that can be created in this software version is out-of-band.

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

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

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

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

- must be currently managed by GigaVUE-FM
- must be running software version 5.1, at a minimum
- must be running the same software version
- must be reachable by GigaVUE-FM, that is, it must be online and not have any authentication failures
- for GigaVUE TA Series nodes that have an Advanced Features License, they must be licensed

For information on clustering concepts, refer to *Manage GigaVUE Nodes and Clusters* on page 187.

Overview of Seed Node

The seed node is a new concept in GigaVUE-FM and is different from the cluster master role.

When a cluster is created, one of the nodes that you have selected for inclusion in the cluster will be deemed as the seed node. The seed node will be used to start the formation of a cluster and will be determined by the cluster master preference settings of the nodes selected for the cluster. For example, if two nodes are selected, and one is a GigaVUE-HC2 with a preference of 60 and the other is a GigaVUE-HD4 with a preference of 80, the GigaVUE-HD4 will be selected as the seed node.

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

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

For creating and managing clusters using GigaVUE-FM, refer to *Create and Manage Clusters* on page 193.

13 Manage GigaVUE Nodes and Clusters

This chapter describes how to add and manage GigaVUE nodes on a GigaVUE-FM. The sections in this chapter specifically apply to H Series and TA Series nodes.

This chapter covers the following:

- *Configure Physical Nodes* on page 188 describes the process to add, configure and manage GigaVUE nodes through GigaVUE-FM.
- *Create and Manage Clusters* on page 193 describes the process to create a cluster using the wizard, add nodes to a cluster, remove nodes from a cluster, edit cluster parameters, and add stack links.
- *Upgrade Software on a GigaVUE Node or a Cluster from GigaVUE-FM* on page 228 describes the process to upgrade standalone nodes and clusters through GigaVUE-FM.
- *Search for Specific Nodes Using Keywords* on page 235 provides information about each of the standalone nodes and clusters, including a visual indication of each nodes status.
- *Overview Page* on page 239 provides the information on each node connected to the GigaVUE-FM. This section covers the following: *Systems Information* on page 239, *Ports Information* on page 240, and *Traffic* on page 240.
- *Workflows* on page 240 describes how to use the workflow wizards to create four different types of maps. With the wizards, you can create the following types of maps:
  - Out-of-band maps
  - Inline maps
  - Basic out-of-band GigaSMART maps
  - Advanced out-of-band GigaSMART maps
- *Chassis Table View* on page 245 describes the Chassis Table View when managing a node with GigaVUE-FM.
- *Safe and Limited Modes* on page 247 describes Safe Mode and Limited Mode.

Configure Physical Nodes

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

To access physical nodes attached to an instance of GigaVUE-FM, log into GigaVUE-FM. Click **Physical** on the top navigation link. From the Sites drop-down list, select a specific site to view only those physical nodes and clusters associated to the selected site or select All Sites to view all the physical nodes and clusters managed by GigaVUE-FM.

The Physical Nodes page displays the following information:

Field	Description
Cluster Name	The name of the cluster.
Host Name	The host name of the box.
Node IP	The IP address of the physical node.
DNS Name	The DNS name of the physical node.
Role	The role of the node in the cluster. The role of the node can be one of Master, Standalone, Slave, or Standby.
Model	The type of the GigaVUE H Series model. **NOTE:** H Series can cluster with TA Series nodes, but G Series nodes can only cluster (stack) with other G Series nodes.
Box Id	The box Identifier of the node.
SW Version	The version number of GigaVUE-FM.
Licensed	The status of the physical node or Advanced Features license.
Device Status	The current health status of the GigaVUE node or cluster. Whenever there is a change in the health status of the node/cluster, the **Status Updates** notification pops-up in the bottom left corner of the page. Click on the host name link to navigate to the Overview page. To know about how the device health status is computed, refer to *Node Health Status* on page 1361. **NOTE:** You can monitor the health status of the device by enabling the SNMP notifications. For more information on configuring the email notifications, refer to *Notifications* on page 1320.
Task Status	The status of the upgrade process. When the upgrade process is in progress, the task status displays the number of steps completed successfully out of the total number of steps to be completed. For example, upgrade: step (2/5) Image Fetch Complete. Once the upgrade process is complete, the upgrade status is displayed as Upgrade Success or Upgrade Failure.
Connection Status	The status of the device discovery. The status displayed can be one of Authentication Failed, Connection Failed, or OK.
Last Connected Time	The timestamp when the physical node or cluster was last connected.

Field	Description
Tag	The tag or site name and value associated to the physical node or cluster. The tag names associated to the physical node or cluster are displayed as separate columns.Under the tag or site name, the respective tag or site value is displayed.

NOTE: The columns in the Physical Nodes page can be customized based on the type of content you want to view in the table. For customizing the columns, refer to *Table View Customization* on page 45.

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

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

To view the last synchronized status, click the cluster and view the status at the top of the Overview page.

Add New Physical Node or Cluster to GigaVUE-FM

You can add physical nodes and clusters to GigaVUE-FM either manually or by importing an Excel spreadsheet. However, before adding a new physical node, ensure that the node credentials are added under the System Page. Refer to Figure 13-1 on page 189.

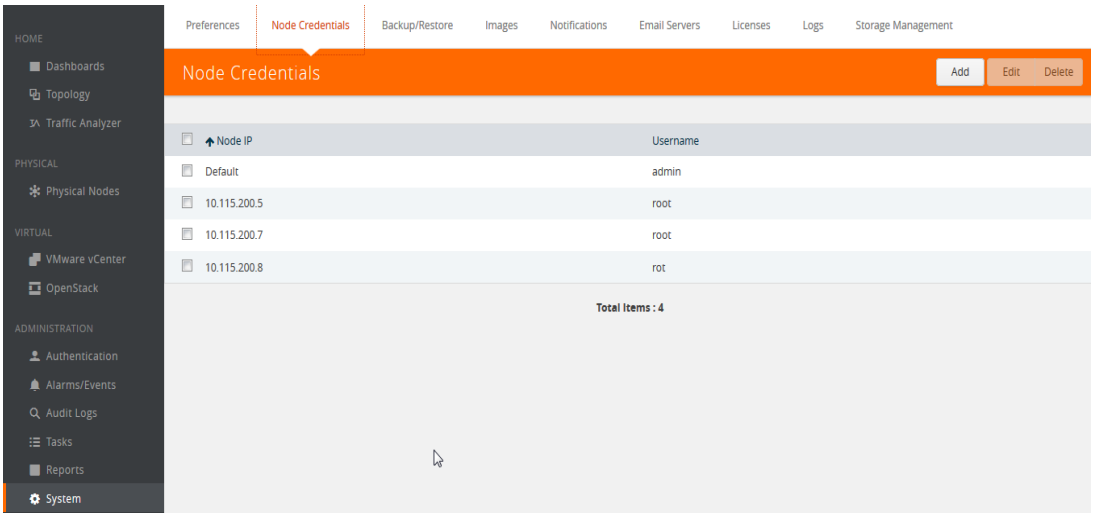

Figure 13-1: Adding Node Credentials

NOTE: In a cluster configuration, the Normal nodes are seen as Slave nodes in GigaVUE-FM.

Add Nodes Manually

To add physical nodes manually, do the following:

1. Click **Physical** on the top navigation link.
2. On the Physical Nodes page, click **Add**. The Add Physical Node page displays.
3. Select **Add Manually**.
4. Enter the DNS name or IP address of the node. Click **+** to add additional nodes or **-** to remove a node.
5. Click **Submit** to add the node or nodes to the list of physical nodes GigaVUE-FM is managing. For standalone nodes, both node IP (Device IP) and cluster ID are the node Addresses (nodeAddress) provided by the user while adding the node.

NOTE: In GigaVUE-FM version 5.5.01, the Node IP was the actual IP and the Cluster ID was either the DNS ID (if the IP address was resolved) or the Node IP (if the IP address was not resolved).

Add Nodes From an Excel Spreadsheet

You can add nodes to GigaVUE-FM by uploading an Excel spreadsheet that contains a list of the physical nodes that you want GigaVUE-FM to manage. You can create the spreadsheet or use a spreadsheet from a previous export of the nodes. The format of the spreadsheet must have at least a column with the node IPs and the header Node_IP.

NOTE: The format of the spreadsheet changed in GigaVUE-FM 3.3. To import a spreadsheet created by GigaVUE-FM prior to GigaVUE-FM 3.3, you can modify the previous spreadsheet or do a new export after upgrading to the current version of GigaVUE-FM. An example of a spreadsheet is shown in Figure 13-2.

	A	B	C	D	E	F	G	H	I	J
1	Cluster	Name	Node_IP	Role	Model	Box_Id	SW_Versio	Licensed	Last_Connected_Time	
2	10.110.150.54	HB1-C03-21	10.110.15-.54	Standalone	HB1	3	4.6.00	Yes	2016-04-01 14:21:36	
3	10.110.150.98	TME-2404-1	10.110.150.98	Standalone	GV2404	1	8.6.10	Yes	2016-04-01 14:21:35	
4	10.110.150.58	HC2-C04-31	10.110.150.58	Standalone	HC2	3	4.5.00	Yes	2016-04-01 14:21:37	
5	10.110.150.59	HC2-CO4-33	10.110.150.59	Standalone	HC2	1	4.5.00	Yes	2016-04-01 14:21:35	
6	10.110.150.57	HD4-C03-08	10.110.150.57	Standalone	HD4	2	4.5.00	Yes	2016-04-01 14:21:36	
7	cluster-demo	HB1-C03-22	10.110.150.55	Slave	HB1	10	4.6.00	Yes	2016-04-01 14:21:39	
8	cluster-demo	HB1-C04-15	10.110.150.62	Standby	HB1	15	4.6.00	Yes	2016-04-01 14:21:39	
9	cluster-demo	HB1-C03-23	10.110.150.56	Slave	HB1	19	4.6.00	Yes	2016-04-01 14:21:39	
10	cluster-demo	HD8-C04-01	10.110.150.60	Master	HD8	1	4.6.00	Yes	2016-04-01 14:21:37	
11	10.110.150.53	HC2-C03-13	10.110.150.53	Standalone	HC2	1	4.6.00	Yes	2016-04-01 14:21:36	
12	10.110.150.50	HC2-C04-29	10.110.150.50	Standalone	HC2	1	4.5.00	Yes	2016-04-01 14:21:35	
13	10.110.150.97	TME-2404-2	10.110.150.97	Standalone	GV2404	1	8.6.10	Yes	2016-04-01 14:21:35	
14	10.110.150.52	TA1-C03-24	10.110.150.52	Standalone	TA1	1	4.5.00	Yes	2016-04-01 14:21:36	
15	10.110.152.51	TA1-C04-35	10.110.150.51	Standalone	TA1	18	4.5.00	Yes	2016-04-01 14:21:35	
16	10.110.150.99	TME-GV-212	10.110.150.99	Standalone	GV212	4	8.6.10	Yes	2016-04-01 14:21:35	

Figure 13-2: Node List Spreadsheet for Import/Export

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

1. Click **Physical** on the top navigation link.
2. On the Physical Nodes page, click **Import**.

The Add Physical Node page displays with the **Import from Excel** option automatically selected. The page displays an area for selecting or dragging and dropping a file.

3. Drop an Excel file onto the page or click **Select File** to upload the file.

   The page displays the list of IP address.

4. Select the nodes to add to GigaVUE-FM.
5. Click Submit to add the node or nodes to the list of Physical Nodes GigaVUE-FM is managing.

Cluster Discovery Behavior

GigaVUE-FM does not detect individual nodes that were part of a cluster, if the cluster was dismantled using the CLI. GigaVUE-FM always reaches the cluster by its virtual IP address (if configured). When a cluster is dismantled through the CLI, the virtual IP address of the cluster will no longer be available and the cluster is therefore marked as unreachable. There is no change in detecting node additions, removals, or membership changes performed from CLI.

ARP/NDP Timer Settings

The Address Resolution Protocol (ARP) or the Neighbor Discovery Protocol (NDP) timer specifies the aging time on the IP interface. The ARP timer is used for IPv4 addresses and the NDP timer is used for IPv6 addresses. The timer is configurable from 3 to 30 seconds. The default is 30 seconds. When an IP interface is configured, ARP/NDP requests are sent to the IP interface to find the gateway MAC address. In response, the gateway sends an ARP/NDP reply. The control card tries to match the IP address of the IP interface with the IP address of the ARP/NDP message received. If a match is found, the ARP/NDP status changes to resolved (otherwise the ARP/NDP status is not resolved).

Once the ARP/NDP status is resolved, the ARP/NDP timer of the IP interface controls the interval at which an ARP/NDP request is sent to the gateway to detect if the gateway is reachable or not.

You must enable the IP Gateway Status SNMP trap to send SNMP notifications when the ARP/NDP status changes. To enable SNMP notifications, refer to the following:
Enable or Disable Events for SNMP Notifications on page 192

Change the ARP/NDP Timer Settings

The default ARP/NDP timer value is 30 seconds. To change the timer setting at the node-level:

1. On the top navigation pane, click **Physical**.
2. In the Physical Nodes page, select the node for which you want to change the ARP/NDP timer setting.
3. Go to **Settings > Global Settings > ARP/NDP**.

4. Click **Settings**.
5. In the ARP/NDP Settings page, choose the required **ARP Refresh Time Interval** or the **NDP Refresh Time Interval** in seconds.
6. Click **OK**.

The ARP Entries table and the IPv6 Neighbor Entries table dynamically refresh to display information such as the IP address and Hardware address mapping, aging, state, and interface details.

Click **Clear > Clear ARP Entries** or **Clear > Clear IPv6 Neighbor Entries** to remove the entries from the tables.

Enable or Disable Events for SNMP Notifications

To enable or disable events for SNMP Notifications:

1. From the device view, go to **System > Settings > Global Settings > SNMP Traps**.
2. Click **Trap Settings**. The Edit SNMP Trap Settings page opens.
3. Select or clear the **Enable** check box for the SNMP trap that you want to enable or disable.
4. Click **OK**.

SNMPv3 Support

Starting with software version 5.4 GigaVUE-FM creates an SNMPv3 user during the upgrade process from version 5.3. The SNMPv3 user is created with same authentication and privacy settings previous held.

NOTE: SNMPv3 support is only available on FM software version 5.4 user and SNMPv3 users created by FM cannot be modified by users.

Enable SNMPv3 on Nodes

You can enable SNMPv3 from the node addition page, If node version 5.4 or higher and SNMPv3 is selected, the FM registers itself as a SNMPv3 trap receiver. If the SNMPv3 is chosen and the node version is 5.3 or below then, FM registers itself as a SNMPv2 trap receiver. When this occurs FM reports an event that FM does not support SNMPv3 on the node with 5.3 or lesser version.

Enabling SNMPv3 During Upgrade

1. Select **Administration> System> Credentials.** The Credentials page displays a listing of the nodes and SNMP versions.
2. Select a node to upgrade, and click **Edit**.
3. Select SNMP version - SMNPv3
4. Click **Save**.

Enabling SNMPv3 on a New Node

1. Select **Physical> Add.** The Add Physical Node displays.

a. Enter **Node Name**

b. Click **Enable**

c. Enter **Username** and **Password**

d. **SNMP Version: SNMPv3**

Click **Submit.**

Create and Manage Clusters

This section describes the GigaVUE-FM clustering. Refer to the following sections for details:

About Cluster

A cluster consists of multiple GigaVUE-OS nodes operating as a unified fabric such that packets entering the cluster on one node can be sent to a destination port on any other node. You set up packet distribution using the standard box ID/slot ID/port ID format, allowing maps to distribute traffic to any port in the cluster.

Figure 13-3 illustrates a sample cluster of GigaVUE-OS nodes.

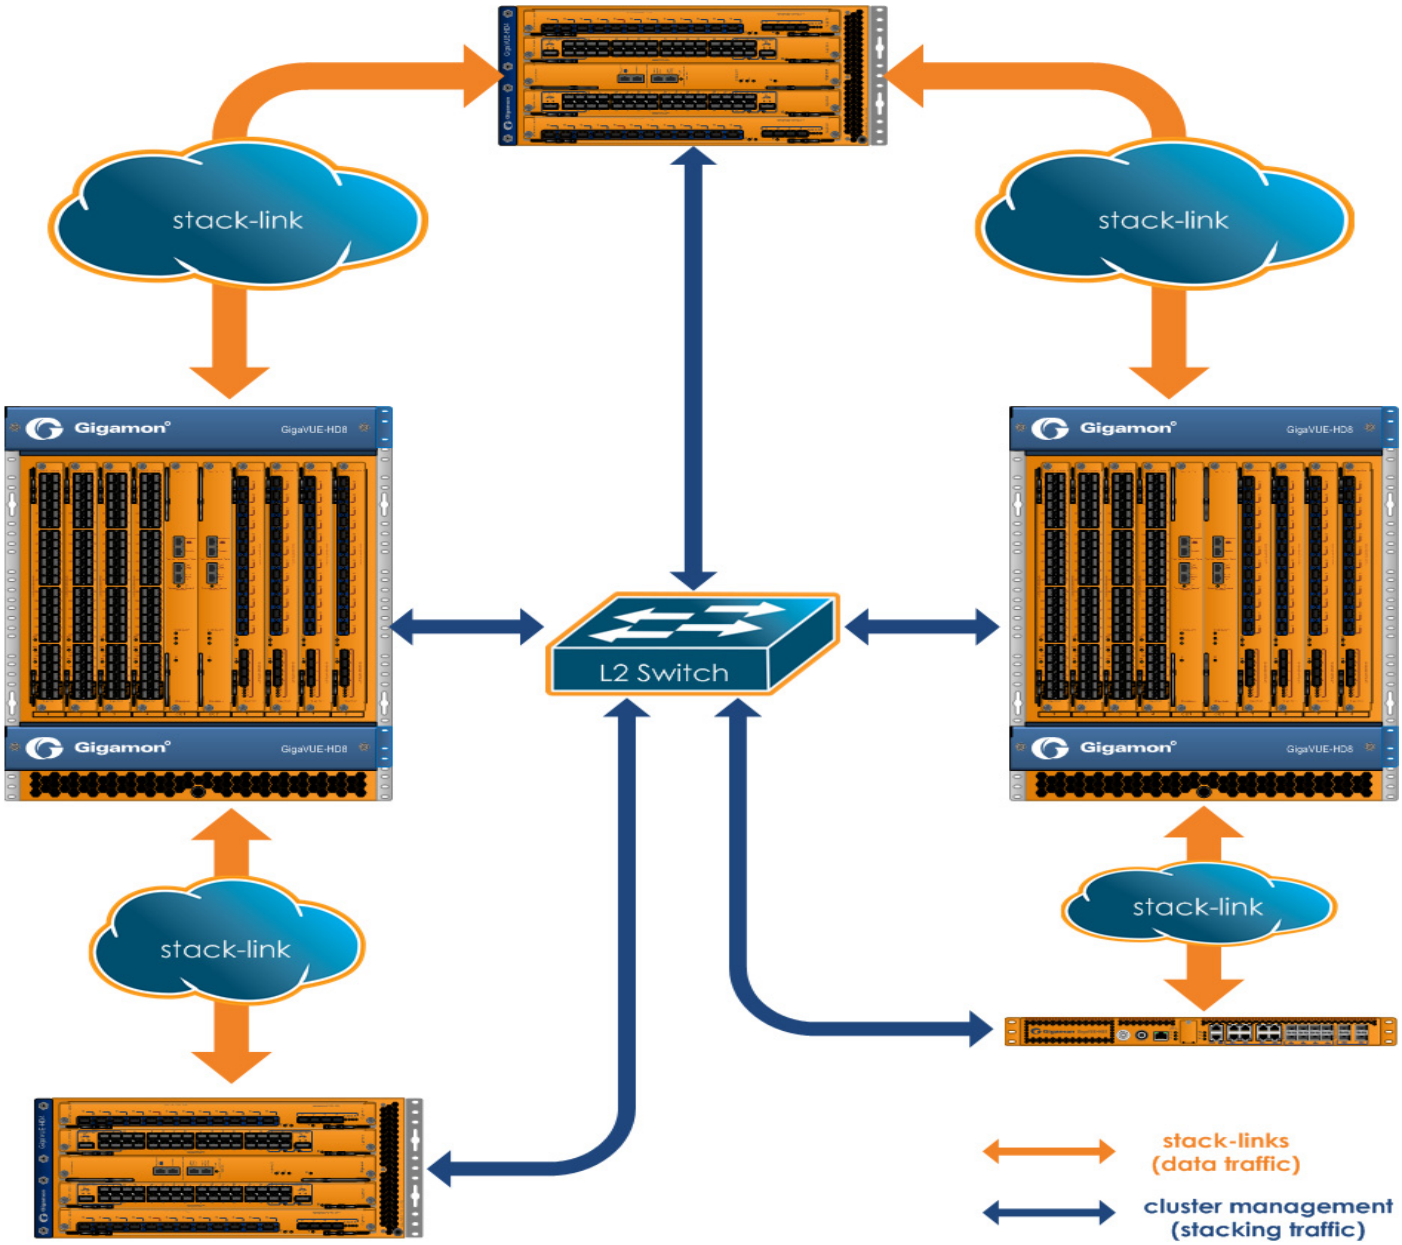

Figure 13-3: Sample Cluster

Cluster Node Limit

Any GigaVUE-OS nodes can join other nodes in a cluster. The GigaVUE-HD8, GigaVUE-HD4, GigaVUE-HC3, GigaVUE-HC2, GigaVUE-HC1, GigaVUE-HB1, and GigaVUE TA Series nodes, including Certified Traffic Aggregation White Box (white box), can be included in the same cluster. Starting in software version 4.5, the maximum number of nodes supported in a cluster is 32. For details, refer to *Cluster Scaling* on page 196.

In addition, there is another independent limit, which is for the maximum number of line cards supported in a cluster, across all of the nodes in the cluster. This limit is

determined by cost units. Cost units are based on the total number of line cards, line card types, and chassis type. Cost units measure the resources that a node needs in a cluster. The higher the cost unit, the more resources are needed to manage the node.

The following table has examples of line card and chassis types, and their cost units.

Line Card Type or Chassis	Cost Unit
PRT-HD0-C02X08 or C02X08A on GigaVUE HD Series	3
SMT-HD0 GigaSMART on GigaVUE HD Series	2
PRT-H00-X04G44 on GigaVUE HD Series	2
All other line cards on GigaVUE HD Series	1
GigaVUE-HB1 Chassis	2
GigaVUE-HC1 Chassis	1
GigaVUE-HC2 Chassis	1
GigaVUE-HC2 Chassis with Control Card version 2 (HC2 CCv2)	1
GigaVUE-HC3 Chassis	1
GigaVUE-TA1	1
GigaVUE-TA10	1
GigaVUE-TA40	1
GigaVUE-TA100	1
GigaVUE-TA200	1
Certified Traffic Aggregation White Box	1

For example, if a GigaVUE-HD8 has one C02X08 line card, one GigaSMART line card, and one Q02X32 line card (other), the cost units are 3 + 2 +1 = 6.

NOTE: The cost unit for a GigaVUE-HC2 is always 1, regardless of how many modules are installed.

The following table displays the cost unit maximum depending on the nodes in the cluster.

Cluster Nodes	Cost Unit Maximum
If any node in the cluster has one or more PRT-H00-X04G44 line cards	127
If there is a GigaVUE-HB1 in the cluster	127
All others	255

Therefore, the largest cluster supported is determined either by the maximum number of nodes (32) or by the cost unit maximum, whichever is reached first.

Cluster Scaling

The maximum number of nodes and map rules supported in a cluster is as follows:

Table 13-1: Maximum Number of Nodes and Map Rules Supported in a Cluster

When a Cluster is Configured with:	Number of Nodes	Maximum Map Rules
Out-of-Band Cluster Management	32	38000
Inband Cluster Management	16	38000

The maximum number of map rules supported in a cluster apply to all nodes in the cluster including GigaVUE H Series nodes: GigaVUE-HD8, GigaVUE-HD4, GigaVUE-HC3, GigaVUE-HC2, GigaVUE-HC1, and GigaVUE-HB1, and GigaVUE TA Series nodes: GigaVUE-TA1, GigaVUE-TA40, GigaVUE-TA100, and GigaVUE-TA200, including Certified Traffic Aggregation White Box (white box).

Cluster Topologies

The following cluster topologies are supported:

- star
- daisy-chain
- tree

Separate Paths for Cluster Control and Stack Traffic

There are two types of clustering: out-of-band and inband.

The nodes in a cluster are constantly communicating with one another, exchanging heartbeats to check on one another's status, exchanging configuration information so that changes made on the master node are propagated to other nodes, and making changes to cluster roles based on changes in status.

GigaVUE-OS separates cluster control traffic from the actual flow of packets from ingress ports on one node to egress ports on another for the two types of clustering, as follows:

Out-of-Band Clustering on Mgmt Ports

With out-of-band clustering, cluster control traffic is carried out-of-band on its own network as follows:

- GigaVUE-HD8/GigaVUE-HD4 nodes (with HCCv2) use the Mgmt port (eth0) or the dedicated cluster Mgmt port (eth2).
- GigaVUE-HC3 nodes use the Mgmt port (eth0) or the two dedicated STACK ports (eth2).
- GigaVUE-HC2 nodes use the Mgmt port (eth0) or the dedicated cluster Mgmt port (eth2).
- GigaVUE-HC1 nodes use the Mgmt port (eth0) or the dedicated cluster Mgmt port (eth2).
- GigaVUE-HB1 nodes use the Mgmt port (eth0).
- GigaVUE-TA1 nodes use the Mgmt port (eth0).
- GigaVUE-TA10 nodes use the Mgmt port (eth0).
- GigaVUE-TA40 nodes use the Mgmt port (eth0).
- GigaVUE-TA100 nodes use the Mgmt port (eth0).
- GigaVUE-TA200 nodes use the Mgmt port (eth0).
- GigaVUE-OS on white box nodes use the Mgmt port (eth0).

Cluster Management port(s). Using the cluster management port(s) lets you route cluster control traffic over a separate network from the network used to access the Mgmt port. This prevents cluster control traffic from overloading the traffic used to access the Mgmt port.

Mgmt Port (eth0). You can also elect to use the standard Mgmt port for cluster control traffic. In this implementation, cluster control traffic uses the Mgmt port's Ethernet connection.

Inband Clustering on Stack-Links

Stack-links are used to create a stacking connection between two GigaVUE nodes in a cluster. Stack-links carry traffic from network ports on one node to tool (or GigaSMART) ports on a destination node.

With inband clustering, cluster control traffic is carried inband through the stack-link.

Stack-links can be constructed out of individual stack ports, for example, a 40Gb port on a PRT-H00-Q02X32 line card for GigaVUE-HD4/8, PRT-HC0-Q06 or PRT-HC0-X24 for GigaVUE-HC2, or stack GigaStream. You decide which to use with the **gigastream** and **ports** arguments in the **stack-link** command. For example, the following command creates a stack-link between the q1 40Gb port on box 1/slot 1 and the q1 port on box 2/slot 7:

(config) #stack-link alias biglink between ports 1/1/q1 and 2/7/q1 comment "40Gb Stack"

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

Keep in mind that because of the 10Gb port density offered by GigaVUE-OS nodes, using only one 10Gb port for a stack-link could cause a serious bottleneck. A stack GigaStream dramatically increases the bandwidth available for stack-links, letting you connect GigaVUE-OS nodes in a cluster and still take advantage of the 10Gb port density. Alternatively, nodes with 40Gb or 100Gb ports can take advantage of their high bandwidth for stack-links.

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

About Cluster Roles

Communication with a GigaVUE-OS cluster is accomplished using a master virtual IP address assigned to the cluster as a whole. Physically, the virtual IP address resolves to only a single master node at any one time. However, the master role on the GigaVUE-OS node is not statically assigned to a single node. Instead, any node (except GigaVUE TA Series and the nodes residing on a different management subnet) in the cluster can take on the master role if the situation requires it (for example, if both the master and the current standby nodes go down).

When a new node becomes the master, it takes ownership of the virtual IP address used for master access to the cluster. Because all of the nodes in the cluster share the same database, this transition takes place seamlessly, ensuring that the cluster survives the failure of one or more nodes.

The virtual IP address is assigned to the primary control card in the configuration jump-start wizard:

Step xx: Cluster mgmt virtual IP address and masklen? [0.0.0.0/0]

Each node in the cluster is performing one of the following roles at any given time:

- **Master** – This node has possession of the cluster's virtual IP (VIP) address and takes responsibility for dispatching commands to the entire cluster.
- **Standby** – This node takes over the master role in the event of a failure on the node currently holding the role.
- **Normal** – These nodes perform normal GigaVUE operations with minimal cluster responsibilities. However, they, too, have a complete copy of the cluster's database. When a master fails and standby is promoted to be the new master, an election process takes place automatically between all normal nodes, ensuring that a new standby is found.

Setting a Node's Priority in the Master Election Process

Clusters of GigaVUE-OS nodes perform a master election in the following situations:

- Cluster reload
- Master or standby node failure

In either of these cases, a new node is selected to perform the necessary role(s). You can set the **cluster master preference** for each individual node in the cluster to specify how likely the node is to claim a master/standby role. Higher values are more likely to claim the master/standby role; lower values are less likely.

Use preference settings from 10 to 100 for master, standby, and normal roles. Use preference settings from 1 to 9 for normal nodes that are excluded from taking the master or standby role.

In software version 4.5, the preference cannot be set to zero. A node with a preference of 0 in an earlier software version will be changed to 1 after upgrading to 4.5 or higher.

GigaVUE-OS sets defaults for the **preference** argument based on the type of control card in use. If you choose to change a node's **preference** setting, it is generally preferable to set higher priorities for nodes with more processing power. GigaVUE HD Series, GigaVUE-HC3, or GigaVUE-HC2 nodes provide the most processing power, followed by GigaVUE-HC1 nodes, followed by GigaVUE-HB1 or GigaVUE TA Series nodes.

NOTE: All GigaVUE TA Series nodes including the white box, will automatically be added to a cluster with preference set to 1 because any Traffic Aggregator can never take the role of, or be eligible to be, the master node.

In addition, in an event of a cluster reboot, any GigaVUE TA Series node in a cluster may show as standby for a couple of minutes while the cluster is coming up from the reboot cycle. However once the cluster is up and running, none of the GigaVUE TA Series nodes can be a standby.

About the "Unknown" Cluster Role

In addition to the standard roles in *About Cluster Roles* on page 199, the system may occasionally report a node operating with an **unknown** cluster role. A node with an unknown cluster role is no longer being actively managed by the master node.

When a node that was formerly part of a cluster transitions to an **unknown** cluster role, its database will typically be out of synchronization with the master node's. You can restore the node to the cluster by using the **reset factory keep-all-config** command, followed by a reboot, and running **configuration jump-start** to rejoin the cluster with a clean local database.

Sample Cluster Control Connections

The GigaVUE-OS provides a flexible approach to cluster control traffic, allowing you to route it over cluster management or Mgmt ports. The ports available and their eth***x*** designations vary by control card version and node type, as summarized in the following table:

Control Card/Node Type	Possible Cluster Control Ports	Deployment Models
HCCv2 Control Card	**Cluster Mgmt (eth2) or Mgmt (eth0)**	• Cluster Mgmt (eth2) and L2 switch • Mgmt (eth0) and L2 switch
GigaVUE-HC3 Node	**STACK (eth2) or Mgmt (eth0)**	• Cluster Mgmt (eth2) and L2 switch • Mgmt (eth0) and L2 switch
GigaVUE-HC2 Node	**Cluster Mgmt (eth2) or Mgmt (eth0)**	• Cluster Mgmt (eth2) and L2 switch • Mgmt (eth0) and L2 switch
GigaVUE-HB1 Node	**Mgmt (eth0)**	• Mgmt (eth0) and L2 switch **NOTE:** When including a GigaVUE-HB1 in a cluster, it is required to use one port for both node and cluster management.
GigaVUE TA Series Nodes	**Mgmt (eth0)**	• Mgmt (eth0) and L2 switch

Sample Cluster Control Configurations

Figure 13-4 shows the cluster control traffic being sent over the Mgmt (eth0) ports to an L2 switch.

Cluster control traffic can also be sent over the cluster Mgmt (eth2) ports to an L2 switch (only for the GigaVUE-HD8, GigaVUE-HD4, GigaVUE-HC3, GigaVUE-HC2, and GigaVUE-HC1 nodes).

Nodes in the same cluster must use the same cluster interface. For example, if there is a GigaVUE-HB1 or a GigaVUE-TA10 in the cluster, all nodes in the cluster must use eth0.

Zeroconf for Cluster Management Ports

By default, cluster management ports use zero configuration networking (zeroconf) to establish networking settings. This eases configuration when establishing clusters using the cluster management port(s).

Keep Cluster Management Ports Connected!

IMPORTANT: Clusters implemented using the cluster management ports for cluster control traffic must ensure that the cluster management ports of all nodes in the cluster are connected at all times. This prevents a situation where multiple masters claim the management VIP address, resulting in the inability to connect to it at all.

Cluster Control Traffic on Mgmt Port (eth0)

Use the Mgmt port (eth0) for cluster control traffic. This is the only supported configuration for a GigaVUE-HB1 node or GigaVUE TA Series nodes– they do not have a dedicated cluster Mgmt port.

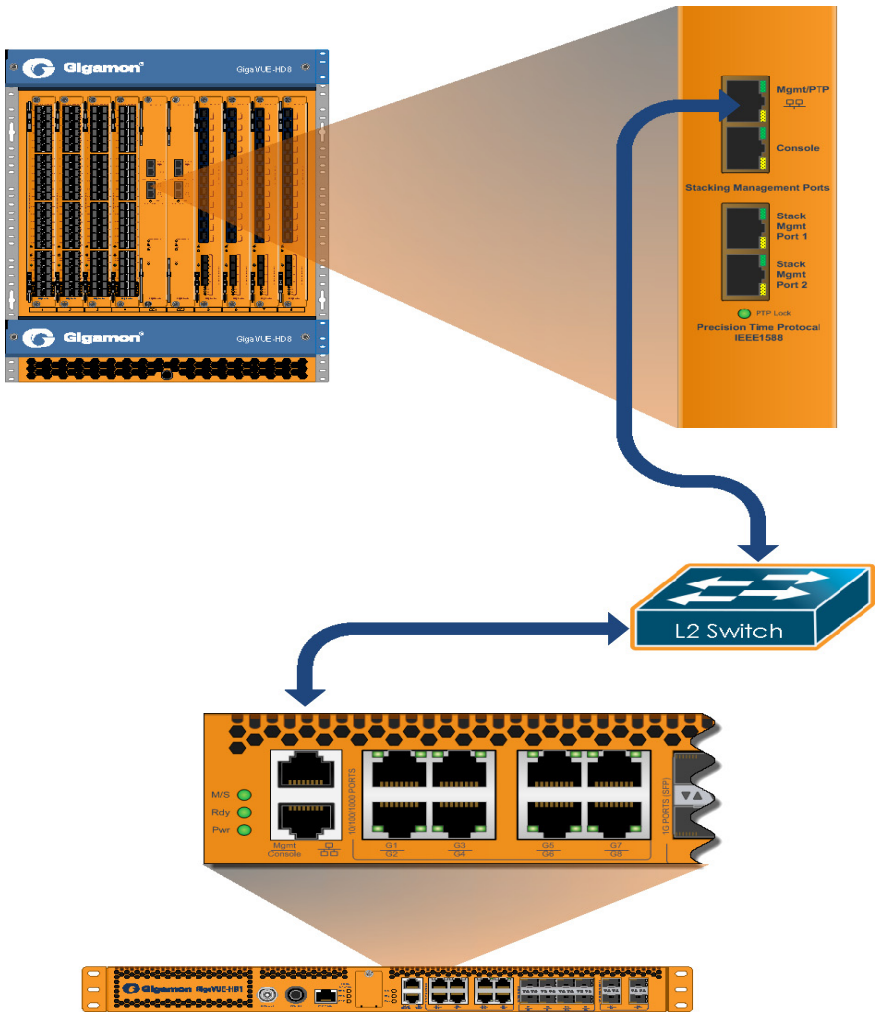

Figure 13-4: Cluster Control Traffic on Mgmt Port (eth0)

Sample Stack-Link Configurations

This section illustrates some sample configurations for the data-carrying stack-links in a cluster, including a star configuration and a daisy-chain using a GigaStream. You can see a combination of star and daisy-chain in *About Cluster* on page 193.

IMPORTANT: Ensure that you do not cable the stack-links in a loop. Use a star or daisy-chain configuration, as follows:

Star Configuration

Use a GigaVUE-OS node as the hub in a star configuration. This makes it easy to create a star configuration that maximizes traffic distribution efficiency. With a star configuration, no destination is further than two hops away. Note that the following image only shows the stack-link connections and not the cluster control connections from the control cards.

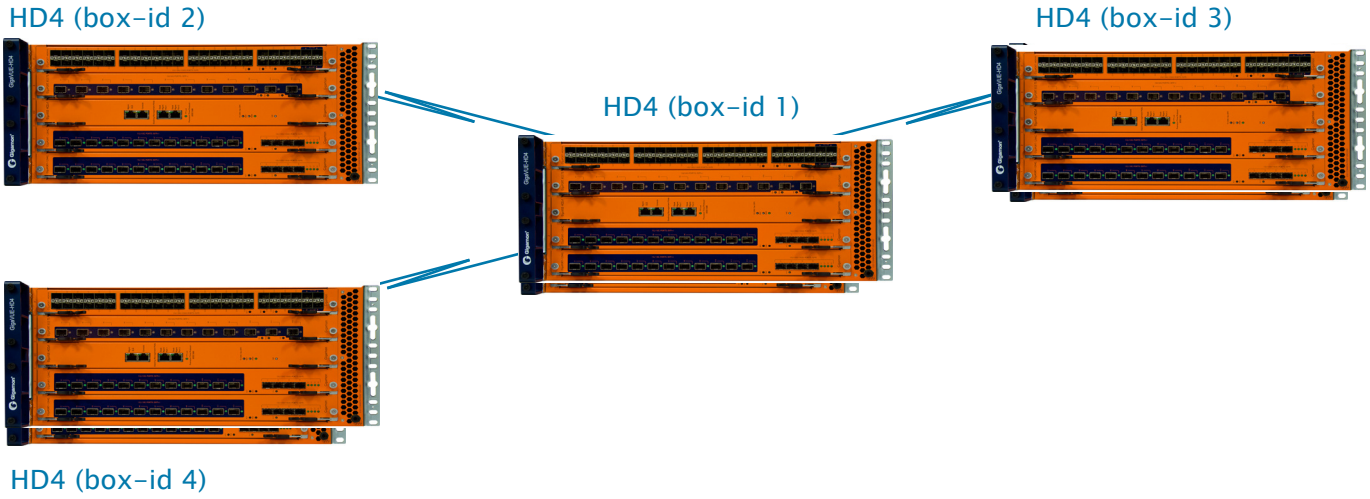

Daisy-Chain Configuration Using GigaStream for Stack-Link

You can connect two GigaVUE-OS nodes together in a daisy-chain, for example, using any 10Gb line card port. Because there can be up to 256 10Gb ports on a single GigaVUE-HD8 node, the stack-link needs enough bandwidth to handle expected cross-node traffic volume. Create a stack GigaStream out of up to 24x10Gb (PRT-HC0-X24) 16x10Gb (PRT-H00-Q02X32) or 8x40Gb (PRT-HD0-Q08) ports to handle expected cross-node traffic loads.

Creating Clusters: A Roadmap

Setting up a cluster consists of the major steps shown in Figure 13-5.

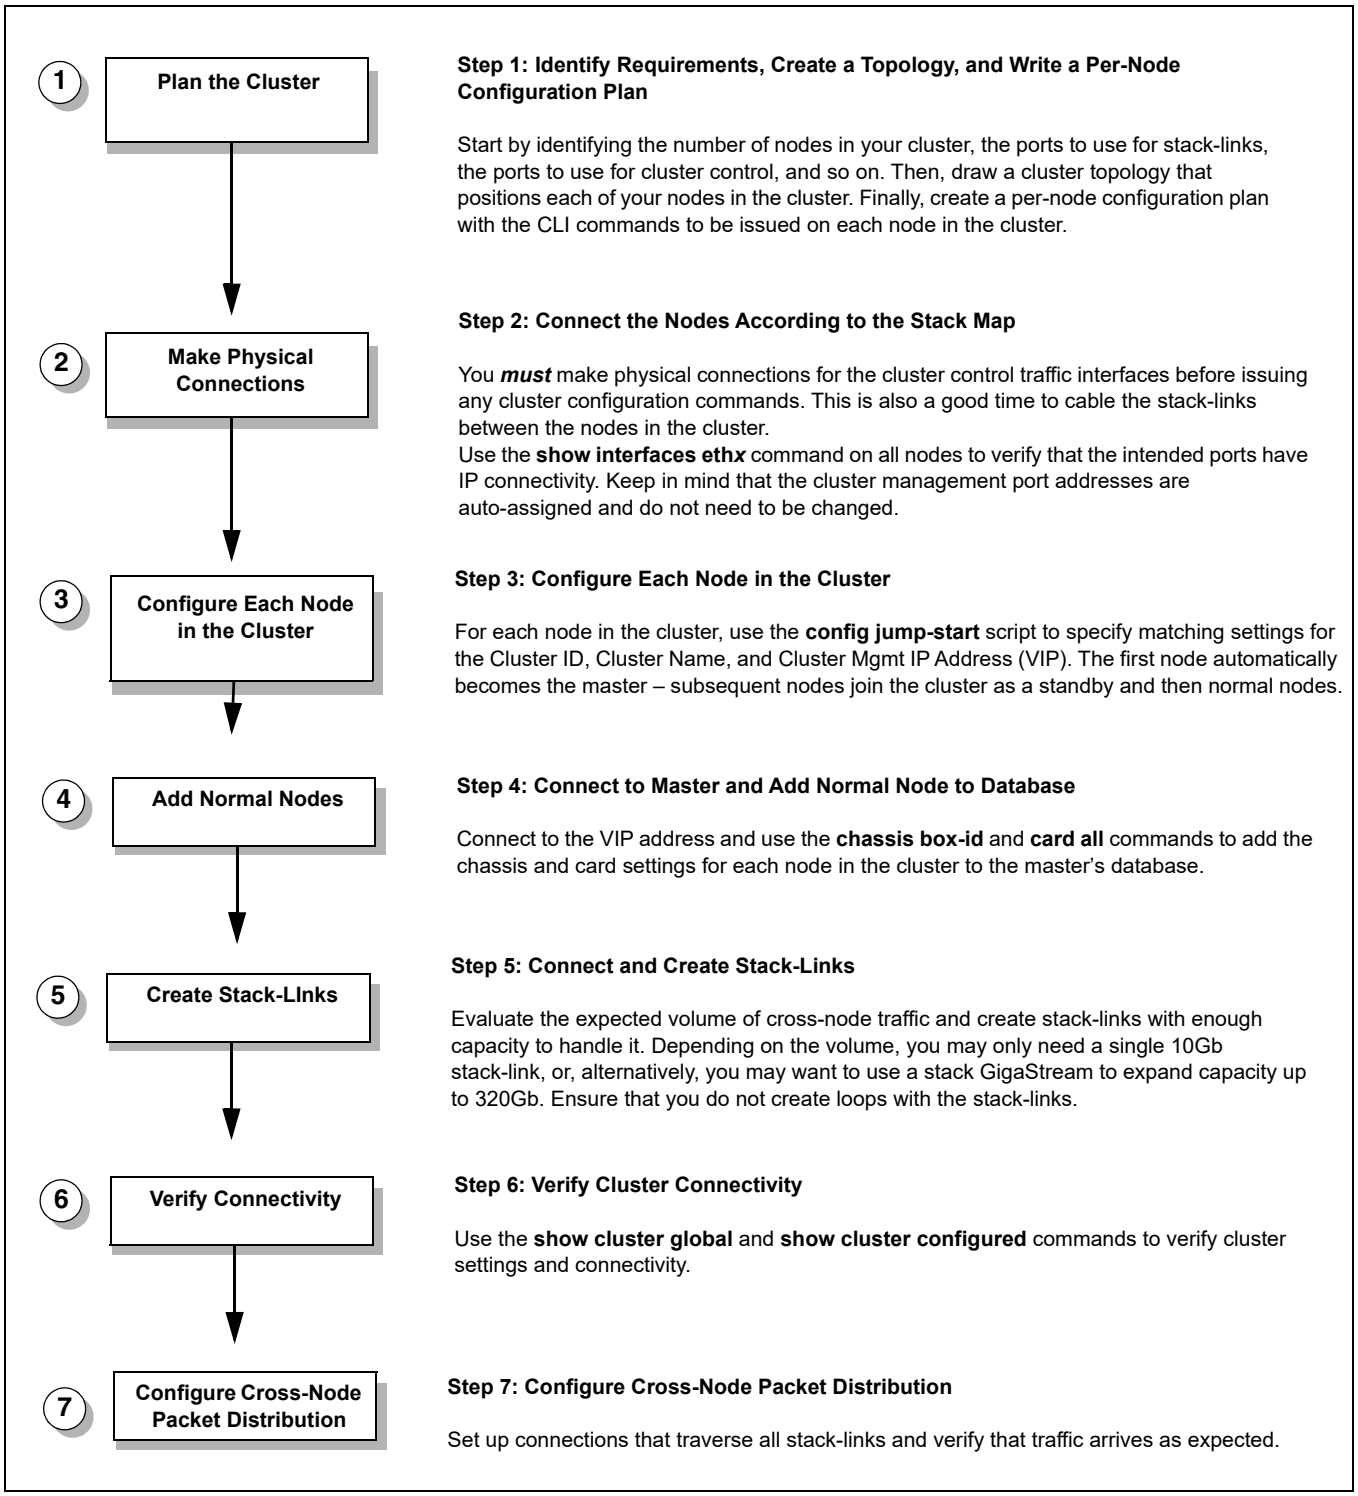

Figure 13-5: Creating Clusters: Major Steps

Cluster Rules and Recommendations

When creating clusters, keep in mind the following rules and recommendations:

- Before joining an existing GigaVUE node to a cluster, it is recommended to use the **no traffic all** or **reset factory** command to clean up existing traffic-related configuration.
- Remove all physical loops before enabling the cluster. An accurate cluster topology will help with this. The GigaVUE-OS node automatically detects and prevents configurations that would cause loops, but its best to avoid them in the first place.
- For GigaVUE-HD8 nodes, leave the cluster management ports on the control card in the CC2 slot (slot 6) disconnected.
- Star configurations offer the most efficient use of bandwidth. In general, use one GigaVUE-OS node at the hub of your star and then connect spokes off of that.
- Create stack-links with enough capacity to match expected cross-node traffic. For example, you can use a 40Gb port on a PRT-HD0-Q08 line card or PRT-H00-Q02X32 line card or create a GigaStream out of stack ports and use that as a stack-link. GigaStream can use up to 16x10Gb (PRT-H00-Q02X32) or 8x40Gb (PRT-HD0-Q08) ports for the GigaVUE HD Series node and 24x10Gb (PRT-HC0-X24) or 6x40Gb (PRT-HC0-Q06) for the GigaVUE-HC2 node.
- Configure only the stack type ports that you will use in the stack-link configuration. Loops can be created if stack type ports are configured but then not used in a stack-link.
- The first node added to the cluster becomes the master. This is important when creating a new cluster using an existing, already-configured node and a new node. If you want to keep the configuration on your existing node, use it as the first node in the cluster. This way, the existing node becomes the master and the new node inherits its configuration, preserving your existing settings. GigaVUE TA Series nodes are an exception since they cannot be the master.
- When joining a new node to an existing cluster, give the new node a lower precedence than the master. Once the database has synchronized to the existing master, you can increase the precedence to make the newly joined node the master, if that is required.
- Connecting two masters to the same cluster is not supported. This is why you must make physical connections for the cluster control traffic interfaces before issuing any cluster configuration commands. Because the first node added to a cluster becomes the master, configuring cluster settings before physically connecting the cluster control network results in a situation with multiple masters attempting to connect to the same cluster.

  NOTE: Merging clusters is not supported.
- GigaVUE TA Series and Certified Traffic Aggregation White Box nodes in a cluster can have tool, network, hybrid, and stack ports.
- A GigaVUE TA Series node cannot be a master node. It can only join a cluster with other node types, such as GigaVUE-HD4/8, GigaVUE-HC3, GigaVUE-HC2, GigaVUE-HC1, or GigaVUE-HB1.

- A GigaVUE TA Series node cannot be a standby node either. If the cluster has one master and all other nodes are GigaVUE TA Series nodes, the cluster will not have a standby.
- Since a GigaVUE TA Series node can never be a master or a standby in a cluster, a database restore is not possible. The best option is a text restore that has the information of the other nodes in the cluster removed from the text backup of the GigaVUE TA Series.

GigaVUE-TA Series and GigaVUE-HC3 Clustering Recommendations

The following recommendations are for GigaVUE-TA Series and GigaVUE-HC3 nodes in a cluster:

- When a GigaVUE-TA Series or GigaVUE-HC3 is connected to a node of a different type, ports may not become operationally up until the stack-links are created between the stack ports. To ensure the ports become operationally up:
  - Configure the specified ports as stack ports.
  - Configure the stack-link between the stack ports.

Cluster Rules

Clusters must adhere to the following rules:

Rule
All GigaVUE-OS nodes in a cluster must run the same version of the GigaVUE-OS software, including the major and minor version numbers.
Each GigaVUE-OS node in a given cluster must share the same **Cluster ID**, **Cluster Name**, and **Cluster Mgmt IP Address.** You can configure these settings in the **config jump-start** script, or, alternatively, use separate **cluster** commands to set them. When adding a node present on the same IP subnet to an existing cluster, so long as you specify the cluster ID correctly, the cluster Mgmt IP address (VIP) will be synchronized from the master automatically.
Cluster management ports must be on the same IP subnet.
Each GigaVUE-OS node in a cluster must have its own unique box ID. The box ID is assigned to a chassis from the master with the **chassis box-id <box ID> serial-num <serial number>** command (). Keep in mind that if you are using GigaSMART trailers to identify ingress ports, only box IDs from 1-64, inclusive, are supported.
You can only connect optical-to-optical stack-links. Stack-links must be at least 10Gb. In addition, they must use the same transceiver types, such as LR-to-LR, or SR-to-SR.
Use a stack-link between different types of GigaVUE-OS nodes so long as the medium, speed, and number of ports involved is the same on both sides.

Best Practices for OOB Clusters with IGMP Snooping

The following are best practices for out-of-band (OOB) clusters if Internet Group Management Protocol (IGMP) snooping is enabled in the cluster.

Clustering relies on the IGMP protocol to discover peer nodes and to communicate with them. Switches often have IGMP snooping enabled by default, which will filter IGMP packets from ports that do not have periodic IGMP membership reports. This can cause IGMP packet drops in out-of-band clusters.

Refer to *About IGMP Snooping in a Cluster* on page 208 for more information. Also refer to the following best practices:

- allow Internet Group Management Protocol (IGMP) traffic by using an IP filter chain. Refer to *Allow IGMP Traffic* on page 209.
- enable an IGMP querier. Refer to *Enable an IGMP Querier* on page 209.

These best practices result in the following:

- hostnames being properly displayed in CLI commands that display cluster information such as **show cluster global brief**
- nodes joining clusters faster, especially nodes that are not capable of becoming a master, such as GigaVUE TA Series nodes
- no multiple masters being created in an out-of-band cluster. This can occur when a node that is capable of becoming a master is not able to see the current master and hence elects itself as a master.

About IGMP Snooping in a Cluster

IGMP snooping is a networking feature that monitors IGMP membership reports received from different ports on a networking switch and learns the ports to which multicast groups belong. When a port stops sending membership reports about a multicast group, the switch will stop forwarding the group's traffic to the port.

An IGMP querier is a router (or switch) feature that periodically queries the network for multicast group interests. If a node on the network belongs to a certain multicast group, it responds to the queries, the router then records or refreshes its record of the node's interest in the traffic for the group, and the router forwards traffic to the network towards the node. The switches on the network with IGMP snooping enabled also learn from the responses and maintain their records about the nodes' interests in groups and forward traffic accordingly.

Hostnames are detected using Multicast Domain Name System (mDNS) packets, which are in multicast group 224.0.0.251.

An IP filter is a chain of rules for the treatment of packets. Refer to the *"Using IP Filter Chains for Security"* section in the *GigaVUE-OS CLI Reference Guide.*

Allow IGMP Traffic

If IP filtering is enabled (and IGMP snooping is enabled):

- Verify that IGMP traffic is allowed.
- For example, issue the following CLI commands:

  (config) # ip filter chain INPUT rule append tail target ACCEPT dup-delete protocol igmp

  (config) # ipv6 filter chain INPUT rule append tail target ACCEPT dup-delete protocol igmpv6
- Verify that mDNS traffic is allowed.

  If IGMP snooping is disabled, you do not need to allow IGMP traffic. However, you must allow UDP multicast traffic that targets 224.0.0.251. For example, issue the following CLI command:

  (config) # ip filter chain INPUT rule append tail target ACCEPT dup-delete dest-addr 224.0.0.251 /32

  where:

  dest-addr specifies the multicast group

Enable an IGMP Querier

If IGMP snooping is enabled:

- Check if there is an IGMP querier on the cluster network. The querier periodically sends queries that trigger the nodes in the cluster to send IGMP membership reports. For example, use a sniffer tool to verify if there is an IGMP querier on the network, such as Wireshark.
- IGMP snooping and IGMP snooping querier settings vary by networking switch. Refer to the respective documentation for how to configure them on your device.

When IGMP traffic is allowed and an IGMP querier is enabled in the network, the switches in the network will be refreshed through the IGMP membership reports.

Cluster Safe and Limited Modes

Starting in software version 4.7, safe and limited modes are introduced to safeguard critical provisioning errors for both standalone nodes and nodes in a cluster.

During provisioning operations such as configuring a map, in rare scenarios there can be unrecoverable system errors that can potentially put the cluster, clustered nodes, or standalone nodes into unsafe or unstable states. Once in such a state, additional operations or configuration changes can cause the node to crash, the cluster to deform, or data traffic to be impacted. For example, due to a node attempting to rejoin a cluster, a chassis can end up in a reboot loop. In previous software versions, there was no way to prevent entering the loop.

These modes provide notification, stop further operations from being performed, and give you time to troubleshoot and plan the recovery of the cluster, the clustered node, or the standalone node.

Two modes are supported. The first is called safe mode and is triggered when the node detects unrecoverable errors, but the existing flow maps are not impacted. The second is called limited mode and is triggered when the node detects continuous system reboots. In this mode, the node will become standalone and only basic configuration will be allowed.

Safe Mode

A node enters safe mode when there are unrecoverable errors. Any node in a cluster can enter this mode. The purpose of this mode is to detect system configuration failures early and avoid future failures, such as system crashes.

Examples of unrecoverable errors are when there are inconsistencies between the system and the running configuration or when the cluster configuration did not merge properly with the existing configuration.

As part of merge error recovery, nodes joining a cluster are automatically restarted so the merge error can be fixed. If the restart cannot correct the merge error, the node will enter safe mode.

Another example is that a TA Series node could enter safe mode when unlicensed cluster ports are used in an offline configured map. (It is recommended to use only licensed ports in map configurations.)

A node will automatically enter safe mode.

To recover from safe mode, reload the node. If safe mode persists, contact Gigamon Technical Support.

Limited Mode

A node automatically enters limited mode when it detects repeated system crashes.

Limited mode is triggered when there are three (3) failures/system crashes within 15 minutes. In limited mode, the cluster configuration is ignored. No cluster configuration or GigaVUE-OS configuration is accepted when the node is in limited mode.

When limited mode has been detected, collect information and report it to Gigamon Technical Support.

Support for Cluster Types

The GigaVUE-FM workflow supports only out-of-band clusters; not inband clusters. To create and manage an inband cluster, refer to the *GigaVUE-OS CLI User's Guide*.

Create Clusters

GigaVUE-FM 5.3 supports workflow-based configurations for forming clusters:

- Refer to *Regular Cluster Formation Workflow* on page 212 for instructions on how to use the regular cluster formation workflow.
- Refer to *Leaf-Spine Cluster Formation Workflow* on page 259 for how to use the leaf-spine cluster workflow

Regular Cluster Formation Workflow

Gigamon's Cluster formation can be done for any number of devices with different combinations of devices.

GigaVUE-FM 5.3 supports workflow-based configurations for forming a cluster. This workflow walks through the required steps to form a complete cluster for a regular cluster.

NOTE: Refer to *Leaf-Spine Cluster Formation Workflow* on page 259 for how to use the Leaf-Spine Cluster workflow

Deployment Checklist

Before forming a Cluster, it is strongly that you familiarize with the relevant documentation and review the deployment checklist to prepare for deployment.

Pre-deployment checklist

- Gigamon Fabric Management must be upgraded to GigaVUE-FM 5.3.00 or later
- Gigamon device must be upgraded to GigaVUE-OS 5.2.00 or later
- Advanced Features License must be installed in TA devices
- Physical connection must be established to create stack links
- Devices must have GDP enabled and be physically connected to create links among devices from GigaVUE-FM.

Create Regular Cluster Formation

To create a cluster:

1. Navigate to **Physical > Physical Nodes**.
2. Click **Create Cluster**.

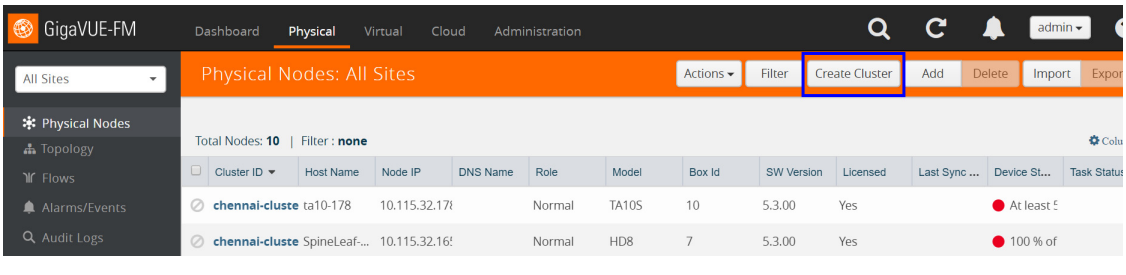

Figure 13-6: Create Cluster

Select Cluster Type

3. The Create a Cluster screen opens with two options:
   - CREATE A CLUSTER
   - CREATE A LEAF SPINE CLUSTER

4. Hover over the CREATE A CLUSTER option and click **Let's Begin** to start the wizard.

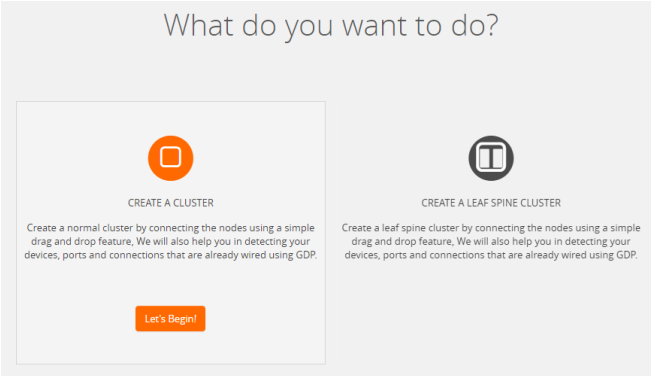

Figure 13-7: Choose CREATE A CLUSTER option

Select Devices

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

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

Figure 13-8: Select the required devices to form a cluster

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

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

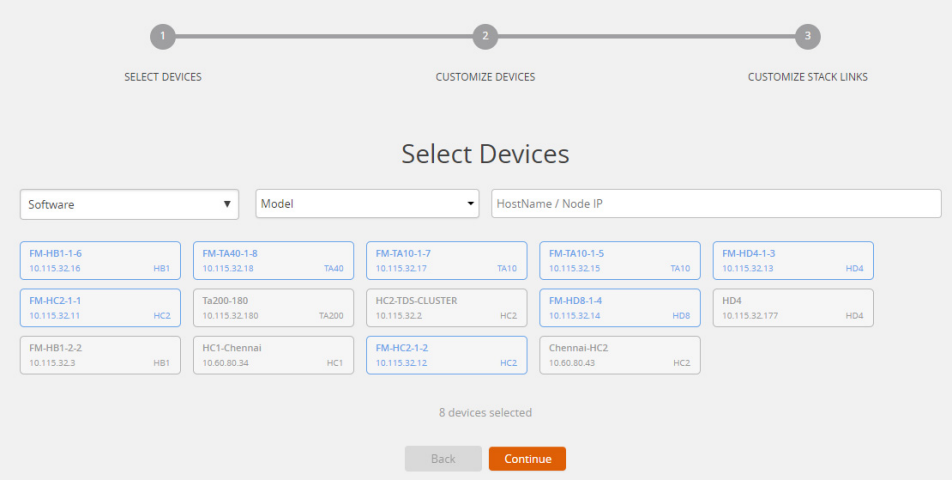

Figure 13-9: Selected devices will be highlighted

Customize Devices

Use the Cluster Configuration window to customize your devices.

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

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

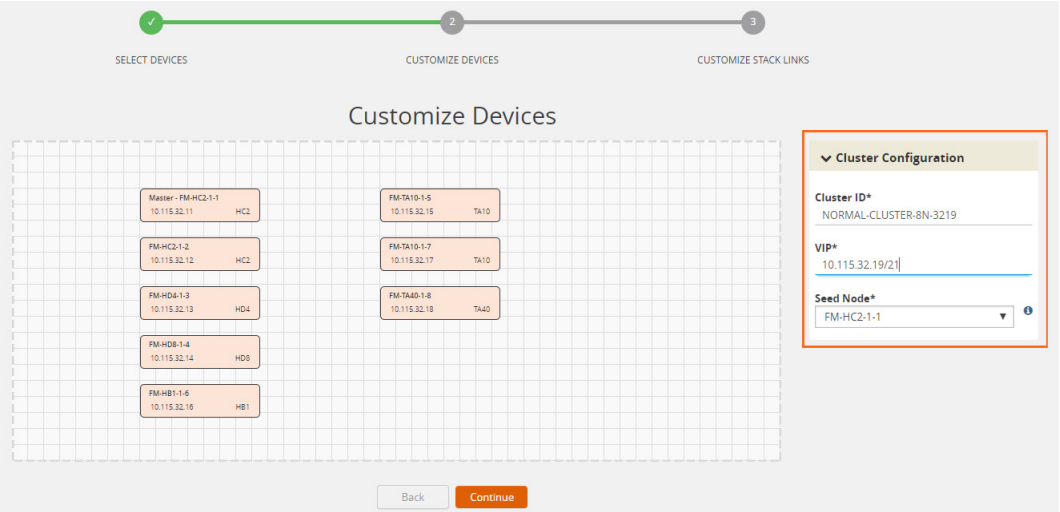

Figure 13-10: Fill up the cluster configuration details

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

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

Customize Stack Links

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

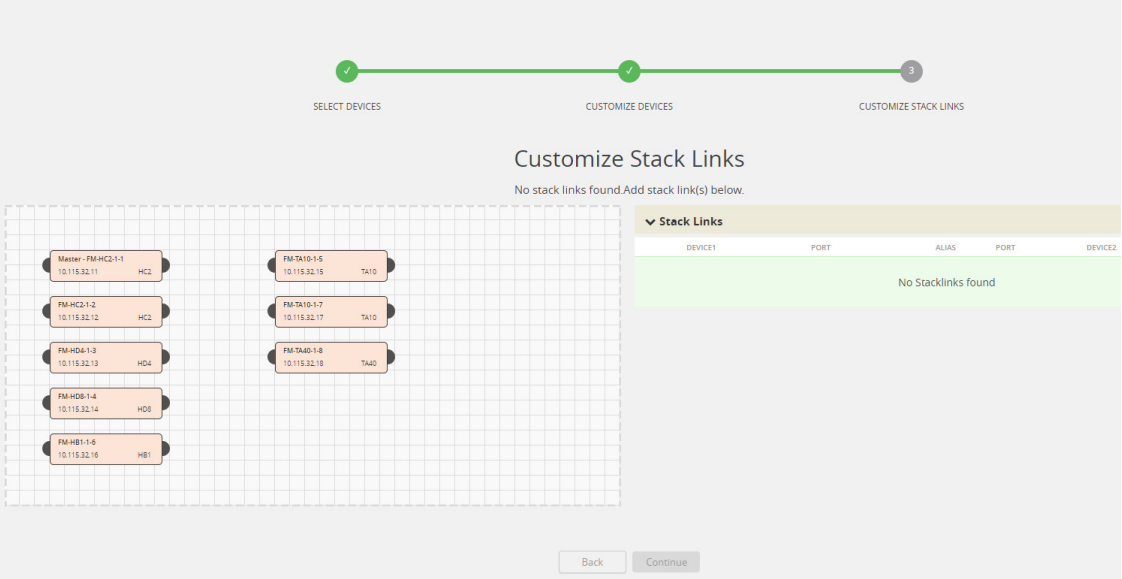

Figure 13-11: Create stack links by manually connecting the devices.

9. Connect any two devices to create a stack links between those two devices.

   Click the tip of the node and drag your cursor to the second node tip to create a link. After you create the link, a dotted line will illustrate the connection.

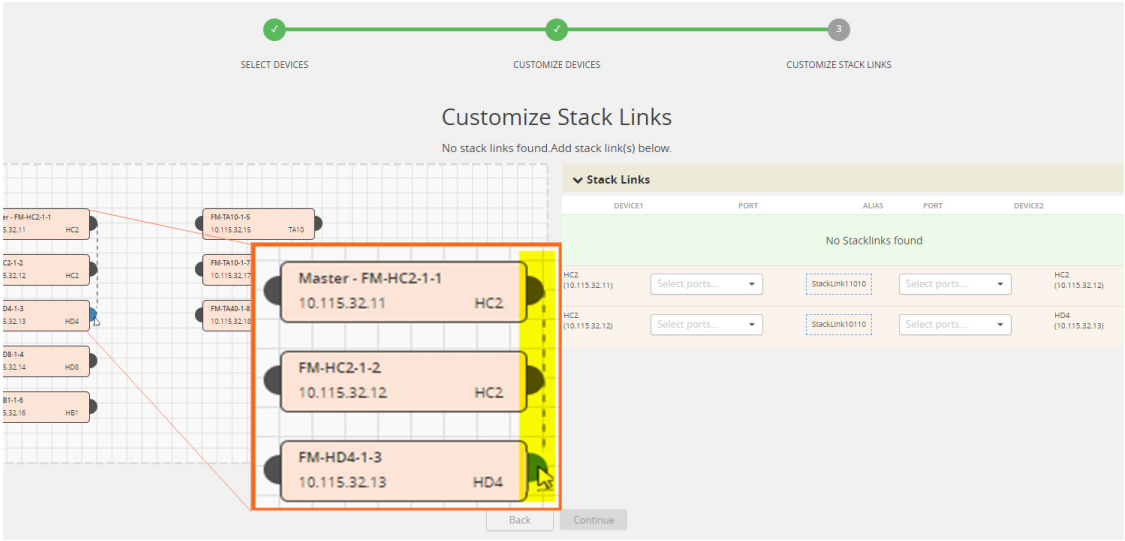

Figure 13-12: Create links between any devices

10. Configure the formed links in the **Stack Links** table as follows:
    - Select ports in each device that are compatible, for example: x-x ports ,x-q ports, q-c ports, x-c ports.
    - Select two or more ports in each device to create a stack GigaStream.

- After selecting the ports, save the stack link by clicking the **Save** button enabled in the right of stack link table.

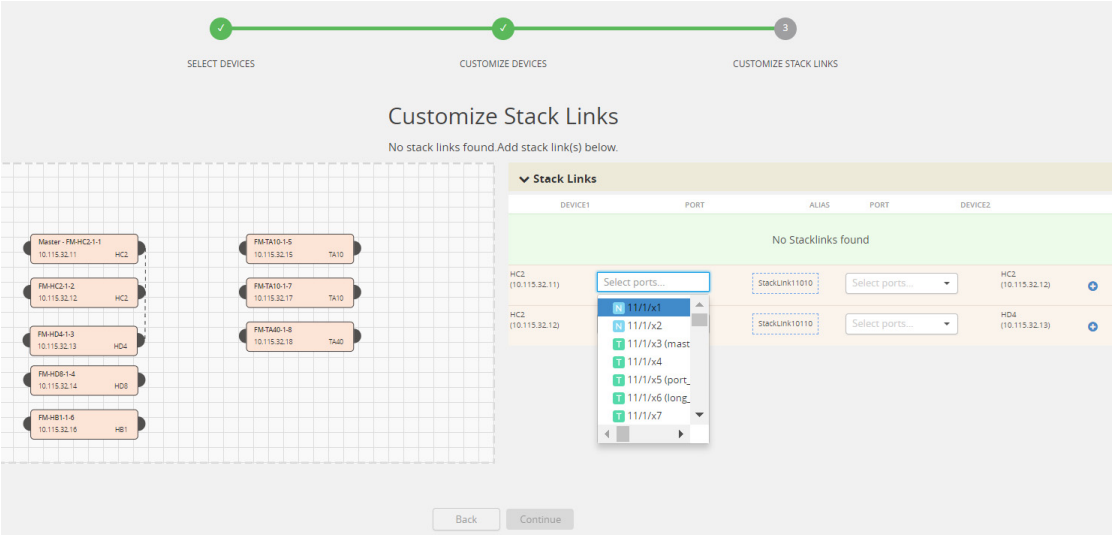

Figure 13-13: Select ports in each device

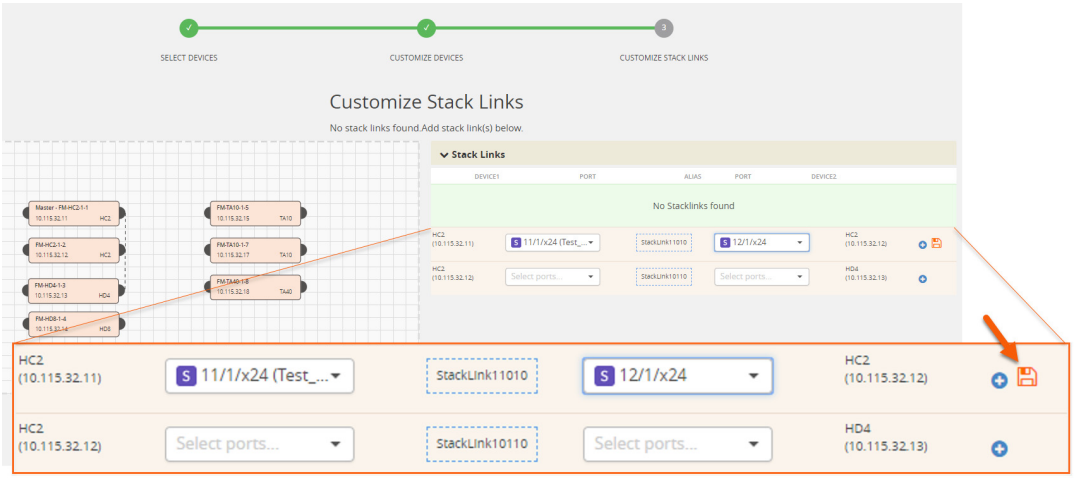

Figure 13-14: Save button will be enabled after selecting ports

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

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

The Creating Cluster page appears as the cluster is being created.

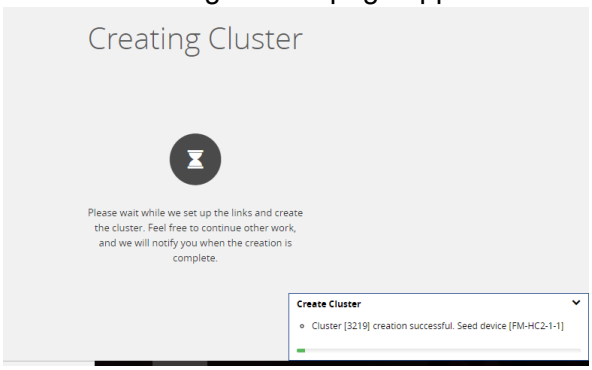

Figure 13-15: Notification to create a cluster

The Create Cluster progress window in the lower right corner of the page shows the status of every node as it joins the cluster. It takes a few minutes for the cluster to form. The cluster creation process involves the following steps:

- Cluster[clusterName] Creation Successful followed by Seed device
- Verifying Nodes[Will display HostName of all devices]
- Adding Node[HostName] to cluster [clusterName]
- Node[HostName] successfully joined to the cluster.
- Configuring cards for cluster[clusterName]
- Rediscovering cluster[clusterName]
- Configuring ports for cluster[clusterName].
- Configuring ports will display the status of each stack link and GigaStream whether the creation is successful or not.

NOTE: Refer to *Check Cluster Status* on page 227 for Alarms/Events.

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

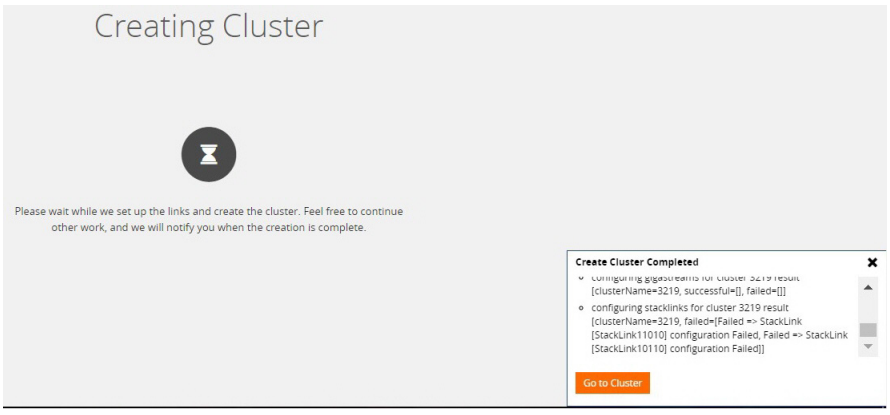

Figure 13-16: After cluster creation

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

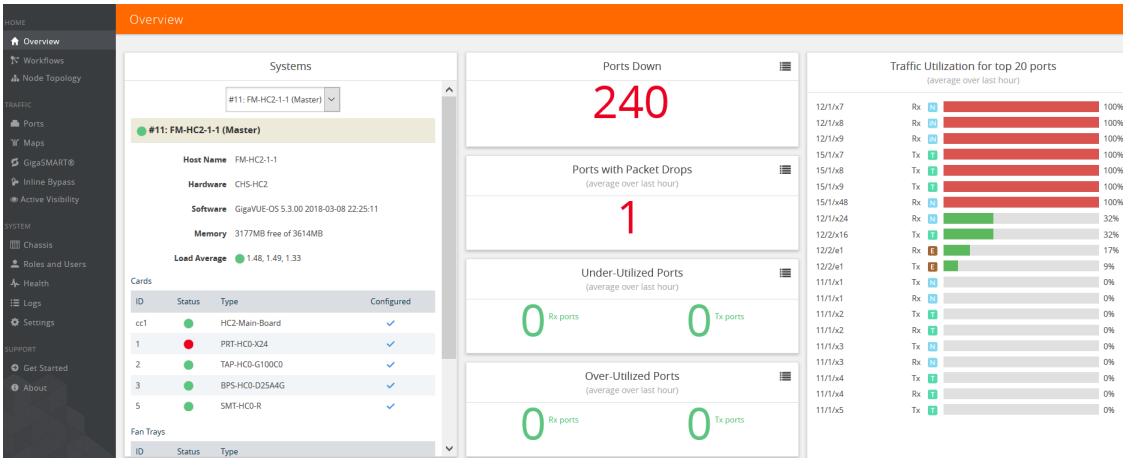

Figure 13-17: Cluster overview page after formation

Edit Cluster

The Edit Cluster action supports the following types of changes:

- Master preferences can be changed for each device.
- Multiple devices can be added to the existing cluster in a single update operation.
- Stack links can be created only from a new device that is being added.
- Stack link alias and GigaStream alias can be edited for newly created links.
- Nodes can be removed from an existing cluster, one at a time.

The following options are not supported by the Edit Cluster action:

- There is no option to remove existing stack links through the cluster wizard.
- There is no option to create links in existing devices.
- Addition and deletion of devices in a single update operation is not supported and is not recommended. If you attempt to add and delete devices in a single update operation, you may get unexpected results.
- There is no option to edit the existing stack link alias and GigaStream alias.

Prerequisites:

You must clear all configurations on a node before adding it to a cluster.

Edit options:

- *Add Nodes to a Cluster* on page 222
- *Remove Nodes from a Cluster* on page 224
- *Edit Cluster Parameters* on page 226

Inband Cluster Management

Inband Cluster Management simplifies traditional network management and maintenance by creating a virtual device to manage multiple physical nodes. This simplified approach makes it possible to oversee large networks by defining policies that span across multiple devices. The Inband Cluster Management feature is designed to reduce operational cost and extend coverage by eliminating a dedicated management network.

Inband Cluster Management is supported on all GigaVUE-OS nodes.

Inband Cluster Management Topologies

The benefits of Inband Cluster Management are to eliminate the Layer-2 Ethernet network and create a virtual management network through the data path where the data traffic is flowing.

Inband Cluster Management supports multiple topologies that include:

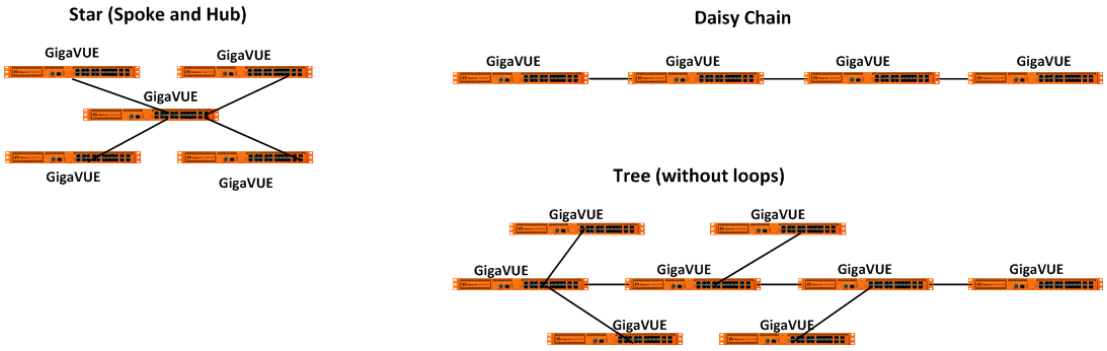

NOTE: Subsets or aggregations of these topologies may be created; however, it is important not to create a loop within these specified topologies.

Loops are typically created in the following scenarios:

- **Two Node Loops** occur between two nodes in a cluster forming two or more stack links and the stack links are not contained in one GigaStream.
- **Multi-Node Loops** occur when multiple nodes form a cluster whereby a link connects between node A and node B, another link connects between node B and node C, and yet another link connects between node C and node A.

Inband Cluster Management Stack Ports

A common Inband Cluster Management topology is configured between the Layer 2 device's Ethernet management port to a GigaVUE-OS node using a stack port configuration.

Two or more GigaVUE-OS nodes may be directly connected in a one-to-many relationship between physical connections. GigaVUE-OS nodes may also be indirectly connected if there is a path of stack ports between the nodes.

Inband Cluster Management Stack Ports Example

Figure visually depicts how Inband Cluster Management uses the grouping of stack ports to connect between GigaVUE-HB1 nodes.

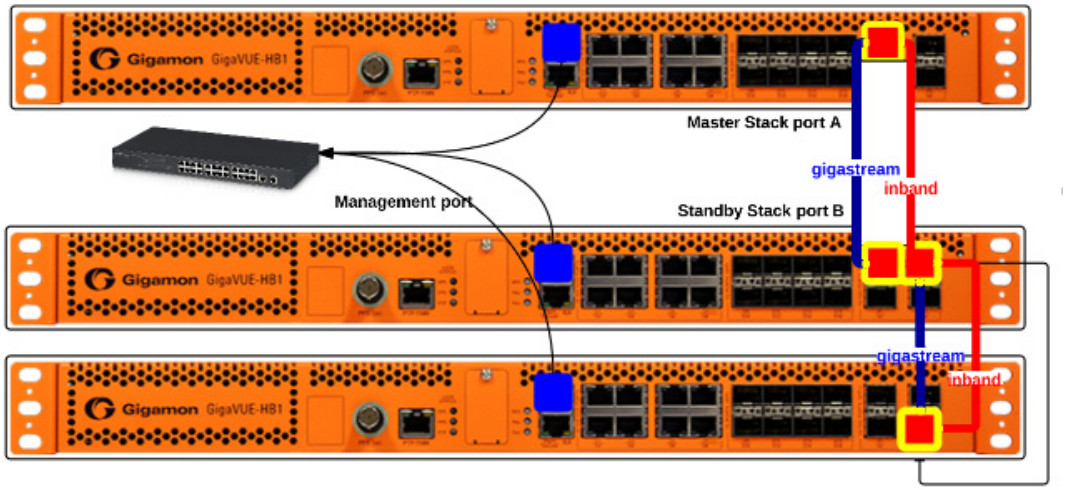

Figure 13-18: Inband Cluster Management Stack Port Configuration on GigaVUE-HB1

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

Inband Cluster Management Configuration Flow Chart

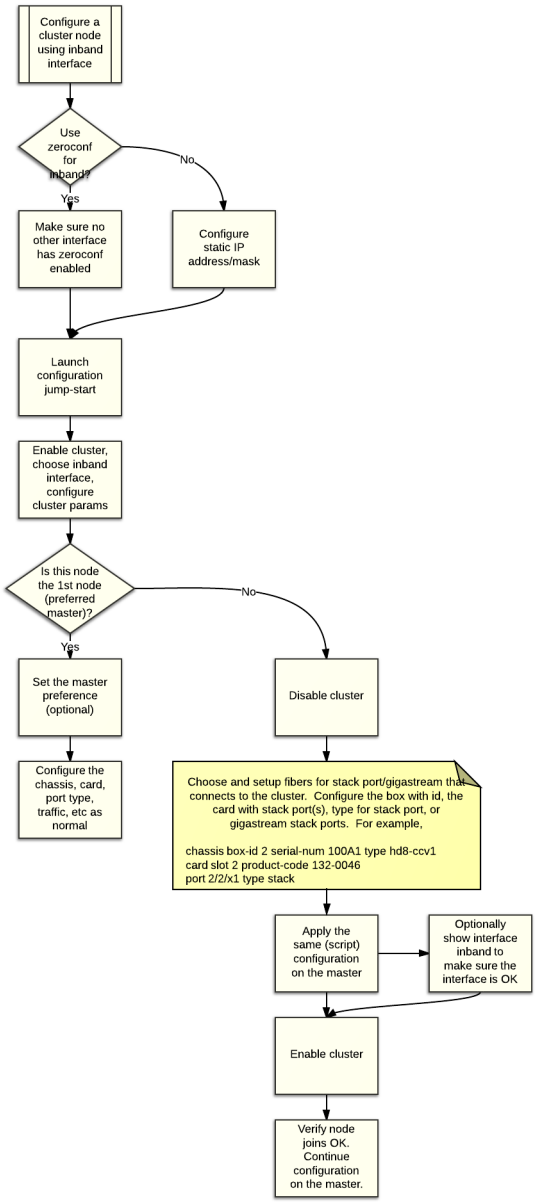

Inband Cluster Management Configuration

An interface called "Inband" has been created upon boot-up to ensure backward compatibility with an existing clustered infrastructure. This interface has similar properties and characteristics of a typical Ethernet interface such as Eth0.

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

Enable Cluster Management for GigaVUE TA Series Nodes

To enable clustering, GigaVUE TA Series nodes require an Advanced Features License. This license can be obtained by contacting Gigamon Sales team. In order to obtain the license for a Gigamon node, have the node serial number available. All licenses are tied to the serial number and cannot be moved.

For licensing the GigaVUE-OS on a white box, users can access the GigaVUE-OS licensing portal and obtain the license key online. In order to generate the license, the following are required: the serial number of the white box, digital footprint, and Gigamon Installation Key (GIK).

Add Nodes to a Cluster

You can manage an existing cluster through GigaVUE-FM by adding nodes to it. The nodes must be standalone nodes that are currently managed by GigaVUE-FM.

When a node joins an existing cluster, all of its existing traffic configuration, including maps, will be replaced by the configuration of the master node.

The following is an example of adding nodes to an existing cluster using GigaVUE-FM.

1. Navigate to **Physical > Physical Nodes**.
2. Select a cluster and choose **Actions > Edit Cluster**. The Edit Cluster - Canvas appears showing the existing stack link configuration details in the cluster wizard. Standalone devices appear in the Devices pane.

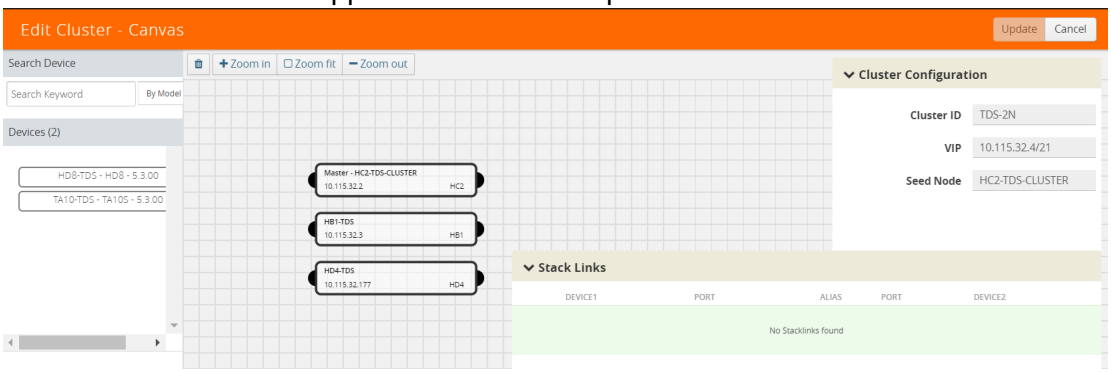

Figure 13-19: Edit cluster page

3. Drag the required devices from the Devices pane into the Edit Cluster canvas under the leaf or spine container.

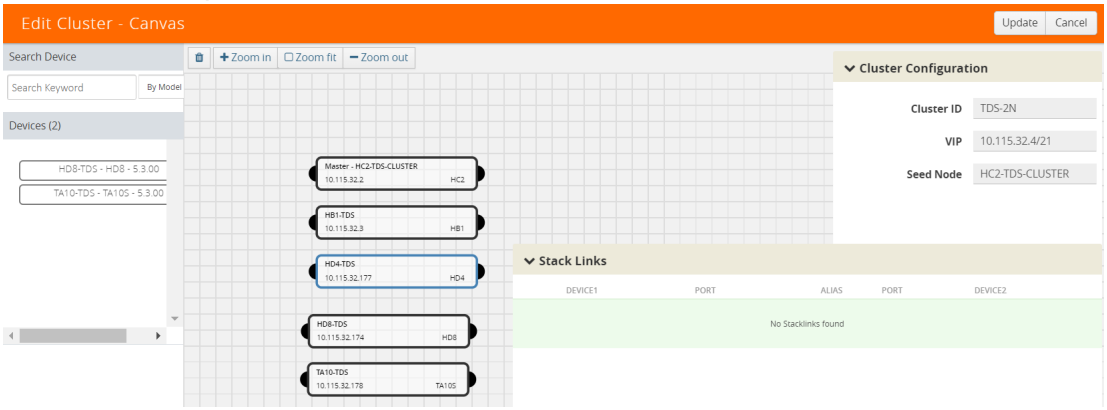

Figure 13-20: Drag the devices from devices column.

4. Connect the newly added devices to other devices to create stack links.

   Click the tip of the node and drag your cursor to the second node tip to create a link. After you create the link, a dotted line will illustrate the connection

   NOTE: No new link is created for existing devices; they need to be added manually.

5. Configure the stack link details in the stack link table and click **Save**.
6. Click **Update** to update the configuration.

   A Confirmation window appears advising that all traffic configurations will be erased on newly added or removed nodes.

7. Click **OK** to continue.
8. The Manage Cluster update notification window appears showing the status of each update activity on the nodes, cards, GigaStreams and stack links.

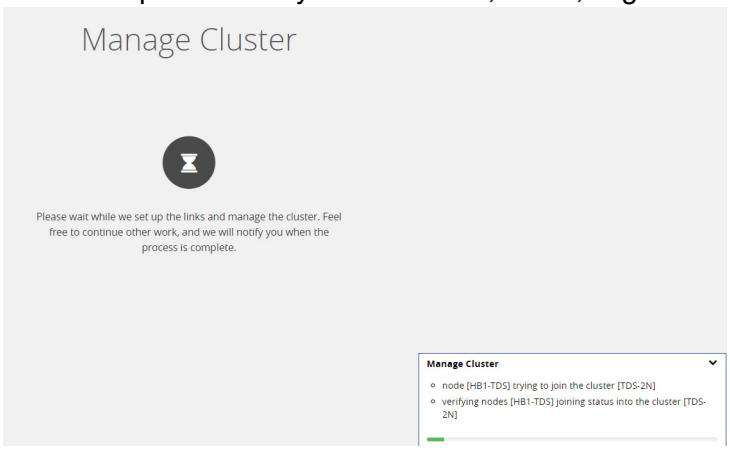

Figure 13-21: Notification pop up after update operation

9. After the cluster update operation completes, a “Manage Cluster Completed” message appears.

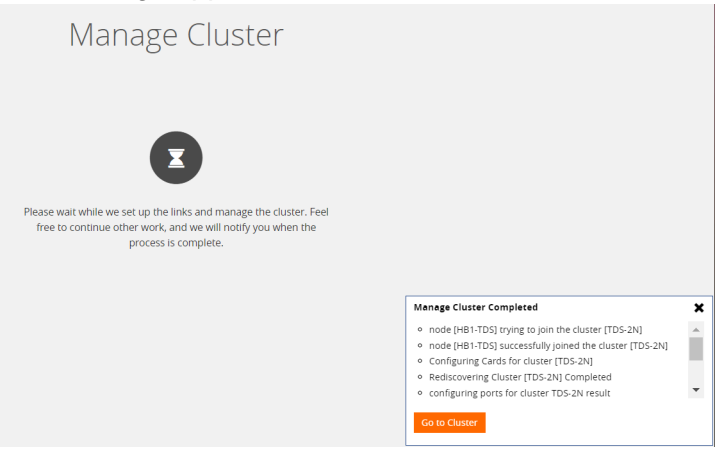

Figure 13-22: Cluster completion event

10. Click **Go to Cluster** to view the updated cluster overview

Remove Nodes from a Cluster

You can manage an existing cluster through GigaVUE-FM by removing nodes from it. After a node is deleted from a cluster, it will become a standalone node. FM will continue to manage it.

For nodes leaving a cluster, the username and password of the **admin** account on the cluster will be used for managing the node after it has been removed from the cluster.

To remove nodes of an existing cluster using GigaVUE-FM:

1. Navigate to **Physical > Physical Nodes**.
2. Select a cluster and choose **Actions > Edit Cluster**.

   The Edit Cluster - Canvas appears showing the existing stack link configuration details in the cluster wizard. Standalone devices appear in the Devices pane.

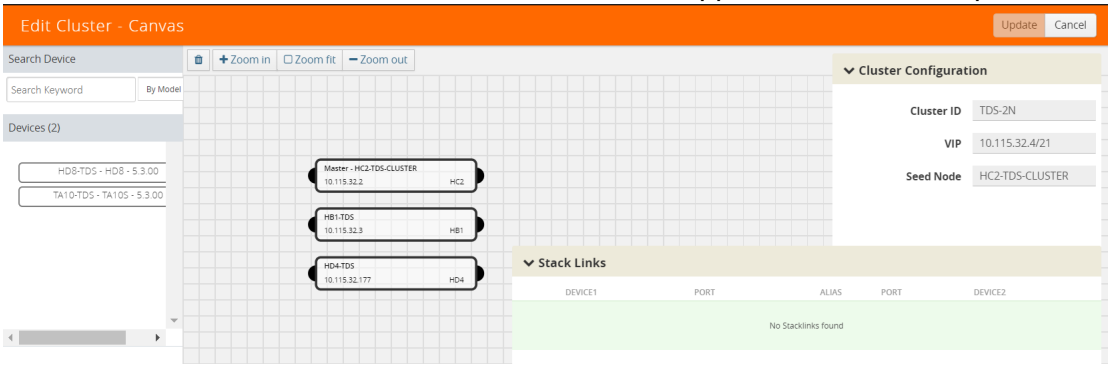

Figure 13-23: Edit cluster page

NOTE: Only one device can be deleted from canvas per operation. It can be any device from the cluster.

3. Right-click the device to be removed from the canvas and click **Remove**.

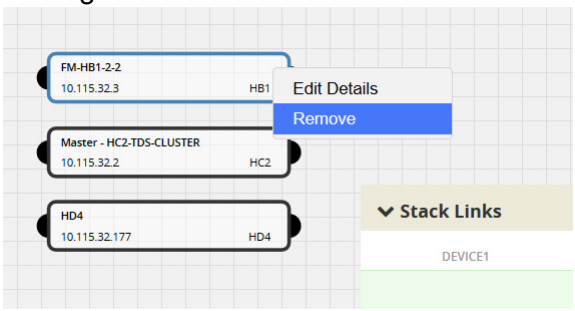

Right click to remove

4. The removed device is deleted from the canvas.
5. Click **Update** to update the configuration.

   A Confirmation window appears advising that all traffic configurations will be erased on newly added or removed nodes.
6. Click **Ok** to continue.
7. The Manage Cluster update notification window appears showing the status of each update activity on the nodes, cards, GigaStreams and stack links.

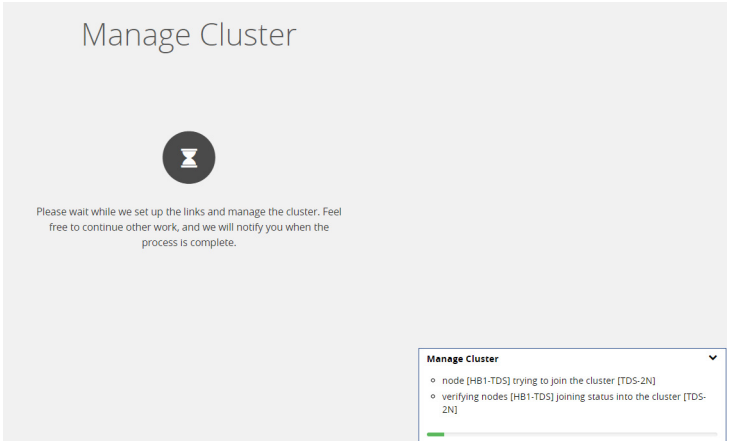

Figure 13-24: Notification pop up after update operation

8. After the cluster update operation completes, a “Manage Cluster Completed” message appears.
9. Click **Go to Cluster** to view the updated cluster overview.

Edit Cluster Parameters

When editing a cluster node, you can only edit the cluster master preference. You can only change the cluster master preference on one node at a time.

For the master preference, higher values are more likely to claim the Master and Standby roles; lower values are less likely.

To edit master preferences:

1. To set the master preference for devices, right-click the required device and click the **Edit Details** options button.

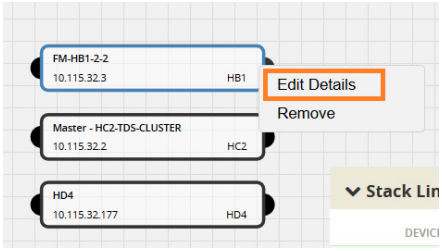

Figure 13-25: Right click to edit the details

2. The Device configuration quick view should appear on the right. Edit the Master Preference in the text box.

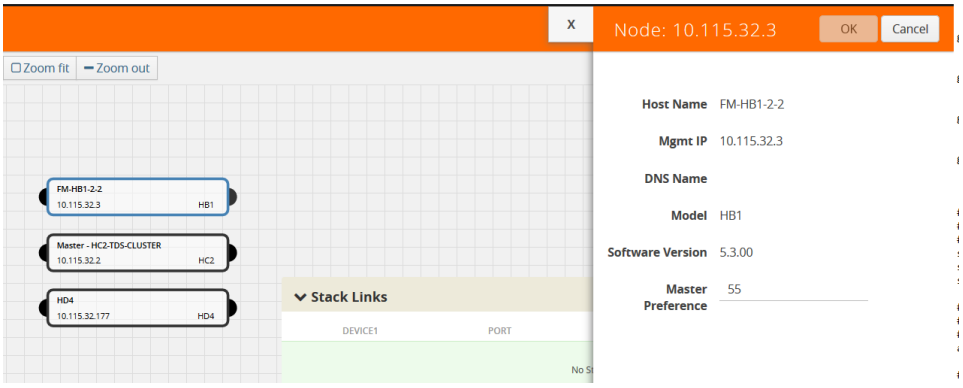

Figure 13-26: Set master preference

3. Click **OK**.

NOTE: Most of the fields are read-only. You can change the cluster master preference, as desired. Use preference settings from 10 to 100 for master, standby, and normal roles. Use preference settings from 1 to 9 for normal nodes that are excluded from taking the master or standby role. GigaVUE TA Series nodes always have a preference of 1.

4. After saving your changes to the nodes, click **Update** to apply the changes to the cluster.

   A Confirmation window appears advising that all traffic configurations will be erased on newly added or removed nodes.

5. Click **Ok** to continue.

6. The Manage Cluster update notification window appears showing the status of each update activity on the nodes, cards, GigaStreams and stack links.

7. After the cluster update operation completes, a “Manage Cluster Completed” message appears.

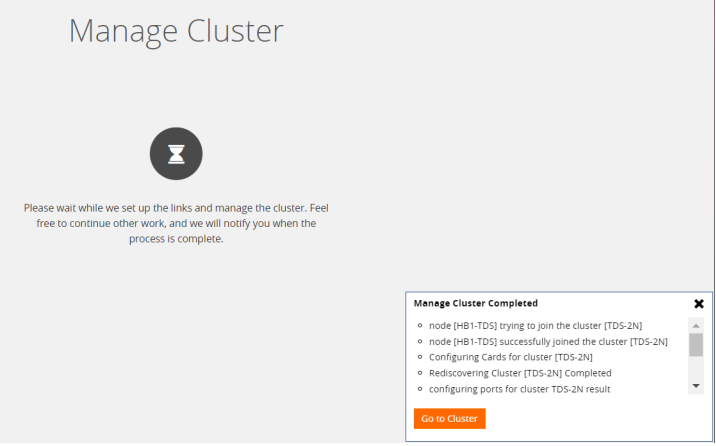

Figure 13-27: Cluster completion event

8. Click **Go to Cluster** to view the updated cluster overview

Check Cluster Status

When a cluster is being created, you can check the status through cluster management events or audit log entries. Refer to the following sections:

- *Cluster Management Events* on page 227
- *Audit Logs* on page 228

Cluster Management Events

On the **Alarms/Events** page, the following event types indicate the progress of the cluster as it is being formed:

- ClusterCreationStarted
- ClusterCreationCompleted
- ClusterCreationFailed

Refer to Figure 13-28 on page 227 for the cluster creation event on the **Alarms/Events** page.

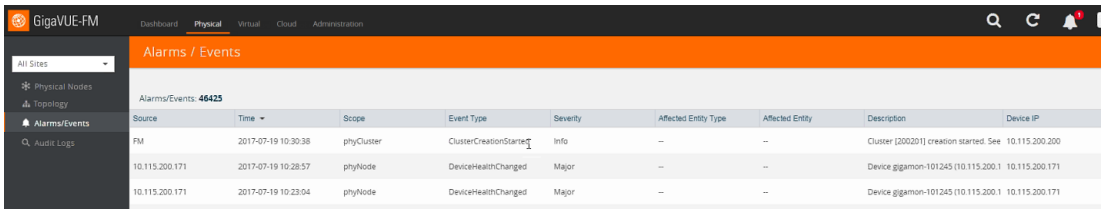

Figure 13-28: Cluster Creation Event

The following events indicate the status of nodes added to or removed from the cluster:

- NodeJoinedToCluster
- NodeFailedToJoinCluster

- NodeRemovedFromCluster
- NodeFailedToRemoveFromCluster

Refer to Figure 13-29 on page 228 for the node joined to cluster event on the **Alarms/ Events** page.

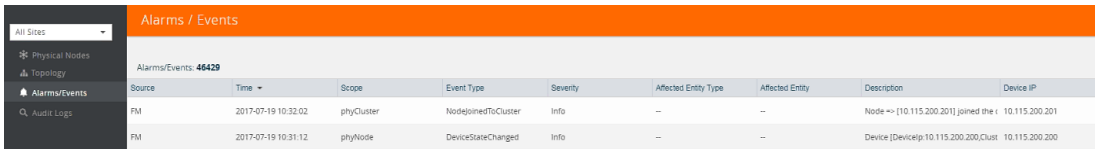

Figure 13-29: Node Joined to Cluster Event

Audit Logs

The following audit logs indicate the cluster management actions triggered by users from GigaVUE-FM:

- User <username> created cluster <clustername>
- User added device <device IP> to cluster <clustername>
- User removed device <device IP> from cluster <clustername>

Export Nodes and Clusters

To export the nodes and clusters:

1. Click **Physical** on the top navigation link.
2. In the physical nodes page, select the nodes and clusters you want to export. You can select a single node or the nodes and clusters associated to the selected site.
3. Click **Export**. The physical nodes table is downloaded with a filename in the format nodes_<yyyymmddhhmmss>; for example, nodes_20161003172336.

Upgrade Software on a GigaVUE Node or a Cluster from GigaVUE-FM

This section describes the steps to upgrade the software on GigaVUE nodes and clusters. You can upgrade software by using an image that is located on an external image server, or you can use GigaVUE-FM as the image server for H Series and TA Series nodes. G Series nodes can only be upgraded with an external image server. Refer to *Upgrade G Series Node* on page 231.

NOTE: When using GigaVUE-FM to upgrade software on the nodes, it is imperative that the minimum version on the node be at version 4.3.xx. If the software is not updated to version 4.3.xx, refer to the *H Series Upgrade Guide* and use either the CLI or GigaVUE-FM at the node level to upgrade the node to the latest revision.

Before upgrading a node or a cluster, it is important to note the following:

- GigaVUE-FM can upgrade up to 32 nodes at a time. For example, if you select a cluster with 8 nodes and 28 standalone nodes, GigaVUE-FM will fail to upgrade as the total number of nodes exceeds the limit.
- Check if the device that you want to upgrade has enough disk space. The disk space is computed based on the last synchronized time. A newly added node may sometimes show low disk space. Wait for approximately 5 minutes for the configuration sync to complete or rediscover the device and then check the disk space.
- When the upgrade process is complete, a post-upgrade sanity check is performed. Physical inventory snapshot of nodes, cards, ports, maps, and GigaSMART operations are captured and the numbers from before and after the upgrade are compared. Sometimes, these numbers do not match immediately after the upgrade, as the device is still being configured. GigaVUE-FM checks every minute whether the nodes have joined the cluster. If the nodes joined do not increment in 30 minutes, then the sanity check will fail.

  After the nodes join the cluster, GigaVUE-FM again checks every minute for other configuration objects to come up. If the configuration objects do not increment in 10 minutes, the sanity check will fail. When the sanity check fails, GigaVUE-FM provides an ability to view the configuration object that failed the sanity check.
- GigaVUE-FM v5.5.00 supports the installation of bootloader on GigaVUE H Series and GigaVUE TA Series nodes during the GigaVUE-OS image upgrade, but it does not automatically activate it.

  NOTE: Depending on your configuration, you may need to enable physical bypass mode before performing the GigaVUE-OS image upgrade, and then restore it after the upgrade is complete.

  During the GigaVUE-OS image upgrade, GigaVUE-FM v5.5.00 first upgrades the GigaVUE-OS image and then reboots the system. After the GigaVUE H Series or GigaVUE TA Series node has rebooted with the new GigaVUE-OS image, the bootloader installation will automatically initiate from GigaVUE-FM.

  NOTE: The bootloader upgrade will apply only if there is a bootloader version change in the specific GigaVUE-OS software release. Refer to the GigaVUE H Series Upgrade Guide and GigaVUE TA Series Upgrade Guide for the upgrade procedures and to check the bootloader version changes.

  When the installation process completes, a message about the bootloader installation will appear in the GigaVUE-FM Task Log Details.

  IMPORTANT: The newer bootloader version only goes into effect after an additional manual reboot. You must explicitly reload the GigaVUE H series or GigaVUE TA series device to update the bootloader version. This additional reboot will cause the optical-protection switches of the relevant Bypass Modules to change states multiple times.

  RECOMMENDATION: To avoid this state-change behavior, enable the physical bypass before performing the GigaVUE-OS image upgrade. Refer to “Configuring Inline Bypass Examples: Protected Inline Bypass Using Bypass Combo Modules” in the GigaVUE-OS CLI User’s Guide for instructions on how to enable the physical bypass.

NOTE: Starting in GigaVUE-OS software version 4.7, the default password admin123A! for the admin user is no longer valid when logging in to the node. When upgrading from

the CLI, the configuration jump-start script requires the user to change the default password for the admin user. However, if the node is upgraded to software version 4.7 or later using GigaVUE-FM, the admin password on the node is not changed.

Upgrade from an External Image Server

This section provides the steps for upgrading the GigaVUE nodes and clusters from an image stored on an external server. The image can be transfered from the server to the node using either SCP or TFTP file protocols.

In a cluster configuration, to upgrade the software on the GigaVUE nodes, all the relevant images must be available to GigaVUE-FM. All nodes should be updated to the same version of software.

To upgrade a node or a cluster from an image stored on an external image server, do the following:

1. Upload the image to the external image server to make it available to GigaVUE-FM.

   To obtain software images, register on the customer portal (https://gigamoncp.forc.com/gigamoncp/) and download the software. You must provide the serial number for each node you want to update.

   To view the chassis serial number, login to H-VUE and select **Chassis** from the navigation pane. Click **List View**.
2. Click **Physical** on the top navigation link.
3. On the Physical Nodes page, select one or more nodes or clusters to upgrade.
4. Select **Actions > Image Upgrade**.

5. Enter the following information:

Table 13-2: Image Upgrade

Menu	Description
Task Name	The name of the upgrade task.
Image Server	The location from where the image can be uploaded. Do the following: 1. For Image Server, choose **External Server**. 2. From the drop-down list, select the external image server added to GigaVUE-FM in *Step 1*. 3. If the external image server is not available, click **Add External Server.** In the Add External Server quick view, enter the following: a. In the **Alias** field, enter the name of the external server. b. In the **Server Address** field, enter the host IP address of the server. c. In the Type drop-down list, select SCP or TFTP as the file transfer protocol. d. If you select SCP, enter the username and password of the server in the **Username** and **Password** fields respectively. e. Click **Add**. 4. In the text box, enter the image path and the image name. **NOTE:** For a cluster, the images can reside on different paths. However, the image server and the protocol for file transfer should be the same for all the nodes.
Backup Config Before Upgrade	The check box to back up the configuration changes prior to performing the upgrade.
Reboot After Upgrade	The check box to reboot the server after performing the upgrade.
Time	The time for performing the upgrade. There are two options: • Immediate—The upgrade is performed immediately. • Scheduled—The upgrade is performed at a scheduled time. Select the date and time.

NOTE: If the disk space of a physical node or a node in a cluster is low, a message is displayed indicating that the node will not be upgraded due to low disk space.

5. Click **Upgrade**.

Upgrade G Series Node

When upgrading a G Series node from GigaVUE-FM, the Image Upgrade page includes a **File Type** field for specifying the type of file to use for the upgrade. Refer to Figure 13-30 on page 232.

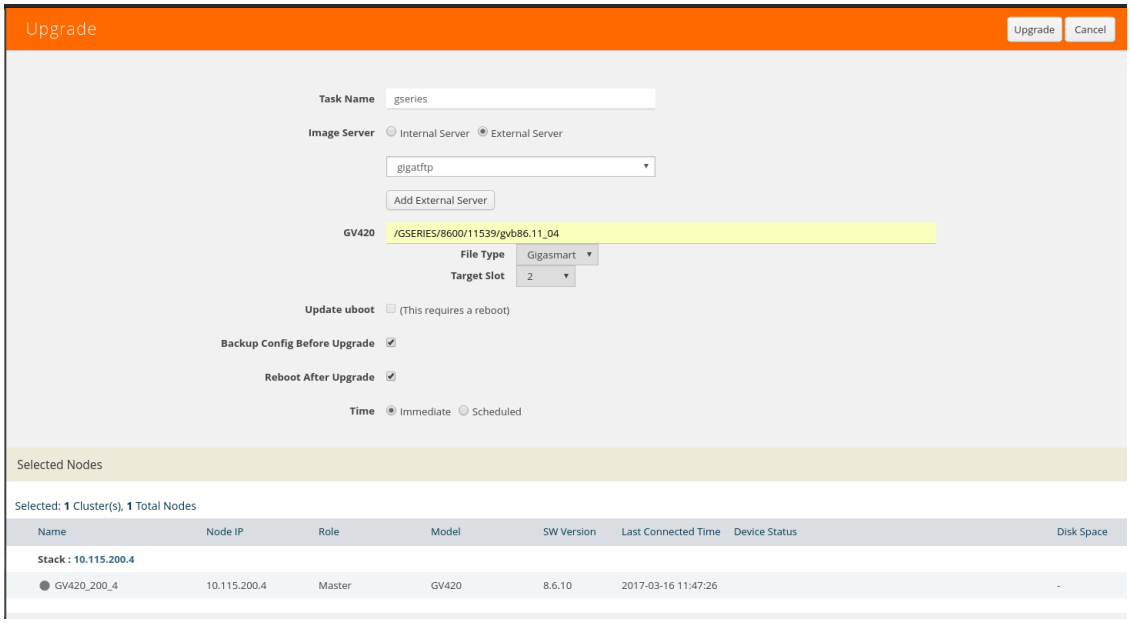

Figure 13-30: Image Upgrade Page for a G Series Node

The possible selections for **File Type** are as follows:

- If you do not make any selection (that is, File Type field is left blank), the CLI image file is used to upgrade the node. The name of the file depends on the model that is being upgraded. The following are example filenames for each model:
  - GV212: gvd8611
  - GV216: gve8611
  - GV420: gvb8611
  - GV2404: gvc8611
- Select **Web** to upgrade the Web server on the node with the Web server image file. Web server image files starts with gvweb. This file is contained in a .tgz file (for example, gvweb8511.tgz) that you must expand on the image serve. After the file is expanded make sure that you have the files and directories shown in the following figure:

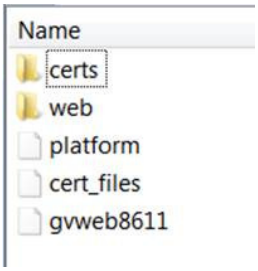

- Select **GigaSMART** to upgrade the GigaSMART card on the node with the GigaSMART image file. The GigaSMART image file name starts with gs. For example, gsc86_08_27.tgz. If you select this option, the **Target Slot** field displays (refer to Figure 13-30 on page 232). Use this field to specify the slot where the GigaSMART card is located.

Upgrade with GigaVUE-FM as the Image Server

This section provides the steps for upgrading the GigaVUE nodes and clusters when GigaVUE-FM is used as the file server instead of an external server.

NOTE: Only H Series and TA Series devices can be upgraded using GigaVUE-FM as the image server. To upgrade G Series devices with GigaVUE-FM, you must use an external server.

To upgrade a node or a cluster using internal image files, do the following:

1. Download the images from the Gigamon website and place them where they can be available for uploading to GigaVUE-FM.

   To obtain software images, register on the customer portal (https://gigamoncp.forc.com/gigamoncp/) and download the software. You must provide the serial number for each node you want to update.

   To view the chassis serial number, login to GigaVUE-OS and select **Chassis** from the navigation pane. Click **List View**.
2. Click **Physical** on the top navigation link.
3. On the Physical Nodes page, select one or more nodes or clusters to upgrade.
4. Select **Actions > Image Upgrade**. Under Selected Nodes, the disk space of all the nodes are displayed.

NOTE: If the disk space of the selected node or a node in the cluster is low, a message will be displayed indicating that the node will not be upgraded due to low disk space.

5. Enter the following information:

Table 13-3: Image Upgrade

Menu	Description
Task Name	The name of the upgrade task.
Image Server	The location from where the image can be uploaded. Do the following: **1.** Choose Internal Server. **2.** From the Version drop-down list, select the version to which you want to upgrade. **3.** Click **Add Image Files**. **4.** In the Add Image File quick view, click **Choose Files** and select the image files that you downloaded in *Step 1*. Click **Add**. **NOTE:** If there is an image file missing in a cluster configuration, the Image file missing error is displayed.
Update uboot	The check box to update the binary bootloader or coreboot included with the active/boot image. Coreboot only applies to a GigigaVUE-TA100, GigigaVUE-TA100-CXP, GigaVUE-HC3, or GigaVUE-HC1 node. A reboot is required after the update.
Backup Config Before Upgrade	The check box to back up the configuration changes prior to performing the upgrade.

Table 13-3: Image Upgrade

Menu	Description
Reboot After Upgrade	The check box to reboot the server after performing the upgrade. If you do not select **Reboot After Upgrade**, you must reload the node manually. To reload the node manually, select the node on the Physical Nodes page and click **Actions > Reboot**.
Time	The time for performing the upgrade. There are two options: • Immediate—The upgrade is performed immediately. • Scheduled—The upgrade is performed at a scheduled time. Select the date and time.

5. Click **Upgrade**.

Problems with SCP?

After upgrading GigaVUE-FM to a new release, under some circumstances you may find that a previously-managed H Series node no longer accepts SCP commands to backup or restore configuration files. This can happen when the SSH keys in use change, causing a mismatch between the keys stored on the H Series node and those presented by GigaVUE-FM. Use the following steps on the H Series node to remove the GigaVUE-FM server from the H Series node's list of addresses, resolving the issue:

1. Log in to the affected H Series node and switch to Configure mode.
2. Use the **ssh client user admin known-host?** command to discover the IP address for the GigaVUE-FM server. For example:

   (config) # ssh client user admin known-host ? 10.150.100.23 10.150.100.77
3. The question mark (**?**) instructs the H Series node to list the known ssh clients. From the list of IP addresses returned by the CLI, identify the one belonging to GigaVUE-FM and remove it using the **remove** argument. For example, if 10.150.100.77 is the GigaVUE-FM server's IP address:

   (config) # ssh client user admin known-host 10.150.100.77 remove

   Return to GigaVUE-FM and attempt the configuration backup again.

Alarms and Events

The Alarms and Events page displays all the alarms and events that occur in the physical nodes or clusters associated to the selected site. An event is an incident that occur at a specific point in time. Examples of events include:

- Device status change
- Stack image install status
- Fan tray changed

An Alarm is a response to one or more related events. If an event is considered of high severity, then GigaVUE-FM raises an alarm. Examples of alarms include:

- High or low port utilization
- High or low CPU utilization
- High exhaust temperature

The alarms and events broadly fall into the following categories: Critical, Major, Minor, or info.

Click **Physical** on the top navigation link. On the Physical Nodes page, select All Sites or a specific site on the left navigation pane and click **Alarms/Events**.

For information about the parameters for each alarm or event, refer to *All Alarm/Event Parameters* on page 1270. For filtering the alarms and events, refer to *Filter Alarms/Events* on page 1272.

NOTE: The alarms and events can be purged or archived only from the All alarms/events page. For more information, refer to *Archive or Purge Alarm/Event Records* on page 1273.

Audit Logs

With Audit Logs, changes and activities that occurred in the physical nodes or clusters due to user actions can be easily tracked for auditing. Audit logs can be displayed for All Sites or a specific site. The logs can also be further filtered to view specific information.

For information about the parameters in the audit log page, refer to *Overview of Audit Logs* on page 1276. Filtering the audit logs allows you to display specific type of logs. For more information, refer to *Filtering Audit Logs* on page 1276.

Search for Specific Nodes Using Keywords

The filter option provides a way for the users to narrow down the display using certain keywords such as Standalone, Clusters, H Series and others. As you click on the Filter button, you will see the quick view window pop-up.

The Filter quick view provides you filter criteria for your search. These options are available in the drop down menu under Criteria. You can further narrow the options using the Model, Software Version #, Cluster ID, Host Name, DNS Name, or Node IP. You are not required to fill in all these options to narrow your search. As you select these options in the quick view, you will see the options narrowing in the main window.

To clear or revert the search, do any of the following:

- To clear a part of the search, use backspace to clear the search entry and re-type a new option.
- To clear all the search criteria, use the Clear button on the top of the quick view window.
- To revert to the main window with the new searches, click on the **X** of the quick view window.

To revert to all nodes visible, you can use the clear filter option on top of the main window or the clear option in the quick view window. When no filters are in place, this option will no longer be visible in the main window.

The following screens show the different dashboards view for H Series and G Series, respectively. Note the changes in the navigation pane. H Series provides more detailed view of the physical nodes. All TA Series nodes are treated as H Series nodes, therefore you will see the same information displayed for all TA Series nodes including the white boxes with GigaVUE-OS.

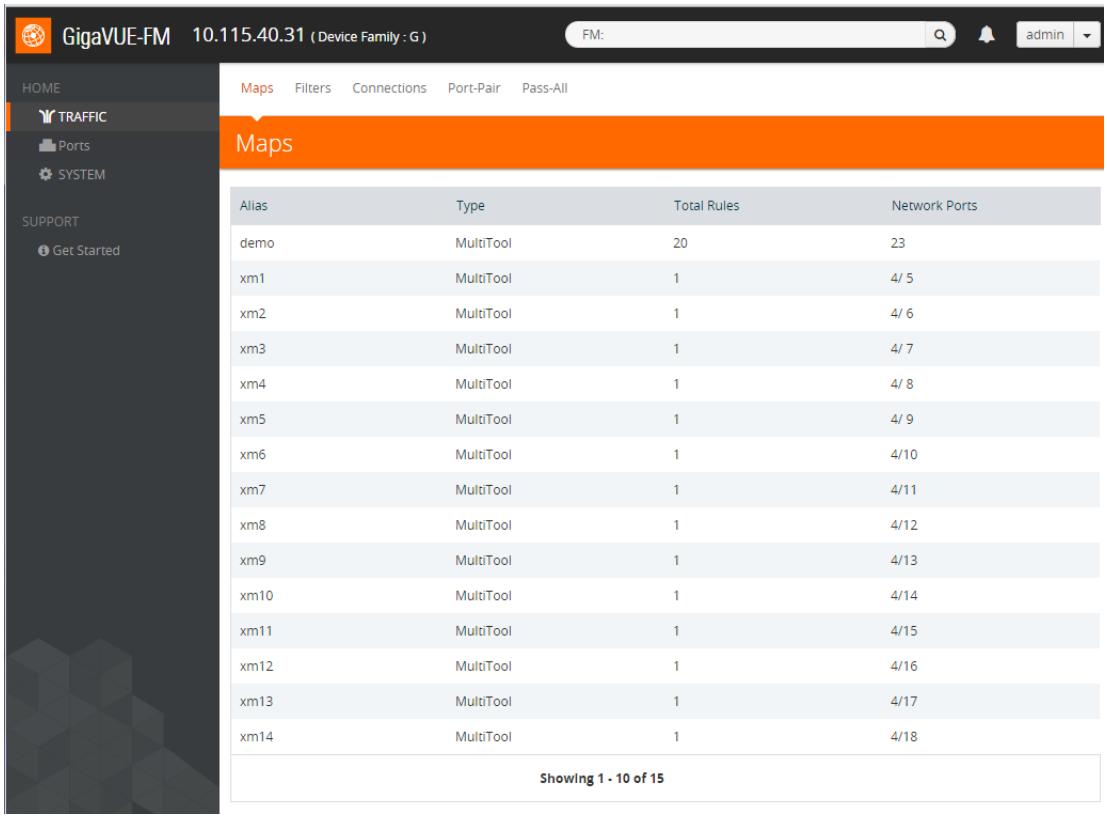

Figure 13-31: H-Series and G-Series Dashboards in GigaVUE-FM

Search for Ports on a GigaVUE Node

When viewing a node from GigaVUE-FM, you can use Port List Filter feature to display only certain ports that match specified criteria, such as only those ports that are used in a map.

To use the port list filter:

1. Select a physical node from the **Physical Nodes** page and then select the node to view.
2. Go to **Ports > Ports > All Ports**.
3. To filter the ports, click **Filter**. The Filter quick view is displayed.
4. Specify the criteria of the ports you want to filter.

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

Criteria	Description
Box/Slot ID	Display only those ports that match the specified box and slot IDs.
Port Alias	Display port with the specified alias.
Port ID	Display ports with specified number in the port ID. For example, if you specify 3 the result will also display ports that include the number 3, 13, 23, 30, and so on.
Type	Display ports with the specified port type. Select one of the following: • Network • Tool • Inline Network • Inline Tool • GigaSMART • Hybrid • Stack
Port Used in Map(s)	Display ports based on their usage in maps. The possible selections are: • All — display all ports either unused or in use by maps. This is the default. • In Use — display ports that are in use by any map. • Unused — display ports that are not use by any maps.
Admin Status	Display ports based on their current admin status. The possible selections are: • All — display ports with a status of Enabled or Disabled. This is the default. • Enabled — display ports with admin enabled. • Disabled — display ports with admin disabled.
Speed	Display ports with the selected port speed. The port speeds available depend on the node.
Transceiver Type	Display ports with the selected transceiver type. The transceivers available selection depend on the type of transceivers connected to the ports.

Figure 13-32 shows an example where filter criteria selected are Network Type and Admin Status Enabled. To remove the filter selections, click **Clear**.

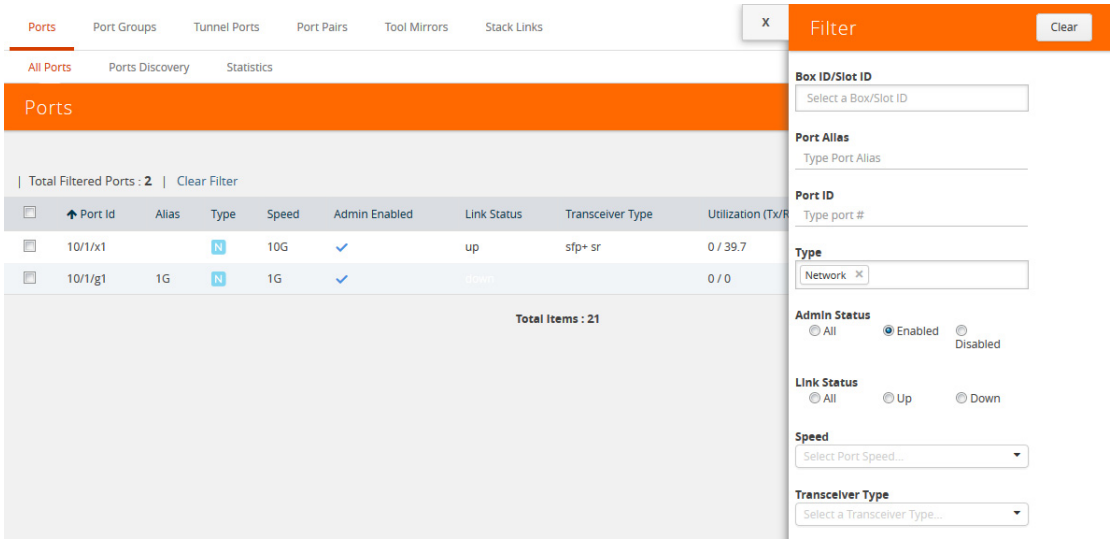

Figure 13-32: Filtering by Network Port Type and Admin Status Enabled

After the filter is applied, the Ports page displays only the ports that correspond to the selected filters and shows the total number of ports that meet the criteria. To clear the filters, select **Clear Filter**. Figure 13-33 shows the Port pages with two ports that correspond to the current filters: Network Type and Admin Status Enabled.

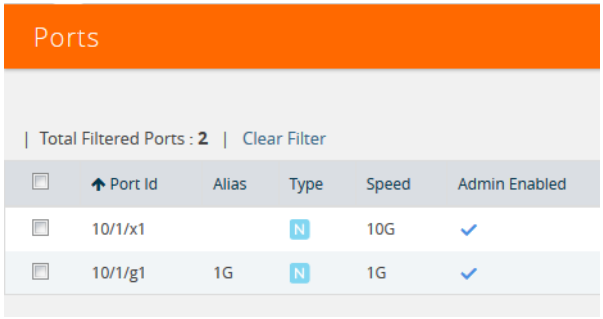

Figure 13-33: Filtered Ports List

Overview Page

The **Overview** page displays general information on the specific H Series node, which includes System, Ports, Maps, and Traffic information.

Systems Information

Systems information is displayed on the System widgets.

The Systems widget displays general information about the specific device that you selected from the drop-down list at the top of the widget. If the system is a cluster, you can select a device in the cluster to display on the widget.This widget gives you a quick status if any issues are present in any of the device's components through color indicators; green (running), amber (warning), or red (alert).

NOTE: Ensure that all the nodes and clusters have a Box ID defined. If the Box ID is missing, the Systems widget may not display any information relating to the node.

NOTE: Red alert appears for cards not present.

Field	Description
Host Name	The host name of the box.
Hardware	The hardware type, (for example, GigaVUE-HD8 or GigaVUE-TA1).
Software	The version of the software running on the device.
Memory	Shows the amount of used and free memory.
Load Average	The average load on the system over the last 1 minute, 5 minutes, and 15 minutes.
Cards	Displays all slots for the specific hardware type including its slot number and the type of card it contains or not. **Note:** When you hover over the card slot, the **temperature** is displayed.
Fan Trays	Indicates that the Fans are On or Off.
Power Supply	Indicates that the power supply is On or Absent. **NOTE:** When one or more power supply units are down, red alert is displayed.

Failure to Authenticate

To view a physical node in the Dashboard System pane, your login credentials must have the appropriate permissions. Otherwise, GigaVUE-FM shows an error message.

There are two possibilities that caused a user authentication error:

- The login user credential for GigaVUE-FM is not "**admin**".
- The password associated with the login user name for GigaVUE-FM is different for the physical device.

NOTE: When a “non-admin” user goes to Physical Nodes page, they are able to view displayed physical nodes with the status as Green, Amber or Red, this is because the physical node information is captured using the default “admin” user role.

Ports Information

The Overview page displays widgets that provide port information for the number of ports down, the number of ports with packet drops, and the number of over utilized receiving (Rx) ports and transmitting (Tx) ports. The ports widgets default to displaying a counter. Clicking on the icon in the upper right-hand corner displays the information as a table. The Ports with Packet drops and Over-Utilized Ports widgets are similar.

Traffic

The Traffic widget the shows most-utilized ports and ordered by traffic count. Each displayed port is labeled with its location, whether it is a transmitting or receiving port, and its percentage of utilization.

NOTE: The Traffic pane is view-only. It reflects traffic activity with port ID at the time of discovery and does not immediately refresh.

Workflows

The Workflows page provides wizards for creating maps. These wizard step you through the work flow to make sure you configure all of the components necessary for configuring out-of-band and inline maps for traffic flow monitoring. The workflows keep track of each step so that you can stop and then return to were you left off in the workflow. However, you can only work on one workflow at a time.

Overview of Workflows

Table 13-4 describes the maps that you can create with the wizards.

Table 13-4: Map Types and Map Wizards

Map Type	Map Wizard	Description
Out-of-Band maps	Map with rules	Walks you through the steps to select source and destinations ports, then create a Regular By Rule map with those ports.
	Pass-all map	Walks you through the steps to select source and destinations ports, then create a Pass All map with those ports.
	Collector map	Walks you through the steps to select source and destinations ports, then create a Collector map with those ports.

Table 13-4: Map Types and Map Wizards

Map Type	Map Wizard	Description
Inline maps	Map with rules	Walks you through the steps to select the destination and source, then use them in an Inline By Rule map. If the Inline Networks or Inline Network Groups have not been created yet, you can create them from the wizard. The wizard selects Symmetric for Traffic Type.
	Pass-all map	Walks you through the steps to select the destination and source port, then use them in an Inline Pass All map. If the Inline Networks or Inline Network Groups have not been created yet, you can create them from the wizard. The wizard selects Symmetric for Traffic Type.
	Collector map	Walks you through the steps to select the destination and source ports, then use them in an Inline By Rule map. If the Inline Networks or Inline Network Groups have not been created yet, you can create them from the wizard. The wizard selects Symmetric for Traffic Type.
	Asymmetric Inline map	Walks you through the steps to select the destination and source ports, then use them in an Inline Pass All map. The source port must be an Inline Network port. The destination port must be an Inline Tool, an Inline Tool Group, or an Inline Serial Tool Group. If the Inline Networks or Inline Network Groups have not been created yet, you can create them from the wizard. The wizard selects Asymmetric for Traffic Type.
Basic out-of-band GigaSMART maps:	Map with GigaSMART Apps	Walks you through the steps to select a GigaSMART Group, GigaSMART operation, select the source and destination ports, and then create a Regular By Rule map. The wizard allows you to create the GigaSMART Group and GigaSMART operation if they do not already exist.
	First level map	Walks you through the steps to select or created the source port and select a or create a virtual port for the destination port, and then create a First Level By Rule map.
	Second level map	Walks you through the steps to select or created the virtual port for the source port and select a or create the port for the destination port, and then create a Second Level By Rule map.
Advanced out-of-band GigaSMART maps	Map with NetFlow	Walks you through the steps to select or create an IP interface; a NetFlow exporter, records, and monitor; a GigaSMART Group and Operation, source ports; and then created a Regular by Rule map.
	SSL-based map	Walks you through the steps to create or select a GigaSMART Group, configure SSL, create or select a GigaSMART Operation, and then create a Regular By Rule map.
	Map with ASF	Walks you through the steps to create or select a GigaSMART Group, virtual port, and GigaSMART Operation; configure GigaSMART Application Session Filter; and create the First and Second Level maps needed for implementing an ASF solution.

How to Use Workflows

To start a workflow, click on one of the links, such as **Map with Netflow**. When you click on the link, the Workflow page is displayed. The Workflow panel on the right shows the tasks in the workflow, indicating the current task.

While using workflows, you can only work on one workflow at a time. Also, you cannot roll back changes made to a node after canceling a workflow.

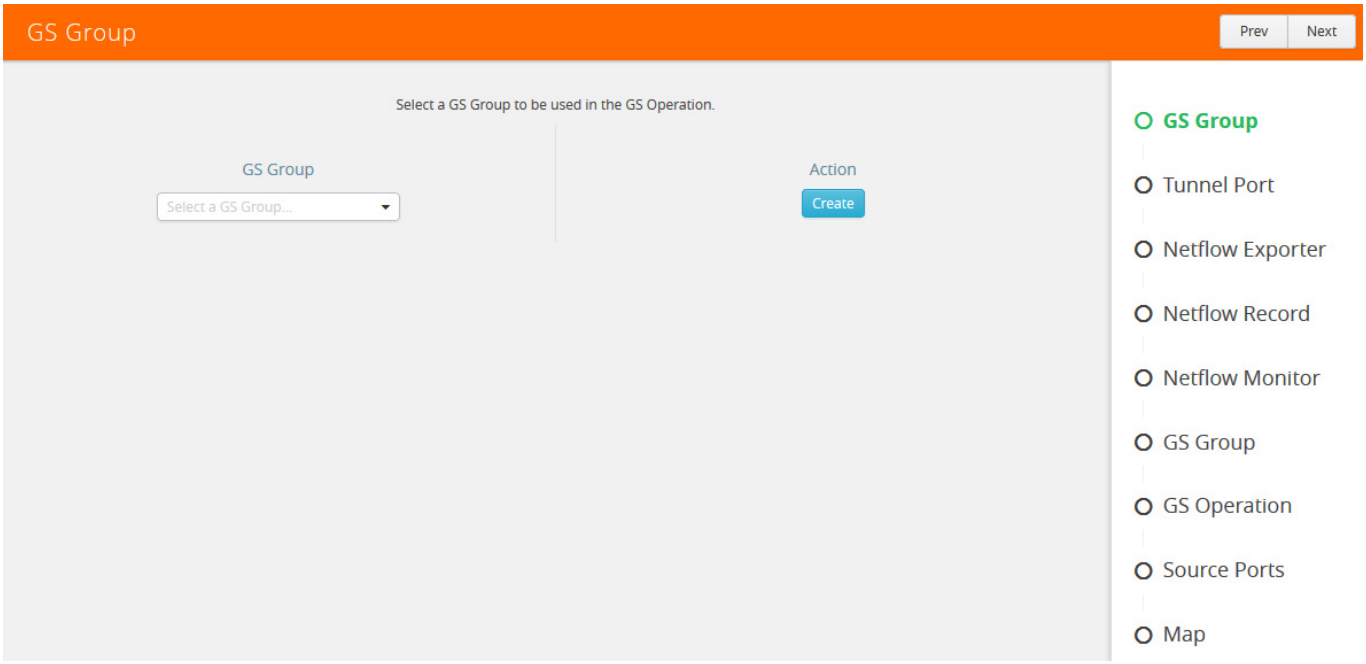

Figure 13-34: Create GS Group in Workflow

A task in the workflow allows you to select an item or create the item if one does not exist. For example, a GigaSMART Group needs to be selected if one does not exist. In this case, click **Create**. Clicking Create takes you to the GS Group configuration page. Configure the GigaSMART Group and click **Save**.

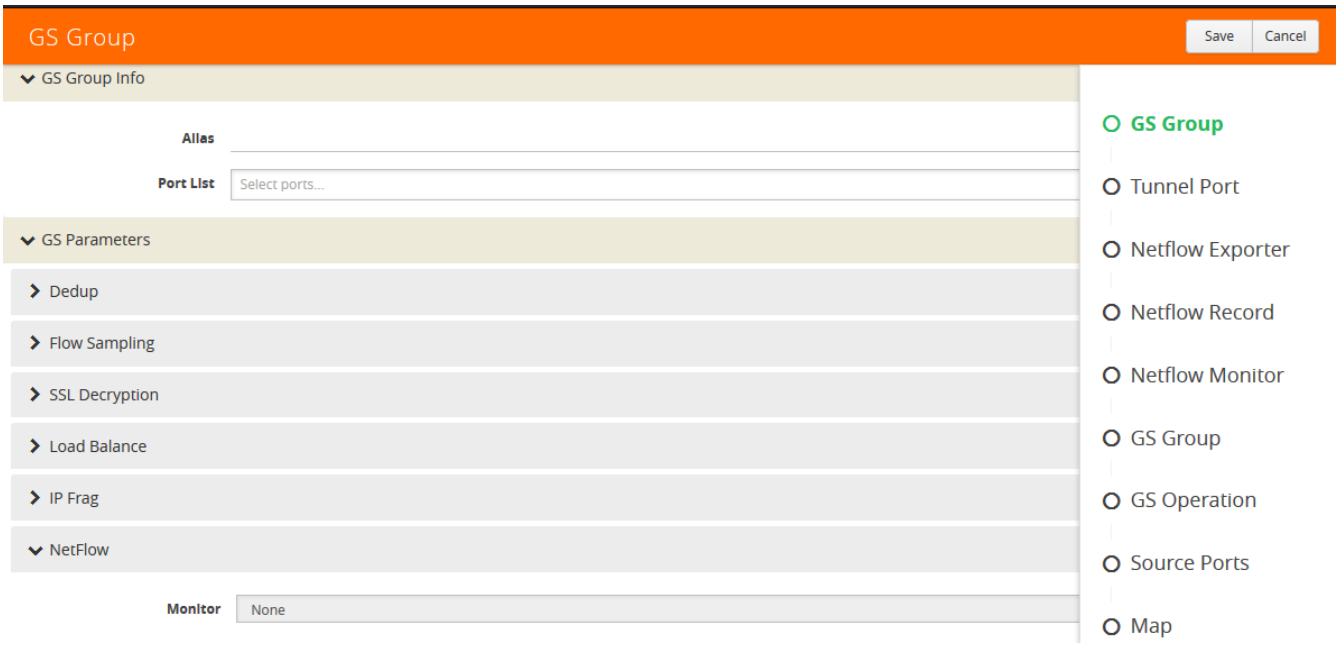

Figure 13-35: GS Group Configuration in Workflow

After saving the configuration, the Workflow moves to the next task and the Workflow panel indicates which tasks have been completed as shown in Figure 13-36.

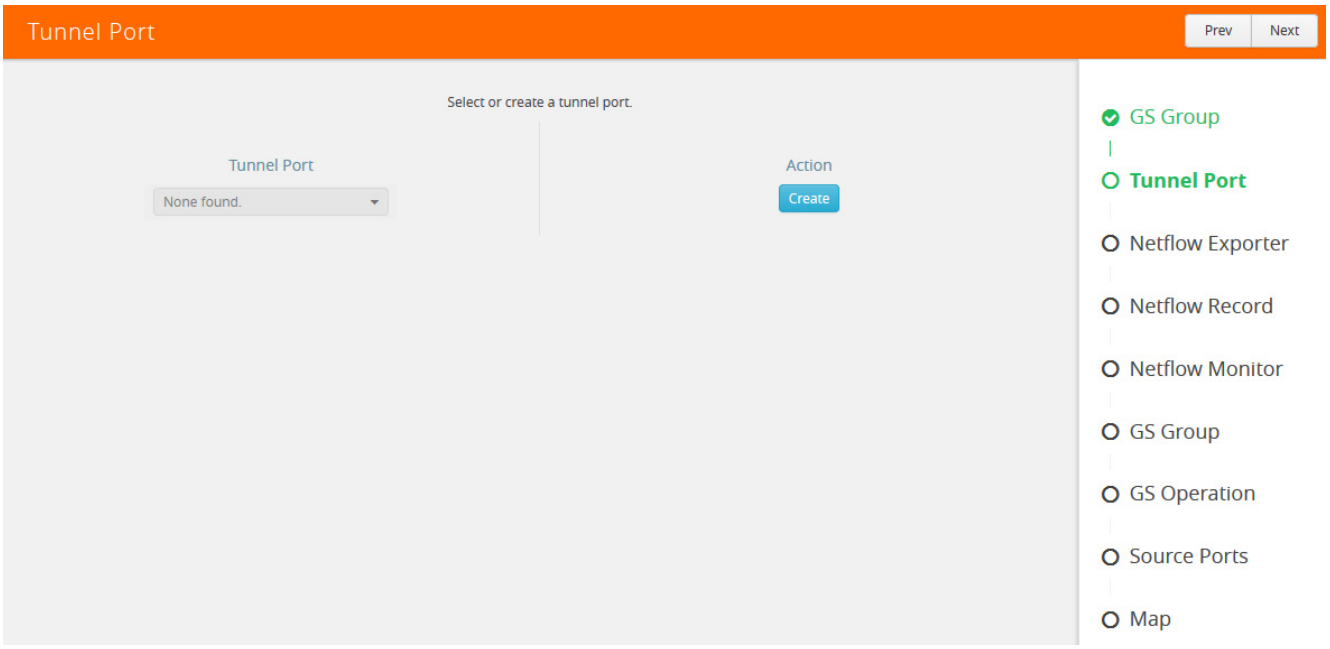

Figure 13-36: GS Group Workflow Task Completed

You can select a task in a different order than shown in the Workflow panel. For example, you can go to the NetFlow Monitor task. After completing the task, the Workflow returns you to the to the GS Group configuration page with the Monitor field completed as shown in Figure 13-37 on page 244.

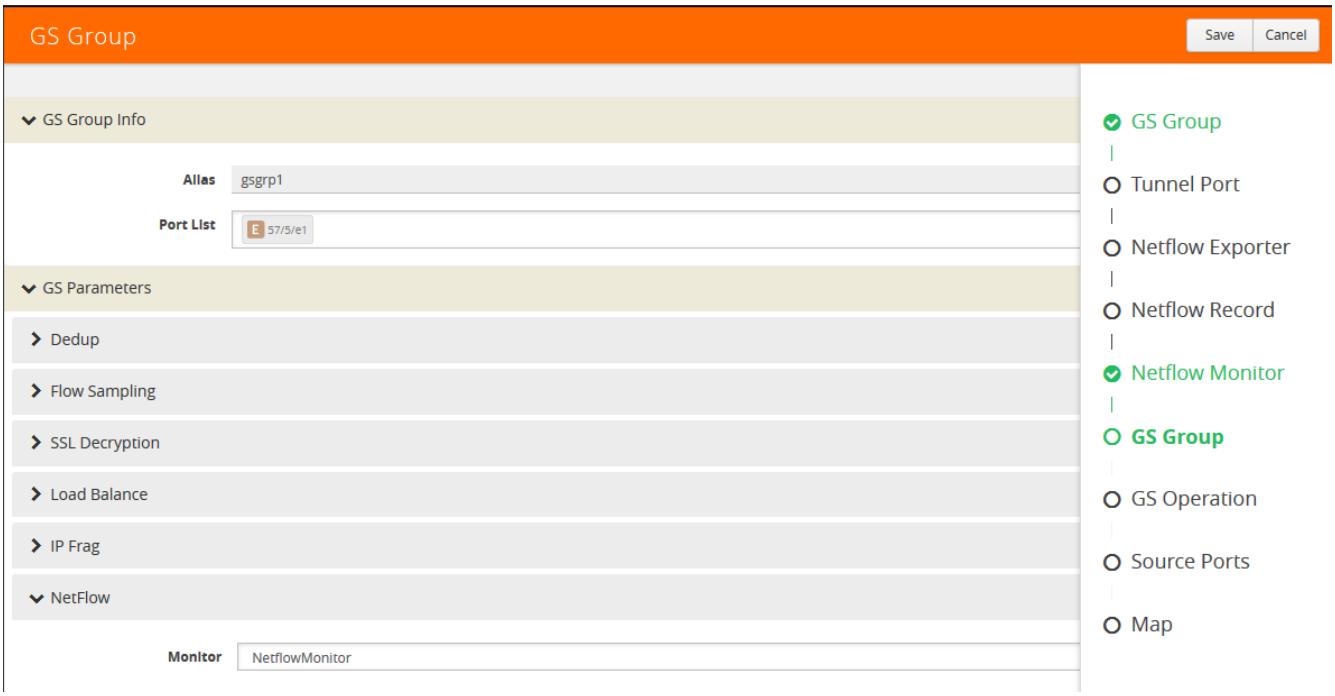

Figure 13-37: NetFlow Monitor Task Completed

Workflow allows you to leave the current workflow and return at anytime during a GigaVUE-FM session. The **In Progress** panel indicates the current workflow and the Workflow panel indicates the competed tasks. Figure 13-38 shows an example of workflow in progress. You can abandon a workflow by clicking the red x in the **In Progress** panel.

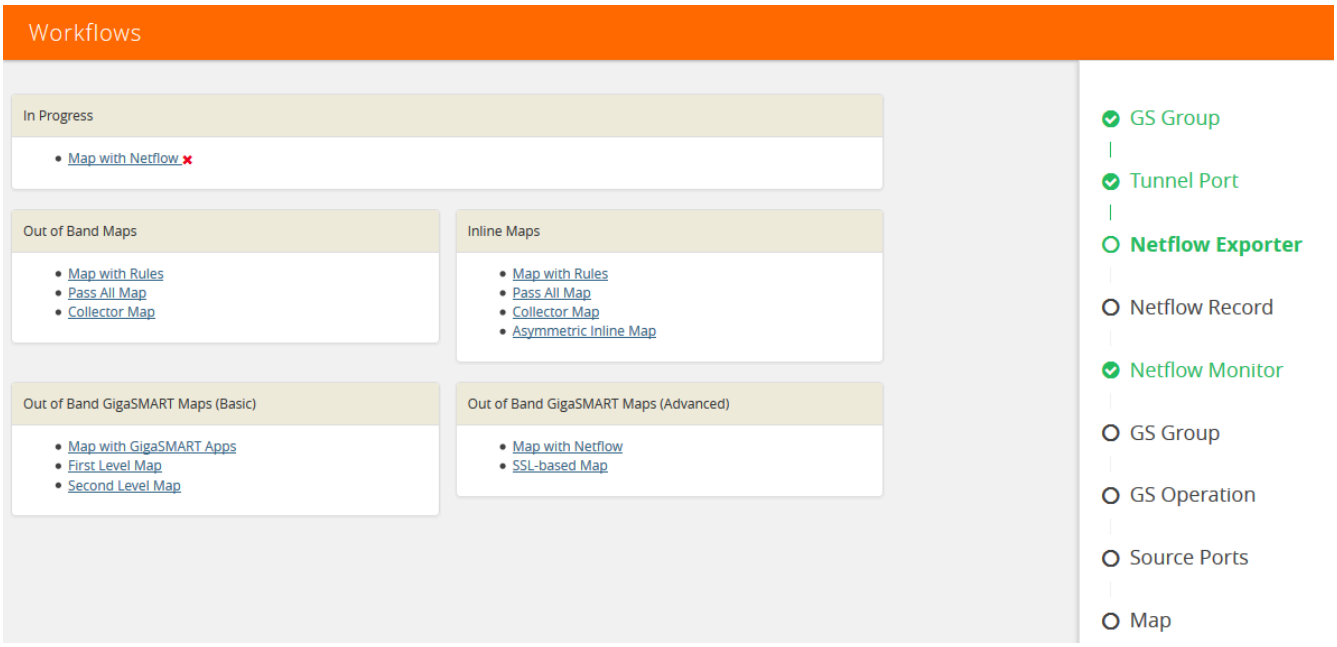

Figure 13-38: Workflow in Progress

When a workflow is completed, a page displays providing you with several choices for the next task. For example, when the Map with Rules workflow is completed (refer to Figure 13-39), you can go to creating a collect map by clicking the **Create a Collector Map**, return to the Workflow page by clicking **To Work Flows**, or go to the Maps page by clicking **To Maps**.

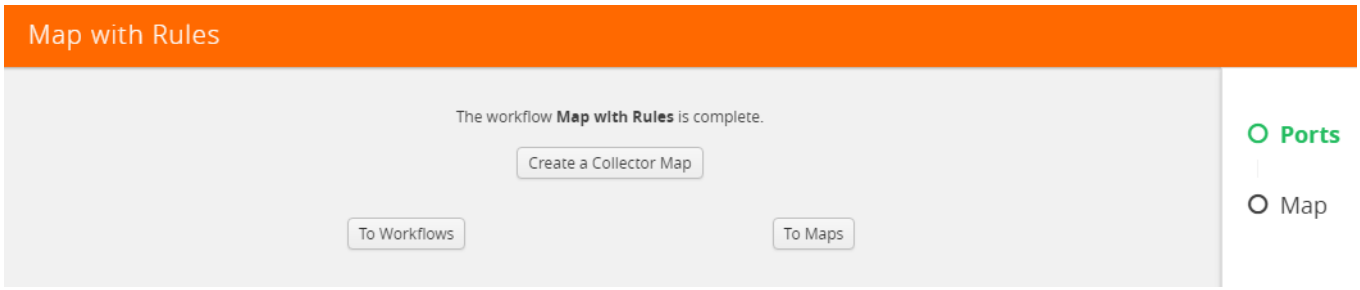

Figure 13-39: Map with Rules Workflow Completed

Chassis Table View

When viewing the Chassis Table View managed from GigaVUE-FM, the Table View includes environment information about the chassis in addition to the other information available when using GigaVUE-HVUE on a device running Gigamon-OS 4.6 or later.

The Table View provides the following information about the chassis and its components:

Chassis Information	Description
Properties	Provides information about the chassis.
Cards	Describes the cards installed in each slot of the chassis, including its current status.
Environment	Provides temperature and voltage information about the chassis.
Power Supplies	Describes the power supplies installed in the chassis, including their current status.
Fan Trays	Describes the fan trays installed in the chassis, including their current status.
Fan RPM	Provides the current RMP of the each fan.

Live Graphing

When viewing ports on a node running GigaVUE-OS 4.6 or later from GigaVUE-FM, you can select to see the graph data display in real time by clicking **Live**. This changes the updating of the information in the graph from the default to every 10 seconds. You can also select the data to display on the graph by selecting an option from the Select Counter list. Figure 13-40 shows an example, where Live is selected and the Data Rate Rx and Packets Rx counters are selected.

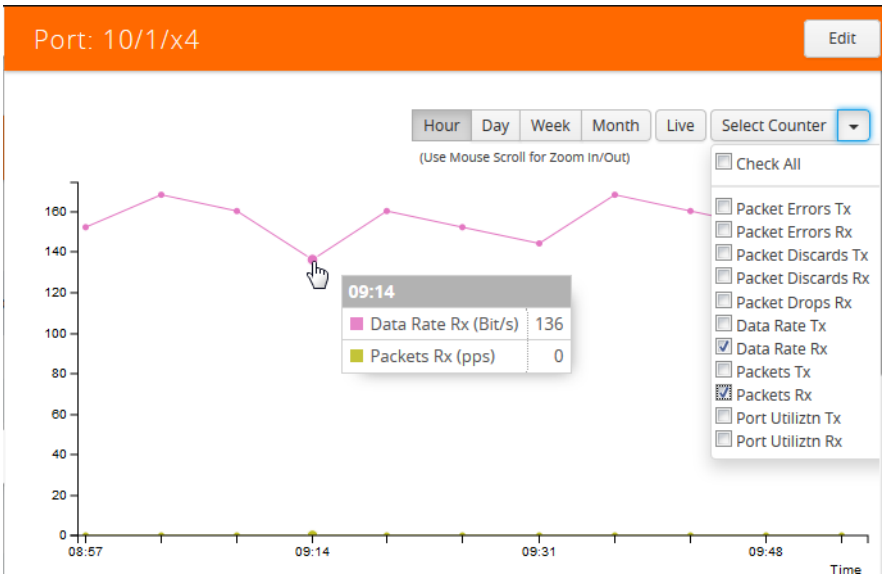

Figure 13-40: Live Graphing and Data Counters Selected

Safe and Limited Modes

Starting in GigaVUE-OS software version 4.7, safe and limited modes are introduced to the cluster environment and standalone nodes.

During clustering operations, in rare scenarios, there can be unrecoverable system errors that can potentially put the cluster or the clustered nodes into unsafe or unstable states. Once in such a state, additional operations or configuration changes can cause the node to crash, the cluster to deform, and the data traffic to be impacted. For example, due to a node attempting to rejoin a cluster, a chassis can end up in a reboot loop. In previous software versions, there was no way to prevent entering the loop.

These modes provide notification, stop further operations from being performed, and give you time to troubleshoot and plan the recovery of the cluster or of any node in the cluster or standalone node.

Two modes are supported. The first is called safe mode and is triggered when the node detects unrecoverable errors, but the existing flow maps are not impacted. The second is called limited mode and is triggered when the node detects continuous system reboots. In this mode, the node will become standalone and only basic configuration will be allowed.

When a cluster is in safe mode, GigaVUE-FM displays a Safe Mode banner and message.

Safe Mode

A node enters safe mode when there are unrecoverable errors. Any node in a cluster can enter this mode.

Examples of unrecoverable errors are when there are inconsistencies between the system and the running configuration or when the cluster configuration did not merge properly with the existing configuration. A node will automatically enter safe mode.

As part of merge error recovery, nodes joining a cluster are automatically restarted so the merge error can be fixed.

When a node is in safe mode:

- The node displays a banner indicating it is in safe mode.
- An SNMP trap is sent to notify the user when the mode changes
- Traffic provisioning is not allowed on the affected node. Any other configuration remains as is.
- Configured traffic continues to be forwarded.
- If the standby node in the cluster is in safe mode, it can still become the master if the current master fails or switches over, but the database on the standby node may not be in sync, so it is not recommended to continue in that state. Instead, take immediate action to recover the node.
- In safe mode, the node does not process any incoming traffic configuration from the cluster master.

When a node is in safe mode and you try do any operations that are not allowed in safe mode, the UI displays the following message:

```
The system has restricted provisioning in safe mode. Contact Gigamon
Support on how to troubleshoot and recover from safe mode.
```

Also, hovering over the status bubble of the nodes on the Physical Nodes page in GigaVUE-FM displays a message that the node is in Safe Mode.

To exit safe mode, reload the node.

Limited Mode

A node automatically enters limited mode when it detects repeated system crashes.

When a node is in limited mode:

- The node displays a banner indicating that it is in limited mode.
- An SNMP trap is sent to notify the user when the mode changes.
- Only basic system provisioning is allowed. Traffic provisioning is not allowed. Only commands that are related to image download, installation, next boot, and reboot are allowed.

Limited mode is triggered when there are three (3) failures/system crashes within 15 minutes. In limited mode, the cluster configuration is ignored. No cluster configuration or GigaVUE-OS configuration is accepted when the node is in limited mode.

When a node is in limited mode, a Limited Mode banner displays in GigaVUE-FM.

Enable SNMP Trap for Safe Mode and Limited Mode

Use the following steps to configure a notification that will be sent to all configured destinations when a node in the cluster changes from operational mode to safe mode or from operational mode to limited mode

The safe mode and limited mode capabilities are enabled through the SNMP trap event Operational Mode Change. To enable the trap on a node, do the following:

1. Click **Physical** on the top navigation link. On the Physical Nodes page, select the node you want to configure.
2. Select **Settings > Global Settings > SNMP Traps**.
3. Click **Trap Settings**.
4. On the Edit SNMP Traps Settings page, select **Operational Mode Change**.
5. Click **Save**.

When the cluster master enters safe mode, the SNMP trap will be sent and the master will be identified as the local node in the trap.

When a node in a cluster enters safe mode, the SNMP trap will be sent and the node will be identified as the local node in the trap. In addition, a notification will be sent to

the cluster master. The node that entered safe mode will be identified by its box ID in the notification to the master.

Log messages also provide information. The following is a sample log:

```
Jun  8 13:46:27 GC-TA10-N6 mgmtd[2400]: [mgmtd.INFO]: SAFE mode: Merge
error detected !! Triggering SAFE mode ...
```

Collect Information for Technical Support

Collecting the following information can help Technical Support:

- sysdumps/debug dumps for all nodes in the cluster
- sysdumps for nodes that observed a crash entering safe or limited mode
- debug dumps for nodes that did not observe a crash
- gslogs for application information
- console logs
- CLI histories
- CLU, H-VUE, or FM screen captures
- SNMP captures

To contact technical support, refer to *Contacting Technical Support* on page 1370.

14 Multi-Path Leaf and Spine

This chapter describes the leaf and spine architecture with multiple paths for achieving high availability in a cluster environment. Refer to the following sections for details:

- *Introduction to Multi-Path Leaf and Spine* on page 252
- *Configuration Overview* on page 256
- *Leaf-Spine Cluster Deployment* on page 258

NOTE: Refer to *Regular Cluster Formation Workflow* on page 212 for how to use the Regular Cluster workflow.

Introduction to Multi-Path Leaf and Spine

The leaf and spine architecture is a two-layer architecture used for network aggregation. There are two kinds of nodes in this architecture, as follows:

- leaf nodes, which are edge nodes and can also have TAPs or tools attached to them
- spine nodes, which are the nodes to which the leaf nodes attach

With multiple paths between the nodes in a cluster, the leaf and spine architecture protects against failures, such as stack link or spine node failures. In the event of a failure, the traffic on one path fails over to the other path.

In this architecture, each leaf node connects to every spine node. This forms a mesh between the leaf and spine nodes. However, no leaf node directly connects to another leaf node and no spine node directly connects to another spine node. An example of a spine and leaf architecture is shown in Figure 14-1 on page 252.

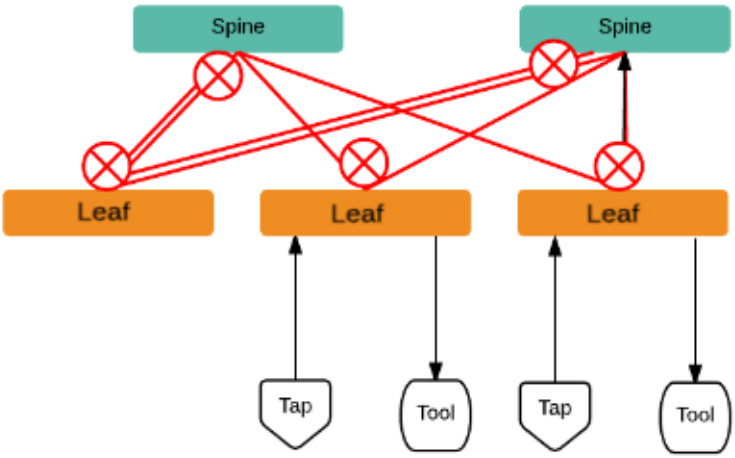

Figure 14-1: Leaf and Spine Architecture

In a cluster, the number of leaf nodes is typically greater than the number of spine nodes. In Figure 14-1, there are three leaf nodes and two spine nodes. The leaf nodes aggregate to a fewer number of spine nodes.

In Figure 14-1, the spine nodes are GigaVUE TA Series nodes, such as GigaVUE-TA100, while the leaf nodes are GigaVUE H Series nodes, such as GigaVUE-HC2, GigaVUE-HC3, or GigaVUE-HD8, which places the traffic intelligence at the edge.

Figure 14-1 shows TAPs or tools connecting to the leaf nodes, and the leaf nodes connecting to the spine nodes. Note that TAPs or tools do not connect to the spine nodes.

Instead of one leaf node connecting to one spine node with a single link, in this architecture there are multiple links from the leaf nodes to the spine nodes. The leaf nodes connect to the spine nodes through at least two paths. Some leaf nodes with higher capacity, such as GigaVUE-TA100, can have more paths, as shown by double red lines in Figure 14-1.

Traffic between ports on a leaf node will be local to that leaf node, but traffic between leaf nodes will go through the spine nodes.

The traffic from a source leaf node to a destination leaf node flows as follows:

- From a TAP, traffic flows to the source leaf node
- From the source leaf node, traffic is load balanced to all spine nodes
- From a spine node, traffic flows to the destination leaf node
- From the destination leaf node, traffic flows to tool ports

Resiliency is achieved when there are multiple paths from the network to the tools across GigaVUE nodes.

Path Protection

The spine leaf architecture in the cluster environment provides failover for the following:

- leaf node failure. Refer to *Leaf Node Failure* on page 253.
- stack link failure on a leaf node (not connected to a tool, but can be connected to a network TAP). Refer to *Stack Link Failure on Leaf (TAP Connected)* on page 254.
- spine node failure. Refer to *Spine Node Failure* on page 254.
- stack link failure on a leaf node (connected to a tool). Refer to *Stack Link Failure on Leaf (Tool Connected)* on page 255.

Leaf Node Failure

Refer to Figure 14-2 on page 253 for a failure in which a leaf node is powered down or rebooted. The leaf node does not have a connected TAP or tool.

NOTE: In the following figures, red arrows indicate traffic direction.

With this type of failure, the stack links connected to the affected leaf node go down, which will be detected by the spine nodes. No action will be required at the spine nodes. Since the links are down, no traffic will be sent to the affected leaf node.

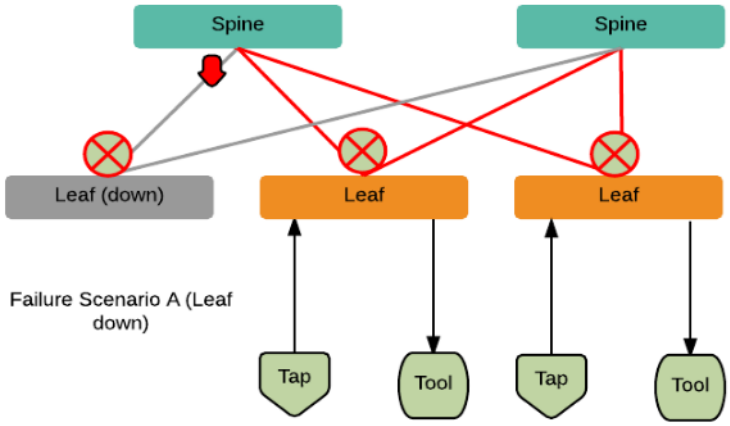

Figure 14-2: Leaf Node Failure

Restoration

Once the leaf node is powered up and booted, it will restore its traffic configuration.

Affected Time

None. Traffic on other leaf nodes will not be affected.

Stack Link Failure on Leaf (TAP Connected)

Refer to Figure 14-3 on page 254 for a failure in which a stack link on a leaf node fails and the leaf node is connected only to a TAP.

With this type of failure, the stack link between the leaf and spine nodes goes down. No action will be required at the spine node. At the leaf node, the affected link will be removed from the stack GigaStream. Traffic will be sent to the other spine node.

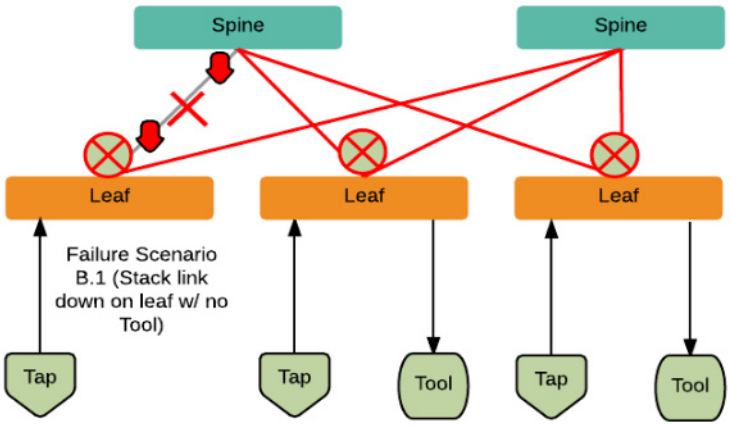

Figure 14-3: Stack Link Failure on Leaf (TAP Connected)

Restoration

When the link comes back up, the leaf node will put the link back into the GigaStream.

Affected Time

When the link is down, traffic recovers in a similar amount of time as a tool GigaStream.

Spine Node Failure

Refer to Figure 14-4 on page 255 for a failure in which a spine node is powered down or rebooted.

With this type of failure, the stack link between the leaf and spine nodes goes down. The leaf nodes will detect that the stack link is down. The affected link will be removed from the stack GigaStream. Traffic will be load balanced to the other spine node.

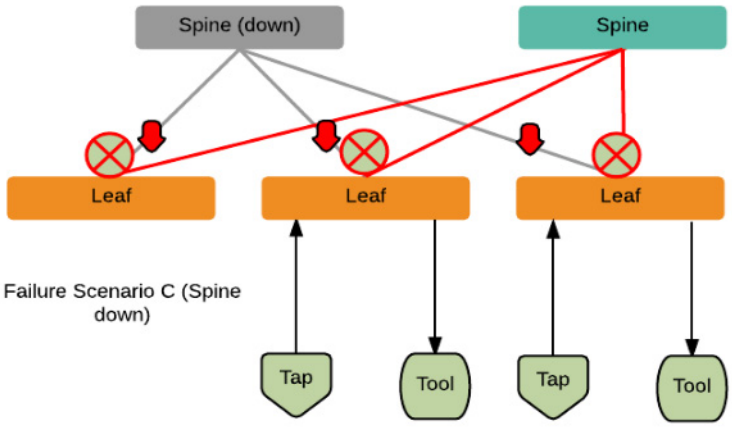

Figure 14-4: Spine Node Failure

Restoration

When the spine node reboots, the cluster will synchronize. When the node converges to the cluster and the configuration synchronizes, traffic will be restored.

Affected Time

When the spine node is powered down or rebooted, traffic recovers in a similar amount of time as a tool GigaStream.

Stack Link Failure on Leaf (Tool Connected)

Refer to Figure 14-5 on page 255 for a failure in which a stack link between the leaf and spine nodes fails and the leaf node is connected to a tool.

In the current software version, this type of failure is not supported.

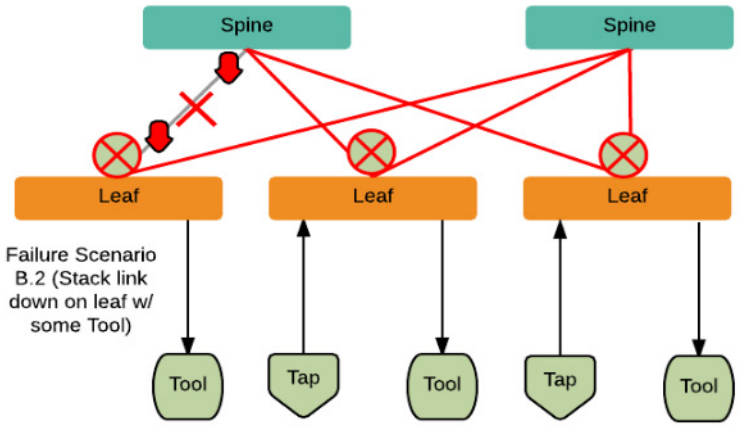

Figure 14-5: Stack Link Failure on Leaf (Tool Connected)

Configuration Overview

This section provides an overview of the configuration. The configuration is done from the master node in the cluster. Follow this configuration sequence to prevent loops.

This configuration connects nodes using multiple paths. For an example of the configuration, refer to Figure 14-6 on page 256.

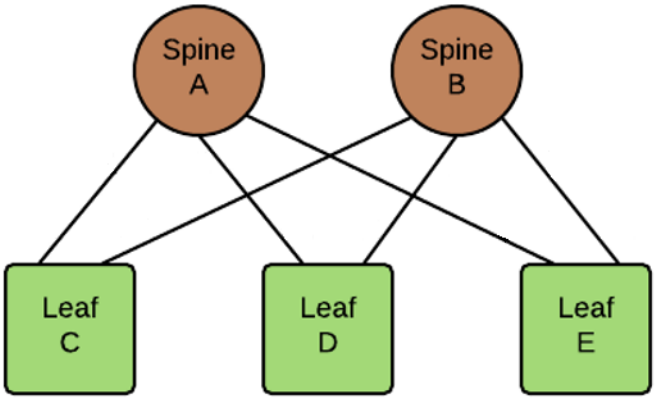

Figure 14-6: Leaf and Spine Configuration

The configuration steps are as follows:

- Configure stack GigaStream. The stack GigaStream connect the spine and leaf nodes. In Figure 14-7 on page 256, the stack GigaStream are: a1, a2, a3, b1, b2, b3, c1, c2, d1, d2, e1, e2. Even if there is only one port that connects the nodes, you must still configure a stack GigaStream. With a configuration of two spine nodes and three leaf nodes, the number of stack GigaStream is 12.

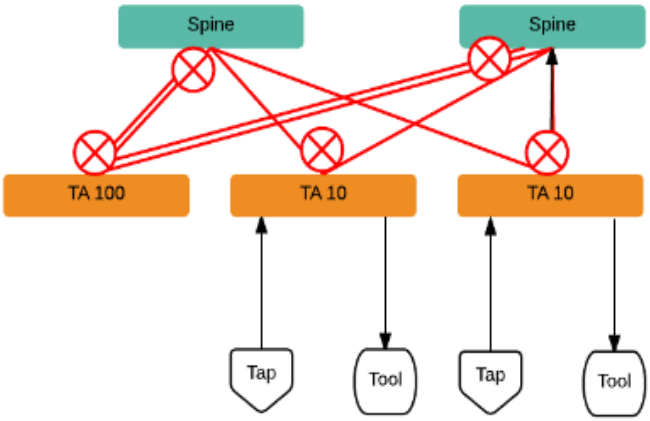

Figure 14-7: Stack GigaStream Configuration

- Configure spine links. On each leaf node, there is one spine link that contains the list of GigaStream connecting the leaf nodes to the spine nodes. The spine links contain multiple stack GigaStream that are bundled together. The spine links are: {c1,c2}, {d1,d2}, and {e1,e2}. The total number of spine links is three for this configuration. The spine links are located at the leaf nodes. Across the spine link members, traffic is load balanced. For this part of the configuration, refer to the circles in Figure 14-8 on page 257.

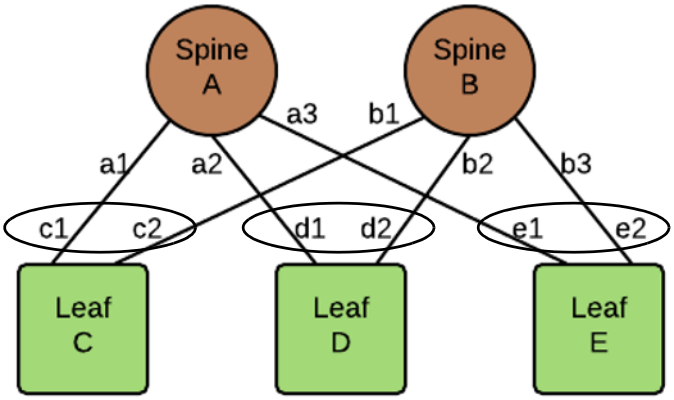

Figure 14-8: Spine Link Configuration

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

- Configure stack links. The stack links are: {a1,c1}, {a2,d1}, {a3,e1}, {b1,c2}, {b2, d2}, and {b3, e2}. The total number of stack links is six for this configuration. For this part of the configuration, refer to the circles in Figure 14-9 on page 257.

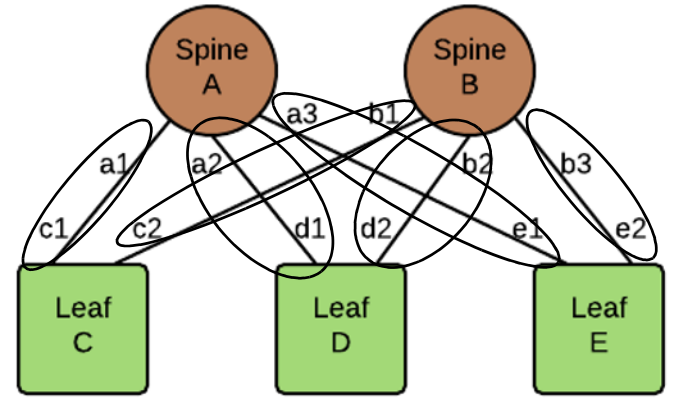

Figure 14-9: Stack Link Configuration

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

Notes and Considerations

Refer to the following notes and considerations:

- The multi-path leaf and spine architecture is only supported in an out-of-band cluster.
- The spine link GigaStream must be of type stack. Stack GigaStream carry bi-directional traffic.
- All spine link GigaStream must have the same parameters, such as the same hash value and failover mode.
- Once a spine link is configured, editing of GigaStream parameters is not supported, except for editing the comment.
- On GigaVUE HD Series nodes, the GigaStream must be on the same line card.
- GigaStream must be configured before spine links are configured.
- Once a GigaStream is configured in a spine link, it cannot be deleted. To delete a spine link, the stack links must first be deleted.
- A spine link cannot be deleted if a map is using the spine link.
- A spine link cannot be created if a map is using the GigaStream.

The number of spine and leaf nodes is not limited. The ratio of spine and leaf nodes are determined by the cluster traffic needed between the leaf nodes. Larger topologies have the same restrictions as the GigaVUE-OS as follows:

- the total number of nodes in a cluster, for example, 32
- the number of links in a GigaStream (which depends on the GigaVUE node and line card or module, for example, the PRT-HC0-X24 module on GigaVUE-HC2 can have 24 stack GigaStream)

Leaf-Spine Cluster Deployment

This section describes the steps and prerequisites to deploy a leaf-spine cluster.

Refer to *Introduction to Multi-Path Leaf and Spine* on page 252 for a conceptual overview of the leaf-spine architecture.

Deployment Checklist

Before forming a Leaf-Spine Cluster, it is strongly recommended that you get familiar with the relevant documentation and review the deployment checklist to prepare for deployment.

Pre-deployment checklist

- Gigamon Fabric Management must be upgraded to Gigamon 5.3.00 or later
- Gigamon device must be upgraded to GigaVUE-OS 5.2.00 or later
- Advanced Features License must be installed in TA devices

- Physical connection must be established to create stack links
- Devices must have GDP enabled and be physically connected to create links among devices from Gigamon.

IMPORTANT: Recommendation is to use TA devices as SPINE Nodes and other devices as LEAF Nodes.

Formation Scenario

The Leaf-Spine cluster can be formed with different combinations of devices with four Spine and six Leaf nodes as a 10-node cluster.

The following configuration creates a leaf spine cluster with two spines and three leafs.

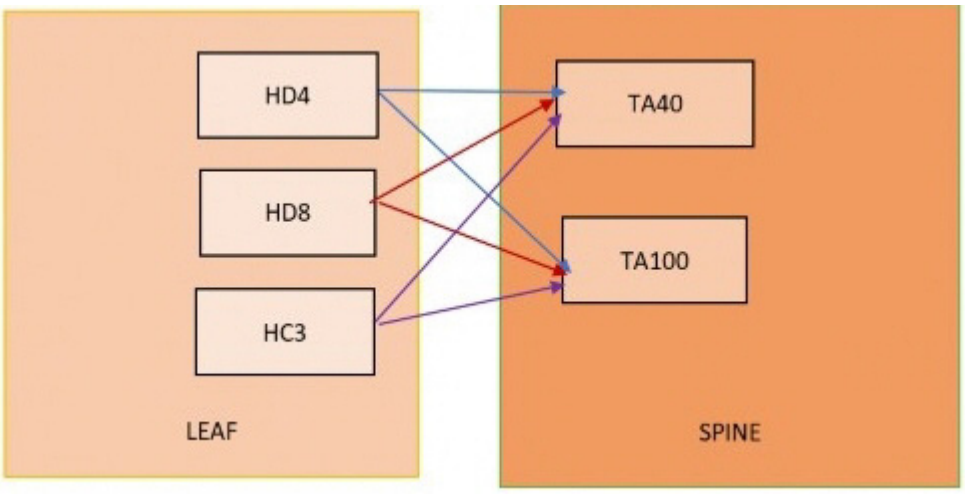

Figure 14-10: Leaf spine cluster overview,

NOTE: GigaStreams support different speeds, as indicated by the different colored connector lines in Figure 14-10 on page 259.

Leaf-Spine Cluster Formation Workflow

GigaVUE-FM 5.3 supports workflow-based configurations for forming a cluster. This workflow walks through the required steps to form a complete Leaf-Spine cluster. Additional procedures for editing and deleting cluster formations are also provided:

- *Create a Leaf-Spine Cluster* on page 259
- *Edit a Cluster* on page 265
- *Delete a Node from a Cluster* on page 268
- How to Change the Master Preference *of a Device* on page 270

Create a Leaf-Spine Cluster

To create a Leaf-Spine cluster:

1. Navigate to **Physical > Physical Nodes**.

2. Click **Create Cluster**.

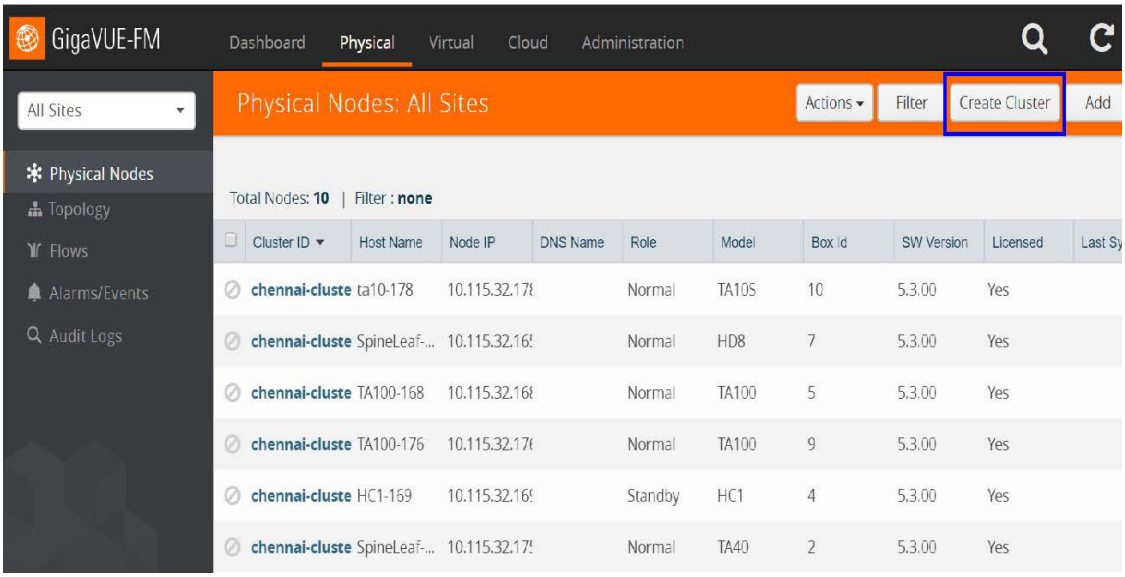

Figure 14-11: Create Cluster

Select the Cluster Type

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

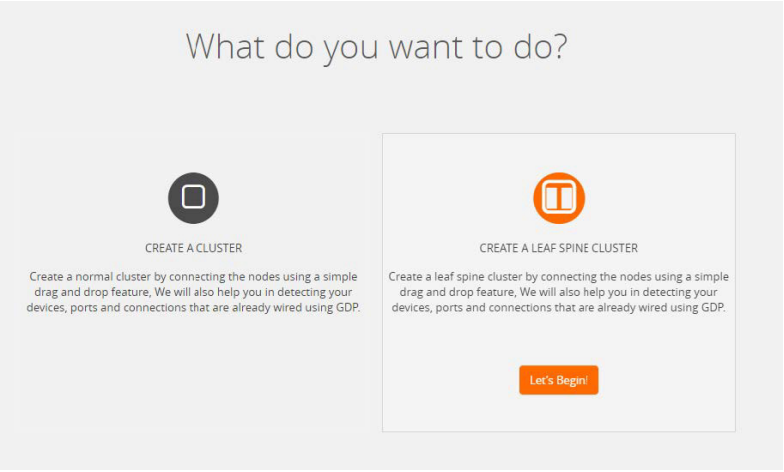

Figure 14-12: Choose CREATE A LEAF SPINE CLUSTER

Select Devices

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

5. The Select Devices page displays a list of standalone devices with filter options:

- *Software*: Filter the nodes based on the software version for which the cluster will be formed.
- *Model*: Filter the nodes based on a Gigamon model.
- *HostName*: Enter the HostName of the Gigamon Nodes to specify a device.

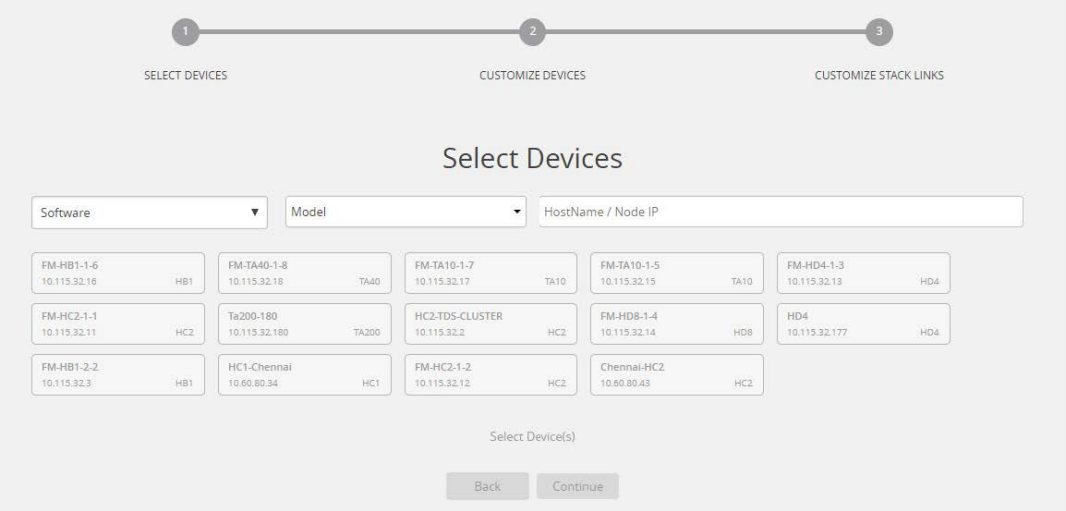

Figure 14-13: Select the required devices to form a cluster

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

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

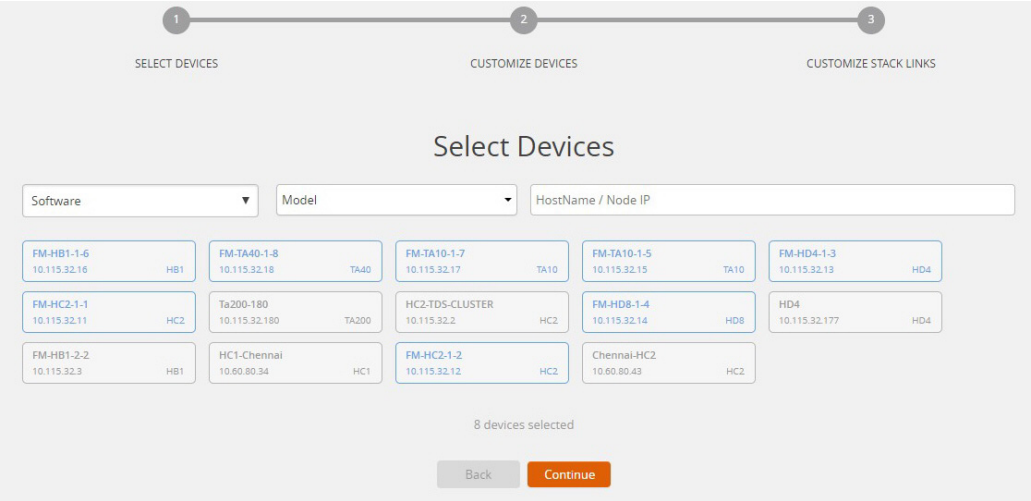

Figure 14-14: Selected devices will be highlighted

Customize Devices

Use the Cluster Configuration window to customize your devices.

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

NOTE: TA devices cannot be a master node.

Figure 14-15: Enter cluster configuration details

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

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

Customize Stack Links for your Leaf-Spine configuration

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

9. If GDP (Gigamon Discovery Protocol) is enabled at the device chassis level, then the corresponding ports used to create links and ports should be admin enabled. If a physical connection exists in the device, then the links will be shown.

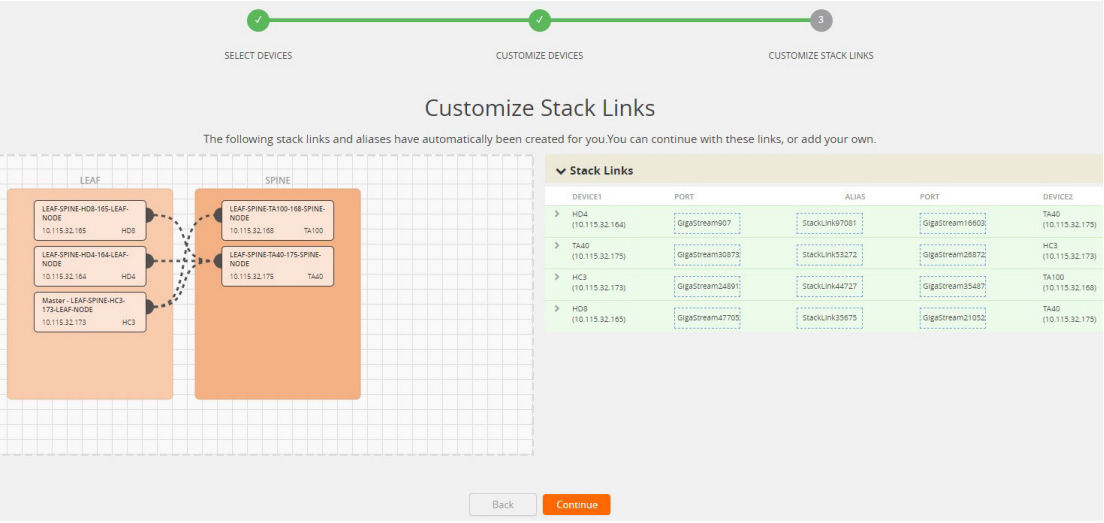

Figure 14-16: If GDP is enabled, the existing links appear in the cluster canvas

10. If GDP is not enabled in the ports, then the links will need to be drawn to connect the devices.

    Connect Leaf devices with Spine devices to create a stack link between them.

Click the tip of the node and drag your cursor to the second node tip to create a link. After you create the link, a dotted line will illustrate the connection.

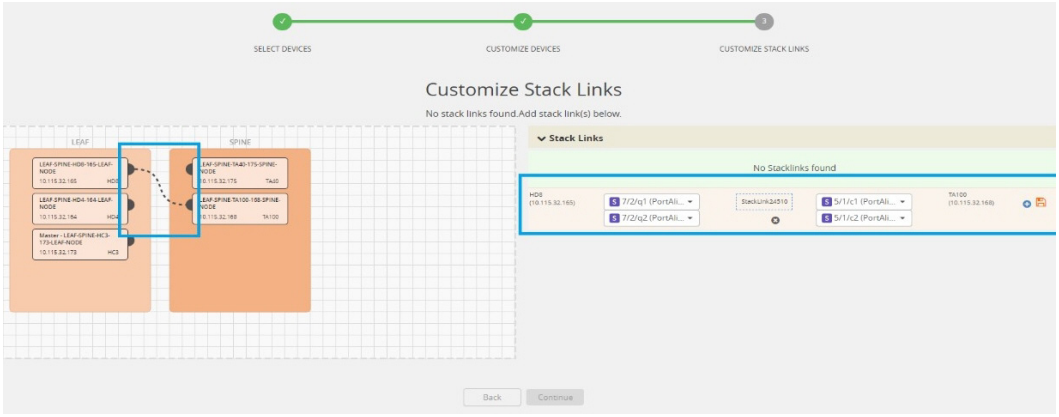

Figure 14-17: If GDP is not enabled, create stack links by connecting the links between leaf and spine

11. Configure the formed links in the Stack Links table as follows:

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

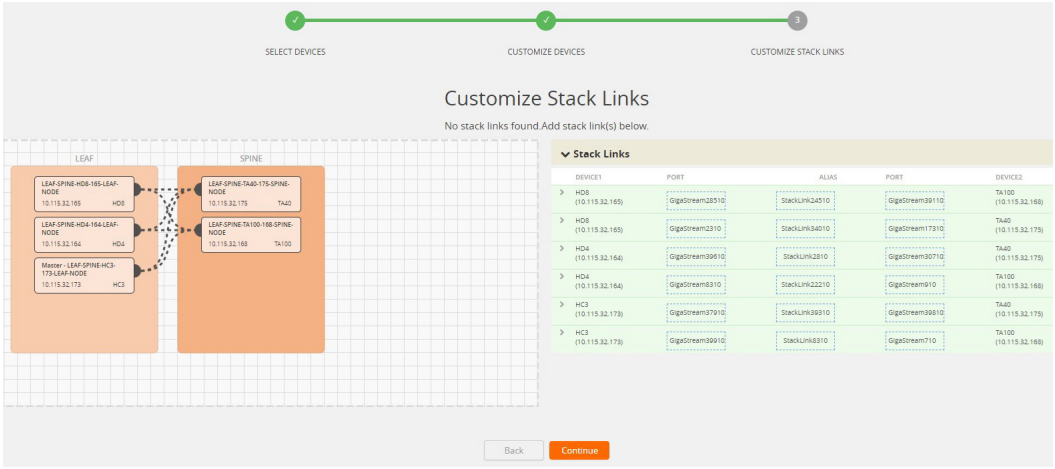

Figure 14-18: Stack link table to create the required GigaStream and stack link

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

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

The Creating Cluster page appears as the cluster is being created.

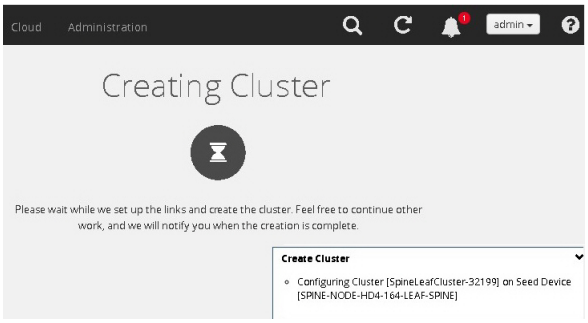

Figure 14-19: Notification window starts once the cluster formation is in progress

The Create Cluster progress window in the lower right corner of the page shows the status of every node as it joins the cluster. It takes a few minutes for the cluster to form. The cluster creation process involves the following steps:

- Cluster[clusterName] Creation Successful followed by Seed device
- Verifying Nodes[Will display HostName of all devices]
- Adding Node[HostName] to cluster [clusterName]
- Node[HostName] successfully joined to the cluster.
- Configuring cards for cluster[clusterName]
- Rediscovering cluster[clusterName]
- Configuring ports for cluster[clusterName].
- Configuring ports will display the status of each stack link and GigaStream whether the creation is successful or not.

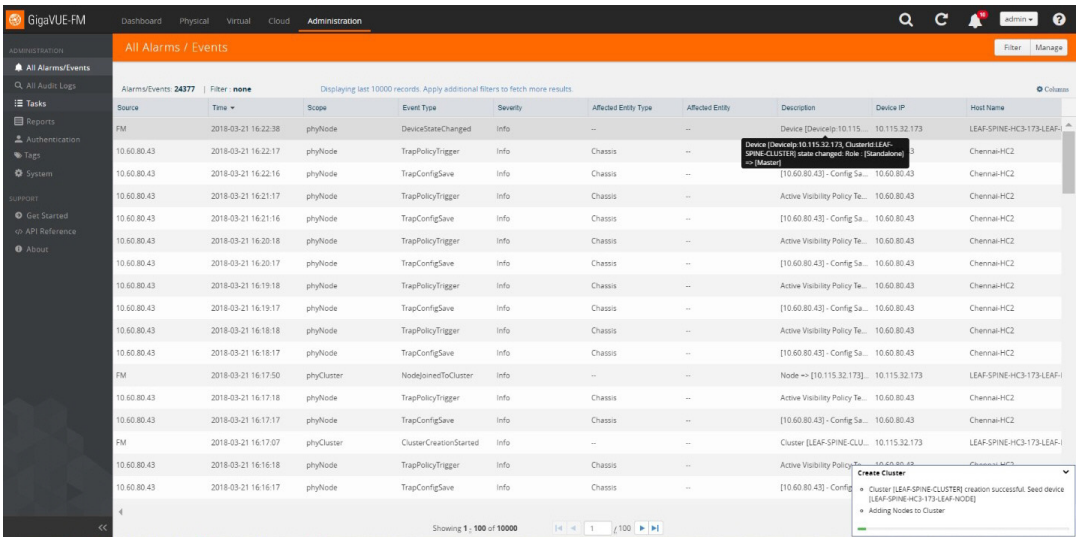

Figure 14-20: Record in Alarms and events page during cluster creation event

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

Figure 14-21: Notification window shows Stack link and GigaStream creation

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

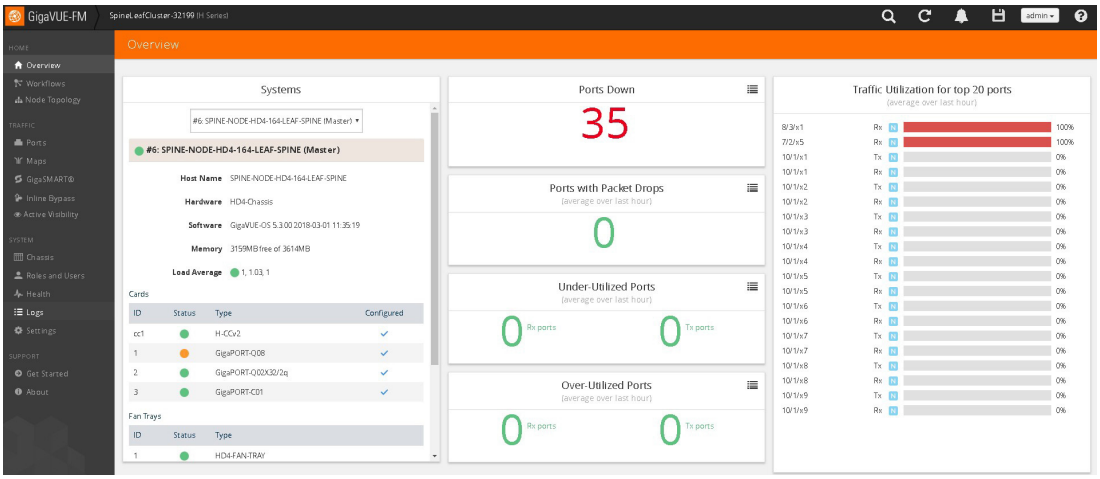

Figure 14-22: Cluster overview after formation

Edit a Cluster

The Edit cluster option supports only the following operations to the existing cluster:

- Multiple devices can be added to the existing cluster in a single update operation.
- Multiple devices can be added as Leaf, Spine, Leaflet.
- Stack links can be created only from the new device which is added into the cluster wizard.
- Master preferences can be changed for each device through edit cluster option.
- Stack link alias and GigaStream alias can be edited for newly created links.

NOTE:

- No option to remove the existing stack links through cluster wizard.
- No option to create links in existing devices.
- Addition and deletion of devices in a single update operation should not be appreciated.
- No option to edit the existing stack link alias and GigaStream alias.

Prerequisites

Standalone devices that have maps cannot be added to cluster if ports used in maps overwrites with the selected ports in stack link table.

This workflow describes how to add a node to a cluster.

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

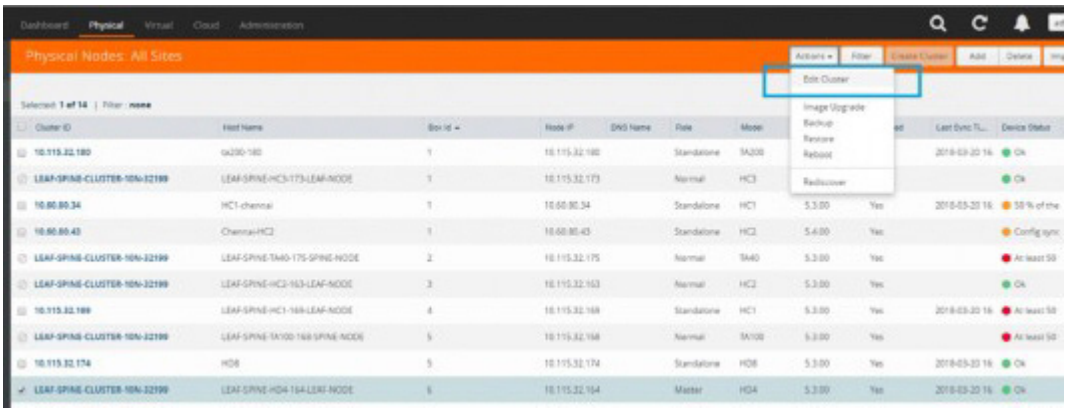

Figure 14-23: Edit cluster

2. The Edit Leaf Spine Cluster Canvas displays the existing stack link configuration details in the cluster wizard canvas. Standalone devices are listed under the Devices pane.

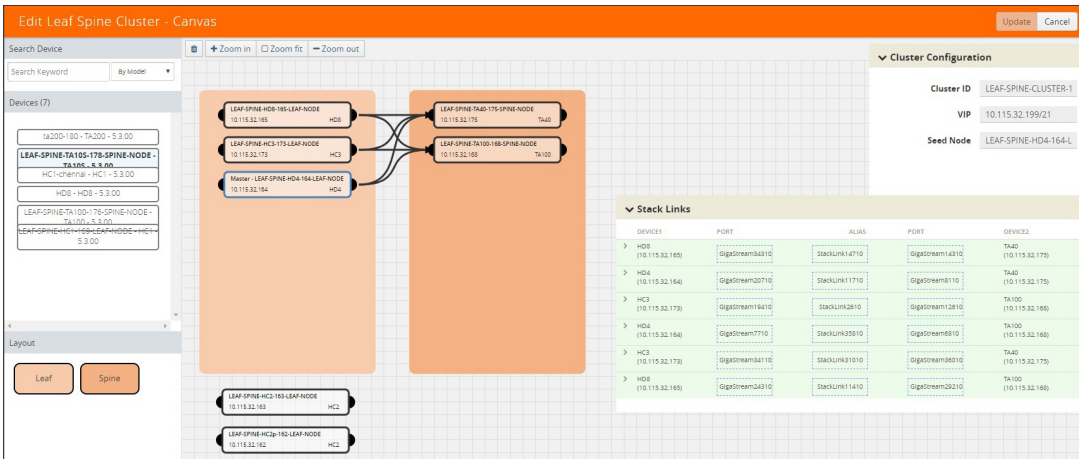

Figure 14-24: Drag required devices into canvas

3. Drag the required devices from the devices pane into the cluster wizard under the leaf or spine container.
4. Draw the links between the newly added device. (**NOTE**: no new link is created for the existing device.)

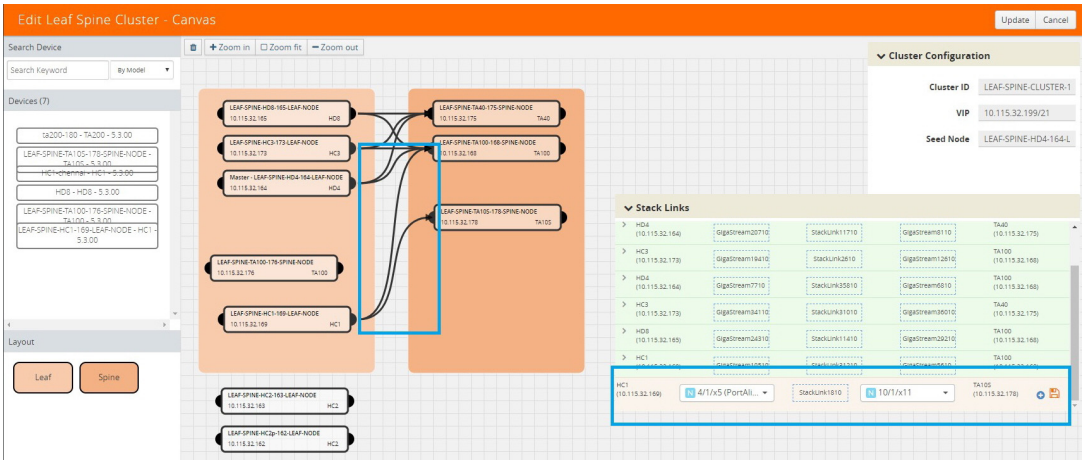

Figure 14-25: Create links between devices

5. Configure the stack link details in the stack link table.
6. After selecting the ports, save the stack link by clicking **Save** in the stack link table
7. Click **Update** to initiate the update process.
8. When prompted by the Confirmation message, click **OK** to run the cluster update.

Figure 14-26: Confirmation message to update the cluster

When the cluster update operation starts, a notification window will appear at the right corner of the GigaVUE-FM window to show the status progression of each node, card, GigaStream and stack link.

When the cluster update operation is completed, “Manage Cluster Completed,”will appear in the in the Manage Cluster notification window.

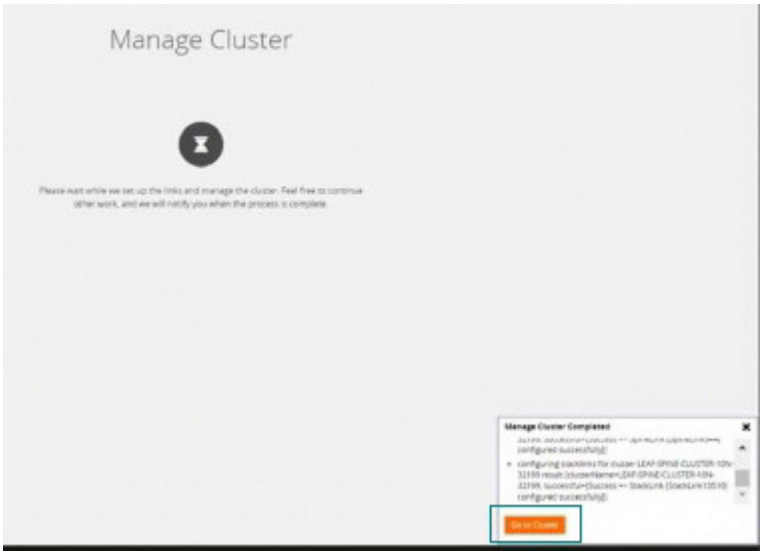

Figure 14-27: Manage Cluster Completed notification

9. Click **Go to Cluster** to go to view the cluster overview.

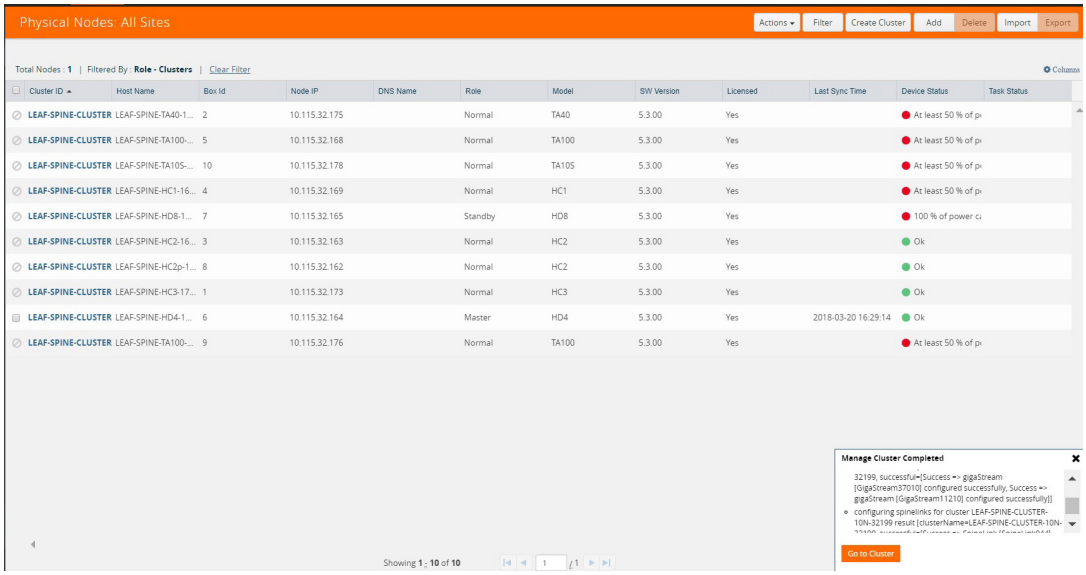

Figure 14-28: Cluster details in the Physical Nodes page

The created GigaStreams will appear in the device Port Groups page, and the created stack links will appear in the device Stack Links page.

Delete a Node from a Cluster

This workflow describes how to remove a device node from a cluster.

NOTE: Only one device can be removed from the cluster per update operation.

The device should not contain any map configurations in a cluster. Those devices cannot be removed until the maps are present.

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

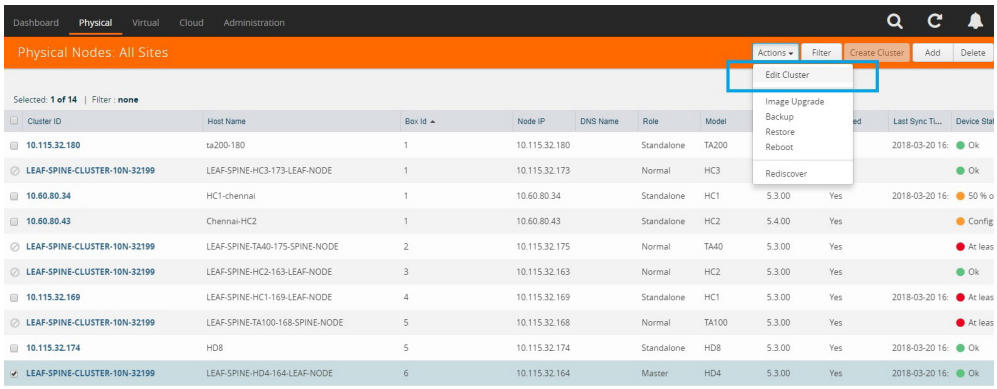

Figure 14-29: Edit cluster

Only one device can be deleted from the canvas. It can be either Leaf, Spine or Leaflet.

2. To remove a device, right-click the device to be removed from the canvas and click **Remove**.

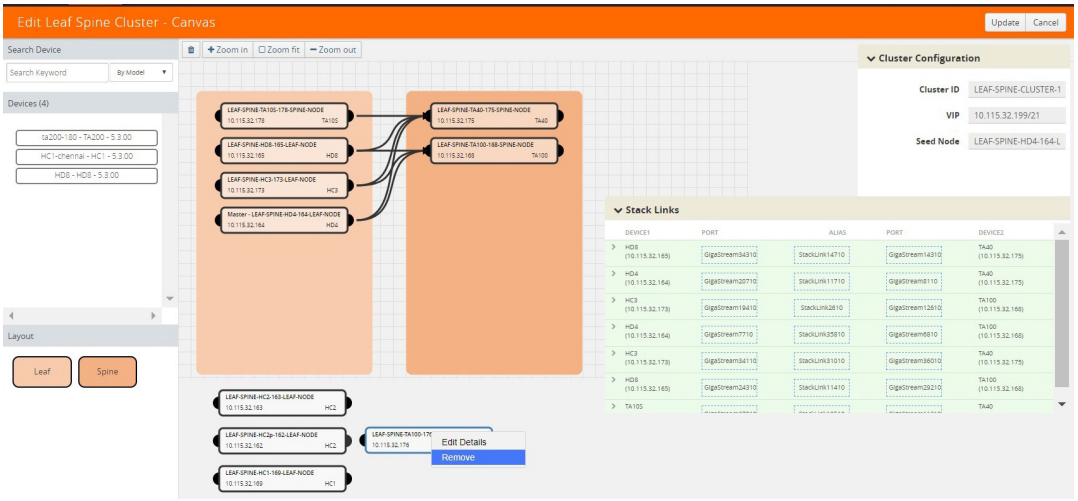

Figure 14-30: Right click the device and remove

The removed device will be deleted from canvas.

3. Click **Update** to initiate the cluster-update operation.

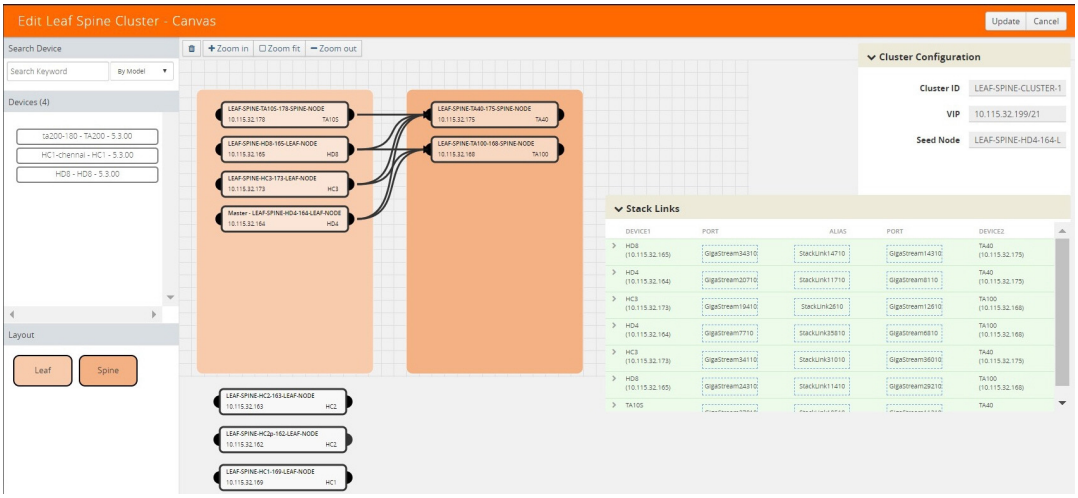

Figure 14-31: Remove will delete the device from canvas

The Manage Cluster notification window shows the progress of nodes being removed from the cluster.

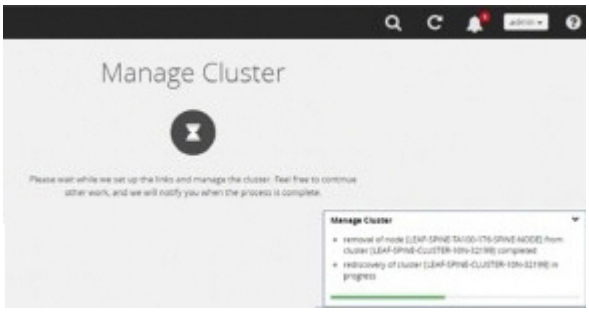

Figure 14-32: Notification for deleting a device from a cluster

When the device is successfully removed from the cluster, a “Manage Cluster Completed,” message will appear in the notification window.

Figure 14-33: Cluster completion status window

4. Click **Go To Cluster** to go into device overview page and see the cluster details.

How to Change the Master Preference of a Device

This workflow describes how to change the device’s master preference.

1. Select a cluster and choose Edit cluster under Actions

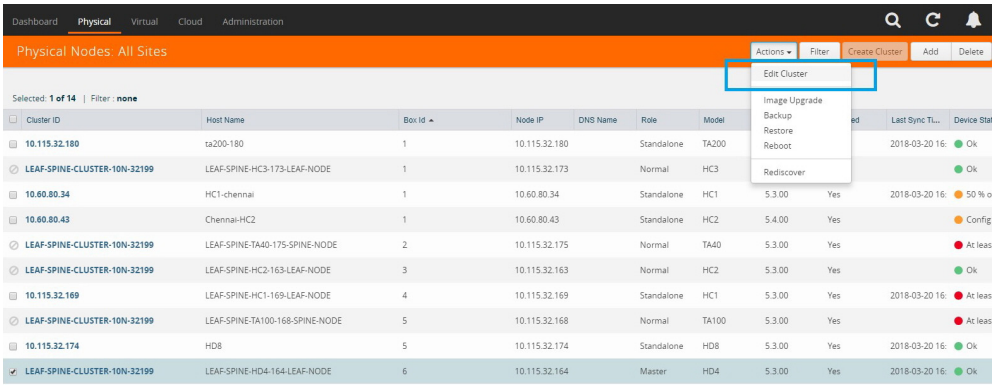

Figure 14-34: Edit cluster

2. To set the master preference for a device, right-click the device and click the **Edit Details** option.

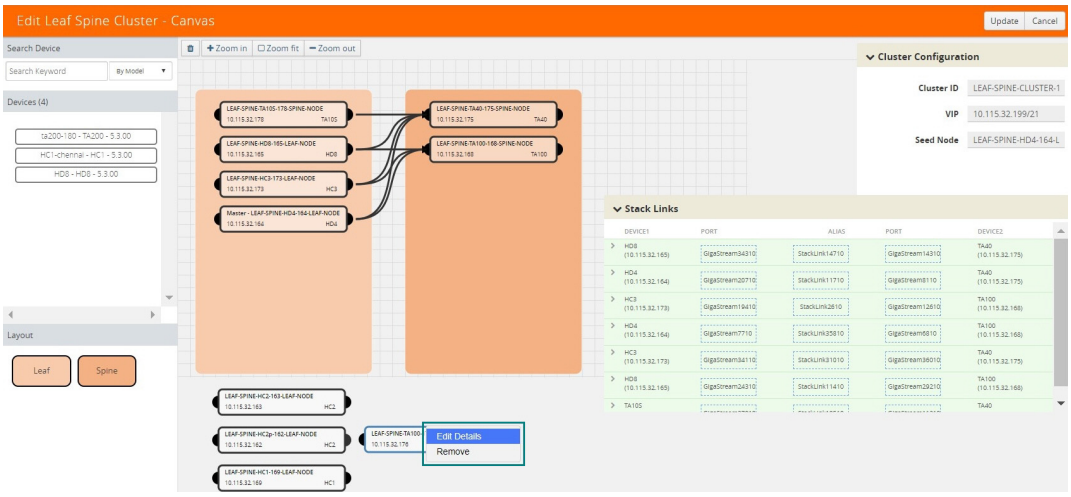

Figure 14-35: Edit details by right clicking the node

3. Click a device in the Edit Leaf Spine Cluster canvas to display the device configuration quick view.

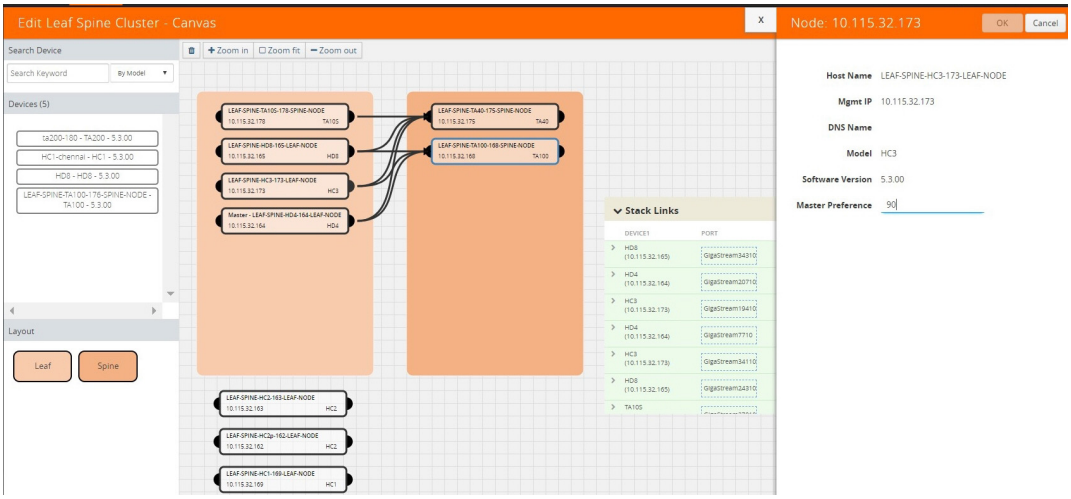

Figure 14-36: Set Master preference

4. Update the master preference text box and click **Update** to proceed.

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

5. Click **Go To Cluster** to go to the device overview page and see the cluster details.

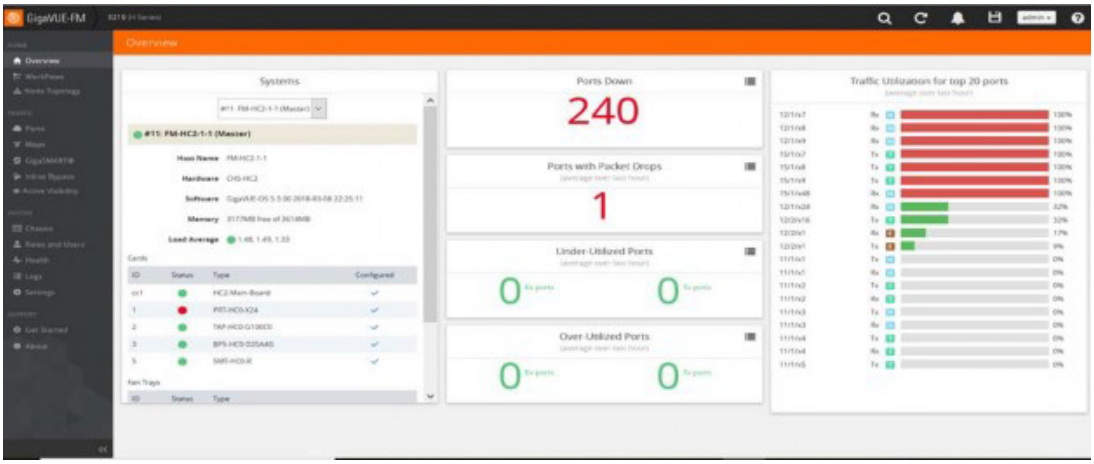

Figure 14-37: Cluster overview page after formation

15 Spine to Spine and Leaf

This chapter describes the Spine to Spine and Leaf architecture for achieving high availability in a cluster environment. Refer to the following sections for details:

NOTE: Refer to *Regular Cluster Formation Workflow* on page 212 for how to use the Regular Cluster workflow.

Introduction to Spine to Spine and Leaf

The Spine to Spine and Leaf architecture is a multi-layer architecture used for network aggregation. This architecture supports leaf nodes and multiple levels of spine nodes. In Spine to Spine and Leaf architecture, connect the leaf and spine as follows:

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

With multiple paths between the nodes in a cluster, the spine to spine and leaf architecture protects against traffic congestion, failures, such as stack link or spine node failures. In the event of a failure, the traffic on one path fails over to the other path. This architecture provides resiliency to the network.

An example of a Spine to Spine and Leaf architecture is shown in Figure 15-1.

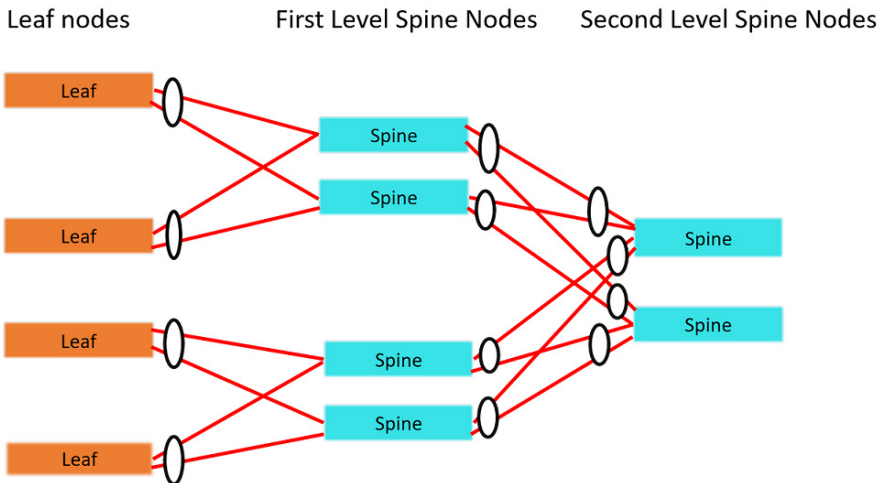

Figure 15-1: Spine to Spine and Leaf Architecture

In a Spine-Leaf cluster, the number of leaf nodes is typically higher than the number of spine nodes. In Figure 15-1, there are four leaf nodes, four first level spine nodes, and two second level spine nodes. The leaf nodes aggregate to a fewer number of spine nodes.

For more information on Spine to Leaf architecture, refer to *Multi-Path Leaf and Spine* on page 251.

In Figure 15-1, the spine nodes are GigaVUE TA Series nodes, such as GigaVUE-TA100, TA40, TA10, or TA200 while the leaf nodes are GigaVUE H Series nodes, such as GigaVUE-HC2, GigaVUE-HC3, or GigaVUE-HD8, which places the traffic intelligence at the edge.

Traffic between ports on the same leaf node will be local to that leaf node, but traffic between different leaf nodes will go through the spine nodes.

The traffic from a source leaf node to a destination leaf node flows as follows:

- From a TAP, traffic flows to the source leaf node.

- From the source leaf node, traffic is load balanced to the connected spine nodes.
- From the spine node, depending on the configuration, traffic flows to the next level of spines or the destination leaf node.
- From the destination leaf node, traffic flows to the tool ports.

Resiliency is achieved when there are multiple paths from the network to the tools across GigaVUE nodes.

Refer to *Path Protection* on page 253 for the leaf node failure, stack link failure on a leaf node or spine node.

Configuration Overview

This section provides an overview of the configuration. You must perform the configuration from the master node in the cluster. Follow this configuration sequence to prevent loops:

1. *Configure Stack GigaStream*
2. *Configure Spine Links*
3. *Configure Stack Links*

This configuration connects nodes using multiple paths. For an example of the configuration, refer to Figure 15-2 on page 275.

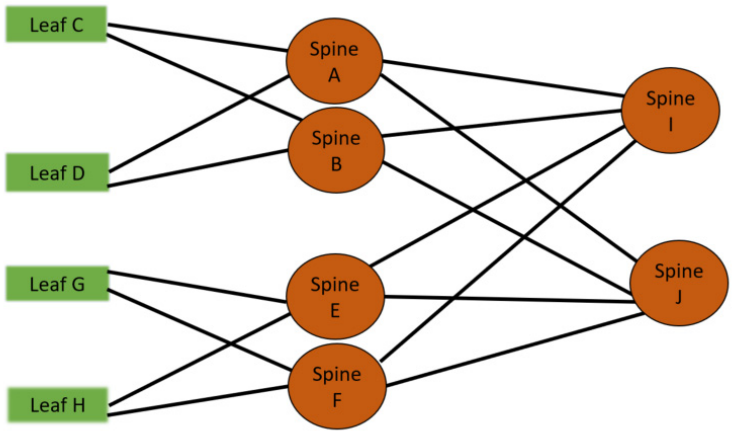

Figure 15-2: Spine to Spine and Leaf Configuration

Configure Stack GigaStream

The stack GigaStream connect the spine and leaf nodes. In Figure 15-3 on page 276, the stack GigaStream are: a1, a2, a3, a4, b1, b2, b3, b4, c1, c2, d1, d2, e1, e2, e3, e4, f1, f2, f3, f4, g1, g2, h1, h2, i1, i2, i3, i4, j1, j2, j3, j4. Even if there is only one port that connects the nodes, you must still configure a stack GigaStream. With a configuration of 6 spine nodes and 4 leaf nodes, the number of stack GigaStream is 32.

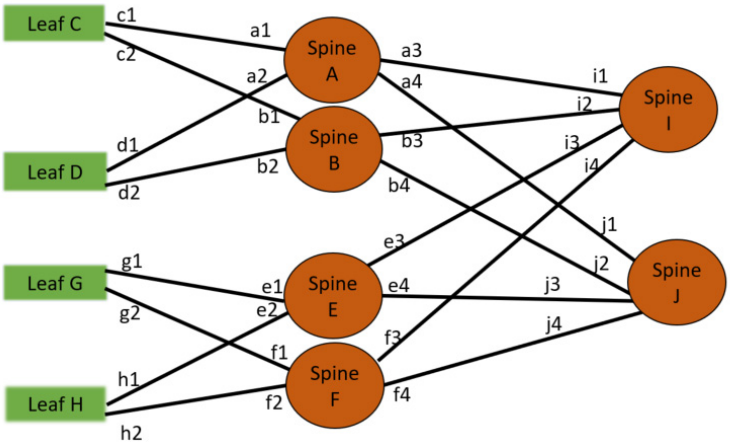

Figure 15-3: Stack GigaStream Configuration

Configure Spine Links

On each leaf node, there is one spine link that contains the list of GigaStream connecting the leaf nodes to the spine nodes. The spine links contain multiple stack GigaStream bundled together. The spine links on the leaf nodes are: {c1,c2}, {d1,d2}, {g1, g2} and {h1, h2}.

The spine node also contains the spine links. A spine node connecting to another spine nodes has a spine link. The spine links on the spine nodes are: {a3, a4}, {b3, b4}, {e3, e4}, {f3, f4}, {i1, i2}, {i3, i4}, {j1, j2}, and {j3, j4}. The total number of spine links is twelve for this configuration. Across the spine link members, traffic is load balanced. For this part of the configuration, refer to the circles in Figure 15-4 on page 276.

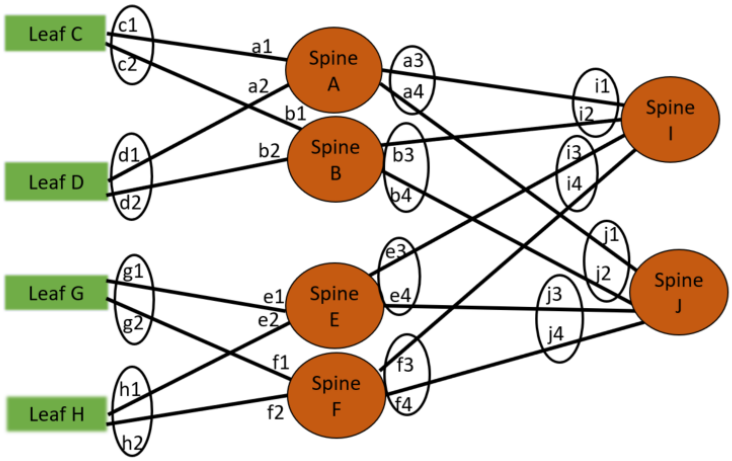

Figure 15-4: Spine Link Configuration

NOTE: You must not configure any spine links from spine nodes to leaf nodes. For example, in the Figure 15-4 on page 276, a1 and a2 should not be configured as spine links.

Configure Stack Links

The stack links are: {a1,c1}, {a2,d1}, {a3,i1}, {a4,j1}, {b1,c2}, {b2, d2}, {b3, i2}, {b4, j2}, {e1, g1}, {e2, h1}, {e3, i3}, {e4, j3}, {f1, g2}, {f2, h2}, {f3, i4}, {f4, j4}. The total number of stack links is sixteen for this configuration. For this part of the configuration, refer to the circles in Figure 15-4 on page 276.

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

Notes and Considerations

For the limitations and considerations, refer the Notes and Considerations in Page 8.

Configuration of Spine to Spine and Leaf Architecture

To configure Spine to Spine and Leaf Architecture, follow these steps:

1. Configure Leaf-Spine Cluster formation workflow. To configure, refer *Leaf-Spine Cluster Formation Workflow* on page 259.
2. Execute the Gigamon Automation GigaVUE-FM SDK to establish the connection between the first level of spines and the second level of spines.
3. View the renewed topology in GigaVUE-FM, and reorder them if required.
4. Verify the configurations performed.

You can view the updated topology in GigaVUE-FM only after the execution of the Gigamon Automation GigaVUE-FM SDK to establish the connection between the first level and second level spines.

Limitations

- It is not recommended to have network and tool in one segment and GSOP in other segment. You must follow any one of the below arrangements:
  - Network and GSOP in one segment of the spine, and the tool in other segment of the spine.
  - Network in one segment of the spine, and GSOP and tool in other segment of the spine.
- If the network port is in first level spine of segment 1, then the leaf should be in the leaf of segment 2.
- If the network port is in first level spine of segment 2, then the leaf should be in the leaf of segment 1.

Leaf-Spine Cluster Deployment

This section describes the steps and prerequisites to deploy a leaf-spine cluster.

Refer to *Introduction to Spine to Spine and Leaf* on page 274 for a conceptual overview of the leaf-spine architecture.

Deployment Checklist

Before forming a Leaf-Spine Cluster, it is strongly recommended that you get familiar with the relevant documentation and review the deployment checklist to prepare for deployment.

Documentation

- Review the *GigaVUE-FM Release Notes* to familiarize yourself with the functionality around creating and managing clusters.
- Review the “Multi-Path Leaf and Spine” chapter to familiarize yourself with leaf and spine architecture.
- Review the *GigaVUE-FM Release Notes* for any known issue that may impact your use case.

Pre-deployment checklist

- Gigamon Fabric Management must be upgraded to Gigamon 5.4.01 or later.
- Gigamon device must be upgraded to GigaVUE-OS 5.4.00 or later.
- Advanced Features License must be installed in TA devices.
- Physical connection must be established to create stack links.
- Devices must have GDP enabled and be physically connected to create links among devices from Gigamon.

IMPORTANT: Recommendation is to use TA devices as SPINE Nodes and other devices as LEAF Nodes.

Formation Scenario

The Spine to Spine and Leaf cluster is formed with different combinations of devices with Spine and Leaf nodes as a node cluster.

The following configuration creates a leaf-spine cluster with four leaves, four first level spines, and two second level spines.

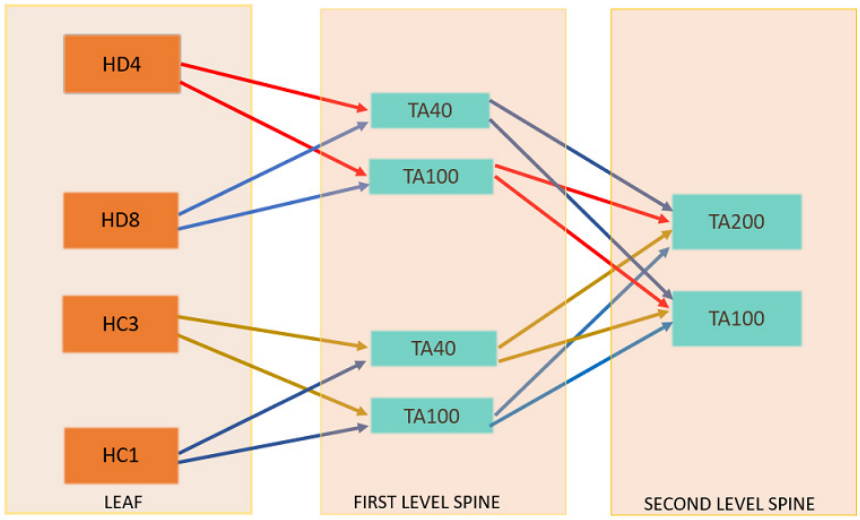

Figure 15-5: Spine to Spine and Leaf cluster overview,

NOTE: GigaStreams support different speeds, as indicated by the different colored connector lines in Figure 15-5 on page 279.

16 Manage G Series Nodes

G Series nodes can be added to GigaVUE-FM. All the G Series node whether standalone or in a stack configuration are visible on the GigaVUE-FM. The nodes are read-only view and cannot be configured from GigaVUE-FM. However, current configuration of G Series nodes from GigaVUE-FM can be backed up or restored to the nodes.

The sections in this chapter include:

G Series Nodes Dashboard

G Series nodes cannot be configured from GigaVUE-FM. You can only view the maps created directly on the nodes through CLI or G-VUE (UI interface for G Series). When you click on a G Series node, you will see a different structure as compared to H Series nodes. (Refer to Figure 16-1.) To review more on setting up G Series nodes, please reference the *G Series CLI User's Guide*. To review G Series nodes in a stack configuration, click on each node to monitor, setup or configure the nodes.

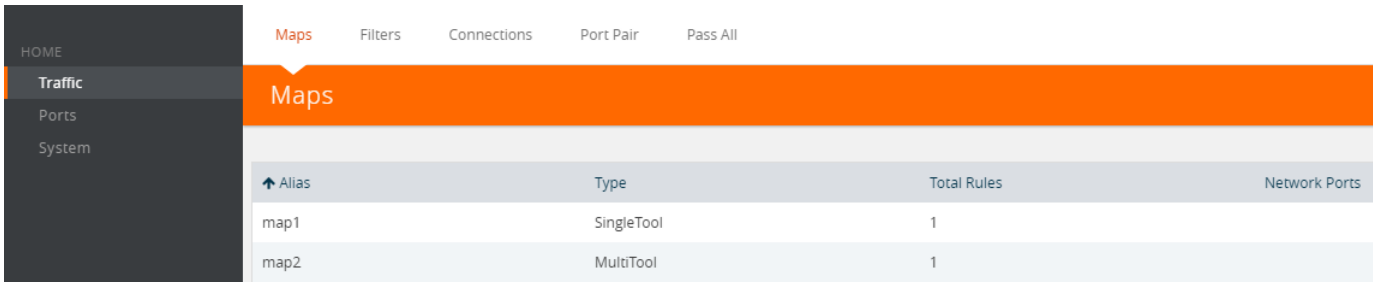

Figure 16-1: G Series Nodes Overview Page

Traffic

The navigation pane for G Series nodes gives you the options to view traffic, ports, and system. Under the traffic tab, there are options to view maps, port filters, port-lists, port-pairs and pass-all maps.

Maps Page

The **Maps** page displays information about the rules that have been set for the node selected in the Physical Nodes. Table 16-1 describes the parameters that display on the Maps page for G Series nodes.

Table 16-1: Parameters on the Maps Page for G Series Nodes

Parameter	Description
Alias	Displays the map's alias.
Type	Indicates the type of map. Either MultiTool or SingleTool.
Total Rules	The number of rules defined in the map.
Network Ports	Lists the network ports to which this map is bound (Box ID/Port ID). Rules send matching traffic to named tool ports. A map typically has multiple rules. But the map as a whole is bound to the same set of network ports. it does not vary by rule.

Filters Page

The **Filters** page displays the filters that have been set for the node selected in the Physical Nodes table.

The Filters table displays the filters available in the filter database on the node. They are not necessarily bound to a port. Once a filter is bound to a port, it becomes a port-filter, and the port is listed in the **Assigned Ports** column.

Table 16-2 describes the parameters that display on the Filters page for G Series nodes.

Table 16-2: Parameters on the Filters Page for G Series Nodes

Parameter	Description
Alias	The filter's alias, assigned when the administrator created the filter.
Action	Indicates the action of the filter. Either Allow or Deny.
Rules	The criteria for the filter.
CLI	The command a user can copy and paste into the node's CLI to create the filter.
Assigned Ports	Lists the ports to which the filter is applied. Either network or tool. If this shows **N/A**, the filter is either not currently applied to a port (that is, it's in the Filter Database but not applied as a port-filter), or you have selected the Master node in an H Series cluster. In the latter case, you can select the individual node where the filter is applied to see the Assigned Ports entry populate correctly.

Connections Page

The **Connections** page displays information about the port connections in place on the node selected in the Physical Nodes table. Connections are simple one-to-one flows between a network port and a tool port or advanced-hash GigaStream. You can set up filters on either end of a connection (pre-filter or post-filter), set up multiple connections on a single network port, or send all the data arriving on a network port to a designated tool port.

Table 16-4 describes the parameters for connections that display on the G Series Connection page for G Series nodes.

Table 16-3: Parameters on the Connections Page for G Series Nodes

Parameter	Description
Alias	The alias for the connection. Assigned when the administrator created the connection.
Network Port List	The network port (Box ID/Port ID). Each entry is a link. Click the link to open the Port Config dialog listing: Port Name, Alias, Owner, Alarm Threshold, Map alias, Filters.
Tool Port List	The tool port (Box ID/Port ID). Each entry is a link. Click the link to open the Port Config dialog listing: Port Name, Alias, Owner, Alarm Threshold, Map alias, Filters.
Comment	A comment entered by the administrator when setting up the connection.

Port-Pair Page

The **Port-Pair** page displays information about the port-pairs that have been set for the node selected in the Physical Nodes table.

A port-pair is a bidirectional connection in which traffic arriving on one port in the pair is transmitted out the other (and vice-versa) as a pass through tap.

Table 16-4 describes the parameters for port pairs that display on the Port Pair page for G Series nodes.

Table 16-4: Parameters on the Port-Pair Page for G Series Nodes

Parameter	Description
Alias	The alias for the port pair. Assigned when the administrator created the port pair.
First Port	The first port in a port pair (Box ID/Port ID). Each entry is a link. Click the link to open the Port Config dialog listing: Port Name, Alias, Owner, Alarm Threshold, Map alias, Filters.
Second Port	The second port in a port pair (Box ID/Port ID). Each entry is a link. Click the link to open the Port Config dialog listing: Port Name, Alias, Owner, Alarm Threshold, Map alias, Filters.
Comment	A comment entered by the administrator when setting up the port pair.

Pass-All Page

The **Pass-All** page displays information about the pass-all maps that have been set for the node selected in the Physical Nodes table.

The pass-all command can be used to send all packets on a network or tool port to one or more tool ports or advanced-hash GigaStreams, irrespective of the connections or maps already in place for the ports.

Table 16-5 *describes the parameters that display on the Pass All page for G Series nodes.*

Table 16-5: Parameters on the Pass-All Page for G Series Nodes

Parameter	Description
Alias	The alias for the Pass All. Assigned when the administrator created the Pass All.
Source Ports	The source port in the Pass All (Box ID/Port ID). Each entry is a link. Click the link to open the Port Config dialog listing: Port Name, Alias, Owner, Alarm Threshold, Map alias, Filters.
Destination Ports	The destination port in the Pass All (Box ID/Port ID). Each entry is a link. Click the link to open the Port Config dialog listing: Port Name, Alias, Owner, Alarm Threshold, Map alias, Filters.
Comment	A comment entered by the administrator when setting up the Pass All.

Ports

Ports link in the navigation pane provides access to the Ports and GigaStreams page.

Ports Page

The Ports Page provides information relating to the Port ID for the node.

Table 16-6 describes the parameters that display on the Ports page for G Series nodes.

Table 16-6: Parameters on Ports Page on G Series Nodes

Parameter	Description
Port ID	Box ID and Port ID for the Port
Alias	Alias name for the Port ID
Type	Port Type: Network, Tool, Stack, Inline tool and Bypass network
Admin Status	Shows the port is enabled or disabled
Link Status	Provides the link status for the ports.
Transceiver Type	Transceiver type as electrical or optical.
Speed	Sets the port speed in Mb/s if autonegotiation is off. Only available for 1Gb copper ports not available for 10Gb capable ports with 1Gb transceiver installed.
Duplex Mode	Ports can be half or full duplex, if autonegotiation is off.
Auto-Negotiation	When autonegotiation is enabled, duplex and speed settings are ignored and set via autonegotiation.
Force Link Up	Forces connection to an optical tool port. Use this option when an optical port is connected to a legacy optical tool that does not transmit light.

GigaStreams Page

The GigaStreams page provides details regarding the G Series node similar to the GigaStreams page for H Series. To learn more about setting up GigaStreams between G Series nodes, refer to the *G Series CLI User's Guide*.

Table 16-7 describes the parameters that display on the GigaStream page for G Series. nodes.

Table 16-7: Parameters for GigaStreams Tab on G Series Nodes

Parameter	Description
Alias	Alias name for GigaStreams.
Algorithm	Hashing Algorithm used for GigaStreams.
Tool Ports	Destination ports for the GigaStreams created.
Type	Port type as Network GigaStreams or Tool GigaStreams.

Table 16-7: Parameters for GigaStreams Tab on G Series Nodes

Parameter	Description
Fail Over	Failover mode configured as T or F. T for failover active to all configured ports.

System

The System option in the navigation pane, displays the Chassis page for the node. The Chassis page provides details regarding the G Series node similar to the Chassis page for H Series. However, there are some limitations. You will not be able to see all the cards available in a GV2404 chassis. To see the details, click on the Cards, Power Modules, or Fan Trays sections.

G Series Node Upgrade from GigaVUE-FM

For information and procedure for upgrading G Series nodes and clusters from GigaVUE-FM, refer to *Upgrade Software on a GigaVUE Node or a Cluster from GigaVUE-FM* on page 228 and *Bulk Configuration* on page 1313.

17 Fabric Statistics

This chapter provides the statistics of all types of ports available in GigaVUE H series and TA series.

Refer to the following sections for details:

About Fabric Statistics

GigaVUE-FM provides the ability to view detailed information about the packets dropped, packets discarded, and packets received and transmitted by the following ports:

- Network ports
- Tool ports
- Hybrid ports
- Inline-network ports
- Inline-tool ports
- GigaSMART engine ports
- Backplane ports
- XAUII ports
- Stack ports

Network ports, tool ports, hybrid ports, inline-network ports, and inline-tool ports are also called front panel ports. These are the ports that are visible on the front view of the GigaVUE node.

The backplane ports on the control card are connected to the backplane ports on the line card or GigaSMART engine ports. They allow packets to move from card to card, for example, from line card to control card and control card to GigaSMART card. If there are packet errors in the backplane ports, this information can be viewed in the Fabric Statistics page. The format

used to represent a backplane port is <box ID>/<slot ID>/<port ID>. For example, 1/1/s1 where s1 refers to port s1.

GigaVUE-FM also collects statistics of XAUII links on the GigaSMART engine ports. Each GigaSMART engine port is made up of 2 to 4 XAUII links, depending upon the platform. They are considered as the child ports of GigaSMART engine ports. The format used to represent a XAUII port is <box ID>/<slot ID>/<port ID>. For example, 1/5/e1x0 where e1x0 refers to the first XAUII port on engine 1.

Display Fabric Statistics for All Ports

Using the Fabric Statistics page, you can view the statistics associated with the port types. If there are packet drops in the ports, click on the port and drill down further to investigate the cause of the packet drops.

To view the Fabric Statistics page:

1. Click **Physical** in the top navigation link.
2. In the Physical Nodes page, click on a GigaVUE node.
3. On the left navigation pane, select **Ports > Ports > Fabric Statistics**.

Table 17-1 describes the columns in the Fabric Statistics table.

Table 17-1: Fabric Statistics Definitions

Counter	Definition	Notes
Port ID	Port ID in the format <box id/slot id/port id>	
Alias	Alias of the port	
Unicast Packets Rx	Total unicast packets received	Excludes multicast packets, broadcast packets, packets with FCS/CRC errors, MTU exceeded errors, oversize packets, and pause packets.
Unicast Packets Tx	Total unicast packets transmitted	Excludes multicast packets, broadcast packets, packets with FCS/CRC errors, MTU exceeded errors, oversize packets, and pause packets.
Packets/sec Rx	Total packets received per second Excludes multicast packets, broadcast packets, packets with FCS/CRC errors, MTU exceeded errors, oversize packets, and pause packets.	Excludes packets with FCS/CRC errors.
Packets/sec Tx	Total packets transmitted per second	Excludes multicast packets, broadcast packets, packets with FCS/CRC errors, MTU exceeded errors, oversize packets, and pause packets.
Octets/sec Rx	Total bytes received per second This indicates the rate of incoming bytes in the last second.	

Table 17-1: Fabric Statistics Definitions

Counter	Definition	Notes
Octets/sec Tx	Total bytes transmitted per second This indicates the rate of incoming bytes in the last second.	
Utilization Rx	Percentage of port utilization by packets received	
Utilization Tx	Percentage of port utilization by packets transmitted	
Link Speed	Maximum link speed of the port	
Type	Type of port	
Octets Rx	Total bytes received Includes all valid and error frames with the exceptions noted in the adjacent columns.	Excludes undersize frames.
Octets Tx	Total bytes transmitted Includes all valid and error frames with the exceptions noted in the adjacent columns.	Excludes undersize frames.
Non-unicast Packets Rx	Total Broadcast and Multicast packets received	
Non-unicast Packets Tx	Total Broadcast and Multicast packets transmitted	
Packet Drops Rx	Total received packets dropped	Packets are dropped when a tool port's bandwidth is exceeded due to oversubscription. Packets are dropped when they reach the tool port but before they are sent out.
Packet Drops Tx	Total transmitted packets dropped	Packets are dropped when a tool port's bandwidth is exceeded due to oversubscription. Packets are dropped when they reach the tool port but before they are sent out.
Discards Rx	Total received packets discarded	This counter increments when a packet is discarded at a tool port due to a tool port map rule.
Discards Tx	Total transmitted packets discarded	This counter increments when a packet is discarded at a tool port due to a tool port map rule.

Table 17-1: Fabric Statistics Definitions

Counter	Definition	Notes
Error Rx	Total error packets received Error packets include undersize, FCS/CRC, MTU exceeded, fragments, and oversize packets.	Excludes oversize packets without FCS/CRC. Packets larger than the MTU setting arriving on a network port are counted twice in the counter. So 1000 oversize packets would show up as 2000. This double-counting only happens with Oversize error packets.
Error Tx	Total error packets transmitted Error packets include undersize, FCS/CRC, MTU exceeded, fragments, and oversize packets.	Excludes oversize packets without FCS/CRC. Packets larger than the MTU setting arriving on a network port are counted twice in the counter. So 1000 oversize packets would show up as 2000. This double-counting only happens with Oversize error packets.

The port types in the Fabric Statistics table are represented as follows:

Table 17-2: Port Types

Port	Port Type
Network Port	N
Tool Port	T
Hybrid Port	H
Inline Network Port	iN
Inline Tool Port	iT
GigaSMART Engine Port	E (bigger size)
Backplane Port	BP
XAUII Port	E (smaller size)
Stack Port	S

Display Fabric Statistics for a Single Port

On the Fabric Statistics page, select any port (except for backplane and XAUII) and click anywhere in the row. The port quick view is displayed. The port quick view provides information about a specific port. Each field is color-coded in the graphical representation as shown in Figure 17-1 on page 292.

Figure 17-1: Port Statistics Quick View

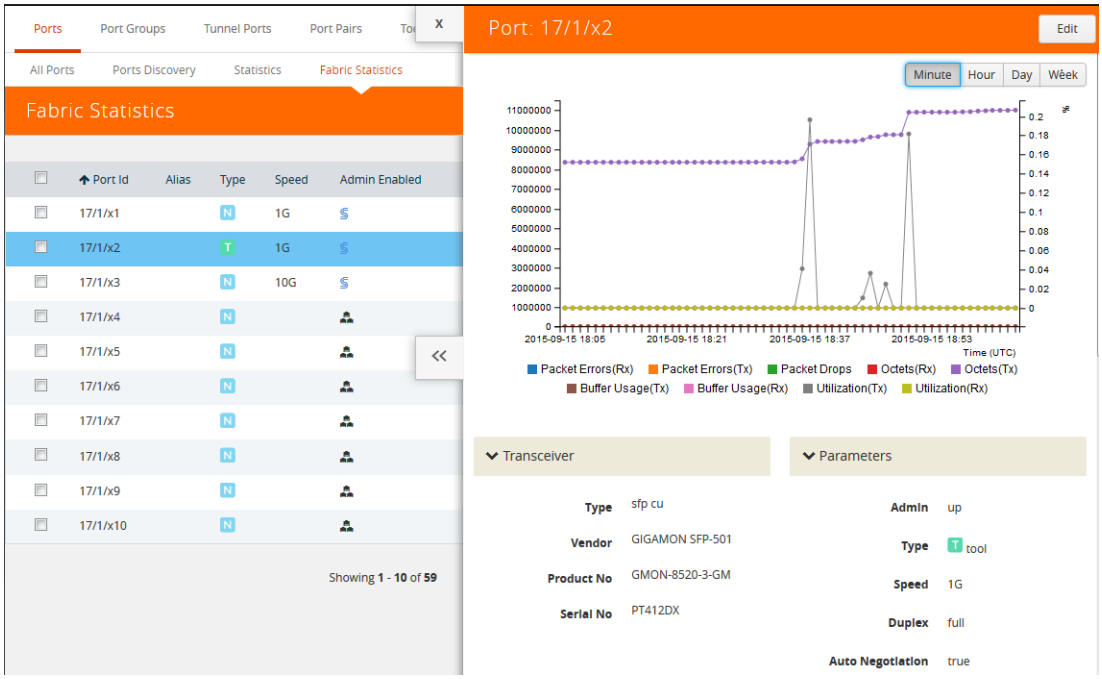

Port Quick View

The port quick view provides a graphical view of the port statistics. You can choose to view the statistics based on an hour, a day, a week, a month, or live to see the real-time data.The quick view also provides information about port properties, statistics about the traffic received (Rx) and transmitted (Tx), and alarms information. Click Counters drop-down list to select other options.

The following table describes the parameters displayed in the Port quick view:

Field	Description
Port Info	Displays the port information such as the name of the port, if any.
Parameters	Displays information about the current link status of the port, admin status, port type, port speed, port's duplex configuration, auto negotiation configuration, and the force link up configuration.
Transceiver	Displays the transceiver type connected to the port, optical input power, vendor from where the transceiver is purchased, product number, serial number, and temperature.
Filters	Displays the number of filters, pass rules, and drop rules applied to the port.
Alarms	Displays the buffer threshold percentage on the ports in the Rx and Tx directions, high and low utilization threshold percentage, and the port utilization in the Rx and Tx directions.
Permissions	Displays the roles and permissions associated with the port, users with permission to lock the port, and users to whom the port is shared.
Related Maps	Displays the related maps associated with the port.

Export Fabric Statistics

To export the fabric statistics, click **Export**. The statistics table is downloaded with a filename in the format Port_Stats_<yyyymmddhhmmss>; for example, Port_Stats_20161003172336.

Filter Fabric Statistics

Using filters, you can search and narrow down the ports displayed in the Fabric Statistics page. To filter the fabric statistics, click **Filter**. A filter quick view appears as shown in Figure 17-2 on page 294. Specify the criteria to filter the ports.

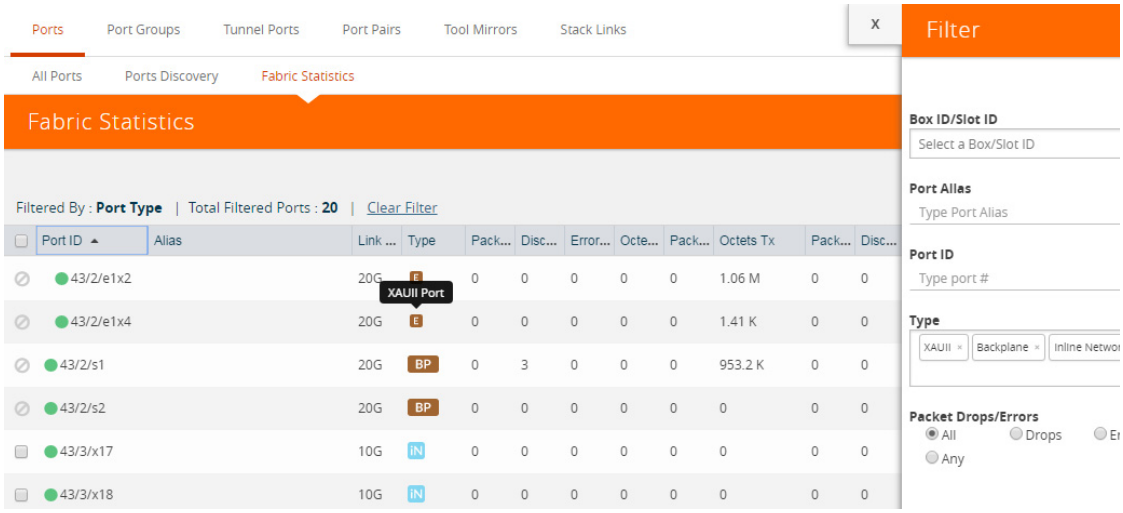

Figure 17-2: Fabric Statistics Filter

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

Criteria	Description
Box/Slot ID	Displays only those ports that match the specified box and slot IDs.
Port Alias	Displays a port with specified alias.
Port ID	Displays ports with specified port ID. For example, if you specify 3 as the port ID, the result will display ports that include the number 3, 13, 23, 30, and so on.
Type	Displays ports with specified port type. Select one of the following: • Network • Tool • Inline Network • Inline Tool • GigaSMART • Hybrid • Stack • Backplane • XAUII

Criteria	Description
Packet Drops/Errors	Displays only those ports with packet drops or errors or both. Select one of the following: • All — displays ports with both packet drops and errors. This is the default. • Drops — displays ports with only packet drops. • Errors — displays ports with only packet errors. • Any — displays ports with either packet drops or packet errors.

18 Topology Visualization

This section describes the Topology view available in GigaVUE-FM. The Topology view displays the nodes and clusters managed by GigaVUE-FM.

This chapter covers the following topics.

Overview of Topology

Topology shows all the physical nodes and clusters that are discoverable.

The GigaVUE H Series and TA Series nodes are capable of snooping Link Layer Discovery Protocol (LLDP) packets and Cisco Discovery Protocol (CDP) packets. Starting in GigaVUE-FM 5.2, they are also capable of sending Gigamon Discovery packets.

GigaVUE nodes receive LLDP and CDP packets from non Gigamon devices. They receive Gigamon Discovery packets from GigaVUE nodes managed by GigaVUE-FM.

If the nodes in your network use either of these protocols, they can identify its immediate neighbors and their capabilities. The LLDP, CDP, and Gigamon discovery information includes the remote port and chassis IDs, as well as other selected information, if it is included by the sender. This information can be used to determine the origin of traffic flows.

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

Gigamon discovery is bidirectional. GigaVUE nodes transmit the LLDP packets with a unique type-length-value (TLV) structure that is identifiable by another GigaVUE node. When a GigaVUE node receives these packets, it detects the TLV structure and identifies the discovered node as a GigaVUE node. These Gigamon discovery packets also include the management IP4 and IP6 address of the node in the unique TLV structure. Using the management IP address, the neighbor information of the discovered nodes is collected to build the entire physical topology.

The Gigamon discovery packets can be transmitted out of tool, stack, network, and hybrid port types. For stack port types, Gigamon Discovery works only when both ends of the link are of type stack.

The hybrid ports and ports configured with force link up are operated in loopback mode. So, when Gigamon discovery is enabled, these ports show themselves as neighbors.

The Gigamon discovery packets are transmitted as soon as the link is up on the relevant port. They are retransmitted when changes are made to the port or chassis configuration such as the host name on the box, port type, IP address of the management interface, and so on. When the transmission or retransmission of Gigamon discovery packets begins, the interpacket time intervals per port is 1, 2, 4, 8, 16, 30, 30, 30, ... seconds. The packets are transmitted every 30 seconds so that the nodes are discovered as soon as possible. Discovery packets are terminated on the receiving GigaVUE node.

GigaVUE-FM constructs the topology by doing the following:

- Groups discovered nodes into clusters by combining nodes with the same cluster ID.

- Queries the master node for the existing stack links.
- Extracts the box IDs from the connected nodes.
  - For port-based stack links, the box ID is extracted from the interconnected stack port IDs.
  - For GigaStream-based stack links, the box ID is extracted from the first port member of each GigaStream.

Topology construction has the following limitations:

- GigaVUE-FM relies on the cluster configuration to construct the topology.
  - Cascading nodes interconnected over network or tool ports or over Port Groups or GigaStream (other than stack links) are not discovered.
  - GigaVUE-FM does not detect miswired cluster cabling.
  - If GigaVUE H and TA Series nodes connected over tool ports are not managed by GigaVUE-FM, then they are not automatically discovered.
- GigaVUE-FM constructs H Series cluster topology by querying nodes for the relevant connectivity information. The topology is constructed based on information from the GigaVUE nodes.
- Gigamon discovery is not supported on inline-tool and inline-network port type.
- For stack ports, Gigamon discovery will work only if both ends of the link are of type stack.
- If a pass-all map is configured for Gigamon discovery enabled ports, the show map statistics will show an increment in the map rule counter for Gigamon discovery packets although the Gigamon discovery packets are not sent to the tool ports.
- If Gigamon discovery is enabled on a port, for example port 9/3/x1, and the port type is changed to an unsupported port types such as inline-network, Gigamon discovery is immediately disabled and an error message is displayed.

Notes:

- Prior to GigaVUE-OS version 4.7, GigaVUE TA Series Traffic Aggregator nodes did not support tool ports. Instead, they supported gateway ports. Starting in GigaVUE-OS version 4.7, all gateway ports on GigaVUE TA Series nodes are tool ports. However, GigaVUE-FM may be managing nodes running a software version lower than version 4.7.
- G-Series nodes are not supported in Topology Visualization.

Enable Discovery Protocols

This section describes the following details:

- *Enable Gigamon Discovery on Chassis* on page 300
- *Enable LLDP, CDP, and Gigamon Discovery on Ports* on page 300

Enable Gigamon Discovery on Chassis

Gigamon discovery is disabled on chassis and ports. The Gigamon discovery packets are transmitted only when Gigamon Discovery is enabled on chassis as well ports.

To enable Gigamon discovery on chassis:

1. Click **Physical** on the top navigation link.
2. In the left navigation pane, click **Physical Nodes**.
3. Click on a node on which you want to enable Gigamon Discovery. The Overview page of the node is displayed.
4. In the left navigation pane, click **Chassis**.
5. Switch the Chassis view to List View.
6. Under Cards, find the line card on which you want to enable Gigamon Discovery and select the card.
7. Click **Actions** and select **Enable Gigamon Discovery**.

   The Enabled option is displayed under Gigamon Discovery for the selected Chassis ID.

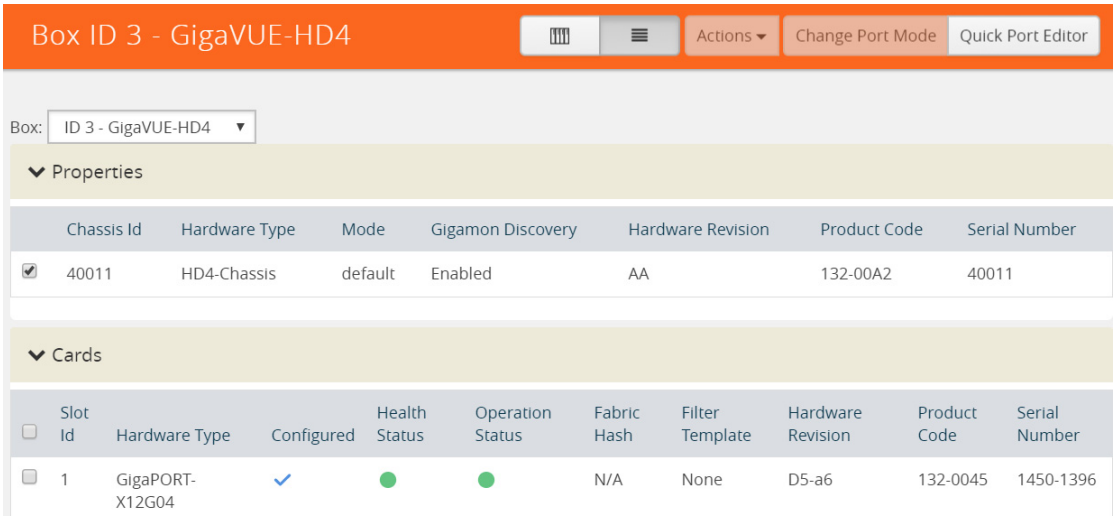

Figure 18-1: Gigamon Discovery Enabled

Enable LLDP, CDP, and Gigamon Discovery on Ports

To enable LLDP, CDP, and Gigamon Discovery on ports:

1. Click **Physical** on the top navigation link.
2. In the left navigation pane, click **Physical Nodes**.
3. Click on a node on which you want to enable LLDP and CDP. The Overview page of the node is displayed.
4. In the left navigation pane, select **Ports > All Ports**.
5. On the Ports page, select the Port ID on which you want to enable LLDP, CDP, and Gigamon Discovery. Click **Edit**.

6. Under **Ports Discovery**, do the following:
   a. To enable network discovery, select the **Enable** check box for Network Discovery**.**
   b. To enable discovery protocols, select one of the following: **All**, **LLDP**, or **CDP**.
   c. To enable Gigamon discovery, select the **Enable** check box for Gigamon Discovery**.**
7. Click **OK**.
8. Click **Save**.
9. Click the **Ports Discovery** tab. For each network port on which the ports discovery is enabled, the neighbor information is displayed.

Limitations of Gigamon Discovery

Following are the limitations of Gigamon discovery:

- Gigamon discovery is not supported on inline-tool and inline-network port type.
- For stack ports, Gigamon discovery will work only if both ends of the link are of type stack.
- If a pass-all map is configured for Gigamon discovery enabled ports, the show map stats command will show an increment in the map rule counter for Gigamon discovery packets although the Gigamon discovery packets are not sent to the tool ports.
- If Gigamon discovery is enabled on a port, for example port 9/3/x1, and the port type is changed to an unsupported port types such as inline-network, Gigamon discovery is immediately disabled and the following message is displayed:

  (config) # ! GDP is not supported with port type inline-network. Disabling gdp on port 9/3/x1

  Similarly, if the port type is changed to inline-tool, the following message is disabled:

  (config) # ! GDP is not supported with port type inline-too. Disabling gdp on port 9/3/x1
- The hybrid ports, tool-mirror source ports, and ports configured with force link up are operated in loopback mode. So, when Gigamon discovery is enabled, these port show themselves as neighbors.

View Clusters and Nodes in the Topology

To see the clusters and nodes managed by GigaVUE-FM, click **Physical** on the top navigation link and then click **Topology** in the left navigation pane. The Gigamon nodes and other devices are represented by graphical icons. Connections between the nodes are represented by lines. Clusters appear as orange containers. Devices can be either discovered or manually added to the topology. Figure 18-2 shows an example star topology with three Gigamon clusters, Gigamon nodes, and other devices.

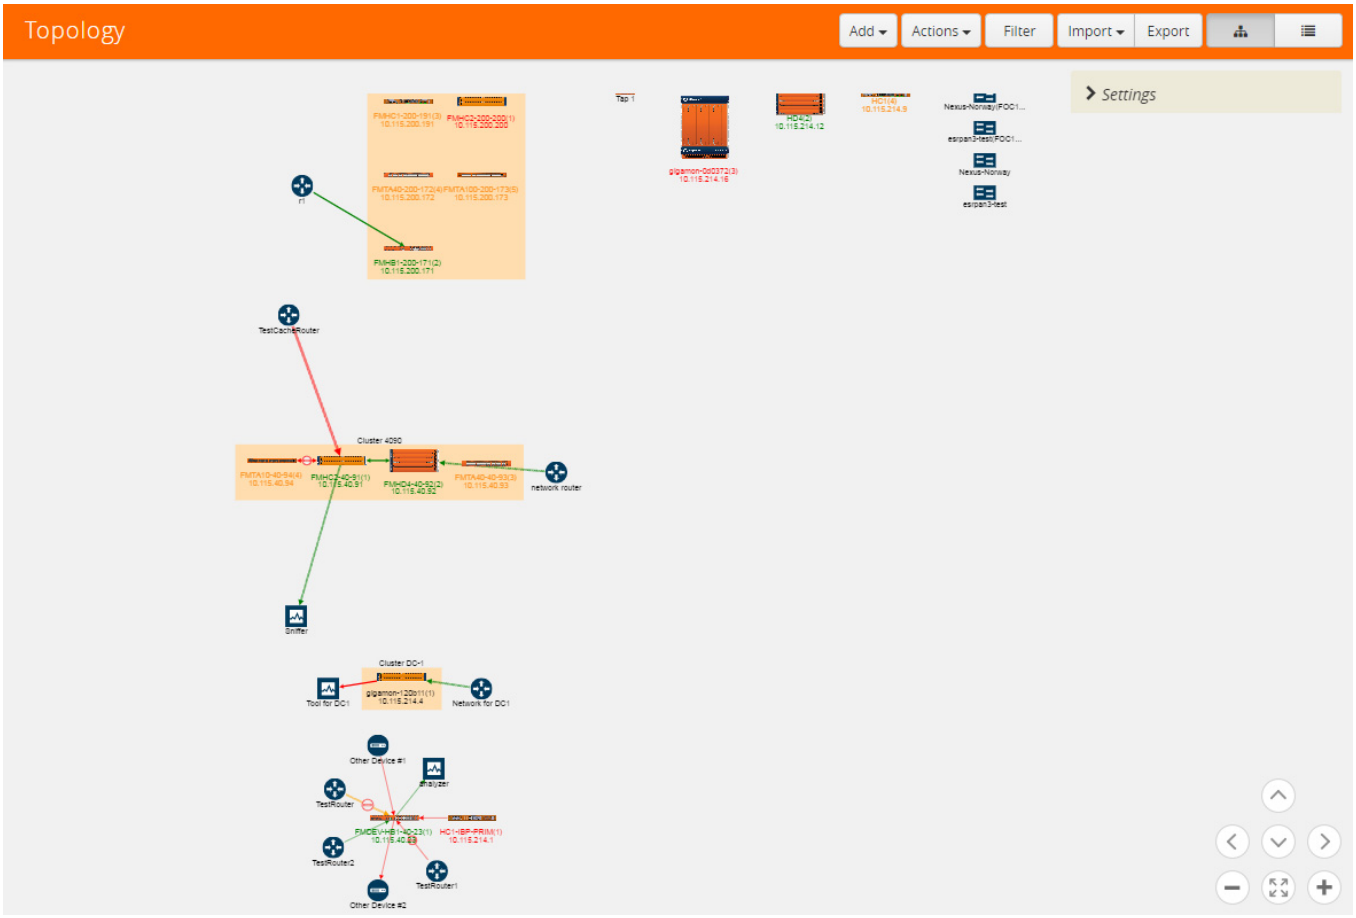

Figure 18-2: GigaVUE-FM Topology

Figure 18-2 shows an example of leaf and spine topology.

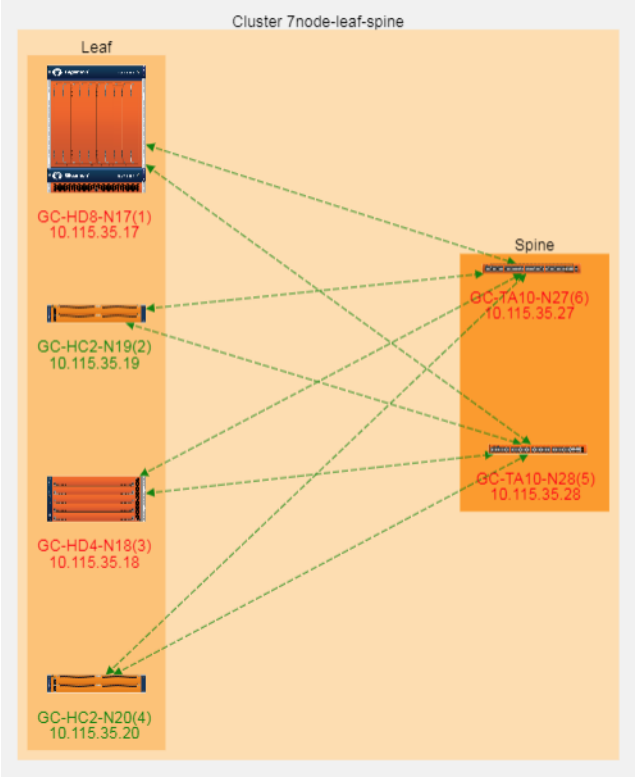

Figure 18-3: Leaf and Spine Topology

There are two kinds of nodes displayed in the leaf and spine topology, as follows:

- Leaf nodes, which are edge nodes and can also have TAPs or tools attached to them. All the leaf nodes are displayed within the Leaf container.
- Spine nodes, which are the nodes to which the leaf nodes attach. All the spine nodes are displayed within the Spine container.

The leaf and spine configuration can be done only by using the CLI commands. For more information about multi-path leaf and spine architecture and configuration, refer to *GigaVUE-OS CLI User's Guide* on the Gigamon Customer Portal.

The Topology page has the following menus and buttons:

Menu/Button	Description
Add	The Add menu has the following options: • Add Device(s)—opens an Add Device(s) page for selecting and manually adding nodes to the topology. For details, refer to *Add Devices* on page 325 • Add Link(s)—opens an Add Link(s) page for adding links between nodes. For details, refer to *Add Links* on page 330 • Add Tap—opens an Add Tap page for adding Gigamon TAPs to the topology. For details, refer to *Add TAP* on page 328.

Menu/Button	Description
Actions	The Actions menu has the following options: • Edit—opens the Edit page for a manually added node or link. The node or link must be selected for this option to be available. • Delete—deletes a manually added node or link. • Delete All—deletes all manually added nodes and links. • Reset Alignment—resets the alignment of clusters and nodes to the default alignment. • View Node Topology—shows the topology of a selected node. • Go to Node Overview—open the Overview page for the selected node.
Filter	The Filter button opens a Quick View for selecting filter criteria for the topology. The filter makes it possible to select only those items that you want to view in the topology, such as nodes of a certain type or specific models of Gigamon nodes.
Import	The Import menu supports the importing of custom nodes and links from an Excel spread sheet. For details, refer to *Export and Import Topology* on page 331.
Export	The Export menu supports the exporting of custom nodes and links to an Excel spread sheet. For details, refer to *Export and Import Topology* on page 331.
Topology View	The Topology View button switches the Topology view to the graphic representation of Gigamon nodes, links, and the various types of nodes. Topology is the default view.
Table View	The Table View button switches the view to a table that lists all the items in the topology. For details about Table View, refer to *Table View* on page 316.

The controls in the lower right-hand corner are used to move the topology up, down, left, right, and to zoom in or out. Clicking the center button makes a *best fit* of the topology on the page.

A context menu is available for clusters and nodes that includes options for realigning elements. Right-click on a cluster or a node to open the context menu, which stays open until you click elsewhere on the page. The options are as follows:

Menu Option	Description
View Node Topology	Opens the Node Topology page for the selected cluster or node. For more information about Node Topology, refer to *Node Topology* on page 333. You can also open the Node Topology by selecting Physical Nodes in the top Navigation pane, selecting the node or cluster, and then selecting Node Topology.
Go to Node Overview	Opens the Overview page for the selected node. You can also open the Overview page for the node by selecting Physical Nodes in the Navigation pane, and then selecting the node or cluster.
Align Neighbor	Aligns all the neighbors of the selected node or cluster. For details, refer to *Align Neighbors* on page 320.
Align Leaf Neighbor	Aligns only the immediate neighbors of the selected node or cluster. For details, refer to *Align Leaf Neighbor* on page 322.
Align Cluster	Aligns the nodes within a cluster container. For clusters with leaf and spine configuration, the leaf and spine nodes are aligned within their respective containers. For details, refer to *Align Cluster* on page 322.

Menu Option	Description
Select Neighbor	Selects the cluster container or node and all of its immediate neighbors. The node or clusters and its neighbors can be moved as a unit. For details, refer to *Select Neighbor* on page 325.

Filter

Clicking the Filter button opens a Filter quick view. For topologies with large numbers of elements, the filter is useful for displaying only those items of immediate interest. The Filter quick view has the following options:

Filter Option	Description
Discovery	Display discovered nodes by selecting one or more of the following (all are selected by default): • Gigamon—displays Gigamon nodes • Manual—displays manually added nodes • CDP/LLDP—displays nodes discovered through CDP or LLDP **NOTE:** If both CDP and LLDP packets are discovered on the same port, the topology construction relies on LLDP packets over CDP. • Taps—display TAPs
Device Types	Display only nodes of the selected type, such as router, switch, virtual switch, load balance, and so on. You can select more than one type.
Vendor	Display manually added nodes that are from the specified vendor. Enter the vendor's name in the Vendor field.
Models	Display only the specified Gigamon model. You can select more than one model.
Cluster IDs	Display the clusters or nodes with the specified IDs. The drop-down list displays the IDs of the clusters and nodes in the topology. You can select more than one node or cluster.
Node Status	Display nodes with the specified status. The possible statuses are OK, Error, and Warning.

In Figure 18-4, all Discovery filters are selected and two clusters are specified in the **Cluster ID(s)** field. Only items that meet the filter criteria are displayed while all other elements are grayed out.

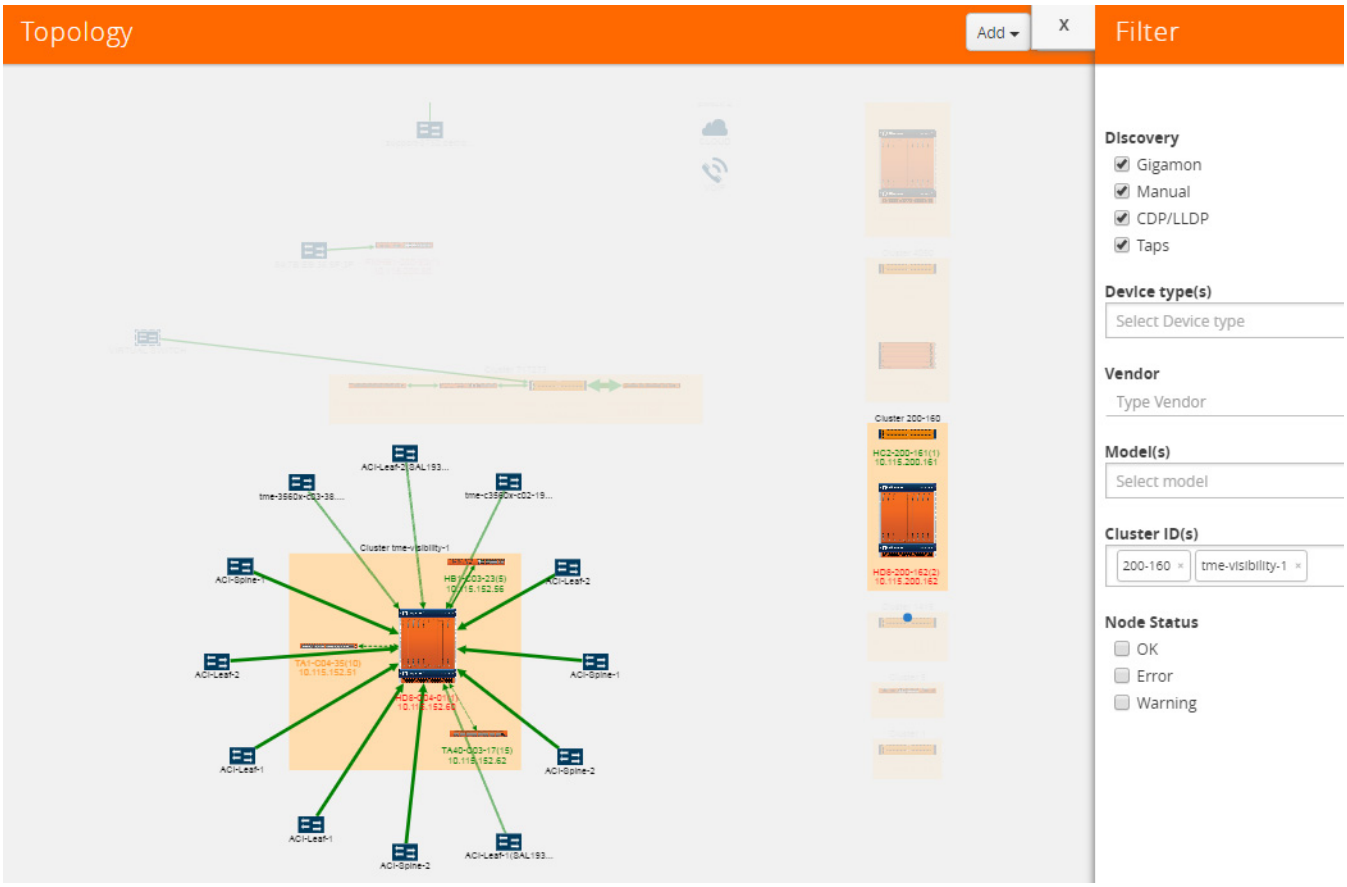

Figure 18-4: Filter on Cluster IDs in Topology

Links

Links between nodes are shown as arrows between nodes. The links are automatically created for nodes in clusters and connections to discovered neighbor nodes. The links are also discovered through LLDP, CDP, or Gigamon Discovery. The LLDP links represent a Gigamon-to-LLDP connection. The CDP links represent a Gigamon-to-CDP connection. The Gigamon Discovery links represent a Gigamon-to-Gigamon connection. The thickness of the link indicates the link's bandwidth. A dashed line represents an aggregated link. Inside a cluster container, a dashed line represents a GigaStream stack link between nodes and the tooltip for the link will provide information about the GigaStream. For more information about link aggregation, refer to *Aggregated Links* on page 310.

Link State

Links between nodes in the topology have a color to indicate the link's current state. The link state is also identified in the tooltip information that displays when clicking on a link. The following are the link state colors and their meanings:

Color	Link State
Green	Up (connected)
Red	Down (disconnected), unreachable, or unknown The link state will be Red (or down) in an aggregated link if all the link states of the associated links of the aggregated link are down. In an aggregated link, the link state is unknown if all the link states of the associated links in the aggregated link are unknown.
Amber	Warning This link state only applies to aggregated links. If any one of the links in the aggregated link is down or unknown, the link state is rendered as amber.

Red indicates that the link is either down or its status is unreachable. When the link is down, a stop icon also displays on the link. For example, in the following figure, the link color (red) and the stop icon on the link from TestRouter1 to the GigaVUE-HB1 indicates that it is down. The link from the GigaVUE-HB1 to Other Device #2 is unknown because it is red and there is no icon on the link.

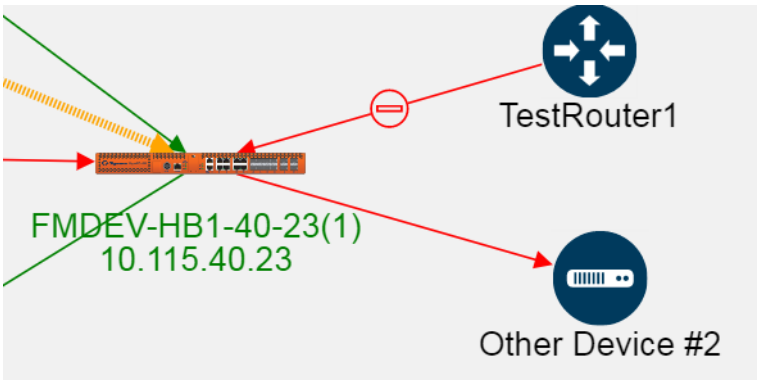

When the link status is between two Gigamon nodes, the link status is determined as follows:

- If the port for the link associated with the node is up, the link status is up or connected (green).
- If the port for the link associated with the node is down or unknown, the link status is down or disconnected (red) and a stop icon is displays on the link. However, if there is no stop icon, the status is unknown.

When the link is between a Gigamon node and an LLDP, CDP neighbor, another GigaVUE node, or a manually created node, the link status is determined by the link status of the Gigamon node as follows:

- If the port status is up, the link status is connected (green).
- If the port status is down or unknown, the link status is disconnected (red).

Link Health State

In addition to a link state, links have a health state based on link utilization. The link's health state is computed only when the link state is up. The health status is indicated in the tooltip when clicking on a link, which shows an LED bubble. The colors indicated the following

Color	Link Health State
Green	Normal utilization.
Red	The link is either underutilized or overutilized
Amber	At least one of the links in an aggregated link is either underutilized or overutilized.
Gray	Health state is not available.

If the health state of the link is not within the normal range of utilization based on the low and high utilization thresholds, a warning icon displays on the link.

Link Utilization Thresholds

The health state of a link is based on the utilization of the link. The utilization thresholds are specified through the settings dialog. The threshold values apply to all nodes managed by GigaVUE-FM. They should not be confused with port level thresholds.

To open the dialog, click **Settings** in the upper-right corner of the topology.

The Link Utilization Thresholds are as follows:

- **Low**—sets the lower limit of the utilization threshold. If the link's utilization falls below the specified percentage, the link is considered underutilized. The default is 40 percent utilization for a low utilization threshold.
- **High**—sets the upper limit of the utilization threshold. If the link's utilization rises above the specified percentage, the link is considered overutilized. The default is 75 percent utilization for a high utilization threshold.
- **Interval**—the time period (in minutes) over which the link utilization is measured. The time interval is used to determine utilization based on an average value of port utilization over the specified interval. The default is 15 minutes.
- **Utilization Overlay**—Turns on/off the utilization warning icons. An icons appear on the link if it is either overutilized or underutilized. To turn on the overlay, slide the switch to **on** and click **OK**. Figure 18-5 on page 310 shows an example where the utilization overlay is turned on. The Utilization Overlay setting only applies to the current session. If you log out of GigaVUE-FM and then log in again, the Utilization Overlay is set to **off**.

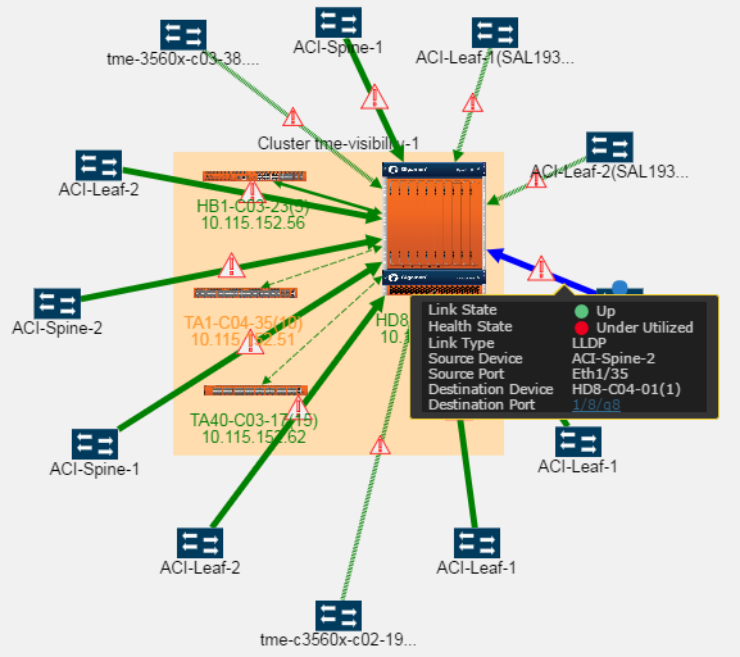

Figure 18-5: Overutilization Overlay

Aggregated Links

A node in the topology can have more than one link to another node. However, these links display in the topology as a single dashed line. In Figure 18-6, the link between the node tme-c3560x-c02-19 and the GigaVUE-HD8 node in the cluster is an aggregated link. The link state or link health state of the links within the aggregation determine the state of the aggregated link. For more information, refer to *Link State* on page 308 and *Link Health State* on page 309.

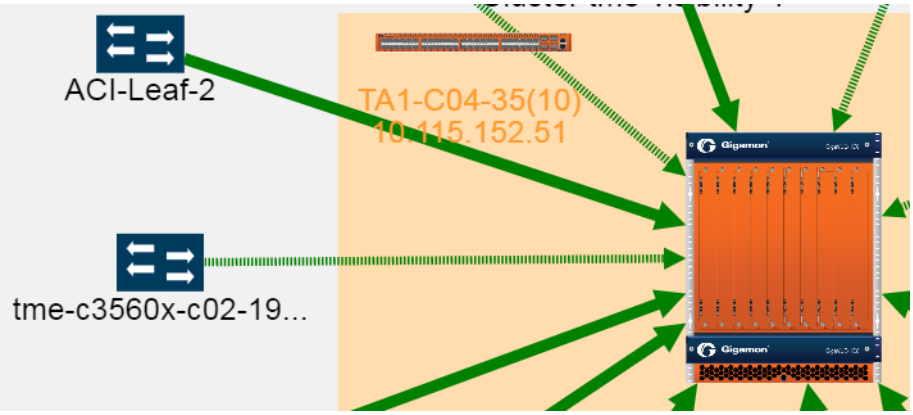

Figure 18-6: Aggregated Link

To view the links in the aggregation, right-click in the cluster container to open the context menu and select **View Node Topology**. Refer to Figure 18-7.

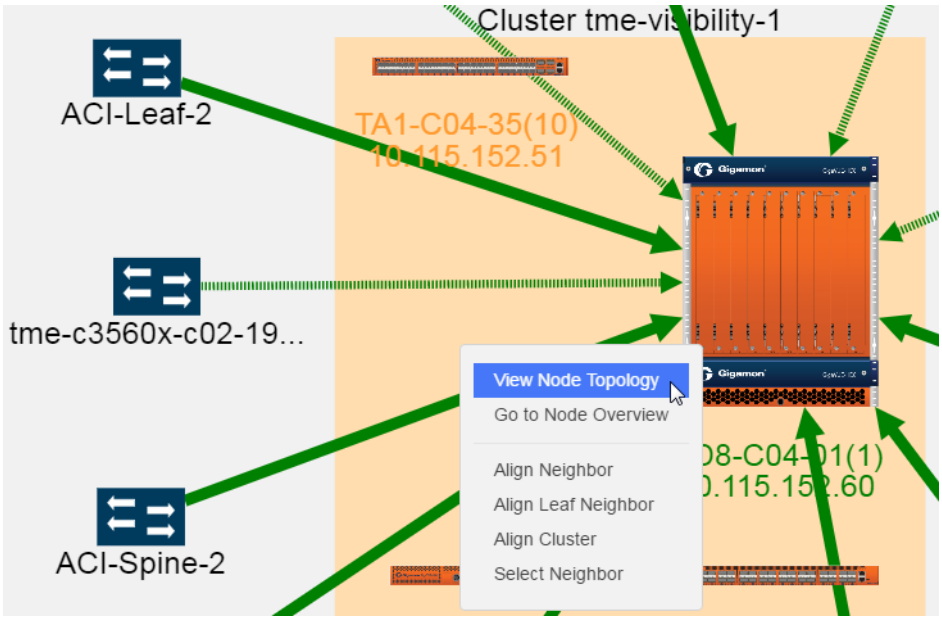

Figure 18-7: View Node Topology Selected

Selecting View Node Topology, opens the Node Topology for the selected cluster, where the links are not aggregated. Figure 18-8 on page 311 shows the cluster in Figure 18-7 on page 311 as it appears in Node Topology with all the links unaggregated. For more information about Node Topology, refer to *Node Topology* on page 333.

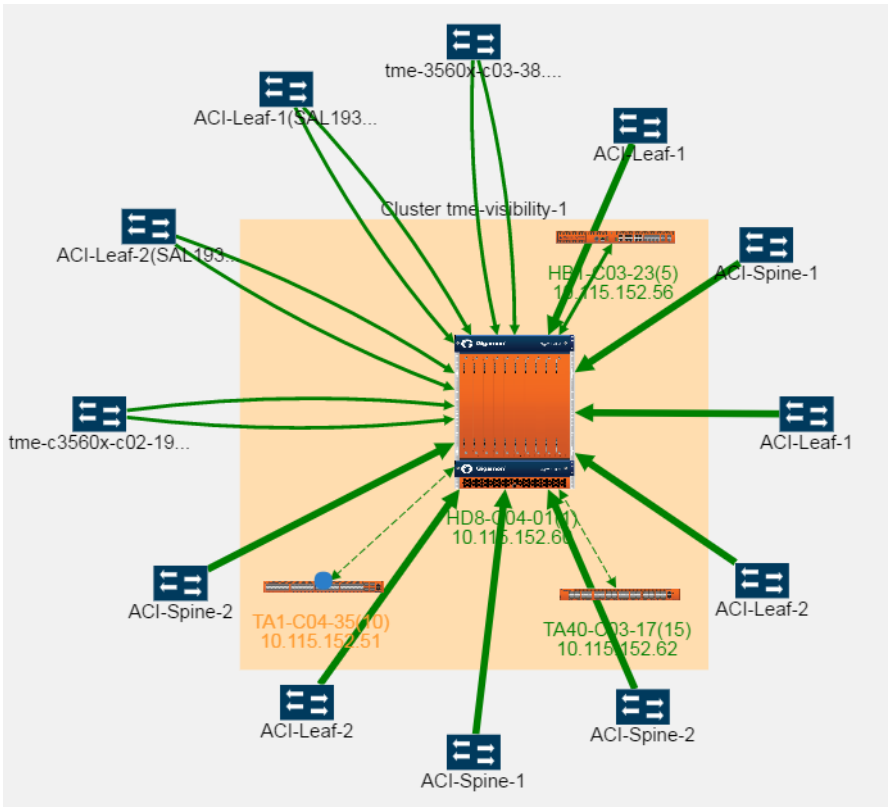

Figure 18-8: Node Topology

To return to the GigaVUE-FM Topology page, right-click on the cluster or node to open the context menu and select **View FM Topology**.

GigaStream

GigaStream links display in the topology as dashed lines within a cluster. Clicking on the link opens a tooltip that provides the source and destination GigaStream of the link. Figure 18-9 shows an example.

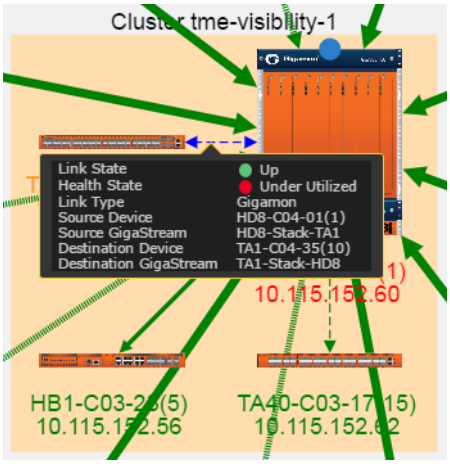

Figure 18-9: Tooltip for a GigaStream

Link Details

To view details about a link, click on the link and select **More Details** from the context menu. A Link Details quick view opens with details about the link. If the link is aggregated, Link Details shows information for each link in the aggregation. Figure 18-10 on page 312 shows the Link Details for an aggregated link with two links.

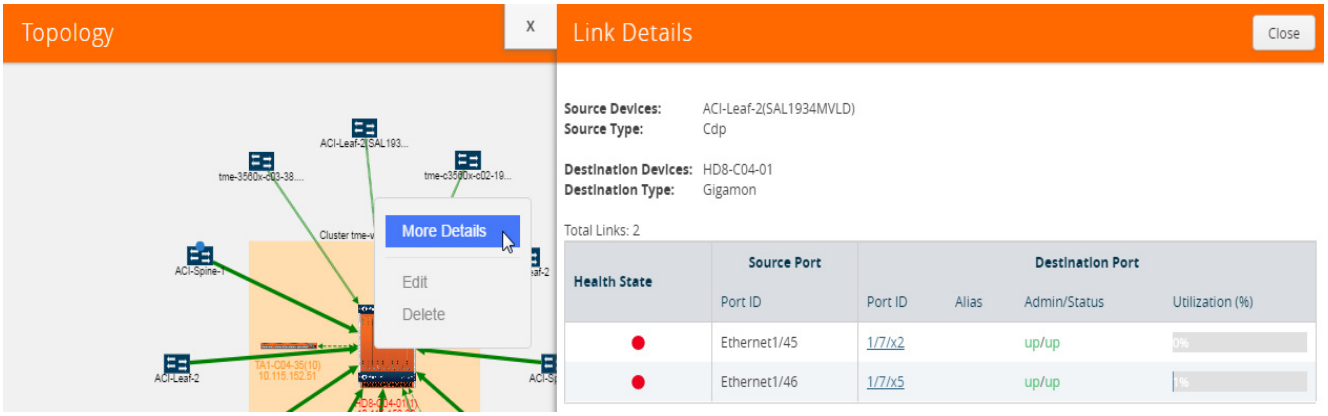

Figure 18-10: Link Details for an Aggregated Link

Link Details provides the following information:

Field	Description
Source Devices	The source nodes for the link.

Field	Description
Source Type	The type of source node: CDP, LLDP, Gigamon, or Manual
Destination Devices	The destination nodes for the link.
Destination Type	The type of destination node: CDP, LLDP, Gigamon, or Manual.
Total Links	Total number of links in the aggregation.
Health State	An LED bubble indicates the current health state of the port. Hovering over the bubble displays a tooltip with the current state of the link. For a description of the health states, refer to *Link Health State* on page 309.
Source Port	The details for the source port, such as the port ID of the link's source port. The exact details vary depending on the type of node.
Destination Port	The details for the destination port, such as the port ID of the link's destination port, alias (if any), admin and link status of the port, and the percentage of utilization displayed as a progress bar. The exact details vary depending on the type of node.

GigaStream Stack Link Details

If the link is a GigaStream stack link connecting a leaf and a spine node, Stack Link quick view is displayed. The quick view provides stack link information and GigaStream status with details such as GigaStream alias and total number of ports.

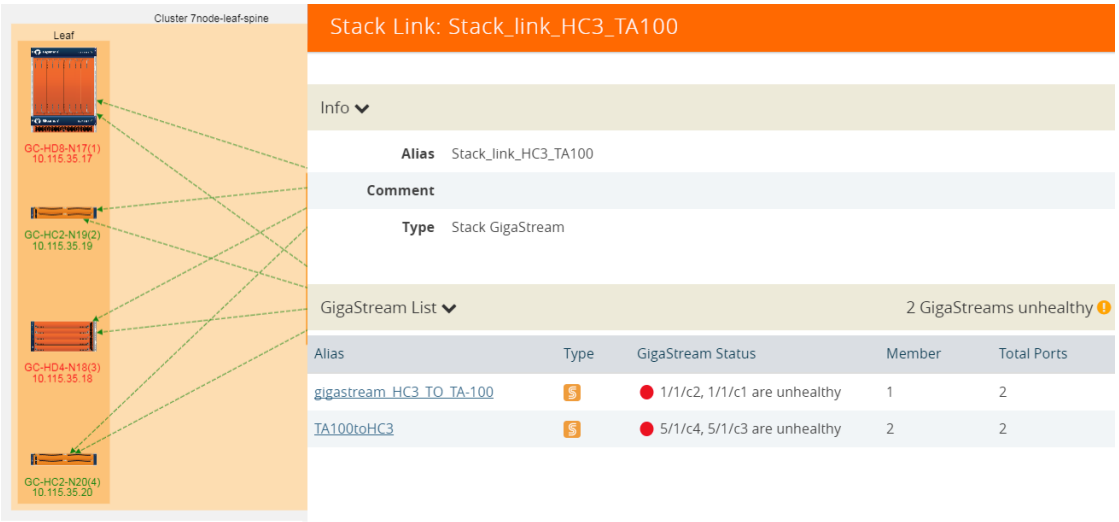

Figure 18-11: Link Details for a Stack Link

The following table describes the parameters displayed in the GigaStream stack link quick-view:

Field	Description
Info	Displays the stack link information such as the name of the stack link, component, and type.
GigaStream List	
Alias	Displays the stack GigaStream name.
Type	Displays the type of the port.

Field	Description
GigaStream Status	Displays the status.
Member	Displays the ranking of the GigaStream member in the stack GigaStream link. The number '1' represents the first GigaStream member configured. The number '2' represents the second GigaStream member configured.
Details	Displays the health status of the stack link.
Total Ports	Displays the total number of ports that connects the nodes.

Discovered Devices

Discovered nodes are displayed in gray containers when there are multiple neighbors. However, only one node displays in the container. The actual number of neighbor nodes is displayed in a label above the container and **More...** displays inside the container. Clicking on the container displays additional information. Figure 18-12 shows a neighbor node container with seven discovered nodes that are neighbors to port 1/4/x3 on the node SHC2C.

Figure 18-12: Discovered Devices

Discovered nodes, including CDP, LLDP, and Gigamon discovery, display in the topology as switch icons. To determine the discovery type, click on the node icon to display a tooltip that provides information about the node. The Discovered Type field will show whether the node was discovered through CDP, LLDP, or Gigamon Discovery. If both CDP and LLDP packets are discovered on the same port, the topology construction relies on LLDP packets over CDP. In Figure 18-13 the selected node shows tooltip that indicates that the node was discovered through LLDP.

NOTE: If the link represents a link to a SPAN port, the link does not display a context menus when clicked.

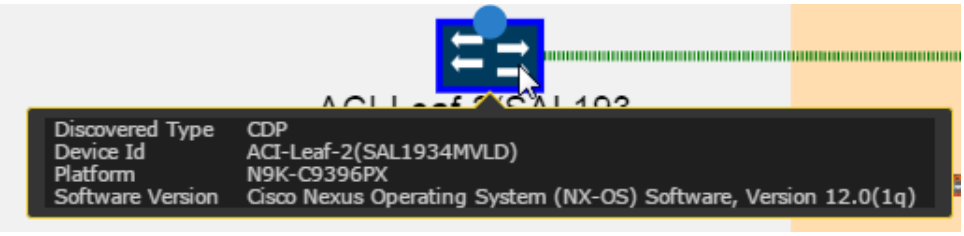

Figure 18-13: Tooltip for a Discovered Device

To view more information about discovered nodes in a discovered node container, click **More**. This opens a quick view that shows the information about the neighbor nodes. Figure 18-14 shows the quick view for the neighbor node container shown in Figure 18-12.

Discovered Devices/Ports on SHC2C (1) Port 1/4/x3

Source Device	Source Port	Source Type	Destination Device	Destination Port	Destination Type
HQ-L3SW01-2	GigabitEthernet0/6	cdp	10.115.94.16	1/4/x3	gigamon
HQ-L3SW01-2	GigabitEthernet0/9	cdp	10.115.94.16	1/4/x3	gigamon
HQ-VL10-2	GigabitEthernet0/45	cdp	10.115.94.16	1/4/x3	gigamon
Core02-1	GigabitEthernet1/2/3	cdp	10.115.94.16	1/4/x3	gigamon
HQ-VL10-1	GigabitEthernet0/47	cdp	10.115.94.16	1/4/x3	gigamon
HQ-WiFi-VL40-8	GigabitEthernet1/0/48	cdp	10.115.94.16	1/4/x3	gigamon
Core02-1	GigabitEthernet2/2/3	cdp	10.115.94.16	1/4/x3	gigamon
HQ-L3SW01-2	GigabitEthernet0/10	cdp	10.115.94.16	1/4/x3	gigamon
HQ-L3SW01-2	GigabitEthernet0/1	cdp	10.115.94.16	1/4/x3	gigamon
HQ-VL10-1	GigabitEthernet0/48	cdp	10.115.94.16	1/4/x3	gigamon
HQ-VL10-1	GigabitEthernet0/46	cdp	10.115.94.16	1/4/x3	gigamon
HQ-WiFi-DMZ	GigabitEthernet1/0/1	cdp	10.115.94.16	1/4/x3	gigamon
HQ-VL10-2	GigabitEthernet0/48	cdp	10.115.94.16	1/4/x3	gigamon
HQ-VL10-1	GigabitEthernet0/45	cdp	10.115.94.16	1/4/x3	gigamon
HQ-L3SW01-2	GigabitEthernet0/5	cdp	10.115.94.16	1/4/x3	gigamon
HQES02(SSI14291AGZ)	Ethernet1/40	cdp	10.115.94.16	1/4/x3	gigamon
HQ-WiFi-VL40-8	GigabitEthernet1/0/47	cdp	10.115.94.16	1/4/x3	gigamon
HQ-VL10-2	GigabitEthernet0/46	cdp	10.115.94.16	1/4/x3	gigamon
HQ-VL10-2	GigabitEthernet0/47	cdp	10.115.94.16	1/4/x3	gigamon

Figure 18-14: Discovered Device Quick View

The container label indicates that there are seven neighbor nodes. However, the quick view lists more than seven nodes. This is because discovery is based on the connections between ports and a port may be connected to more than one node, resulting in the node being listed more than once in the quick view.

Table View

The default view is Topology Visualization. You can switch to Table View by clicking the Table View button. In Table View, you can select different views from the **View** menu. Table 18-1 describes each of the menu options.

Table 18-1: Table Views

View By	Description
All	Lists all the items in the topology. The table is read-only.This is the default table.
CDP	Lists only nodes in the topology discovered through CDP. The information includes Device ID, Platform, and Software Version. This table is read-only.
Gigamon	Lists only Gigamon nodes. The information includes the Device ID (IP address of the node), Host Name, Box ID, Model, Cluster ID, and Cluster Mode (Standalone, Master, Slave, or Standby). Clicking on the IP address opens that Overview page for that node. This table is read-only.
LLDP	Lists only nodes in the topology discovered through LLDP. The information includes Chassis ID, System Name (if available), and Description (if available). This table is read-only.
Manual	Lists only manually added nodes in the topology. The information includes the Name, Vendor, Type, Mode, and Comment (if any). Each item has a checkbox. Selecting a checkbox for an item enables **Edit** and **Delete** in the **Actions** menu.
Links	Lists only the links in the topology. The information includes the Source Device, Source Port on the node, Source Type (CDP, LLDP, Gigamon Discovery, Gigamon, or Manual), Destination Device, Destination Port on the node, and Destination Type. Manually added links have a checkbox. Selecting a checkbox for an item enables **Edit** and **Delete** in the **Actions** menu.

How to Customize Topology

The Topology page automatically displays the Gigamon nodes managed by GigaVUE-FM and neighbor ports that are discoverable through LLDP, CDP, or Gigamon Discovery. However, you can customize the topology by adding custom nodes and links or modifying the alignment of items on the page.

Default Alignment

The first time Topology opens, GigaVUE-FM presents the elements in a predetermined alignment. The elements are aligned in general columns according to the following order from left to right:

1. Clusters and standalone nodes with links.
2. Manually added network nodes
3. Manually added TAPs
4. Clusters without links.
5. Standalone Gigamon nodes. These are ordered by dimension:
    d. GigaVUE TAPs
    e. GigaVUE-HD8
    f. GigaVUE-HD4
    g. GigaVUE-HC3
    h. GigaVUE-HC2
    i. GigaVUE-HD2
    j. GigaVUE-HD1
    k. GigaVUE TA Series
6. Manually added tool nodes

You can restore the default alignment at any time by selecting **Actions > Reset Alignment**.

Move Elements in Topology

The Topology provides a context menu that makes it possible to quickly move elements on the page. To open the context menu, select an element and then right-click. For a traditional cluster configuration, Figure 18-15 on page 318 shows the context menu with a number of options for aligning the elements in the topology. These options are:

- Align Neighbor
- Align Leaf Neighbor
- Align Cluster
- Select Neighbor

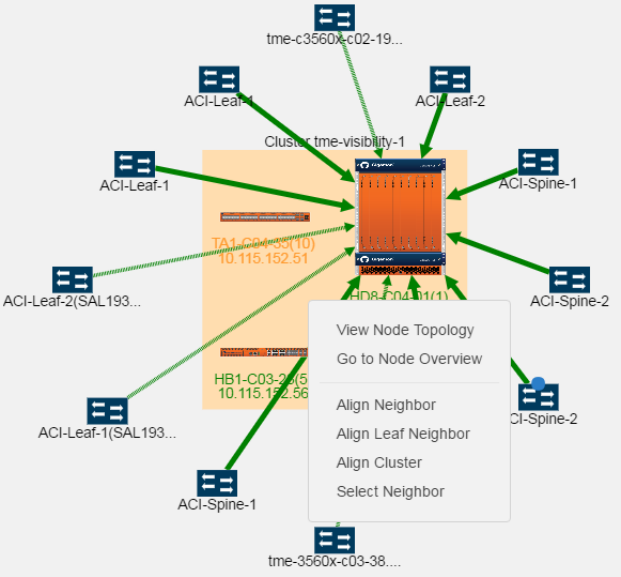

Figure 18-15: Context Menu Displayed for Traditional Configuration

For a leaf and spine topology, following options are displayed in the context menu (refer to Figure 18-16 on page 319):

- Align Neighbor
- Align Leaf Neighbor
- Align Cluster
- Align Cluster Spine
- Align Cluster Leaf

- Select Neighbor

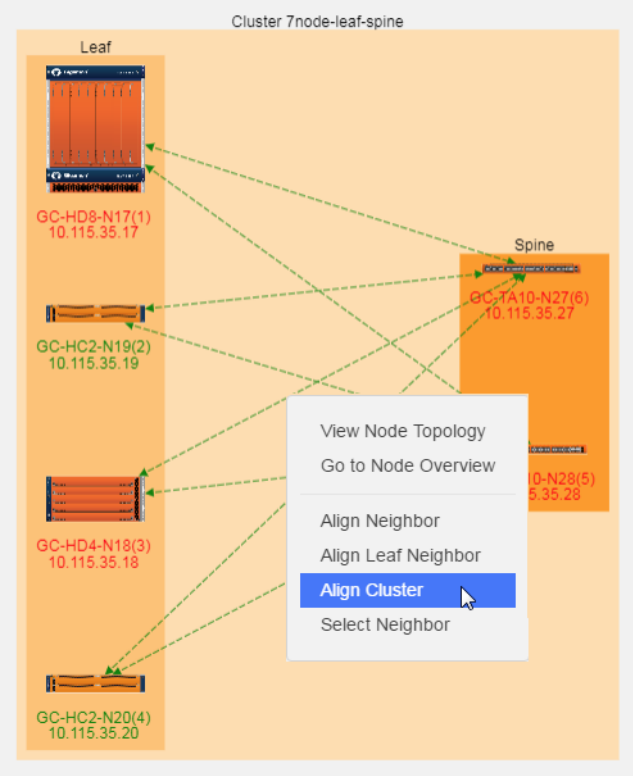

Figure 18-16: Context Menu Displayed for Leaf and Spine

Resize Links

When the topology is first created, a link may appear too long or too short. You can change the length of the link by selecting the node at either end of the link and dragging the node so that the link is the length that you want. For example, in Figure 18-17, the link between the node named Packet Capture #1 and the node named HC2-Longevity in the cluster is too long. To make it shorter, select the node and drag it closer to the cluster.

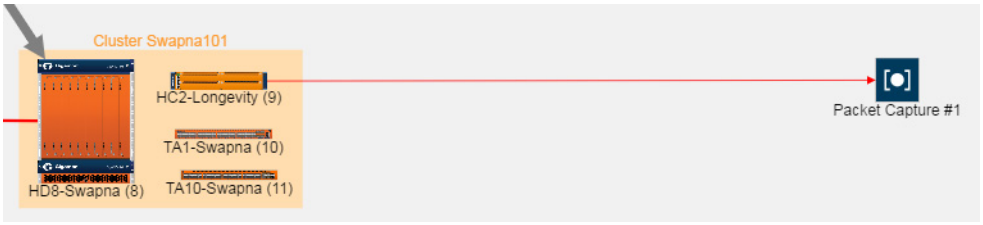

Figure 18-17: Long Link Between Devices

Figure 18-18 show the topology after dragging the node to adjust the length of the link.

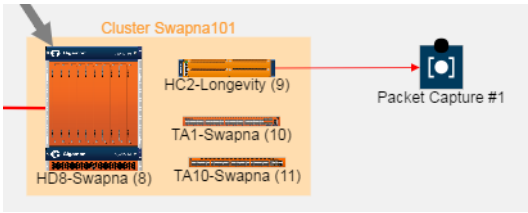

Figure 18-18: Length of Link Adjusted

In addition to dragging individual items on the page, you can use the Align Neighbor and Align Leaf Neighbor options from the context menu. For details, refer to *Align Neighbors* on page 320 and *Align Leaf Neighbor* on page 322.

Align Neighbors

The **Align Neighbor** option aligns all the neighbors of a selected node and any neighbors of those neighbors. The alignment follows the path of all the linked neighbors, aligning neighbors equidistant from the central neighbor based on a 360 degree arc. In Figure 18-19 on page 320, the GigaVUE-HD8-C04-01 has 12 neighbors. The neighbors are arranged at 30 degree angles around GigaVUE-HD8. If any of the neighbors have linked neighbors, the alignment of those neighbors are also calculated in the same manner. In Figure 18-19 on page 320, the linked neighbors are on one side of the cluster and all neighbors are on the left side of the cluster.

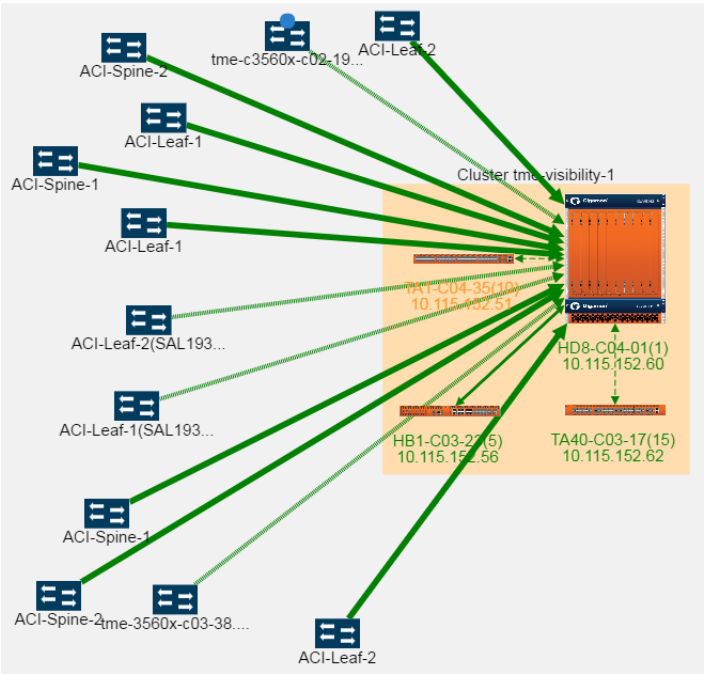

Figure 18-19: Aligning Neighbors

To realign the neighbors, do the following:

1. Select the cluster
2. Right-click to open the context menu and select **Align Neighbor**.

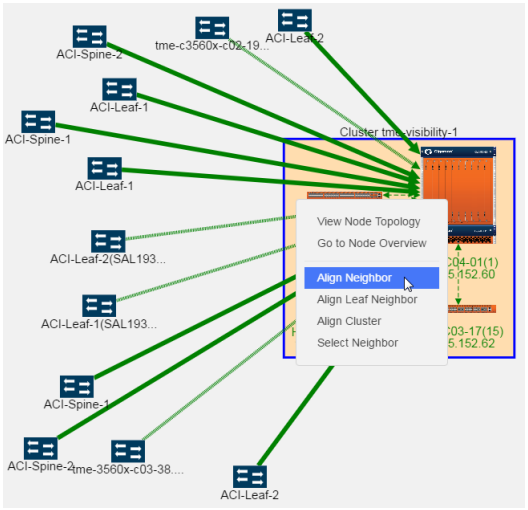

Figure 18-20: Displaying Align Neighbor Option

The neighbors are realigned around the selected element as shown in the following figure.

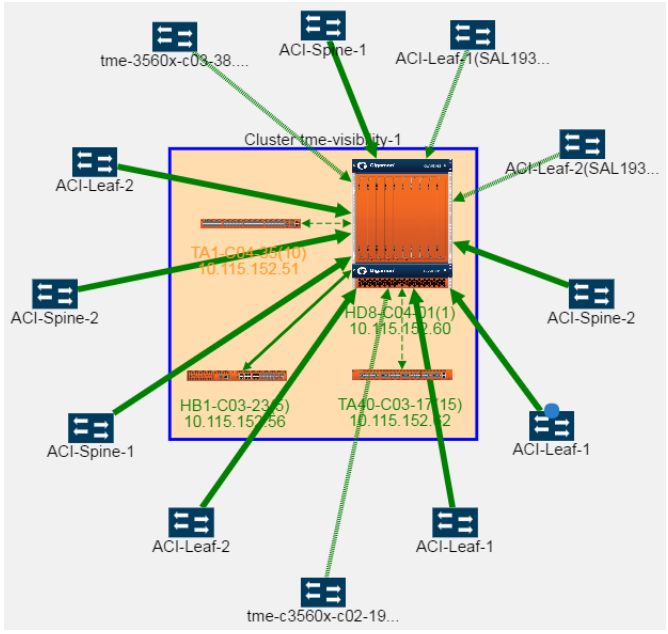

Figure 18-21: Nodes Aligned after Selecting Align Neighbor Option

When there are multiple neighbors, the alignment propagates through the topology. The neighbors of the selected element are aligned, then the neighbors of those neighbors, and so on throughout the topology. This makes it possible to realign an entire topology very quickly. In Figure 18-22, the right-most element was selected and alignment propagation flows to the left and throughout the topology.

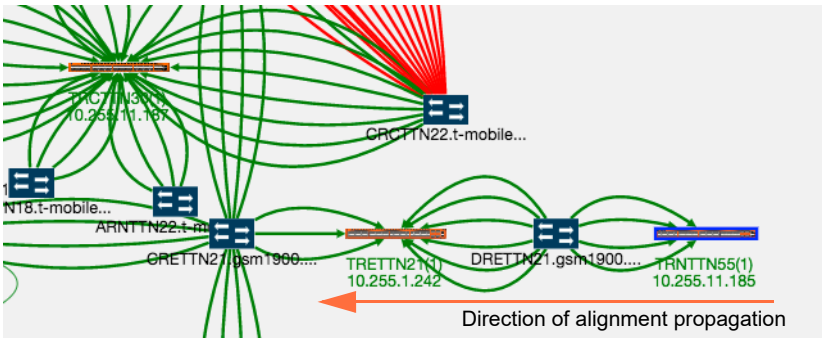

Figure 18-22: Neighbor Alignment Propagation

Align Leaf Neighbor

The **Aligning Neighbor** realigns all neighbors linked to the selected node and alignment is propagated throughout the topology. To align only the immediately adjacent neighbors, use **Align Leaf Neighbor**. The alignment will occur with the immediate neighbors and alignment will not propagate throughout the entire topology.

Align Cluster

Nodes within a cluster container can be repositioned by selecting and dragging individual nodes or using the **Align Cluster** option. To select this option, right-click on the cluster container and select **Align Cluster Spine** as shown in Figure 18-23.

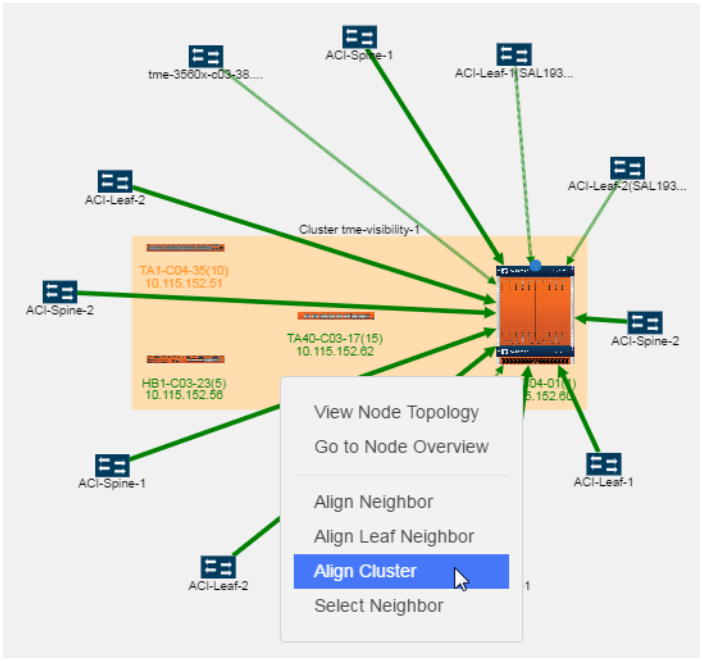

Figure 18-23: Align Cluster Option Selected

After selecting the Align Cluster option in Figure 18-23, the nodes in the cluster are rearranged. Figure 18-24 on page 323 shows the results of applying **Align Cluster** to the cluster in Figure 18-23.

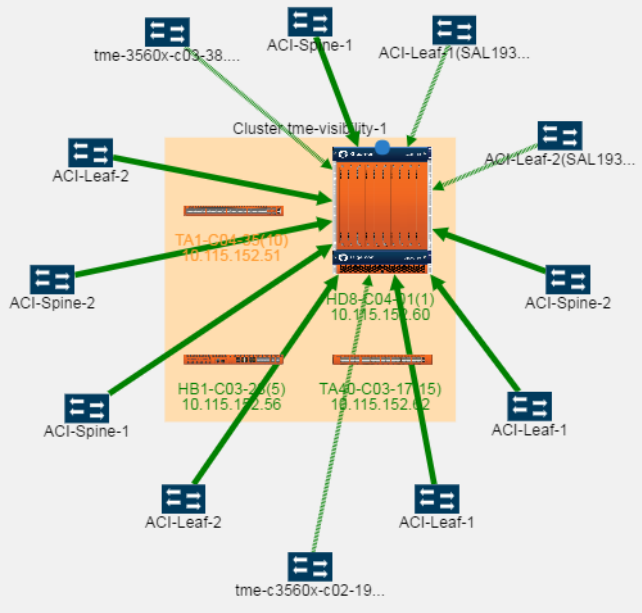

Figure 18-24: Nodes in a Cluster Container Arranged Horizontally

For leaf and spine cluster configuration, the Align Cluster option aligns the spine nodes within the spine container and leaf nodes within the leaf container.

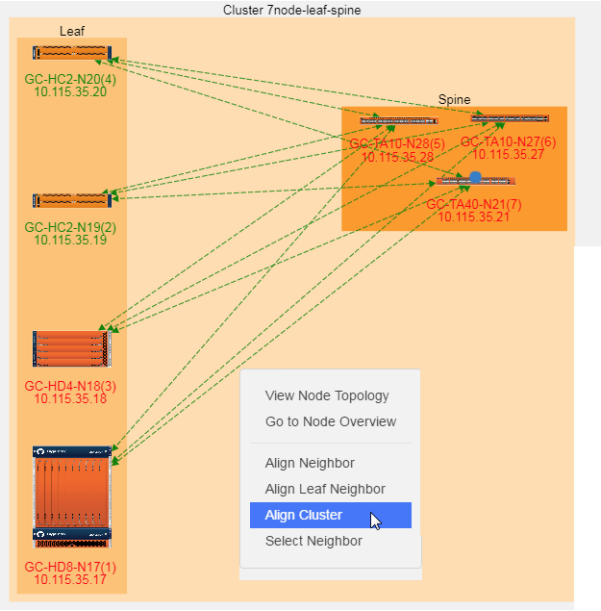

Figure 18-25: Displaying Align Cluster Option

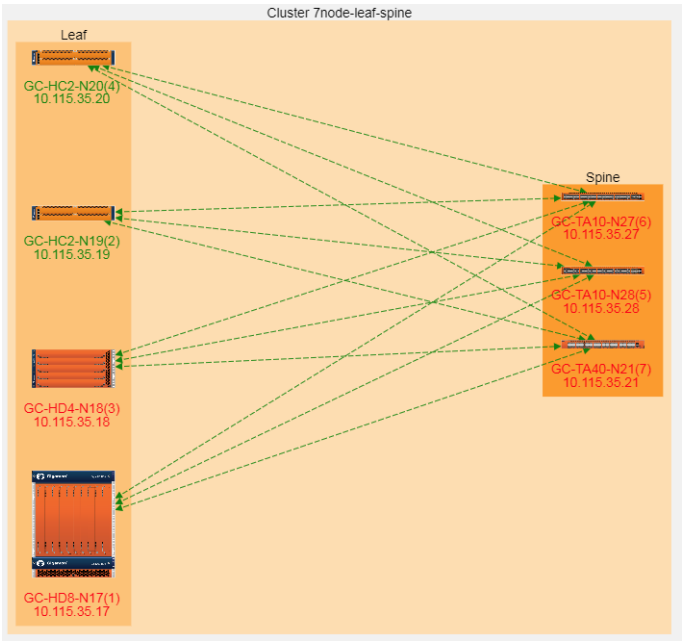

Figure 18-26: Spine and Leaf Nodes after Selecting Align Cluster Option

Select Neighbor

The **Select Neighbor** option is used to move an entire neighbor in the topology. This option selects a node or cluster and its linked nodes so you can move all of the elements at one time. To move a cluster or node and all its neighbors, do the following:

1. Select the cluster or node with linked neighbors.
2. Right-click to open the context menu and select **Select Neighbor**. The node or cluster and its neighbors are highlighted in blue. In the following figure, the GigaVUE-HB1 node and its linked neighbors are selected.

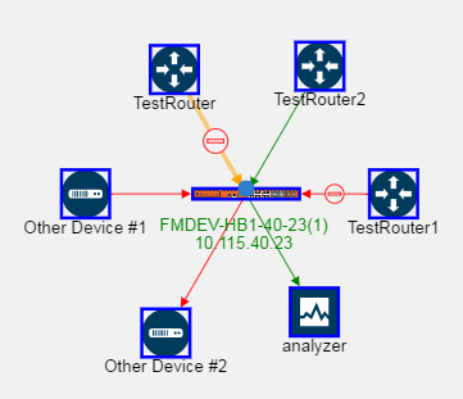

3. Drag the selected cluster and node to the new position. The linked neighbors will move with the node or cluster.

Manage Devices

This section describes the following:

Add Devices

To add a node or nodes to the topology, do the following:

1. Select **Add > Add Device(s)**.

   The Add Device page shown in the following figure displays.

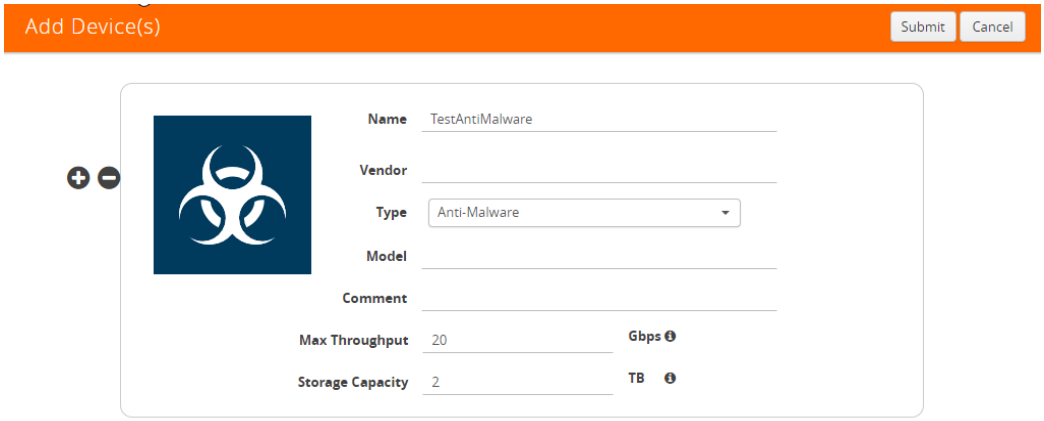

You can add additional nodes by clicking the **+** button. To remove a node, click the **-** button.

2. Enter the following information for the node:
   - **Name**—the name of the node.
   - **Vendor**—the node's manufacturer. For Gigamon nodes, you cannot edit this field.
   - **Model**—the model number of the node.
   - **Comment**—optional description or additional information about the node.
   - **Max Throughput** - the maximum throughput traffic currently being sent to this tool.
   - **Storage Capacity** - the total processing capacity dedicated to all Gigamon ports physically connect to the tool.

   NOTE: This field is not displayed when selecting a device type **Network**.

   The information entered for the node displays when hovering over the icon on the topology. Also, the type selected determines the icon.
3. Click in the **Type** field to select the icon for this node. The possible types are as follows:

Type	Icon
Router	
Switch	
Virtual Switch	
Load Balancer	

Type	Icon
Firewall	
VOIP	
Cloud	
Analyzer	
Anti-Malware	
Customer Experience Management	
Application Performance Management	
Data Loss Prevention (DLP)	
Forensics	
Intrusion Detection System (IDS)	
Intrusion Prevention System (IPS)	
Next Generation Firewall (NGFW)	
SIEM	
Other	

4. Click **Next**. The system returns to the Topology, where the new node appears on the page.

Edit Devices

To edit a node, do the following:

1. Select the node.
2. Select **Actions > Edit**.

   The Edit page opens, which displays information about the node and the links to the node.
3. On the Edit page, make the changes to the node or link information or both.
4. Click **Submit**.

NOTE: You cannot edit Gigamon nodes.

Delete Devices

To delete a single node from the topology, do the following:

1. Select the icon for the node in Topology View or select the checkbox in Table View.
2. Select **Actions > Delete**.
3. On the confirmation dialog, click **OK**.

To delete all nodes, select **Actions > Delete All**. All nodes and links are removed from the topology.

NOTE: You cannot delete Gigamon nodes from the topology.

Add TAP

In addition to nodes and links, you can add Gigamon TAPs to the topology with the TAP wizard or manually add the TAP and links.

To add a TAP with the wizard, do the following:

1. Select the links for the TAP.

   To select the links, press Shift and select the two node links for the TAP.
2. Select **Add > Tap**.

   The Add Tap page displays, which includes the dialogs for configuring the network and monitor links for network A and the network and monitors links for network B. The default names for the links are Network A, Network B, Monitor A, and Monitor B. Figure 18-27 shows the Add Tap Wizard with a Gigamon G-TAP A-SF selected.

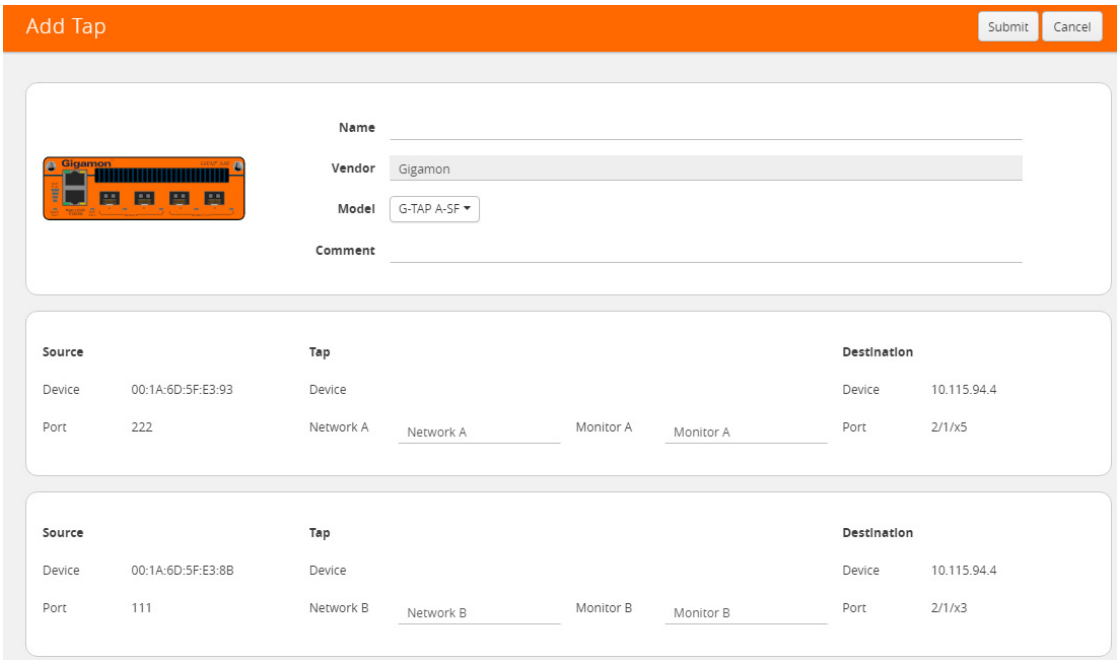

Figure 18-27: Add Tap Wizard

3. Enter type a name for the TAP in the **Name** field.
4. Select a Gigamon TAP model from the list of **Models**.
5. (Optional) Add a comment about the TAP in the **Comment** field.
6. Click **Submit**.

   The wizard adds the TAP and links to the topology as shown in the following figure, where TapDemo2 is the newly added TAP. When the TAP is added, it is initially placed in the lower right-hand corner of the page.

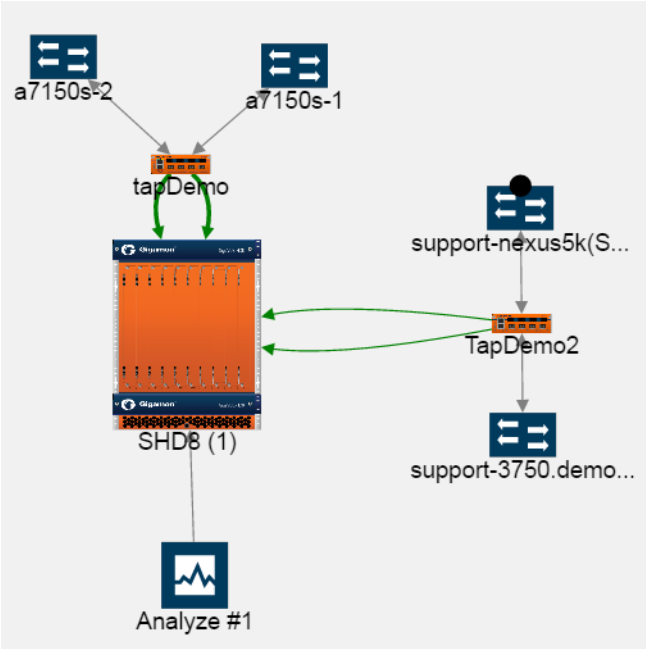

Manage Links

After adding nodes to the topology, you can add one or more links between nodes. However, there are some restrictions regarding links:

- CDP and LLDP links are read only.
- Only CDP and LLDP nodes can be the source for a TAP.
- Only a Gigamon node can be the destination for a TAP.

NOTE: When Gigamon discovery is enabled, the manual links that already exists between the two GigaVUE nodes are suppressed and only the Gigamon discovery links are shown.

Add Links

To add a link or links between nodes, do the following:

1. Click **Add > Link(s)** or click on the source node and drag a line from the dot that appears on the icon to the destination node as shown in the following figure.

The Add Link page shown in the following figure displays.

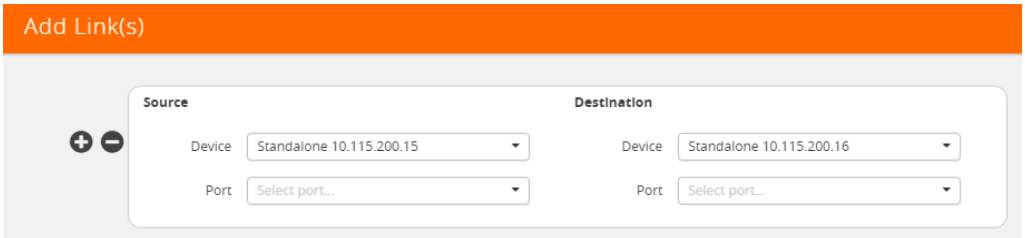

2. Select the source and destination ports for the link.

   The **Device** fields for the Source and Destination nodes are automatically populated with the nodes that you selected for the link if you used the drag and drop method for creating the link. Otherwise, select the nodes from the drop-down lists.

   a. Under **Source,** click in the **Port** field and select the source port from the drop-down list.

   b. Under **Destination**, click in the **Port** field and select the destination port from the drop-down list.

   The Add Links page should look similar to the following figure.

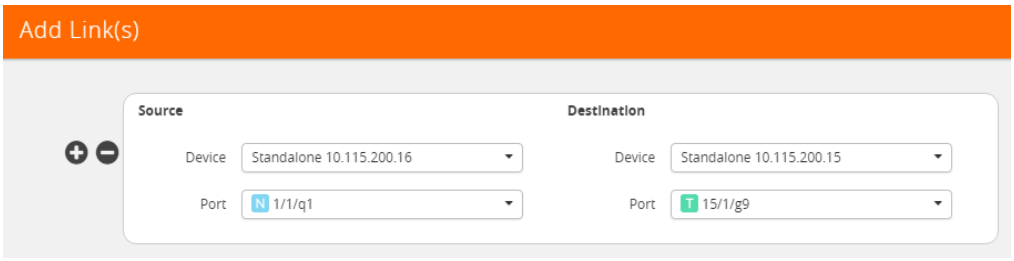

3. (Optional) Add another link by clicking the **+** button, and then repeat steps 2 and step 3.

   NOTE: When adding more than one link between two nodes, both Source and Destination must contain information in the **Port** field.

4. Click **Submit**.

   The topology shows a link between the two nodes.

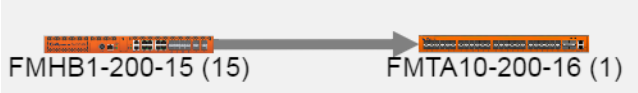

Edit Links

To edit links, do the following:

1. Select the link to edit.

   When you select a link with neighbors, all the neighbors are automatically selected.

2. Select **Actions > Edit**.
3. On the Edit Link page, make changes to the link information.
4. Click **Submit**.

Export and Import Topology

You can export the description of the topology to an Excel spreadsheet, which you can modify, and then import.

Export Topology

To export the Topology, click **Export**, which downloads an Excel spreadsheet. To identify the file, the filename ends in a timestamp string that is in the format yyymmddhhmmss.

The spreadsheet contains four sheets: Devices, Links, Gigamon Devices, and Discovered Links. The Gigamon Devices and Discovered Links pages are informational only. You can use the Devices and Links sheets for modifying and then import the information back to the GigaVUE-FM Topology. The format of the Gigamon Devices page is the same as the Gigamon Table View in the UI (refer to *Table View* on page 316).

Export exports all the manual links and manual nodes in a topology. Aggregated links are exported as the individual links of the aggregation. When exporting from the Cluster view, the expanded view of the aggregated links is exported.

Import Topology

Devices and links can be added to the Topology by importing an Excel spreadsheet that describes the custom nodes or links. When importing a spreadsheet, the spreadsheet must include a Device sheet and a Links sheet. The Devices sheet must have the following columns:

- Node Name
- Vendor
- Type
- Model
- Comment

The Type column can be any of the following:

- Router
- Switch
- Virtual Switch
- Load Balancer
- Firewall
- VOIP
- Cloud
- Analyzer
- Anti-Malware
- Customer Experience Management
- Application Performance management (APM)
- Packet Capture
- Data Loss Prevention (DLP)
- Forensics
- Intrusion Detection System (IDS)
- Intrusion Prevention System (IPS)
- Next Generation Firewall (NGFW)
- SIEM
- Other

The Links sheet must have the following columns:

- Source_Device
- Source_Port

- Source_Type
- Destination_Device
- Destination_Port
- Destination_Type

To import an Excel spreadsheet with node and link information, do the following:

1. Click **Import**.
2. Select one of the following:
   - Import Device(s)
   - Import Link(s)
3. On the Import page, do either of the following:
   - Click **Select File** and navigate to the file you want to import.
   - Drag and drop the file onto the page.
4. Click **Submit** to apply the spreadsheet information to the topology.

FabricVUE Topology Views

The FabricVUE Topology views allow you to visualize the complete underlying network topology graphically from GigaVUE-FM. There are two topology views available when viewing a node from GigaVUE-FM:

- *Node Topology* on page 333
- *Map Topology* on page 337

Node Topology

The Node Topology provides the same functionality as the Topology view displayed for all clusters nodes, and nodes discovered by GigaVUE-FM or manually added. For details, refer to *Overview of Topology* on page 298. Node Topology displays the topology of individual clusters and standalone nodes and their immediate neighbors, whereas Topology displays all clusters nodes, and nodes discovered by GigaVUE-FM or manually added. Nodes in a cluster are displayed inside a yellow container. Figure 18-28 is an example of a cluster with four nodes and a number of connected nodes.

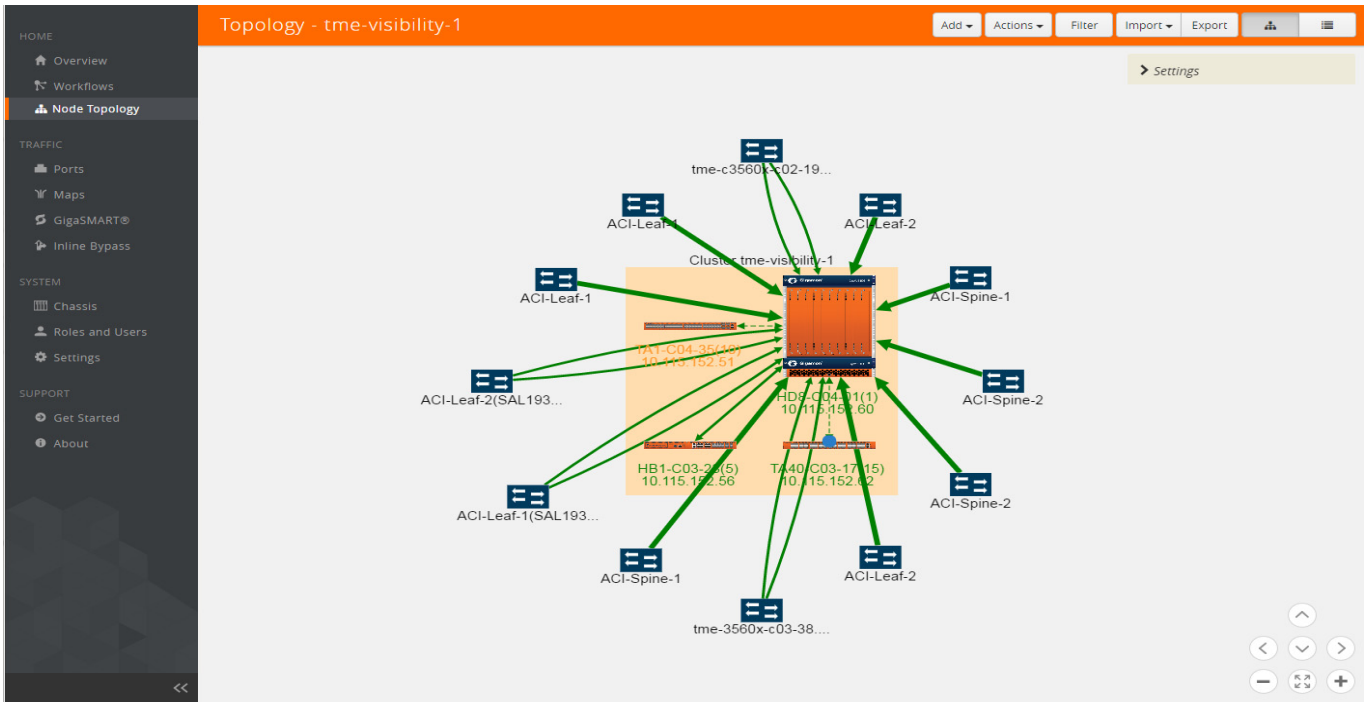

Figure 18-28: Physical Node Topology of a Cluster

In Node Topology, only the immediate neighbors of the node is displayed in the topology. Figure 18-29 on page 334 shows a topology where the node with the host name HC2-03-13 was selected. It has two immediate neighbors, so the topology shows the node and its two neighbor nodes.

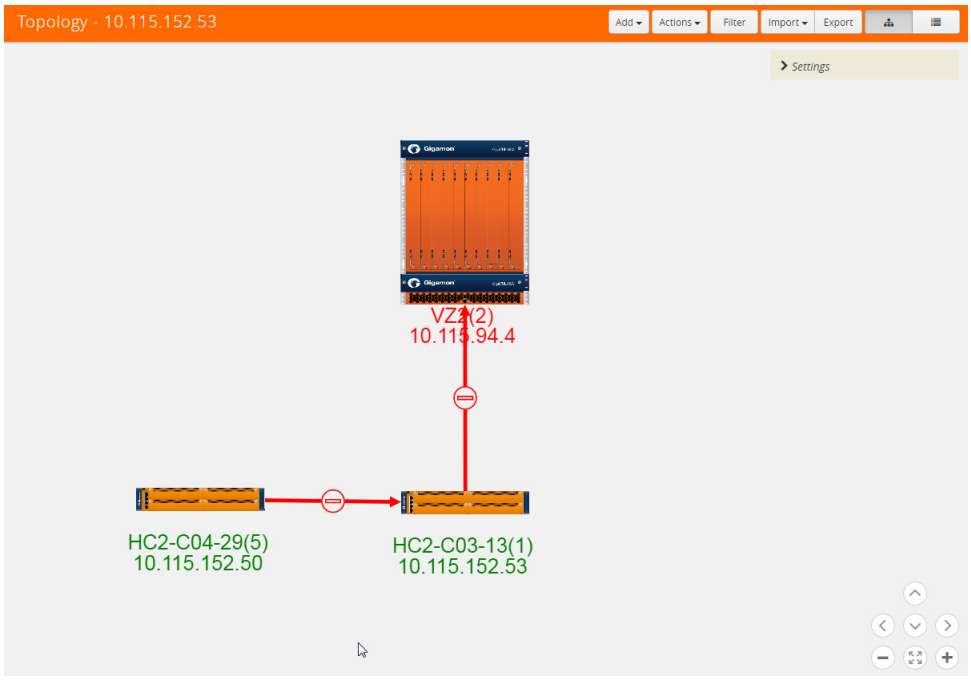

Figure 18-29: Node with Two Immediate Neighbors

However, if the node with the hostname HC2-C04-29 is selected, it has only one immediate neighbor HC2-C03-13, so the Node Topology shows only those two nodes.

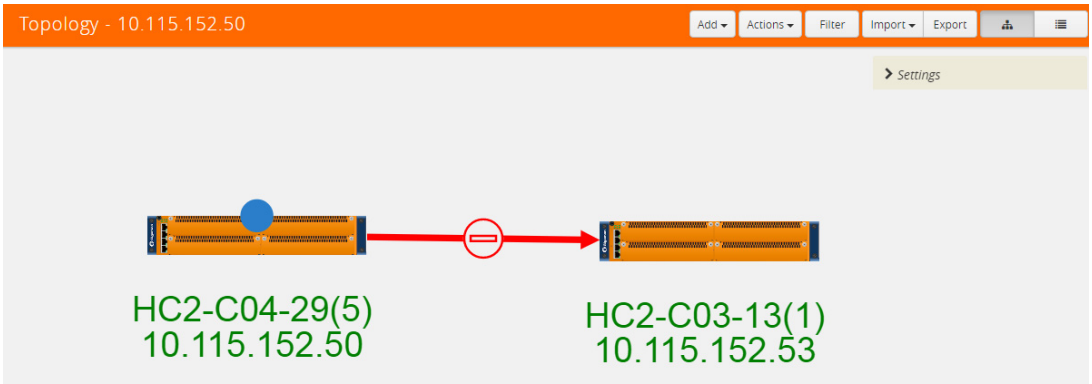

Figure 18-30: Node with One Immediate Neighbor

The labels for each node provide basic information about the node: hostname, box ID, IP address. The color of the label indicates the node's health. The node health colors are defined in Table 18-2.

Table 18-2: Node Health Status

Node Health Status	Meaning
Green	Node is healthy
Amber	Node is in a warning state
Red	Node id down
Gray	Node status is unknown or node is not reachable.

The orientation of the nodes in the cluster container can be changed with the **Align Cluster** option in the context sensitive menu that appears upon right-clicking the cluster. You can also rearrange the chassis icons by clicking on the icon and dragging it to a new location. For more information about repositioning and aligning element in the topology, refer to *Align Cluster* on page 322.

Clicking on the graphic of a chassis shows the details about the chassis in a tool tip. In Figure 18-31, the details for a GigaVUE-HD8 is displayed.

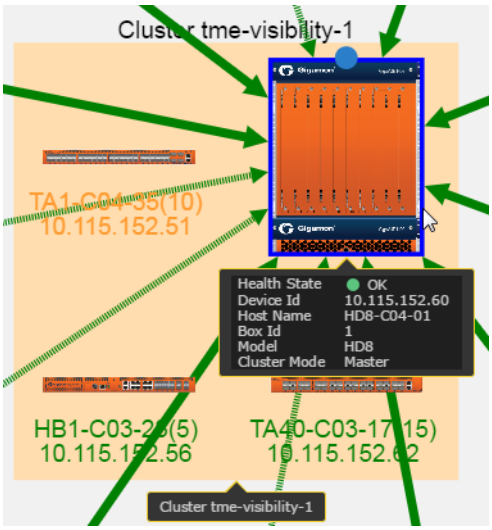

Figure 18-31: Details in Node Topology

The color of the links between nodes and nodes connected to the cluster indicate the link status between the nodes and cluster. Stack links between nodes in the cluster appear as dashed lines. In Figure 18-28 on page 334, the stack links are between the GigaVUE-HD8 node, GigaVUE-TA1, and GigaVUE-TA10. The thickness of the lines indicates the bandwidth of the link. The color of the links indicated its status. The colors are defined in Table 18-3.

Table 18-3: Link Status

Link Status	Meaning
Green	Link is up
Amber	Link is in a warning state
Red	Link is down
Gray	Link is in an unknown state.

To display information about a link, click on the link to display a tool tip. The tool tip provides information about the link's state, health state, link type (LLDP, CDP, or Gigamon Discovery), source and destination nodes, and source and destination ports

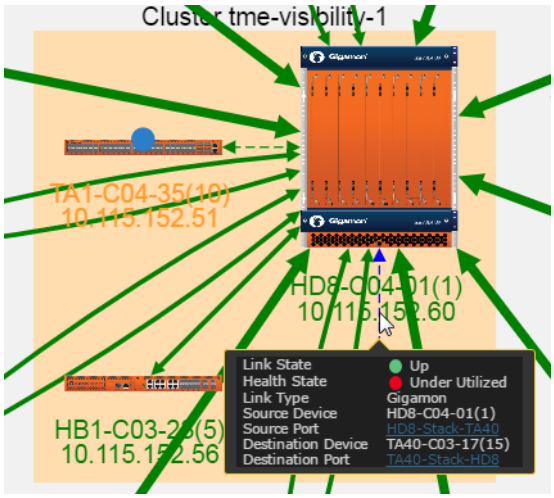

Figure 18-32: Link Information

To display details about a link, right-click on a link and select **More Details** from the context menu. A Link Details quick view opens (refer Figure 18-33 on page 337). For more information about the Link Details quick view, refer to *Link Details* on page 312.

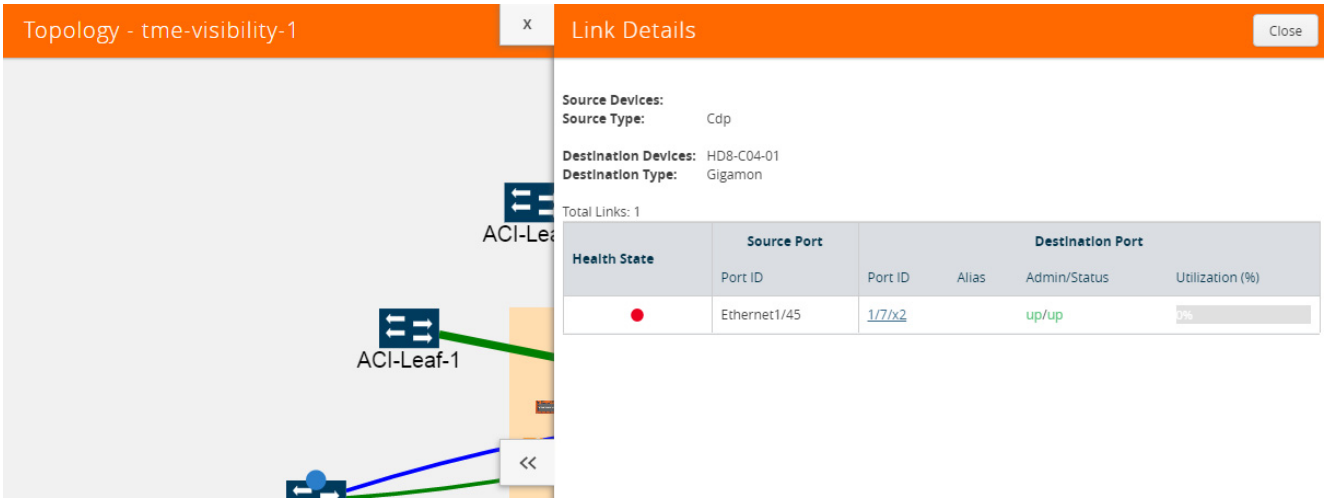

Figure 18-33: Link Details

For details about GigaStream stack link details in leaf and spine cluster configuration, refer to *GigaStream Stack Link Details* on page 313.

Map Topology

When you access a physical node from GigaVUE-FM, and then select **Maps**, you can toggle the page between a list view and a topology view. Clicking the topology button changes the page to the topology view for the maps on the node. Figure 18-34 shows an example with network, hybrid, tools ports, and maps. A legend in the top-right corner of the page describes the icons shown in the diagram. The number of items in the legend depends on the icons used in the current graph. Use the controls in the lower right-hand corner of the page to position the map and zoom in or zoom out.

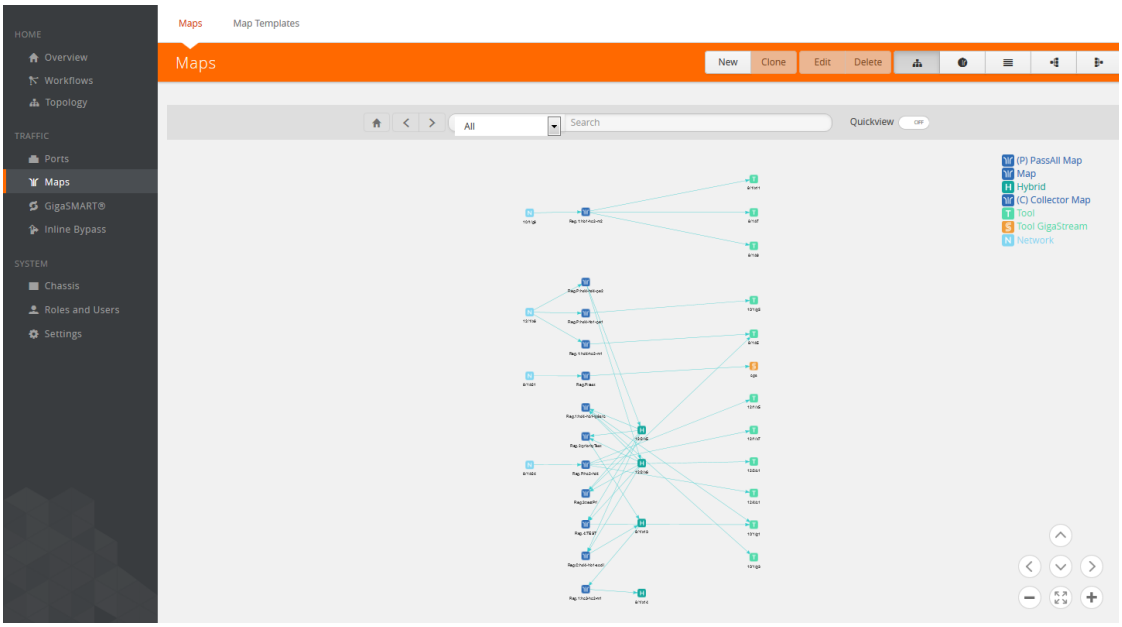

Figure 18-34: Map Topology View

Hovering over a node in the topology highlights the connections between ports and maps. For example, in Figure 18-35, hovering over the hybrid port highlights the path from the network port through the maps to the hybrid port and the destination ports.

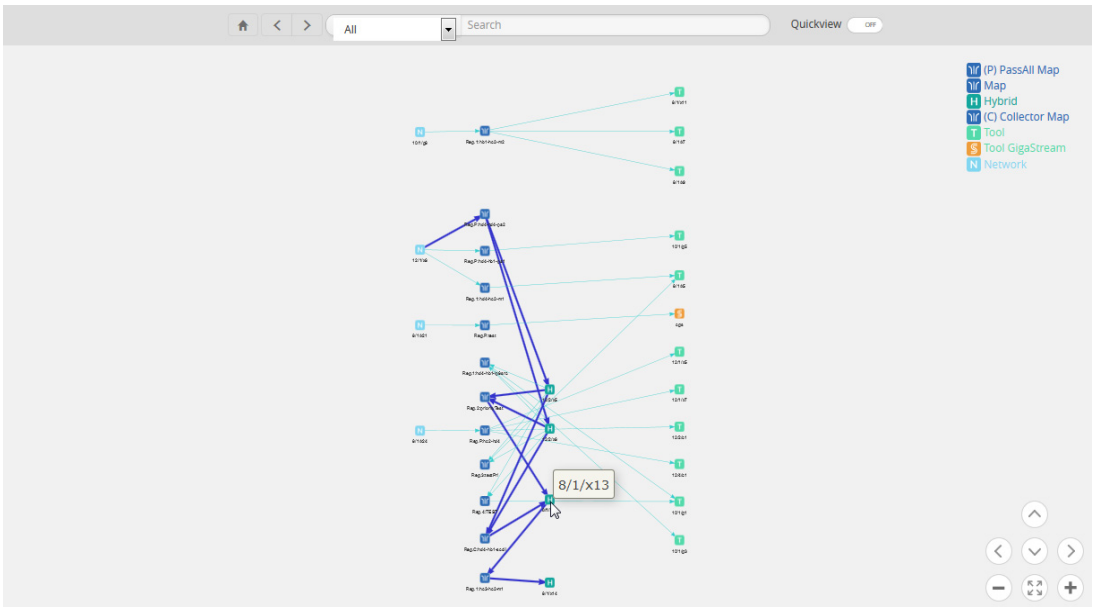

Figure 18-35: Source Port Selected and Paths Highlighted

Clicking on the item in the diagram displays only the connections for that item as shown in Figure 18-36.

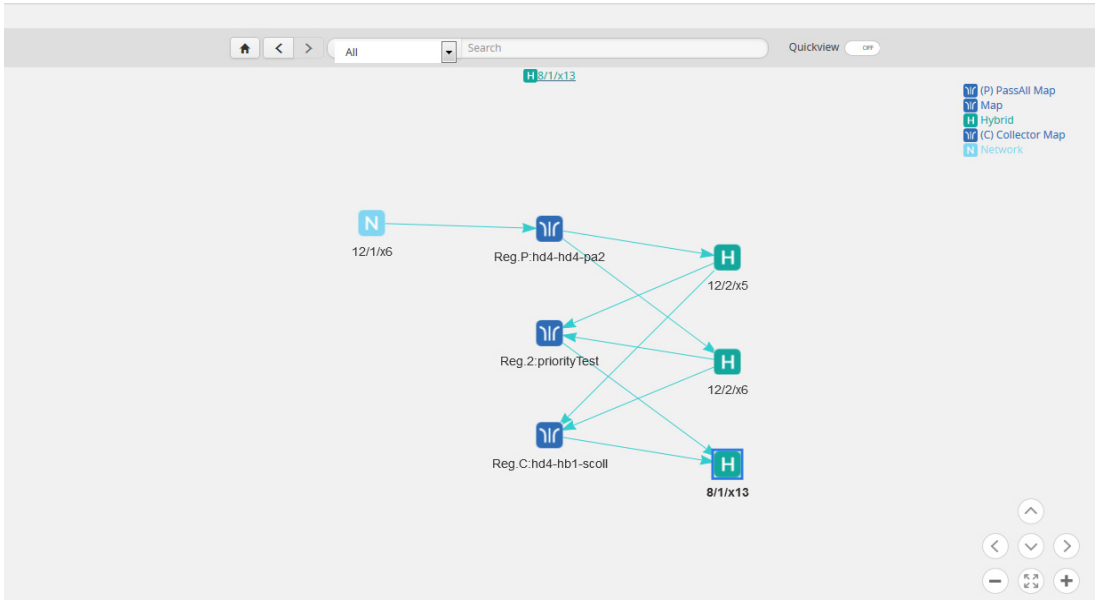

Figure 18-36: Map Selected

When an item in the diagram is selected, a link is displayed at the top of the page. Clicking on the link displays a quick view for that item. For example, if the item is a map, a Map quick view displays and if the item is a port, a Port quick view displays. Figure 18-37 shows an example, where port 8/1/x13 is selected and the quick view for the port is displayed.

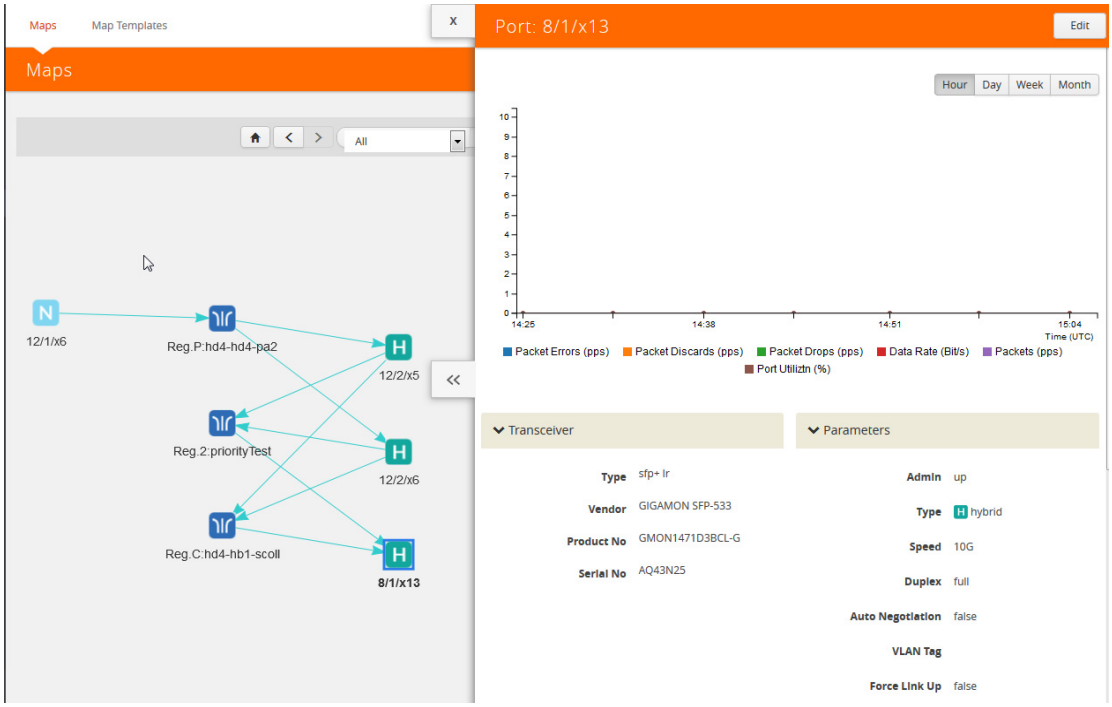

Figure 18-37: Map Quick View from Map FabricVUE

Redistribute Traffic Flows

GigaSMART software version 5.4 provides support for redistributing traffic flows from a particular tool port to other available tool ports with available capacity. When traffic flows from tools ports to end tool ports that exceed their traffic bandwidth, those end tool ports become problematic, as shown in the illustration below.

To use this feature the users have to create a tool device by entering throughput and storage capacity they wish to allocate for Gigamon ports. After the device is created, the user can create inks between Gigamon devices and manual devices. Based on the number of incoming links to tool, total traffic per port, max throughput and storage capacity, the wrap around time and current throughput percentage % is calculated.

When a tool device is linked to Gigamon device(s), it is included as a part of the traffic flows. if a Gigamon device with maps and ports is linked to a tool device, then the terminal vertex of the flow is the tool device. Flow health for a flow connected to a tool device is then be calculated by taking into consideration the current throughput percentage% of the tool device along with existing parameters

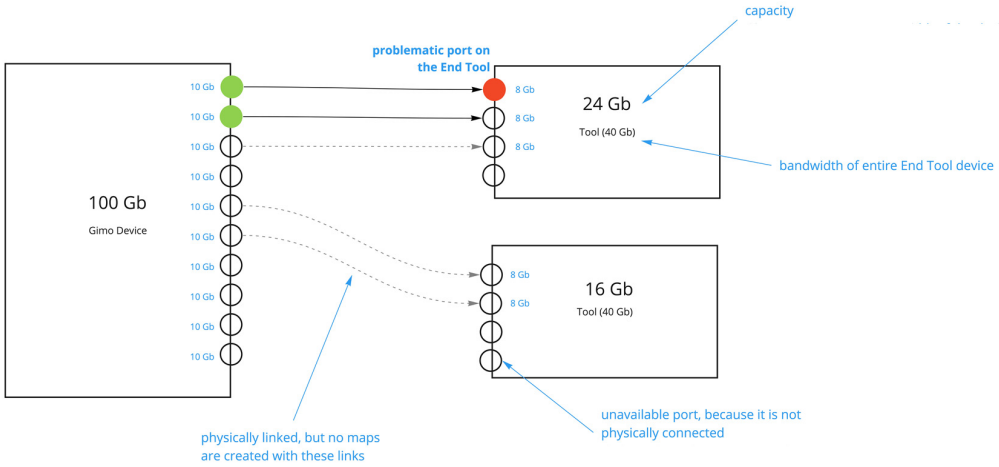

Support for this functionality includes the following:

- Listing all the Tools that are present in Fabric Manager
- Troubleshooting if there is any problem with the port that is connected to tool
- Listing the total traffic flowing to the tool from the Gigamon ports that are connected to the tool.
- Analyzing the traffic by sending or duplicating it to another port or tool.
- Redistributing the traffic to one or more available tool ports.
- Moving all traffic from the problematic port to another tool.

Redistribute traffic to one or more available tool ports

This section describe the tasks associated with redistributing traffic flow from one tool port to another available tool port.

Add Devices

1. Add a new tool and indicate the **Max Throughput** and **Storage Capacity** the tool will dedicate to Gigamon ports.

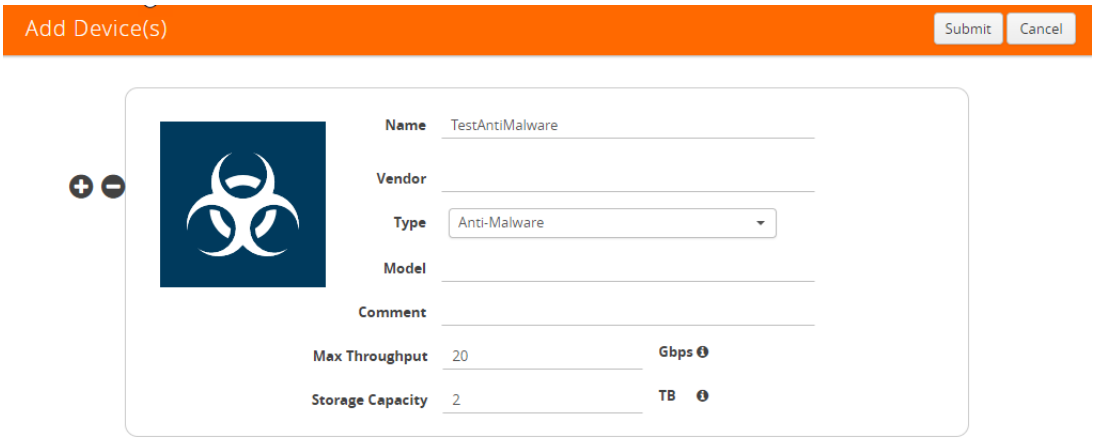

2. Enter the following information for the node:
   - **Name**—the name of the node.
   - **Vendor**—the node's manufacturer. For Gigamon nodes, you cannot edit this field.
   - **Model**—the model number of the node.
   - **Comment**—optional description or additional information about the node.
   - **Max Throughput** - the maximum throughput traffic currently being sent to this tool.
   - **Storage Capacity** - the total processing capacity dedicated to all Gigamon ports physically connect to the tool.

   NOTE: This field is not displayed when selecting a device type **Network**.

   The information entered for the node displays when hovering over the icon on the topology. Also, the type selected determines the icon.

3. Create a link from the Gigamon device to the new tool you create in **Step 1** and provide the Gigamon ports that are connected to the new tool you created.

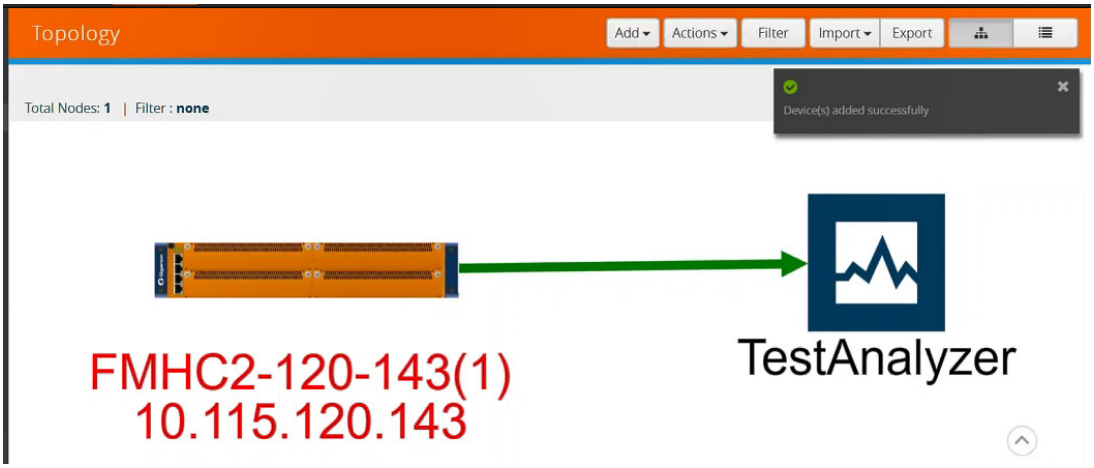

4. Add the Device links **Source** and **Destination port** details.

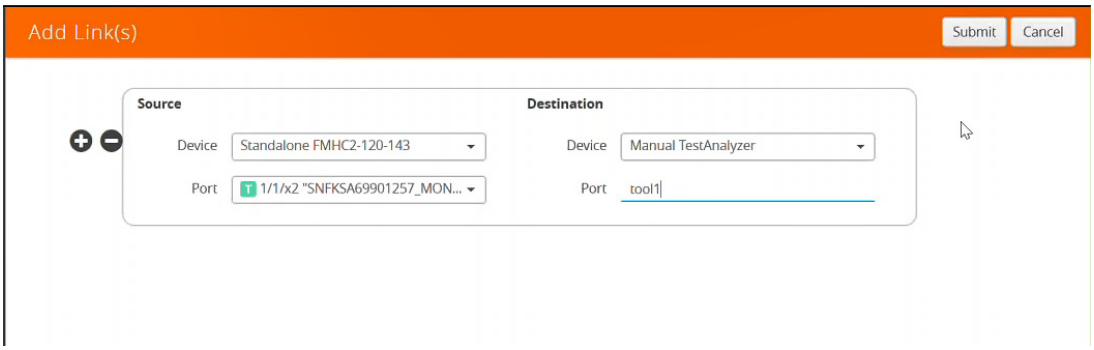

5. Click **Submit**.
6. Click **Physical > Tools** to display the Tools page with new tool you created in first Step**.**

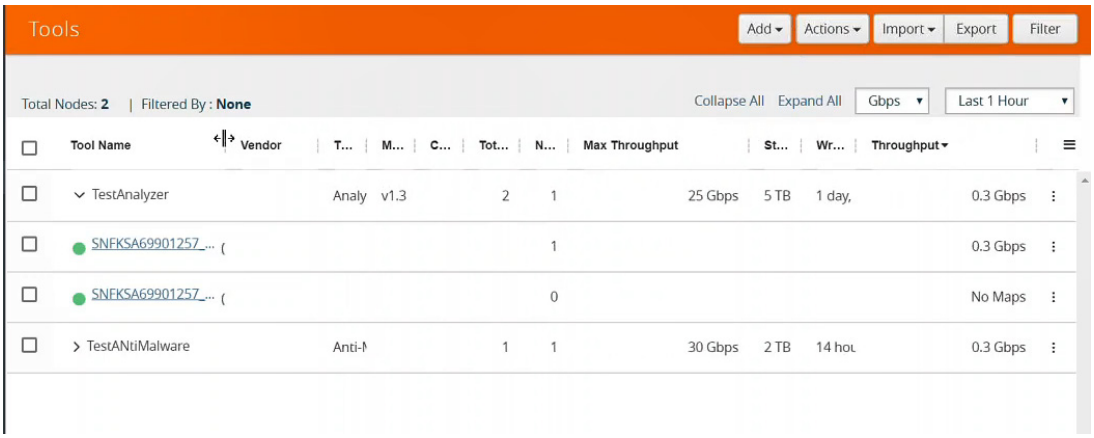

The Tools page displays the available Tools in your network topology, including the tool you created and the associated traffic flows. Using this screen you can view the traffic

being handled by each tool port. The columns show details about the tool port. The columns can be changed using the **Columns** menu in right corner of the screen.

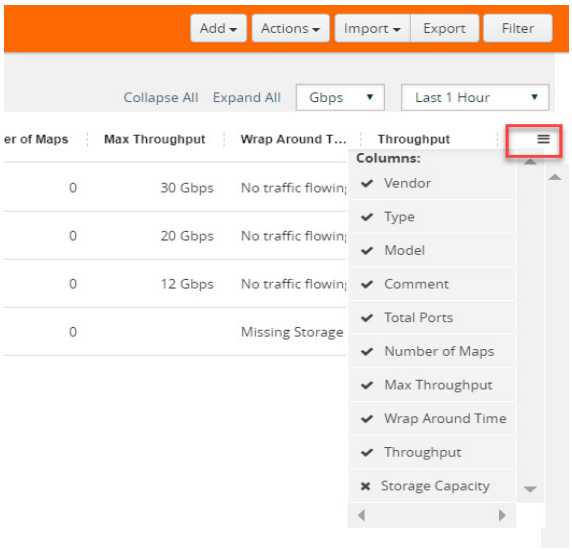

- Use the metrics drop-down to view the traffic throughput in **Mbps** or **Gbps.**
- Use the **Last time** drop-down to select and view the tool port traffic for a specific time period.

Redistribute Traffic to Another Port

Traffic that may be problematic for one tool port can be moved and redistributed to another tool port. Using the Tools page you can list the total traffic flowing to the tool from the Gigamon ports that are connected to the tool. The Tools page will help you analyze the traffic for redistribution by sending or duplicating it to another port or tool.

Example: Redistribute Traffic to Another Port

Below is an example of Redistributing Traffic to Another Port.Viewing Traffic Flow

Task	UI Steps
Add a new tool and indicate the Max Throughput and Storage Capacity the tool will dedicate to Gigamon ports.	**Select Add > Add Device(s).** You can add additional nodes by clicking the + button. To remove a node, click the - button. Enter the following information for the node: • Name—the name of the node. • Vendor—the node's manufacturer. For Gigamon nodes, you cannot edit this field. • Model—the model number of the node. • Comment—optional description or additional information about the node. • Max Throughput - the maximum throughput traffic currently being sent to this tool. • Storage Capacity - the total processing capacity dedicated to all Gigamon ports physically connect to the tool. **Click** in the **Type field** to select the icon for this node. **Click Next.** The system returns to the Topology, where the new node appears on the page.

Task	UI Steps
Create links between Gigamon devices and manual device. Based on the number of incoming links to tool, total traffic per port, max throughput and storage capacity, the wrap around time and current throughput % is calculated.	**Click Add > Link(s) or click on the source node and drag a line from the dot that appears on the icon to the destination node.** **Select the source and destination ports for the link.** • The Device fields for the Source and Destination nodes are automatically populated with the nodes that you selected for the link if you used the drag and drop method for creating the link. Otherwise, select the nodes from the drop-down lists. • Under **Source, c**lick in the Port field and select the source port from the drop-down list**.** • Under **Destination,** click in the Port field and select the destination port from the drop-down list. • (Optional) Add another link by clicking the + button, and then repeat steps 2 and step 3. **NOTE:** When adding more than one link between two nodes, both Source and Destination must contain information in the Port field. **Click Submit.**
Use the Tools page detail views to analyze current traffic flows to tool ports. Use the Max Throughput and Storage Capacity data to help determine if any tool port receiving traffic is near or at capacity. Users can identify any available tool port to redistribute traffic.	**Click Physical > Tools. T**he Tools page list all ports available in your network receiving traffic from devices. **Select one of the ports receiving traffic.** Expand a port to see which maps are sending traffic to that port**.** 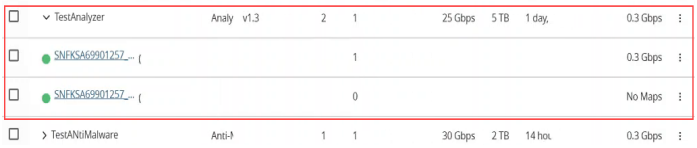 **Select Actions> View Details.** 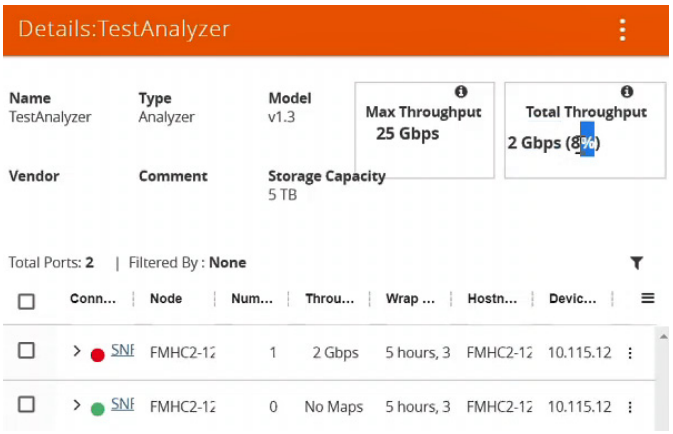 From the **Details** pane you can view the traffic flow for each tool port. Here you can determine if a tool port is at or near capacity and decide if you want to move it to another port with availability and linked to the device.

Task	UI Steps
Redistribute traffic flow from one tool port to another available tool port. Replace an existing port with another port that is currently linked to the end tool.	From the **Tools page** expand a port to view which maps are sending traffic to that port. **Select a Device** Select a **different map** from the **current map** that is at or near capacity to send traffic. **Click Tools> Actions> Edit** - Enter information for the node as you did in the first task. - Add the Device links Source and Destination port details. Click **Submit**. The tools page shows the available tool port with the new traffic flow data.

View Traffic Flows

To view tool port traffic flow, do the following.

1. Navigate to Physical >**Flows.** Flows for all nodes in your network topology displays based on the filter parameters. You can change the display of the flow details by clicking the **Filter** button and modifying the filter parameters.

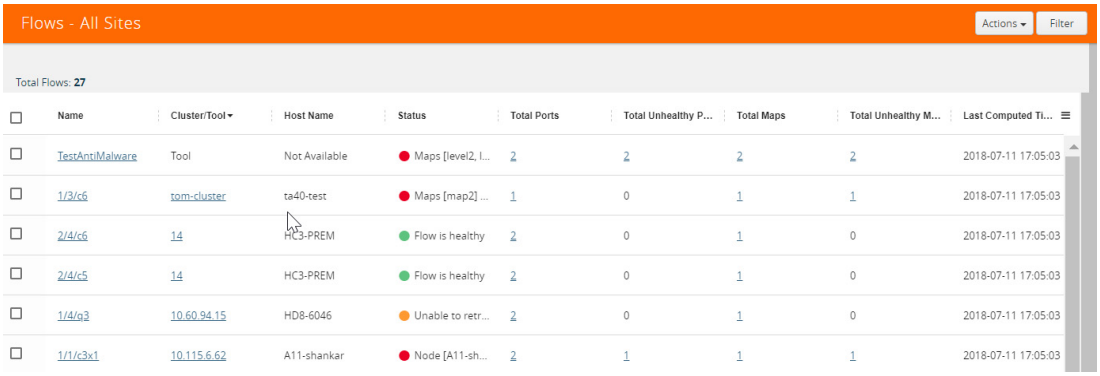

2. Click the name of the node to display the flow.

View Tool Port Details

To view Tool Port Details:

1. Click **Physical > Tools.** A listing of all the tools displays.

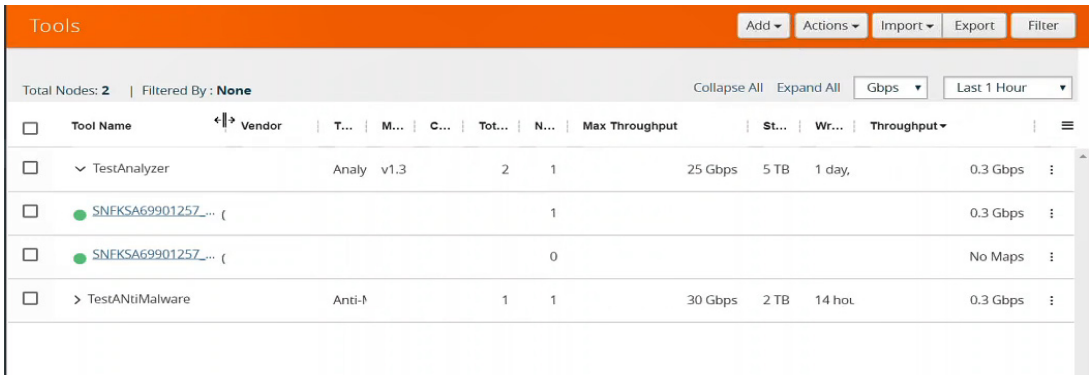

2. Select a **device**.
3. Click **Actions > View Details.**

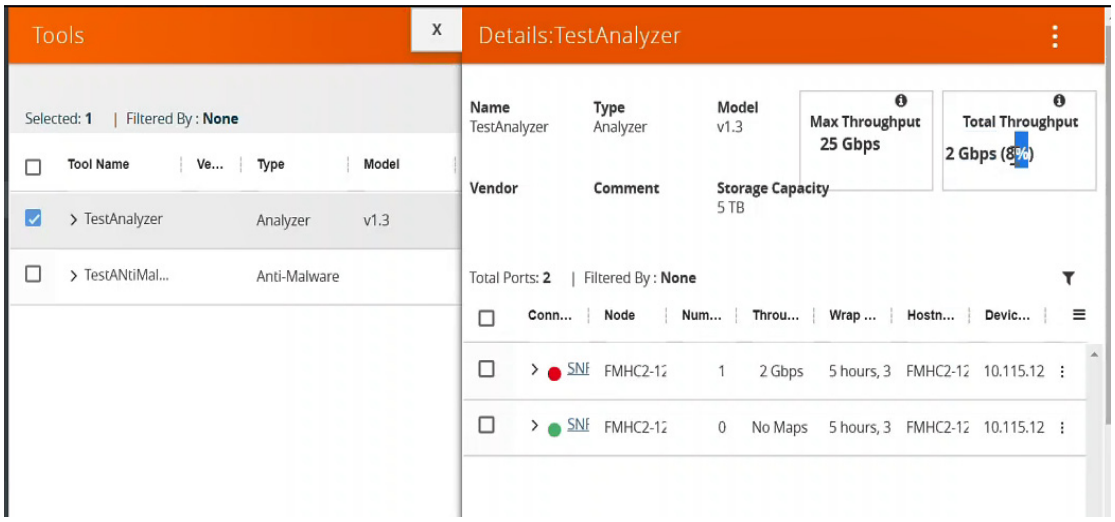

- **Max Throughput** - Sum of all port's throughput connected to the tool device.
- **Current Throughput%** - Total throughput / Maximum Through (User entered data for that Tool).
- **Wrap Around Time** - Total Storage (User entered data for the Tool)/Current Throughput.

View Tool Port Statistical Graph

To view Tool Port Statistical data:

1. Click **Physical > Tools**.
2. Select a **device**.
3. Under the Actions drop down, select View Statistical Graph. The statistics is an aggregated graph of all the ports connected to the tool device.

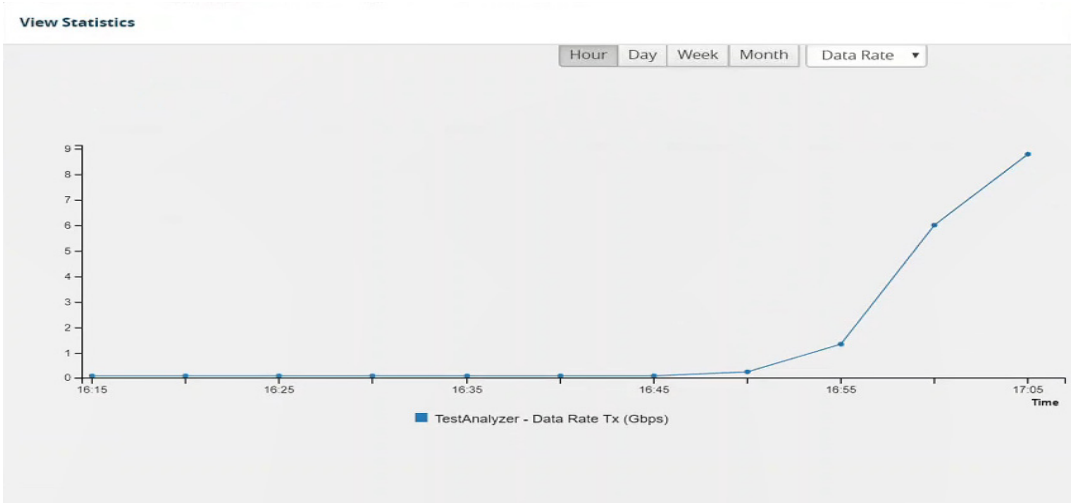

19 Flows

This chapter describes what is a flow and how to quickly view the packet drops and packet errors that are causing the flow to be unhealthy.

Refer to the following sections for details:

About Flows

Flows provide the ability to view the traffic flowing from a network port that receives the traffic from a network TAP to the tool port that sends the traffic to the tools.

A flow is constructed by traversing backward starting from the egress tool port connected to the monitoring tools all the way up to the network ports that receive the traffic from a network TAP. If there are five egress tool ports sending the traffic out to the monitoring tools, there will be five flows displayed in the Flows page.

Figure 19-1 on page 348 shows an example of a flow. It illustrates the network ports that receive the packets from different sources, a set of map rules that filter the packets, and the egress tool port that sends the packets to the monitoring tools. A flow name is determined by the egress tool port ID or alias. In this example, the flow name is SPEGE0070_Te5, which is the name of the egress tool port.

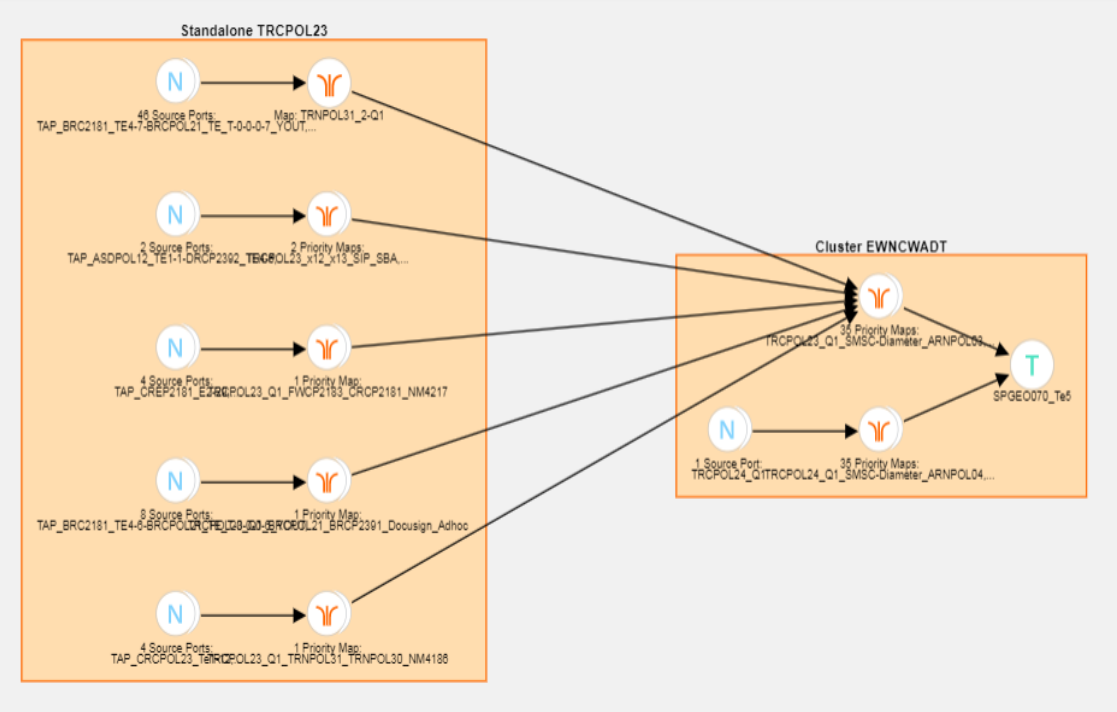

Figure 19-1: Flow - All Sites

Figure 19-2 on page 349 shows a number of GigaVUE nodes in the Topology page that are not connected to each other.

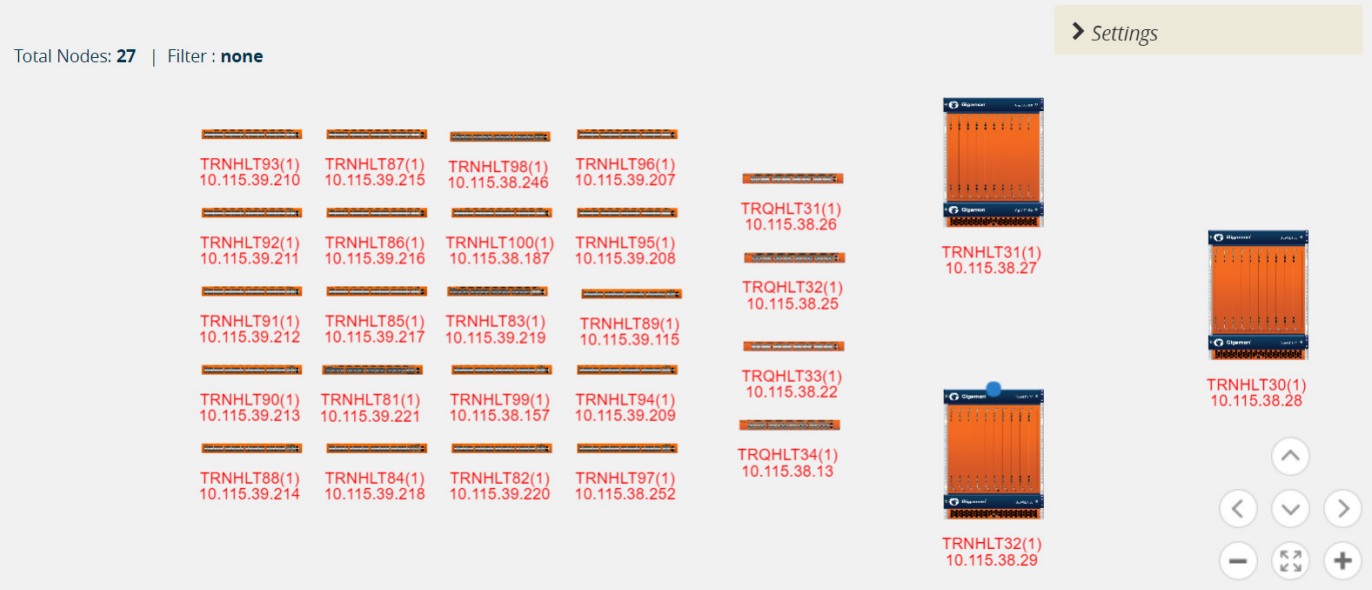

Figure 19-2: Standalone GigaVUE Nodes

Each node seen in Figure 19-2 on page 349 is represented as a separate flow in the Flows page.

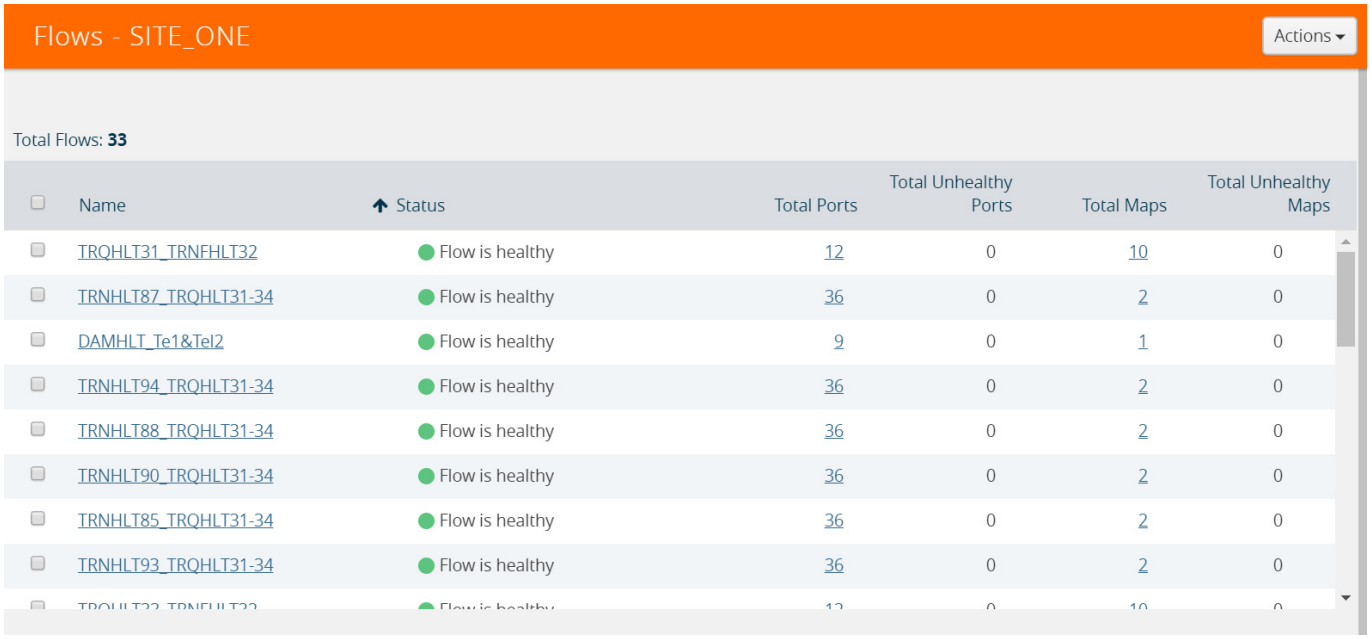

Name	Status	Total Ports	Total Unhealthy Ports	Total Maps	Total Unhealthy Maps
TRQHLT31_TRNFHLT32	Flow is healthy	12	0	10	0
TRNHLT87_TRQHLT31-34	Flow is healthy	36	0	2	0
DAMHLT_Te1&Te12	Flow is healthy	9	0	1	0
TRNHLT94_TRQHLT31-34	Flow is healthy	36	0	2	0
TRNHLT88_TRQHLT31-34	Flow is healthy	36	0	2	0
TRNHLT90_TRQHLT31-34	Flow is healthy	36	0	2	0
TRNHLT85_TRQHLT31-34	Flow is healthy	36	0	2	0
TRNHLT93_TRQHLT31-34	Flow is healthy	36	0	2	0

Figure 19-3: Flows Representing the Standalone Nodes

When those GigaVUE nodes are connected using manual links or Gigamon Discovery links, the number of flows created depends on the number of egress tool ports sending

the traffic out to tools. In this example, three flows are created as shown in Figure 19-2 on page 349.

Flows - All Sites

Actions

Total Flows: **3**

Name	Status	Total Ports	Total Unhealthy Ports	Total Maps	Total Unhealthy Maps
DAMHLT_Te1&Te2	Maps [TRNHLT81_1Q1Q2_TRNHLT31_1Q7Q8_GTP-DCPCF001, TRNHLT81_1Q1Q2_TRNHLT31_1Q7Q8_Permanent] in the flow are unhealthy	843	2	125	2
DAMHLT02_Te1	Maps [TRNHLT81_1Q1Q2_TRNHLT31_1Q7Q8_GTP-DCPCF001, TRNHLT81_1Q1Q2_TRNHLT31_1Q7Q8_Permanent] in the flow are unhealthy	426	2	63	2
DAMHLT01_Te1	Flow is healthy	426	0	63	0

Figure 19-4: Healthy Flows - sample illustration of flows after connections are made

NOTE: Gigamon Discovery is supported only from GigaVUE-FM 5.2 and above.

Figure 19-10 on page 356 shows how unhealthy priority maps are illustrated in a flow view page.

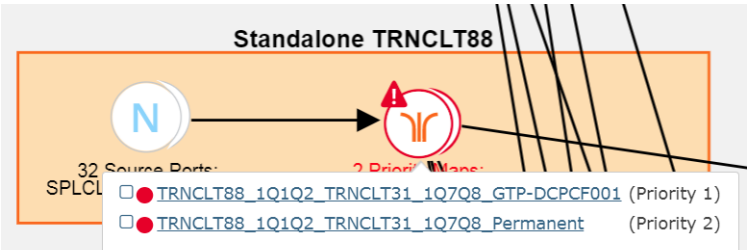

Figure 19-5: Unhealthy Maps

Using flows, you can perform the following:

- Quickly identify the maps and ports that are unhealthy.
- Quickly drill down the unhealthy map to investigate the cause of failure.
- Export a list of unhealthy ports and maps.
- View the data traffic across the flow, starting from the egress tool port that sends the traffic to the monitoring tool to the ingress network ports that receives the traffic from the network taps or span/mirror ports. The arrows in the flow indicate the path of the traffic flow.
- View the pass all maps and the priority maps. A priority map set contains multiple maps configured with the same source ports in the port list.
- Select multiple maps and view the statistics to verify if the traffic is flowing as expected.

- Filter flows by Flow Name, Status or Cluster ID.
- Update flows for the selected site.

A flow is automatically constructed every 24 hours, and flow health is calculated every 5 minutes.You can also manually trigger flow and graph calculations on demand.

The following components affect the health of a flow:

- Manual links or Gigamon Discovery links that connect the GigaVUE nodes.
- Maps that are participating in the flow.
- Health status of all the underlying components of a map such as GigaStream, port group, GigaSMART group, and vport.

For information on flow health status, refer to *Flow Health Status* on page 1366. For instructions on updating flows, see *How to Update Flows* on page 362.

A flow can be viewed per site. All the nodes participating in the flow must belong to the same site.

View Flows

To view a flow:

1. Click **Physical** in the top navigation link.
2. (Optional) From the sites drop-down box, select a site.
3. In the left navigation pane, click **Flows**.

   The Flows page provides a summary of the flow name, cluster, host name, status, total ports, total unhealthy ports, total maps, total unhealthy maps and last computed time.

Table 19-1: Flows Summary Page

Option	Description
Name	Name of the flow.
Cluster	Where a specific flow is ending. Cluster ID
Host Name	Host name of the cluster.
Status	The Heath status of the flow.
Total Ports	The total number of ports involved in the flow are displayed in the Total Ports link.
Unhealthy Ports	Total number of ports that are unhealthy. If there are unhealthy ports, click the Unhealthy Ports link to view the related ports that are unhealthy.
Total Maps	Total number of maps participating in the flow. Click the Related Maps link to view the related maps in the flow.

Option	Description
Unhealthy Maps	Total number of maps that are unhealthy. If there are unhealthy maps, click the Unhealthy Maps link to view the related maps that are unhealthy.
Last Computed Time	

4. Click the numbers under Total Ports, Total Unhealthy Ports, Total Maps, and Total Unhealthy Maps to view detailed information about the maps and ports participating in the flow. Refer to *View Maps and Ports* on page 356.
5. To open a flow, select a flow and click **Actions** > **Open Flow**. Alternatively, click a flow.

View the Flow Summary and Statistics

A flow can display all maps that are created for managing the packet distribution in the GigaVUE nodes. A flow summary provides information only about the total number of maps and ports participating in the flow along with the total number of unhealthy maps and ports in the flow. You can also select multiple maps and view the statistics to check how exactly the packets are flowing.

To view the flow summary:

1. Follow steps 1 to 3 as described in *View Flows* on page 351.
2. In the Flows page, click a flow that you want to view. Alternatively, select a flow and click **Actions** > **Open Flow.** The flow view page is displayed.

   The following table describes the information provided in the Summary tab:

Table 19-2: Flow Summary

Option	Description
Related Ports	Total number of ports participating in the flow. Click the Related Ports link to view the related ports in the flow. Refer to *View Total Ports* on page 356. **NOTE:** V ports and GigaSMART ports are not considered.
Unhealthy Ports	Total number of related ports that are unhealthy. If there are unhealthy ports, click the Unhealthy Ports link to view the related ports that are unhealthy. Refer to *View Total Unhealthy Ports* on page 357. To know more about how the health of a port is calculated, refer to *Port Health Status* on page 1362.

Table 19-2: Flow Summary

Option	Description
Related Maps	Total number of maps participating in the flow. Click the Related Maps link to view the related maps in the flow. Refer to *View Total Maps* on page 357.
Unhealthy Maps	Total number of related maps that are unhealthy. If there are unhealthy maps, click the Unhealthy Maps link to view the related maps that are unhealthy. Refer to *View Unhealthy Maps* on page 358. To know more about how the health of a map is calculated, refer to *Map Health Status* on page 1363.

3. To view the statistics, select the maps. To select the maps, follow one of the following methods:

   Method 1:

   a. Click on a map. A list of priority maps are displayed. Refer to Figure 19-6 on page 353.

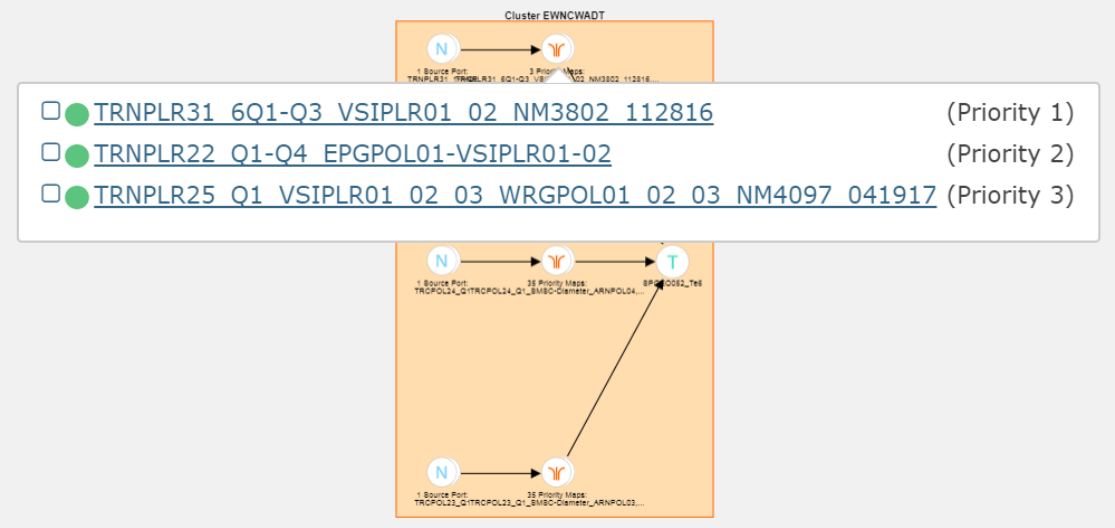

Figure 19-6: Selecting Maps From Maps List

b. Click the check box next to the priority map. The selected map is displayed in the **Items Selected** pop-up. Refer to Figure 19-7 on page 354.

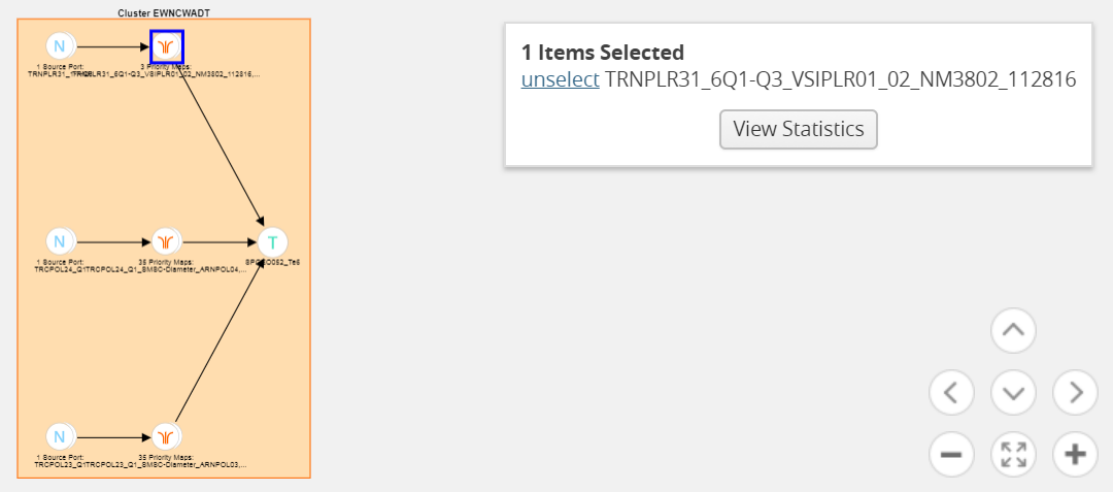

Figure 19-7: Selected Maps in Items Selected Pop-up

c. Repeat step a and b for selecting multiple maps.

Method 2:

a. In the flow view page, double-click on a map. A list of priority maps are displayed. Refer to Figure 19-8 on page 355.

Figure 19-8: List of Priority Maps

b. Click on a map. The map is selected and is simultaneously displayed in the **Items Selected** pop-up. Refer to Figure 19-9 on page 355.

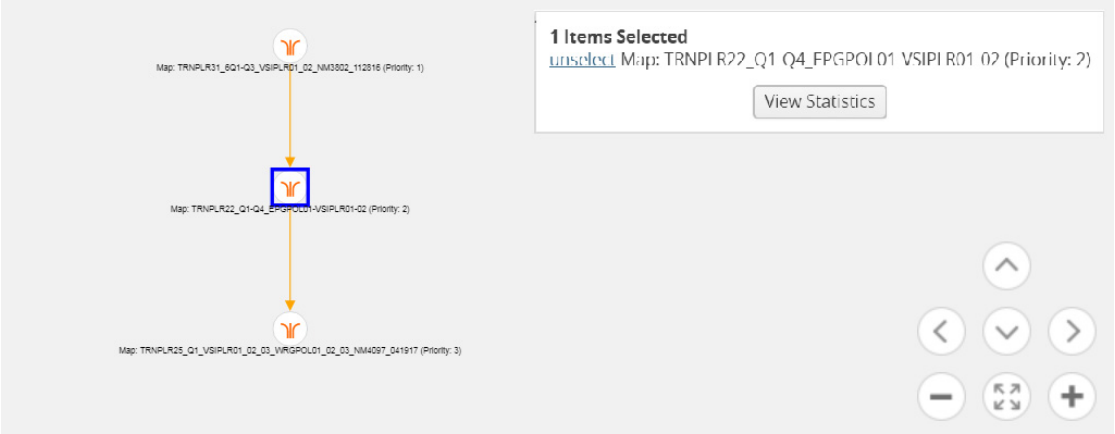

Figure 19-9: Maps Selected

c. Repeat step b to select multiple maps.

4. In the **Items Selected** pop-up, click **View Statistics**. Alternatively, select the **Statistics** tab and click **View Statistics**. The **Statistics** tab displays a graph to show how the packets are flowing from the selected maps.

View Maps and Ports

Figure 19-10 on page 356 shows how unhealthy priority maps are illustrated in a flow view page.

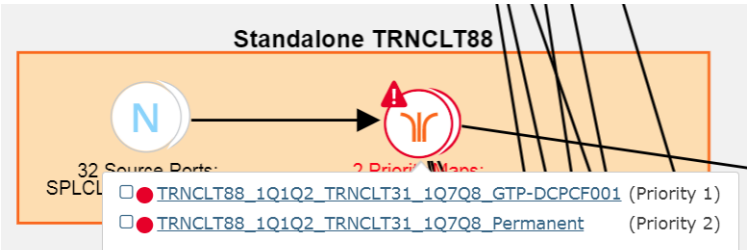

Figure 19-10: Unhealthy priority Maps

Figure 19-11 on page 356 shows how a pass-all map is illustrated in a flow view page.

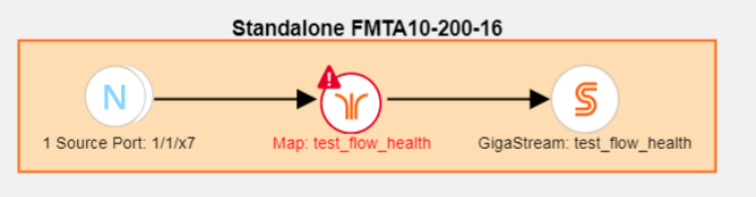

Figure 19-11: Unhealthy Pass-all Map

A priority map set as illustrated in Figure 19-10 on page 356 can contain one or more maps configured with the same source ports. The health of this map is determined by the aggregated health of the priority maps in the map set. Sometimes, this map might be unhealthy because a map that is not participating in the flow is unhealthy. The overall health status of a flow is determined based on the aggregated health of all the maps that are involved in the flow.

There are two ways to view the total number of maps and ports that are healthy and unhealthy.

- In the Flows page, click the total number of ports, unhealthy ports, total maps, and total unhealthy maps to view detailed information about the maps and ports involved in the respective flow.
- Click a flow in the Flows page. In the Summary tab, click the Related Ports, Total Unhealthy Ports, Related Maps, and Total Unhealthy Maps links to view detailed information about the maps and ports involved in the selected flow.

View Total Ports

The total number of ports involved in the flow are displayed in the Total Ports link.

To view the total ports involved in the flow:

1. In the Flows page, click the **Total Ports** link. The All Ports quick view displays a list of all ports that are involved in the flow. It also provides information such as port ID or alias, type, port health status, and the node on which the port is configured.
2. In the All Ports quick view, click a Port ID to open the Port quick view.

3. To return to the All Ports quick view, click **Back**.
4. In the All Ports quick view, click a node under the Node column to open the Overview page.

View Total Unhealthy Ports

The total number of unhealthy ports involved in the flow are displayed in the Total Unhealthy Ports link.

To view the unhealthy ports:

1. In the Flows page, click the **Total Unhealthy Ports** link. The Unhealthy Ports quick view displays a list of unhealthy ports that are involved in the flow. It also provides information such as port ID or alias, type, port health status, and the node on which the port is configured.
2. In the Unhealthy Ports quick view, click a Port ID to open the Port quick view.
3. Click a node under the Node column to open the node's Overview page.

View Total Maps

The total number of maps involved in the flow are displayed in the Total Maps link.

1. In the Flows page, click the **Total Maps** link. The All Maps quick view displays a list of all maps that are involved in the flow. It also provides information such as map ID or alias, type, map health status, and the priority node on which the map is configured. Refer to Figure 19-12 on page 357.

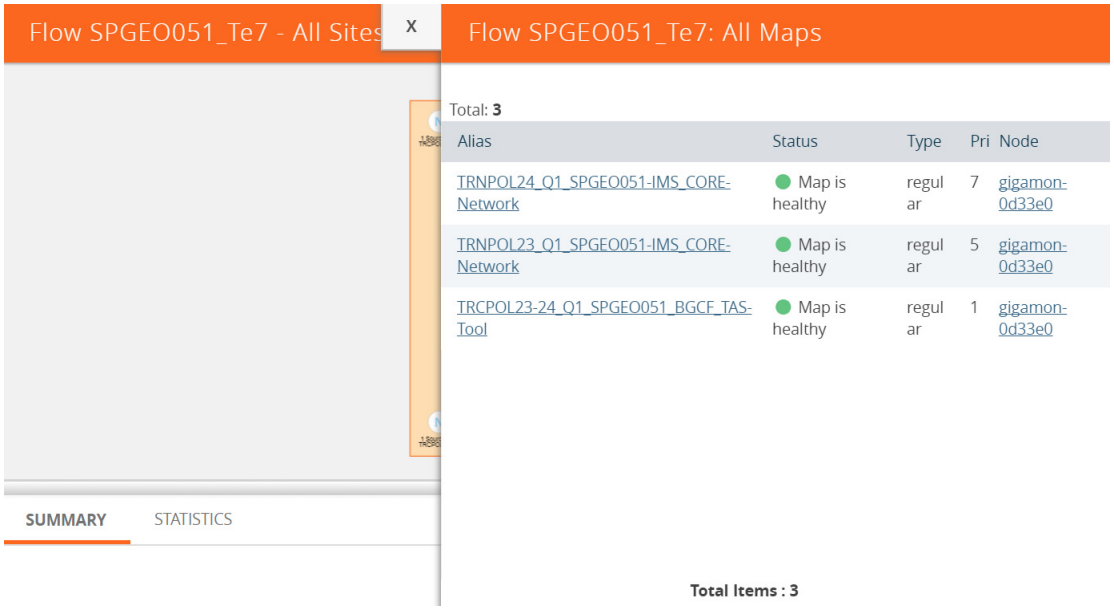

Figure 19-12: Flow - All Maps

2. In the All Maps quick view, click a Map alias to open the Map quick view. Refer to Figure 19-13 on page 358.

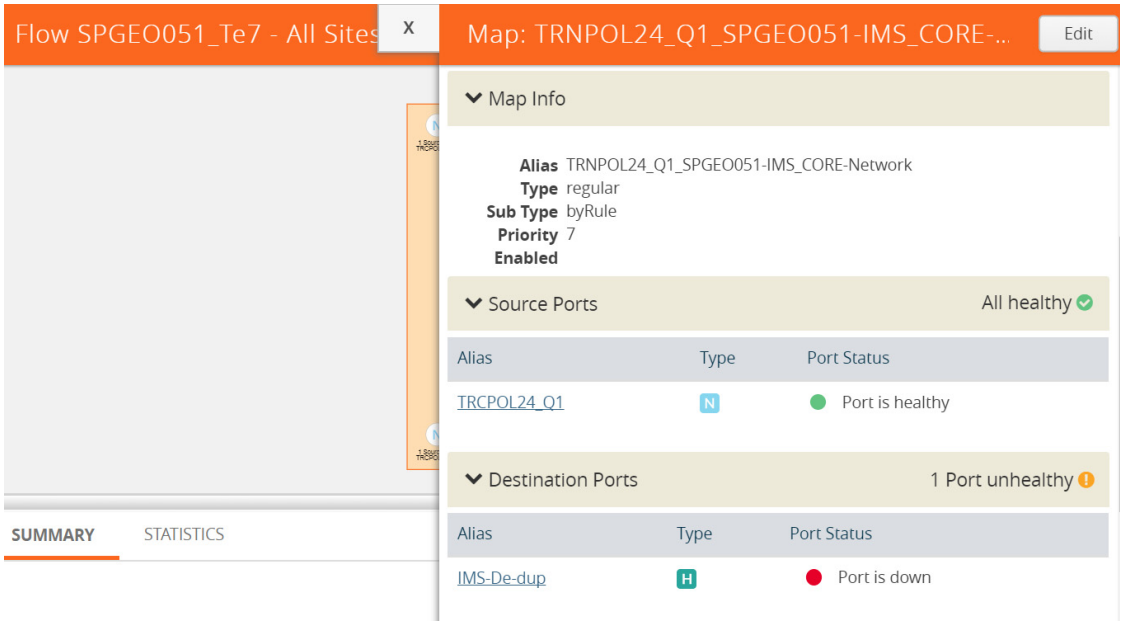

Figure 19-13: Flows - Map Quick View

View Unhealthy Maps

The total number of unhealthy maps involved in the flow are displayed in the Total Unhealthy Maps link.

In the Flows page, click the **Total Unhealthy Maps** link.

The Unhealthy Maps quick view displays a list of all unhealthy maps that are involved in the flow. It also provides information such as map alias, type, map health status, and the priority node on which the map is configured. Refer to Figure 19-14 on page 358.

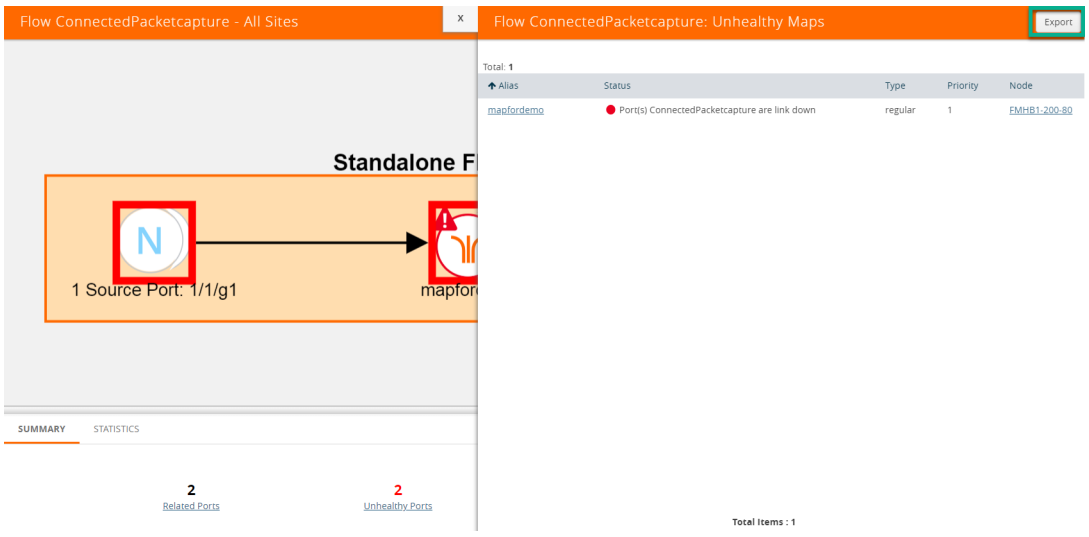

Figure 19-14: Flow - Unhealthy Maps

In the Unhealthy Maps quick view you can:

- Click **Export** to export an Excel report listing the unhealthy maps.
- Click a a map alias link in the Alias column to open the Map quick view. Refer to Figure 19-13 on page 358.
- Click a node under the Node column to open the Overview page.

Filter Flows

To filter flows by name, do the following:

1. Click **Physical** in the top navigation link.
2. (Optional) From the sites drop-down box, select a site.
3. In the left navigation pane, click **Flows**. The Flows page is displayed.
4. Click the **Filter** button at the top of the pane. The Filter panel displays.

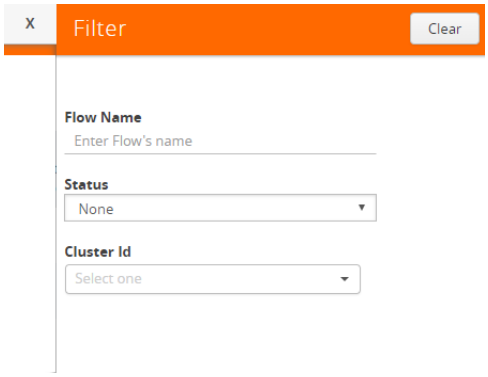

Figure 19-15: Flow Filter Pane

5. Start typing the first few characters of the name of a flow you wish to filter on. The flow list is sorted and all the flows containing the characters you entered displays in order at the top on the page.

Figure 19-16: Flow Filter - Names

To filter flows based on their status:

6. Click the **Status** drop down and select one of the Status colors to display only those flows that correspond to the color selected.

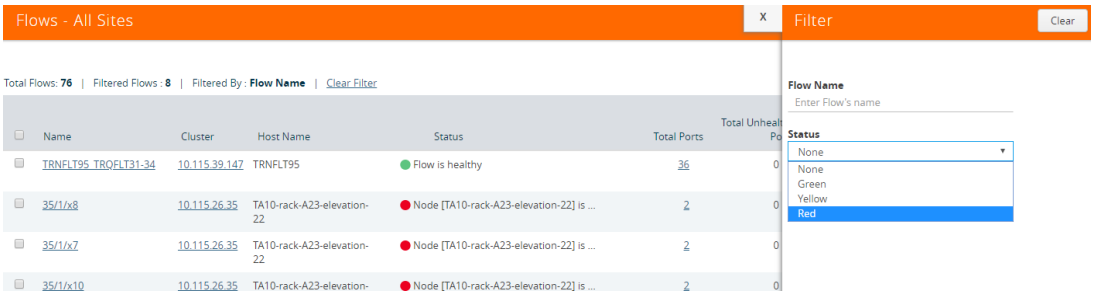

Figure 19-17: Flow Filter - Status

To filter flows based on the Cluster ID:

7. Click the **Cluster ID** drop down and select one of the Cluster IDs to display only those flows that correspond to the Cluster ID you selected.

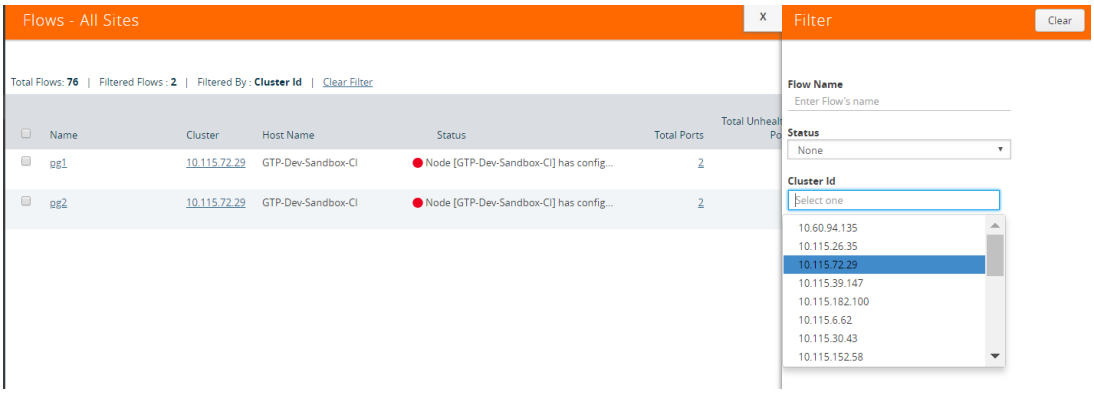

Figure 19-18: Flow Filter - Cluster ID

How to Change the Flow Layout

By default, a flow layout is arranged in a directional flow. The ingress network port receiving the traffic from the network TAP is aligned to the left and the tool port sending the traffic to the monitoring tool is aligned to the right. The maps configured on the standalone nodes and clusters are displayed within the respective containers. The arrows in the layout shows the direction in which the packets are moving from network ports to tool ports. You can drag and drop the cluster or node containers to change the alignment. You can also drag and drop the ports and maps within the container to change the layout.

To reset, save, or restore the layout:

1. In the Flow view page, click and drag a port or a map to change the layout as desired.

2. To save the new flow layout, click **Actions > Save current layout**. The current alignment of the flow is saved.

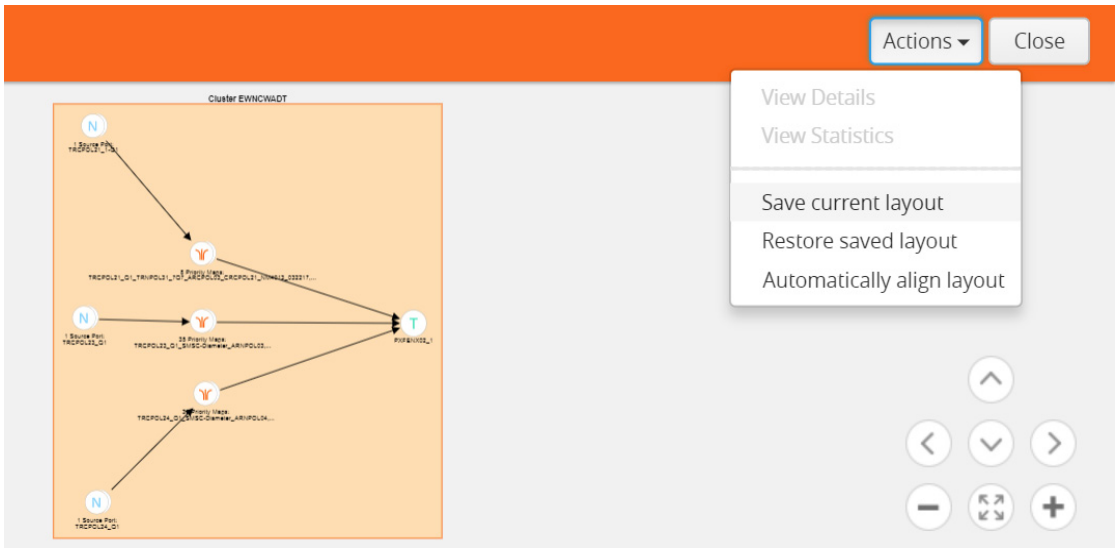

Figure 19-19: Save Current Layout

3. To align the layout to default, click **Actions > Automatically align layout**. The layout is automatically set to the default.

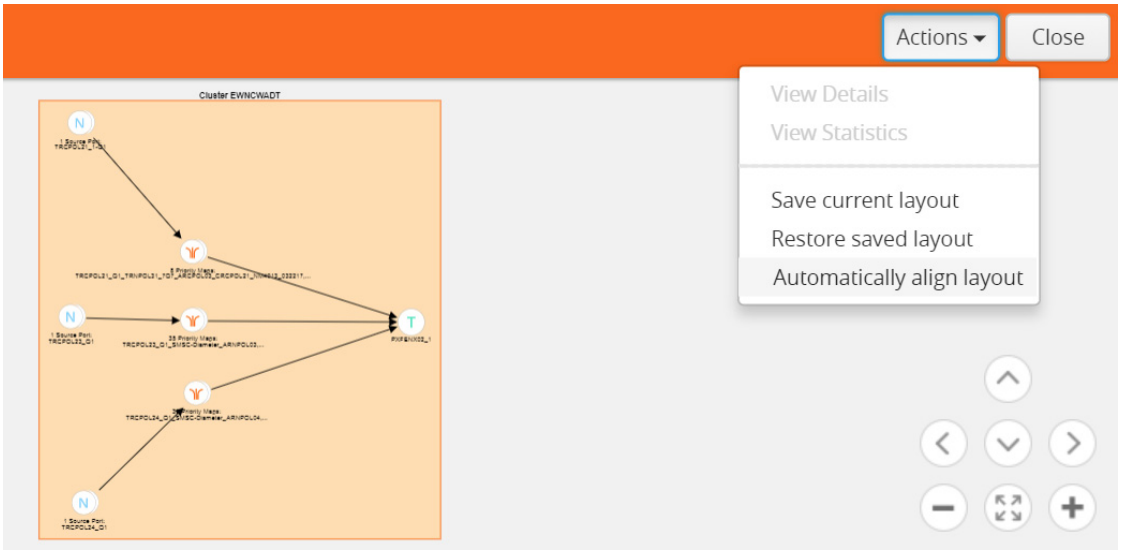

Figure 19-20: Automatically Align to Default Layout

4. To restore the saved layout, click **Actions > Restore saved layout**.

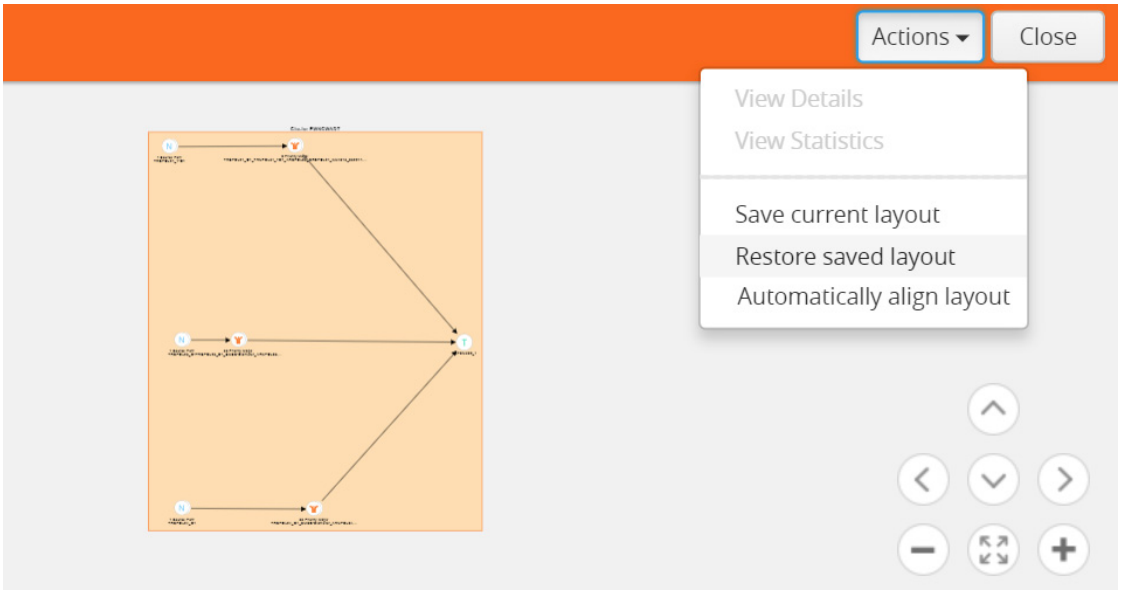

Figure 19-21: Restore to Saved Layout

How to Update Flows

Flows are computed based on the Gigamon discovery or manual links connecting the GigaVUE nodes and the maps participating in the flow. The changes made to the flow are not seen unless they are automatically updated after every 24 hours or manual updated from the Flows page. In release 5.3, flows can be updated

NOTE: It is recommended to update flows on a per-site basis for best results. If you update flows without selecting a site, the action will update all flows across all sites, which will take a very long time and is not a recommended practice. Refer to *Update Flows for a Selected Site* on page 363.

To manually update a flow:

1. Click **Physical** in the top navigation link.
2. (Optional) From the sites drop-down box, select a site.
3. In the left navigation pane, click **Flows**. The Flows page is displayed.

   Important: if you update flows without selecting a site, all flows across all sites will be updated, which will take a long time. This is supported but is not recommended. To update a flows on a per-site bases, refer to *Update Flows for a Selected Site* on page 363.
4. Click **Actions** > **Update**. A message is displayed to indicate that the flow calculation is in progress. The check boxes to select a flow disappear when the flow

is being updated. They slowly start appearing when the update is completed for that particular flow.

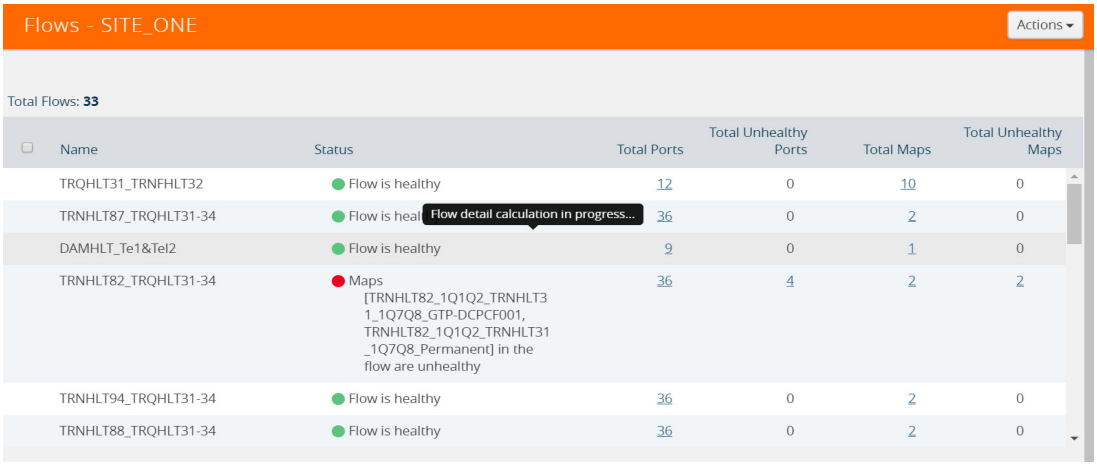

NOTE: All components of a flow must belong to the same site. If a GigaVUE node does not belong to the same site as compared to the node connected to the tools, then an error message is displayed.

Update Flows for a Selected Site

This process shows how to update flows on a per-site basis. This is the recommended method for updating flows.

Flows are automatically updated every 24 hours or can be manually triggered if you believe the underlying components have changed, such as map source or destination ports, GDP links or manual links. The following steps describe how to perform a manual update at a specific site.

To manually update flows for a selected site, do the following:

1. Click **Physical** in the top navigation link.
2. From the sites drop-down box, select a site where you wish to update the flows.

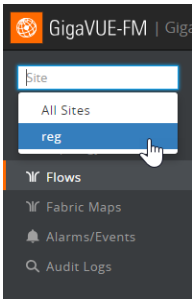

Figure 19-22: Updating Flows - Selecting Specific Sites

3. Click the Update link under the action menu drop down in the right corner of the screen.

The flows you selected are updated and the **Last Compute Time** column displays the time and date of the update.

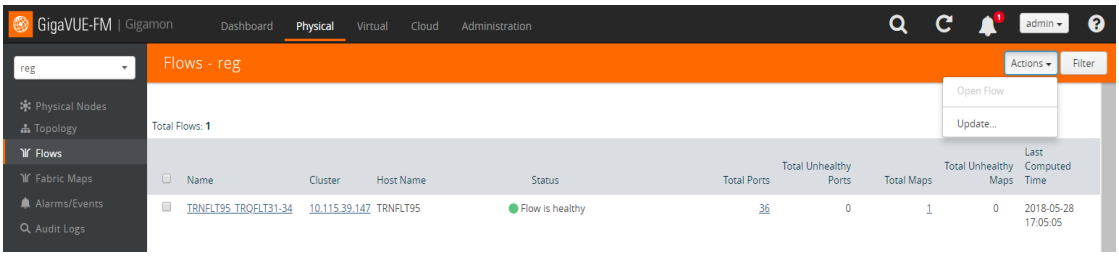

Figure 19-23: Updating Flows - Specific Sites

View Alarms and Events

On the Alarms and Events page, the FlowHealthStateChange event type indicates that there is a change in the health status of a flow. For more information about Alarms and Events, refer to *All Alarms/Events* on page 1269.

Refer to Figure 19-24 on page 364 for the flow health change event on the **Alarms/ Events** page.

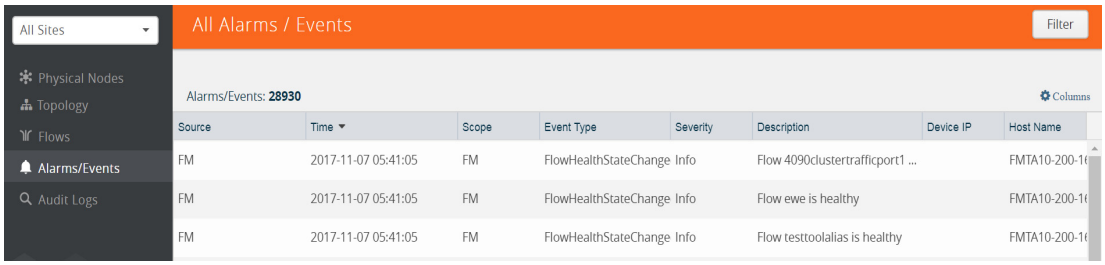

Figure 19-24: Flow Health Event Type

Set Notifications

On the Notifications page, the email notification for the **Flow Health State Changed** event type can be configured to automatically send emails to the specified addresses when there is a change in the health of any of the flows managed by GigaVUE-FM. For more information about configuring notifications, refer to *Notifications* on page 1320.

Refer to Figure 19-25 on page 364 for the flow health state changed notification on the **Notifications** page.

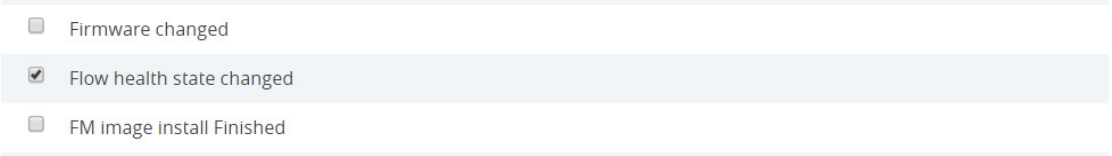

Figure 19-25: Flow Health State Changed Notification

Limitations of Flows

If Gigamon Discovery is enabled on GigaVUE nodes and a link is created between a source and a destination node that are of hybrid port types, the flow ends at the source node hybrid port. As hybrid port acts as an indirect traffic source port and as a tool port, the flow construction fails to identify the direction of the traffic flow. The workaround is to add a manual link to indicate the direction of the traffic flow.

20 Device Logs and Event Notifications

This chapter describes how to view the logs for the selected node.

Refer to the following sections:

Stream Device Logs to GigaVUE-FM

GigaVUE H Series device/cluster nodes provide comprehensive logging capabilities to keep track of system activity. Logging is particularly useful for troubleshooting system issues, as well as maintaining an audit trail. You can specify what types of events are logged, view log records by priority, date, or name, and upload log files to a remote host for external troubleshooting.

When H Series devices are added to GigaVUE-FM, the default settings ensure that log records stream to the GigaVUE-FM server and are written on the device's file system in a messages file. Log records can be used to analyze the system behavior directly from the GigaVUE-FM interface instead of needing to sign in to each device.

External third-party servers can also be added to host the streamed logs. Refer to *Add an External Logging Host Server to a Node* on page 374.

In this section:

Cluster Behavior

Logging configuration is local to the node. Each node has its own logging configuration.

Important: This is a different behavior than in software versions prior to 5.3 in which the logging configuration on the master node was synchronized to the other nodes in the cluster.

Standardized Logs

Standardized log messages can be streamed to GigaVUE-FM. The log messages follow the industry standard described in RFC5424. The format includes structured data, timestamp, version, and message ID. The timestamp includes milliseconds for increased accuracy.

In GigaVUE-FM, the log information is displayed in a table with configurable, sortable columns and extensive filter options.

Refer to *View Device Logs* on page 371 for information about viewing and filtering logs in GigaVUE-FM.

Refer to the *GigaVUE-OS-CLI User's Guide* for information about the standardized log message format in the CLI view.

Device Log Categories

Table 20-1 describes the categories in device/cluster log messages as per the standardized log format. Refer to *Standardized Logs* on page 368 for standardization information.

Table 20-1: Device Log Categories

Message ID	Description
GENERAL-ERR	Errors that are common across applications
HIGH-TEMPERATURE	High Temperature
PACKETDROP	Packet Drop
LINK	Link
INIT	Initialization
RESOURCE-UTIL	Resource Utilization
REQUEST	Query system information

Device Log Message Types

Table 20-2 describes the message types in device/cluster log messages as per the standardized log format. Refer to *Standardized Logs* on page 368 for standardization information.

Table 20-2: Device Log Message Types

Message ID	Description
memoryAllocError	Memory allocation error
memoryAccessError	Unexpected NULL access
fileAccessError	File access error
switchCpuHighTemperature	High Switch CPU Temperature
opticsHighTemperature	High Optics Temperature
gigasmartCpuHighTemperature	High GigaSMART CPU Temperature
ambientCpuHighTemperature	High Ambient CPU Temperature
exhaustCpuHighTemperature	High Exhaust CPU Temperature
egressPacketDrop	Packet Drop at Egress
ingressPacketDrop	Packet Drop at Ingress
linkChangeNotify	Link Change Notification
mgmtModule	Failed to load mgmt module
systemCpuUtil	High System CPU Utilization
systemMemoryUtil	High System Memory utilization
processCpuUtil	High CPU Consumption by a process
processMemoryUtil	High Memory Consumption by a process
processAccessError	Process access errors
systemAccessError	Process access errors
queryFail	Failed to fetch system data
moduleInitFail	Failed to load a cli module

Device Logging Levels

Table 20-3 shows the standard logging levels that are used to rank logged events by degree of severity. When configuring the device/cluster logs to stream, the highest severity level will capture fewer logs than the lowest severity. For example, if the lowest level (info or debug) is selected in your configuration, all levels of logs, from the

selected level up to the highest level, will be streamed. If the highest level (emergency) is selected, only emergency logs will be streamed.

Table 20-3: Logging Levels

Log-Level	Description
emergency	Emergency – the system is unusable. The severity level with the least logging – only emergency level events/commands are logged.
alert	Action must be taken immediately.
critical	Critical conditions.
error	Error conditions.
warning	Warning conditions.
notice	Normal but significant condition.
info	Informational messages.
debug	Debug-level messages. Authorized for factory use only.

Device Logging Processes

Logs that are specific to a process are logged with the process ID. Table 20-4 provides a list of logging processes and their descriptions.

Table 20-4: Logging Processes

Logging Process	Description
acctd	AAA Accounting daemon
avd	Active Visibility daemon
cli	Command Line Interface
clusterd	Cluster daemon
debuggabilityd	Debuggability daemon
frm	Foreign Resource Manager
gsd	GigaSMART daemon
gprof	Profiler
httpd	HTTP daemon
licd	License daemon
mgmtd	Management daemon
ndiscd	Network Discovery daemon
netdevd	Netdev daemon
notf_mgr	Notification Manager
ntpd	Network Time Protocol daemon
peripd	Peripheral daemon
persistd	Persistence daemon

Table 20-4: Logging Processes

Logging Process	Description
pm	Process Manager
ptpd	PTP Protocol daemon
restapid	REST API daemon
sched	Scheduler daemon
snmpd	SNMP daemon
statsd	Statistics daemon
syncd	Sync daemon
syshth	System Health
ugwd	Unified Gateway daemon
wizard	Wizard
wsmd	Web session Manager daemon
xd	XML Gateway
xinetd	Extended Internet Service daemon

View Device Logs

GigaVUE-FM provides the ability to view logs for each device/cluster node. These logs list all the user events and enable Gigamon Technical Support to troubleshoot in case of any problems.

NOTE: Log streaming is supported on H Series devices only.

In addition to describing how to view logs, this section describes the following supported functionality related to understanding and configuring the Logs view.

- *Arrange Columns in the Logs View* on page 373

To view the logs of a single node:

1. Click **Physical** in the top navigation link.

   The Physical Nodes page displays the list of physical nodes managed by GigaVUE-FM.

2. Select a GigaVUE node.

   The Single Node view displays overview statistics about the selected node.

3. In the left navigation pane, click **Logs**.

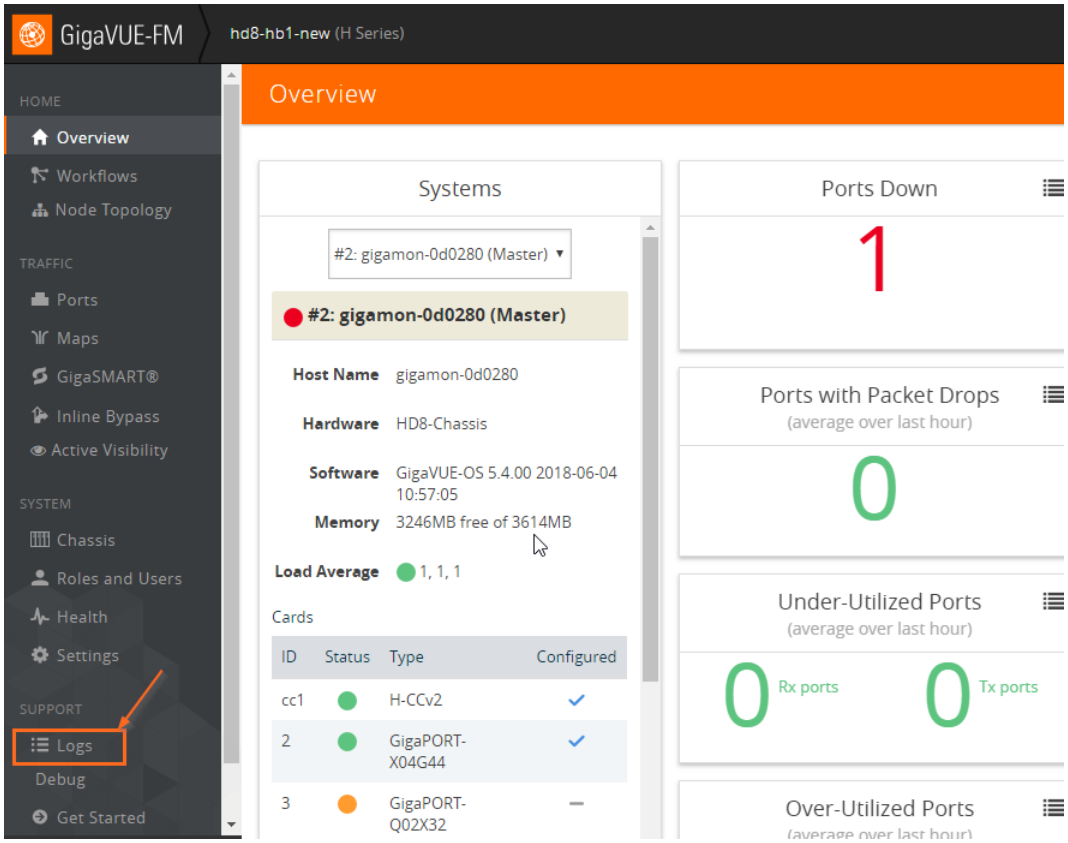

Figure 20-1: Overview of a Single Node

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

4. (Optional) Click the heading cell for any column to sort the list by that attribute. Refer to *Arrange Columns in the Logs View* on page 373 for details.

5. (Optional) Click **Filter** to refine the list of logs by any of the available attributes. The Filter quick view displays the attributes of the logs.

   a. Enter a value in any of the available attribute fields to narrow the list of logs according to that value.

Filterable Attributes	Description
Host Name	Enter a hostname or partial host name.
Device IP	Enter a device IP address or partial address.
Time	Specify a start date and time and an end date and time to define a time period. **IMPORTANT:** Timestamps are shown in the time zone of the client browser's computer and not the timezone of the node reporting the event. The timestamp is based on the correctly configured clock on the GigaVUE-FM server and converted from UTC to the client computer's configured timezone.

Filterable Attributes	Description
Process	Specify a process to target, such as clusterd, httpd, ntpd, restapid, snmpd, sshd, ugwd, or wsmd. Partial text entries are acceptable. Refer to *Device Logging Processes* on page 370
Log Level	Select one from the list of log levels, which are ranked by degree of severity. Refer to *Device Logging Levels* on page 369.
Category	Select a category of issue from the list of values, such as: High Temperature, Link, Packet Drop, or Resource Utilization. Refer to *Device Log Categories* on page 368.
Priority	Enter a priority number as defined by RFC 5424.
Version	Enter a version number.
Type	Select an event type from the list of values, such as CPU utilization high, memory utilization high, egress or ingress packet drop, and so on. Refer to *Device Log Message Types* on page 369.
Affected Entity Type	Select an affected entity type from the list of values. Options are: CPU, Memory, Port, or Optics
Affected Entity	Enter a specific affected entity

b. Click **Apply Filter** to apply the filter on the current list of logs.

   The list of logs is now limited to the set matching your filter criteria.

c. To clear the filter, click the “Clear Filter” link above the table.

Arrange Columns in the Logs View

The Logs view has sortable, configurable columns. The options below describe how to use these options to personalize your view of the logs:

- Click the heading cell for any column to sort the list by that attribute.
- To move the order of a column, click the heading cell for that column and drag it to the location you prefer, then release the mouse click to drop it in the new location.
- Configure the columns to display in this view:

  a. Click **Columns** on the right above the table.

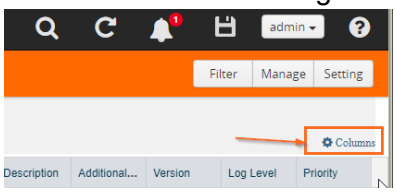

  b. The COLUMNS window appears.

  c. Amend your column settings:

     To remove a column, select the column label in the selected columns table and click **Remove**.

     To add a column, select the column label in the Available Columns table and click **Add**.

To to resort the order in which the columns appear, select any label and click **Move Up** or **Move Down** to reposition the column.

To revert to the GigaVUE-FM default setting, click **Reset to Defaults**.

d. When you are done configuring the columns, click **OK**.

Device Log Host Servers

All device/cluster log messages for all H Series nodes are streamed to the GigaVUE-FM server by default and are written on the device's file system in a messages file.

External third-party servers, like Splunk, can be added to host the streamed logs. You may want to have certain types of log messages streamed to different servers for different purposes. The Log Settings enable you to add additional host servers for the log streams and to configure the type of logs each server will host.

When an external server is specified, the GigaVUE H Series node will send logged events through UDP, TCP, or SSH to the specified destination.

In this section:

NOTE: For a description of log levels, refer to *Logging Levels* on page 370. For a description of Syslog Server configuration options, refer to *Host Server Options* on page 375.

Add an External Logging Host Server to a Node

This topic describes how to configure a system log server as a destination for logging in GigaVUE-FM.

To add a host server:

1. Navigate to **Physical Nodes** and select a node.
2. Click **Logs** from the Single Node view.
3. Click **Settings**.

   The Log Settings view appears.
4. Click **Add Server**.

   The Add Log Server quick view appears.
5. Specify the server attributes of the log server you wish to add. The parameters vary based on the protocol you select.

   Select the logging protocol: **UDP (default)**, **TCP**, or **SSH**.

   For **UDP**, do the following:

   a. Enter the external server's IP address in the **IP Address** field.

b. Select the logging level from the **Log Level** list. For a description of the logging levels, refer to Table 20-3 on page 370.

For **TCP**, do the following:

a. Enter the external server's IP address in the **IP Address** field.

b. Select the logging level from the **Log Level** list. For a description of the logging levels, refer to Table 20-3 on page 370.

c. Enter the port number in the **Port** field.

For **SSH**, do the following:

a. Enter the external server's IP address in the **IP Address** field.

b. Select the logging level from the **Log Level** list. For a description of the logging levels, refer to Table 20-3 on page 370.

c. Enter the port number in the TCP Port field.

6. Click **OK** to save the settings and add the server.

   The additional server appears in the server list.

7. Specify the optional settings for the server. Refer to *Host Server Options* on page 375 for a description of the options.

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

Host Server Options

The following table describes the optional host server settings for device logs.

Table 20-5: Syslog Server Settings

Setting	Description
Logging	This option enables or disable the server. • Click **Enable** to enable the server. (Default) • Click **Disabled** to disable the server. • Click **Delete** to remove the server from this node.
Server Log Level	Select the level of alert you wish to tack on this server: • Info • Notice • Warning (default) • Error • Alert • Critical • Emergency **NOTE:** Refer to *Device Logging Levels* on page 369.
Protocol	Select the protocol for communicating with this server: • TCP • UDP (default) **NOTE:** GigaVUE-FM can receive logs in TCP as well as UDP on port 5672. However by default it is UDP.
Port	Enter the Port Number.

Table 20-5: Syslog Server Settings

Setting	Description
Ssh Enabled	Check the **Ssh Enabled** check box to enable Ssh on this server.

Edit Host Server Settings

To edit log settings:

1. Navigate to **Physical Nodes** and select a node.
2. Click **Logs** from the Single Node view.
3. Click **Settings**.

   The Log Settings view appears.

   The top portion of the Logs Setting page shows the global Storage Management settings. The table in the lower portion lists the servers
4. Edit the settings for the server directly in the server list table. Refer to *Host Server Options* on page 375 for a description of the options.
5. Click **OK** to save the settings.

Storage Management for Device Logs

Device/cluster log messages for H Series nodes in GigaVUE-FM are continuously being recorded. As a result, the logs can take up a lot of storage space over time. You may want to delete old records on a regular basis to clear-up storage space. You may want to export log records as a back-up before performing a delete operation or to preserve for external analysis.

GigaVUE-FM Storage Management allows you to define how the stored logs are managed. You can specify a schedule for purging old device logs. You can also specify an SFTP server to export the log records prior to purging.

NOTE: GigaVUE-FM Storage Management is used for all storage settings, including device logs, alarm/event notifications, and statistics. This topic, however, describes the storage management for device logs, in particular.

Access Storage Management

There are two ways to access the Storage Management settings for device logs. Aside from the path, the instructions are the same. Both paths allow you define the settings across nodes. Neither setting allows you to define device-specific storage settings. These paths are described in Table 20-6 on page 376.

Table 20-6: Accessing Storage Management

Path
Click to **Administration** from the top navigation link, then select **System** > **Storage Management**. Refer to *Storage Management* on page 1334.

Table 20-6: Accessing Storage Management

Path
Click **Physical** on the top navigation link and select **Physical Nodes** from the left navigation pane. Select a node, then select **Logs** > **Settings**. From the Log Settings page, click the “Edit” link next to the export or delete setting summary.

NOTE: To perform a one-time manual clean-up of old logs on a specific node, refer to *Manage Device Log Output* on page 377.

To access Storage Management settings from the node’s Logs page:

1. Click **Physical** on the top navigation link and select **Physical Nodes** from the left navigation pane.
2. Select a node.
3. Click **Logs** from the left navigation pane in the node view.
4. Click **Settings**.

   The Log Settings view appears.
5. To edit the Storage Management settings for the device logs, click the “Edit” link next to the export or delete setting:

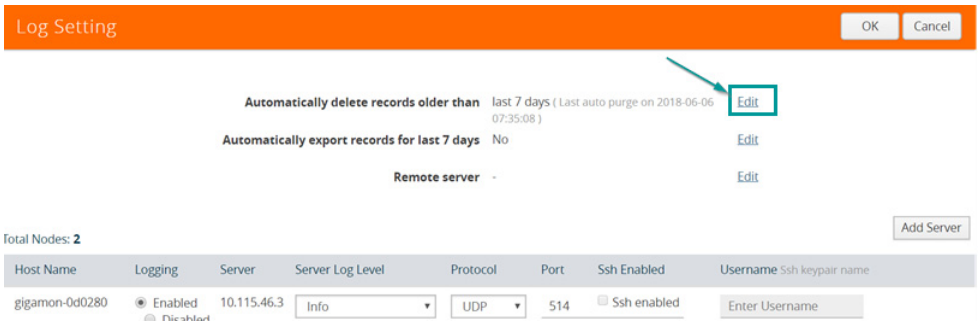

   The Storage Management page under Administration appears.

   NOTE: This setting is not node-specific. This setting applies to all logs in GigaVUE-FM across devices. Refer to *Storage Management* on page 1334.

Manage Device Log Output

Log management allows you to export and delete logs that are stored in GigaVUE-FM. Use this process to manually clean-up old logs on a specific node before a specified cut-off date and time. You can also specify an SFTP server and file path to export the records prior to purging.

NOTE: Refer to *Storage Management for Device Logs* on page 376 for global Storage Management settings that control purging of old logs on a scheduled basis across all devices.

To edit the log management settings:

1. Navigate to **Physical Nodes** and select a node.
2. Click **Logs** from the node view.

3. Click **Manage**.

   The Manage Logs view appears.

4. Complete the fields on this page to specify how to manage the logs:

Setting	Description
Time Range	
Select records older than	Specify a cut-off date and time for deleting log records. Log records that were created prior to the specified date and time will be deleted. The time is based on the current timezone in GigaVUE-FM.
Export Records To	
SFTP Server Address	The records are exported to a CSV file. Specify the ftp/sftp location to send the CSV files. For example: `sftp://username@121.0.0.1/path/directory`
Username and Password	If this is a secure server, which is recommended, specify the username and password for accessing the server.
File Path	Specify the path on the server to store the file. For example: `/root/dir/archive.zip`
Purge Selected Records	
Check this check box to enable purging the selected records. **Important**: If this option is selected, records will be deleted immediately when you click **OK**.	

5. Click **OK** to save the settings.

Device Event Notifications

Events can be streamed to GigaVUE-FM. Each device (GigaVUE node) can directly stream events. In software version 5.3, the following events are transmitted:

- port up and port down
- GigaSMART packet drops
- GigaSMART application core crash when a back trace trigger is initiated
- GigaSMART application core crash when a soft reset is initiated
- CPU utilization in the GigaSMART cores when the high threshold is exceeded

NOTE: The events described in this section are not SNMP traps.

Cluster Limitation

Event Notifications require that network connectivity be available from each node in the cluster to its managing GigaVUE-FM, not just to the master node.

Configure Device Event Notifications

When H Series devices are added to GigaVUE-FM, GigaVUE-FM is added as a notification target on the attached device. This ensures that event notifications from the device stream to the GigaVUE-FM server. The events that are streamed are stored in the GigaVUE-FM database, and are available for querying using the filters in the Alarms/Events page. Refer to *All Alarms/Events* on page 1269.

A list of the configured event notification targets for the device/cluster can be seen on the Event Notification page in the physical node settings. To access this page, click **Physical** on the top navigation link. On the left navigation pane, select **Physical Nodes**. Double-click on a node ID to open it. From the node, navigate to **Settings > Global Settings > Event Notification**.

The Event Notification page lists all target hosts for Alarm/Event notifications from this device/cluster. This page allows you to control (enable/disable) whether a target host receives notifications from this device/cluster. You can also view the Alarm/Event purge and export settings defined under Storage Management.

To manage notifications for a device/cluster:

1. Select **Physical > Physical Nodes**.
2. Click a Cluster ID link to open the node.
3. From the node, navigate to **Settings > Global Settings > Event Notification**.

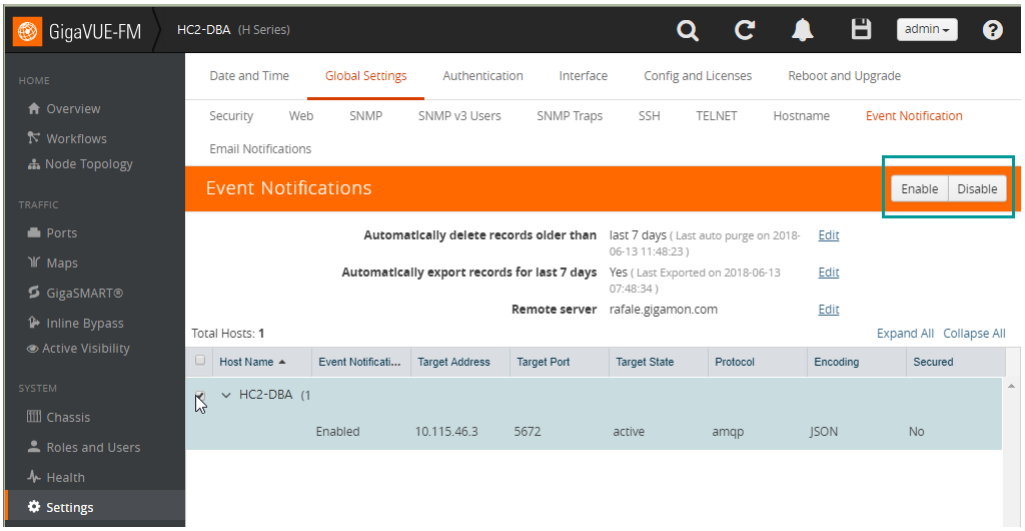

Figure 20-2: Device/Cluster Event Notification Settings

Storage Management settings for all system events are summarized above the table.

4. To access Storage Management settings for all system events:
   a. Click the “Edit” link next to any of the summary purge and export settings.
   b. Refer to *Storage Management* on page 1334.
5. To control notifications from this device:
   a. Click the check box next to the host name.

b. Click **Enable** or **Disable** to control whether notifications from this device are sent to the selected host.

System Notification Settings

To access the Notifications page that spans all nodes, click **Administration** on the top navigation link. On the left navigation pane, select **System > Notifications**.

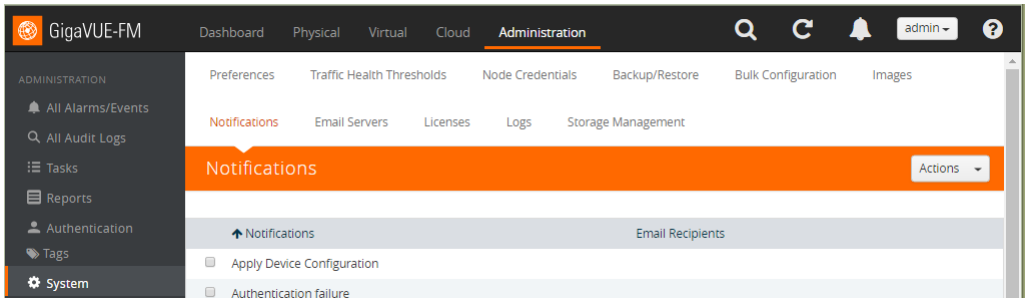

Figure 20-3: System Event Notification Settings

Refer to *Notifications* on page 1320 and *Configure Email Notifications* on page 1320.

REST API

NOTE: Events and logs go into a database on GigaVUE-FM called Elastic Search. You can use the REST API to access the database and become a client for logs or events.

21 Backup/Restore

This section describes the Backup/Restore feature of GigaVUE-FM, which provides the ability to backup the current configuration or to restore from a list of available configurations on GigaVUE-FM.

GigaVUE-FM lets you back up and restore the configuration of all of the managed GigaVUE nodes, including H Series, TA Series, and G Series. This section also provides the steps for backup and restore step that you can use for each of these platforms.

Backing up and restoring nodes is a time consuming process. At the end of the backup or restoring process an event is posted that indicates a success or failure of the backup. For G Series nodes, you can also use the Bulk Configuration feature. For more information, refer to *Bulk Configuration* on page 1313.

This section covers the following:

Nodes and Cluster Backup

This section describes how to backup H Series, TA Series, and G-Series nodes. When a node is backed up, the backup file is saved in local storage on the machine where Fabric Manager is installed. The filename is the timestamp of the backup. Starting from 5.5, backups are in text based and binary formats. For security reasons, text configuration files do not include plain text passwords, such as SMTP passwords, AAA keys (RADIUS or TACACS+), or private keys in RSA/DSA identities. When a cluster is backed up, a backup file is created for the master only.

You can schedule node or nodes and node clusters for immediately backup or schedule backups to occur at a specified time. For example, you can schedule a backup for a particular day, week, month, or date.

Notes:

- Prior to GigaVUE-FM 3.2, backup file for physical nodes were in a binary format. Starting with GigaVUE-FM 3.2 backup and restore files use a text based format and binary backup or restore on physical nodes is not supported. If you are upgrading from a version lower than 3.2, you can backup your configuration prior to upgrading to the current version of GigaVUE-FM if you desire, but the files will be in a binary format. Existing binary backups are not visible to the current version of GigaVUE-FM. For binary backups, you must back up the node using the CLI commands rather than GigaVUE-FM. For more information about the CLI commands, refer to the *GigaVUE-OS CLI User's Guide*.
- Clusters can be backed up only if the Master node in the cluster is licensed.
- For clusters with software version 4.6 or lower and a nat-enabled setup, GigaVUE-FM does not support the backup/restore operation. For nat-enabled clusters with software version 4.7 or higher, GigaVUE-FM supports backup/restore.
- For a cluster, the cluster name is the actual cluster's name. For example, Gigamon-Cluster.
- For a restore operation on a cluster, cluster name changes are not supported. The cluster name must be the same as when the backup was made.
- For standalone devices the cluster name is the IP of the device.

Enable Events for Backup

If you want to see fine-grained events on the node during the backup process, you need to enable the *configuration save* SNMP trap. To enable the trap, do the following:

1. Click **Physical** on the top navigation link. On the **Physical Nodes** page, select the node on which you want to enable the trap.
2. Select **Settings > Global Settings > SNMP Traps**. The SNMP Trap page is displayed.
3. Click **Trap Settings.**
4. On the Edit SNMP Trap Setting page, select **Configuration Save**.
5. Click **Add**.

The system returns to the SNMP Traps page and displays an event message that the SNMP trap is enabled as shown in Figure 21-1 on page 383.

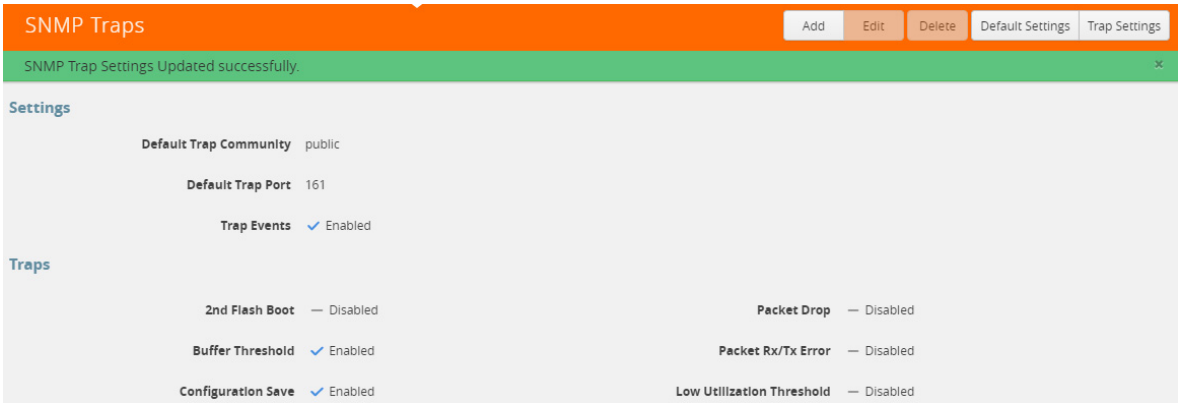

Figure 21-1: SNMP Trap Configuration Save Enabled

BackUp Nodes and Clusters

To backup a node, nodes, or clusters, do the following:

1. Click **Physical** on the top navigation link.
2. On the Physical Nodes page, select the node, nodes, or clusters that you want to backup.
3. Select **Actions > Backup**. The Backup page displays, showing the nodes selected for backup.
4. Select one of the following:
   - **Immediate—** Allows the back up to occur immediately.
   - **Scheduled—**Allows you to schedule a time for the backup or have reoccurring backups. For information about scheduled backup, refer to *How to Schedule Backups* on page 384 for details on how to create a schedule.
5. Click **OK**.

   If you selected **Immediate** in *Step 4*, the system returns to the Physical Nodes page and displays an event message about the start of the backup process as shown in Figure 21-2. You can also use the **Alarms/Events** to monitor progress.

   If you selected **Scheduled**, the next backup occurs according to the schedule.

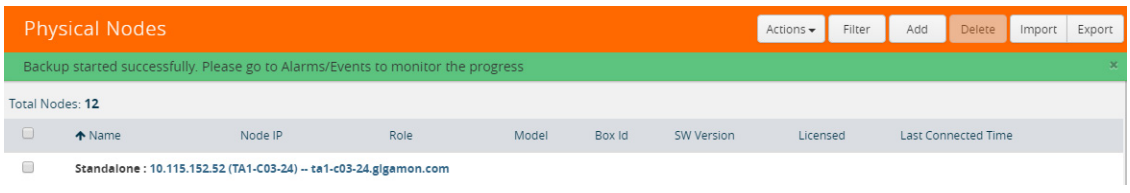

Figure 21-2: Backup Started Successfully

How to Schedule Backups

When creating a backup of nodes and clusters, you can create a schedule for performing regular backups of selected nodes and clusters. This allows you to backup the devices managed by GigaVUE-FM at best times, such as when you expect network traffic to be the least.

To set a schedule for backing up a nodes and clusters, do the following:

1. Click **Physical** on the top navigation link.
2. On the Physical Nodes page, select the **Node IP** for each node that you want to backup.
3. Click **Actions > Backup**.

   The Backup page shows the number of nodes selected for backup.
4. Select **Scheduled**.

   The GigaVUE-FM time clock is based on ESX host time. Make sure that you have synchronized clock before any scheduling operation.

   Figure 21-3 shows an example of nodes with scheduled backups. In this example, the weekly backups start on December 03 and occurs every Saturday at 8:30 pm until December 28.

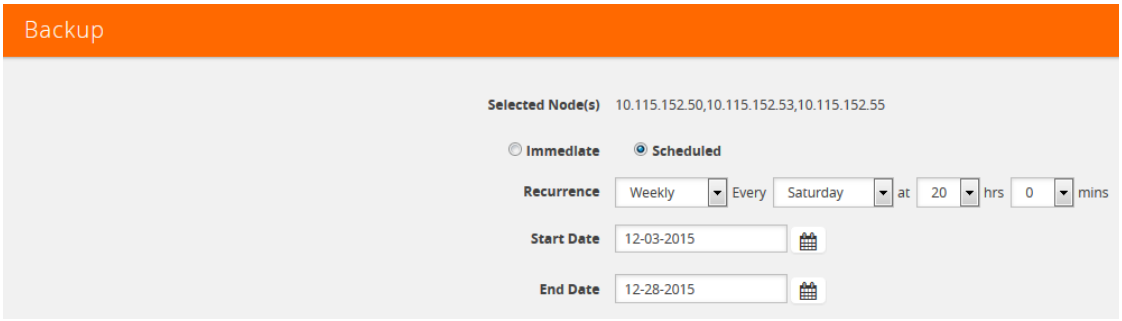

Figure 21-3: Nodes Selected for Scheduled Backups

5. From the **Recurrence** drop-down list, select one of the following:

Table 21-1: Recurrence Options

Option	Description
Once Only	Select this option for scheduling one time backup. Set a start date and start time for the backup to begin.
Daily	Select this option for scheduling daily backups. Set a start date and time for the backup to recur once a day. Set an end date to determine until when the backup must recur.
Weekly	Select this option for scheduling weekly backups. Set a day, time, start date and end date for the weekly backup to recur.

Table 21-1: Recurrence Options

Option	Description
Monthly	Select this option for scheduling backups once a month. Set a specific day of the month for the backup to recur. For example, if you want the backup to occur on every 15th day of the month, select 15th. Set a time, start date, and end date for the monthly backup to recur.
Yearly	Select this option for scheduling backups once a year. Set a specific day, month, time, start date, and end date for the yearly backup to recur.

6. Click **OK**. To monitor the progress of the event select Alarms/Events in the main navigation pane.

Once you have scheduled a recurring backup, the scheduled backup will appear as a scheduled task on the Scheduled Tasks page. To view tasks, select **Tasks > Scheduled Tasks**.

Download Backup Files

Because backup files are text-based and binary format, you can edit them in a text editor. You can restore the device configuration in binary format and view the configurations in text format. This is useful when you want to make modification before restoring the backup such as an error occurring during restore.

After creating a backup file as described in *BackUp Nodes and Clusters* on page 383, you can down load the file by doing the following:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **System > Backup/Restore > Physical Nodes** to open the Backup Files page.

   The Backup Files page lists the backup files created from a scheduled or immediate backup.
3. Click the “Show Config” link on the backup record row to view the file contents of the file you wish to restore.
4. From the preview panel, click **Download**.
5. The backup file will be downloaded to the local environment.

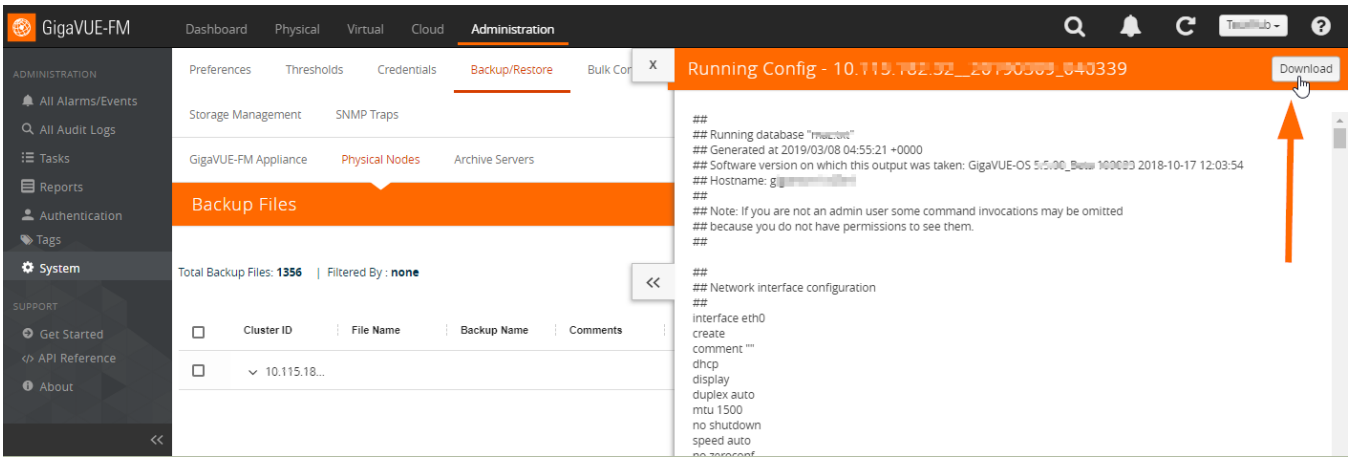

Figure 21-4: Backup File Selected for Download

Add Comments to Backup File

Starting in GigaVUE-FM 3.4, you can add a comment to the backup file that displays on the Backup Files page. A comment is useful for identifying particular backup files. To add a comment, do the following:

1. On the Backup Files page, select the file to which you want to add a comment.
2. Select **Edit**.
3. On the Edit page, enter a comment about the backup file in the **Comment** field.
4. Click **OK**.

   The comment is added to the comment field for the backup file. The following figure shows an example.

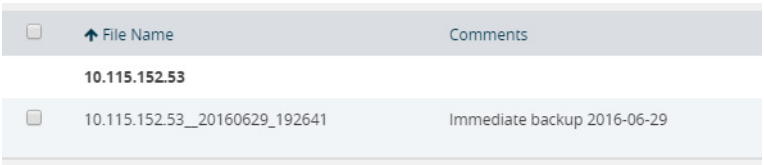

Set Do Not Purge Flag

NOTE: GigaVUE-FM runs a background task every 12 hours that purges the backup files and restore logs if the number of backup files for a node is greater than 10. The oldest backup files are purged first. You can set a **Do Not Purge** flag so that backup files are not removed when a purge occurs. Files with the Do Not Purge flag set are not included in the automatic purge.

To set Do Not Purge for a backup file, do the following:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **System > Backup/Restore > Physical Nodes**.
3. On the Backup Files page, select the backup file that you do not want to be purged.
4. Select **Actions > Enable Do Not Purge**.

A check mark appears in the Do No Purge field for the selected backup file.

To remove Do Not Purge for the backup file, select **Actions > Disable Do Not Purge**.

Delete Backup Files

To delete a backup file, do the following:

1. Click **Administration** on the top navigation link.
1. On the left navigation pane, select **System > Backup/Restore > Physical Nodes**.
2. Select one or more filenames on the Backup Files page.
3. Click **Delete**.

   The system displays a dialogue to confirm that you want to delete the file.
4. Click **OK**.

Backup files obtained from standalone nodes before they joined the cluster are called orphaned backup files. GigaVUE-FM does not allow you to delete an orphaned device-backup file. You must first delete the files from /var as well as from the database. Contact Gigamon Customer support to delete the files from the database. You can also refer to the knowledge base article for more details.

Node and Cluster Restore

Starting with GigaVUE-FM 3.2, backup files are text based and binary format. GigaVUE-FM restores the device configurations in binary format and allows you to view the configurations in text format. During the restore process, the commands listed in the configuration file are executed. If any error occur during the restore process, the text-based file makes it possible to edit the file and attempt to restore the configuration again by uploading and applying the modified file.

When restoring clusters, you can modify the file before uploading and applying it to the cluster. However, the modified file must have the same name as the backup file that was downloaded. If you change the file name, GigaVUE-FM will reject the file during the upload operation. The configuration file is applied to the current master node in a cluster and this node could be a different node than the one when the backup was done.

Notes:

- GigaVUE-FM does not support the restore operation if the cluster name changes. The cluster name should have the same name at the time of the restore operation as it did at the time of the backup operation.
- Text-based and Binary format backed up configurations created directly on the node, using either the CLI or H-VUE, are also available for restoring from the GigaVUE-FM.

Restore Nodes and Clusters

To restore nodes or clusters, do the following:

1. Click **Physical** on the top navigation link.
2. On the Physical Nodes page, select the IP address for each node or cluster that you want to restore.
3. Select **Actions > Restore**.

   The Restore From File page displays, showing the file names from which to restore.
4. Select the configuration to restore by clicking the check box next to the file name. Only one configuration can be selected with an restore action.
5. Click **OK**.

View Restore Logs

Restores are a binary-based restore and use a fail-continue option during the restore process. If any errors occur, they are logged to the a restore log file. You can download and view the restore logs by doing the following:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **System > Backup Files > Physical Nodes**. The Backup Files page shows the list of existing backup files.
3. Click the **Restore Log Files** link.

   The Restore Logs page displays the restore logs currently available. If no restore action has occured, the restore logs page will be empty.

Part 5: Traffic

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

To access traffic operations from the GigaVUE-FM interface:

1. Select **Physical** from the top navigation menu.

   This displays the list of Devices/Cluster Nodes managed by this instance of FM.
2. Click the Cluster ID of any node to open the node.

   The Traffic option is displayed on left navigation pane.

The following traffic management options are described in this section:

- *Ports* on page 391
- *Maps* on page 485
- *Inline Bypass Solutions* on page 563
- *Work With Inline SSL Decryption* on page 645
- Flexible Inline Arrangement (refer to *Work With Flexible Inline Arrangements* on page 707)
- GigaSMART (refer to *About GigaSMART Applications* on page 742)
- Application Intelligence (refer to *Application Intelligence* on page 727)
- Active Visibility (refer to *Configure Active Visibility* on page 542)

22 Ports

This chapter provides the following information:

About Ports

This section provides an overview of the various port types, describes the steps involved in configuring ports, and provides details about port filters and port status. This section includes the following major topics:

About Network and Tool Ports

Packets arrive at the Gigamon Visibility Platform at ***network ports*** and are directed to monitoring and analysis tools connected to ***tool ports*** by flow maps. Figure 22-1 illustrates the concept of data flows between network and

tool ports. Data arrives from different sources at the network ports on the left and is forwarded to different tools connected to the tool ports on the right.

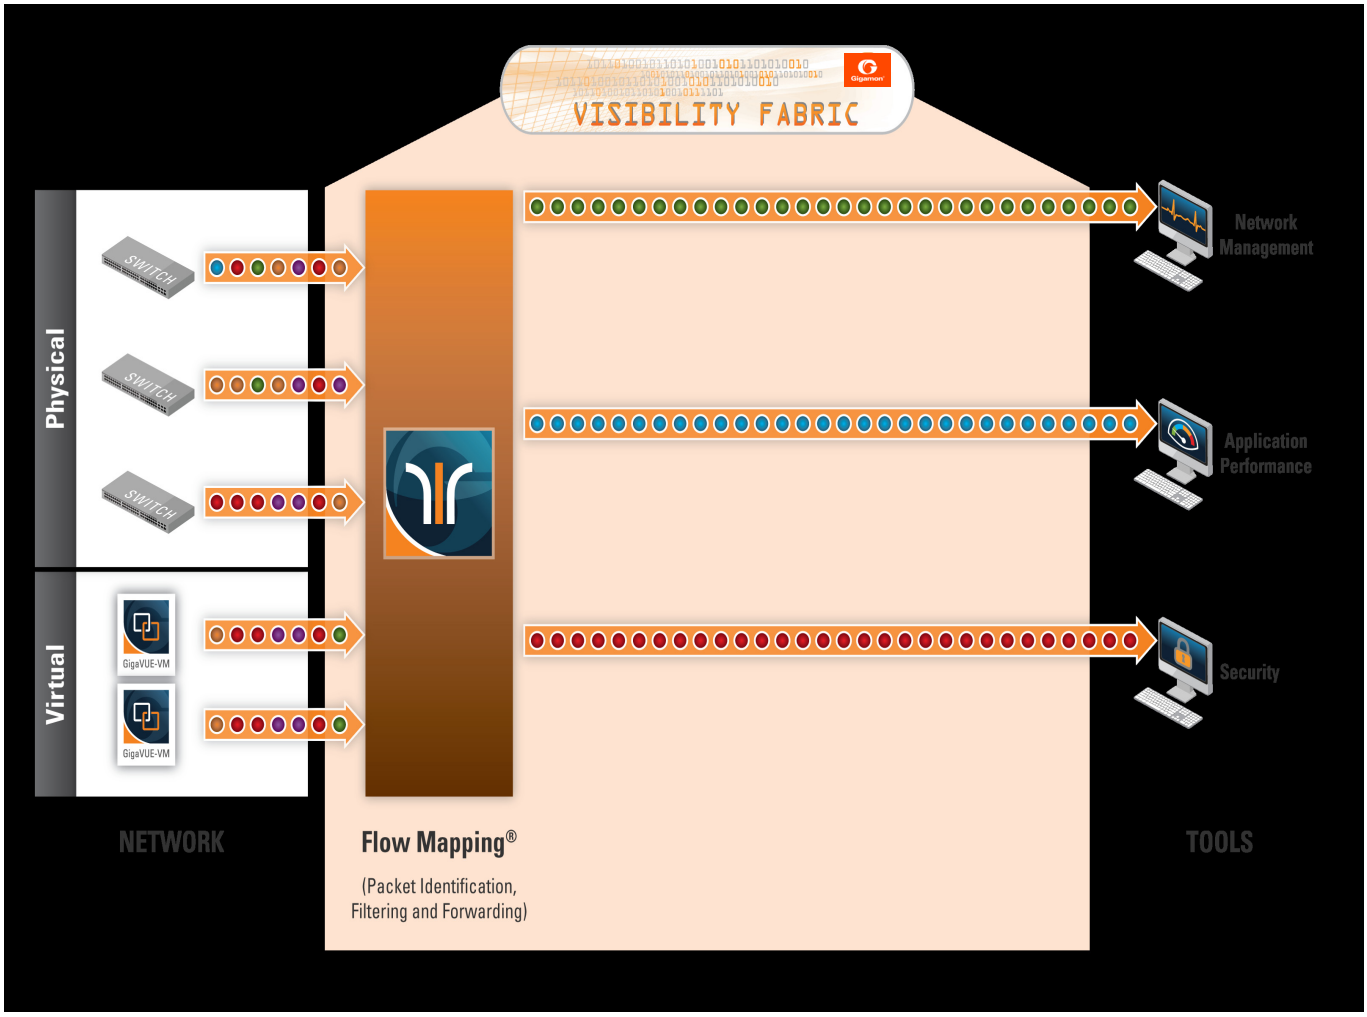

Figure 22-1: GigaVUE-OS Packet Distribution

Network (Ingress) Ports Defined

Network ports are where you connect data sources to GigaVUE nodes. For example, you could connect a switch's SPAN port, connect an external TAP, or simply connect an open port on a hub to an open port on a line card. Regardless, the idea is the same – network ports are where data arrives at the GigaVUE node.

NOTE: In their standard configuration, network ports only accept data input – no data output is allowed.

Tool (Egress) Ports Defined

Tool ports are where you connect destinations for the data arriving on network ports on GigaVUE nodes. For example, an IT organization could assign one set of tool ports to its Security Team for an intrusion detection system, a forensic data recorder, and a traditional protocol analyzer while a separate set of tool ports assigned to the

Application Performance Management team is used for a flow recorder and a long-term packet capture device. Regardless of the specific tool connected, the idea is the same – tool ports are where users select different portions of the data arriving on network ports.

NOTE: Tool ports only allow data output to a connected tool. Any data arriving at the tool port from an external source will be discarded. In addition, a tool port's **Link Status** must be **up** for packets to be sent out of the port. You can check a port's link status on the Ports page by selecting **Ports > Ports > All Ports** and looking at the Link Status field. Figure 22-2 shows an example where the link status is up for ports 1/1/x1, 1/1/x2, and 1/1/x3 bu down for port 1/1/x4.

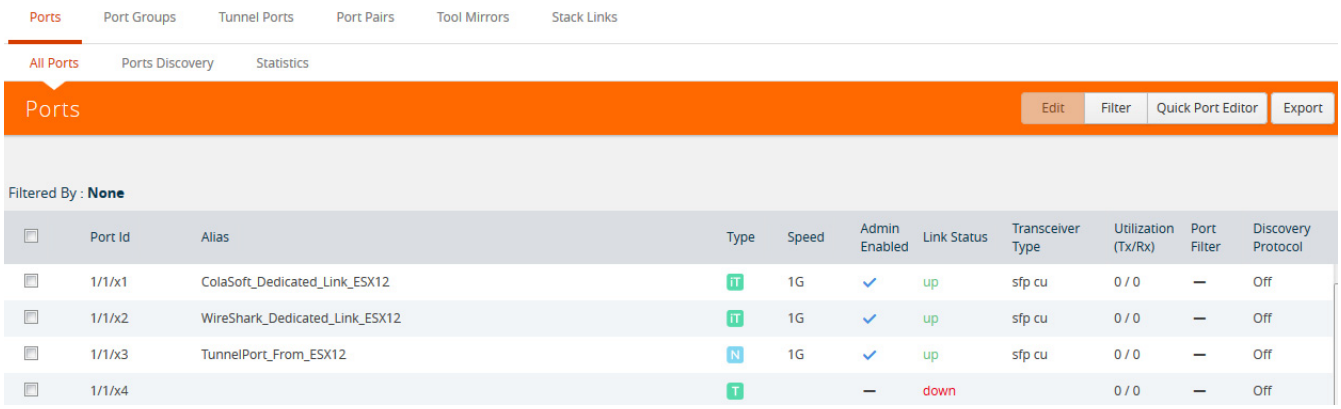

Figure 22-2: Port Link Status

Ports on GigaVUE TA Series Traffic Aggregator Nodes

Prior to software version 4.7, GigaVUE TA Series Traffic Aggregator nodes did not support tool ports. Instead, they supported gateway ports as displayed in Figure 22-3 and described in *Concepts Illustrated in* Figure 22-3 on page 394.

Starting in software version 4.7, all gateway ports on GigaVUE TA Series nodes are tool ports. For details, refer to *Tool Ports on GigaVUE TA Series Nodes* on page 395.

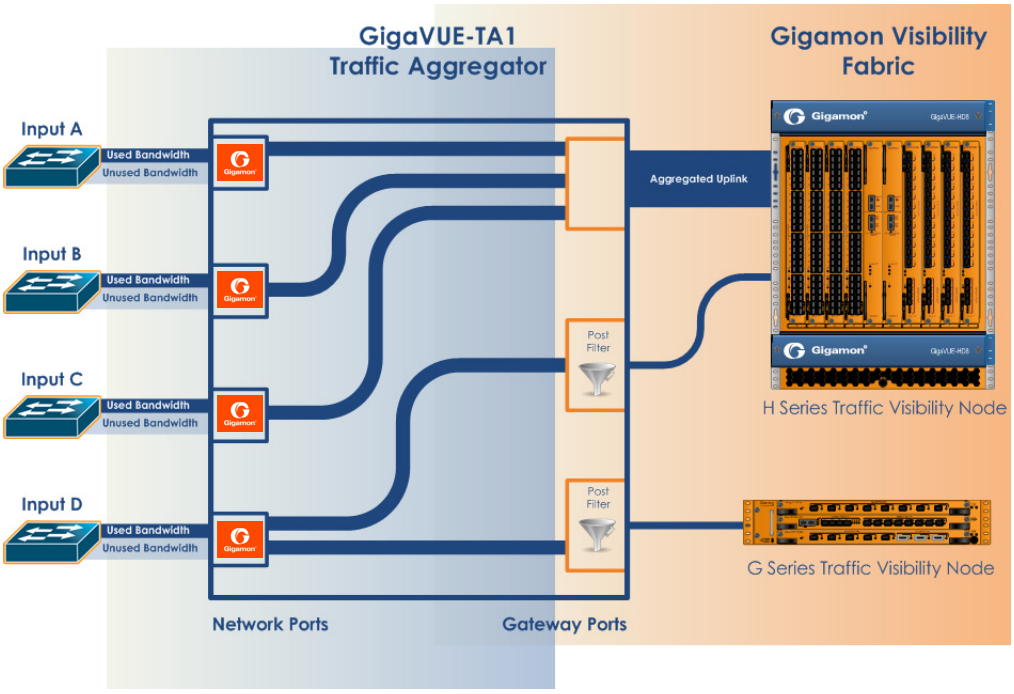

Figure 22-3: GigaVUE-TA1 Packet Distribution

Concepts Illustrated in Figure 22-3

Figure 22-3 illustrates the concept of data flows. Data arrives from different sources at the network ports on the left and is forwarded to different Gigamon nodes connected to the tool ports (formerly gateway ports) on the right.

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

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

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

- Traffic arriving at a single network port can be sent to multiple destination tool ports.

  Note that in Figure 22-3, the traffic arriving on **Input D** is sent to two different tool ports.

- Filters can be applied to tool ports:

  Filters applied to tool ports are called ***egress-filters***. Egress-filters are useful if you want to send the same traffic to multiple tool ports and have each one allow or deny different packets based on specified criteria. You can use up to **20 egress-filters** at a time on GigaVUE TA Series nodes.

NOTE: In Figure 22-3, egress-filters are set to focus on different parts of the data stream arriving at **Input D** – traffic on a VLAN range, a subnet range, and so on.

Tool Ports on GigaVUE TA Series Nodes

Refer to the following notes and considerations for GigaVUE TA Series nodes starting in software version 4.7 (including GigaVUE-TA1, GigaVUE-TA10, GigaVUE-TA40, GigaVUE-TA100, and Certified Traffic Aggregation White Box):

- Gateway ports on GigaVUE TA Series nodes are removed and converted to tool ports. In addition, gateway mirrors are removed and converted to tool mirrors.
- During an upgrade to 4.7 (from a prior software version), gateway ports and gateway mirrors are automatically converted to tool ports and tool mirrors.
- Tool ports on GigaVUE TA Series nodes can continue to be used to aggregate traffic (as displayed in Figure 22-3 and described in *Concepts Illustrated in* Figure 22-3 on page 394).
- Tool ports on GigaVUE TA Series nodes can also be used to directly connect to tools, such as firewalls, Intrusion Prevention Systems, or Application Performance Monitors.
- GigaVUE TA Series nodes support network, tool, stack, and hybrid port types.
- Hybrid ports are fully supported in both standalone and cluster mode on GigaVUE TA Series nodes. When a GigaVUE TA Series node is in a cluster, hybrid ports can be configured.
- GigaVUE TA Series nodes can continue to be clustered with GigaVUE H Series nodes. Tool ports are supported on GigaVUE TA Series nodes in a cluster.
- When GigaVUE TA Series nodes are in a cluster, bidirectional traffic flow is enabled on the stack links of GigaVUE TA Series nodes.
- Map rules using GigaVUE TA Series tool ports in the egress direction are supported.

Hybrid Ports

Hybrid ports are created by creating a dual function tool port. A physical tool port is set as a virtual network port which can then send traffic to other tool ports using secondary maps. A hybrid port is operated in loopback mode. This is only available if the H node is upgraded to minimum of 4.2 release. For more details on how to setup hybrid ports and the caveats, refer to the *GigaVUE-OS CLI User's Guide.*

Stack Ports

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

Inline Network Ports

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

NOTE: Inline network ports are supported only on GigaVUE-HC1, GigaVUE-HC2 and GigaVUE-HC3.

Inline Tool Ports

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

NOTE: Inline tool ports are only supported on GigaVUE-HC2 and GigaVUE-HC1.

Circuit Ports

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

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

Circuit ports are supported on the following:

- All GigaVUE H Series and TA Series nodes.
- As a source port in a regular map and as a destination port in a regular collector map.
- GigaStreams, port filter, and port groups.

GigaSMART Engine Ports

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

Port Lists

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

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

Port Aliases

GigaVUE-OS lets you configure textual aliases for network and tool ports. Aliases can be used in place of the numerical **bid/sid/pid** identifier required in many packet distribution.

To set up Port Aliases in H-VUE:

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

   Figure 22-4 shows an example of an port alias for network port 17/1/x1.

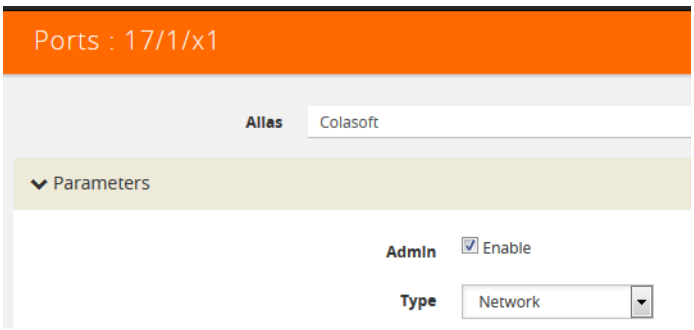

Figure 22-4: Creating a Port Alias

Work with Hybrid Ports

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

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

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

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

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

Using hybrid ports, you can create maps in a daisy chain.

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

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

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

- as a map source and destination (for regular maps, as well as map-passall, and map-scollector)
- in a GigaStream
- in a port group
- with an egress port filter. The hybrid port has the same limitation as a tool port (100 filters per line card or module). In the GigaVUE-HC2 equipped with Control Card Version 2 (HC2 CCv2) or the GigaVUE-HC3 node, the limit is 400 filters. The GigaVUE TA Series can only support 20 tool port-filters. When the GigaVUE-TA100 or GigaVUE-TA200 are in a cluster, they can support 400 filters.

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

When using hybrid ports, consider the following:

- Be aware not to configure traffic loops, such as such as H1 → H2, H2 → H3, H3 → H1.Do not use the same hybrid port as ingress as well as egress on all maps, such as such as H1 → H1.
- Once a hybrid port is used in a map or other traffic object, the port type cannot be changed.
- Hybrid ports cannot be used in any inline objects, such as inline-network or inline-tool.
- Hybrid ports do not support ingress port VLAN tagging.
- Hybrid ports cannot be used in a tool mirror because that is for tool ports only.
- Hybrid ports cannot be used in a port pair because that is for network ports only.
- Hybrid ports are not supported on 100Gb ports with CFP2 transceivers.

Hybrid ports are supported on GigaVUE H Series nodes. In a cluster environment, hybrid ports can be configured across nodes.

To configure a hybrid port, do the following

1. Select **Ports** > **Ports** > **All Ports**.
2. Click **Quick Port Editor**.
3. Enter the port ID, the Quick Search field. For example, 1/1/x10.

4. Click in the **Type** field and select **Hybrid** as shown in the following figure.

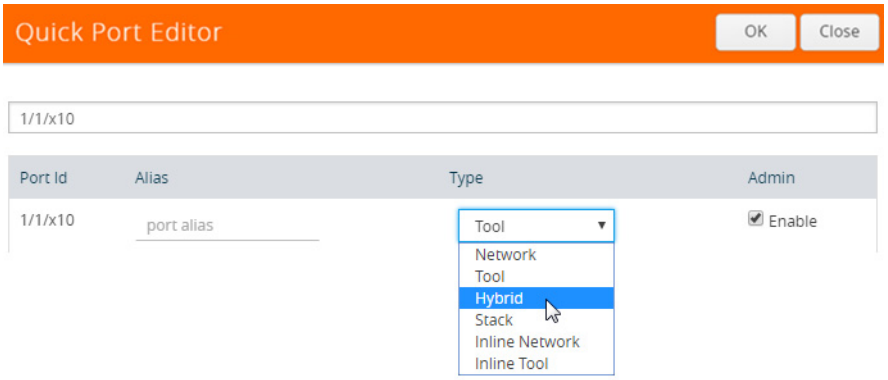

5. Click **Save**.

Hybrid ports can be used in the following:

- Regular map
- Regular map with GigaSMART operation
- First level and second level maps with vports

When configuring a map, use a hybrid port as follows:

- In the **Source** field when it is used as an indirect traffic source port
- In the **Destination** argument when it is used as a tool port

NOTE: You cannot use the same hybrid port in one map as both **Source** and **Destination**, or create a loop from multiple maps.

There is no limitation to the number of maps that can be used as second level maps to which packets can be forwarded.

Port Filters

Flow Mapping provides the ability to apply filters to tool ports, passing or dropping traffic after it has been forwarded from a network port.

Tool port-filters provide a convenient way to narrow down the traffic seen by tools without having to change an entire map. However, they are less efficient and scalable than flow maps – focus on using flow maps as your first packet distribution technique.

How to Apply Port Filters

To apply a port filter, do the following:

1. Select **Ports > Ports > All Ports.**
2. Select the egress port that to which you want to apply a filter.
3. Click **Edit**.
4. Under Filters on the Ports page, click **Add Rule**.
5. Select and configure the rule.

Add a new port-filter using the specified criteria as follows:

- Use a **drop** rule to deny packets matching the specified criteria.
- Use a **pass** rule to allow packets matching the specified criteria. All other packets are denied.

6. Click Save.

Port Filter Notes

Keep in mind the following notes when managing port-filters:

- The **filter** is only supported for egress ports – network ports use maps to direct traffic.
- You can only configure egress port filters on a single port at a time. The **filter** argument is blocked when used the with multiple tool ports or port groups.

Port-Filter Maximums

Each GigaVUE-HC2 or GigaVUE-HC3 module, or GigaVUE-HB1 node supports 100 combined tool port-filters. In the GigaVUE-HC2 equipped with Control Card Version 2 (HC2 CCv2) or the GigaVUE-HC1 or GigaVUE-HC3 node, the limit is 400 filters. The GigaVUE TA Series can only support 20 tool port-filters. When the GigaVUE-TA100 or GigaVUE-TA200 are in a cluster, they can support 400 filters.

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

Status of Line Cards/Nodes and Ports

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

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

How to Check Port Status with Ports Page

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

1. Select **Ports > Ports > All Ports** to open the Ports page.
2. Locate the port to check by entering the port ID or port alias in the search field.
3. If you need to change the port type or enable the port:
   a. Click **Quick Port Editor**.
   b. In the Quick Port Editor, enter the port ID or port alias in the Quick search field to find the port.
   c. Set the port type by selecting type from the drop-down list in the **Type** field. enable the by selecting **Enable** as needed.

How to Check Port Status with Chassis Page

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

1. Select **Chassis** from the Navigation pane to open the Chassis view shown in Figure 22-5.
2. Use the view buttons on the Chassis page to check the status of the cards as well as the ports. When viewing a node in cluster, there is a drop down option to select a specific node in a cluster configuration.

For details about the Chassis page, refer to “Chassis” in the *GigaVUE-OS H-VUE Administration Guide*.

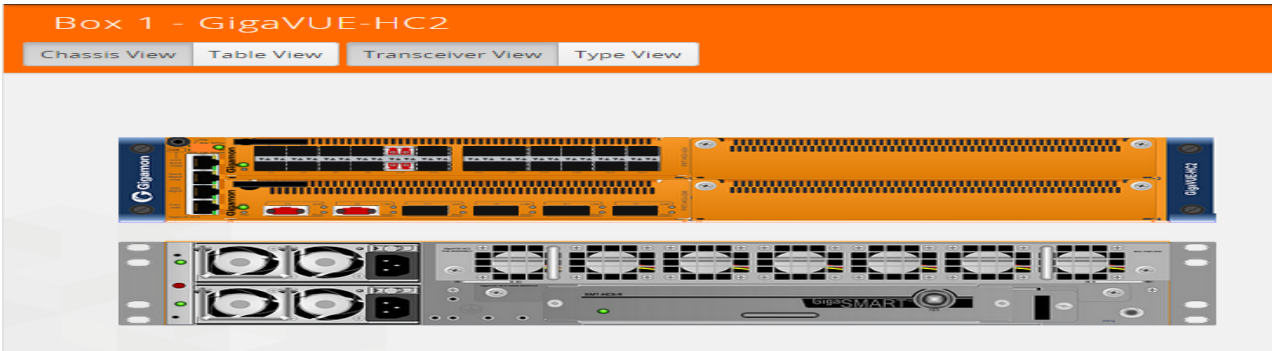

Figure 22-5: Chassis Page

Managing Ports

The Ports pages allows you to manage and configure ports for various functions. In case of an GigaVUE-HC2, the pages for managing inline bypass ports are available by selecting Inline Bypass from the main navigation pane.

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

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

NOTE: Starting in software version 5.5.01, any change in the port health status is indicated by a **Status Update** notification pop-up that appears in the bottom-left corner of the following port pages:

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

The link in the notification opens the port quick view.

Ports

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

All Ports

The All Ports page displays when you select All Ports. The Ports page shows a table with detailed information about each port ID on a specific device. Only the GigaVUE H Series and GigaVUE TA Series devices are presented in the Port Page view as shown in Table 22-1. You can control which ports display on the page by selecting a set of filters or configure the ports through the Quick Port Editor or selecting Edit for a selected port. For details about filtering ports, refer to *Port List Filter* on page 404. For details about the Quick Port Editor, refer to *Quick Port Editor* on page 406.

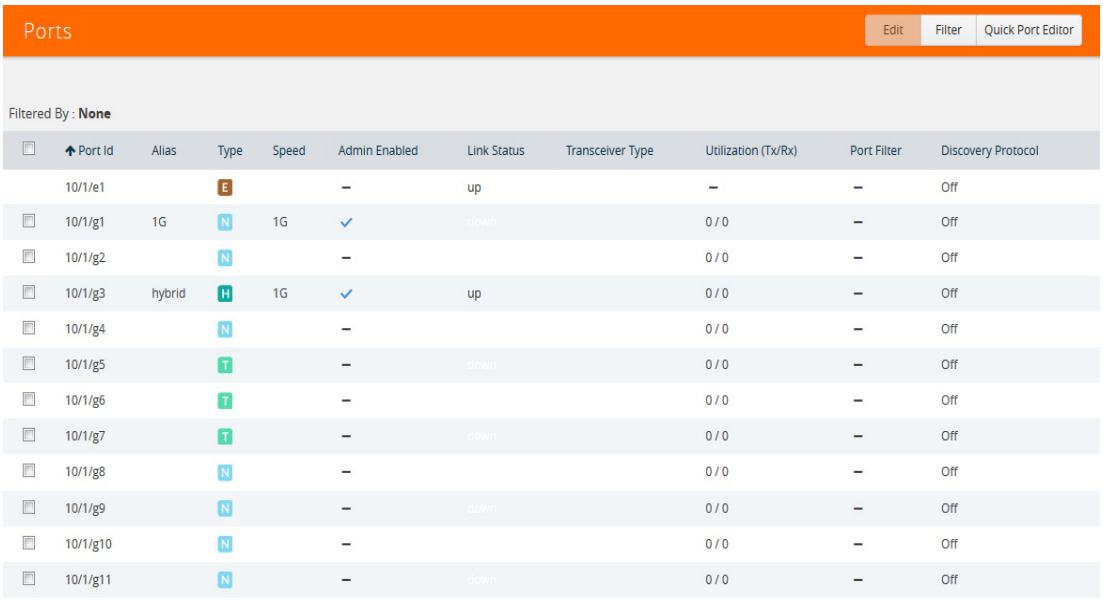

Figure 22-6: Ports Page

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

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

NOTE: Not all port types are supported on all platforms. Inline network and inline tool ports are only supported on GigaVUE HC Series nodes.

Table 22-1 provides descriptions of the columns on the ports page.

Table 22-1: Descriptions of Ports Page Columns

Column	Description
Port ID	The port number is in <box ID>/<slot ID>/<port_ D> format in the CLI and H-VUE. For the GigaVUE TA Series and the GigaVUE-HB1 slot ID is always 1 as they are not modular. For the GigaVUE-HC1, GigaVUE-HC2, and GigaVUE-HC3, the line cards or modules are identified by the slot number. In a standalone (default) configuration, box ID is always designated as 1 but can be changed through the CLI (it cannot be changed through H-VUE). In a cluster configuration, the box ID can vary.
Alias	Alias name of the port, if any.
Speed	Current setting for the port's speed.
Type	List the type of port such as network, tool, stack, inline network, inline tool, circuit, or hybrid. You can set the port type through the Quick Port Editor or selecting the port on the Port page and clicking **Edit**. The port type is set by selecting the type in the **Type** field.
Enabled	Indicates whether the port is administratively enabled or disabled.
Force Link Up	Indicates the 'force link up' setting for the port. When enabled, this option forces connection on the optical port.
Ude	Indicates whether the port is enabled for unidirectional (Ude) or bidirectional traffic. Enabled means Ude; Disabled means bidirectional.

Table 22-1: Descriptions of Ports Page Columns

Column	Description
Link	The current status of the link connected to the port, either port link up or port link down.
Transceiver Type	The type of transceiver installed in this port.
Port Filter	Indicates if a tool port filter is associated with this port.
Discovery Protocol	Protocol used to discover neighboring nodes using CDP or LLDP. This feature is only enabled for network ports.

Port Quick View

The Quick View for ports displays when you click on a row in the Ports page to quickly get more information about a specific port. The quick view shows the port properties, statistics information on receiving (Rx) and/or Transmitting (Tx) ports and alarms information. The quick view also shows a graphical representation of port statistics. Refer to Figure 22-7 for an example.

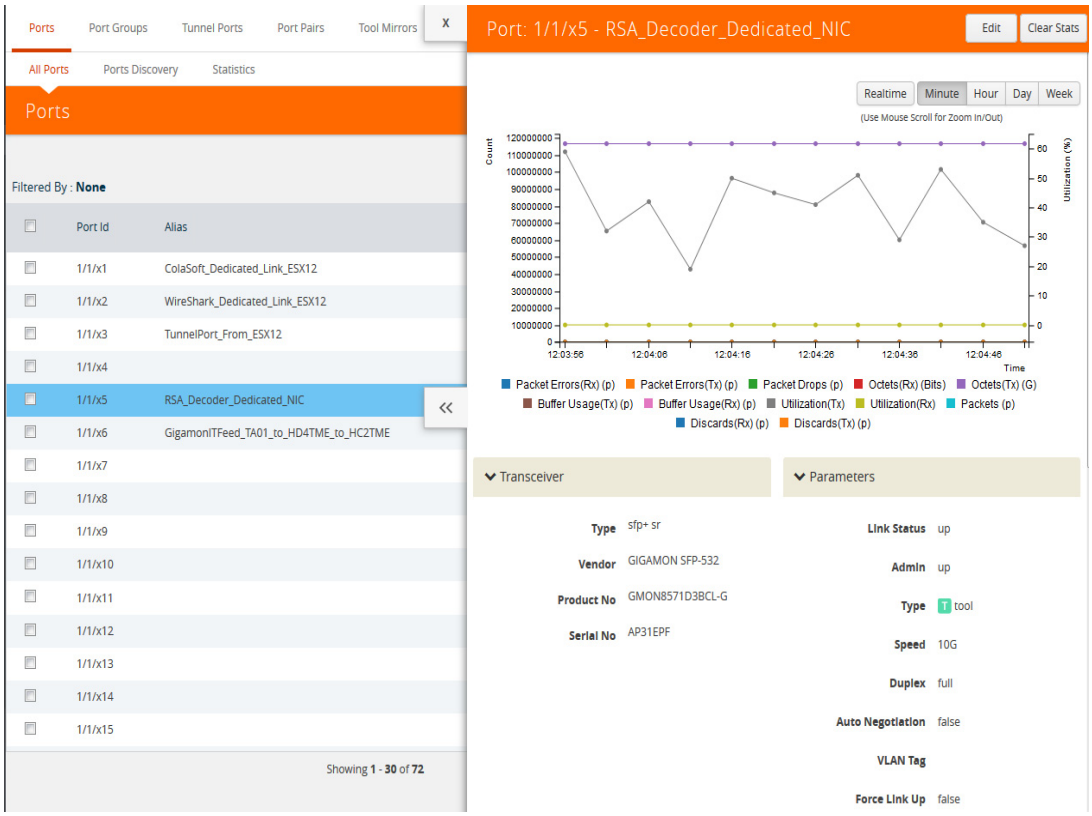

Figure 22-7: Ports Quick View

Port List Filter

The ports that display on the Ports page can be filtered so that only ports that meet certain criteria display on the page, such as port type and admin status. To filter the ports, select **Filter**. This opens the Filter view shown in Figure 22-8 where you can specify how to filter ports displayed on the Ports page.

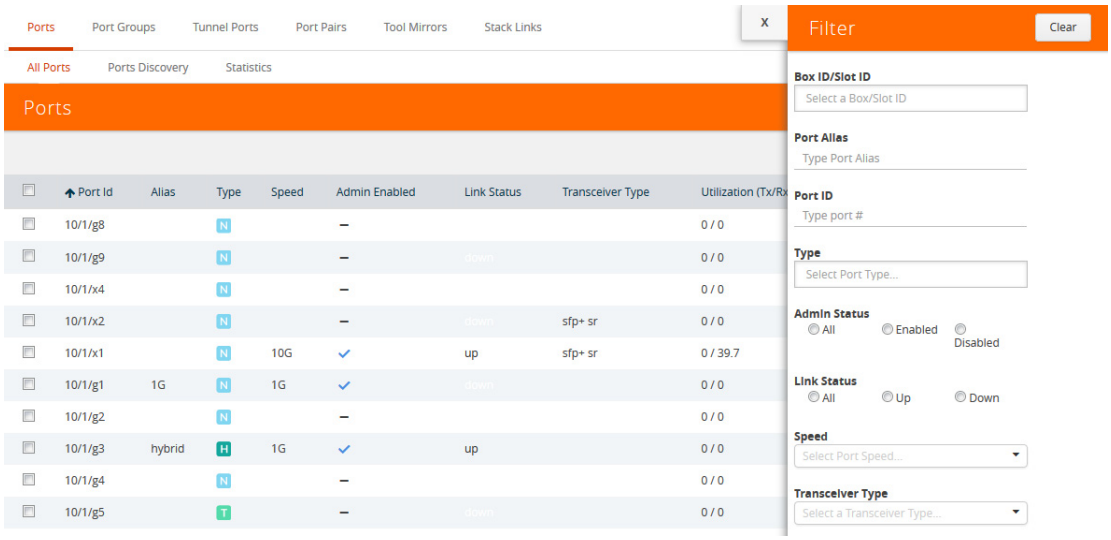

Figure 22-8: Port List Filter

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

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

Criteria	Description
Speed	Display ports with the selected port speed. The port speeds available depend on the node.
Transceiver Type	Display ports with the selected transceiver type. The transceivers available selection depend on the type of transceivers connected to the ports.

To filter the ports, enter the information to use for filtering the ports and select the radio buttons. For example, in Figure 22-9, the filters selected are Network Type and Admin Status Enabled. Click the **Clear** button to remove the filter selections.

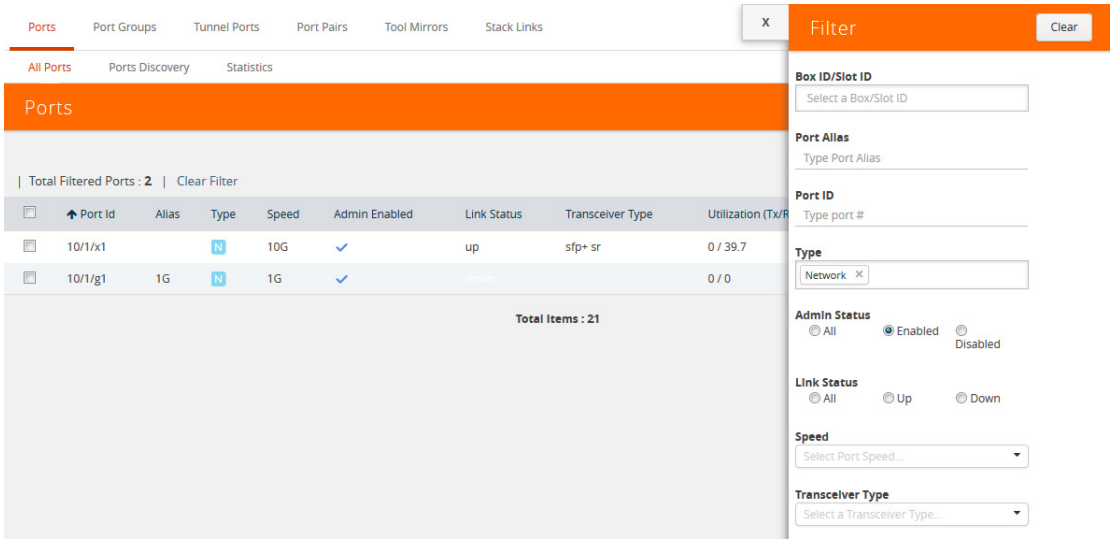

Figure 22-9: Filtering by Network Port Type and Admin Status Enabled

After the filter is applied, the Ports page displays only the ports that correspond to the selected filters and shows the total number of ports that meet the criteria. To clear the filters, select **Clear Filter**. Figure 22-10 shows the Port pages with two ports that correspond to the current filters: Network Type and Admin Status Enabled.

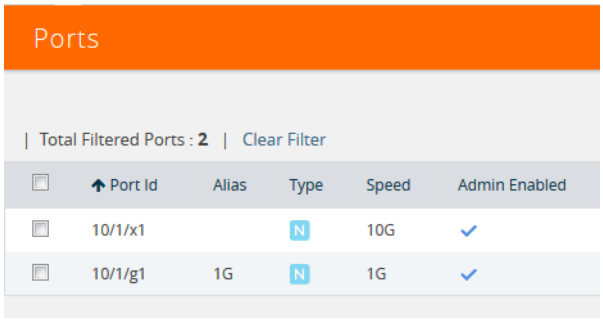

Figure 22-10: Filtered Ports List

Quick Port Editor

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

1. Click **Quick Port Editor**.
2. For each port on which you want to set the port type, select the type from the drop-down list. In Figure 22-11, port 17/1/x2 is being change from a network port to at tool port.

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

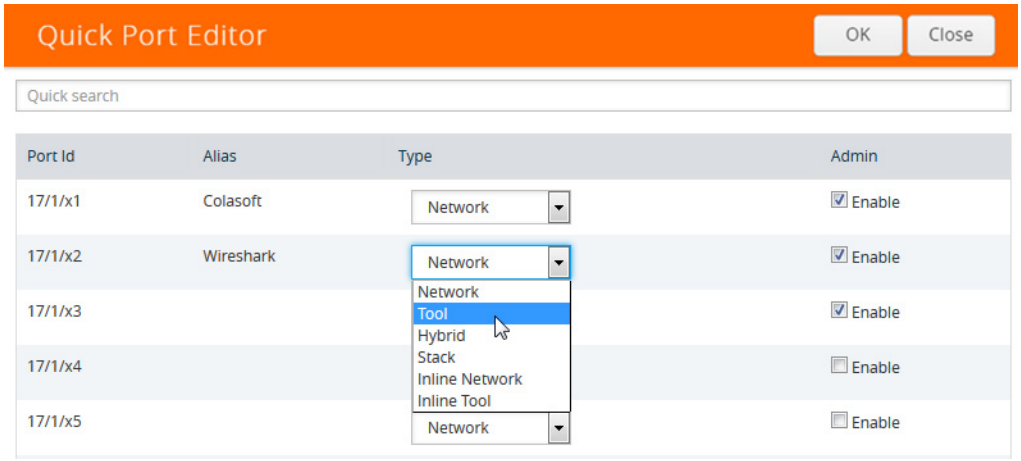

Figure 22-11: Port Type Selection

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

Figure 22-12 shown an example where port 17/1/x1 is configured as a network port, port 17/2/x2 is configured as a tool port, and port 17/1/x3 is configured a hybrid port. Each port can also be assigned an alias. Any port types set in the CLI or through the GigaVUE-FM APIs are reflected on this page. Any changes made in H-VUE are reflect in the CLI and responses to an API. For more information, refer to *Port Aliases* on page 397 for port aliases and to the *GigaVUE-FM REST API Getting Started Guide* and the *GigaVUE-FM Reference* for APIs.

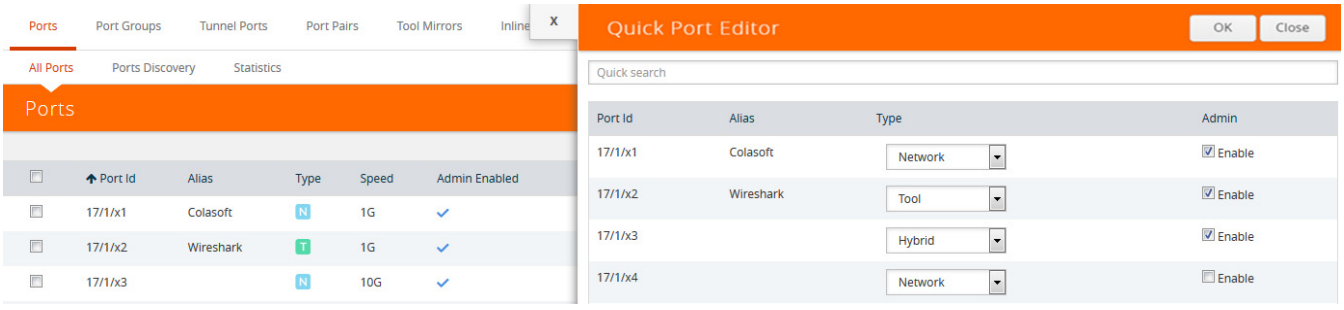

Figure 22-12: Port Type Editor

Configure Ports

From the Ports page, you can either configure or edit a specific port by selecting a port and clicking the **Edit** button on the Ports page or on the Quick Port Editor. Table 22-2 describes the options on the configuration page.

Table 22-2: Port Configuration Options

Field	Description
Alias	The alias configured for this port, if any. Aliases can be used in place of the numerical bid/sid/pid identifier required in many packet distribution commands in the CLI. For example, instead of configuring a connection between, say, 1/1/x1 and 1/2/x4, you could connect Gb_In to Stream-to-Disk. Note that aliases can only be applied to single ports. They cannot be applied to groups of ports. Port alias can be up to 128 characters long including special characters. Aliases are case sensitive.
Admin	Check to enable the port.
Type	Specifies whether the port is configured as an Inline Network port, Inline Tool port, Network port, Tool port, Stack port, Circuit port, or Hybrid port.
Speed	Specifies the speed for the selected port. For copper ports, you can click to change the speed as long as Auto Negotiation is disabled.
Duplex	Specifies the port's duplex configuration. Only full duplex is supported. Starting in software version 5.2, half duplex support is removed from all GigaVUE nodes. If half duplex was configured in a previous software version, it will remain intact following the upgrade to 5.2 or higher release. Update to full duplex, if required.
Auto Negotiation	Select to enable autonegotiation for the selected port. When autonegotiation is enabled, duplex and speed settings are ignored. They are set through autonegotiation.
Force Link Up	When enabled, this option forces connection on an optical port. Use this option when an optical GigaPORT tool port is connected to a legacy optical tool that does not transmit light. This option is not available for 10Gb capable ports with a 1Gb SFP installed.
Ude	When selected, this option indicates the port is unidirectional (UDE). When deselected (disabled), the port is bidirectional. UDE is enabled by default. **Note**: This option is available for GigaVUE-HC2 (CCv2), GigaVUE-HC3, GigaVUE-TA100, GigaVUE-TA100-CXP, and GigaVUE-200 platforms with 100Gb BiDi (QSB-512) transceiver. If used with passive taps, ports used for monitoring should be set to Network port with UDE enabled. **IMPORTANT:** If you clear the UDE check box, the laser will start to transmit, which may affect the remote connectivity.

Table 22-2: Port Configuration Options

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

Table 22-2: Port Configuration Options

Field	Description
Utilization Threshold	Sets the utilization percentage for this port at which the GigaVUE H Series node will generate high or low utilization alarms for the port. For more information about port utilization, refer to *Monitor Port Utilization* on page 468. By default, the threshold is set to **0**, which disables the threshold.
Lock Port	Restricts use of the port for only your user account as follows: • Users with the admin role can lock any port in the system. Users with the Default/Operator role assigned can only lock ports to which their account has been granted access. • Administrators can lock a port for another user by including the optional user. • You can optionally share a locked port by specifying users in the **Lock shared with Users** field or selecting users to share the lock with through their assigned roles. For more information about who to set lock sharing, refer to *Managing Ports* on page 401.

Ports Discovery

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

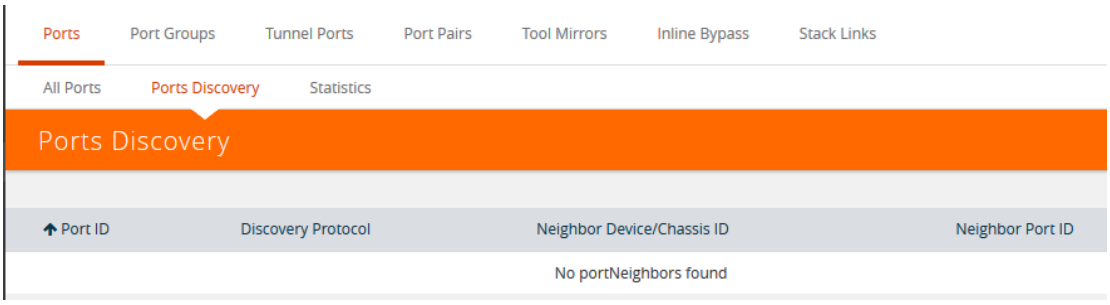

Figure 22-13: Ports Discovery Page Displaying Neighbors Found Using Discovery

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

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

Neighbor information is removed or replaced as follows:

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

- If there are currently less than two discovered neighbors for a port, the newly discovered neighbor is added (actually exceeding the 2K limit to guarantee a minimum of two neighbors per port).

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

For information about the discovery protocols and enabling port discovery, refer to *Port Discovery* on page 420

Statistics

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

Column	Definition	Notes
Octets (Rx/Tx)	The count of packets/bytes received and transmitted by this port.	Error packets are not transmitted, therefore they are not counted. Excludes undersize frames.
Octets/sec (Rx/Tx)	The count of packets/bytes received and transmitted by this port per sec.	Error packets are not transmitted, therefore they are not counted.
Unicast Packets (Rx,/Tx)	The count of packets/bytes received and transmitted by this port.	Excludes multicast packets, broadcast packets, packets with FCS/CRC errors, MTU exceeded errors, oversize packets, and pause packets.
Non-Unicast Packets (Rx/Tx)	Total Non-unicast packets received or transmitted.	
Packets/sec (Rx/Tx)	The rate which packets are received or transmitted.	
Packet Drops (Rx)	Total Dropped Packets	Packets are dropped when a network port's bandwidth is exceeded due to oversubscription. Packets are dropped when they reach the port but before they are sent out.
Discards (Rx/Tx)	Total received and transmitted packets discarded. This counter increments when a packet is discarded at the tool port due to a tool port filter.	Discards are counted in the following cases: • Traffic arriving at a network port that is not logically connected using a map or map passall. • Map rules applied on a network port. • In packets on a tool port. • Pause frames.
Errors (RX/TX)	Total Error Packets Received or Transmitted. Error packets include undersize, FCS/CRC, MTU exceeded, fragments, and oversize packets.	Excludes oversize packets without FCS/CRC. Packets larger than the MTU setting arriving on a network port are counted twice in the counter. So 1000 oversize packets would show up as 2000. This double-counting only happens with Oversize error packets.
Utilization (Rx/Tx)	Percentage of port utilization by packets received or transmitted	

Port Groups

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

All Ports Groups

Selecting **All Port Groups** under **Port Groups** opens the **Ports Group** page by default. This page is used to create a port group, which can simplify administration of GigaVUE ports. Administrators can create groups of ports that can then be quickly assigned to different user groups. With clustered ports potentially numbering into the hundreds, port groups provide a useful shorthand when assigning multiple ports to different user groups. (To create user groups, select **Roles and User** from the navigation pane, and refer to *Managing Roles and Users*" in the *GigaVUE-OS H-VUE Administration Guide* for more details.) The following are the different types of port groups:

- Network Port Group—contains only network ports.
- Tool Port Group—contains only tool ports or tool GigaStream, which is a combination of multiple tool ports.
- Hybrid Port Group—contains only hybrid ports or hybrid GigaStream, which is a combination of multiple hybrid ports.
- Circuit Port Group—contains only circuit ports or circuit GigaStream, which is a combination of multiple circuit ports.
- Load Balancing Port Group—contains tool ports for load balancing. The maximum number ports allowed for port balancing is 16.

However, port groups that include GigaStream can only be used with GTP Overlap Flow Sampling maps. For more information about GTP Flow Sampling, GTP Whitelisting and GTP Overlap Flow Sampling maps, refer to *GigaSMART GTP Whitelisting and GTP Flow Sampling* on page 913.

GigaStreams

Selecting **GigaStreams** under **Port Groups**, opens the GigaStreams page. A GigaStream is a bundle of multiple ports on a GigaVUE-OS node used for stacking, tool ports. The following are the types of GigaStream you can create:

- Tool GigaStream—It is a bundle of multiple tool ports used as a single logical group. This type of GigaStream as a single addressable destination, allowing you to overcome tool port oversubscription issues.
- Hybrid GigaStream—It is a bundle of multiple hybrid ports that are combined into a single logical group in all H Series nodes.
- Stack GigaStream—It is a bundle of multiple stack ports used as a single logical group. Stack-links can use GigaStream to distribute data between multiple H Series nodes operating in a cluster. With the number of 10Gb/40Gb/100Gb ports possible in a GigaVUE H Series chassis, using only one 10Gb port for a stack-link could cause a serious bottleneck. A GigaStream dramatically increases the bandwidth

available for stack-link connections, letting you connect H Series nodes in a cluster and still take advantage of the 10Gb port density.

- Circuit GigaStream—It is a bundle of multiple circuit ports that are combined into a single logical group. The circuit ports send or receive traffic that is tagged with the circuit ID.
- Controlled GigaStream

  Controlled GigaStream provides more control of the traffic stream by specifying the size of a hash table and allowing the assignment of hash IDs to the ports in a GigaStream. This makes it possible to keep the hashing algorithm from reapplying the algorithm to the ports if one of the ports in the GigaStream goes down.

For more detailed information about GigaStreams, refer to *How to Use GigaStream* on page 430

Port Pairs

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

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

To configure a port pair, do the following:

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

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

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

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

4. Click **Save**.

Tool Mirrors

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

Tool-mirror can be created from:

- Tool port to tool port or ports on the same node.
- Tool port to GigaStream or GigaStreams on the same node.

  The destination for a tool-mirror must always be either a tool port or a GigaStream.
- Tool Mirrors can cross line cards/modules – they can start on one line card and end on another in the same node. However, they cannot cross nodes in a cluster. On GigaVUE-HC2, tool-mirrors can exist between ports on separate modules on the same node.
- Tool Mirrors on GigaVUE-HB1 can be created on tool ports or GigaStream ports.
- Tool Mirrors are not allowed from Tool GigaStream to tool port.
- Tool Mirrors are not supported on tool ports with copper SFPs installed.

Create Tool Mirror

To create a Tool Mirror, do the following:

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

   The Tool Mirrors page displays a list of the currently configured Tool Mirrors. The following figure shows an example.

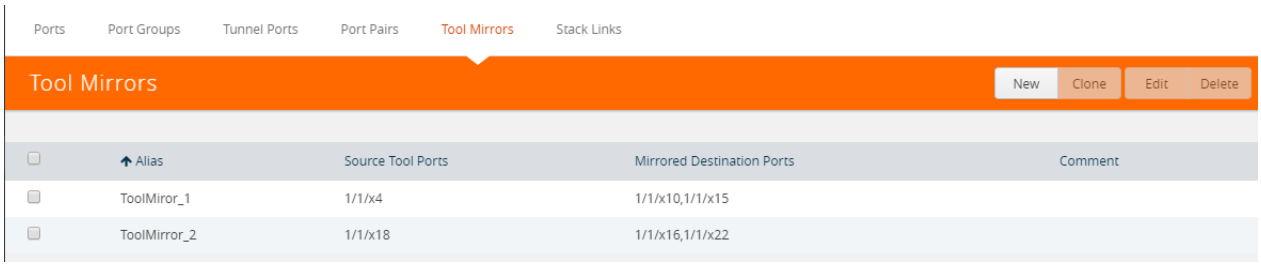

2. Click **New**.

   The Tool Mirror configuration page displays as shown i the following figure

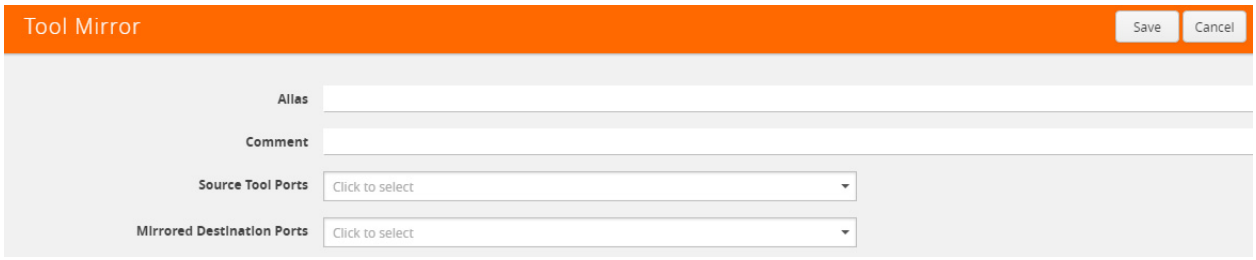

3. Configure the Tool Mirror:
   a. Enter an alias in the **Alias** field.
   b. (Optional) Enter a comment in the **Comment** field.
   c. Click in the **Source Tool Ports** field and select the source tool ports for this tool mirror.
   d. Click in the **Mirror Destination Ports** field and select the destination tool ports for this tool mirror.
4. Click **Save**.

Edit Tool Mirror Comments

Comments in a Tool Mirror configuration are optional. However, you can add comments at any time, or edit existing comments. To add or edit comments, do the following:

1. Select **Ports > Tool Mirrors**. The Tool Mirrors page displays a list of the currently configured Tool Mirrors. The following figure shows an example.

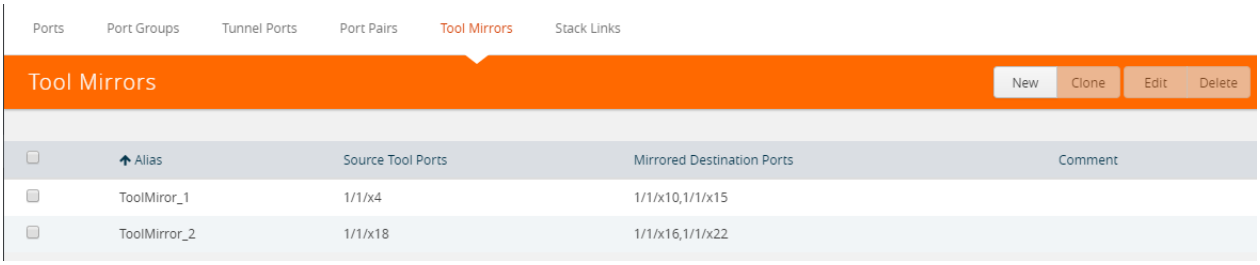

2. Select a Tool Mirror in the list of Tool Ports, and then click **New**. In the following figure, ToolMirror_1 is selected for edit.

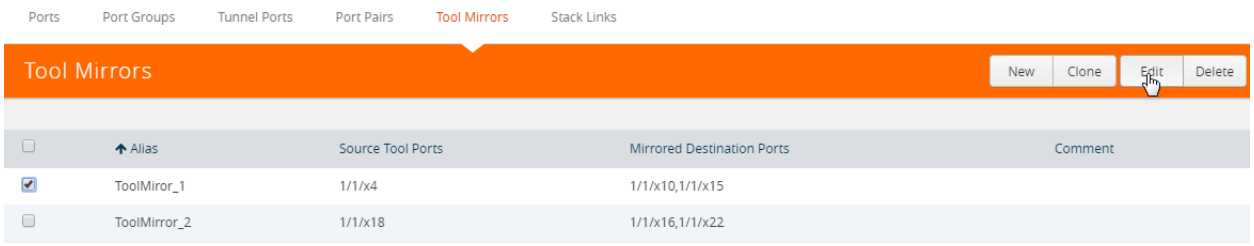

3. Enter or change a comment in the **Comment** field. (You cannot make any other changes to the Tool Mirror.)
4. Click **Save**.

Clone Tool Mirror

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

1. Select **Ports > Tool Mirrors**.
2. Select the Tool Mirror that you want to copy, and the click **Clone**.
3. Enter a new alias in the **Alias** field.
4. (Optional) Add or update comments in the **Comments** field.
5. Make change to the **Source Tool Ports** and **Destination Tool Ports** as needed.
6. Click **Save**.

Stack Links

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

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

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

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

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

To create a stack link, do the following:

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

IP Interfaces

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

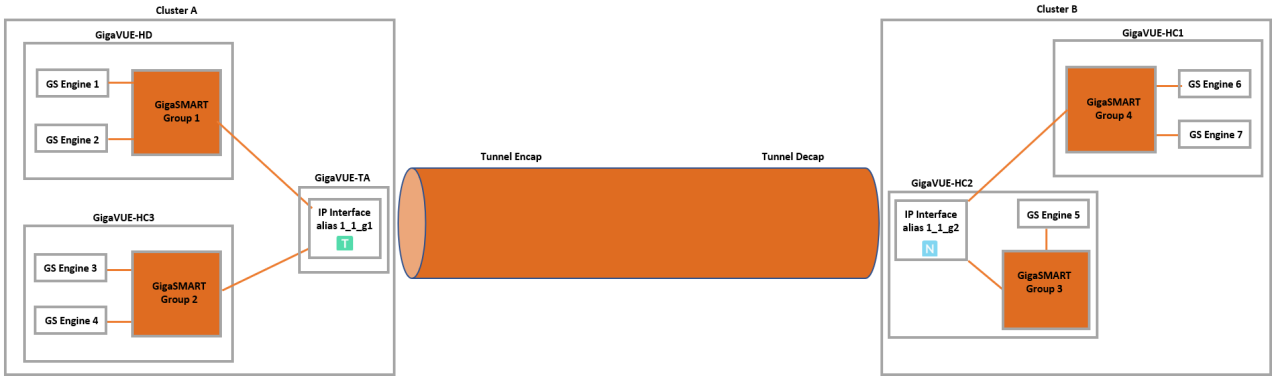

Figure 22-14: Tunnel Centralization

In this example, the GigaVUE-HD, GigaVUE-HC3, and GigaVUE-TA nodes reside in cluster A. In GigaVUE-HD, GigaSMART engines 1 and 2 are associated with the GigaSMART group 1. Similarly, in GigaVUE-HC3, GigaSMART engines 3 and 4 are associated with GigaSMART group 2. An IP interface with alias 1_1_g1 is configured with a tool port in GigaVUE-TA. Both the GigaSMART groups, 1 and 2 are associated with the IP interface. The IP interface 1_1_g1 is the originating point for the tunnel, where encapsulation happens.

Now, let us look at the termination point of the tunnel. The GigaVUE-HC1 and GigaVUE-HC2 nodes reside in cluster B. In GigaVUE-HC2, 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-HC2. 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.

Upgrade from Release 5.4.xx

When you upgrade from release 5.4.xx to either 5.5.xx or 5.6.xx, the tunnel ports that were configured prior to the upgrade will be converted to IP interface. The IP interfaces that are converted during the upgrade will have a standard naming convention, “giga_auto_tunnel_<*bid*>_<*sid*>_<*pid*>”, where *bid* is the box ID, *sid* is the slot ID, and *pid* is the port ID.

For example, the tunnel port, 1/1/g1 will be converted to IP interface with the alias, “giga_auto_tunnel_1_1_g1”.

Moreover, the ARP/NDP timer settings will be moved from GigaSMART to control card. The ARP/NDP timer value is 3-30 seconds. If you had configured the ARP/NDP timer settings for more than 30 seconds, it will be decreased to 30 seconds after the upgrade.

Configure IP Interface

Before you configure an IP interface, you must:

- Configure a network and a tool port.
- Create a GigaSMART group and NetFlow exporter.

  NOTE: You can associate multiple GigaSMART engines to a GigaSMART group.

To create an IP interface associated with a network port, do the following:

1. On the top navigation pane, click **Physical**.
2. In the Physical Nodes page, select the node for which you want to configure the IP interface.
3. From the left navigation pane, go to **Ports > IP Interfaces**.
4. In the IP Interfaces page, click **New**. The IP Interface page opens as shown in Figure 22-15.

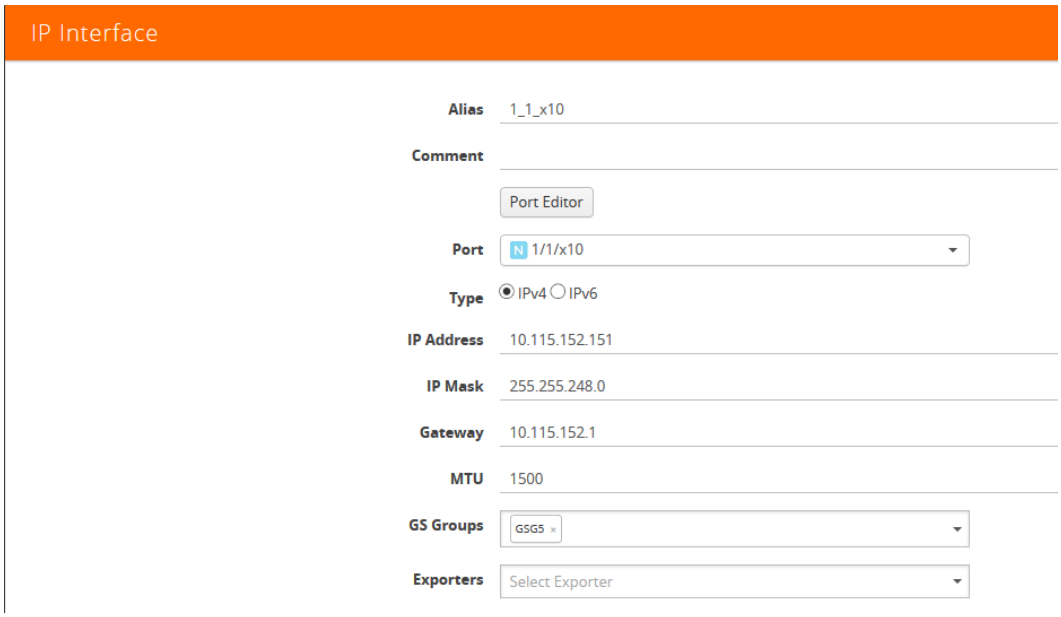

Figure 22-15: Configuring IP Interfaces

The following table provides a description of the fields on the IP Interfaces page.

Field	Description
Type	Tool port address type. Options: IPv4 and IPv6
IP Address	Specify the IP address for the IP interface.
IP Mask	Specify the IP Mask for the IP interface using the format: **255.255.255.255**
Gateway	Specify the IP address of the Gateway for the IP interface.
MTU	Specify the MTU for the IP interface (100 - 9600 bytes). The MTU for ports is fixed at 9600 for all network/tool ports on the following platforms: • GigaVUE-HB1 • GigaVUE-TA1, GigaVUE-TA10, and GigaVUE-TA40 • Certified Traffic Aggregation White Box The MTU is fixed at 9400 for all network/tool ports on the following platforms: • GigaVUE-HC2 and GigaVUE-HC2 equipped with Control Card version 2 (HC2 CCv2) • GigaVUE-HC1 • GigaVUE-HC3 • GigaVUE-TA100, GigaVUE-TA100-CXP, and GigaVUE-TA200 **RECOMMENDATION**: Set the MTU to 9400 on all platforms.
GS Group	Use the drop-down menu to assign decapsulation for this IP interface to one of the configured GS Groups.

5. In the **Alias** and **Comment** fields, enter a name and description for the IP interface.

6. From the **Ports** drop-down list, select the tool or the network port that you had configured.
   a. In the **Port** field, start by using the Box ID and Slot ID fields to select the line card with the port you want to use for the tunnel.
   b. From the list of available ports, selected the network port. You can select a maximum of one port.
7. Enter the **IP Address**, **IP Mask**, **Gateway**, and **MTU** for the IP interface.

   You can specify the subnet mask using either of the following formats:
   - netmask – For example, 255.255.255.248
   - mask length – For example, /29
8. From the **GS Groups** drop-down list, select the GigaSMART groups that you have created.

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

9. From the **Exporters** drop-down list, select the NetFlow exports that you have created.
10. Click **OK** to create the IP interface associated with a network or tool port and add it to the list of currently configured IP interfaces.

Circuit Tunnels

Circuit tunnels are used to route traffic between two clusters. The traffic is tapped and sent through network ports on the TAP landing nodes in a cluster. Based on the flow map configuration, traffic is filtered at the TAP landing nodes and sent to the circuit ports. The circuit ports encapsulate the traffic with a Circuit ID and routes the encapsulated traffic through a circuit tunnel. At the receiving end, the traffic is decapsulated and sent to the tool ports. The circuit tunnels are bidirectional. For more information about circuit tunnels, refer to *About Circuit-ID Tunnels* on page 473.

Port Discovery

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

Port Discovery with LLDP and CDP

The GigaVUE H Series is capable of snooping Link Layer Discovery Protocol (LLDP) packets and Cisco Discovery Protocol (CDP) packets. If the devices in your network use either of these protocols, a GigaVUE H Series node can identify its immediate neighbors and their capabilities. Snooped LLDP and CDP information includes the

remote port and chassis IDs, as well as other selected information, if it is included by the sender. This information can be used to determine the origin of traffic flows.

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

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

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

- Chassis ID
- Port ID
- Time-to-Live (TTL)
- End of TLVs
- Port description
- System name
- System description
- System capabilities available
- System capabilities enabled
- VLAN name
- Management address
- Port VLAN ID
- Management VLAN ID
- Link Aggregation port ID
- Link Aggregation status
- Maximum Transmission Unit (MTU)

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

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

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

Notes:

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

Enable Port Discovery

Port discovery is disabled by default. It can only be enabled on network type ports and only on ingress.

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

1. Select **Ports > All Ports**.
2. On the Ports page, click on the Port ID for port on which you want to enable port discovery. Ensure that this port is set as network port.

   The Quick View window displays for the port ID.
3. Select **Edit** from the top right corner of the Quick View Window.
4. To enable ports discovery do the following under **Ports Discovery**:
   a. Select **Enable**
   b. For Discovery Protocols, select one of the following: **All**, **LLDP**, or **CDP**.

   Figure 22-16shows ports discovery enabled using the LLDP protocol for network port 17/1/x1.
5. Click **Save**.

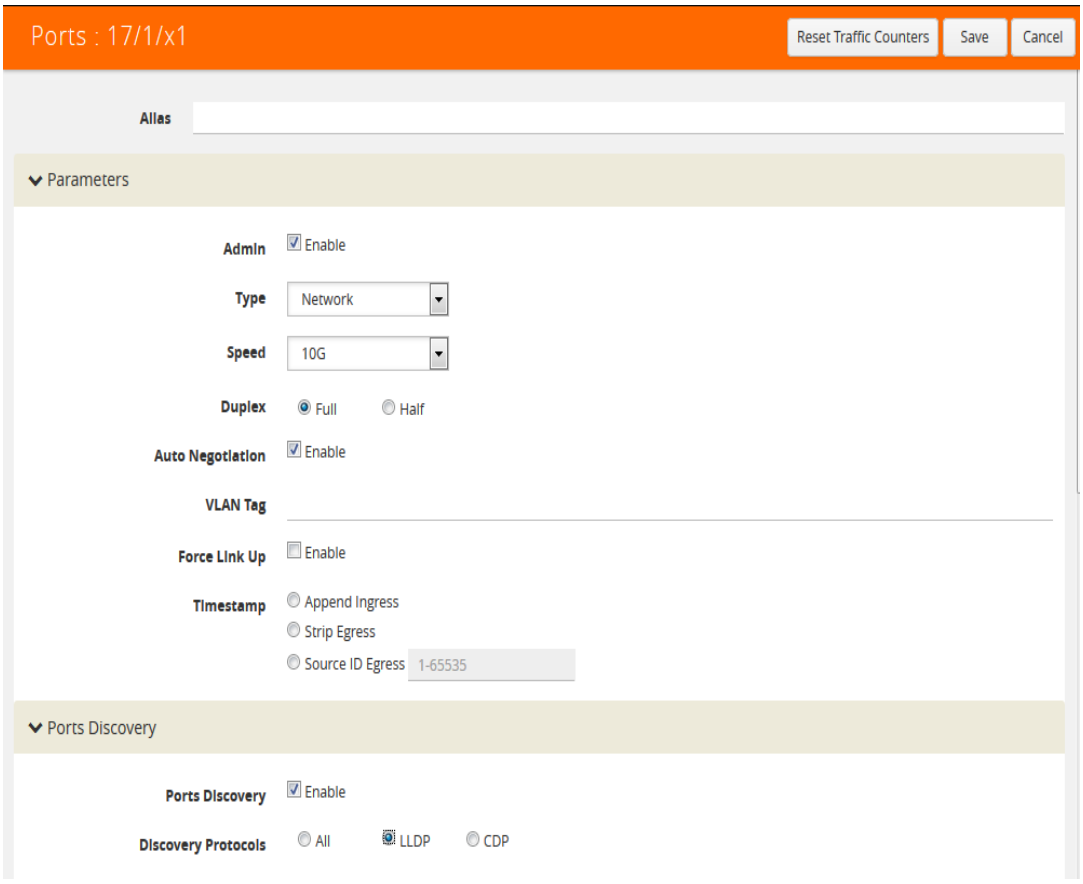

Figure 22-16: H-VUE Page to Enable and Configure Discovery Protocol

Limits of Discovery Information

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

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

Neighbor information will be removed or replaced as follows:

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

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

Port Discovery Support

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

Port Discovery for a Cluster

LLDP and CDP discovery can be enabled on any ingress port in the cluster. The discovery information will be aggregated and available on the cluster master.

Port Discovery Supported for SNMP

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

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

- LLDP-MIB

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

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

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

Ingress and Egress VLAN

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

About Ingress Port VLAN Tagging

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

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

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

Each port can only have one VLAN tag. The same VLAN tag can be assigned to multiple network ports or to both ports in an inline network port pair. For details on VLAN tagging for inline network groups, refer to *Configurable VLAN Tagging* on page 585.

VLAN tagging is supported in a cluster.

Refer to Figure 22-17 for an example. In the example, traffic from San Jose is tagged with VLAN 1001 and traffic from San Francisco is tagged with VLAN 1002.

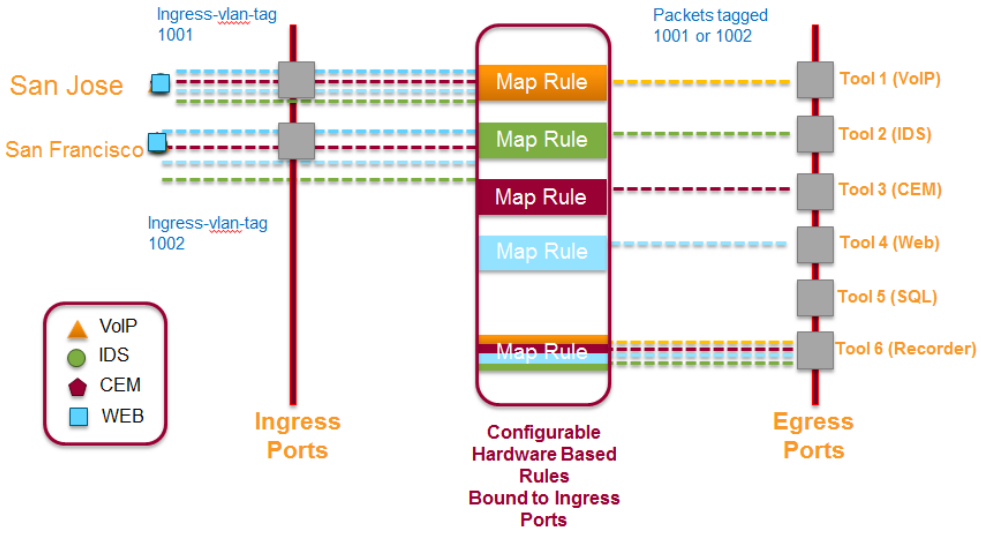

Figure 22-17: Using Ingress Port VLAN Tagging

Ingress Port VLAN Tagging

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

The VLAN ID is specified in the **VLAN Tag** field of the port configuration page for an network or inline network port. The value of the VLAN ID is specified as a number between 2 and 400.

Adding VLAN Tags

To add/modify VLAN tags, follow these steps:

1. Select **Ports** in the main navigation pane.
2. Select **Ports > All Ports** in the top navigation bar.
3. Click on the **Port ID**. Ensure that this Port ID is set as a network port.

   The Quick View window for the Port ID displays.
4. Select **Edit** from the top right corner of the Quick View Window.
5. Add the **VLAN ID** to the parameter field **VLAN Tag** and click **Save**.

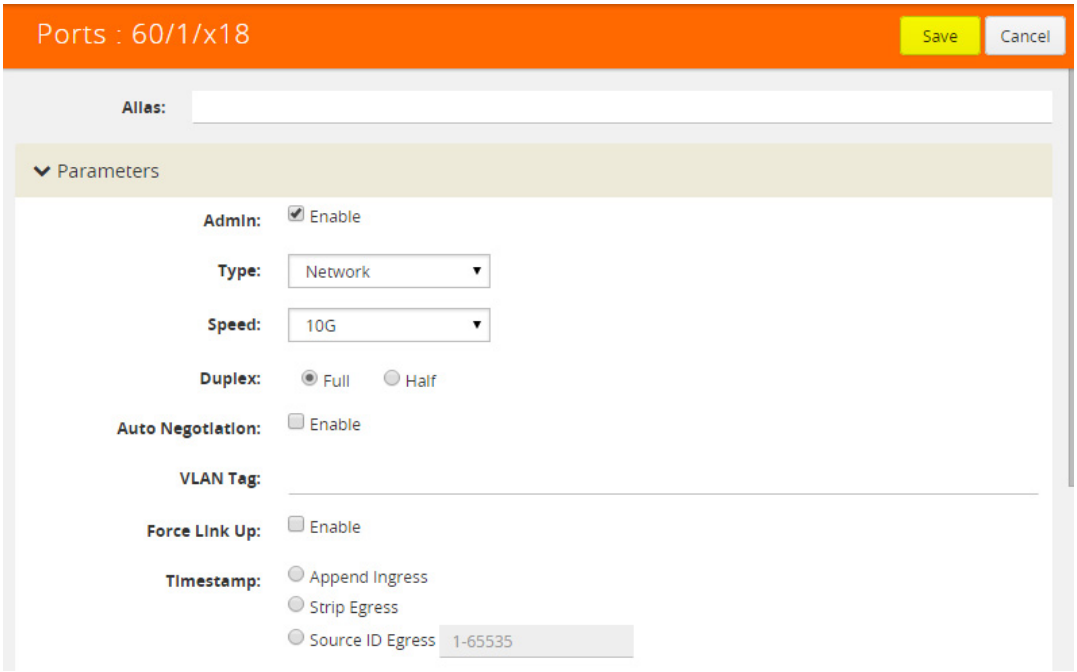

Figure 22-18: H-VUE Screen for Port-ID Configuration

Deleting VLAN Tags

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

Using VLAN Tags in Maps

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

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

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

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

Ingress Port VLAN Tag Limitations

The following sections describe limitations of ingress port VLAN tagging:

- *Second Level Maps* on page 427
- *Double-Tagged Packets* on page 427
- *IP Interfaces* on page 427

Second Level Maps

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

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

Double-Tagged Packets

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

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

IP Interfaces

For IP interfaces, a VLAN tag added at the network port of the encapsulation path (n1 in Figure 22-19) will become part of the payload going to the decapsulation path. But a VLAN tag added at the network port of the decapsulation path (n2 in Figure 22-19) will be available at the end tool port for filtering (t2 in Figure 22-19).

Refer to Figure 22-19 on page 428. VLAN tag (vlan1) added at the encap network port (n1) is encapsulated in the tunnel payload and cannot be used for filtering at the decap side. VLAN tag (vlan2) added at the decap network port (n2) can be used in a filter rule to send packets to tool port (t2).

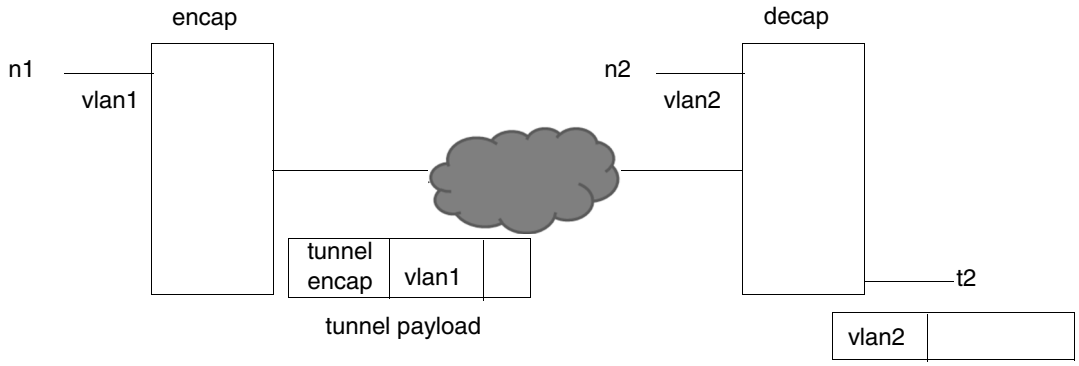

Figure 22-19: IP Interfaces

Configure Egress Port VLAN Stripping

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

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

Enable Egress Port VLAN Stripping

To enable egress port VLAN stripping:

1. Select a tool or hybrid port on the Ports page. Refer to Figure 22-20 on page 428.

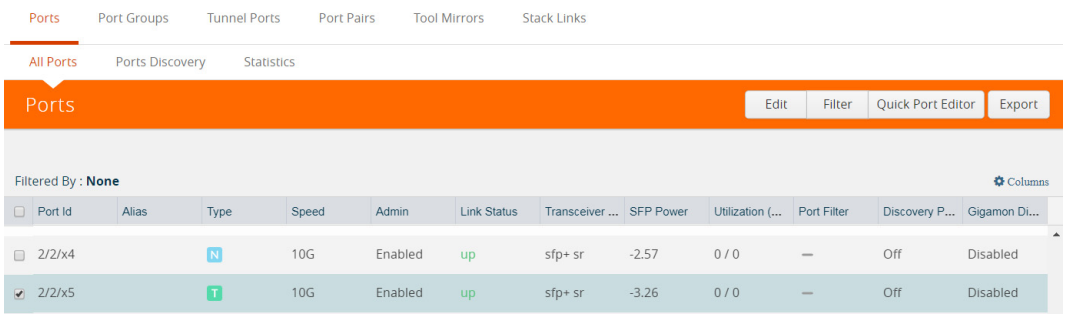

Figure 22-20: Select Tool or Hybrid Port

2. Click **Edit**.
3. Under Parameters, for Egress Vlan Tag, select Strip. Refer to Figure 22-21 on page 429.

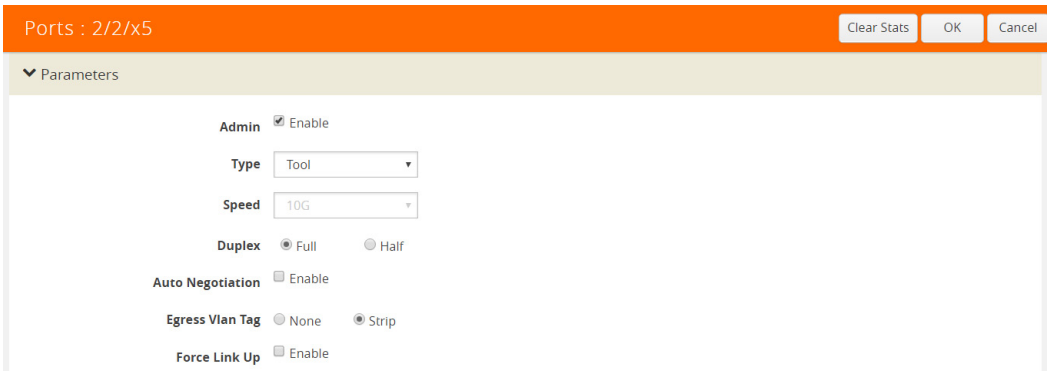

Figure 22-21: Select Strip Egress Vlan Tag

4. Click **OK**.

Disable Egress Port VLAN Stripping

Once egress port VLAN stripping is enabled, it can be disabled. In Figure 22-21 on page 429, for Egress Vlan Tag, select None.

Display Egress Port VLAN Stripping

To display egress port VLAN stripping configuration:

1. Double-click a tool or hybrid port on the Ports page. View the configuration under Port Info. Refer to Figure 22-20 on page 428.

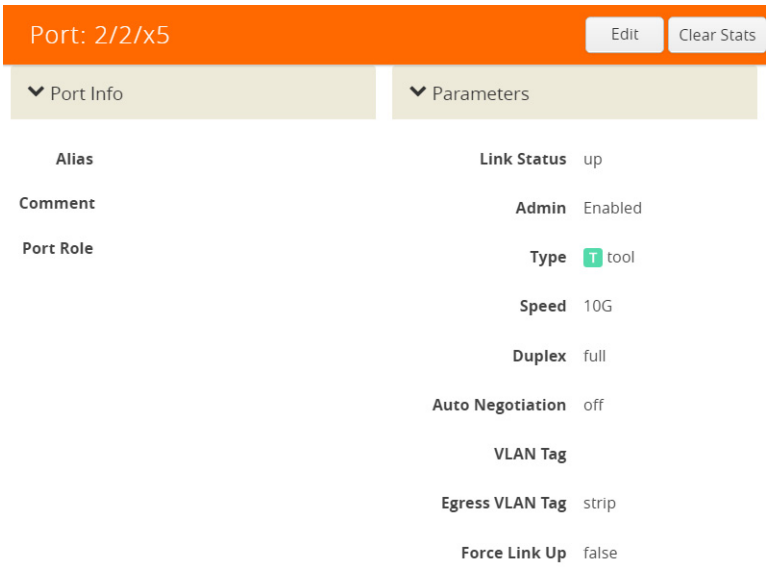

Figure 22-22: Double-Click Tool or Hybrid Port

Egress Port VLAN Stripping Limitations

The following are limitations of egress port VLAN stripping:

- Enabling both ingress port VLAN tagging and egress port VLAN stripping on the same port is not supported.
- Egress port VLAN stripping does not support inline tool ports or stack ports.
- Egress port VLAN stripping is not supported on the following: GigaVUE-HB1 nodes.
- If a port is configured for egress port VLAN stripping, configuring a port filter with either pass or drop VLAN rules is not recommended.
- Egress port VLAN stripping with outer tag ethertype of 0x88A8 or 0x9100 is supported if the tool port is local to the network port node (in the case of clustering). An outer tag of ethertype 0x8100 works without any limitation when egress port VLAN stripping is used in conjunction with ingress port VLAN tagging on the upstream GigaVUE node.

How to Use Both Ingress Tagging and Egress Stripping

When ingress port VLAN tagging is enabled on a network port and egress port VLAN stripping is enabled on a tool port on the same GigaVUE node, refer to the Table 22-3:

Table 22-3: VLAN Stripping Table

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

NOTES:

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

How to Use GigaStream

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

- *About GigaStream* on page 431
- *Regular GigaStream* on page 431
- *Controlled GigaStream* on page 438
- *Advanced Hashing* on page 446

About GigaStream

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

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

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

For details on regular GigaStream, refer to *Regular GigaStream* on page 431.

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

For details on controlled GigaStream, refer to *Controlled GigaStream* on page 438.

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

Regular GigaStream and controlled GigaStream differ in the following ways:

- how traffic is distributed based on hashing
- how traffic fails over when a port goes down
- how the configuration can be edited, such as adding ports on the fly

Regular GigaStream

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

All ports in a GigaStream must be running the same speed, such as 10Gb or 40Gb, and must use the same port type, either tool or stack. All ports in a GigaStream can be on different modules of the same GigaVUE-HC2 or GigaVUE-HC3 node.

With regular GigaStream, the hashing is computed based on traffic. Incoming packets arriving through network ports are processed with various map rules and then directed

to ports. The result of the hash distributes traffic equally across the GigaStream members. Refer to Figure 22-23 on page 432.

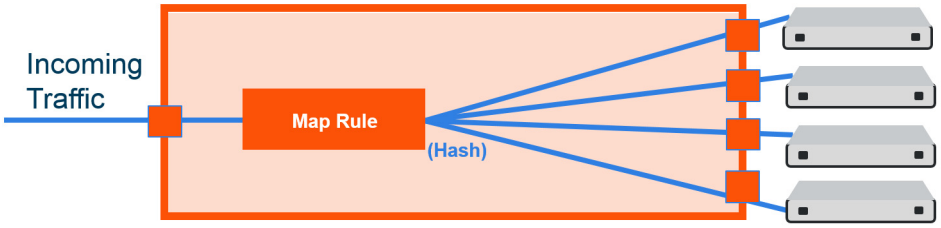

Figure 22-23: Regular GigaStream Overview

With the advanced hashing mode, the hash is a result of multiple parameters such as source MAC address, destination MAC address, source IP address, destination IP address, protocol, or other criteria. The hashing algorithm determines the destination tool port for a particular packet. All packets matching a particular set of hashing criteria will be sent to the same port. Sessions are maintained within a stream.

For example, a regular tool GigaStream is configured with ports x1 to x4. The hash table of size 4 is evenly divided among the 4 ports, and the traffic is distributed accordingly. For more information on how traffic is distributed with regular GigaStream, refer to *Traffic Distribution Across Controlled GigaStream* on page 444.

Regular Tool GigaStream

A regular tool GigaStream can be used as a single addressable destination, allowing you to overcome tool port oversubscription issues.

NOTE: A regular tool GigaStream can consist of tool ports or hybrid ports.

Refer to Figure 22-24 on page 433.

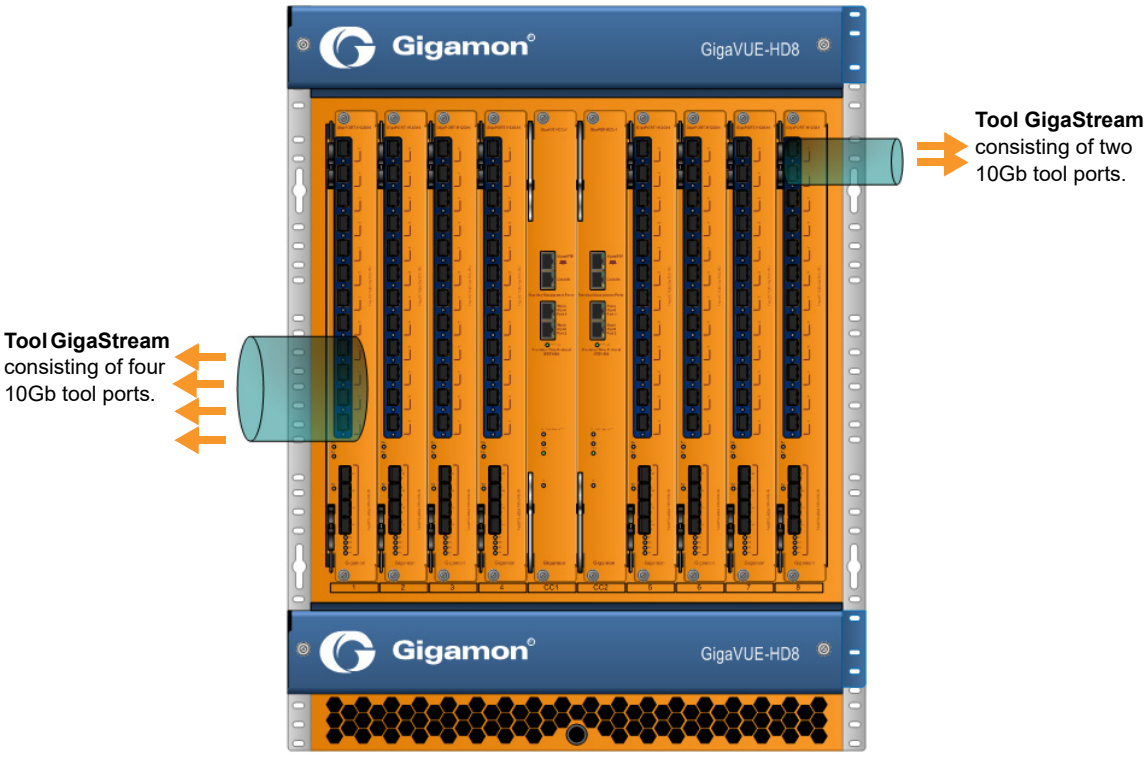

Figure 22-24: Regular Tool GigaStream Illustrated

Regular Stack GigaStream

A regular stack GigaStream can use stack-links to distribute data between GigaVUE-OS nodes operating in a cluster. With the terabits of throughput possible in a GigaVUE H Series node, using only one 10Gb port for a stack-link could cause a bottleneck. A regular stack GigaStream dramatically increases the bandwidth available for stack-link connections, providing greater flexibility and throughput within a cluster.

Refer to Figure 22-25 on page 434.

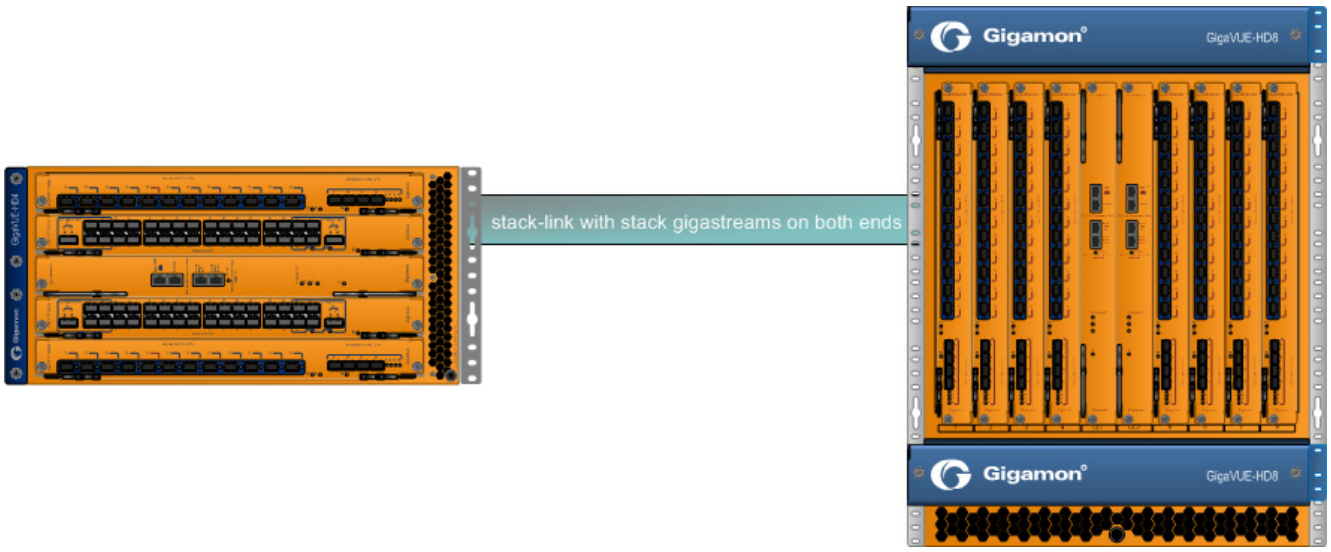

Figure 22-25: Regular Stack GigaStream Illustrated

Configure Regular GigaStream

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

Also, refer to *Advanced Hashing* on page 446 for optional advanced hashing settings and *Weighted GigaStream* on page 455 for optional weighting settings.

Before you configure a regular GigaStream, ensure that you have configured the required ports—tool, hybrid, stack, or Circuit. For information about configuring a port, refer to *Configure Ports* on page 408.

To configure a regular GigaStream:

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

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

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

9. In the **Drop Weight** field, enter the relative weight to drop the traffic. For example, if you enter 2 in this field, 2% of the total traffic entering the regular GigaStream will be dropped.

   NOTE: The **Weighting** and the **Drop Weight** fields are not available when you configure a regular stack GigaStream.

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

The configured regular GigaStream appears in the table in the GigaStream™ page.

Edit Regular GigaStream

Starting with software version 5.4, GigaSMART provides support for editing a regular GigaStream. You now have the benefit of adding and deleting tool ports from GigaStreams without the need to recreate the GigaStream with a new map.

Support is available to edit Tool and Hybrid GigaStream's attached to the following maps types:

1. Level one GS map
2. 2nd level GS map (includes GTP overlap sampling maps)
3. Regular map (including port mirroring, collector, etc.)

NOTE: Regular GigaStream and controlled GigaStream are interchangeable. You can change the type of GigaStream from regular to controlled in real-time.

To edit a regular tool GigaStream attached to a map:

1. Select **Ports > Port Groups > GigaStreams**.
2. Select a GigaStream and click **Edit** to open the GigaStream configuration page.

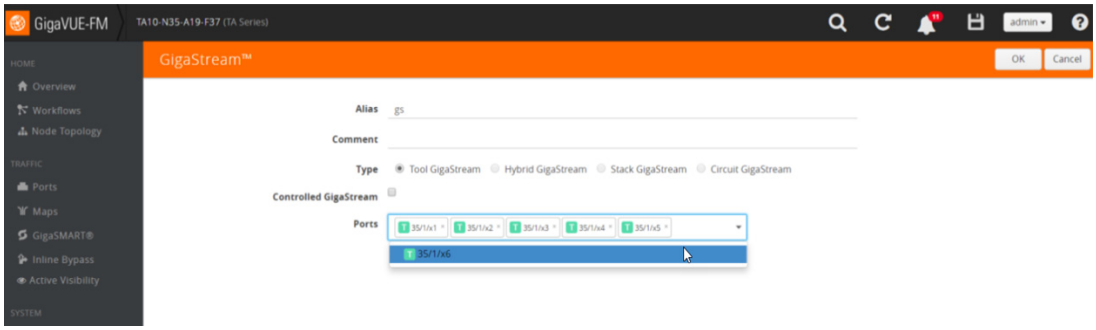

3. Click in the **Ports field** and **Change**, **Add** or **Delete** tool ports associated with the GigaStream as needed.

4. Click **OK**.

Edit Regular Stack GigaStream

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

To edit a regular stack GigaStream:

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

Traffic Distribution Across Regular GigaStream

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

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

For more information on hashing, refer to *Advanced Hashing* on page 446.

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

- The criteria configured using the Advanced Hash Setting page for the selected line card or chassis. (Click **Advanced Hash Setting** on the GigaStream page to open the Advance Hash Setting.) Because traffic is hashed across member ports rather than divided evenly, the bandwidth available for a regular GigaStream is not a straight multiple of the number of ports in the bundle – some flows will use more bandwidth than others.

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

The distribution described in this section applies to GigaVUE-HC2 and GigaVUE-HC3 modules, GigaVUE-HC1 nodes, and GigaVUE TA Series nodes for regular tool GigaStream and regular stack GigaStream.

- Weighting mode and hash weights assigned to the different ports in the regular GigaStream. For more information about Weighted GigaStream, refer to *Weighted GigaStream* on page 455.

Regular GigaStream Failover Protection

Regular GigaStream has built-in failover protection. When a tool port goes down, all the hash values in the GigaStream are redistributed among the remaining available ports. Failover is automatic.

The reassignment of hash values is applied to all the remaining ports in the GigaStream. This will result in sessions being reassigned to different tools, even if the tools remained healthy.

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

Resilient GigaStream

Resilient GigaStream Failover

In a GigaVUE Operating system, when there is a failover of a port that is part of a Resilient GigaStream, the traffic is redistributed to the other tool ports without disturbing the session continuity.

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

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

Resilient Gigastream is supported on all GigaVUE HC Series, and TA Series nodes and on Hybrid GigaStream, Tool GigaStream, and Circuit GigaStream.

Port Down in a Resilient GigaStream

When a tool port in the resilient GigaStream fails, the sessions allocated to that port alone are redistributed to other ports.

Port Up in a Resilient GigaStream

When a tool port resumes after a failure, automatically the system reasserts the originally assigned sessions before the failure.

Add or Remove Port in Resilient GigaStream

When a port is removed from a Resilient GigaStream, only the sessions allocated to that particular port is redistributed to other available ports in the GigaStream.

When a new tool port is added to the Resilient GigaStream, the system does not rehash all the sessions. Instead, a fraction of the hash values is moved from the original tools to the new tool.

When you add a healthy port to a Resilient GigaStream, which has few ports that are in the down state, the hash buckets of the down ports need to be assigned to the newly added port. When the down ports resume after the failure, the hash buckets need not be reassigned to the ports.

Add Down Port in GigaStream

When a down port is added to a resilient GigaStream, no sessions are assigned to the new port until the port is up. This prevents the session loss to any existing tools, when the port is added to the GigaStream.

When the port becomes active, the sessions are allocated to the new port.

Remove Down Port in GigaStream

When a port goes down in a resilient GigaStream, the sessions corresponding to the down port are distributed to the remaining ports. When you remove the down port from the GigaStream, then the hash values added from this port of the GigaStream becomes the original values of other ports.

Controlled GigaStream

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

All GigaVUE H Series and TA Series nodes support controlled GigaStream, with the following distinctions:

- GigaVUE-HB1 is not supported.

GigaVUE nodes with controlled GigaStream are supported in a cluster environment.

Controlled GigaStream can only be used as a packet egress destination (tool GigaStream). All port speeds are supported.

With controlled GigaStream, the hashing is computed based on traffic. There is a configurable number of hash buckets, from 1 to 256. The hash size of a controlled GigaStream specifies the number of logical tools the traffic will be distributed across. Refer to Figure 22-26 on page 439.

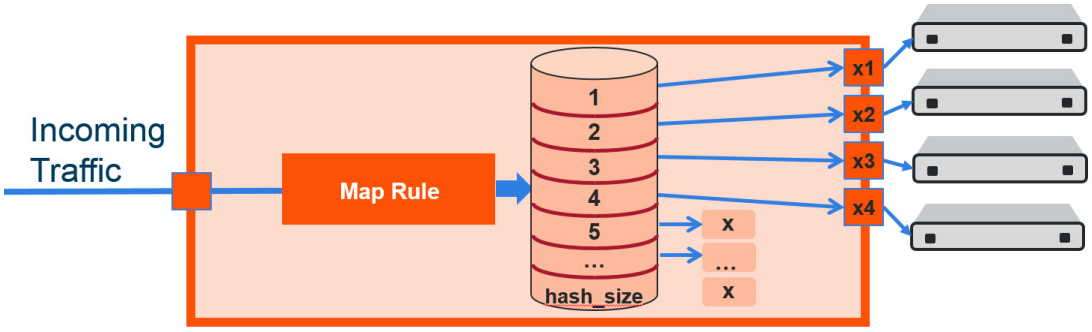

Figure 22-26: Controlled GigaStream Overview

Controlled GigaStream can manage network port bandwidth hashed to GigaStream tool ports. For example, if there is 10Gb of distributed traffic coming in on network ports directed to a GigaStream, and the tools connected to each tool port of the GigaStream can handle only 2Gb of bandwidth, the GigaStream can distribute the streams to 5 tool GigaStream ports. The ingress bandwidth divided by the number of tools determines the number of hash buckets.

Not all hash buckets need to be mapped to ports. In Figure 22-26 on page 439, four buckets are mapped to ports, while the remaining buckets are black holed. This provides a form of sampling, that is, only a sample of traffic is sent to the tools.

To determine the best hash size to use for your monitoring needs, divide the maximum bandwidth being sent to the tools by the bandwidth that can actually be consumed by the tools. For example, if you have 150 Gb of traffic, but the tools can only process 3 Gb, the recommended hash size is 150/3 = 50. To have completely even distribution across the logical tools, round up to the nearest power of 2. In this example, round up a hash size to 64.

When you have more bandwidth than the tools can process, you can use controlled GigaStream to restrict the amount of traffic sent to each tool. The hash size is determined by:

- the amount of traffic to be monitored, for example 300Gb
- the maximum bandwidth of the monitoring tools, for example 2.5Gb

Then divide (300/2.5=120), and round up to a power of 2 (for example, 16, 32, 64, 128). In this case, the hash size would be 128.

Another use for a controlled GigaStream is to increase the reliability of tool ports. For example, a trunk size of 5 is configured on 4 ports with 1 hash bucket each, port x1 is allocated or mapped to hash bucket ID 1, port x2 is mapped to hash bucket ID 2, port x3 is mapped to hash bucket ID 3, and port x4 is mapped to hash bucket ID 4. Hash bucket ID 5 is not mapped to a port. It can be reserved to be mapped to a port later. Until then, any traffic hashed to bucket 5 will be black holed. Refer to Figure 22-27 on page 440.

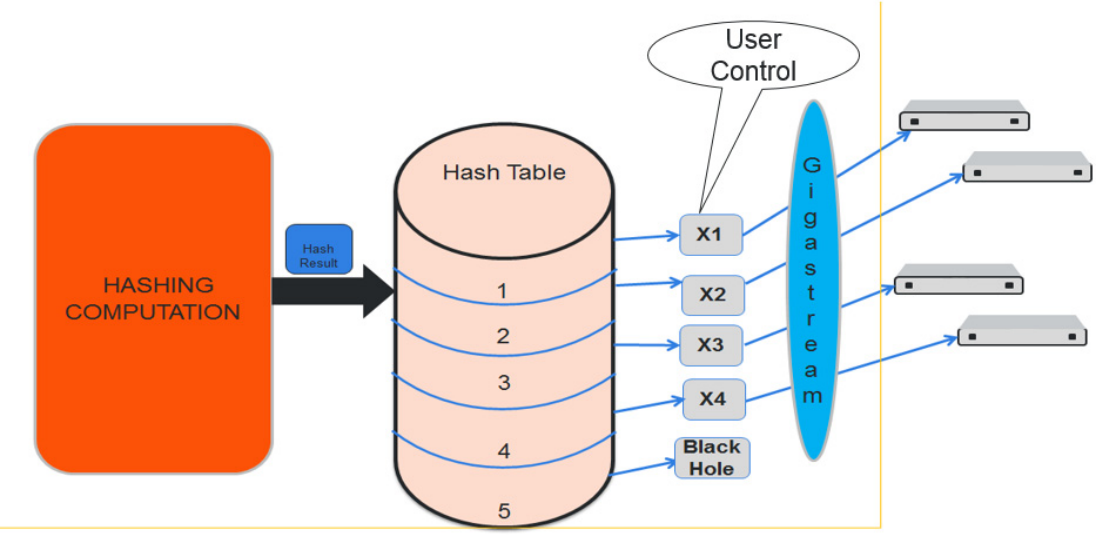

Figure 22-27: Controlled GigaStream Example

Generally, if there are only four tools available, with controlled GigaStream, a GigaStream trunk size of 5 can be configured and allocated to the available four tool ports of the GigaStream. The fifth tool port can be reserved and you can attach that port to the GigaStream whenever it is needed. The existing traffic streams are not impacted.

GigaStream controlled traffic distribution provides enhanced control of traffic hashed across the trunk ports as compared to regular GigaStream. The GigaStream trunk size is configurable, and ports can be dynamically added and deleted.

To configure controlled GigaStream, there are two parameters needed as follows:

- the hashing trunk size, which defines the number of hash buckets to be configured. It defines the maximum number of hash bucket IDs, from 1 to 256.
- the hash bucket ID, which specifies the mapping to ports. Each member of the trunk is mapped to a hash bucket ID. Mapping a port to a hash bucket ID makes it part of a GigaStream. The mapping to ports is static. When a port goes down, traffic is not re-hashed to the remaining ports.

In general, a controlled GigaStream is defined with a hash size equal to the number of trunk ports expected. Mapping a hash bucket ID to each trunk port will evenly distribute the traffic among the ports.

A particular trunk member can be mapped to multiple hash bucket IDs. If one tool can handle 4Gb, 2 hash bucket IDs can be mapped to that tool port. A trunk member that is configured to two hash bucket IDs will be two times more likely to receive hashed traffic as compared to a trunk member with one hash bucket ID. Thus, more traffic can be sent to the higher capacity tools in the GigaStream.

NOTE: A hash bucket ID cannot be mapped to multiple ports.

For more information, refer to *Traffic Distribution Across Controlled GigaStream* on page 444.

Notes and Considerations for Controlled GigaStream

Refer to the following notes and considerations for controlled GigaStream:

- Controlled GigaStream can be used with a regular map, map-passall, or map-scollector.
- Controlled GigaStream supports tool ports, but not inline tools or inline tool groups.
- Controlled GigaStream does not support stack or hybrid port types.
- Controlled GigaStream does not support a GigaSMART operation (gsop), or the first and second level maps associated with it.
- The maximum hash size is 256 per trunk.
- All the tool ports participating in the GigaStream must be on the same node. GigaStream can be created across GigaVUE-HC2 or GigaVUE-HC3 modules.
- All participating ports in the GigaStream must be running the same speed and must use the same port type.
- A **Controlled GigaStream** checkbox is used on the GigaStream configuration page for controlled GigaStream, enabling the prefix mode for specifying hash size and hash bucket IDs.
- Before attaching a controlled GigaStream to a map, it should be configured with at least one port.
- Controlled GigaStream can be modified on the fly, even after it is attached to a map. Refer to *Edit Regular GigaStream* on page 435.
- If the GigaStream is already attached to a map, the last mapped hash bucket ID cannot be deleted. That is, do not delete all the ports from a controlled GigaStream.

Controlled GigaStream Configuration

To configure a controlled tool GigaStream, specify hash size and hash bucket ID.

The following are the steps to configure Controlled GigaStream:

1. Use the Quick Port Editor to configure ports as type tool for the controlled GigaStream.
2. Select **Ports > Port Groups > GigaStreams** and click **New** to open the GigaStreams configuration page.
3. Enter an name for the GigaStream in the **Alias** field. For example stream 2.
4. (Optional) Enter a comment in the **Comment** field. For example, controlled GigaStream.
5. For **Type**, select **Tool GigaStream**.
6. Select **Controlled GigaStream**. The the **Port** field changes to **Hash Size**.

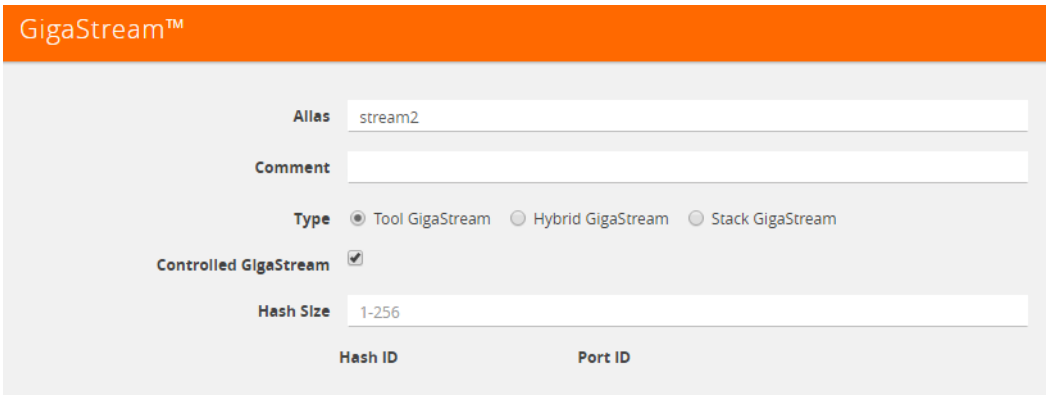

Figure 22-28: Controlled GigaStream Selected

7. In the **Hash Size** field, specify a hash size value. The range is 1 through 256.

   The hash size value determines the number of hash bucked IDs and ports available for assigning to the GigaStream. For example, in Figure 22-29 on page 442, the Hash Size is set to 5 so the GigaStream page displays five Hash IDs and five Port ID fields.

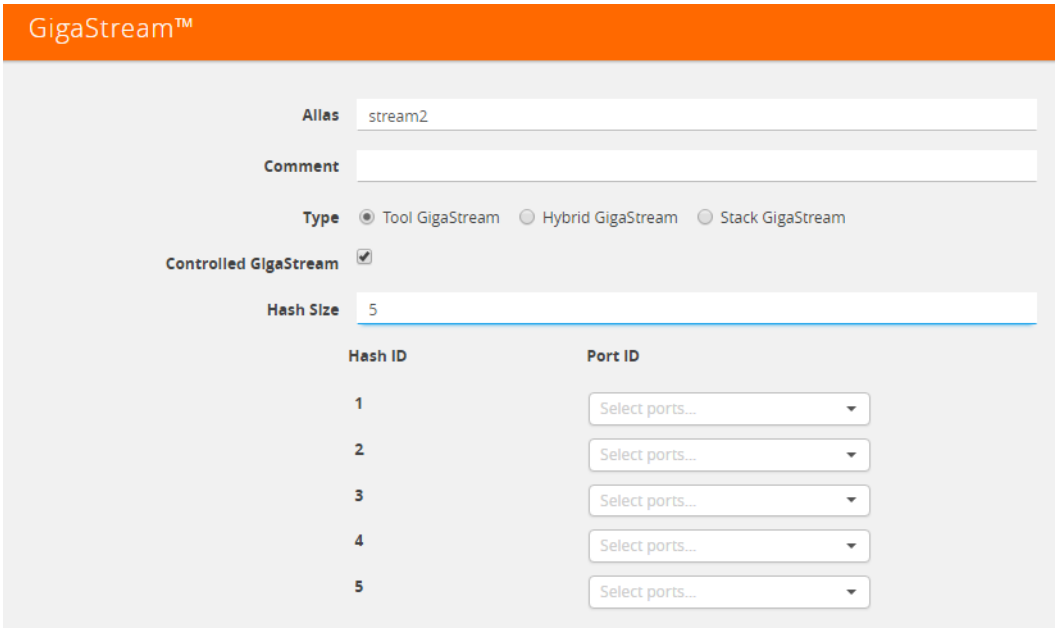

Figure 22-29: Controlled GigaStream with Five Hash IDs and Port IDs

8. Assign ports to the hash bucket IDs by clicking in each **Port ID** field and selecting a tool port.

   The example in Figure 22-30 on page 443 assign tool ports to hash buckets 1 though 4. Hash bucket 5 has no port assigned to it.

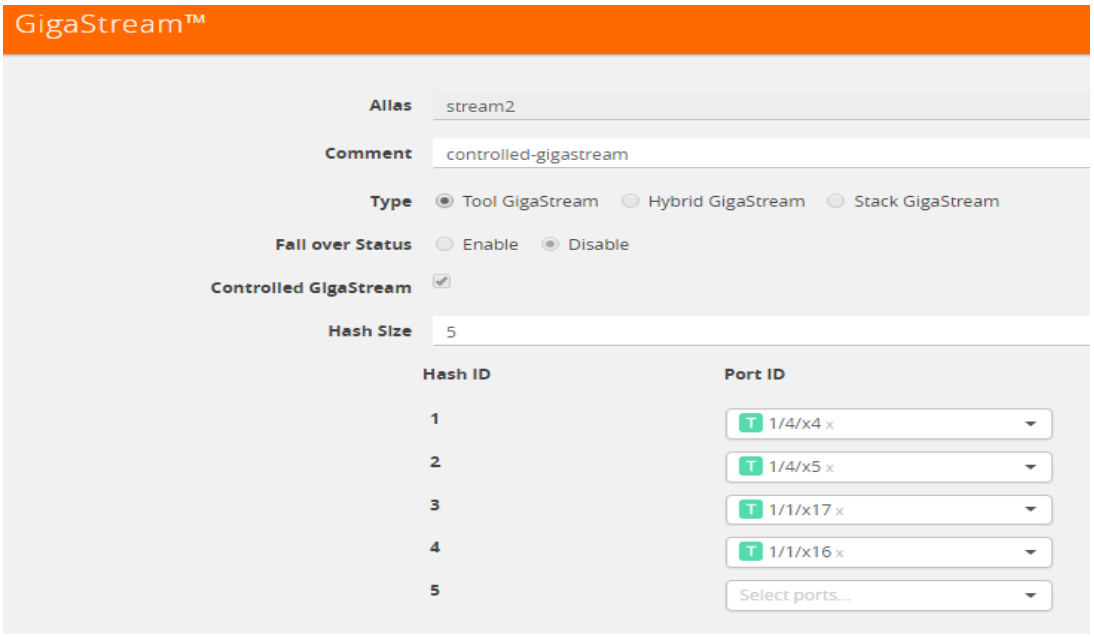

Figure 22-30: Ports Assigned to Hash Bucket IDs

9. Click **Save**.

After saving the controlled GigaStream, it appears on the GigaStreams page. Refer to Figure 22-31 on page 443.

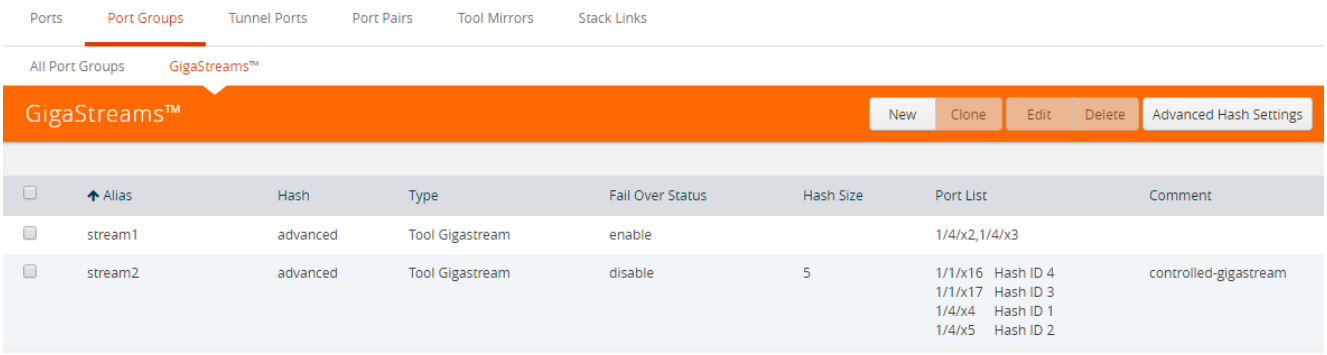

Figure 22-31: GigaStream Page with Controlled GigaStream

NOTE: For controlled GigaStream, the GigaStream page shows a Failover Status of disabled. When a port goes down, traffic is not re-hashed. Refer to *Failover and Controlled GigaStream* on page 446.

Edit Controlled GigaStream

A controlled GigaStream can be edited, even when the GigaStream is attached to a map. Unlike regular GigaStream, you can make changes without deleting the map or the GigaStream.

You have the control to map unused hash bucket IDs to any tool port dynamically, without deleting the trunk. This modification of a tool port mapping to a hash bucket ID

will not affect the streams flowing on the hash bucket IDs that are mapped to other ports. In addition, you can replace the mapping of any hash bucket ID to a port, dynamically.

If one of the GigaStream ports goes down, all the hash bucket IDs mapped to that port will be black holed until they are re-mapped to a new port, or until the port comes back up. This means that the packets sent to the remaining tools are unaffected.

If one port is receiving a lesser amount of bandwidth, the traffic can be reallocated to it. For example, if port x4 is underutilized, you have the control to reconfigure hash bucket ID 5 to also map to port x4. Then port x4 receives all the traffic that is hashed to hash bucket IDs 4 and 5.

You also have the flexibility to change the size of the trunk anytime, but this will require reprogramming of the whole hash table, so that might impact the existing streams.

Increasing the size of the trunk creates new hash buckets, which can be mapped to new or existing GigaStream tool ports. You can increase the bucket size per GigaStream. For example, if the bucket size is 4, you can increase it to 5.

If the size of the trunk has to be decreased, you have to take extra caution when releasing the hash bucket IDs gracefully, since they are mapped to GigaStream tool ports.

If you decrease the bucket size, empty out the bucket by unmapping buckets to ports. Also, do not reduce the hash size to less than the last occupied hash bucket ID.

NOTE: There is some packet drop associated with the following type of controlled GigaStream editing:

- adding a new port
- deleting an existing port
- changing the hash size

The Port Statistics page may display Discards in these cases, but not when additional buckets are added to the same port, or when a port goes down.

Traffic Distribution Across Controlled GigaStream

Controlled GigaStream has N buckets (where N is from 1 to 256) distributed across one or more ports, logical or physical.

Controlled GigaStream uses advanced hashing with 1 to 256 buckets. With controlled GigaStream, you define the number of buckets first, unlike with regular GigaStream.

Controlled GigaStream can manage network port bandwidth hashed to GigaStream tool ports. For example, if you are monitoring 500Gb traffic and have 10 tools, 50Gb per tool would be required. But if the tools cannot handle 50Gb, packets will be lost randomly.

With controlled GigaStream, first determine how much traffic the tools can process. For example, perhaps each tool can process 5Gb of traffic.

The formula is ingress bandwidth divided by tool capability. For example, 500Gb/5Gb = 100 tools. But if you only have 10 tools, you create a controlled GigaStream of 100 logical tools or 100 logical hash buckets, and then map only 10 of them.

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

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

Hash buckets IDs are mapped to ports as follows:

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

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

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

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

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

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

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

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

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

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

Failover and Controlled GigaStream

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

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

Advanced Hashing

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

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

- Tool port
- Hybrid port
- Circuit port

To open the Advanced Hash Settings page, select **Ports > Port Groups > GigaStream** and click **Advanced Hash Settings**. (For more details, refer to *Advanced Hash Settings* on page 447.) On the GigaVUE nodes, GigaStream hashing is per chassis, not per line card.

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

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

How to Change Advanced Hash Criteria

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

Advanced Hash Settings

The Advanced Hash Settings page is where criteria for the advanced hash algorithm is set. To open the page, select **Ports > Port Groups > GigaStream** and click **Advanced Hash Settings**. Refer to Figure 22-32 on page 447.

Figure 22-32: Advanced Hash Settings

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

Field	Description
Box	Identifies chassis to which the advanced algorithm will be applied.
Slot	Identifies the line card to which the advanced hash algorithm will apply. Each line card in certain GigaVUE H Series nodes has its own individual advanced hash algorithm. On GigaVUE-HC3, GigaVUE-HC2, GigaVUE-HC1, and GigaVUE-HB1, GigaStream hashing is per chassis, not per line card. For example, the slot field will only shown cc1 when configuring an GigaVUE-HC2.
Default	Sets the advanced hash algorithm to its default settings. By default, the advanced hash algorithm includes source/destination IPv4/IPv6 addresses and ports.
Custom	Clears the field from the advanced has and allows you to select specific own criteria.
All	Selects all criteria.
None	Clears all fields from the advanced hash.
IPv4	This area of the page lets you select the following criteria: • IPv4 Source Address—Adds IPv4 source IP • IPv4 Destination Address—Adds IPv4 destination IP • IPv4 Protocol—Adds IPv4 protocol • IPv4 Source Port—Adds IPv4 source port • IPv4 Destination Port—Adds IPv4 destination port
IPv6	This area of the page lets you select the following criteria: • IPv6 Source Address—Adds IPv6 source IP • IPv6 Destination Address—Adds IPv6 destination IP • IPv6 Next Header—Adds IPv6 next header field. • IPv6 Source Port—Adds IPv6 source port • IPv6 Destination Port—Adds IPv6 destination port
Layer2	This area of the page lets you select the following criteria: • Source MAC Address—Adds L2 source MAC • Destination MAC Address—Adds L2 destination MAC • Ether Type—Adds L2 ethertype field.
MPLS	This area of the page lets you select the following criteria: • MPLS Hash—Adds MPLS labels (up to three)
GTP TEID	This area of the page lets you select the following criteria: • GTP TEID—Adds GTP tunnel endpoint identifier
Ingress Port	This area of the page lets you select the following criteria: • Ingress Ports—Adds ingress port

Advanced Hash Examples

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

The example in Figure 22-33 on page 449 sets a **Custom** advanced hash method for slot 3 in box ID 5 that distributes traffic based on matching IPv4 source and destination addresses.

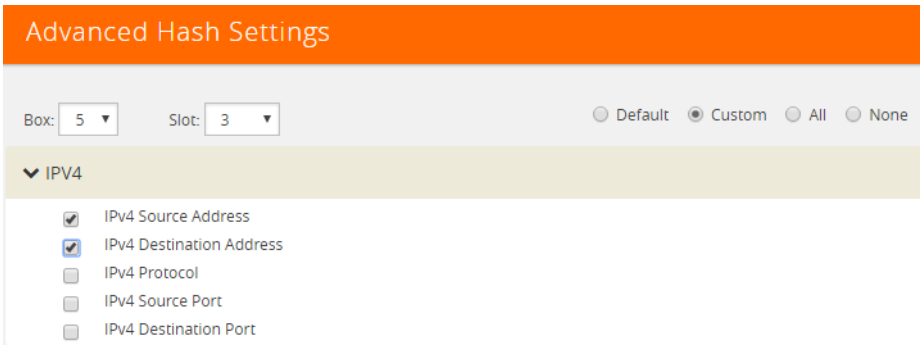

Figure 22-33: Advanced Hash with IPv4 Source and Destination Addresses

The example in Figure 22-34 sets the advanced hash for slot cc1 in box ID 2 to the **Default** criteria.

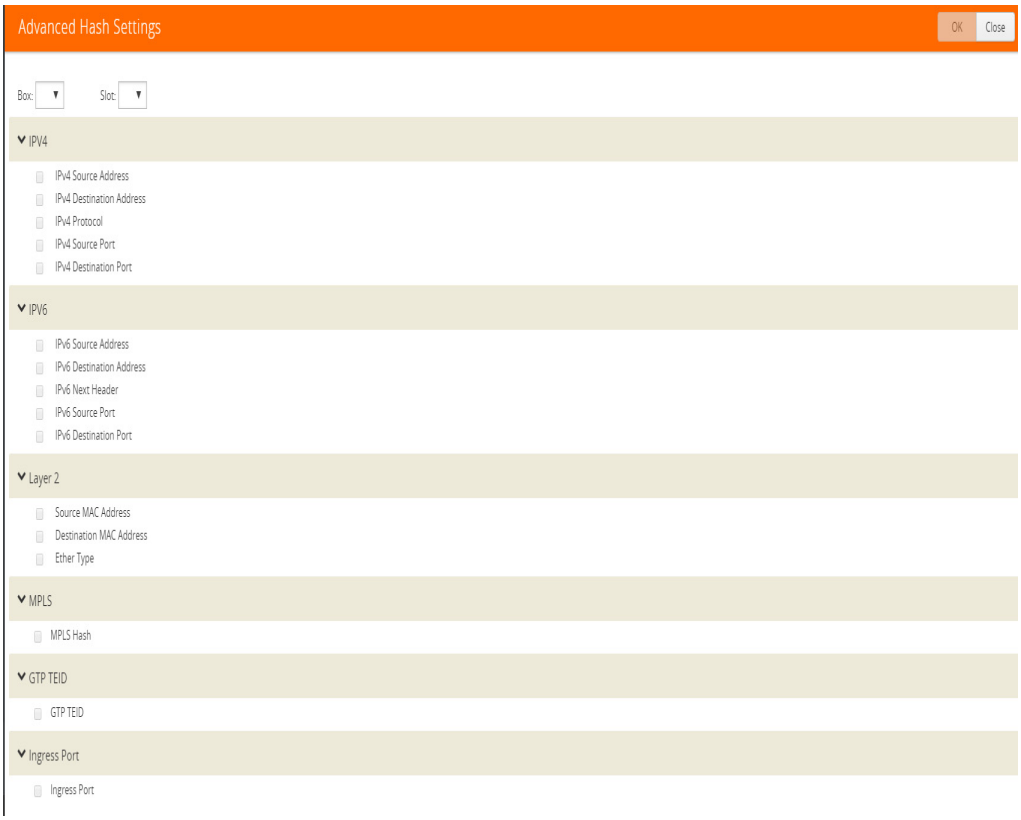

Figure 22-34: Advanced Hash Default Criteria

The example in Figure 22-35 on page 450 sets a **Custom** advanced hash for slot 5 in box ID 5 that distributes traffic based on matching source and destination MAC addresses.

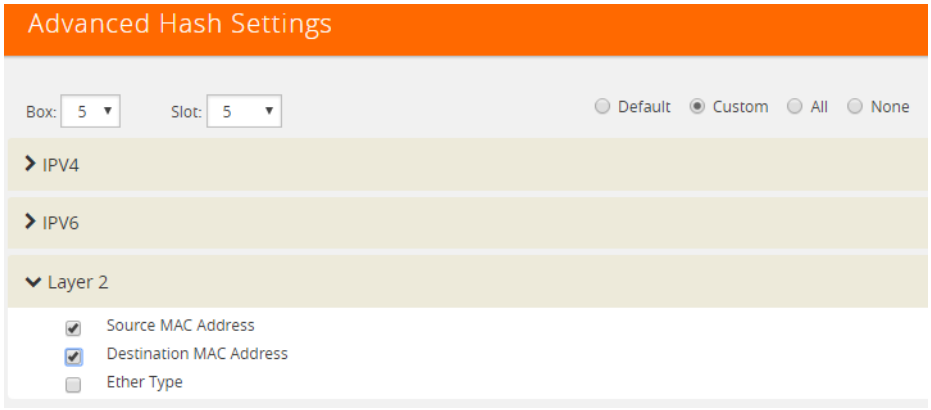

Figure 22-35: Advanced Hash with Source and Destination MAC Address

Hashing Behavior

Table 22-4 shows the possible hash criteria field combinations and the corresponding hashing behavior based on packet type for advanced hashing for non-MPLS packets. (Refer to Table 22-5 for MPLS packets.)

Table 22-4: Hashing Behavior Based on Hash Criteria Field Combinations

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

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

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

Notes and Considerations for Advanced Hashing

Refer to the following notes and considerations for advanced hashing:

- With symmetric hashing, packets with their source and destination IP addresses and Layer 4 (L4) ports interchanged will go to the same GigaStream port. It is recommended to enable source and destination IPv4 or IPv6 pairs and L4 source and destination ports.
- On GigaVUE-HC3 nodes, symmetric hashing is enabled for all GigaStream.
- On GigaVUE-HC2 nodes, symmetric hashing is enabled for stack GigaStream.
- For stack GigaStream on GigaVUE-HC2 (with control card version 1 or 2) and GigaVUE-HC3, the following limitation applies:
  - For non-MPLS IPv4 and IPv6 packets, the hashing is fixed to the following 3-tuple: ipsrc, ipdst, and protocol or ip6src, ip6dst, and protocol.
  - All other traffic follows the advanced hash settings.
- ASICs used in Gigamon devices do not support hashing of the IP header fields when there is a PPPOE header. Only Layer 2 (L2) fields can be used for such packets.

Advanced Hashing with MPLS

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

Advanced hashing with MPLS is supported on all GigaVUE H Series and TA Series nodes, with the following distinctions:

- GigaVUE-HB1 is not supported.

Use the **Advanced Hash Settings** to specify MPLS, which can detect up to three MPLS labels. Packets with one to three MPLS labels can be hashed, along with IP address fields, if present. If a packet has more than three MPLS labels, IP address fields after the third MPLS label cannot be hashed. Refer to Table 22-5 on page 452 for details of the hashing behavior.

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

MPLS hashing applies to the following:

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

Table 22-5 shows the possible hash criteria field combinations and the corresponding hashing behavior based on packet type for advanced hashing with MPLS packets. (Refer to Table 22-4 for non-MPLS packets.)

Table 22-5: Hashing Behavior Based on Hash Criteria Field Combinations

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

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

NOTES:

- No hash means that the packets will be sent to the first port in the GigaStream.
- If an MPLS packet has a router alert label as one of its labels, the router alert label is skipped and the other available labels are used for hashing. For example, if a packet has four labels and the second label is the router alert, the first, third, and fourth labels are used for hashing.

Advanced Hashing with GTP TEID

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

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

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

- The following nodes are supported: GigaVUE-HC1, GigaVUE-HC2 (with control card version 1 or 2), GigaVUE-HC3, GigaVUE-TA40, and GigaVUE-TA100.
- The following nodes are not supported: GigaVUE-HB1, GigaVUE-TA1, and GigaVUE-TA10.

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

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

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

Refer to Table 22-6 on page 454 for details of the hashing behavior.

GTP TEID hashing applies to the following:

- regular GigaStream
- controlled GigaStream
- stack GigaStream (except GigaVUE-HC2 and GigaVUE-HC3)
- inline tool groups

NOTE: GTP TEID hashing is not supported for stack GigaStream on GigaVUE-HC2 and GigaVUE-HC3 in this software version due to the 3-tuple limitation listed in *Notes and Considerations for Advanced Hashing* on page 451.

Table 22-6 shows the possible hash criteria field combinations and the corresponding hashing behavior based on packet type for advanced hashing with GTP packets.

Table 22-6: Hashing Behavior Based on Hash Criteria Fields: GTP TEID

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

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

Packet Distribution and the Advanced Hash Algorithm

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

Weighted GigaStream

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

Weighted GigaStream is supported on the following:

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

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

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

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

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

GigaStream Rules and Maximums

The following rules apply to regular GigaStream and controlled GigaStream:

GigaStream Rule	Description
Port Location	All participating ports must be on the same GigaVUE node. On the GigaVUE-HC2 and GigaVUE-HC3, GigaStream can be across modules.
Speed Requirements	• All participating ports must be running the same speed (1Gb, 10Gb, 40Gb. or 100Gb) and must use the same port types (for example, all g, x, q, or c). • A stack GigaStream must consist of ports with 10Gb speed or higher. • The system will not let you change the speed of any port participating in a tool GigaStream. Keep in mind that the only ports that allow speed changes through are g***x*** ports. • You can use g***x*** ports in a regular tool GigaStream, but only with ports running at the same speed (and no slower than 1000Mb).
Addressing	Once a port belongs to a GigaStream, it must be addressed by its GigaStream alias. It can no longer be addressed as an individual port. For example, if tool port 1/1/x4 is part of a tool GigaStream, the GigaVUE H Series node prevents you from using it as the destination for a map rule.
Stack Ports	Stacking ports must be 10Gb or higher. Therefore, the Maximum Stack Ports per GigaStream for any 1Gb port is N/A in Table 22-7 on page 457 to.Table 22-15 on page 461
SFP+	For SFP+ ports that can operate at 10Gb or 1Gb, refer to the values in the Maximum Tool Ports per GigaStream column for the 10Gb Ports rows in Table 22-7 on page 457 to Table 22-15 on page 461.

For GigaStream maximums, refer to *Maximum Ports per GigaStream* on page 457, which have been updated in software version 5.1.

Maximum Ports per GigaStream

Table 22-7 to Table 22-15 list the maximum ports per GigaStream for GigaVUE nodes.

Table 22-7: GigaVUE-HC3: Maximum Ports per GigaStream

GigaVUE-HC3	Ports Allowed	Maximum Tool Ports per GigaStream	Maximum Stack Ports per GigaStream
100Gb Ports			
PRT-HC3-C08Q08 (QSFP28)	c1..c8	8	8
PRT-HC3-C16 (QSFP28) * Must be equipped with CCv2	c1..c16	16	16
SMT-HC3-C05 (QSFP28)	c1..c5	5	5
40Gb Ports			
PRT-HC3-C08Q08 (QSFP+)	c1..c8	8	8
PRT-HC3-C08Q08 (QSFP+) (2x40G mode)	c1q1..c1q2 to c8q1..c8q2	16	N/A
PRT-HC3-C16 (QSFP+) * Must be equipped with CCv2	c1..c16	16	16
SMT-HC3-C05 (QSFP+)	c1..c5	5	5
25Gb Ports * Must be equipped with Control Card version 2 (CCv2)			
PRT-HC3-X24 (SFP28)	x1..x24	24	N/A
PRT-HC3-C08Q08 (QSFP28) (4x25G mode)	c1x1..c1x4 to c8x1..c8x4	8	N/A
PRT-HC3-C16 (QSFP28) (4x25G mode)	c1x1..c1x4 to c15x1..c15x4 odd-numbered ports only	32	N/A
SMT-HC3-C05 (QSFP28) (4x25G mode)	c1x1..c1x4 to c5x1..c5x4	20	N/A
BPS-HC3-C25F2G (SFP28)	x1..x16	16	N/A
10Gb Ports			
PRT-HC3-X24 (SFP+)	x1..x24	24	N/A
PRT-HC3-C08Q08 (QSFP+) (4x10G mode)	c1x1..c1x4 to c8x1..c8x4	32	32
PRT-HC3-C16 (QSFP+) (4x10G mode) (8 ports)	c1x1..c1x4 to c15x1..c15x4 odd-numbered ports only	32	N/A
SMT-HC3-C05 (QSFP+) (4x10G mode)	c1x1..c1x4 to c5x1..c5x4	20	N/A
BPS-HC3-C25F2G (SFP+)	x1..x16	16	N/A

Table 22-7: GigaVUE-HC3: Maximum Ports per GigaStream

GigaVUE-HC3	Ports Allowed	Maximum Tool Ports per GigaStream	Maximum Stack Ports per GigaStream
BPS-HC3-C25F2G (4x10G mode)	c1x1..c1x4 to c4x1..c4x4	16	N/A

NOTE: GigaStream can be created across GigaVUE-HC3 modules. For example, with four PRT-HC3-C08Q08 modules in **4x10G** mode, the maximum tool ports per GigaStream on a GigaVUE-HC3 is 8 x 4 x 4 = 128, or with three PRT-HC3-X24 modules and one PRT-HC3-C08Q08 module in **4x10G** mode, the maximum tool ports per GigaStream on a GigaVUE-HC3 is (24 x 3) + (8 x 4) = 104.

You can also create stack GigaStream across GigaVUE-HC3 modules. The maximum number of ports in a stack GigaStream across the GigaVUE-HC3 node is 32.

Table 22-8: GigaVUE-HC2 Modules: Maximum Ports per GigaStream

GigaVUE-HC2 Module	Ports Allowed	Maximum Tool Ports per GigaStream	Maximum Stack Ports per GigaStream
100Gb Ports			
PRT-HC0-C02	c1..c2	2	N/A
40Gb Ports			
PRT-HC0-Q06	q1..q6	6	6
10Gb Ports			
PRT-HC0-X24	x1..x24	24	24
BPS-HC0-D25A4G BPS-HC0-D25B4G BPS-HC0-D35C4G	x1..x16	16	16
BPS-HC0-Q25A28	x1..x8	8	8
SMT-HC0-X16	x1..x16	16	16

NOTE: GigaStream can be created across GigaVUE-HC2 modules. For example, with four PRT-HC0-X24 modules, the maximum tool ports per GigaStream on a GigaVUE-HC2 is 24 x 4 = 96.

You can also create stack GigaStream across GigaVUE-HC2 modules. The maximum number of ports in a stack GigaStream across the GigaVUE-HC2 node is 48.

Table 22-9: GigaVUE-HC1: Maximum Ports per GigaStream

GigaVUE-HC1	Ports Allowed	Maximum Tool Ports per GigaStream	Maximum Stack Ports per GigaStream
10Gb Ports			
GigaVUE-HC1 base chassis	x1..x12	12	12
TAP-HC1-G10040	N/A	N/A	N/A
1Gb Ports			
GigaVUE-HC1 base chassis	g1..g4	4	N/A
BPS-HC1-D25A24	x1..x4	4	4

Table 22-10: GigaVUE-HB1: Maximum Ports per GigaStream

GigaVUE-HB1	Ports Allowed	Maximum Tool Ports per GigaStream	Maximum Stack Ports per GigaStream
10Gb Ports			
GigaVUE-HB1	x1..x4	4	4
1Gb Ports			
GigaVUE-HB1	g1..g16	8	N/A

Table 22-11: GigaVUE-TA1/GigaVUE-OS on a White Box: Maximum Ports per GigaStream

GigaVUE TA Series	Ports Allowed	Maximum Tool Ports per GigaStream	Maximum Stack ports per GigaStream
40Gb Ports			
GigaVUE-TA1/GigaVUE-OS on a white box (48x mode)	q1..q4	4	4
GigaVUE-TA1/GigaVUE-OS on a white box (56x mode)	q3..q4	2	2
GigaVUE-TA1/GigaVUE-OS on a white box (64x mode)	N/A	N/A	N/A
10Gb Ports			
GigaVUE-TA1/GigaVUE-OS on a white box (48x mode)	x1..x48	48	32
GigaVUE-TA1/GigaVUE-OS on a white box (56x mode)	x1..x56	56	32
GigaVUE-TA1/GigaVUE-OS on a white box (64x mode)	x1..x64	64	32

Table 22-12: GigaVUE-TA40: Maximum Ports per GigaStream

GigaVUE-TA40	Ports Allowed	Maximum Tool Ports per GigaStream	Maximum Stack Ports per GigaStream
40Gb Ports			
GigaVUE-TA40 (0x mode)	q1..q32	32	32
GigaVUE-TA40 (16x mode)	q5..q32	28	28
10Gb Ports			
GigaVUE-TA40 (0x mode)	N/A	N/A	N/A
GigaVUE-TA40 (16x mode)	x1..x16	16	N/A

Table 22-13: GigaVUE-TA100: Maximum Ports per GigaStream

GigaVUE-TA100	Ports Allowed	Maximum Tool Ports per GigaStream	Maximum Stack Ports per GigaStream
100Gb Ports			
GigaVUE-TA100 (QSFP28) (16 ports)	c1..c16	16	16
GigaVUE-TA100 (QSFP28) (24 ports)	c1..c24	24	24
GigaVUE-TA100 (QSFP28) (32 ports)	c1..c32	32	32
40Gb Ports			
GigaVUE-TA100 (QSFP+) (16 ports)	c1..c16	16	16
GigaVUE-TA100 (QSFP+) (24 ports)	c1..c24	24	24
GigaVUE-TA100 (QSFP+) (32 ports)	c1..c32	32	32
10Gb Ports			
GigaVUE-TA100 (QSFP+) (4x10G mode) (16 ports)	c1x1..c1x4 to c16x1..c16x4	64	64
GigaVUE-TA100 (QSFP+) (4x10G mode) (24 ports)	c1x1,,c1x4 to c24x1..c24x4	96	64
GigaVUE-TA100 (QSFP+) (4x10G mode) (32 ports)	c1x1..c1x4 to c32x1..c32x4	128	64

Table 22-14: GigaVUE-TA100-CXP: Maximum Ports per GigaStream

GigaVUE-TA100-CXP	Ports Allowed	Maximum Tool Ports per GigaStream	Maximum Stack Ports per GigaStream
100Gb Ports			
GigaVUE-TA100-CXP (CXP ports) (20 ports)	c1..c20	20	N/A
GigaVUE-TA100-CXP (QSFP28 ports) (8 ports)	c21..c28	8	N/A
40Gb Ports			
GigaVUE-TA100-CXP (QSFP28 ports) (8 ports)	c21..c28	8	N/A
10Gb Ports			
GigaVUE-TA100-CXP (QSFP+) (4x10G mode) (8 ports)	c21x1..c21x4 to c28x1..c28x4	32	N/A

Table 22-15: GigaVUE-TA200: Maximum Ports per GigaStream

GigaVUE-TA200	Ports Allowed	Maximum Tool Ports per GigaStream	Maximum Stack Ports per GigaStream
100Gb Ports			
GigaVUE-TA200 (QSFP28)	c1..c64	64	16
25Gb Ports			
GigaVUE-TA200 (QSFP28) (4x25G mode) (32 ports)	c1x1..c1x4 to c10x1..c10x4 and c23x1..c23x4 to c32x1..c32x4	80	N/A
GigaVUE-TA200 (QSFP28) (4x25G mode) (64 ports)	c33..c64	N/A	N/A
10Gb Ports			
GigaVUE-TA200 (QSFP+) (4x10G mode) (32 ports)	with no restrictions: c1x1..c1x4 to c10x1..c10x4 and c23x1..c23x4 to c32x1..c32x4 with restrictions: c11x1..c11x4 to c22x1..c22x4	128	N/A

Port Statistics and Counters

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

The major sections in This section include:

- *Display Port Statistics* on page 462
- *How to Reset Traffic Counters* on page 467

Display Port Statistics

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

- A single port or view. For details, refer to *Display Port Statistics for a Single Port* on page 462
- All the ports. For details refer to *Display Statistics for All Ports* on page 464

Display Port Statistics for a Single Port

From the **Ports > Ports > All Ports**, select any port by clicking on the row. The quick view window that appears provides port statistics for that particular port. Each field is color-coded in the graphical representation as shown in Figure 22-36.

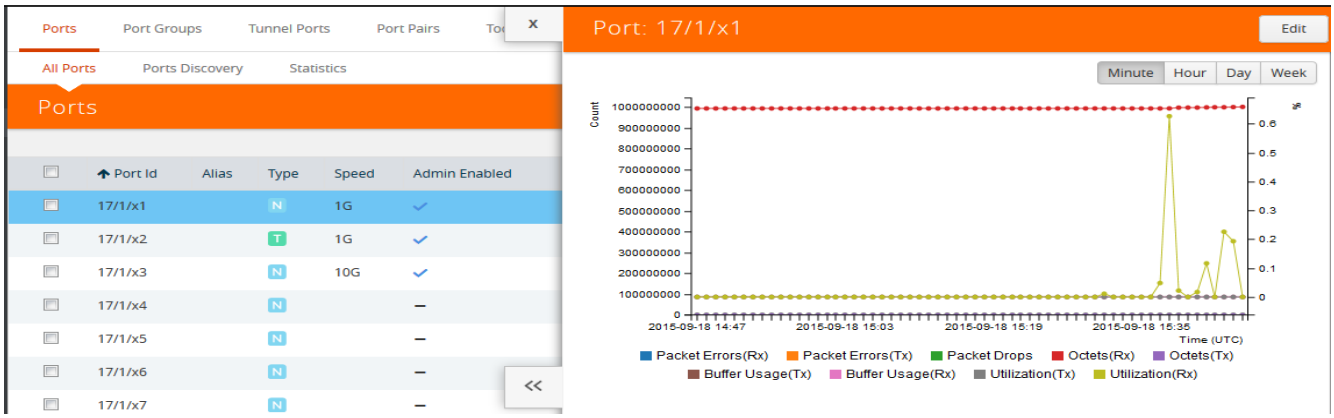

Figure 22-36: Quick View Window for Port Statistics

By hovering over the graph, numerical value for each of the data points is visible as shown in Figure 22-37.

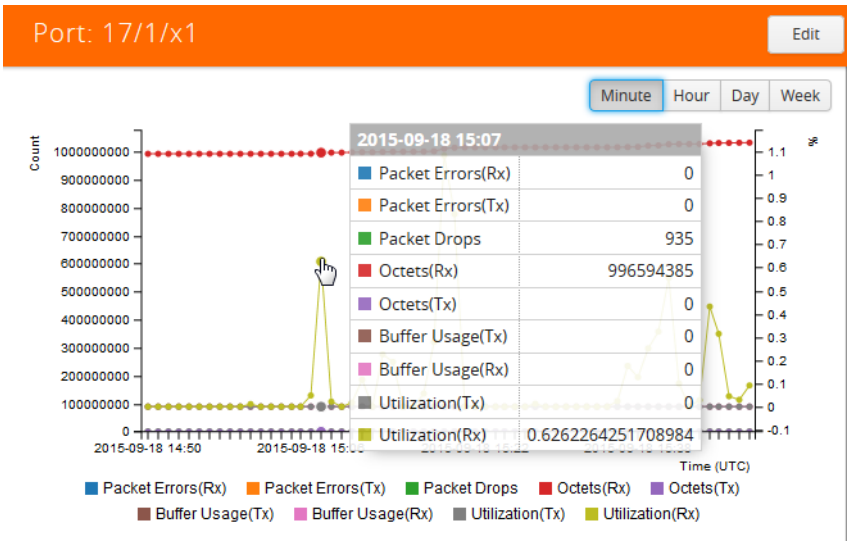

Figure 22-37: Numerical Values for Data Points on the Statistics Graph

You can modify the time lapse for measuring various data points by selecting the Minutes, Hour, Day, or Week button. Figure 22-38 shows the data points measured by hour.

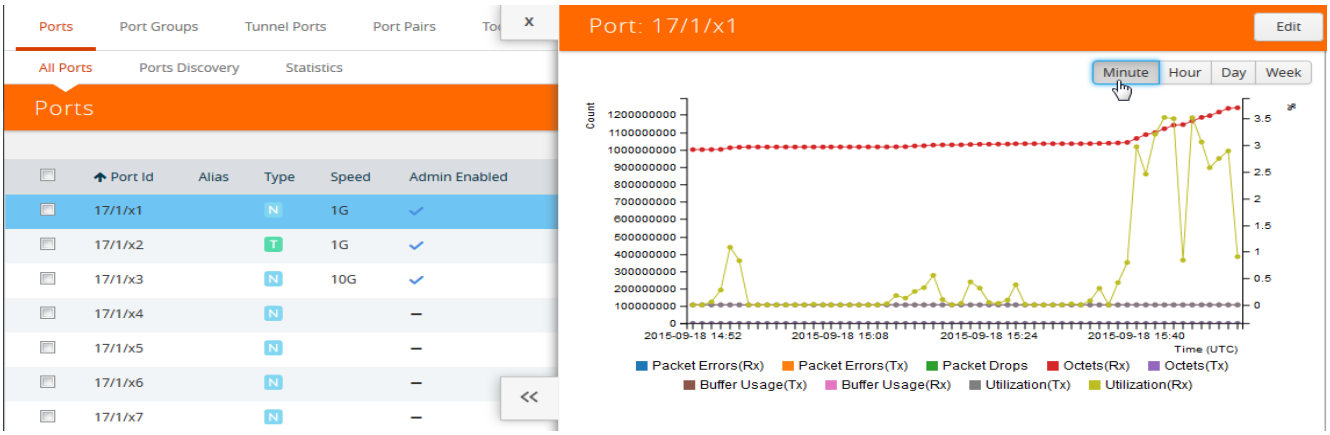

Figure 22-38: Numerical Values for the Data Points Measured by Hour

Table 22-16 describes the port statistics available.

Table 22-16: Port Statistics Definitions

Counter	Definition	Notes
Packet Errors	Total Error Packets Received or Transmitted This indicates hardware detected errors. Error packets include undersize, FCS/CRC, MTU exceeded, fragments, and oversize packets.	Excludes oversize packets without FCS/ CRC. Packets larger than the MTU setting arriving on a network port are counted twice in the counter. So 1000 oversize packets would show up as 2000. This double-counting only happens with Oversize error packets.
Discards	Discards Received or Transmitted	Discards are counted in the following cases: • Traffic arriving at a network port that is not logically connected using a map or map passall. • Map rules/map rules applied on a network port. • In packets on a tool port. • Pause frames.
Packet Drops	Total Dropped Packets Received or Transmitted	Packets are dropped when a port's bandwidth is exceeded due to oversubscription. Packets are dropped when they reach the port but before they are sent out.
Octets	Total Bytes Received or Transmitted Includes all valid and error frames with the exceptions noted in the adjacent columns.	Excludes undersize frames.
Buffer Usage	Percentage of buffer space used by packets transmitted or received	The buffer is used when the port reaches 100 percent utilization during a microburst. If the buffer reaches 100 percent utilization, packets may be dropped.
Utilization	Percentage of port utilization by packets received or transmitted	
Packets	Total Packets Received or Transmitted Excludes multicast packets, broadcast packets, packets with FCS/CRC errors, MTU exceeded errors, oversize packets, and pause packets.	Excludes packets with FCS/CRC errors.

Display Statistics for All Ports

To view the statistics for all ports select, **Ports > Ports > Statistics**. The **Ports Statistics** page displays, which shows a table with the statistics for each port as

shown in Figure 22-39. For the definitions of the statistics shown in the table, refer to Table 22-16.

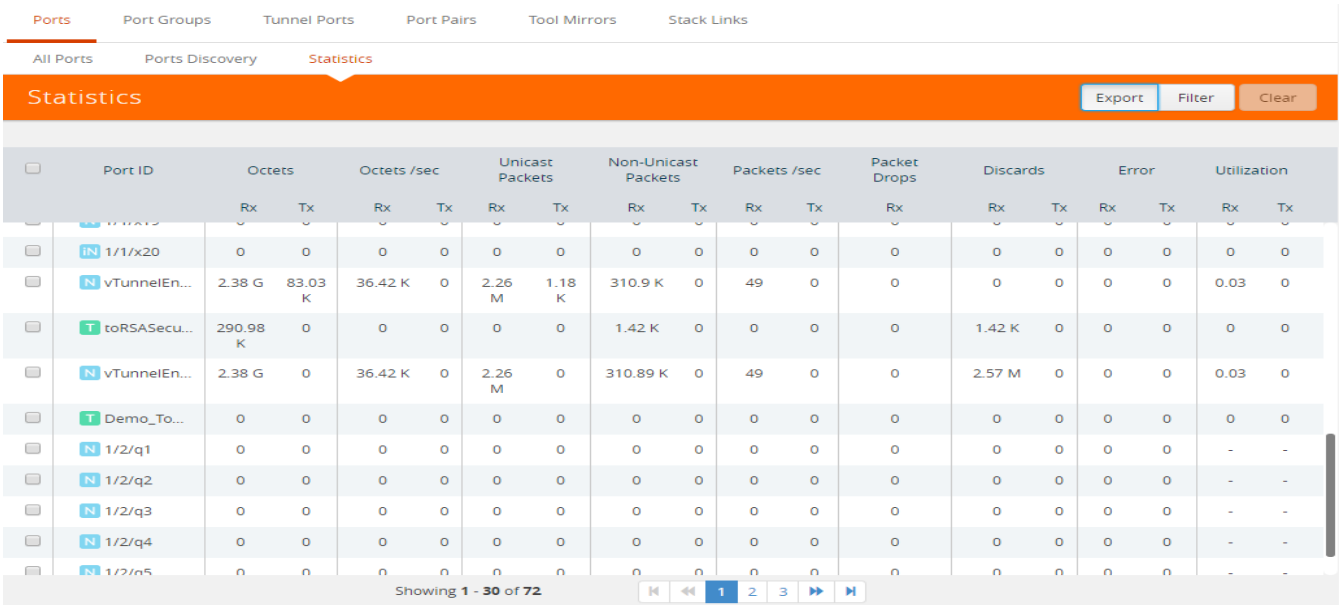

Figure 22-39: Port Statistics for All Ports

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

To filter the port statistics, click **Filter**. A Filter Quick View opens (refer to Figure 22-40 on page 466, where you can specify how to filter ports displayed on the Statistics page.

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

Criteria	Description
Box/Slot ID	Display only those ports that match the specified box and slot IDs.
Port Alias	Display port with the specified alias.
Port ID	Display ports with specified number in the port ID. For example, if you specify 3 the result will also display ports that include the number 3, 13, 23, 30, and so on.
Type	Display ports with the specified port type. Select one of the following: • Network • Tool • Inline Network • Inline Tool • GigaSMART • Hybrid • Stack

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

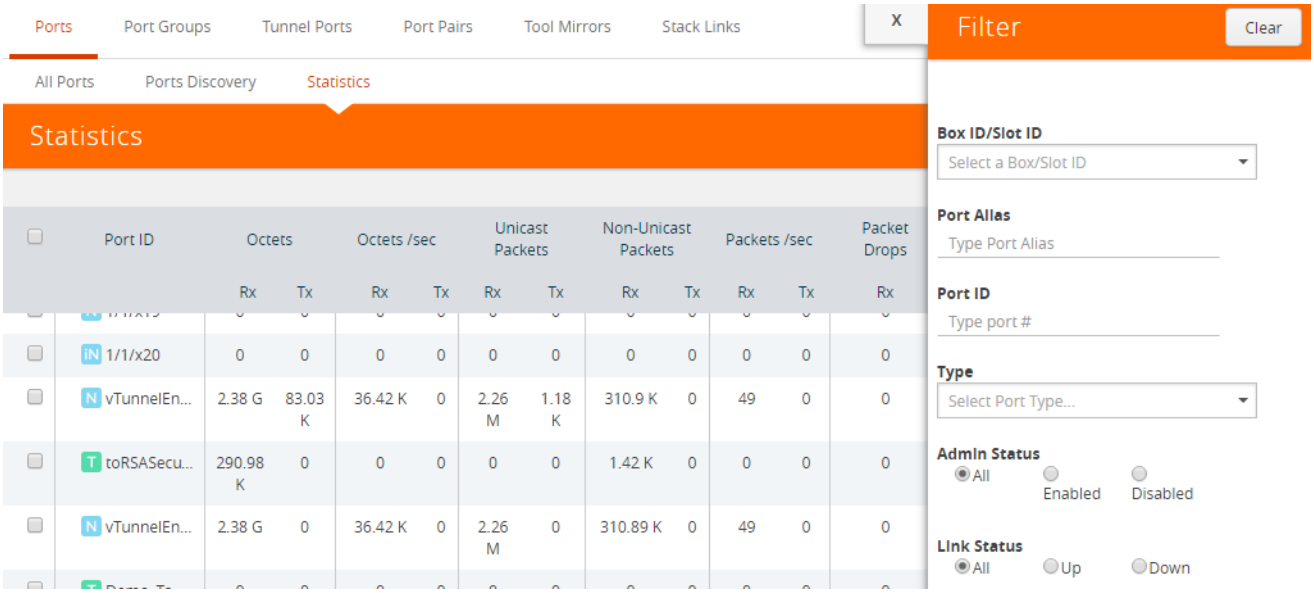

Figure 22-40: Port Statics Filter

How to Reset Traffic Counters

To reset the traffic counters, do the following:

1. Select a port on the All Ports page. The Quick View page opens for the port as shown in Figure 22-41.

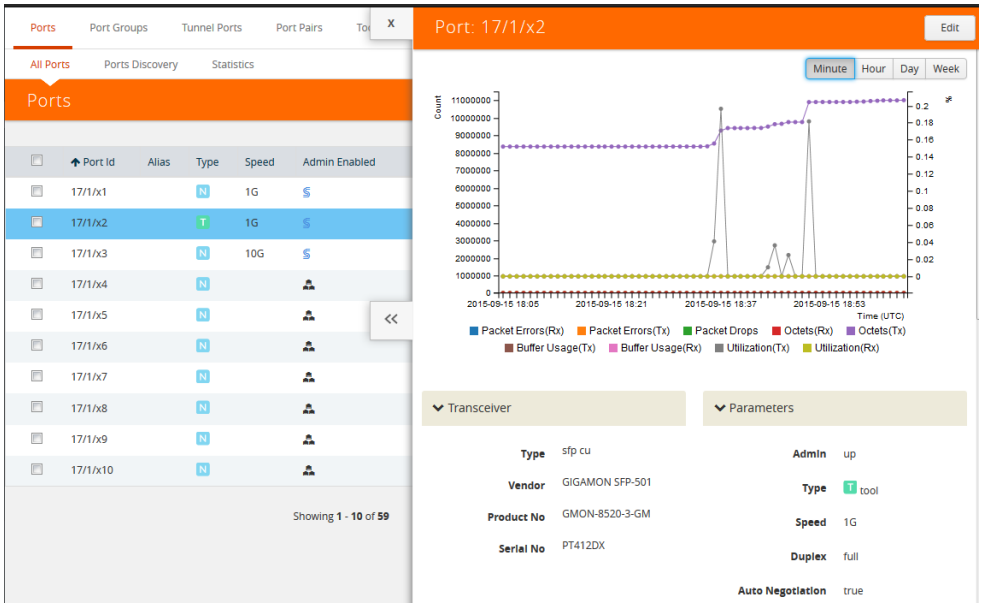

Figure 22-41: Port Quick View

2. On the Quick View page, click **Edit**.
3. On the ports page, click **Reset Traffic Counters**.

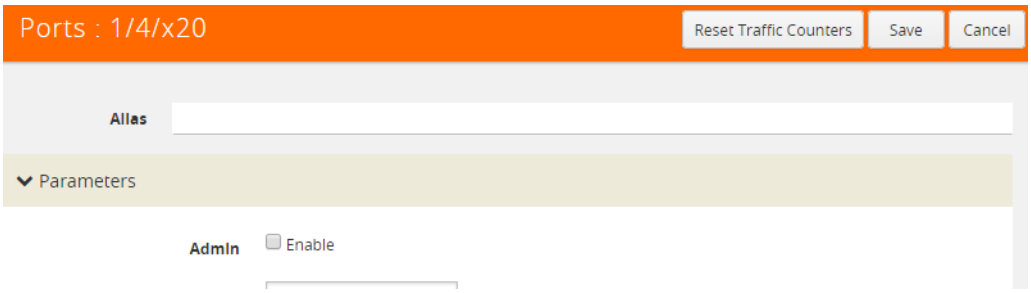

Figure 22-42: Reset Traffic Counters From Port Editor Page

Monitor Port Utilization

This section describes how to monitor port utilization and buffer thresholds on the GigaVUE H Series and GigaVUE TA Series nodes. Refer to the following sections for details:

Port Utilization Availability by Port Type

You can view port utilization for all network, tool, hybrid, and stack link ports on the GigaVUE H Series or GigaVUE TA Series node by selecting **Ports > All Ports** and looking at the Utilization (Tx/Rx) column in the table as shown in Figure 22-43.

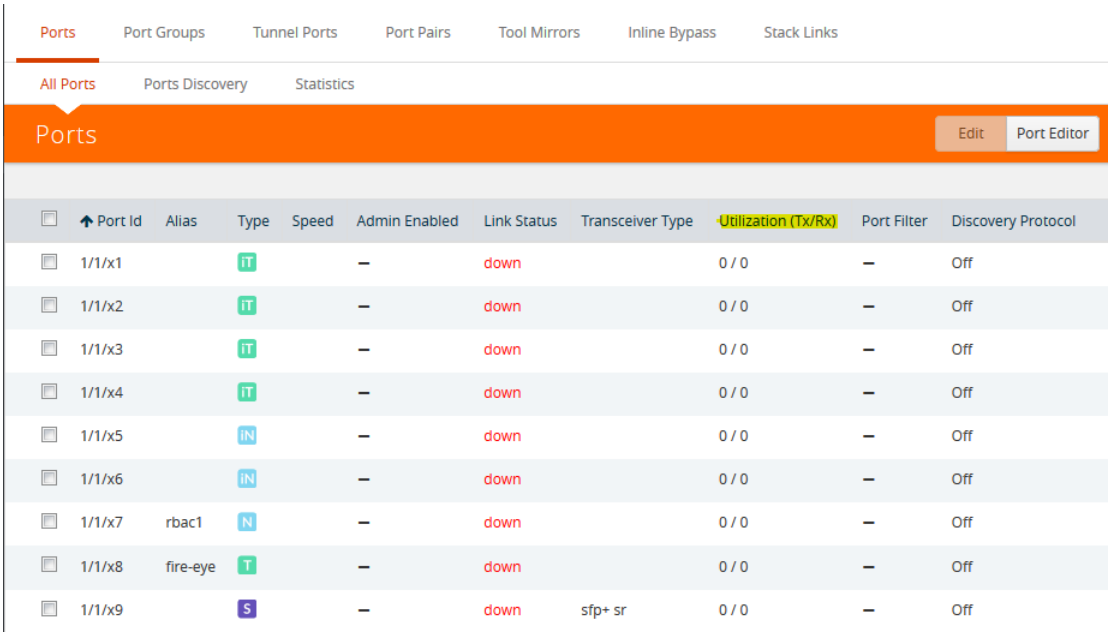

Figure 22-43: Port Utilization

The utilization is the utilization for all requested ports with the port number, port type, port speed, receive (Rx) utilization percentage (network and stack ports), transmit (Tx) utilization percentage (tool, hybrid, and stack ports), alarm threshold, and the last time the threshold was exceeded on either the transmit or receive channel.

Set Port Utilization Thresholds

To set the Alarms for port utilization, do the following:

1. Select **Ports > Ports > All Ports**.
2. Select Port ID of the port on which you want to set the utilization threshold.
3. Click **Edit** to open the port editor.

4. Under **Alarms**, in the **Utilization Threshold** field, enter the percentage at which the GigaVUE H Series node logs an alarm for the port. By default, the thresholds are 0, which means disabled.

   NOTE: Network ports always use an Rx threshold. Tool ports always use Tx. Stack ports and hybrid ports use both Rx and Tx, and the same threshold is used for each.

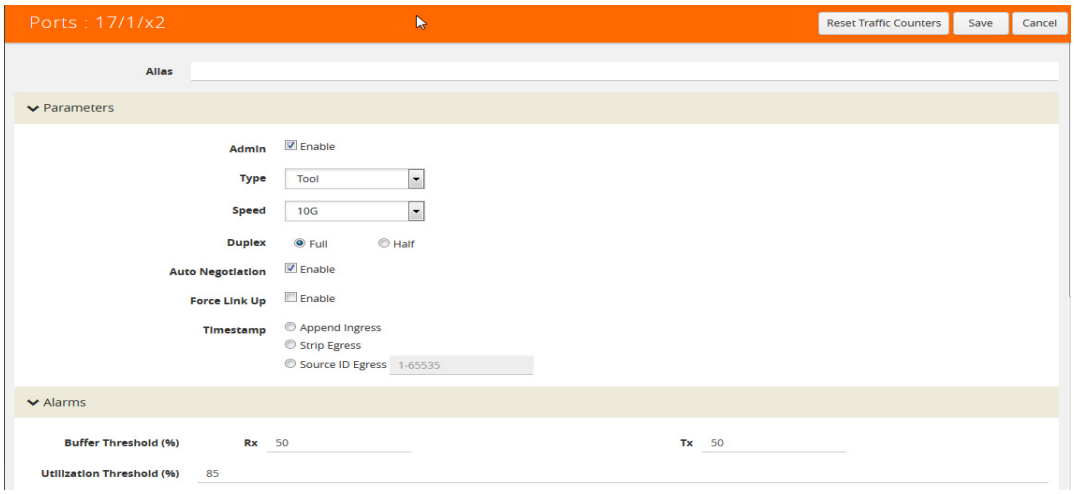

Utilization Alarm/SNMP Trap Generation

The GigaVUE H Series or GigaVUE TA Series node generates a utilization alarm each time the configured threshold is exceeded for more than six consecutive seconds. Once the percentage utilization falls below the configured threshold for at least six consecutive seconds, a second alarm is generated to indicate that utilization has returned to normal. Once utilization has returned to normal (six consecutive seconds below the threshold), a new utilization alarm can be generated once the measured rate again remains above the threshold for six consecutive seconds.

Utilization alarms are written to syslog and forwarded to all SNMP management stations configured as notification destinations. For SNMP traps to be generated, forwarded, and displayed correctly in your SNMP management station, all of the following must be true:

Configure Alarm Buffer Thresholds

Often network ports are utilized at rates below 50%. If several network ports are aggregated, there is a risk of oversubscribing the tool ports. Alarm buffer thresholds are used to monitor the congestion within the GigaVUE node caused by microbursts or by oversubscription of tool ports.

The buffer usage on any port remains at zero until the maximum line rate of the port is reached. When the usage crosses 100% either instantaneously, in the microburst case, or prolonged, in the oversubscription case, there is congestion.

The internal buffer on the GigaVUE node can absorb a certain number of packet bursts. During congestion, packets are buffered in the chassis and the buffer usage is

reported on the corresponding ports and in the corresponding direction: rx (ingress) and tx (egress).

Reporting the buffer usage provides a trend of how the microbursts are causing congestion, so more tool ports can be added before packets are dropped. Buffer usage is measured in intervals of 5 seconds. The peak buffer usage within a 5-second interval is reported.

When buffer usage is less than or equal to zero, there is no congestion, so no packets are dropped due to buffer unavailability.

When buffer usage is greater than zero, there is congestion. When buffer usage is greater than zero on any port in any direction, there is a chance that the packets (that caused the buffer usage to increase) are dropped due to unavailable buffers. However, it is unlikely to see packet drops due to buffer unavailability when the buffer usage on a port is less than 5%.

The buffer usage feature is supported on all ports and module types on the GigaVUE-HC3 and GigaVUE-HC2 (equipped with Control Card version 1 only).

Refer to *Set Alarm Buffer Thresholds* on page 470 for configuring buffer thresholds and for configuring a notification that can be sent when a threshold is exceeded.

Set Alarm Buffer Thresholds

Use the Alarms section of the Ports configuration page to set rx (ingress) and tx (egress) alarm buffer threshold on a port and utilization threshold. You can specify the alarm buffer threshold in the rx and tx directions on network and stack type ports and in the tx direction on tool type ports. By default, the threshold is set to 0, which disables the threshold.

When a buffer usage threshold has exceeded its configured percentage, a message is logged, and optionally, an SNMP trap is sent to all configured destinations.

The SNMP trap will be sent when a threshold is exceeded in any 5-second interval. Once the trap is sent, there is a 30 second hold-off time before the trap is sent again.

For information about how to set SNMP traps, refer to “Using SNMP” in the*GigaVUE-OS H-VUE Administration Guide*.

To set the alarm buffer threshold and the usage thresholds on a port do the following:

1. Select **Ports** > **Ports** > **All Ports**.
2. Select the Port on the Ports page.
3. In the Alarms section of the Port configuration page, enter the **Rx** and **Tx** values for the **Buffer Threshold** and the **High** and **Low Values** fir the **Utilization Threshold**. Figure 22-44 on page 471 shows an example.
4. Click **Save**.

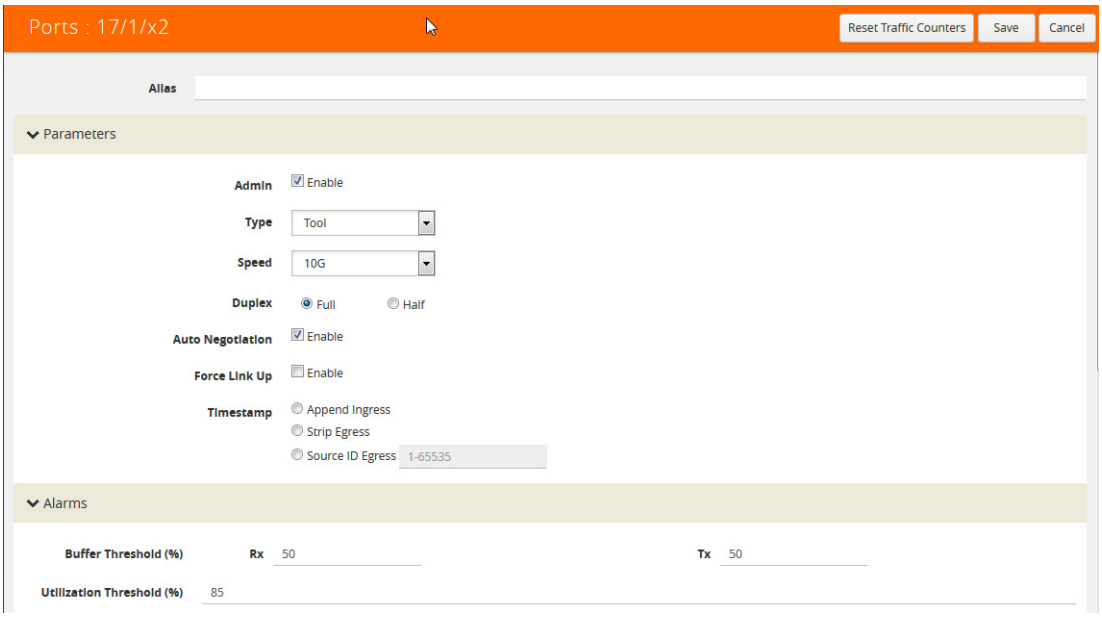

Figure 22-44: Alarm Buffer Thresholds Set

23 About Tunnels

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

Required License: Advanced Feature License on GigaVUE TA Series Nodes

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

About Circuit-ID Tunnels

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

Figure 23-1 illustrates the circuit flow mapping between two clusters using circuit-ID.

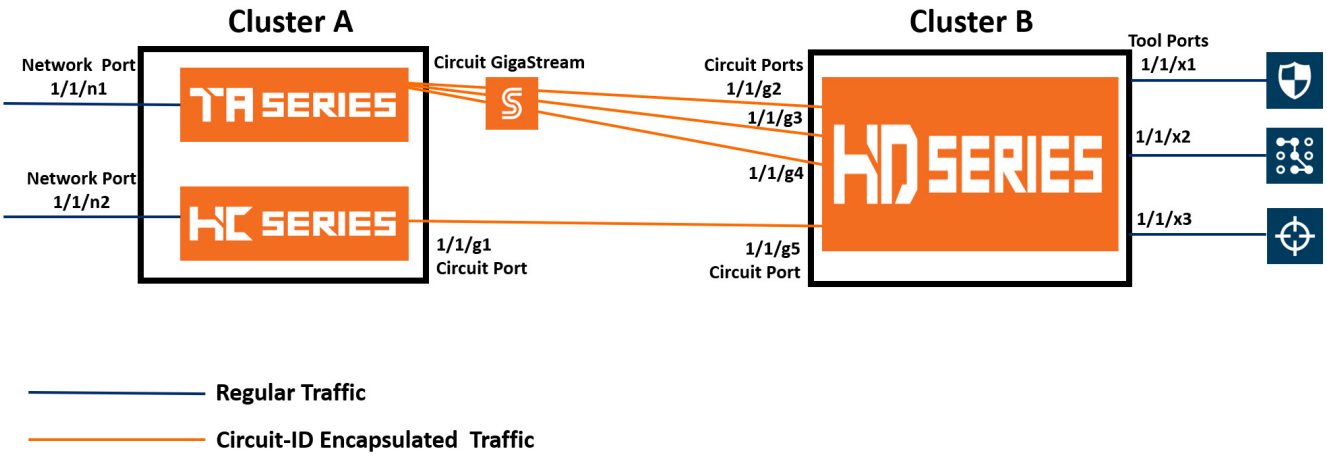

Figure 23-1: Circuit Flow Mapping

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

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

- *Circuit-ID Tunnels—Rules and Notes* on page 474
- *Circuit-ID Tunnel Encapsulation* on page 475
- *Circuit-ID Tunnel Decapsulation* on page 475

Circuit-ID Tunnels—Rules and Notes

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

- A maximum of 512 circuit-IDs are supported within a cluster for encapsulation and decapsulation.
- If a network port receives a double-tagged packet that is encapsulated with a circuit-ID, the five tuple hashing will not work only in the second cluster, that is the cluster in the decapsulation side over stack GigaStream or tool gigastream. Hence, traffic cannot be filtered using the IP/L4 parameters. After decapsulation, flow mapping filters the traffic based on circuit-ID.
- It is not supported for inline scenarios.

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

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

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

- A maximum of 512 circuit-ID tunnels can be created for decapsulation.
- Circuit-ID tunnel decapsulation is not supported on Pass All and Shared Collector maps.
- A circuit-ID must be paired with a circuit port, only in one circuit-ID tunnel.

Circuit-ID Tunnel Encapsulation

Before creating a Circuit-ID tunnel for encapsulation, refer to the *Circuit-ID Tunnels—Rules and Notes* on page 474.

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

S.No	Task	Refer to...
1.	Configure the required circuit ports and circuit GigaStream.	• *Configure Ports* on page 408 • *Configure Regular GigaStream* on page 434
2.	Configure a circuit-ID tunnel for encapsulation. Ensure that you select the mode as **Encap**.	*Create Tunnel* on page 481
3.	Configure a map to encapsulate the traffic and attach the circuit-ID tunnel to the map.	*Create a New Map* on page 519

Circuit-ID Tunnel Decapsulation

Before creating a Circuit-ID tunnel for encapsulation, refer to the *Circuit-ID Tunnels—Rules and Notes* on page 474.

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

S.No	Task	Refer to...
1.	Configure the required circuit ports and circuit GigaStream.	• *Configure Ports* on page 408 • *Configure Regular GigaStream* on page 434
2.	Configure a circuit-ID tunnel for decapsulation. Ensure that you select the mode as **Decap** and attach the circuit ports or circuit GigaStream that you configured in step 1 to the circuit-ID tunnel.	*Create Tunnel* on page 481

S.No	Task	Refer to...
3.	Configure a map to decapsulate the traffic and ensure that you specify the circuit-ID as a pass/ drop rule in the map.	*Create a New Map* on page 519

About Layer 2 Generic Routing Encapsulation (L2GRE) Tunnels

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

The encapsulated packet is sent out of the 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. The packet is ingressed at the circuit port of the GigaVUE device at the main office. The received packet's destination IP is checked against the IP configured for the circuit port. If they match, decapsulation is applied. The Ethernet + IP + GRE header is stripped and the remaining packet is sent to the tool port.

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

- *About L2GRE Tunnel Termination* on page 476
- *L2GRE Tunnel Termination—Rules and Notes* on page 477
- *Configure L2GRE Tunnel Termination* on page 478

About L2GRE Tunnel Termination

Figure 23-2 illustrates the L2GRE tunnel termination on a GigaVUE device located in a site that is remote to the device from where the traffic was routed across the cloud.

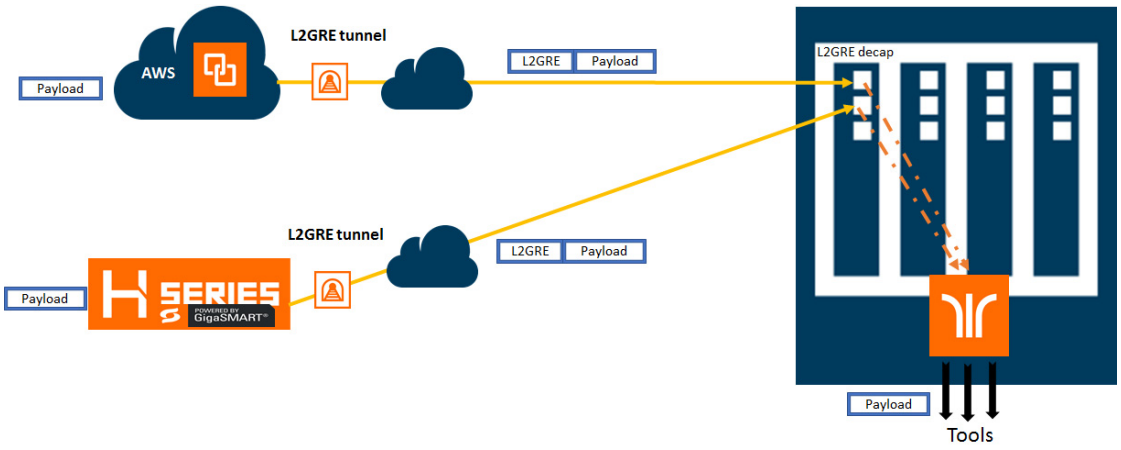

Figure 23-2: L2GRE Tunnel Termination

In this diagram, traffic is tapped on a GigaVUE H Series device at a remote site, and then it is tunneled through L2GRE encapsulation across the network before it reaches the GigaVUE H Series device at the main office site, which is connected to the actual tools. The traffic from a GigaVUE V Series appliance running on an AWS platform is tunneled through L2GRE encapsulation across the cloud, and the L2GRE tunneled traffic hits the GigaVUE H Series device at the main office site. The tunnel termination is executed on an ingress circuit port (IP interface). After tunnel termination the packet is presented to the flow mapping module to filter based on map rule parameters.

L2GRE Tunnel Termination—Rules and Notes

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

- L2GRE tunnel termination is supported only on GigaVUE-HC1, GigaVUE-HC2 CCv2, GigaVUE-HC3, GigaVUE-TA40, GigaVUE-TA100, and GigaVUE-TA200 devices.
- A maximum of 1500 L2GRE IDs are supported.
- Flow mapping that is configured on the circuit port used for L2GRE decapsulation will filter only the inner packet attributes along with L2GRE-ID. Any other non-tunneled packets that ingress on this circuit port will not be filtered or redirected to tool ports, even if it matches the rules configured on the map.
- IPv6 is not supported with L2GRE tunnels.
- L2GRE tunnel termination do not support reassembly of packets.
- L2GRE tunnel termination is supported only on encapsulated packets that are not tagged.
- Map-passall is not supported for the circuit port that decapsulates the L2GRE packet.
- Inner VLAN qualifier is not supported on the port in which the L2GRE tunnel termination is enabled.

- L2GRE ID qualifier is available as part of existing static templates. Following table provides details about the platforms for which the static templates are available:

Template	Platform	
	GigaVUE-HC2 (CCv2)/ GigaVUE-HC1/ GigaVUE-TA40	GigaVUE-HC3/ GigaVUE-TA100
IPv4	No	Yes
IPv6	Yes	Yes
IPv4+UDA	No	Yes
IPv4+MAC	Yes	Yes
UDA	Yes	Yes

Configure L2GRE Tunnel Termination

Before creating a L2GRE tunnel termination, refer to the *L2GRE Tunnel Termination—Rules and Notes* on page 477.

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

S.No	Task	Refer to...
1.	Configure the required circuit port.	*Configure Ports* on page 408
2.	Configure an IP interface and attach the circuit port that you created in task 1.	*Configure IP Interface* on page 418
3.	Configure a L2GRE tunnel for decapsulation and attach the IP interface that you created in task 2. Ensure that you select the Type as **L2GRE**.	*Create Tunnel* on page 481
4	Create a L2GRE group for a device and add all the L2GRE IDs that are specific to the device.	*Create VXLAN / L2GRE Group* on page 482
5.	Configure a map to decapsulate the traffic. Ensure that you add the IP interface that you created in task 2 as the source of the map and specify the required pass/drop rule in the map to filter the traffic based on inner packet attributes or L2GRE ID for the template configured.	*Create a New Map* on page 519

About Virtual Extensible LAN (VXLAN) Tunnels

VXLAN is a simple tunneling mechanism that allows overlaying a Layer 2 (L2) network over a Layer 3 (L3) underlay with the use of any IP routing protocol. It uses MAC Address-in-User Datagram Protocol (MAC-in-UDP) encapsulation. A remote device, such as the Gigamon cloud or a customer-specific device, encapsulates the filtered

traffic, adds an encapsulation header that consists of Layer 2 + IP + UDP + VXLAN headers. The encapsulated packet is sent out of the 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. The packet is ingressed at the circuit port configured in the GigaVUE-H Series or GigaVUE-TA Series device at the main office. After validation of the source port, destination port, and VXLAN Network Identifier (VNI) of the packet, the VXLAN tunnel header will be removed and the inner payload will be sent to the tools based on the map rules configured.

Refer to the following sections for details about the VXLAN tunnel termination:

- *About VXLAN Tunnel Termination* on page 479
- *VXLAN Tunnel Termination—Rules and Notes* on page 479
- *Configure VXLAN Tunnel Termination* on page 480

About VXLAN Tunnel Termination

Figure 23-3 illustrates the VXLAN tunnel termination on a GigaVUE device located in a site that is remote to the device from where the traffic was routed across the cloud.

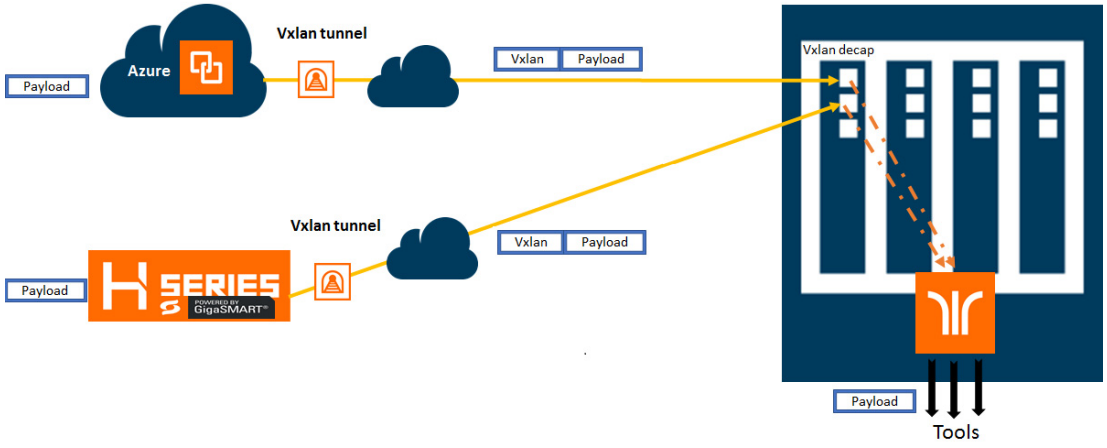

Figure 23-3: VXLAN Tunnel Termination

In this diagram, traffic is tapped on a GigaVUE H Series device at a remote site, and then it is tunneled through VXLAN encapsulation across the network before it reaches the GigaVUE H Series device at the main office site, which is connected to the actual tools. The traffic from a GigaVUE V Series appliance running on an AWS platform is tunneled through VXLAN encapsulation across the cloud, and the VXLAN tunneled traffic hits the GigaVUE H Series device at the main office site. The tunnel termination is executed on an ingress circuit port (IP interface). After tunnel termination the packet is presented to the flow mapping module to filter based on map rule parameters.

VXLAN Tunnel Termination—Rules and Notes

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

- VXLAN tunnel termination is supported only on GigaVUE-HC1, GigaVUE-HC2 CCv2, GigaVUE-HC3, GigaVUE-TA40, GigaVUE-TA100, and GigaVUE-TA200 devices.
- A maximum of 1500 VXLAN IDs are supported.
- Flow mapping that is configured on the circuit port used for VXLAN decapsulation will filter only the inner packet attributes along with VXLAN-ID. Any other non-tunneled packets that ingress on this circuit port will not be filtered or redirected to tool ports, even if it matches the rules configured on the map.
- IPv6 is not supported with VXLAN tunnels.
- VXLAN tunnel termination do not support reassembly of packets.
- VXLAN tunnel termination is supported only on encapsulated packets that are not tagged.
- Map-passall is not supported for the circuit port that decapsulates the VXLAN packet.
- When a circuit port is configured for VXLAN tunnel termination, you cannot use the port in any other regular map in which a network port is configured as the source port.
- Inner VLAN qualifier is not supported on the port in which the VXLAN tunnel termination is enabled.
- VXLAN ID qualifier is available as part of existing static templates. Following table provides details about the platforms for which the static templates are available:

Template	Platform	
	GigaVUE-HC2 (CCv2)/ GigaVUE-HC1/ GigaVUE-TA40	GigaVUE-HC3/ GigaVUE-TA100
IPv4	No	Yes
IPv6	Yes	Yes
IPv4+UDA	No	Yes
IPv4+MAC	Yes	Yes
UDA	Yes	Yes

Configure VXLAN Tunnel Termination

Before creating a VXLAN tunnel termination, refer to the *VXLAN Tunnel Termination—Rules and Notes* on page 479.

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

S.No	Task	Refer to...
1.	Configure the required circuit port.	*Configure Ports* on page 408

S.No	Task	Refer to...
2.	Configure an IP interface and attach the circuit port that you created in task 1.	*Configure IP Interface* on page 418
3.	Configure a VXLAN tunnel for decapsulation and attach the IP interface that you created in task 2. Ensure that you select the Type as **VXLAN**.	*Create Tunnel* on page 481
4	Create a VXLAN group for a device and add all the VXLAN IDs that are specific to the device.	*Create VXLAN / L2GRE Group* on page 482
5.	Configure a map to decapsulate the traffic. Ensure that you add the IP interface that you created in task 2 as the source of the map and specify the required pass/drop rule in the map to filter the traffic based on inner packet attributes or VXLAN ID for the template configured.	*Create a New Map* on page 519

Create Tunnel

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

1. From the device view, go to **Ports > Tunnels**.
2. Click **New**. The Tunnel configuration page appears.
3. In the **Alias** and **Comment** fields, enter the name and description of the tunnel.
4. From the **Type** drop-down list, select one of the following option:
   - **Circuit-ID—**to create Circuit-ID tunnels.
   - **VXLAN—**to create VXLAN tunnels.
   - **L2GRE—**to create L2GRE tunnels.
5. Select one of the following modes:
   - **Encap**—Select this mode to send the encapsulated traffic to the destination node that resides in another cluster.
   - **Decap**—Select this mode to decapsulate the traffic received from the source node that resides in another cluster.

   NOTE: The **Mode** field is available only when you select **Circuit-ID** in the **Type** drop-down list.
6. In the **Circuit-ID** field, enter the circuit ID used to encapsulate or decapsulate the traffic. You can enter multiple circuit-IDs when you create a circuit-ID tunnel for decapsulation.

   NOTE: This field is available only when you select **Circuit-ID** in the **Type** drop-down list.
7. From the **Attached entity** drop-down list, select the required entity based on the tunnel type you are creating:
   - If you are creating a Circuit-ID tunnel, this field is available only when you select the **Mode** as **Decap**. Ensure that you attach the required circuit ports or circuit GigaStreams.

- If you are creating a VXLAN tunnel or a L2GRE tunnel, ensure that you select the IP interface to which you have attached the circuit port.

8. Click **OK**.

The newly added tunnel appears in the Tunnels listing page.

Create VXLAN / L2GRE Group

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

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

The newly added group appears in the VXLAN / L2GRE Groups listing page.

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

View VXLAN / L2GRE ID Statistics

You can view the statistics such as the number of packets and bytes that were decapsulated using a VXLAN ID or a L2GRE ID for a device. To view the VXLAN / L2GRE ID Statistics page, go to **Ports > Tunnels > Statistics**. Figure 23-4 illustrates the statistics.

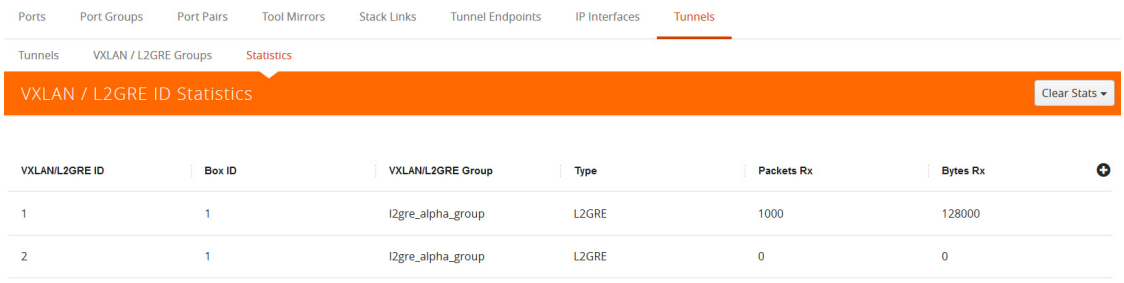

Figure 23-4: VXLAN / L2GRE ID Statistics

The following table describes the VXLAN or L2GRE ID statistics:

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

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

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

24 Maps

This chapter provides the following information about flow mapping:

About Flow Mapping

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

Flow Mapping Overview

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

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

traffic streams, applying Flow Mapping to tunneled traffic, specialized applications and even higher-layer protocols.

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

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

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

Get Started with Flow Mapping

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

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

NOTE: These tasks can also be performed with the GigaVUE-FM APIs. For more information about the APIs, refer to *GigaVUE-FM REST API Getting Started Guide* and the *GigaVUE-FM API Reference*.

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

- *Example: How to Create a Simple Map* on page 513
- *Example: How to Handle Overlaps when Sending VLANs and Subnets to Different Tools* on page 516

Check Status of Nodes and Ports

Before configuring maps, check the status of line cards, modules, and ports. To view the status of line cards and modules, select **Chassis** from Navigation pane and select Table view. For more information about the Chassis page, refer to “Chassis” in the*GigaVUE-OS H-VUE Administration Guide*. To view the status of ports, select **Ports > Ports > All Ports** to see a table with details about each port. For more information about ports, refer to *About Ports* on page 391 and *Managing Ports* on page 401.

Designated Port Types

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

- network
- tool
- stack
- inline-network
- inline-tool
- hybrid

NOTE: Not all port types are supported on all platforms. Inline-network and inline-tool port types are only supported on GigaVUE HC Series nodes.

About Shared Collectors

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

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

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

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

In cases, where multiple network ports are sharing the multiple maps, packets that do no fit any of the maps can be sent to the shared collector.

No Map Statistics for Shared-Collector Only

A shared collector is intended to be used with other maps. For example, use a shared collector with a regular map containing at least one rule. If there is a shared collector but no other map, there will be no map statistics.

Shared Collector Configuration

A shared-collector is a special type of map configured with only a set of **Source** ports shared-collector ports or GigaStream. Rules, priority settings, GigaSMART operations and destination ports are not allowed in shared-collector maps. In H-VUE, the collector ports can be selected from a list of tool or hybrid ports.

To create a shared-collector map, do the following:

1. Select **Maps > Maps > Maps**.
2. Click **New**.
3. Type an alias in the Alias field to identify this map. For example, shared_collector.
4. For **Type**, select **Regular**.
5. For **Subtype**, select **Collector**.
6. Click in the **Source** field and select a network port.
7. Click in the **Destination** field and select the tool or hybrid ports that will be the shared-collector ports for the map. The map should look similar to the map shown in the following figure.

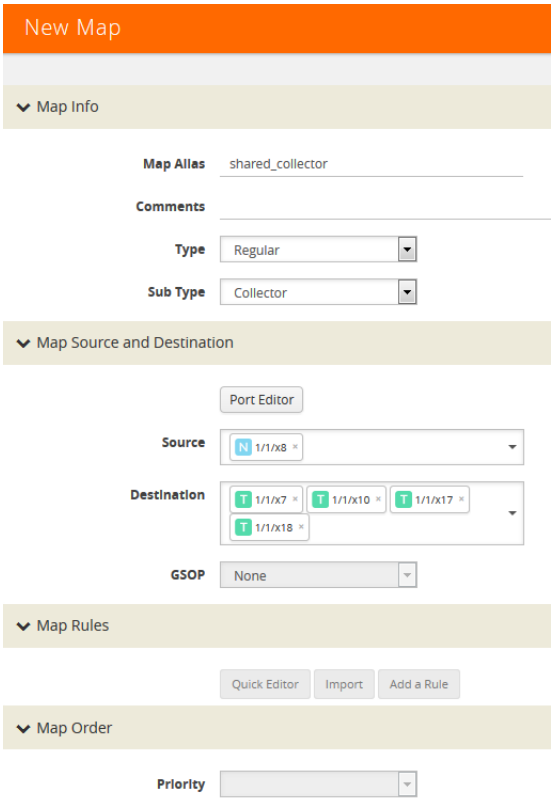

NOTE: Shared-collector maps do not include any rules, priority settings, or GigaSMART operations. These are grayed out in the UI when Collector is select for the map's subtype.

8. Click **Save**.

In Figure 24-1, the Maps page shows the shared-collector map shared_collector with the standard components. The **Source** ports match those used by a set of normal flow maps. The **Destination** ports are the collector ports are where you want to send any packets not matching the normal flow maps.

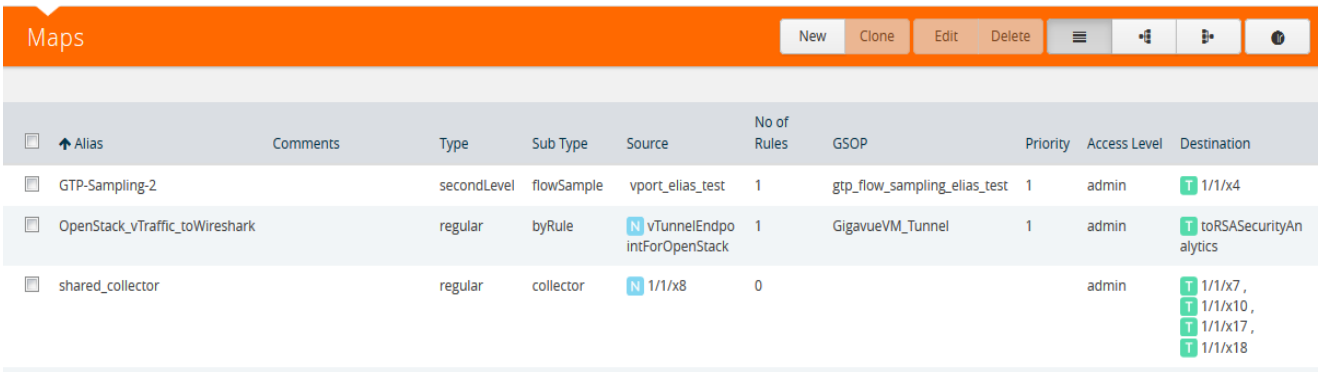

Figure 24-1: Shared Collector Map

About Map-passall Maps

Map-passall maps provide a way to specify a destination for packets without any filtering on a set of network ports. As indicated by the name, all traffic is passed through. A map-passall may share the network ports with a map which filters using map rules.

The same logic as set for Shared Collector can be set for map-passall. That is, that Map-passall can be set to a GigaStream alias or multiple GigaStream aliases or a single tool port or multiple tool ports. To set the **Destination,** use the same map range configuration.

Map-Passall and Shared-Collector Only

If a map-passall and a shared-collector both use the same network source port and there are no other maps, such as a regular map containing at least one rule, all traffic will be passed to the shared-collector.

Map-passalls Configuration

A map-passall map is a special type of map configured with only a set of **Source** ports and **Destination ports** or **GigaStream**. Map rules and GigaSMART operations are not allowed in passall maps.

Map-passalls in H-VUE

The web-based H-VUE interface for GigaVUE-OS nodes provides an all traffic **Pass All** subtype selection for regular maps that performs the same function as a **map-passall** in the CLI. Making this selection in H-VUE turns the map into a map-passall, delivering all traffic from the selected network ports to another tool port or GigaStream on any line card in the same node, irrespective of the other packet distribution. Although the names are different in H-VUE and the CLI, the two features are identical.

Define Map Source Port Lists

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

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

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

- map2 **Source** ports 1/1/x5, and 1/1/x6 can also be configured
- map3 **Source** ports 1/1/x3, 1/1/x4, 1/1/x5 cannot be configured because ports x4 and x5 are in both maps (they overlap)

Share Network Ports Between Maps

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

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

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

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

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

Share Tool Ports Between Maps

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

Map Priority

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

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

Maps sharing the same source port list are grouped together for the purpose of prioritizing their rules. Traffic is subjected to the rules of the highest priority map first and then the rules of the next highest priority map and so on. Within a map, drop rules

are applied first and then pass rules, in other words, drop rules always have higher priority than pass rules. Currently when a map's source port list is defined the map is grouped/prioritized with other maps sharing the same source port list. Newly configured maps are added as the lowest priority map within the group when initially configured unless changed by the user.

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

Adjust Map Priority in GigaVUE-HVUE

Before you get started adjusting map priority, start by reviewing the current map priorities in place by opening the Maps page and viewing the priority of the maps in the Priority field. For example, Figure 24-2 shows three maps MyMap1, MyMap2, and MyMap3 with the same source port 1/1/x8. The Priority column in the table shows the current priority of each map.

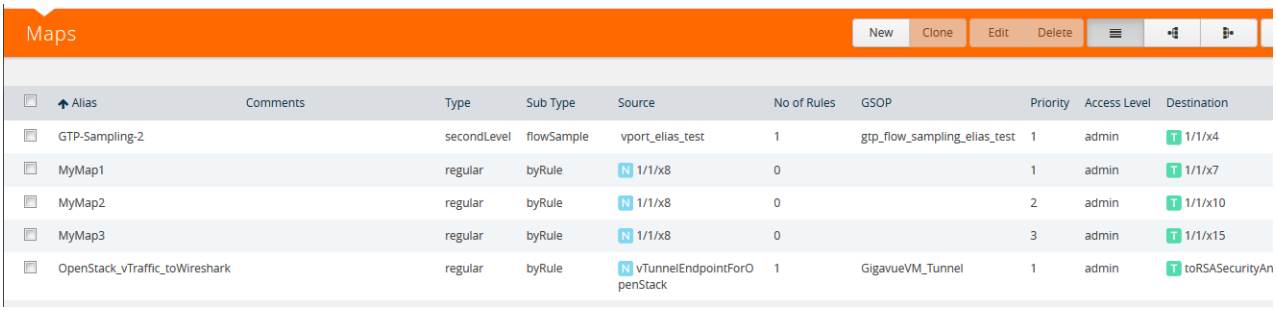

Alias	Comments	Type	Sub Type	Source	No of Rules	GSOP	Priority	Access Level	Destination
GTP-Sampling-2		secondLevel	flowSample	vport_elias_test	1	gtp_flow_sampling_elias_test	1	admin	T 1/1/x4
MyMap1		regular	byRule	N 1/1/x8	0		1	admin	T 1/1/x7
MyMap2		regular	byRule	N 1/1/x8	0		2	admin	T 1/1/x10
MyMap3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

Figure 24-2: Map Priorities

Then, once you have reviewed the existing hierarchy of map priorities, you can fine-tune the priority of maps by using the **Priority** field in the map to select one of the following:

- Highest (top)—set the map to the highest priority
- After map - <map-alias> — set the map priority after the map with the specified alias.
- Lowest (bottom)—set the map priority to the lowest priority

Flow Map Syntax and Construction

This section provides information about map types, map rules, and working with map passalls.

Map Types

Map configuration consists of several parameters, including the following:

- **Source**—Specifies the source ports for the map.
- **Destination**—Specifies the destination ports for the map.
- **Type**—Specifies the type of map.
- **Subtype**—Specifies map subtype.
- **GSOP**—Specifies a GigaSMART operation.

Starting in software version 4.4, there are four types of maps, with the map type being determined by the **Source** and **Destination** as well as the map rules as follows:

- **Regular**—Specifies a regular map type, with the **Source** parameter specifying network or hybrid ports, or single inline-network or single inline-tool ports (for out-of-band maps) and the **Destination** parameter specifying tool, hybrid, or GigaStream ports.

  When the subtype for a **Regular** map type is **By Rule**, a **Pass Traffic** option for **No Rule Matching** is available. Selecting Pass Traffic specifies what to do with traffic that does not match any rule in a map that only has drop rules. This argument changes the default behavior of drop to pass in a drop-only map. If you do not use this argument and there are only drop rules in a map, the default behavior is that all traffic not matching the rules will be dropped, or, if a shared collector is configured, traffic will be sent to the shared collector. However, if you use the Blacklisting option and there are only drop rules in a map, traffic will be passed rather than dropped.
- **inline**—Specifies an inline map type, with the **Source** parameter specifying inline-network pairs or inline-network-groups and the **Destination** parameter specifying inline-tool pairs, inline-tool-group, inline-serial, or bypass.
- **First Level**—Specifies a first level map type, with the **Source** parameter specifying network or hybrid ports and the **Destination** parameter specifying virtual ports, used with GigaSMART operations.

  When a First Level map type is selected, a **Traffic Type** option is available. This option is for GTP applications to pass GTP control traffic (GTP-c) to all GigaSMART engines in a GTP engine group. To enable GTP-c, select **Control**.

Starting in software version 5.3, a fifth map type is added as follows:

- **flexinline**—Specifies a flexible inline map, with the **from** parameter specifying inline-network and the **a-to-b** or **b-to-a** parameters specifying an ordered list of inline-tool or inline-tool-group.

Define Map Source Port Lists for First Level Maps

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

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

For example:

- map1 has Source ports 1/1x1,1/1x2 already configured.
- map 2 with Source ports 1/1x1,1/1x2 can also be configured.
- map3 with Source port 1/1/x3 can be configured. If you try to add source ports 1/1x1,1/1x2 as part of map3 then it is not possible because, 1/1x1,1/1x2 are already part of map 1.

NOTE: This is also applicable for regular maps.

- **Second Level**—Specifies a second level map type, with the **Source** parameter specifying virtual ports, used with GigaSMART operations, and the **Destination** parameter specifying tool, hybrid, or GigaStream ports.

NOTE: There is also a **Template** map type for creating map templates.

Map types are described in Table 24-1.

Table 24-1: Matrix of Map Types

Type	Subtype	Map Rule	GSOP	Source	Destination
Regular	By Rule, or Pass All, or Collector	Rule	yes	network ports, or hybrid ports, or single inline-network ports or single inline-tool ports (for out-of-band maps)	tool ports, or hybrid ports, or GigaStream
Inline	By Rule, or Pass All, or Collector	Rule	no	inline-network pairs, or inline-network-groups	inline-tool pairs, or inline-tool-groups, or bypass, or inline-serial tools
First Level	By Rule	Rule	no	network ports, or hybrid ports	virtual ports only (collector is not allowed)
Second Level	By Rule, or Collector	gsrule	yes	virtual ports	tool ports, or hybrid ports, or GigaStream, or port gro
	Flow Filter	flowrule	yes		
	Flow Sample	flowsample	yes		
	Flow Sample Overlap	flowsample	yes		
	flowSample-sip	flowsample	yes		
	Flow Whitelist	whitelist	yes		
	Flow Whitelist Overlap		yes		

Table 24-1: Matrix of Map Types

Type	Subtype	Map Rule	GSOP	Source	Destination
flexinline	inline-network	ordered list of inline-tool or inline-tool-group in a-to-b or b-to-a direction or both	rule	byRule	no
			-	collector	no

Map Subtypes

The map subtype describe in Table 24-1 on page 493 is optional. It specifies the following:

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

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

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

Map Type and Subtype Modification

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

Backwards Compatibility

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

Minimum Requirements for Map Creation

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

Map Rules

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

Other Types of Map Rules for GigaSMART Operations

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

- Adaptive Packet Filtering (APF) and Adaptive Session Filtering (ASF). For details, refer to *GigaSMART Adaptive Packet Filtering (APF)* on page 1003 and *GigaSMART Application Session Filtering (ASF) and Buffer ASF* on page 1054
- GTP correlation. For details, refer to *GigaSMART GTP Correlation* on page 885.
- GTP whitelisting and GTP flow sampling. For details, refer to *GigaSMART GTP Whitelisting and GTP Flow Sampling* on page 913.

IPv4/IPv6 and Map Rules

GigaVUE-OS provides a variety of criteria for pass/drop rules specific to IPv6 traffic, including:

IPv6 Entity	Rule Condition
IPv6 Source/Destination Addresses	**IPv6 Source/IPv6Destination**
IPv6 Flow Labels	**IPv6 Flow Label**
IPv6 Traffic	**IP Version and Version it set to v6**

In addition to the explicit IPv6 filters listed in the table, you can use the **IP Version** condition to change how some of the other attributes are interpreted.

When **IP Version** is used by itself in a map rule, it returns all traffic matching the specified IP version, **4** or **6**. However, when **IP Version** is set to **6**, several of the other arguments are interpreted differently when used in the same rule, as follows:

Condition	IP Version set to 4 (or not specified)	IP Version set to 6
Port Destination/ Port Source	Matches all IPv4 traffic on the specified port number.	Matches all IPv6 traffic on the specified port number.
IPv4 Protocol	When used with the **<1-byte-hex>** argument, matches against the protocol field in the standard IPv4 header.	When used with the **<1-byte-hex>** argument, matches against the Next Header field in the standard IPv6 header.
	NOTE: These fields perform essentially the same service in both versions, specifying what the next layer of protocol is. However, they have different names and are found at different locations in the header. Refer to *Protocol Map Rules and IPv6* on page 496 for a list of the useful values for the <1-byte-hex> field.	
IPv4 TTL	Matches against the standard TTL (time-to-live) field in the IPv4 header.	Matches against the standard Hop Limit field in the IPv6 header.
	NOTE: These fields perform essentially the same service in both versions, specifying how long a datagram can exist.	

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

Protocol Map Rules and IPv6

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

To address this, the GigaVUE H Series node provides the <1-byte-hex> option to match against the standard hex values for these protocols in the Next Header field. The standard 1-byte-hex values for both IPv4 and IPv6 are set by selecting the option from

the **Value** list or specifying a decimal value when **IPv4 Protocol** is selected. The following are options or values that you can select from the **Value** list.

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

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

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

Set Map Rules for TCP Control Bits

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

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

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

Examples

The map rule shown in Figure 24-3 matches packets with only the SYN bit set:

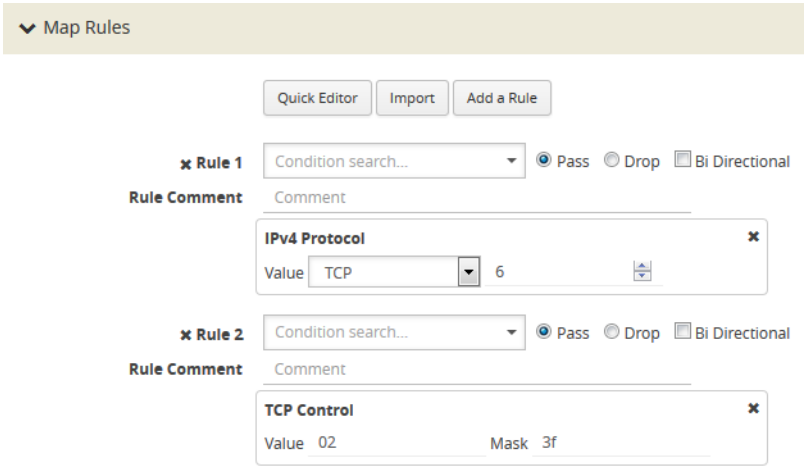

Figure 24-3: Map Rule with SYN Bit Set

Many packets will have some combination of these bits set rather than just one. So, for example, the map rule in Figure 24-4 matches all packets with both the ACK and SYN bits set.

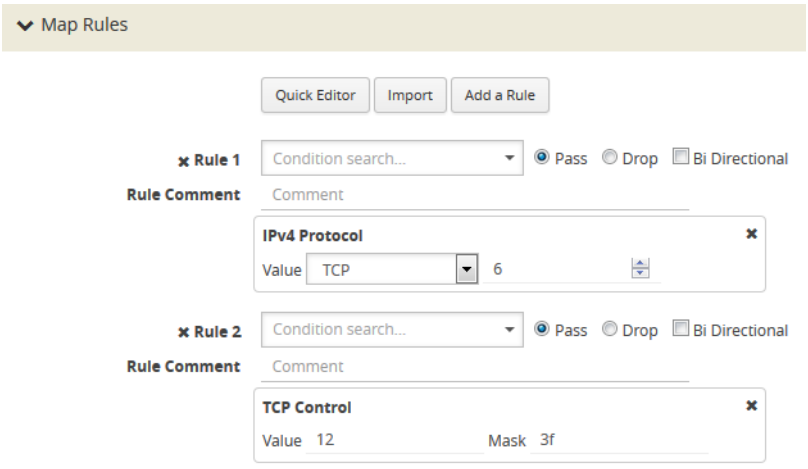

Figure 24-4: Map Rule Matching All Packets with Both ACK and SYN Bits set

How to Use Bit Count Netmasks

The following table summarizes the bit count netmask value for standard dotted-quad IPv4 netmasks. You can enter IP netmasks in the bit count format by using the **/nn** argument.

Bit count netmasks are easier to visualize for IPv6 addresses, specifying which portion of the total 128 bits in the address correspond to the network address. So, for example, a netmask of /64 indicates that the first 64 bits of the address are the network address

and that the remaining 64 bits are the host address. This corresponds to the following hexadecimal netmask:

ffff:ffff:ffff:ffff:0000:0000:0000
ffff:ffff:ffff:ffff:0000:0000:0000

Standard Netmask	Bit Count Netmask
255.255.255.255	/32
255.255.255.254	/31
255.255.255.252	/30
255.255.255.248	/29
255.255.255.240	/28
255.255.255.224	/27
255.255.255.192	/26
255.255.255.128	/25
255.255.255.0	/24
255.255.254.0	/23
255.255.252.0	/22
255.255.248.0	/21
255.255.240.0	/20
255.255.226.0	/19
255.255.192.0	/18
255.255.128.0	/17
255.255.0.0	/16
255.254.0.0	/15
255.252.0.0	/14
255.248.0.0	/13
255.240.0.0	/12
255.226.0.0	/11
255.192.0.0	/10
255.128.0.0	/9
256.0.0.0	/8
254.0.0.0	/7
252.0.0.0	/6
248.0.0.0	/5
240.0.0.0	/4
226.0.0.0	/3

Standard Netmask	Bit Count Netmask
192.0.0.0	/2
128.0.0.0	/1
0.0.0.0	/0

How to Combine Rules and Rule Logic

When working with maps, you can easily combine multiple criteria into a single rule. GigaVUE-OS processes rules as follows:

- Within a single rule, criteria are joined with a logical **AND**. A packet must match each of the specified criteria to satisfy the rule.
- Within a map, rules are joined with a logical **OR**. A packet must match at least ONE of the rules to be passed or dropped.

NOTE: When used in a map rule with multiple criteria, the **ipver** argument changes the interpretation of some map rule arguments. Refer to *IPv4/IPv6 and Map Rules* on page 496 for details.

Examples of Map Rule Logic

For example, the rules shown in the following table are both set up with criteria for **vlan 100** and **portsrc 23**.

- The first example combines the two criteria into a single rule. This joins the criteria with a logical **AND**.
- The second example creates two separate rules – one for each of the criteria. This joins the criteria with a logical **OR**.

How to Mix Pass and Drop Rules

GigaVUE-OS lets you mix pass and drop rules on a single port. Mixing pass and drop rules can be useful in a variety of situations. Figure 24-5 shows a pass rule set up to include all traffic matching a particular source port range combined with a drop rule configured to exclude ICMP traffic.

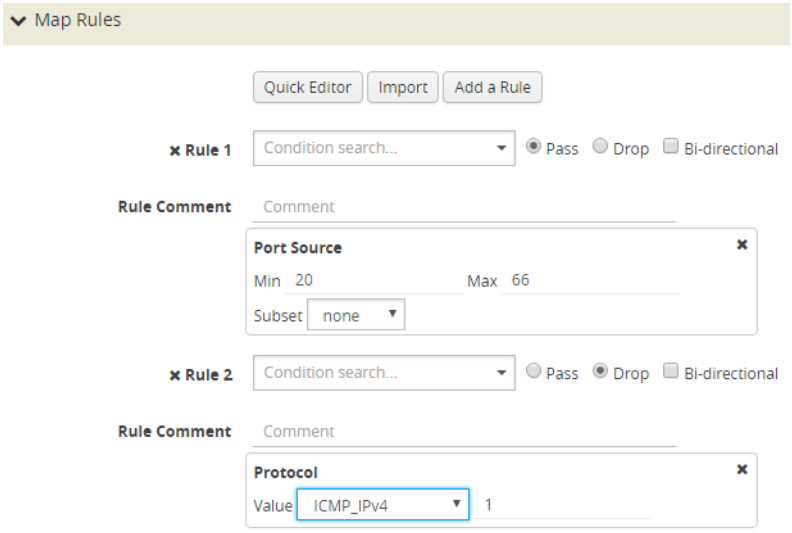

Figure 24-5: Pass and Drop Rules in a Map

Drop Rules Have Precedence!

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

Work with User-Defined Pattern Match Rules

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

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

User-defined pattern matches consist of the following components:

Step	Description
Pattern	Use the **UDA1 Value** and **UDA2 Value** fields for map rule to set up the actual bit patterns you want to search for. Refer to *User-Defined Pattern Match Examples* for details.
Mask	Use the **UDA1 Mask** and **UDA2 Mask** fields for map rules to specify which bits in the pattern must match to satisfy the map rule.
Offset	Use the **UDA1 Offset** and **UDA2 Offset** fields for map rules to specify where in the packet bits must match. **NOTE:** The GigaVUE H Series node accepts a maximum of two offsets per line card or GigaVUE-HB1 node. When both of the available offsets for a line card are in use with existing map rules, you will not be able to add a new rule with a different value for **UDA*x* Offset** until at least one of the **UDA*x* Offset** is freed up from all existing map rules.

User-Defined Pattern Match Syntax

The user-defined pattern match syntax is as follows:

- Both the **UDA*x* Value** and **UDA*x* Mask** arguments are specified as 16-byte hexadecimal sequences. Specify the pattern in four 4-byte segments separated by hyphens. For example:

  0x01234567-89abcdef-01234567-89abcdef
- Masks specify which bits in the pattern must match. The mask lets you set certain bits in the pattern as wild cards – any values in the masked bit positions will be accepted.
  - Bits masked with binary 1s must match the specified pattern.
  - Bits masked with binary 0s are ignored.
- You can set up the two global offsets allowed per line card or GigaVUE-HB1 node at 4-byte boundaries beginning at frame offset 2 and ending at offset 110. The resulting data range for pattern matches is from byte 3 through byte 126.
  - Multiple offsets must be set either equal to one another, or set beyond the boundaries of each other. For example, if **UDA1 Offset** starts at byte 2, the **UDA2 Offset** can only start either at byte 2 or at any point beginning with byte 18 (which would be the next 4-byte boundary after the 16-byte pattern used at **UDA1 Offset**).
  - Offsets are always frame-relative, not data-relative.
  - In many cases, you will be looking for patterns that do not start exactly on a four-byte boundary. To search in these position, you would set an offset at the nearest four-byte boundary and adjust the pattern and mask accordingly.

User-Defined Pattern Match Rules

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

- Offsets are specified in decimal; patterns and masks are specified in hexadecimal.
- All hexadecimal values must be fully defined, including leading zeroes. For example, to specify 0xff as a 16-byte value, you must enter 00000000-00000000-00000000-000000ff.
- User-defined pattern-match criteria are not allowed in tool port-filters.
- You can use user-defined pattern matches as either standalone map rules or in tandem with the other available predefined criteria for map rules (for example, port numbers, IP addresses, VLAN IDs, and so on).
- You can use up to two separate user-defined pattern matches in a single map rule. When two user-defined pattern matches appear in the same map rule, they are joined with a logical AND. However, the two patterns cannot use the same offset.
- User-defined pattern matches are combined in map rules using the same logic described in *How to Combine Rules and Rule Logic* on page 500.
- Avoid using user-defined pattern matches to set map rules for elements that are available as predefined criteria (for example, IP addresses, MAC addresses, and so on).

- GigaVUE H Series nodes accept a maximum of two offsets per line card or GigaVUE-HB1 node. When both of the available offsets for a line card are in use with existing map rules, you will not be able to add a new rule with a different value for **UDA*x* Offset** until at least one of the **UDA*x* Offsets** is freed up from all existing map rules.

User-Defined Pattern Match Examples

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

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

- Send all traffic to and from the 10.218.0.0 IP address range inside the GTP tunnel to tool port 1/5/x4.
- Send all traffic to and from the 10.228.0.0 IP address range inside the GTP tunnel to tool port 1/5/x9.

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

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

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

Description	UI Step
Map #1 – GTP_Map218	
Our first map will send traffic to and from the 10.218.0.0 IP address range inside the GTP tunnel to tool port 1/5/x4.	
Create a map with the alias **GTP_Map218**.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Enter GTP-MAP218 in the Alias field.
Specifies the map type and subtype.	**4.** Select Regular for Type. **5.** Select By Rule for Subtype.
Specify that this map will match packets arriving on network port 1/5/x1.	**6.** Select 1/5/x1 for Source.
Specify that packets matching this map will be sent to tool port 1/5/x4.	**7.** Select 1/5/x4 for Destination.
Next, add the map rules for our first address range – 10.218.0.0. This IP address translates to 0ada in hex. The first rule matches the 10.218.0.0 address at the source address offset of 62 in the GTP tunnel.	**8.** Click Add a Rule to add the first rule. **9.** Select Pass. **10.** Select UDA1 and set the values: - Value: 0ada0000-00000000-00000000-00000000 - Mask: ffff0000-00000000-00000000-00000000 - Offset: 62

Description	UI Step
The second rule matches the same address range (10.218.0.0) but at the destination address offset of 66 in the GTP tunnel. Notice that we have still specified the offset as 62 and have simply masked out to the correct location of the destination address. This way, we have still only used one of the two possible offsets in place for the GigaVUE H Series node at any one time.	**11.** Click Add a Rule to add the second rule. **12.** Select UDA1 and set the values: - Value: 00000000-0ada0000-00000000-00000000 - Mask: 00000000-ffff0000-00000000-00000000 - Offset: 62
Save the map.	**13.** Click Save.

Map #2 – GTP_Map228

Our second map will send traffic to and from the 10.228.0.0 IP address range inside the GTP tunnel to tool port 1/5/x9.

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

How to Handle Q-in-Q Packets in Maps

In software versions prior to 4.7, for traffic that matched the map pass rule shown in Figure 24-6, only Q-in-Q packets of TPID ethertype 0x8100 were passed and all tagged packets with a TPID ethertype other than 0x8100, such as 0x88A8 and 0x9100, were dropped:

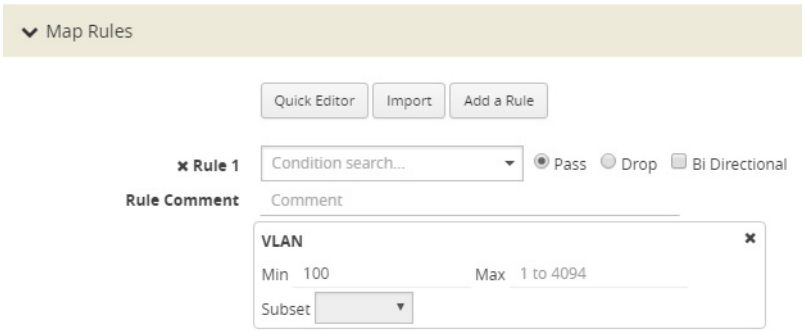

Figure 24-6: VLAN Pass Rule

Conversely, for traffic that matched the map drop rule shown in Figure 24-7, only Q-in-Q packets of TPID ethertype 0x8100 were dropped:

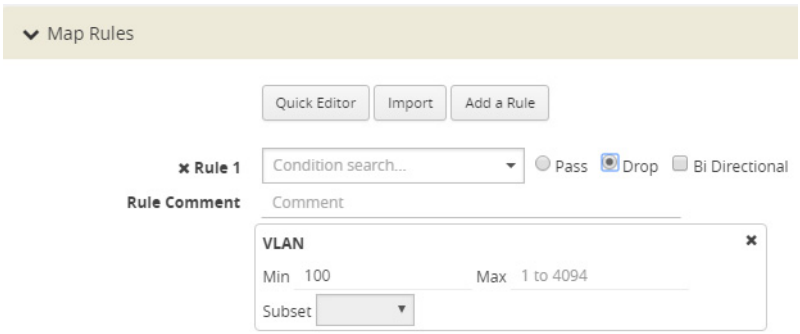

Figure 24-7: VLAN Drop Rule

Starting in software version 4.7, the rules shown in Figure 24-6 and Figure 24-7 will pass (or drop) TPID ethertypes 0x8100, 0x88A8, and 0x9100.

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

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

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

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

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

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

For handling of priority tagged packets, refer to *Priority Tagged Packets* on page 508.

For filtering of Q-in-Q packets on inner VLAN tag, refer to *Flow Mapping on Inner VLAN Tags* on page 508.

Upgrade Note

If you had previously defined TPID **Ether Types** in the earlier software version, during an upgrade to 4.7, they will be converted (or removed) as follows:

- In earlier software version, rules with both a TPID ether type and VLAN such as the rules shown in Figure 24-8 will be converted in 4.7 to the rules shown in Figure 24-9 on page 507 (with the TPID Ether Types removed from the rule).

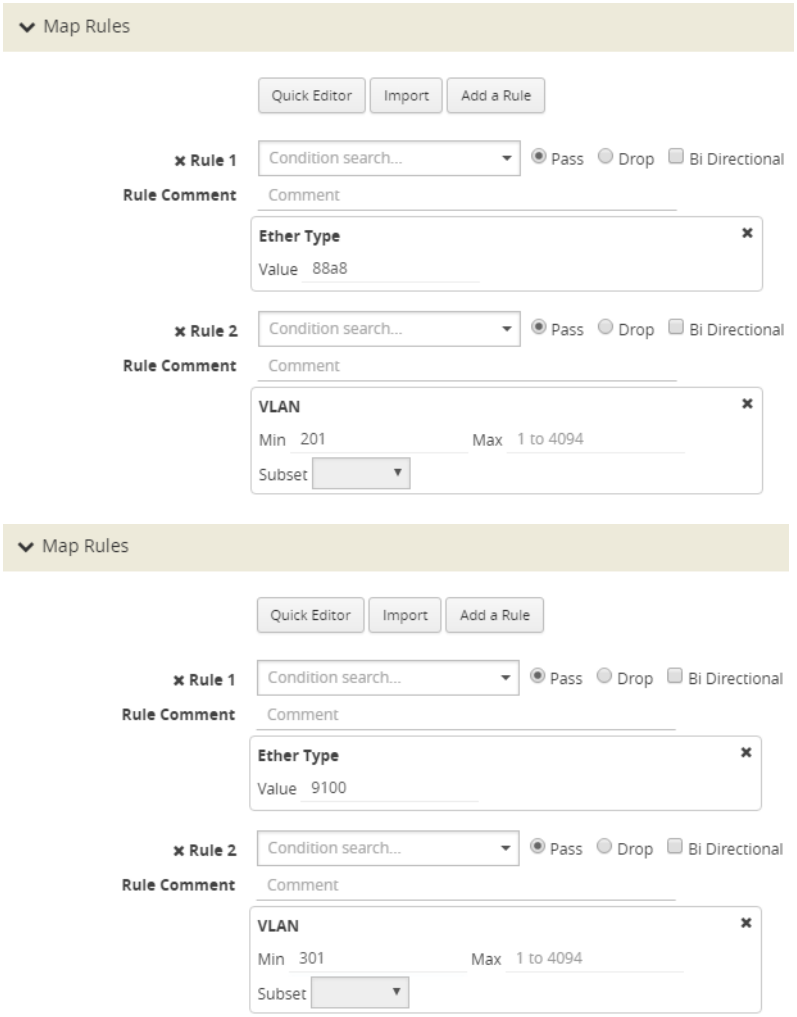

Figure 24-8: Rules with both TPID Ether Type and VLAN in Earlier Software Version

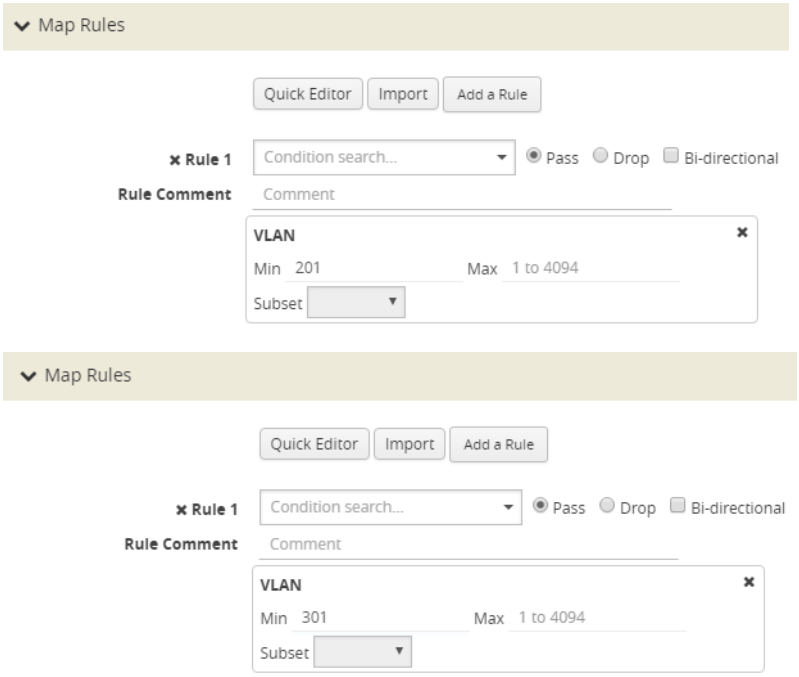

Figure 24-9: Rules Converted in version 4.7

- In the earlier software version, rules with only a TPID Ether Type Value, such as 88a8 or 9100, will be removed from the map in version 4.7.

It is recommended that you revisit your configuration after the upgrade.

Comparison of Q-in-Q Tagging

The following table details the various combinations and corresponding behaviors for software versions 4.7 and higher compared to 4.6.01 and earlier releases.

	4.7 and Higher			4.6.01 and Earlier		
Packet Content	**Rule: pass vlan 100**	**Rule: pass ether type 0x8000**	**Rule: pass vlan 100 ether type 0x0800**	**Rule: pass vlan 100**	**Rule: pass ether type 0x8000**	**Rule: pass vlan 100 ether type 0x0800**
No tags, ethertype 0800	drop	pass	drop	drop	pass	drop
One tag: TPID 8100, VID 100, ethertype 0800	pass	pass	pass	pass	pass	pass
One tag: TPID 9100, VID 100, ethertype 0800	pass	pass	pass	drop	drop	drop
One tag: TPID 88a8, VID 100, ethertype 0800	pass	pass	pass	pass	pass	pass
Two tags: outer TPID 8100 VID 100, inner TPID 8100 VID 200, ethertype 0800	pass	pass	pass	drop	drop	drop
Two tags: outer TPID 9100 VID 100, inner TPID 8100 VID 200, ethertype 0800	pass	pass	pass	drop	drop	drop

Packet Content	4.7 and Higher			4.6.01 and Earlier		
	Rule: pass vlan 100	Rule: pass ether type 0x8000	Rule: pass vlan 100 ether type 0x0800	Rule: pass vlan 100	Rule: pass ether type 0x8000	Rule: pass vlan 100 ether type 0x0800
Two tags: outer TPID 88a8 VID 100, inner TPID 8100 VID 200, ethertype 0800	drop	pass	drop	drop	pass	drop
Two tags: outer TPID 8100 VID 200, inner TPID 8100 VID 100, ethertype 0800	drop	pass	drop	drop	drop	drop
Two tags: outer TPID 8100 VID 100, inner TPID 88a8 VID 100, ethertype 0800	drop	pass	drop	drop	drop	drop
Two tags: outer TPID 88a8 VID 200, inner TPID 8100 VID 100, ethertype 0800	drop	pass	drop	drop	drop	drop
Two tags: outer TPID 8100 VID 100, inner TPID 88a8 VID 100, ethertype 0800	pass	drop	drop	pass	drop	drop
Two tags: outer TPID 8100 VID 100, inner TPID 9100 VID 100, ethertype 0800	pass	drop	drop	pass	drop	drop

Priority Tagged Packets

Starting in software version 5.3, priority tagged packets are handled by the GigaVUE node. These packets have a user priority of 0 to 7 in the packet. Single tagged packets or double tagged packets with a VLAN ID of zero or a non-zero value will be sent accordingly to the tool ports.

Flow Mapping on Inner VLAN Tags

Starting in software version 5.2, flow mapping on inner VLAN tags is supported for filtering on Q-in-Q traffic.

For packets that have both an inner and an outer VLAN tag, the outer tag is detected when the ethertype is 0x8100, 0x88A8, or 0x9100. The inner tag is detected only when the ethertype is 0x8100.

To specify an inner VLAN tag, add a new map rule (pass or drop) of type Inner VLAN.

Select a VLAN (Min) or a range of VLANs (Min and Max). Subset, even or odd, is optional.

The inner VLAN range is supported with any other qualifier with a range, such as VLAN or portsrc.

NOTE: There is no filtering after the two VLAN tags (inner and outer).

Filtering on inner VLAN uses application filter resources. To track resource usage, go to **Chassis > Quick Port Editor** for a particular box ID, card and slot.

Each map rule uses a number of entries. A single inner VLAN uses one entry per map rule. A range of inner VLANs uses two or more entries per map rule. For the same map source, identical inner VLAN or inner VLAN range spread across different rules will consume the same map rule resources.

A maximum of 454 application filter resource entries is available if no other application filters are using resources. The number of entries in the output of **Application Filter Resources** might be impacted by the other applications listed, such as GSD or Discovery.

The application filter resources are as follows:

- GSD—for GigaSMART tunnels
- Map Src—for network port source local to the node or slot (one entry per unique network port source). Note that 50 is always reserved per node or slot.
- Map Rule—for each inner VLAN rule
- Discovery—for LLDP/CDP

The following GigaVUE nodes have a maximum limit of 454 entries (the limit of 504 minus the 50 reserved):

- GigaVUE TA Series—per node
- GigaVUE-HB1—per node
- GigaVUE-HC1—per node
- GigaVUE-HC2—per node
- GigaVUE-HC3—per slot
- GigaVUE HD Series—per slot

Inner VLAN Limitation

Overlapped inner VLAN range is not supported within a map or set of maps that has the same network source. An identical VLAN range (and values) is supported.

For example, the following two rules are not supported because the inner VLAN range overlaps:

- Rule1: rule add pass inner-vlan 100 portsrc 1000
- Rule2: rule add pass inner-vlan 100..110 portsrc 1100

To overcome this, specify the rules as follows:

- Rule1: rule add pass inner-vlan 100 portsrc 1000
- Rule2: rule add pass inner-vlan 100 portsrc 1100
- Rule3: rule add pass inner-vlan 101..110 portsrc 1100

NOTE: You cannot use map rule editing to change an existing inner VLAN range to a range that overlaps with the original range. To edit an inner VLAN range, delete the rule and create a new rule with the new range.

Work with Map-Passalls and Port Mirroring in H-VUE

In addition to regular maps, H-VUE also makes it possible to create map-passall maps and configure tool-mirror ports for packet distribution.

A map-passall map lets you send all packets on a network port one or more tool ports or tool GigaStream on the same node or between the nodes in a cluster, irrespective of the packet distribution already in place for the ports.

Tool-mirror ports let you configure all a pass-all between two tool port or tool port and a too GigaStream on the same node, irrespective of the packet distribution already in place for the ports.

These map-passall maps and tool mirror ports are particularly useful in the following situations:

- Redirecting all traffic to IDS monitors regardless of any map rules applied to network ports.
- Temporary troubleshooting situations where you want to see all traffic on a port without disturbing any of the maps already in place for the port.

This section includes the following topics:

- *Syntax for Maps-passalls and Port Mirroring* on page 510
- *Rules for Map-Passalls and Port Mirroring* on page 510

Syntax for Maps-passalls and Port Mirroring

Refer to the following sections for details on map-passalls and port mirrors:

- *About Map-passall Maps* on page 489
- *Managing Ports* on page 401

Rules for Map-Passalls and Port Mirroring

Keep in mind the following rules for the map-passalls and tool mirrors:

- You can set up a map passall from:
  - Network port(s) to tool port(s) on the same node.
  - Network port(s) to one or more GigaStream.
- You can set up a tool mirror from:
  - Tool port to tool port(s) on the same node.
  - Tool port to GigaStream(s) configured with the advanced-hash algorithm on the same node.

  NOTE: The destination for a map-passall or tool-mirror can be a tool port, a hybrid port, a circuit port, a tunnel, a tool GigaStream, or a hybrid GigaStream.
- You cannot set up a map-passall or tool mirror from network port to network port. To be able to create such functionality, refer to *Port Access and Map Sharing* on page 511.

- A map-passall can cross line cards or modules – they can start on one line card/ module and end on another in the same node. Also, they can cross nodes in a cluster.
- A tool mirror can cross line cards or modules – they can start on one line card or module and end on another in the same node. They cannot, however, cross nodes in a cluster.
- Tool mirrors are not allowed from tool GigaStream to tool port.
- Tool mirrors are not supported on tool ports with copper SFPs installed or on 100Gb ports with CFP2 transceivers.
- A map-passall cannot be used with a port that is part of a port-pair.

View and Delete Map-passalls

Map-passalls are created by selecting **Subtype Pass All** for a **Type Regular** map. You can view the map by clicking on the map alias on the Maps page to open the Quick View for the map. To delete a pass-all map, select the map on the Map page and click **Delete**.

Port Access and Map Sharing

There are two ways to define a user's access to ports and maps:

- Port-based access levels
- Map sharing

Both methods assign permissions to user roles, as defined by the user groups, rather than specific user accounts.

Port-based Access Levels

Users are assigned roles based on their user group. Each user group is given permission to specific ports on the node. There are four port-based permission levels:

- Level 1—Can view the port but cannot make any changes to port settings or maps. When applied to a network port, can view maps attached to the network port. This level is used for users who only need to monitor the activities of the port.
- Level 2—Can use the port for maps, create tool-mirror to/from port, and change egress port filters. Can configure port-lock, lock-share, and all traffic objects except port-pair. Also includes all Level 1 permissions.
- Level 3—Can configure port parameters (such as administrative status of the port, speed, duplex, and autonegotiation, as well as create port pairs. Also includes all Level 2 and Level 1 permissions.
- Level 4—Can change the port type. Also includes all Level 3, 2, and 1 permissions.

Table 24-2 summarized the permissions for each of the levels.

Table 24-2: Port-based Permission Levels

Permissions	Level 1	Level 2	Level 3	Level 4	Admin
View port	✓	✓	✓	✓	✓
View maps attached to network port	✓	✓	✓	✓	✓
Create/edit map attached to port	×	✓	✓	✓	✓
Create tool-mirror to/from port	×	✓	✓	✓	✓
Change egress filters	×	✓	✓	✓	✓
Edit port parameters	×	×	✓	✓	✓
Create port pairs	×	×	✓	✓	✓
Change port type	×	×	×	✓	✓

How to share Maps

Maps can be shared with one or more user groups. When sharing a map, the map owner or Admin designates which user groups have which permissions. There are four map-sharing permission levels:

- **Read Only** – Can view the map but cannot make any changes.
- **Listen** – Can add or remove tool ports they own*. This is equivalent to "subscribing" to a map.
- **Read/Write** – Can delete and edit the map, can remove any network ports, can add network ports they own*, and can add or remove tool ports they own*.
- **Read/Write/Owne**r – Can perform all the Read/Write functions and assign map sharing permission levels.

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

Table 24-3 summarizes the permission levels for map sharing.

Table 24-3: Permission Levels for Map Sharing.

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

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

NOTE: In Table 24-3, tool port includes ports of type tool and inline-tool. Network port includes ports of type network and inline-network.

The admin user can also assign map sharing permissions.

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

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

Map Examples

This section provides the following map examples:

- *Example: How to Create a Simple Map* on page 513
- *Example: How to Handle Overlaps when Sending VLANs and Subnets to Different Tools* on page 516

Example: How to Create a Simple Map

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

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

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

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

The packet distribution of for this scenario is illustrated in Figure 24-10.

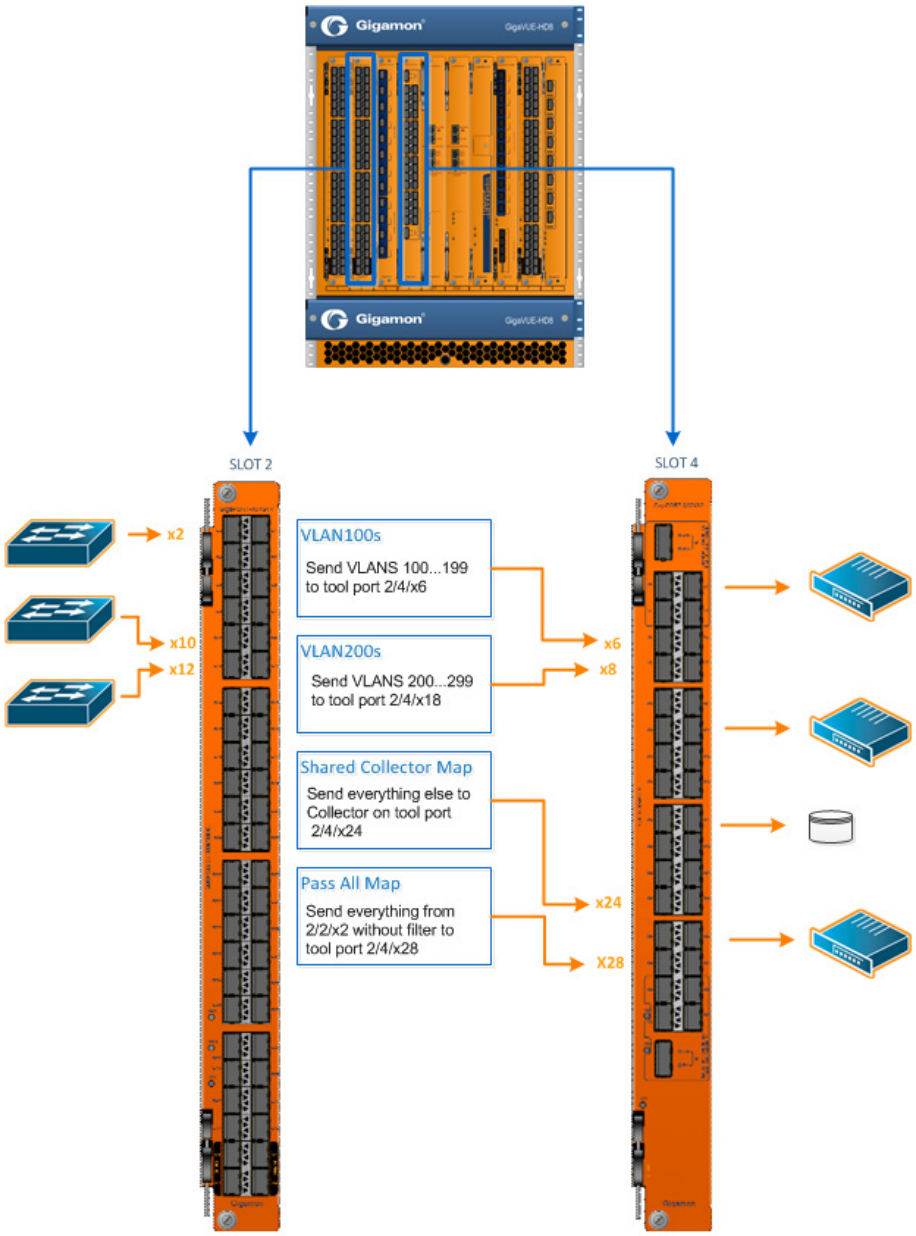

Figure 24-10: Graphical representation of Sample Maps

In the example shown in Figure 24-10, the configuration is as follows:

- BoxID for the GigaVUE-HD8 is set as 2.
- The network ports are set on the second blade in the.
- The tool ports are set on the fourth blade in the node.

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

- The network port IDs are set as 2/2/x2, 2/2/x10, and 2/2/x12.
- The tool port IDs are set as 2/4/x6, 2/4/x8, 2/2/x24, and 2/4/x28

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

- The maps with the VLAN rules specified have the subtype set as **By Rule**.
- The shared collector maps are set as subtype **Collector** and no rules added.
- The passall maps are defined as subtype **Pass All** and no rules added.

For details about the different map types, refer toTable 24-1 on page 493 and *Map Subtypes* on page 494.

Except for the following map types, the map type defaults to **Regular**

- Only maps with inline bypass solutions can be set as type **Inline**.
- Only maps that have GigaSMART operations defined for second level maps can be set as **First Level** or **Second Level** maps.

The following are the steps to create the simple maps as shown in Figure 24-10:

1. Check the port types for the each of the ports that the maps will use and set them if necessary.

   To set the port types, select Port > Ports > All Ports and use the Port Type Editor to set the port types. For more information about port types, refer to

2. Create a **Regular** map with a **By Rule** subtype to pass the VLAN100 traffic as shown in Figure 24-11.

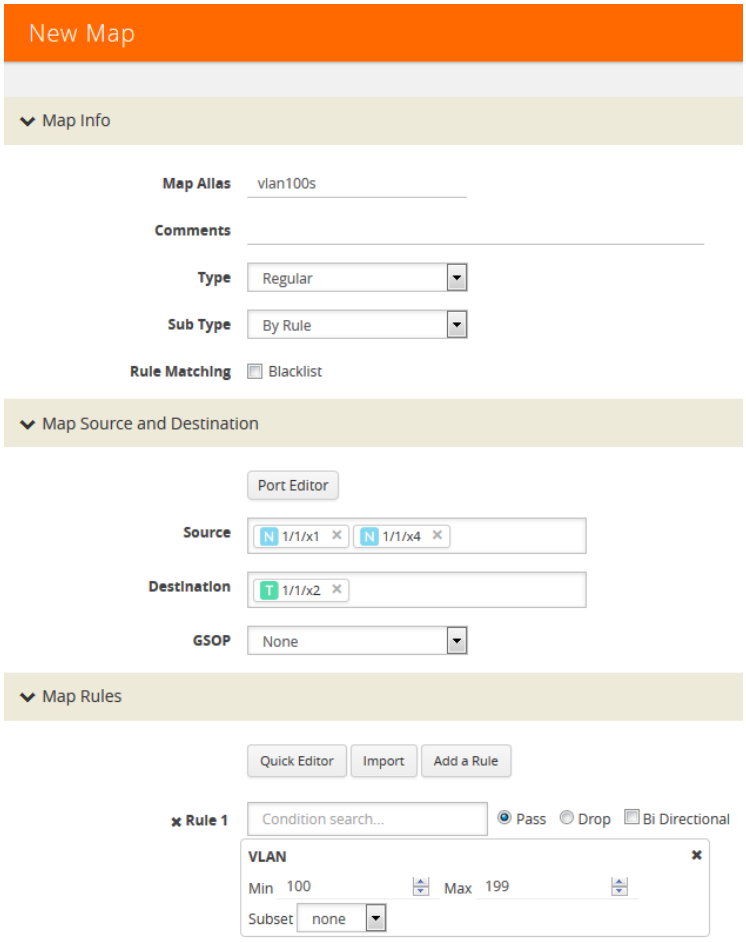

Figure 24-11: Map for VLAN 100..199

3. Create a **Regular** map with a **By Rule** subtype to pass the VLAN200 traffic.
4. Create a map with type **Regular** and subtype **Collector** to create a shared collector map. For more information about shared collector maps, refer to *About Shared Collectors* on page 487.
5. Create a map with type **Regular** and subtype **Pass All** to create the passall map.

Example: How to Handle Overlaps when Sending VLANs and Subnets to Different Tools

Figure 24-12 shows how to use map priority when handling packets matching criteria in multiple maps. In this example, we want to achieve the following results:

- Send packets on the 172.16.0.0 subnet to 1/2/x1
- Send packets on the 172.17.0.0 subnet to 1/2/x2
- Send packets on VLAN 100 to 1/2/x3

The trick is in how to handle packets on either 172.16.0.0 or 172.17.0.0 ***and*** VLAN 100. In this example, we use map priority to ensure that packets such as this are sent to both of their desired destinations.

Notice that the first two maps configured in Figure 24-12 are set up to handle this situation. For example, **map1** has a pass rule that accepts packets on 172.16.0.0 and VLAN 100. It sends matching packets to both 1/2/x1 (the destination we wanted for the 172.16 subnet) and 1/2/x3 (the destination we wanted for VLAN 100). Because this map was entered before **map3**, it has higher priority, ensuring the packet goes to both 1/2/x1 and 1/2/x3 and not just the 1/2/x3 destination specified by **map3**.

The same principle is applied in **map2** for packets on 172.17.0.0 and VLAN 100.

NOTE: If we did not observe the order of map entry shown in Figure 24-12, we could always adjust the priority as needed using the instructions in *Example: How to Create a Simple Map* on page 513.

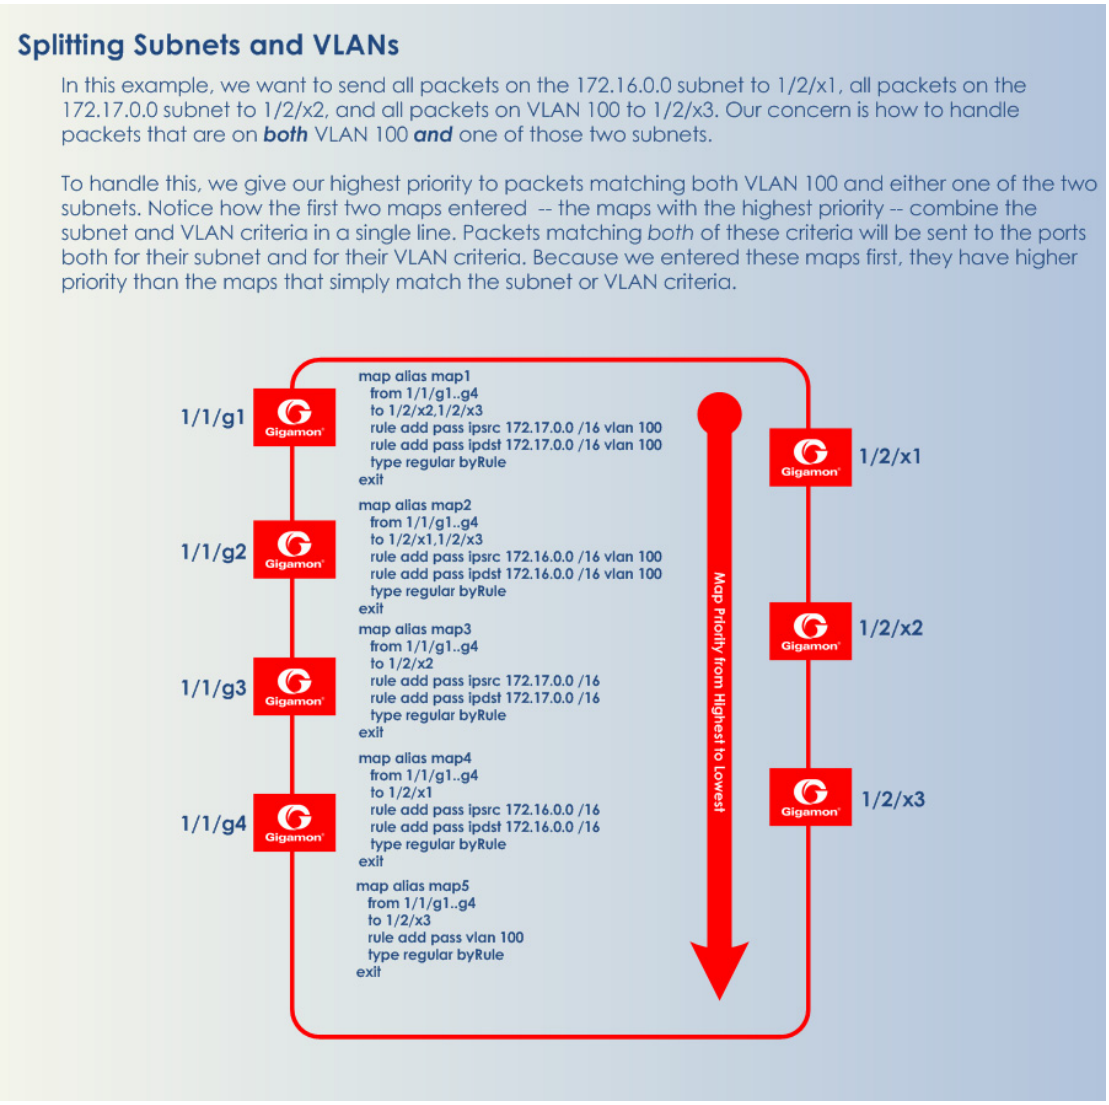

Figure 24-12: Sending Subnets and VLANs to Different Ports

Manage Maps

This section provides a description of the Maps pages in the H-VUE UI. It covers the following topics:

Map Views

The Maps page displays the maps created using the CLI, H-VUE, or GigaVUE-FM APIs. The maps can be displayed in following types of views:

- *List View*
- *Network View*
- *Tool View*

NOTE: Starting in software version 5.5.01, any change in the map health status is indicated by a **Status Update** notification pop-up that appears in the bottom-left corner of the Maps page. The link in the notification opens the map quick view.

List View

The List View is the default view of the maps when the Maps page is opened after selecting **Maps > Maps**. This view shows the basic information about each map, such as its alias, type, subtype, and source and destination ports.

Network View

The Network View displays the maps grouped by source ports then destination ports.

Tool View

The Network View displays the maps grouped by destination ports then source ports.

Manage Maps

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

- *Delete Maps* on page 521
- *Create Map Groups* on page 521

It also includes information about the following:

- *About Comments to Map Rules* on page 522
- *How to Use the Quick Editor for Pass and Drop Rules* on page 523

Map Aliases

A map alias specifies the name of the map. The alias must be unique and can contain up to 128 alphanumeric characters. Aliases are case-sensitive.

Most special characters are supported in map aliases. However, map aliases that are only one period (.) or two periods (..) should not be created. These aliases cannot be accessed for editing.

Create a New Map

The following are the steps for creating a map:

1. Check the status of the nodes and ports that you plan to use with the map.

   For information about how to check the status of the nodes and ports, refer to *Status of Line Cards/Nodes and Ports* on page 400.
2. From the device view, go to **Maps > Maps > Maps** to open the Maps page.
3. Click **New**.
4. Enter the Map Information:
   a. Enter an alias for the map.

      Use an alias that helps identify the task and destination. For example netflix_traffic_to_wireshark.
   b. (Optional) Enter a comment about the map. When adding comments, consider the following:
      - Use up to 128 characters, including spaces.
      - Enclose the comment in quotation marks, if the comment is longer than one word.
      - To include double quotation marks (") inside the quotation marks, precede it with a backslash (\).

      See also *About Comments to Map Rules* on page 522.
   c. Select the **Type**.

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

      For detailed information about the types of maps, refer to *Map Types* on page 491
   d. Select the map **Subtype**.

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

For detailed information about Pass All, refer to *About Map-passall Maps* on page 489. For detailed information about Collector, refer to *About Shared Collectors* on page 487.

5. Specify the **Map Source and Destination**.

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

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

   NOTE: For details about port types that are supported for the different types of maps, refer to *Port Lists* on page 396.

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

   NOTE: For details about circuit tunnels, refer to *About Circuit-ID Tunnels* on page 473.

   c. If the map is used to redirect the decapsulated traffic to the required tool ports, ensure that you select the IP interface in the **Source** drop-down list. You must have attached the IP interface to the VXLAN or L2GRE tunnel. For details about VXLAN or L2GRE tunnels, refer to *About Virtual Extensible LAN (VXLAN) Tunnels* on page 478 and *About Layer 2 Generic Routing Encapsulation (L2GRE) Tunnels* on page 476. From the **Destination** drop-down list, select the required tool ports.

   d. If the map will use a GigaSMART operation, select the operation from the **GigaSMART Operations (GSOP)** drop-down list.

6. Add rules to the map.

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

   - Use the **Quick Editor**. For details, refer to the *How to Use the Quick Editor for Pass and Drop Rules* on page 523.
   - Import a map template by clicking **Import**.
   - Create a rule by clicking **Add a Rule**.

   For detailed information about map rules, refer to *Map Rules* on page 495.

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

   For details about map priority, refer to *Map Priority* on page 490.

8. Set the **Map Permissions**.

   For details about map permissions, refer to *Port Access and Map Sharing* on page 511.

Clone Map

You can create a copy of an existing map by doing the following:

1. Select **Maps > Maps > Maps** to open the Maps page.
2. Select the check box of the map that you to clone.
3. Click **Clone**.

4. Make changes to the map as necessary, such as specifying an alias.
5. Click **Save**.

Edit Maps

To edit an existing map:

1. Click the List view button.
2. Select the map on the Maps page by either selecting the check box and then clicking **Edit**, or click on the row in the table and clicking **Edit** in the Map Quick View.
3. Make the changes to the map.

   When editing a map, you can only modify the following:

   - Comments
   - Change the source and destination.
   - Select a different GS Operation.
   - Modify rules.
   - Add rules.
   - Delete rules.
   - Permissions as allowed.
4. Click **Save**.

Delete Maps

To delete a map or maps, do the following:

1. Select **Maps > Maps > Maps** to open the Maps page.
2. Select the check box for the map or maps to delete, and then Click **Delete**.
3. When the message appears, asking if you want to delete the selected maps, click **OK**.

NOTE: In the GRIP configuration, when you delete a map on the primary node, irrespective of the inline-network traffic-path, the traffic is switched to the secondary node. The port utilization must be 0% on the primary node and active on the secondary node.

Create Map Groups

Map Groups are a collection of maps that are used with GTP Flow Sampling Overlap and GTP Whitelisting Overlap GigaSMART solutions.

Use the Map Groups page to create a group of maps for GTP whitelisting and GTP flow sampling. All the maps in a map group receive traffic according to map rules, rather

than map priority. Thus, multiple copies of a GTP packet can be sent to more than one tool. This functionality is referred to as overlapping maps.

The virtual port for specified as the source for GTP whitelisting and GTP flow sampling must have **GTP Overlap** enabled.

When creating Map Groups keep the following in mind:

- A map group can be associated with only one GigaSMART group (gsgroup).
- All maps within a map group must be connected to the same vport.
- A map group can consist of only one GTP whitelisting map or only one GTP flow sampling map but it cannot contains two maps of the same type.
- Once a map group is created, it cannot be edited to change the type or subtype of the map. However, you can add and edit the map rules for a map while it is configured in a map group.
- If multiple map groups are configured, the maps within each map group must point to the same port groups as the other map groups.

To create a Map Group, do the following:

1. Select **Maps > Maps > Map Groups** to open the Maps page.
2. Click **New**.
3. Enter an alias for the map.

   Use an alias that helps identify the map group.
4. (Optional) Enter a comment about the map group. Refer to *About Comments to Map Rules* on page 522 for the considerations regarding comments.
5. Click in the **Maps** field and select the maps to add to the map group.
6. Click **Save**.

About Comments to Map Rules

Use comments to label the purpose of a rule or the type of traffic covered by a rule. To add a map rule comment to a map select **Maps > Maps > Edit**.

Consider the following when adding map rule comments:

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

Error Messages

Error messages are displayed when a comment is invalid, for example:

- if the comment is longer than one word and does not include double quotation marks
- if the comment is longer than 128 characters

- if the rule with which the comment is included is not valid.

Edit Map Rule Comment

To edit a map rule comment, do the following:

1. Select **Maps > Maps**
2. Select the map to edit.
3. Click **Edit**.
4. Change the comments the **Comments** field
5. Click **Save** to recreate the with a different comment.

How to Use the Quick Editor for Pass and Drop Rules

When creating a map, you can use the Quick Rule Editor to quickly select custom port numbers for a map rule or add a range of IP address.

How to Use the Quick Editor to Add Port Numbers

To use the Quick Rule Editor, do the following:

1. While on an Edit Map or New Map page, click **Quick Editor** under Map Rules. This opens the **Quick Rule Editor**.

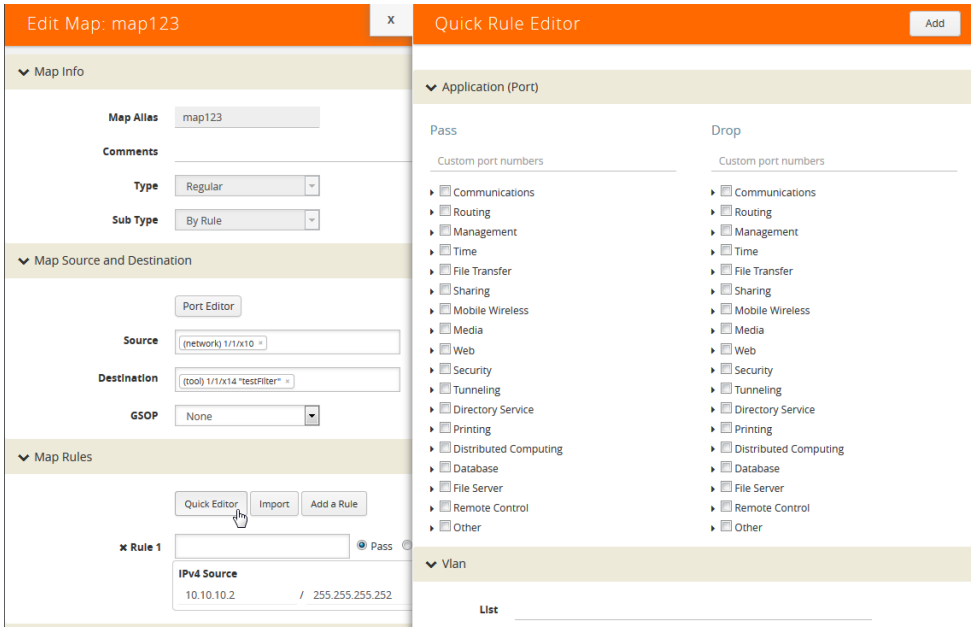

Figure 24-13: Selecting Quick Editor

2. On the Quick Rule Editor view, select the port number or numbers to add for a pass or drop rule or both.

   The Quick Rule Editor has two columns of custom port numbers, one for pass rules and one for drop rules. In each column, the ports are categorized by type. For

example, **Web** provides a list of HTTP ports as shown in Figure 24-14, where HTTPS port 443 is selected for pass and HTTP port 80 is selected for drop.

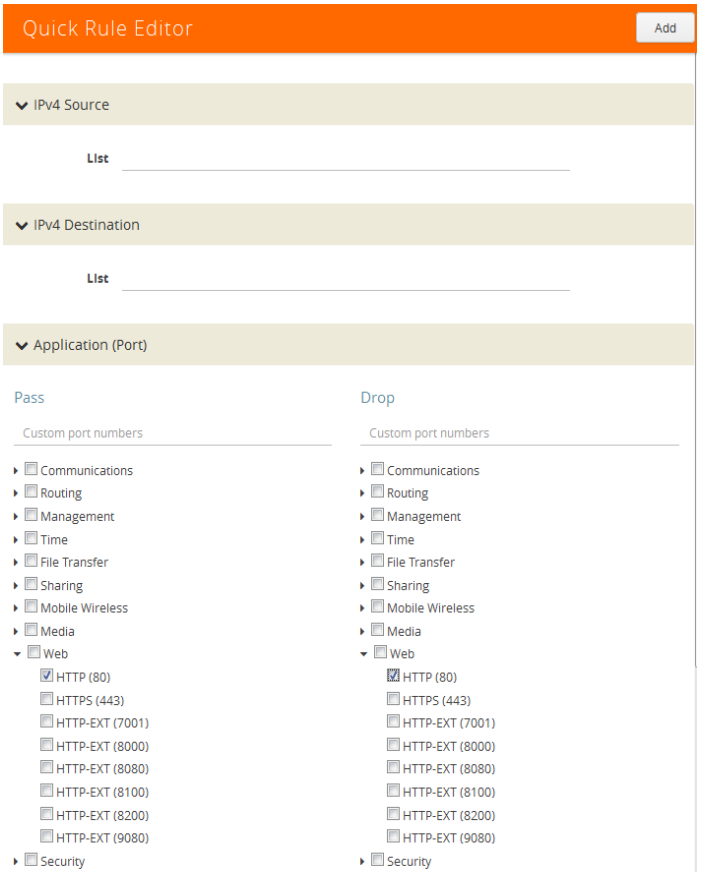

Figure 24-14: Custom Port Number Selected in Quick Rule Editor

3. Click **Add**.

   A rule with a port source is added for each custom port number selected in the Quick Rule Editor. If the port was selected from the custom port numbers under Pass, the rule is a pass rule. If the port was selected from the port numbers under Drop, the rule is a drop rule. Figure 24-15 shows two rules added by the example shown in the previous step.

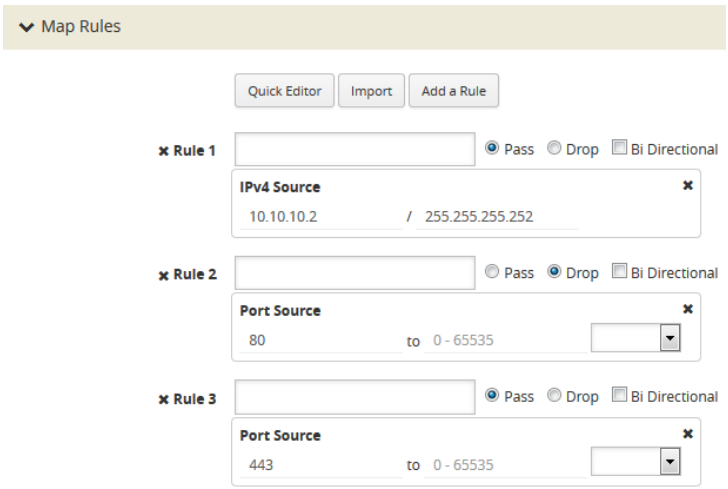

Figure 24-15: Rules Added with the Quick Rule Editor

How to Use the Quick Port Editor to Add IP Address

Rather than enter IP address for map rules one at a time, the Quick Rule Editor makes it possible to enter a range to quickly add the IP addresses, saving time.

To enter a range of IP address with the Quick Rule Editor, do the following:

1. While on an Edit Map or New Map page, click **Quick Editor** under Map Rules. This opens the **Quick Rule Editor**.
2. Enter an IP address range in the **List** field under IPv4 Source or IPv4 Destination or both. For example, 10.10.10.9..11/32 in the **List** field under IPv4 Source as shown in the following figure.

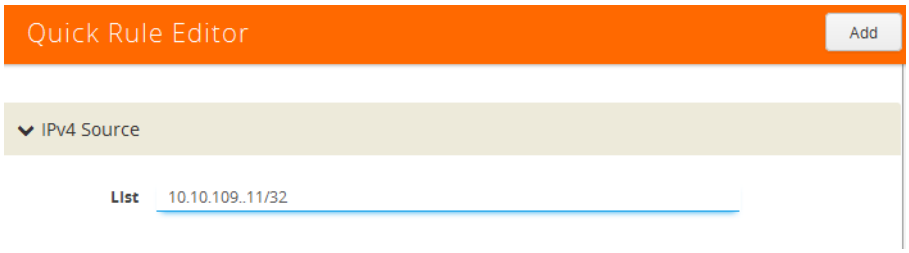

3. Click **Add**.

   The Quick Rule Editor adds the IPv4 Source rules with the IP addresses. For example, if you entered 10.10.9..11 for the IPv4 Source, the editor adds three Ipv4 Source rules with the IP addresses 10.10.10.9/32, 10.10.10.10/32, 10.1010.11/32 as shown in the Figure 24-15.

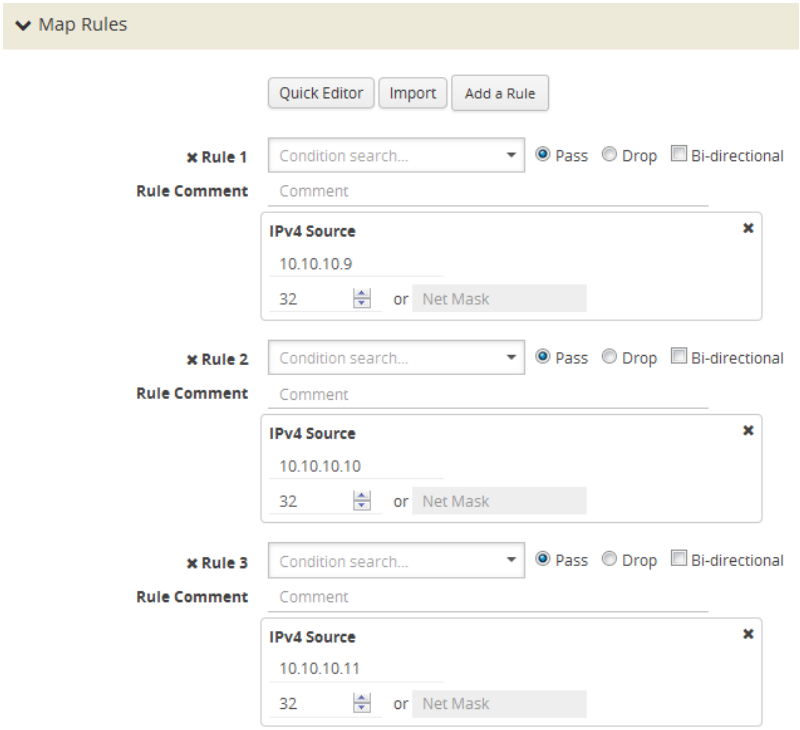

Figure 24-16: IPv4 Source Rules Added Through Quick Rule Editor

Map Templates

Map templates can be created by admin users. Once created, any user creating a map can use any template.

Admin users can define standardized traffic flows, applications, and rules that will be convenient for users when creating their maps. To do this, the admin creates map templates; later, users can use one or more templates as the basis for their maps.

Templates are created using the same rules and parameters as regular maps, but they do not have any network or tool ports. GigaSMART operations are also not included in templates.

The rules defined in the template become the starting point for the map. The rules can be edited or removed and new rules can be added to the map.

No changes made to the map will be reflected back in the original template. Once the map is created, it has no association with the original template from which it was created. Any changes to a template will not be reflected in any maps created with the previous version of the template.

Create Map Templates

To create a map template, do the following:

1. Select **Maps > Map Templates**.

2. Click **New**. The New Map Template opens.

Figure 24-17: New Map Template

3. Enter an alias in the **Map Template Alias** field.
4. (Optional) Enter comments about the template.
5. Add map rules by clicking **Add a Rule** for each rule that you want to add.
6. After you are done creating rules, click **Save**.

   The new map is included on the Map Templates page.

Edit Map Templates

To edit a map template, do the following:

1. Select **Maps > Map Templates** to open the Map Templates page.
2. On the Map Template page, select the template to edit.
3. Click **Edit**.
4. Modify the map template by adding or deleting rules or comments. (You cannot change the alias.)
5. Click **Save**.

Map Template Quick View

When you click on the alias of a map template on the Map Templates page, a Quick View displays that shows the comments (if any) and rules. An example is shown in Figure 24-18.

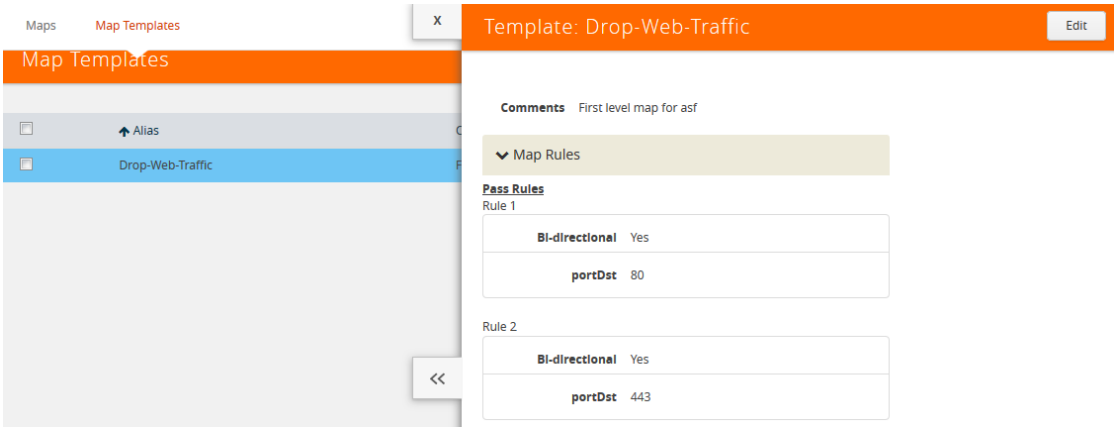

Figure 24-18: Map Template Quick View

Manage Map Rule Resources

The resources available on a GigaVUE H Series line card or node changes depending on the combination of criteria in place on the line card or node as a whole. In general, adding or removing MAC address, UDA pattern match, or IPv6 criteria in the map rules

bound to a line card or node changes the type of filter template used on the line card or node. This can result in a brief interruption of traffic as the new template is applied.

Template Groups

The template groups are listed sequentially from least resource-intensive to most resource-intensive:

- **IPv4 Only** – This is the default filter template, including all IPv4 arguments without any MAC addresses, UDA data patterns, or IPv6 arguments. This template can support the IPv4 and related filter criteria, including VLAN tags, source/destination ports, protocol criteria, and so on.
- **IPv4 and MAC Addresses** – This template combines MAC address criteria with the standard IPv4-related criteria. When MAC address criteria are in use, map rule resources are decreased.
- **IPv4 and Single UDA Data Pattern** – This template combines one of the two available UDA data patterns with the standard IPv4-related criteria. Using a single UDA criteria does not affect the total number of drop map rules available, but it does decrease the number of pass map rules available.
- **Both UDA Data Patterns** – This template uses both UDA data patterns but removes the ipv4 argument. Drop map rule availability is not affected by adding a second UDA data pattern, but pass map rules are decreased again from what was available when only a single UDA was used.
- **IPv6 Arguments** – This template adds the use of the IPv6 argument. IPv6 criteria are resource-intensive, significantly decreasing both drop and pass map rule capacity, as shown in the following table. Note also the changes in available criteria and available resources.

The map rule criteria available in each filter template (or “group”) is shown in the following table.

Table 24-4: Map Rule Criteria for Default Templates

Rules - Pass and Drop	IPv4	IPv4 + MAC	IPv4 + UDA	UDA	IPv6
IPv4	✓	✓	✓		
IPv6					✓
MAC		✓			
UDA1			✓	✓	
UDA2				✓	
VLAN	✓	✓	✓	✓	✓
Inner VLAN	✓	✓	✓	✓	✓
L4 Port dst and src	✓	✓	✓	✓	✓
Ethertype		✓	✓		
IP Ver	✓	✓	✓	✓	✓
Protocol	✓	✓	✓	✓	✓

Table 24-4: Map Rule Criteria for Default Templates

Rules - Pass and Drop	IPv4	IPv4 + MAC	IPv4 + UDA	UDA	IPv6
DSCP	✓	✓	✓	✓	✓
ToS	✓	✓	✓	✓	✓
TCP Ctl	✓	✓	✓	✓	✓
IP Frag	✓	✓	✓	✓	✓
TTL	✓	✓	✓	✓	✓
IPv6 Flow Label		✓	✓		

The number of rule entries in a cluster is shown in the following table.

Table 24-5: Rule Entries in a Cluster

# Rule Entries in a Cluster	IPv4	IPv4 + MAC	IPv4 + UDA	UDA	IPv6
GigaVUE-HD4/8 Line Cards, except PRT-HD0-C06X24	2048	1024	1024	512	512
PRT-HD0-C06X24 on GigaVUE-HD4/8	16384	8192	8192	4096	4096
GigaVUE-HC3	4096 (1024 per slot)	4096 (1024 per slot)	4096 (1024 per slot)	4096 (1024 per slot)	4096 (1024 per slot)
GigaVUE-HC2 CCv1 Node	4096	2048	2048	2048	2048
GigaVUE-HC2 CCv2 Node	16384	8192	8192	8192	8192
GigaVUE-HC1 Node	16384	8192	8192	8192	8192
GigaVUE-HB1 Node	2048	1024	1024	512	512
GigaVUE-TA10/TA40 Node	2048	1024	1024	512	512
GigaVUE-TA100 Node	1024	1024	1024	1024	1024
GigaVUE-TA200 Node	1024	1024	1024	1024	1024

NOTE: In addition, you can use flexible filter templates on GigaVUE-HC3 to support up to 6K rules per slot. On GigaVUE-TA100 and GigaVUE-TA200 they support up to 6K rules per pseudo-slot, or 24K rules per node. Refer to *Flexible Filter Templates* on page 530 for details.

Flexible Filter Templates

Flexible filter templates maximize the number of map rules, optimize filter resources, and enhance the scalability and flexibility of flow mapping. Flexible filter template is supported in GigaVUE-HC1, GigaVUE-HC2 CCv2, GigaVUE-HC3, GigaVUE-TA100, and GigaVUE-TA200.

Refer to *Manage Map Rule Resources* on page 527 for template groups on other GigaVUE nodes.

Flexible filter templates increase the number of map rules and also eliminate current restrictions on map rule combinations, such as ipv6+MAC or ipv6+UDA.

Refer to the section *Flow Mapping FAQ* on page 538 for the number of map rules supported.

Flow mapping uses filter templates to determine the traffic to filter based on qualifiers specified in the template. A filter template has a specific set of qualifiers used to apply to map rules. You can control the template that you apply to a specific slot on GigaVUE-HC3 or a specific pseudo-slot on GigaVUE-TA100 or GigaVUE-TA200. For GigaVUE-HC1 and GigaVUE-HC2 CCv2, you can apply the filter template only at the control card level which will be applied across all the line cards.

Flexible filter templates offer five default templates. Custom templates can also be created that have a qualifier set selected from the list of available qualifiers.

Refer to the following sections for details:

- *Filter Template Qualifiers and Defaults* on page 531
- *Custom Filter Template Configuration* on page 531
- *Filter Template Limits* on page 533
- *Filter Template Rules and Recommendations* on page 534
- *Filter Template Best Practices* on page 534
- *Filter Templates in a Cluster* on page 534
- *Filter Templates Formulas* on page 535

Filter Template Qualifiers and Defaults

Refer to the rows in Table 24-4 on page 528 for the list of qualifiers for filter templates. Refer to the columns in Table 24-4 for the default templates and the qualifiers that are predefined for the defaults.

NOTES:

- The default templates cannot be deleted.
- The **ipver** qualifier is implicitly included in all default and custom templates.

Custom Filter Template Configuration

To configure filter templates:

1. Access the GigaVUE node using a Web browser and log in with administrator user credentials.
2. Select **Maps > Filter Templates**. The Filter Templates page shown in Figure 24-19 on page 531 displays the default templates.

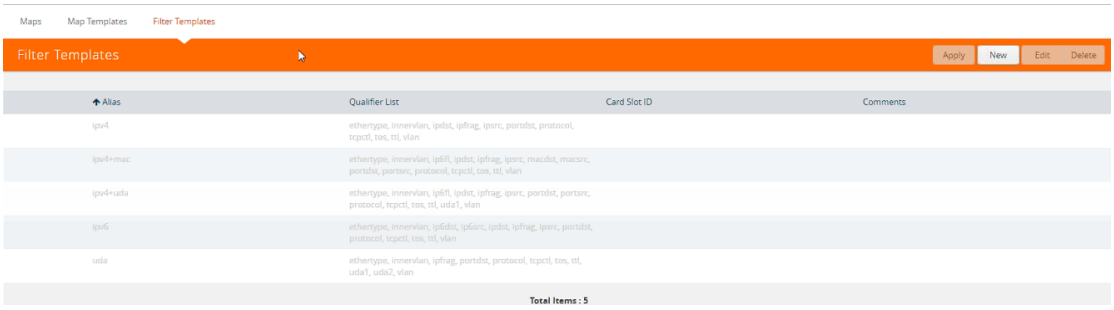

Figure 24-19: Default Filter Templates

3. To add a custom template, click **New**. Refer to Figure 24-20 on page 532.

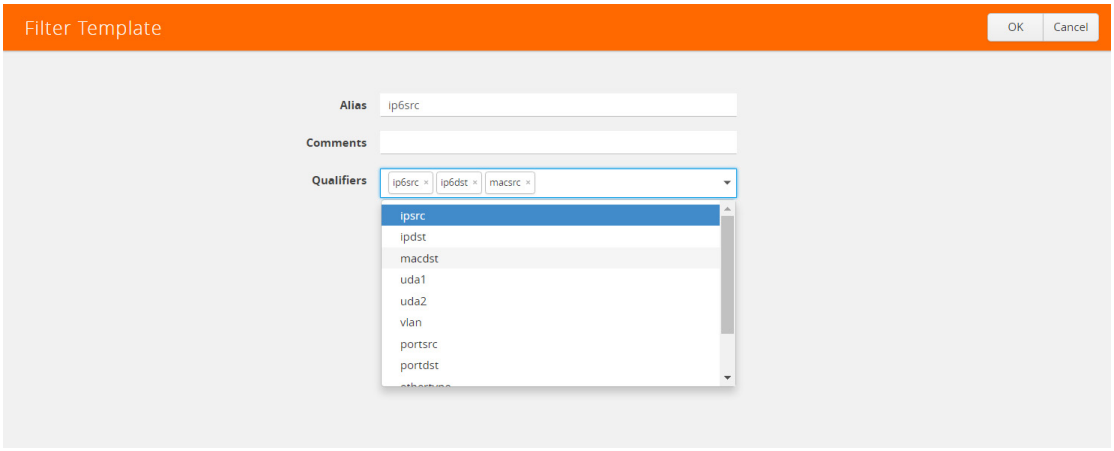

Figure 24-20: Custom Filter Templates

4. Specify an alias, an optional comment, then select qualifiers. Click **OK**.
5. To apply a custom template to a slot or pseudo-slot, select it and click **Apply**. Refer to Figure 24-21 on page 532.

NOTE: For GigaVUE-HC1 and GigaVUE-HC2 CCv2, you can apply a filter template only at the control card level which will be applied across all the line cards.

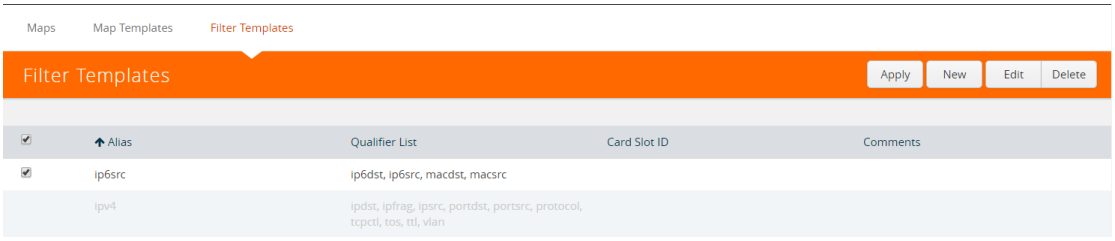

Figure 24-21: Apply Filter Template

6. Select the slot or pseudo-slot and click **OK**. Refer to Figure 24-22 on page 532.

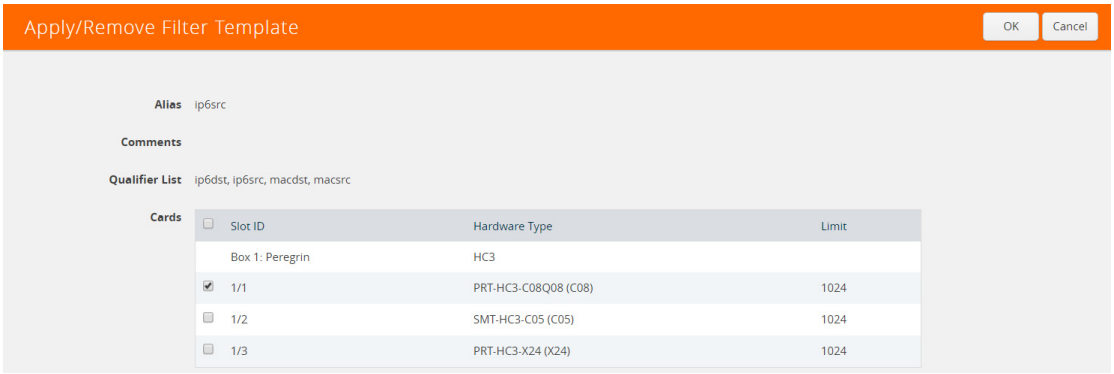

Figure 24-22: Apply Filter Template to Slot

The Filter Templates page displays the applied slot or pseudo-slot. You can edit an existing custom filter or delete it. A template can be deleted if it is not currently in use, meaning that it has not been applied. Refer to Figure 24-23 on page 533.

Figure 24-23: Filter Template Edit or Delete

7. To display filter templates, click on a row in the Filter Templates page. Refer to Figure 24-24 on page 533.

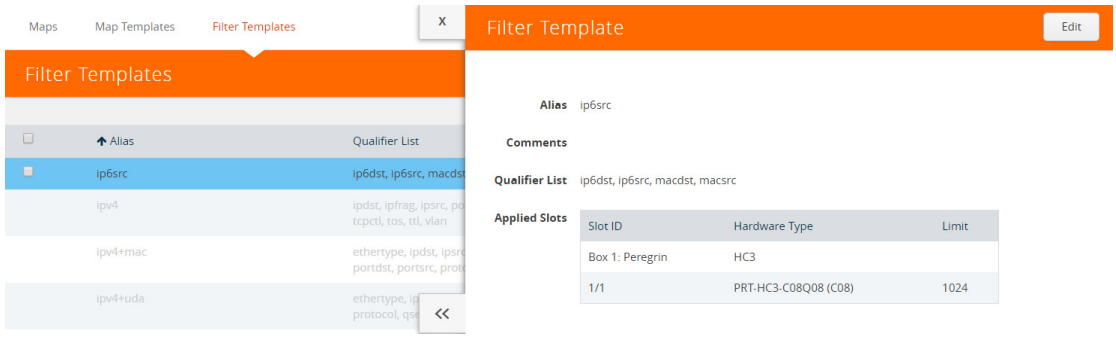

Figure 24-24: Filter Template Display

Filter Template Limits

The number of qualifiers in a template limits the total number of rules that can be defined. The maximum rule limit on the GigaVUE-HC3, GigaVUE-TA100, or GigaVUE-TA200 is 1K (1024) per slot or pseudo-slot when using the default templates.

Custom templates allow the creation of templates with only those qualifiers needed for the rules that you plan to use in flow maps. The qualifiers specified in a flexible template can increase or decrease the maximum rule limit, depending on the qualifiers selected. With flexible filter templates, it is possible to reach a maximum limit of 6K rules per slot on the GigaVUE-HC3 node and 6K rules per pseudo-slot on the GigaVUE-TA100 or GigaVUE-TA200 node, or 24K total rules.

Figure 24-24 on page 533 displays a Limit.

How to Understand Map Filter Resources

Starting in software version 5.0, when a filter template is applied, filter resources display the total number of map rules used in a map as well as the limit. If the limit is 1024, 1023 is displayed, even though the actual limit is 1022, or two less than the limit. This discrepancy is due to extra resources needed for internal usage.

Filter Template Rules and Recommendations

When creating flexible filter templates, keep the following rules and recommendations in mind:

- Filters are applied to a specific slot or pseudo-slot, not to the node.
- By default, all slots will be in the pre-defined **ipv4** template.
- There is a limit of 512 custom templates.
- Custom templates can have duplicate sets of qualifiers.
- The filter limit is calculated when the template is created. In most cases, a higher-cost qualifier set (for example, IPv6, UDA, or MAC are higher cost) consumes more resources and leads to a lower filter limit.
- Flexible filter templates have no effect on existing flow mapping behavior, including pass versus drop map rules, map priority, network port sharing, GigaSMART operations, or first level and second level maps.

Filter Template Best Practices

The following are best practices for optimizing filter resources using filter templates.

First determine all the needed qualifiers, then create a template, apply the template, and configure the map rules.

- Connect network ports of a slot to flows of the same application.
  For example, if you have two flows:
  - one is filtered on **macsrc** and **macdst**
  - the second one is filtered on **ipdst** and **ipsrc**
- In case both flows connect to ports on the same slot, that slot will have to have a template of **macsrc**, **macdst**, **ipsrc**, and **ipdst**, with a limit of 1024 rules.
- However, filter resources can be optimized by connecting these two flows to ports on different slots with one template for **macsrc** and **macdst** and the other template for **ipsrc** and **ipdst**. Both templates will have a limit of 3072 rules.

The following are best practices for adding more rules if a limit has been reached:

- Create a new template with all the qualifiers that are in use on a specified slot.
- Issue the **show filter-resource slot** command to obtain the list of qualifiers in use.
- Issue the **filter-template alias <alias> qualifiers add** command with that list of qualifiers.
- Issue the **show filter-template limit** command to check if the new template allows a higher limit. If it does, apply the filter using the **card slot <slot ID> filter-template** command.

Filter Templates in a Cluster

Filter template configuration is synchronized across the cluster. However, a cluster can have different GigaVUE nodes, so one set of qualifiers may or may not be valid on all nodes.

Filter Templates Formulas

The formulas in this section can help you determine the number of map rules that are supported, based on the qualifiers specified in the filter template. Use the formulas as guidelines.

The number of map rules depends on the number of qualifiers a template can support. The more qualifiers, the lower the limit.

The cost of each qualifier depends on the number of bits it consumes. The following table lists the number of bits each qualifier consumes.

Table 24-6: Bits Consumed per Qualifier

Qualifier	Bits
ipdst	32
ipsrc	32
ip6dst	128
ip6src	128
macdst	48
macsrc	48
uda1	128
uda2	128
vlan	16
inner-vlan	16
portdst	16
portsrc	16
ethertype	16
protocol	8
qset1	58*

* qset1 is made up of the following: tos: 8, ipfrag: 2, tcpctl: 8, ttl: 8, ip6fl: 32

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

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

Examples:

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

- For the portdst qualifier only—portdst is 16 bits, so the total is 16+54 bits, which is a cost less than 80 bits, so 6K rules per slot are supported.

The maximum cost supported is 480 bits/template.

Review Map Statistics with Map Rule Counters

Map Statistics can be viewed in the following ways:

- The Statistics page.

  For details, refer to *Viewing Map Statistics with the Statistics Page* on page 536
- The Map Quick View.

  For details, refer to *Viewing Map Statistics with Quick View* on page 537

A single packet may match multiple rules in the map and will not cause multiple rule counters to increment. Only the last rule which is the highest priority in the order will increment. The flow map rule priority is based on the order it was created. Thus, the sum of the map rule counters across all the map rules may be higher than the total number of packets received and transmitted by the map.

NOTE: Drop rules have a higher priority than pass rules.

For example, consider the following three map rules:

- **Rule 1** – ipsrc 10.10.0.0 /24 bidir
- **Rule 2** – ipsrc 10.10.0.100 /32 bidir
- **Rule 3** – portsrc 80

A packet with ipsrc 10.10.0.100 and portsrc 20 will match Rule 1 and Rule 2, which will be forwarded to the tool port or ports. The counters for Rule 2 will only be incremented.

There are several reasons a map rule counter may not increment:

- Traffic matching the rule is not currently present in the production network.
- The network port is not monitoring the network at the proper location to see the traffic specified by the map.
- A higher-priority map is forwarding the packet before it can be inspected by this particular map.
- The map rule itself may be specified incorrectly.

Viewing Map Statistics with the Statistics Page

To review map statistics indicating the total packets and total octets handled by maps and the number of rules in the map, select **Maps > Maps > Statistics**. The Statistics page displays the map statistics in a table format listing the maps by their alias as shown in Figure 24-25. Clicking on a map alias opens a Quick View for that map. To clear map counters, click the **Clear** button.

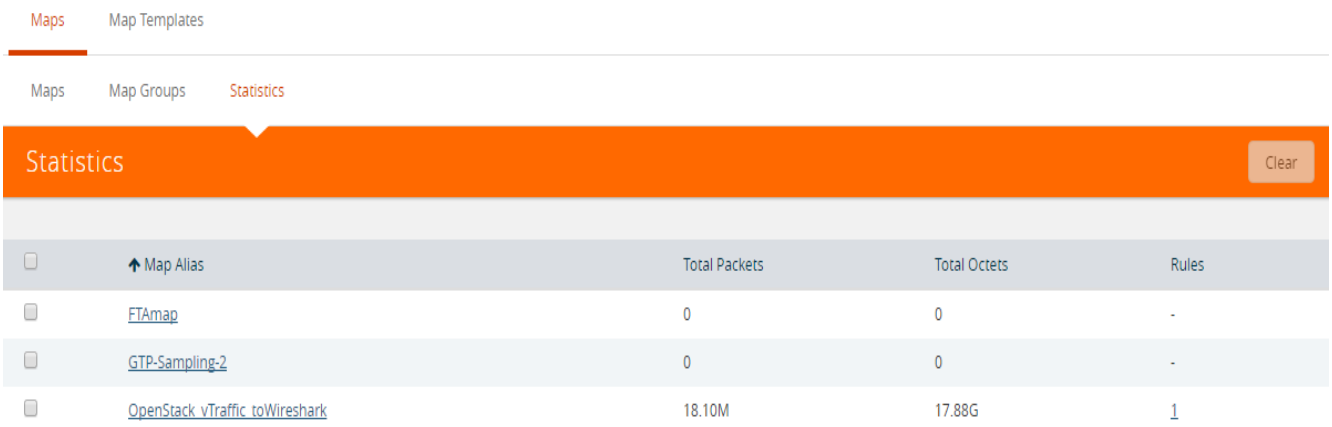

Figure 24-25: Map Statistics Page

Viewing Map Statistics with Quick View

To review map rule counters indicating the number of rule matches for a map as packets are inspected and forwarded, select the map and view the information in the Quick View window as shown in Figure 24-26.

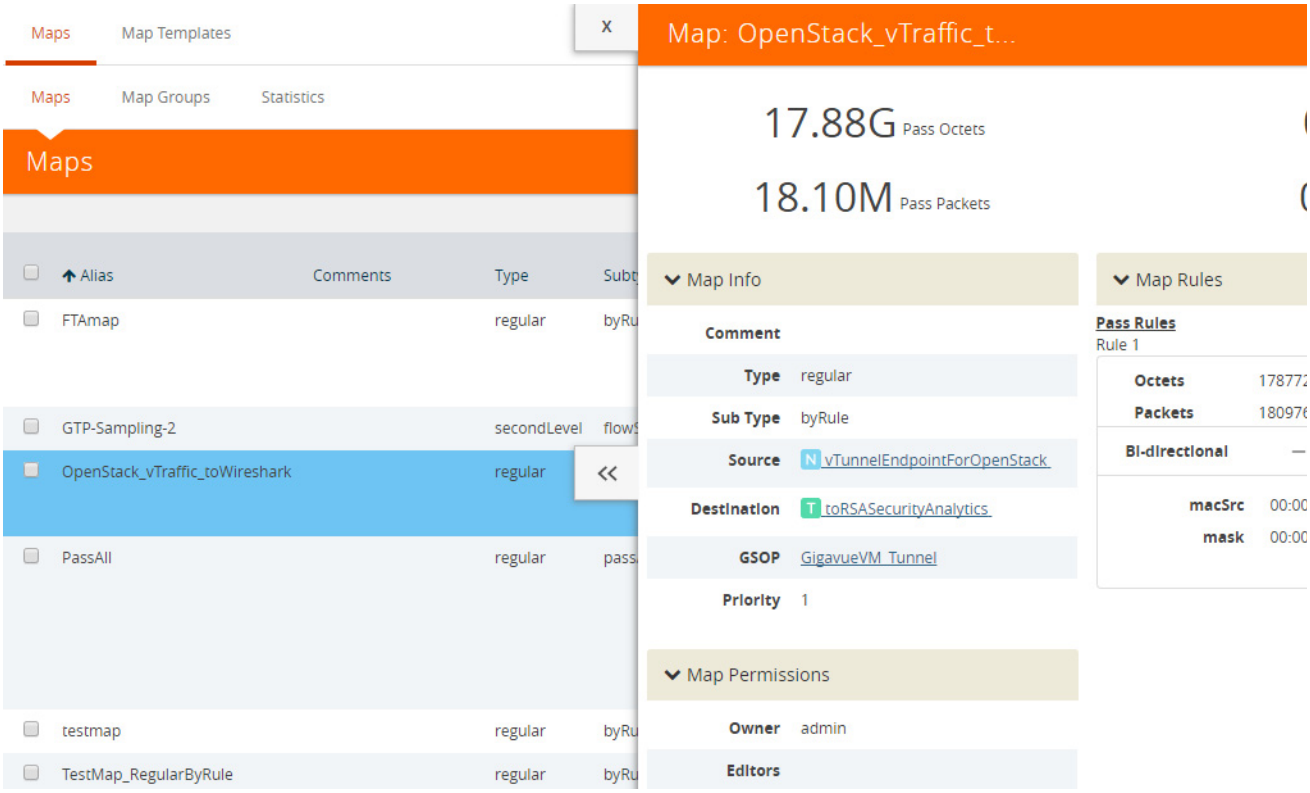

Figure 24-26: Map Counters Available in Quick View Page

Flow Mapping FAQ

This section answers frequently asked questions by users migrating to the Flow Mapping model.

How Many Map Rules are Supported?

The maximum number of map rules supported per line card or standalone node are shown in Table 24-8 on page 539.

Table 24-7: Maximum Map Rules

Node Type	Maximum Combined Rules	Maximum with Flexible Filter Templates	Maximum with Advanced Features License
GigaVUE-HD8/GigaVUE-HD4 Line Cards, except PRT-HD0-C06X24	2K (2048)	N/A	N/A
GigaVUE-HD8/HD4 PRT-HD0-C06X24 Line Card	16K (16383)	N/A	N/A
GigaVUE-HB1	2K (2048)	N/A	N/A
GigaVUE-HC2	4K (4096)	N/A	N/A
GigaVUE-HC2 with Control Card version 2 (HC2 CCv2)	16K (16383)	16K (16383)	N/A
GigaVUE-HC3 Node **GigaVUE-HC3 Module**	4K (4096) 1K (1024)	24K (24576) 6K (6144)	N/A
GigaVUE-HC1	16K (16383)	16K (16383)	N/A
GigaVUE TA Series	256	The Flexible Filter Templates are available only for GigaVUE-TA100 and GigaVUE-TA200 nodes.	The Advanced Features License is available only for the GigaVUE-TA Series. The extended max rule limit varies based on the GigaVUE-TA Series node.
GigaVUE-TA10	256	N/A	2K (2048)
GigaVUE-TA40	256	N/A	4K (4096)
GigaVUE-TA100	256	256	6K (6144) per pseudo-slot with Flexible Filter Template and Advanced Feature License. 1K (1024) per pseudo-slot with only Advanced Feature License.
GigaVUE-TA100-CXP	1K (1024)	N/A	N/A
GigaVUE-TA200	256	256	6K (6144) per pseudo-slot with Flexible Filter Template and Advanced Feature License. 1K (1024) per pseudo-slot with only Advanced Feature License.

The limit for GigaVUE TA Series standalone nodes is 256 combined pass/drop rules but is up to 2048 with the Advanced Features License installed.

Refer to *Manage Map Rule Resources* on page 527 for managing map rules.

The maximum number of nodes and map rules supported when in a cluster is as follows:

Table 24-8: Maximum Number of Nodes and Map Rules Supported in a Cluster

When a Cluster is Configured with:	Number of Nodes	Maximum Map Rules
Out-of-Band Cluster Management	32	38000
Inband Cluster Management	16	38000

The maximum number of map rules supported in a cluster apply to all nodes in the cluster including GigaVUE H Series nodes: GigaVUE-HD8, GigaVUE-HD4, GigaVUE-HC3, GigaVUE-HC2, GigaVUE-HC1, and GigaVUE-HB1, and GigaVUE TA Series nodes: GigaVUE-TA1,GigaVUE-TA40, and GigaVUE-TA100, including Certified Traffic Aggregation White Box (white box).

How Many Rules Can Each Map Have?

The maximum number of rules per map is 4K (4096), except on products that only support a total of 2K map rules. GigaVUE-HC3 supports 6K rules per map. Refer to Table 24-8 on page 539.

For example on a GigaVUE HD Series line card, a single map can have the maximum number of rules allowed. However, such a map would consume all of the mapping resources for the line card.

How Many Maps Can Run at Once?

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

What Criteria can be Filtered in Q-in-Q Packets?

Maps on GigaVUE nodes can match Layer 3/Layer 4 criteria in packets using Q-in-Q with up to two tags. For more information refer to *How to Handle Q-in-Q Packets in Maps* on page 504.

Which Line Card Do Map Rules Count Against?

The limits for map rules are based on the locations of the network ports for the map and count equally against each line card in GigaVUE HD Series nodes with ports in the map. For example, a map with 500 pass rules provisioned against ports 1/4/x2 and 1/5/x8 counts as 500 rules against the maximum available for the line card in both slot 1/4 and 1/5.

For optimal allocation of map rules, maps using different sets of multiple network ports should keep those sets on different line cards, modules, or nodes. However, this is not a requirement.

Example

Consider the following scenario:

- Line card in slot 2/8 already has 1600 pass rules allocated.
- Attempting to provision a map consisting of 500 pass rules on network ports 1/1/x1 and 2/8/x9 will result in an error message stating that the map cannot be created.
- Reconfiguring the map to use network ports 1/1/x1 and 1/1/x2 (instead of 1/1/x1 and 2/8/x9) prevents this exposure and allows the map to be created.

How Many Maps Can Share a Network Port?

There is no limit to the number of maps that can share a network port.

How Many Network Ports and Tool Ports Can Be in a Map?

If the ports are not in a GigaStream, the number of individual map ports in the **Source** or **Destination** field of a map is limited to 64 on all GigaVUE H Series nodes. The individual ports can be any of the following port types: network or tool.

On GigaVUE-HC2 and GigaVUE-HC3, if the ports are in a GigaStream, the limit is 95.

Are Port-Filters Supported?

Yes.

Egress port-filters are less efficient and scalable than flow maps, but they do provide a convenient way to narrow down the traffic seen by the tools/GigaVUE H Series nodes without having to change an entire map. Refer to *Port Filters* on page 399 for details.

Each GigaVUE HD Series line card, GigaVUE-HC2 or GigaVUE-HC3 module, or GigaVUE-HB1 node supports 100 combined tool port-filters. A single filter applied to multiple tool ports counts multiple times against the 100-filter limit.

In the GigaVUE-HC2 equipped with Control Card Version 2 (HC2 CCv2), or the GigaVUE-HC1 or GigaVUE-HC3 node, the limit is 400 filters.

The GigaVUE TA Series can only support 20 tool port-filters. When the GigaVUE-TA100 or GigaVUE-TA200 are in a cluster, they can support 400 filters.

Does Flow Mapping Support Passalls?

Yes.

Flow Mapping supports passalls with the following:

- The map **Subtype Pass All** option for network to tool port passalls.

- A tool-mirror connection between to tool ports. The Tool Mirrors page replaces the tool-to-tool port passalls. To create a Tool Mirror, select **Ports > Tool Mirrors** to go to the Tool Mirror page shown in Figure 24-27.

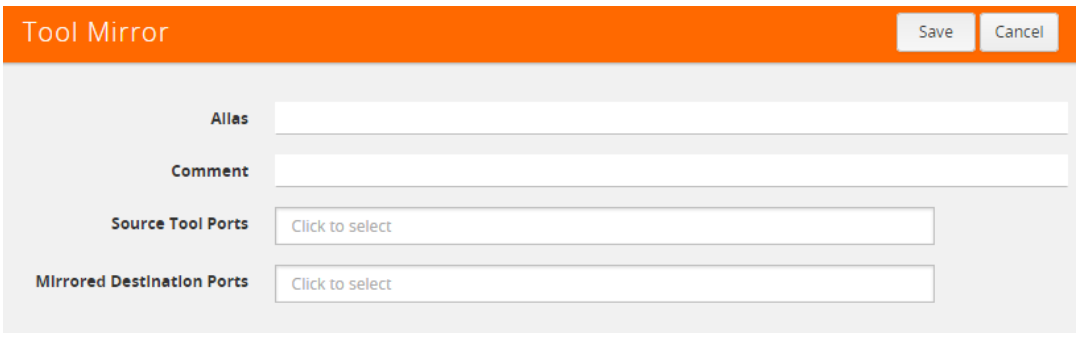

Figure 24-27: Tool Mirror Page

Does Flow Mapping Support port-pairs?

Yes.

Select **Ports > Port Pairs** to go to the Port Pairs page shown in Figure 24-28. For more information about port pairs and details on configuring two ports as an inline TAP, refer to *Port Pairs* on page 413.

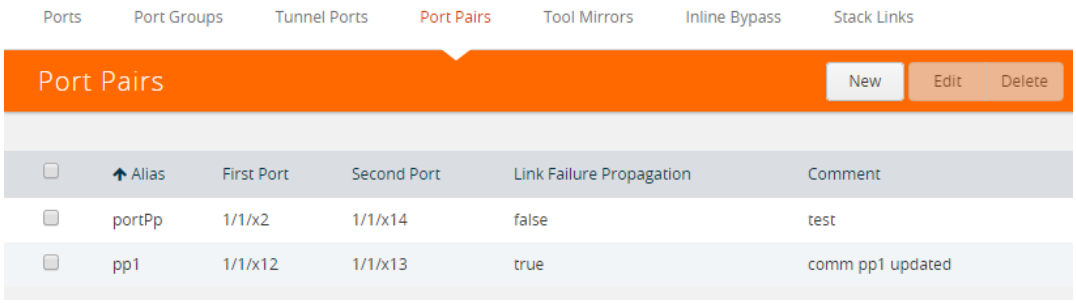

Figure 24-28: Creating Port Pairs

Does Flow Mapping Support UDA Pattern Matches?

Yes.

Pattern match rules are still supported in map rules. However, they can only be used in pass rules. UDA pattern matches are not supported in drop rules. Refer to *Work with User-Defined Pattern Match Rules* on page 501 for details.

Are Maps Supported Across Nodes in a Cluster?

Yes.

Clusters of GigaVUE nodes operate as a unified fabric. Use the standard box ID/slot ID/port ID syntax to create packet distribution, just as on a standalone node. Maps can have network ports and tool ports on different physical nodes within the cluster.

Similarly, a map does not need to keep its network and tool ports on the same physical node in order to take advantage of GigaSMART operations – a GigaVUE TA Series node in a cluster can take advantage of the GigaSMART processing available on a GigaVUE HD Series line card, GigaVUE-HC2 or GigaVUE-HC3 module, or GigaVUE-HC1 or GigaVUE-HB1 node elsewhere in the same cluster.

Does Flow Mapping Support GigaSMART Operations?

Yes.

The wizards in H-VUE make it easier than ever to include GigaSMART operations – rather than having to create a GigaSMART operation separately before including it in a map, you can now create and use it all within the same mapping wizard. Refer also to *Work with GigaSMART Operations* on page 741 for examples of flow mapping.

Can a GigaStream Act as a Shared Collector?

Yes.

Multiple individual ports for a Shared Collector can be setup or a GigaStream tool group. Refer to *About Shared Collectors* on page 487 for details.

What Are the GigaStream Maximums?

The number of GigaStream per line card, module, or node varies by product. Refer to *GigaStream Rules and Maximums* on page 456 for the details.

Does H-VUE Provide the Same Features as the GigaVUE-OS CLI?

The H-VUE provides all of the mapping features of the CLI in an intuitive and highly-usable setting. The map creation wizard speeds map creation.

Configure Active Visibility

This section describes active visibility. Refer to the following sections for details:

Overview of Active Visibility

Active visibility is a framework that allows your visibility network to adapt to dynamic events. The framework is designed to react to evenfts and take actions in response to events in your visibility network.

An active visibility policy defines conditions and actions. When conditions are met, actions are executed. The policy specifies both the conditions and the actions and ties them together.

The conditions and actions are pre-defined. Conditions are events that are used to trigger changes to configuration. Actions can notify users of certain events and/or modify the configuration in response to certain events. Conditions can be port-based or time-based.

For example, if a tool port is overloaded, you might want to reduce the traffic sent to the tool. You can configure two maps, one targeted for your high priority traffic and the other targeted for your low priority traffic. Then you can specify conditions to monitor tool port utilization. When traffic is below a threshold, both maps can be enabled, thus all traffic will be sent to the tools. But when traffic is above the threshold, you can specify an action so that only the map targeted for the high priority traffic is enabled and thus only that traffic will be sent to the tool.

Another example is that you might want to use a different set of maps and map rules to provide visibility during different times, such as during working hours, or on weekends. If there is going to be a backup on the weekend, you can specify a policy to disable a map at a specific time or day.

For details, refer to “Configuring Active Visibility” in the *GigaVUE-OS CLI User’s Guide*.

Configure Active Visibility

Refer to the following sections for details:

- *Display Conditions* on page 544
- *Display Actions* on page 544
- *Configure Policies* on page 544

Display Conditions

Go to **Active Visibility > Conditions** to display the pre-defined conditions, their description, and comment. Refer to Figure 24-29 on page 544.

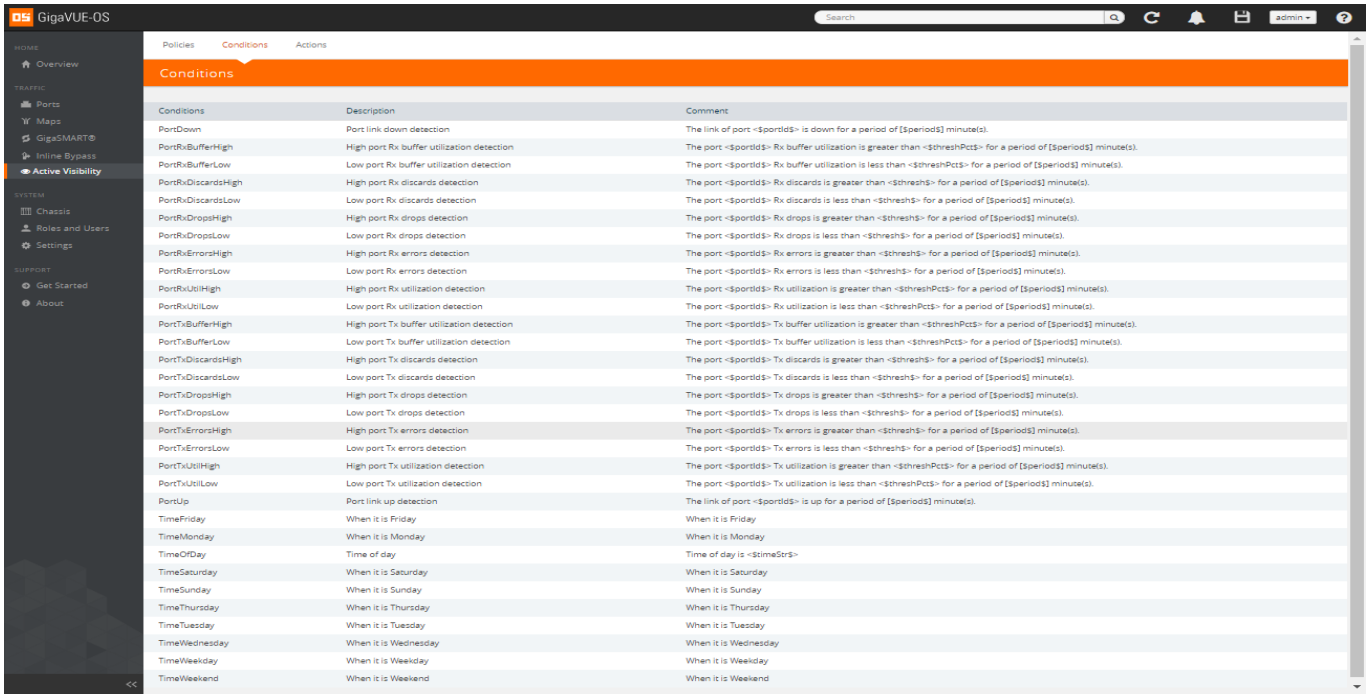

Figure 24-29: Conditions, Description, and Comment

Display Actions

Go to **Active Visibility > Actions** to display the pre-defined actions, their description, and comment. Refer to Figure 24-30 on page 544.

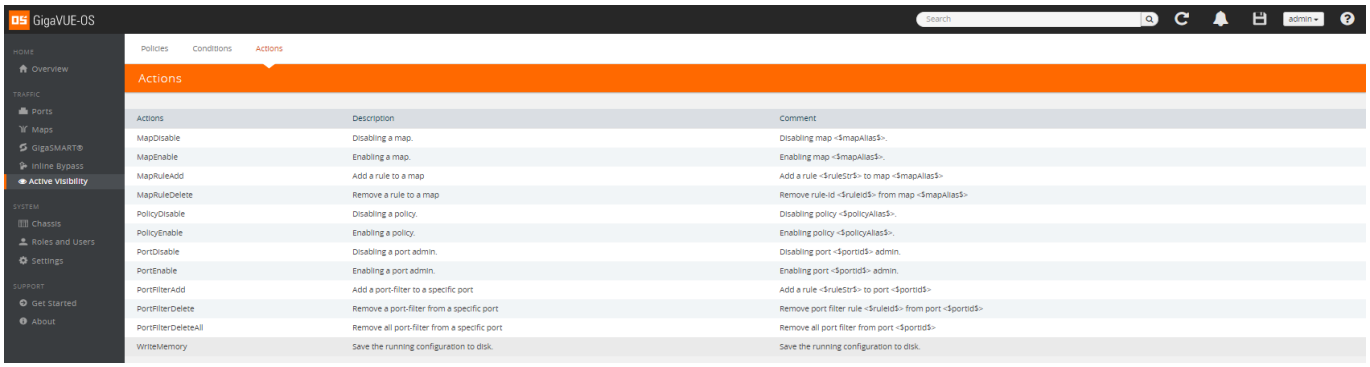

Figure 24-30: Actions, Description, and Comment

Configure Policies

An active visibility policy defines conditions and actions. When all conditions are met, all actions are executed. The policy ties the conditions and actions together.

A policy must have at least one condition. An action is not required.

Up to five (5) conditions can be specified in a policy. The parameters and values that are specified in the condition depends on the condition. The policy is executed only when all conditions are met.

Within a policy, there is one unique condition. For example, there can only be one **PortUp** condition, not multiple **PortUp** conditions within a policy.

Up to five (5) actions can be specified in a policy. The parameters and values that are specified in the action depends on the action.

Within a policy, there can be multiple actions, including the same action. For example, there can be multiple **PortEnable** actions within a policy.

A policy is triggered if all the conditions are met. Then all the actions are executed. If there are five conditions in the policy, they all have to be met before the action or actions are executed. If there are multiple actions, they are executed in sequence, as specified in the policy. The policy specifies the priority of the actions.

When there are multiple actions, they will continue to be executed even if there is an error. For example, if there are three actions and there is an error with the second action, the first and third actions will be executed. When the policy executes, each action will have their status reported (success or failure).

Multiple policies can be in effect at the same time. The policies can monitor different conditions and take different actions. Up to 100 policies per cluster can be defined.

A policy must be enabled for it to become active. However, when you first create a policy, you might want to disable it so that it does not become active right away.

When a policy is triggered, an SNMP event and email notification (policytrigger) can optionally be generated.

An active visibility policy defines conditions and actions. When all conditions are met, all actions are executed. The policy ties the conditions and actions together.

1. Go to **Active Visibility > Policies** to start configuring a policy.
2. Click **New**. Refer to Figure 24-31 on page 545.

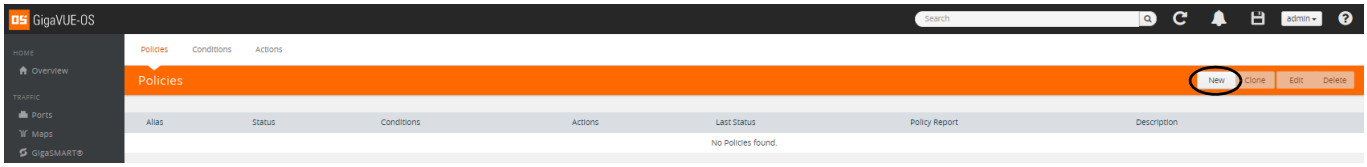

Figure 24-31: Policies, New

3. The Policy page has sections for Policy Info, Conditions, and Actions. Refer to Figure 24-32 on page 546.

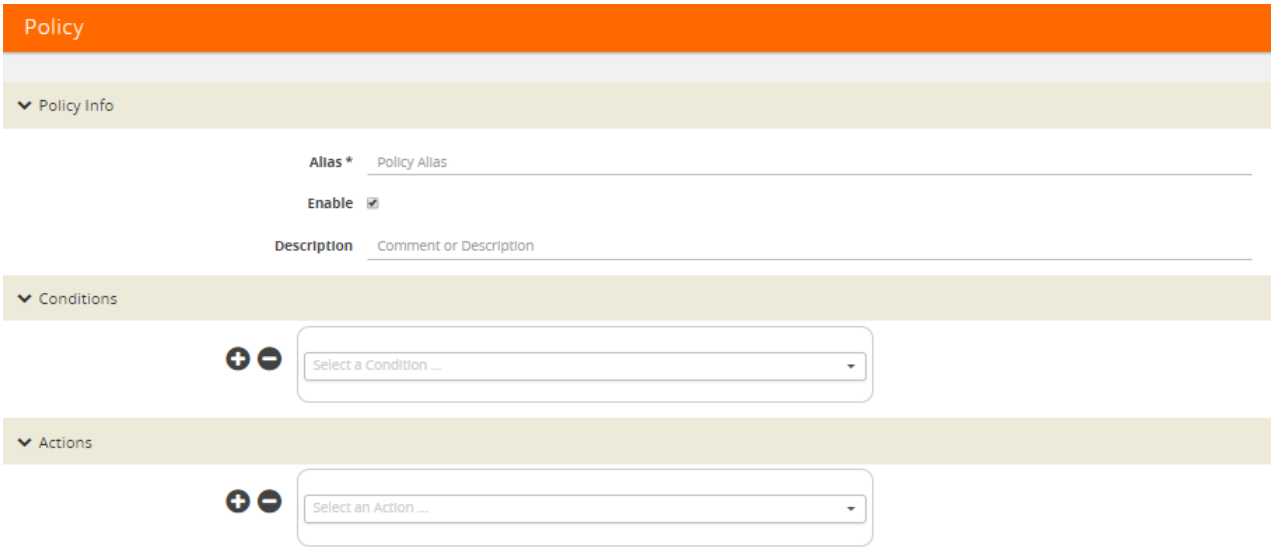

Figure 24-32: Policy Page

4. Enter an alias for the policy and a description, if desired. When you first define a policy, you will probably want the policy to be disabled, so clear the Enable checkbox. Refer to Figure 24-33 on page 546.

Figure 24-33: Policy Info on Policy Page

5. Select a condition from the drop-down menu. The parameters and values that are populated for each condition depend on the condition. For example, in Figure 24-34 on page 547, the condition is PortTxUtilLow, so there is an Any checkbox and fields for Port ID, Threshold (%), and Period (seconds). For details on configuring conditions, refer to *Configure Conditions* on page 550.

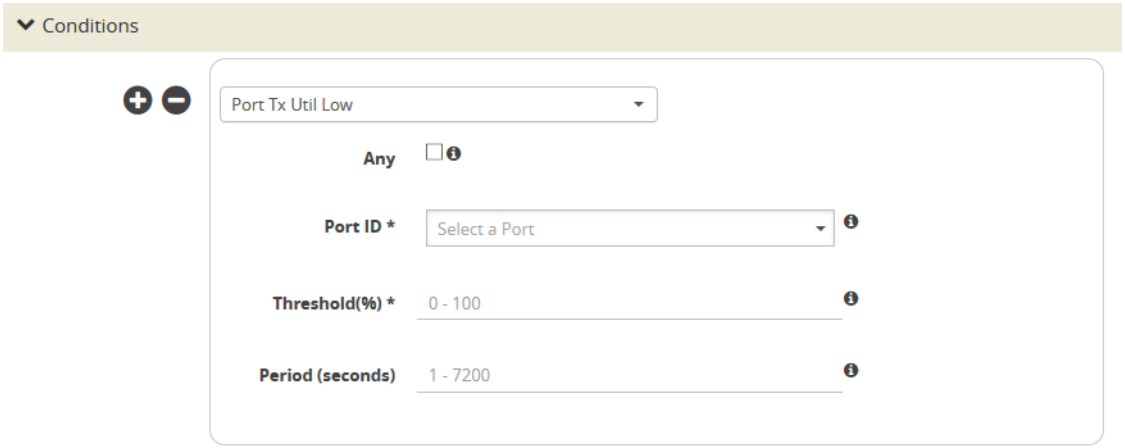

Figure 24-34: Conditions on Policy Page

6. Select or type the parameters and values for each condition. Refer to Figure 24-35 on page 547.

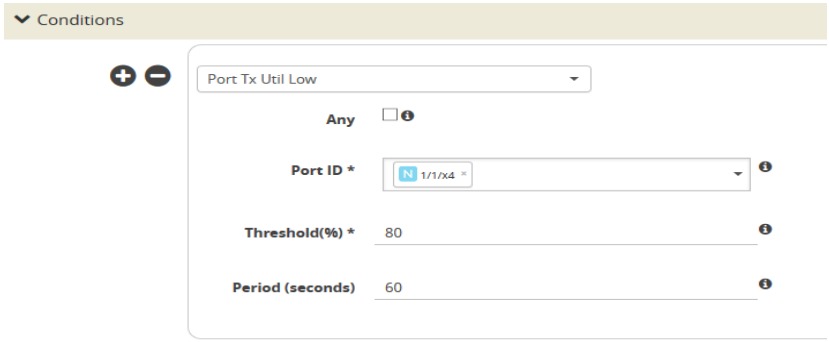

Figure 24-35: Conditions on Policy Page, Entered

7. For multiple conditions (up to five), click + or -.
8. Select an action from the drop-down menu. The parameters and values that are populated for each action depend on the action. For example, in Figure 24-36 on page 547, the action is Port Disable, so there is a Porting checkbox and a field for Port ID. For details on configuring actions, refer to *Configure Actions* on page 553.

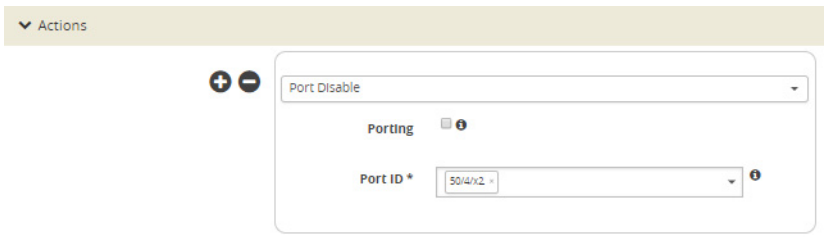

Figure 24-36: Actions on Policy Page

9. For multiple actions (up to five), click + or -.
10. When you have specified the conditions and actions, click **OK**.
11. The newly defined policy is displayed. Refer to Figure 24-37 on page 548.

Figure 24-37: Policy Defined

12. When you are ready to enable a policy, select the policy, click **Edit**, select the Enable checkbox, and click **OK**. Refer to Figure 24-38 on page 548.

Figure 24-38: Enable a Policy

13. Another way to enable a policy is to select the policy, click **Actions**, and select Enable. Refer to Figure 24-39 on page 548.

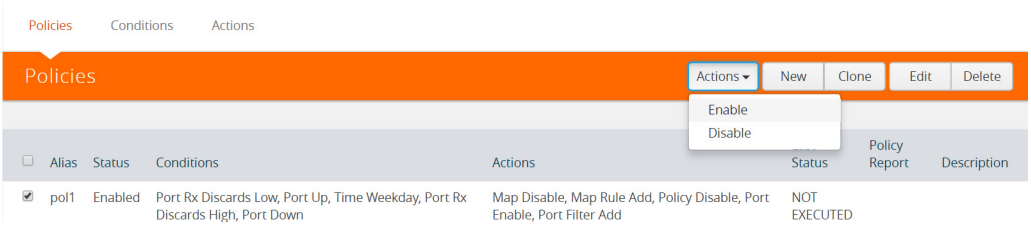

Figure 24-39: Using Action to Enable a Policy

14. When a policy has not been executed, it will show a status of NOT EXECUTED. When a policy has been executed, it will show a status of SUCCESS or FAILURE. Refer to Figure 24-40 on page 548.

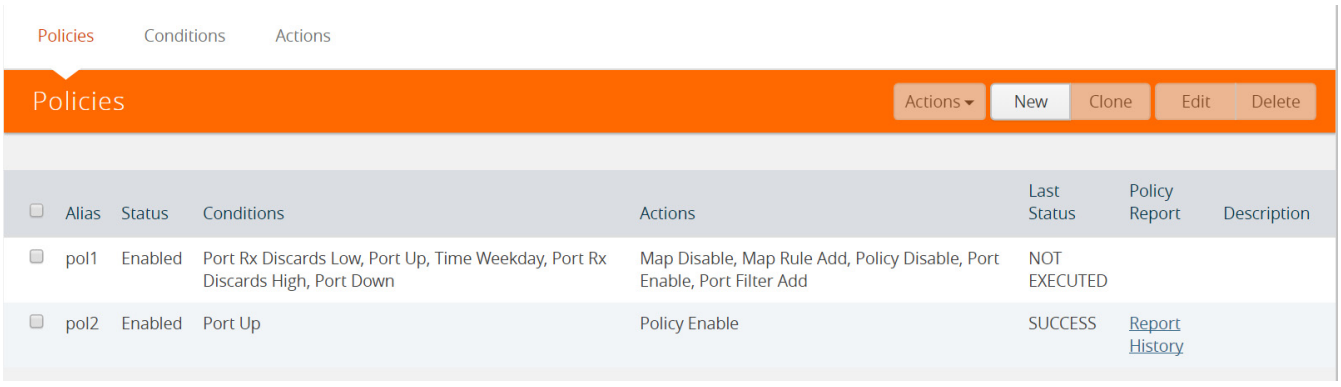

Figure 24-40: Policy Statuses

15. To view information on a policy that has executed, click **Report History**. The policy report history is displayed. Refer to Figure 24-41 on page 549.

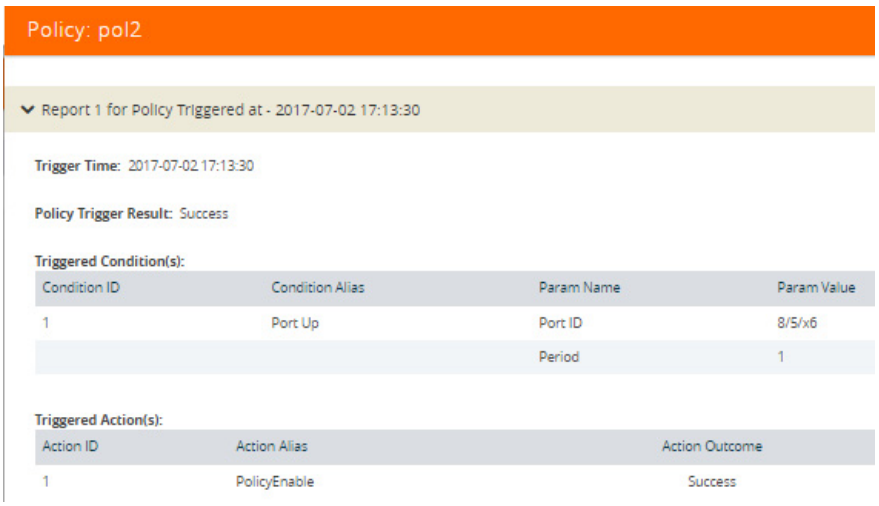

Figure 24-41: Policy Report History

16. To edit a policy, select a policy and click **Edit**. Refer to Figure 24-42 on page 549.

Figure 24-42: Edit a Policy

17. The Policy Info, Conditions, and Actions are displayed.
18. Make the changes and click **OK**.
19. To clone a policy, select a policy and click **Clone**. Refer to Figure 24-43 on page 549.

Figure 24-43: Clone a Policy

20. Change the policy alias, make any changes to conditions or actions, and click **OK**.
21. To delete a policy, select a policy and click **Delete**. Refer to Figure 24-44 on page 550.

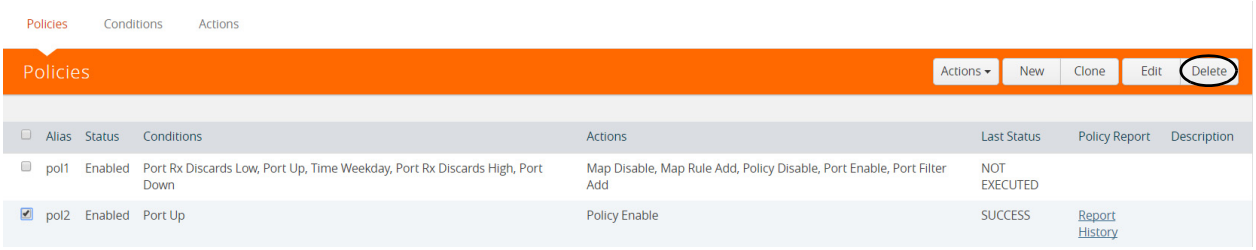

Figure 24-44: Delete a Policy

22. A confirmation displays. To confirm, click **OK**.
23. To view details of a policy, click a policy in the Policies page. Refer to Figure 24-45 on page 550.

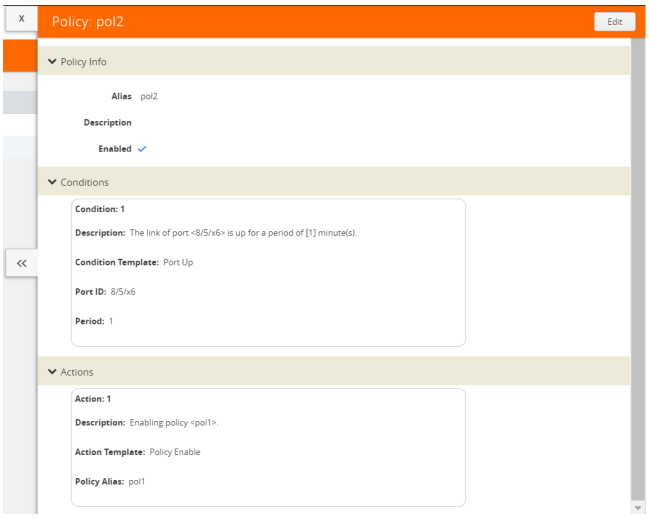

Figure 24-45: Policy Details

24. To edit the policy, click **Edit** or to exit, click **X**.

Configure Conditions

The parameters and values that are populated for each condition depend on the condition. Refer to the following sections for details:

- *Configure Port Up/Port Down Conditions* on page 551
- *Configure Port Rx/TxConditions* on page 551
- *Configure Time Conditions* on page 552
- *Configure Time of Day Conditions* on page 552

Configure Port Up/Port Down Conditions

To configure port up and port down conditions:

1. Select Port Up or Port Down. Refer to Figure 24-46 on page 551.

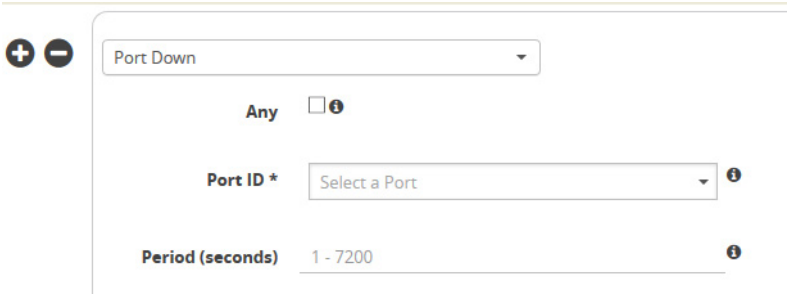

Figure 24-46: Port Up or Port Down Conditions

2. Specify the following:
   - Any—Click if any of the selected ports match the condition.
   - Port ID—Select one or more ports to match the condition.
   - Period—Specify the number of seconds to wait for the condition to be true, in integers from 1 to 7200.

Configure Port Rx/TxConditions

To configure port Rx/Tx conditions:

1. Select Port Rx or Port Tx conditions. Refer to Figure 24-47 on page 551.

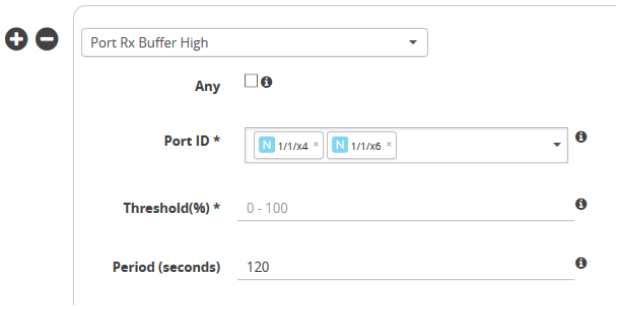

Figure 24-47: Port Rx or Tx Conditions

2. Specify the following:
   - Any—Click if any of the selected ports match the condition.
   - Port ID—Select one or more ports to match the condition.
   - Threshold—Specify the threshold as a percentage from 0 to 100% for Port Rx/Tx Buffer and Util, or as a value between 0 and 2^64, for example, 1024, for Port Rx/Tx Errors, Drops, and Discards.
   - Period—Specify the number of seconds to wait for the condition to be true, in integers from 1 to 7200.

Configure Time Conditions

To configure Time conditions:

1. Select a Time condition. Refer to Figure 24-48 on page 552.

Figure 24-48: Time Conditions

Configure Time of Day Conditions

To configure time of day conditions:

1. Select Time Of Day. Refer to Figure 24-49 on page 552.

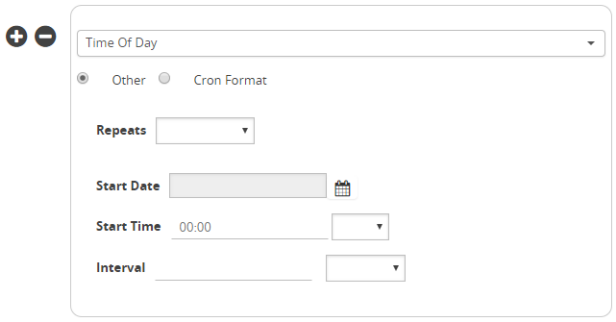

Figure 24-49: Time of Day Conditions

2. Click Cron Format. Refer to Figure 24-50 on page 552.

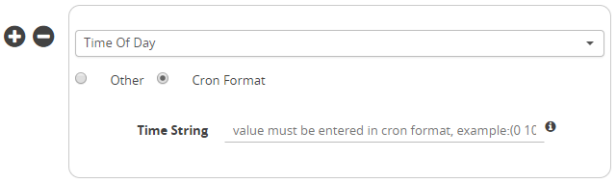

Figure 24-50: Time of Day Conditions, Cron Format

3. Enter the value in Cron format. For more information on Cron, refer to: http://www.nncron.ru/help/EN/working/cron-format.htm

4. Or, click Other. Refer to Figure 24-51 on page 553.

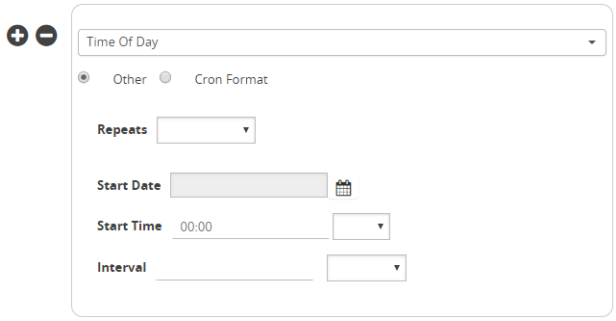

Figure 24-51: Time of Day Conditions, Other Format

5. The Other format allows you to specify the following:
   - Repeats—Specify Once Only, or a recurrence of Minute, Hourly, Daily, Weekly, Monthly, or Yearly.
   - Start Date—Specify the Start Date using the calendar.
   - Start Time—Specify the Start Time in hours and minutes, then select AM or PM.
   - Interval—For repeat occurrences, specify the interval. For occurrences longer than Hourly, specify the interval as Minute, Hour, or Day.
   - Every—For occurrences longer than Hourly, specify the number of days, weeks, or months. For example, if Repeats is Daily, Every can be every 2 days.
   - End Date—For occurrences longer than Hourly, specify an end date that is After a specified number of occurrences or By a date specified using the calendar. The default is None.
   - On—For Weekly repeats, specify the day of the week. For Monthly repeats, specify the First, Second, Third, Fourth, or Last day of the month, or specify the day of the month as a number from 1 to 31. The default is None.

Configure Actions

The parameters and values that are populated for each action depend on the action. Refer to the following sections for details:

- *Configure Port Actions* on page 554
- *Configure Policy Actions* on page 554
- *Configure Map Actions* on page 554
- *Configure Map Rule Add Actions* on page 555
- *Configure Map Rule Delete Actions* on page 555

Configure Port Actions

To configure port actions:

1. Select Port Disable or Port Enable. Refer to Figure 24-52 on page 554.

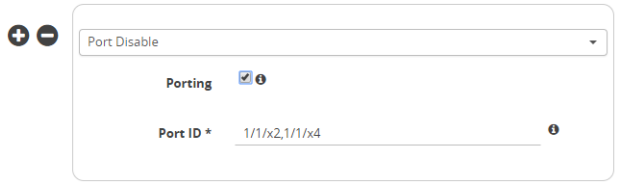

Figure 24-52: Port Actions

2. Specify the following:
   - Porting—Select if you want to pass the port, defined in a condition, to the action.
   - Port ID—Select one or more ports to match the action.

Configure Policy Actions

To configure policy actions:

1. Select Policy Disable or Policy Enable. Refer to Figure 24-53 on page 554.

Figure 24-53: Policy Actions

2. Select a policy alias.

Configure Map Actions

To configure map actions:

1. Select Map Disable or Map Enable. Refer to Figure 24-54 on page 554.

Figure 24-54: Map Actions

2. Select a map alias.

Configure Map Rule Add Actions

To configure map rule add actions:

1. Select Map Rule Add. Refer to Figure 24-55 on page 555.

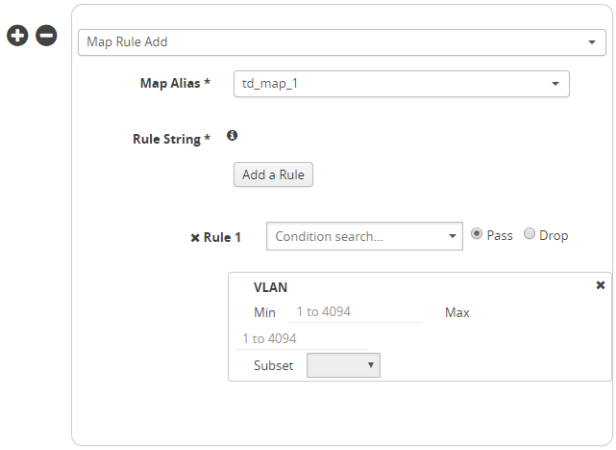

Figure 24-55: Map Rule Add Actions

2. Specify the following:
   - Map Alias—Select a map alias.
   - Add a Rule—Click to add one or more map rules.
   - Rule String—Specify a map rule, such as pass VLAN 100. The parameters and values that are populated for each rule depend on the rule.

Configure Map Rule Delete Actions

To configure map rule delete actions:

1. Select Map Rule Delete. Refer to Figure 24-56 on page 555.

Figure 24-56: Map Rule Delete Actions

2. Specify the following:
   - Map Alias—Select a map alias.
   - Rule ID—Specify a map rule ID.

GigaSECURE Security Delivery Platform

This section is an overview of the GigaSECURE® Security Delivery Platform (SDP), which transforms the way security applications are deployed and leveraged, enabling them to be more effective at protection and remediation, less complex, and more cost-effective. This section also provides links to detailed information for the pillars of GigaSECURE and to other documents.

GigaSECURE solutions provide pervasive visibility of network traffic, user, application, and suspicious activity, which bolsters security effectiveness, eliminates blind spots, and enables protection against threats. With GigaSECURE, you look inside networks to detect where compromise has occurred.

Figure 24-57 displays the pillars of the Security Delivery Platform along the bottom of the figure, consisting of GigaVUE Visibility Platform nodes (H Series and TA Series, as well as standalone and embedded TAPs) with GigaVUE-VM providing access to virtual traffic, and applications for NetFlow/IPFIX Generation, SSL Decryption, and Application Session Filtering (ASF), as well as inline bypass.

On the right of Figure 24-57, the Application Programming Interface (API) enables a degree of automation, allowing security tools to control the traffic feeds they receive from the SDP.

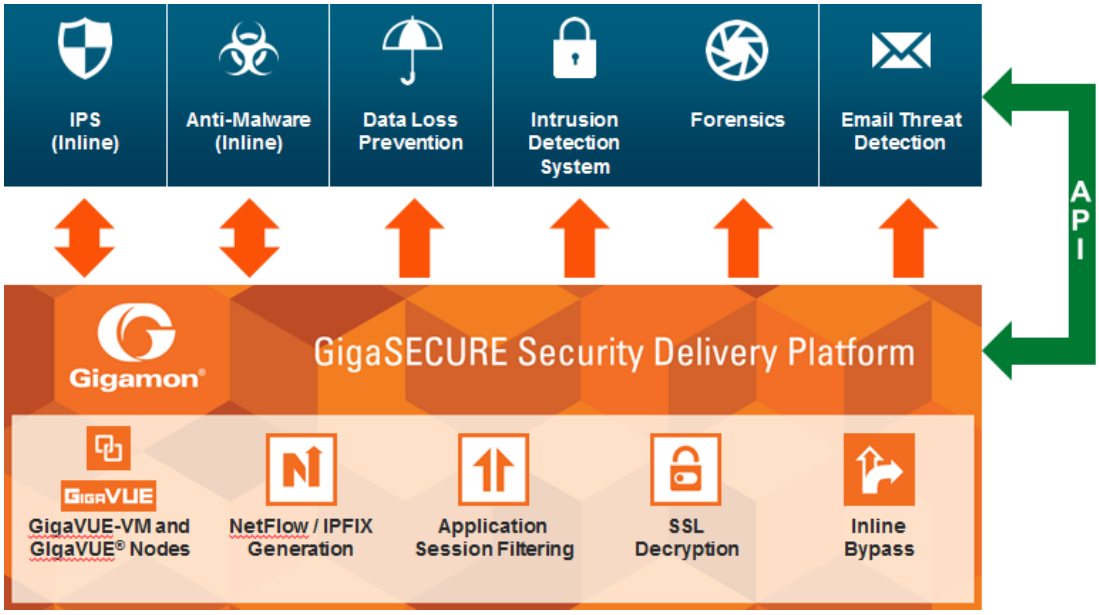

Figure 24-57: GigaSECURE Security Delivery Platform

GigaSECURE SDP connects into the network, across physical and virtual infrastructures, and delivers traffic to the applications that require it. Security tools connect into the SDP at whatever interface speeds they are capable of connecting. Consequently, they will receive a relevant traffic stream from across the network infrastructure.

Through SDP, organizations can unify their security applications (displayed along the top of Figure 24-57), eliminating blind spots and partial coverage. Security devices can

get a complete view of the network flows and the metadata that they require in order to detect threats or network compromise.

GigaSECURE Security Delivery Platform provides the following:

- a complete network-wide reach, physical and virtual. Refer to *GigaSECURE and GigaVUE nodes, TAPs, GigaVUE-VM and FM* on page 557.
- scalable metadata extraction for improved forensics. Refer to *NetFlow Generation* on page 558 under *GigaSECURE and GigaSMART Applications*.
- isolation of applications for targeting inspection. Refer to *Application Session Filtering* on page 559 under *GigaSECURE and GigaSMART Applications*.
- visibility to encrypted traffic for threat detection. Refer to *SSL Decryption* on page 560 under *GigaSECURE and GigaSMART Applications*.
- inline bypass for connected security applications. Refer to *GigaSECURE and Inline Bypass* on page 561.

GigaSECURE and GigaVUE nodes, TAPs, GigaVUE-VM and FM

To provide complete, network-wide reach, GigaSECURE consists of Visibility Platform nodes running GigaVUE-OS software, a virtualized node (GigaVUE-VM), and a centralized fabric controller (GigaVUE-FM). Traffic visibility is delivered from physical and virtual environments even when users, devices, and applications move around.

Refer to the following sections:

- *GigaVUE H Series and TA Series Nodes and TAPs* on page 557
- *GigaVUE-VM* on page 558
- *GigaVUE-FM and APIs* on page 558

GigaVUE H Series and TA Series Nodes and TAPs

GigaVUE H Series and TA Series are Visibility Platform nodes with the ability to cluster multiple nodes. Features include traffic aggregation, intelligent filtering, and the ability to replicate traffic to multiple security tools without performance impacts.

The GigaVUE TA Series provide a cost-effective way to provide scale-out traffic visibility. These nodes provide aggregation, filtering, and replication capabilities at a cost-effective price. This enables traffic from across the infrastructure to be channeled back to selective security tools.

Standalone TAPs or embedded TAPs are also available to TAP network traffic from 10Mb to 100Gb links.

GigaVUE H Series, TA Series, and Certified Traffic Aggregation White Boxes all run GigaVUE-OS software.

The GigaVUE H Series products are described in the following documents: *GigaVUE HD Series Hardware Installation Guide*, *GigaVUE-HC3 Hardware Installation Guide*, *GigaVUE-HC2 Hardware Installation Guide*, *GigaVUE-HC1 Hardware Installation Guide*, and *GigaVUE HB Series Hardware Installation Guide*.

The GigaVUE TA Series products are described in *GigaVUE TA Series Hardware Installation Guide*.

The Certified Traffic Aggregation White Box is described in *GigaVUE-OS Installation Guide on a White Box*.

TAPs are described in *G-TAP M Series Hardware Guide*.

GigaVUE-VM

GigaVUE-VM is a virtualized node that provides the ability to deliver traffic visibility into virtualized workloads. This enables a physical security tool to extend the security function to virtual traffic.

GigaVUE-VM also provides the ability to track virtual machines as they move from server to server, and enforce Follow-the-VM policies to ensure that application traffic is always sent to the security tools even if the VMs move.

GigaVUE-VM is described in the *GigaVUE-VM User's Guide*.

GigaVUE-FM and APIs

GigaVUE-FM serves as the centralized controller that provides the ability to unify the different components of the GigaSECURE Security Delivery Platform. It serves as a centralized policy definition point for the virtualized and physical Visibility Platform nodes.

GigaVUE-FM exposes a set of northbound APIs that allow security solutions to fine-tune in near real-time the traffic feeds that they are receiving, so as to adjust their visibility into the network infrastructure based on what real-time anomalies, threats, and conditions they are seeing.

GigaVUE-FM is described in the *GigaVUE-FM User's Guide*.

GigaVUE-FM APIs are described in *GigaVUE-FM REST API Getting Started Guide.*

GigaSECURE and GigaSMART Applications

The GigaSMART applications in the GigaSECURE® Security Delivery Platform provide the ability to act on traffic streams and perform a series of functions that serve to offload and optimize a variety of security solutions.

The three GigaSMART applications in the GigaSECURE Security Delivery Platform are NetFlow Generation, SSL Decryption, and Application Session Filtering (ASF).

NetFlow Generation

NetFlow generates detailed flow and session intelligence based on actual traffic, not just a sample of traffic.

IPFIX is a powerful standards-based technology that is gaining momentum in the network security space for forensics, trend analysis, and anomaly detection. IPFIX looks at raw network packets and derives sophisticated flow-based metadata such as records of conversations between endpoints, duration of conversations, and channels of communications.

GigaSECURE centralizes the function of generating these flow records so that this can be done consistently across heterogeneous and disparate infrastructure. The flow records can be served up to a variety of security solutions that analyze flow metadata. The flow metadata generation is done at very high throughput so as to generate high-fidelity records that are essential for good security analytics.

The solution also enables custom templates to be defined so that the information that can be gleaned from the traffic can be highly tailored to the specific deployment environment.

The GigaSECURE Security Delivery Platform with NetFlow Generation:

- provides unsampled NetFlow/IPFIX record generation to detect “low-and-slow” attacks
- filters records based on configurable parameters to predetermined tools
- offloads NetFlow/IPFIX record generation from the overloaded network infrastructure
- enables end-to-end security enforcement with visibility into every flow
- provides advanced information elements

NetFlow Generation is described in this document. Refer to *GigaSMART NetFlow Generation* on page 1077 for details.

Application Session Filtering

Application Session Filtering (ASF) with or without buffering provides the ability to deliver just the relevant traffic streams to the specific types of security tools. For example, an email security solution need not see YouTube traffic. Sending only relevant traffic allows the security solutions to function more effectively and waste less bandwidth and resources processing irrelevant information.

Many security solutions do not need to look at entire flows that are either trusted or that they have no ability to process. ASF provides the ability to look deep into the packet at the application layer, identify application flows based on patterns within the packets, and steer entire sessions to a specific security solution (for example, all packets belonging to a session, even if subsequent or preceding packets for that session do not match the pattern) or to discard the entire session.

This powerful capability allows precise control of the types of traffic data that are sent to security tools based on Layer 4 to Layer 7 and more sophisticated content matching, thereby ensuring that security solutions are focused on working off network traffic that is most relevant to them while simultaneously offloading those tools from having to process large volumes of irrelevant data.

Application Session Filtering is described in this document. Refer to *GigaSMART Application Session Filtering (ASF) and Buffer ASF* on page 1054 for details.

Also refer to the *Application Session Filtering Cook Book*.

SSL Decryption

There are two SSL Decryption applications as follows:

- *Out-of-Band SSL Decryption* on page 560
- *Inline SSL Decryption* on page 560

Out-of-Band SSL Decryption

SSL decryption for out-of-band tools provides a solution to decrypt encrypted communications so that security tools can detect malware that leverages encrypted communication channels and ensures that sensitive information is not compromised.

As the volume of malware that leverages encrypted communication channels increases, the need to peek into those encrypted channels of communication increases. Decrypting those encrypted channels of communication is best done within the GigaSECURE Security Delivery Platform so that this is done once, at very high performance thereby eliminating this blind spot simultaneously for multiple security tools that do not have the ability to deal with encrypted communications. For those security tools that have the ability to do this, it offloads a computationally intensive task from being repetitively done in each security tool.

Out-of-band SSL decryption delivers decrypted traffic to out-of-band tools that can then detect threats entering the network. When a threat is detected, the tools can send a notification.

Out-of-band SSL decryption is described in this document. Refer to *GigaSMART Out-of-Band SSL Decryption* on page 1169 for details.

Inline SSL Decryption

SSL decryption for inline tools provides visibility into encrypted traffic. Inline SSL decryption delivers decrypted packets to tools that can be placed inline or out-of-band. The tools look into decrypted packets for threats, such as viruses or other malware.

The amount of Internet traffic that is encrypted is increasing, and much of it is encrypted with Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols.

Malware increasingly uses encrypted SSL traffic, thus a significant percentage of attacks hide in SSL. Inline SSL decryption delivers a complete view of encrypted applications and hidden threats in your organization.

Many applications, such as email, also use SSL. Encryption protects data from being viewed in transit over the Internet such as in an exchange of emails. Encryption also keeps the data private. But when data is encrypted, packets are not inspected, which can create blind spots in your network.

Providing visibility into encrypted traffic eliminates this blind spot. SSL/TLS blind spots in your network can be eliminated across any port or application, for example, port 443, or email, Web, or VoIP applications.

Inline SSL decryption inspects SSL encrypted traffic inline. The advantage of this solution is that when SSL decryption is inline, tools can act when a threat is detected.

Inline SSL decryption is described in the *Inline SSL Decryption Guide*.

GigaSECURE and Inline Bypass

Inline bypass in the GigaSECURE® Security Delivery Platform supports inline and out-of-band network security deployments from the same platform. It provides the ability to load balance both inline and out-of-band security tools as well as to bypass inline security tools in the event of failure.

Many security tools work inline with the network traffic to prevent malware and malicious activities in real-time. Many other security tools work out-of-band for detection and incident generation purposes. The GigaSECURE Security Delivery Platform provides a common platform to serve traffic feeds to both inline and out-of-band security deployments.

When serving inline security deployments, the GigaSECURE platform provides the ability to load balance traffic across multiple inline security solutions, as well as the ability to guide traffic serially to different inline security tools, each providing different levels of protection. Traffic can be distributed to the security tools based on a variety of criteria, while ensuring that forward and reverse traffic for a given flow always goes to the same security tool.

The platform also provides resiliency and protection in the event that any of the inline security tools experiences a failure, both in load balanced mode as well as when inline tools are connected in a serial fashion, thereby ensuring that network traffic forwarding is not disrupted in the event of a failure.

Security tools can also be moved from out-of-band to inline and vice versa with no disruption to the network.

The GigaSECURE Security Delivery Platform with inline bypass:

- maximizes tool efficacy
- increases scale of security monitoring
- provides seamless add, remove, and upgrade of tools
- consolidates multiple points of failure into a single, bypass-protected solution
- integrates inline, out-of-band, and flow-based tools

Inline bypass is described in this document. Refer to *Configure Inline Bypass Solutions* on page 573 for details.

25 Inline Bypass Solutions

This chapter provides the following information about inline bypass solutions:

About Inline Bypass Solutions

Security tools such as firewalls and intrusion protection systems (IPSs) are often connected inline on production networks, with traffic flowing from the network segment through the tool and then back onto the production network.

Inline bypass solutions involve bidirectional traffic between two networks, intercepted by a Gigamon node, and guided through one or more inline tools.

Inline bypass is a pillar of the GigaSECURE Security Delivery Platform. For an overview of GigaSECURE, refer to *About Inline Bypass Solutions* on page 563.

Inline bypass is supported on all GigaVUE HC Series nodes: GigaVUE-HC3, GigaVUE-HC2, and GigaVUE-HC1.

This section describes inline bypass solutions. Refer to the following sections for details:

NOTE: The configuration of inline bypass solutions can be complex. Follow the order of configuration outlined in *Configuration Steps* on page 617 and demonstrated in *Inline Bypass Solution Examples* on page 633.

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-HC2, and GigaVUE-HC1 (bypass combo modules equipped with optical protection switches)
- Supports physical bypass on GigaVUE-HC2 and GigaVUE-HC1 (TAP modules equipped with electrical relays).

- Distributes traffic across multiple inline tools, providing load sharing and traffic redistribution should a tool fail.
- Supports out-of-band maps for selective forwarding of inline traffic to monitoring tools.
- Supports sharing of inline tools by multiple inline networks.
- Supports N+1 and 1+1 inline tool redundancy for inline tool groups.
- Supports guiding inline traffic through inline tools in a serial fashion.
- Supports inline flow mapping through rule-based and shared collector inline maps.
- Provides a high availability solution, Gigamon Resiliency for Inline Protections (GRIP).

Logical Bypass and Physical Bypass

Logical bypass lets traffic bypass the inline tool should it experience a failure. A failure is declared if the GigaVUE node either loses connectivity to the inline tool or fails to receive a heartbeat from the tool.

Logical inline bypass does not require any specialized hardware and can be facilitated on regular ports on GigaVUE-HC3, GigaVUE-HC2, and GigaVUE-HC1 modules. On GigaVUE-HC1, logical inline bypass can be configured on the base module.

On the GigaVUE-HC2, fiber TAP modules and copper TAP modules (except TAP-HC0-G100C0) cannot be used to configure logical inline bypass.

Physical bypass provides protection against failure of the GigaVUE node, such as if power is lost. On GigaVUE HC Series nodes, it is implemented with specialized hardware called bypass combo modules.

The specialized hardware triggers a bypass when power is lost to the node. When the physical bypass is activated, traffic flows from one side of the network to the other, but without monitoring.

The GigaVUE-HC3 bypass combo module is a bypass switch (BPS) module, as follows:

- Bypass Combo Module with two 100Gb/40Gb SR4 MPO inline network port pairs and sixteen regular SFP+ (10Gb) port cages (BPS-HC3-C25F2G)

The GigaVUE-HC2 bypass combo modules are bypass switch (BPS) modules, as follows:

- Bypass Combo Module with four SX/SR (50/125µm multimode) inline network port pairs and sixteen regular SFP+ (1Gb/10Gb) port cages (BPS-HC0-D25A4G)
- Bypass Combo Module with four SX/SR (62.5/125µm multimode) inline network port pairs and sixteen regular SFP+ (1Gb/10Gb) port cages (BPS-HC0-D25B4G)
- Bypass Combo Module with four LX/LR (singlemode) inline network port pairs and sixteen regular SFP+ (1Gb/10Gb) port cages (BPS-HC0-D35C4G)
- Bypass Combo Module with two 40Gb SR4 (multimode) inline network port pairs and eight regular SFP+ (1Gb/10Gb) port cages (BPS-HC0-Q25A28)

NOTE: The 40Gb BPS module, BPS-HC0-Q25A28, is only supported on GigaVUE-HC2 with Control Card version 2 in this release.

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)

Physical bypass is also supported on the following copper TAP modules on GigaVUE-HC2 and GigaVUE-HC1:

- On the GigaVUE-HC2 copper TAP module, TAP-HC0-G100C0
- On the GigaVUE-HC1 copper TAP module, TAP-HC1-G10040

The GigaVUE-HC3, GigaVUE-HC2, and GigaVUE-HC1 offer physical and logical inline bypass. Physical bypass provides automatic failover protection in the case of a power failure. On the GigaVUE HC Series nodes, the bypass combo modules provide the physical bypass function. Or on GigaVUE-HC2 and 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 bypass combo modules or the TAP-HC0-G100C0 module on GigaVUE-HC2, refer to the *GigaVUE-HC2 Hardware Installation Guide*. This document also contains procedures for installing, removing, and replacing modules in the GigaVUE-HC2.

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.

Refer also to Figure 25-1 and *Protected Inline Network* on page 568.

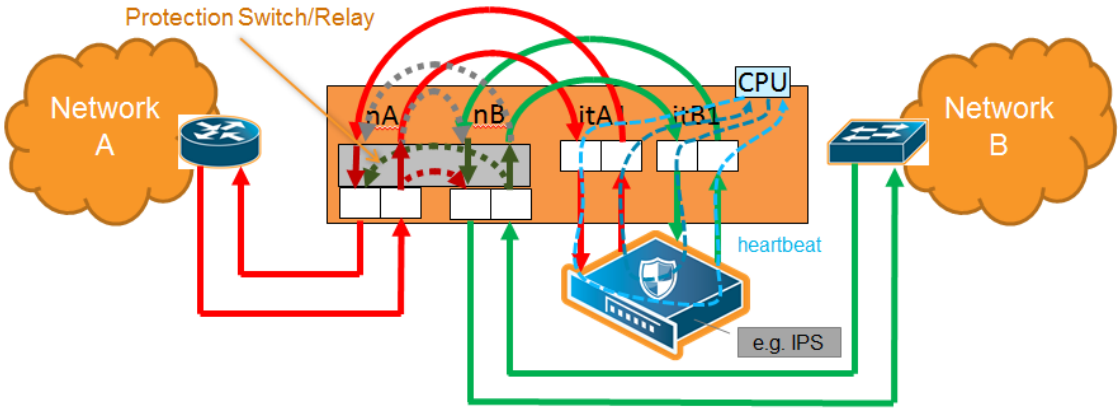

Figure 25-1: Physical Bypass Protection

Types of Inline Networks

The types of inline networks are described in the following sections:

- *Unprotected Inline Network* on page 567
- *Protected Inline Network* on page 568
- *Mix of Protected and Unprotected* on page 568

Unprotected Inline Network

An unprotected inline network consists of two ports of the inline-network type, which facilitate access to a bidirectional link between two networks (or more precisely, two far-end network devices).

Any available network type ports on a GigaVUE HC Series node can be configured to be inline-network type ports and combined to form an unprotected inline network. An unprotected inline network is shown in Figure 25-2.

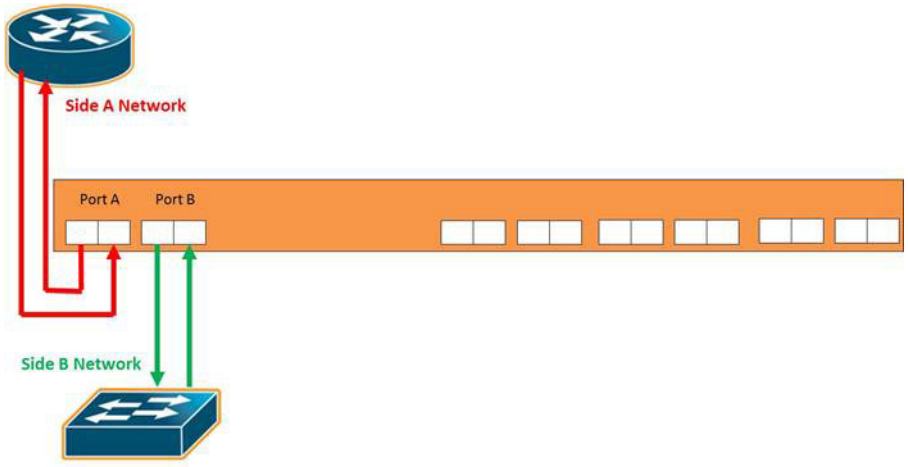

Figure 25-2: Unprotected Inline Network

Protected Inline Network

A protected inline network uses the port pairs associated with optical protection switches located on bypass combo modules on GigaVUE-HC3, GigaVUE-HC2, and GigaVUE-HC1. The specialized hardware triggers a bypass when power is lost to the node. On GigaVUE-HC2, protected inline networks for copper are supported on the TAP-HC0-G100C0 module. On GigaVUE-HC1, protected inline networks for copper are supported on the TAP-HC1-G10040 module.

Physical bypass provides protection against failure of the Gigamon node due to a loss of power. When physical bypass is activated, traffic flows from one side of the network to the other, but without monitoring.

Because of the specialized hardware, when a bypass combo module initializes, the system creates inline-network type ports for each of the protected port pairs on the module. This does not happen automatically on the TAP-HC0-G100C0 module on the GigaVUE-HC2 or on the TAP-HC1-G10040 module on the GigaVUE-HC1.

A protected inline network is shown in Figure 25-3.

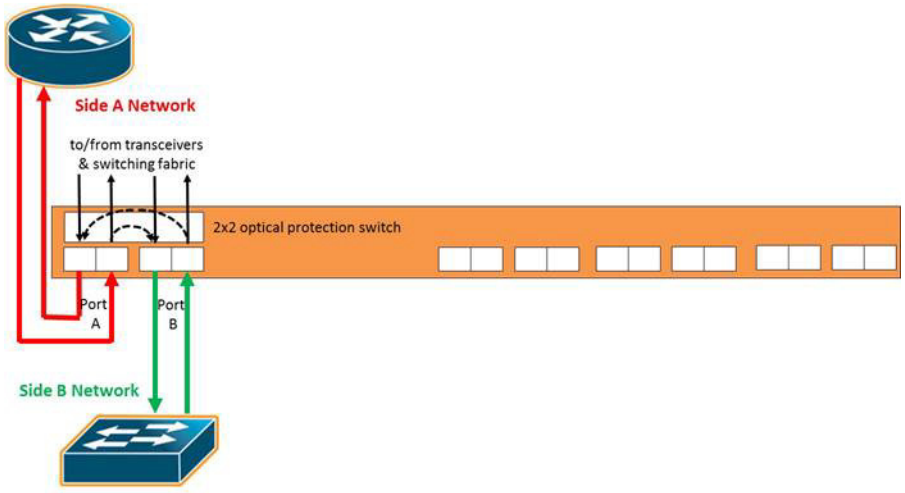

Figure 25-3: Protected Inline Network

Mix of Protected and Unprotected

Any combination of protected and unprotected inline networks is supported in inline bypass solutions.

Simple and Complex Inline Bypass Solutions

Simple and complex inline bypass solutions are described in the following sections:

- *Typical Configuration* on page 569
- *Distribution to Multiple Inline Tools* on page 569
- *Inline Tools in a Series* on page 570
- *Multiple Inline Networks* on page 571

- *Inline Flow Mapping* on page 571
- *Send Traffic to Out-of-Band Tools* on page 572

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 25-4 on page 569. The physical protection switch is optional and present only when using bypass combo modules or copper TAP modules.

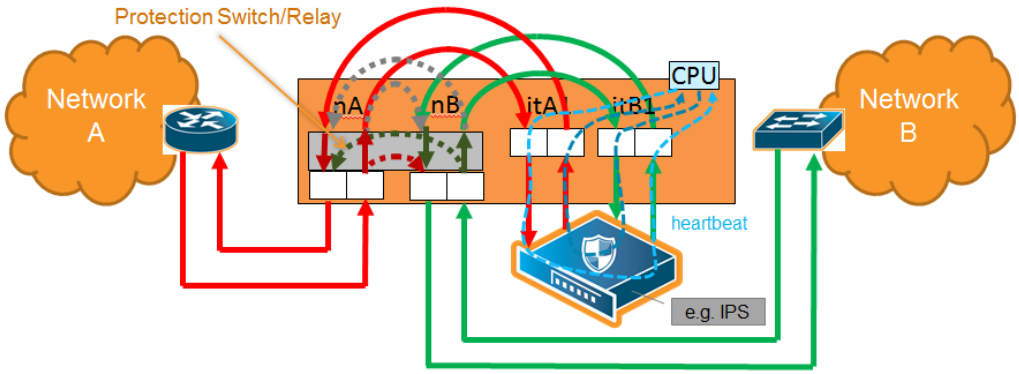

Figure 25-4: Simple Configuration

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 25-5.

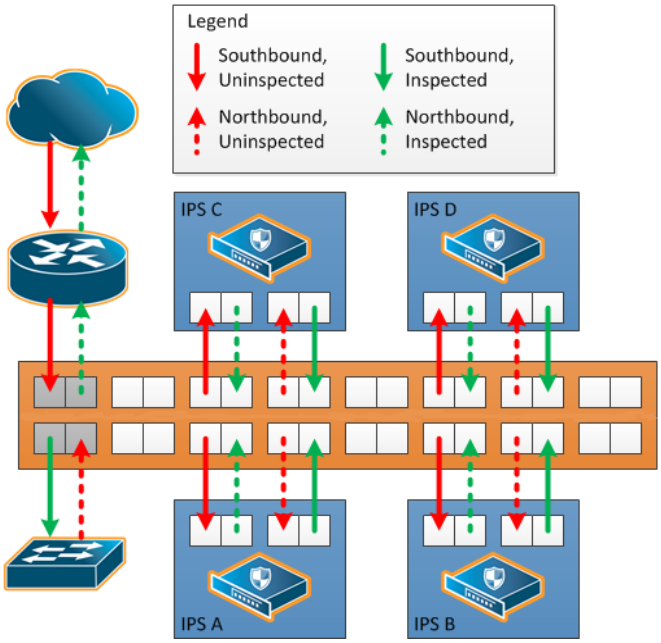

Figure 25-5: Multiple Tool Arrangement

Multiple tool distributions can be non-redundant or have 1+1 or N+1 redundancy. *About Inline Tool Groups* on page 595

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 25-6 for an inline tool series. In Figure 25-6, traffic is only shown from A-to-B.

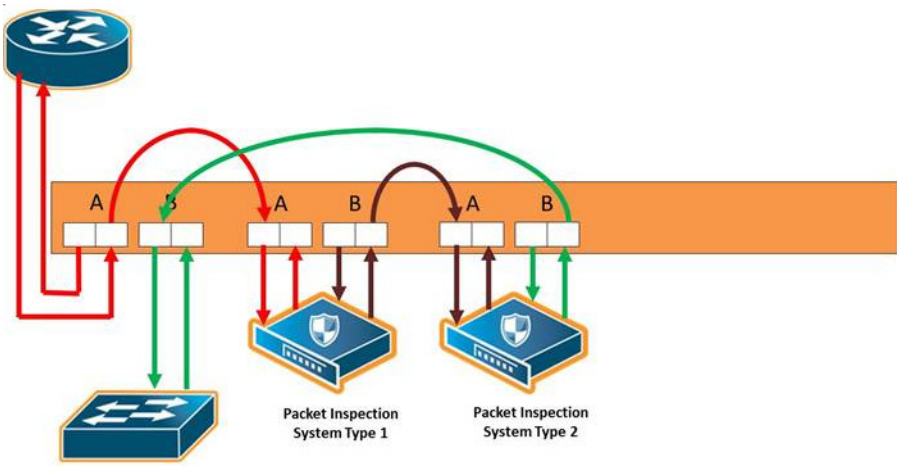

Figure 25-6: Inline Tool Series

For more information on inline tool series, refer to *Inline Tools in a Series* on page 570.

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 25-7 or many-to-many. Traffic is guided to a particular inline network through internal VLAN ID tagging.

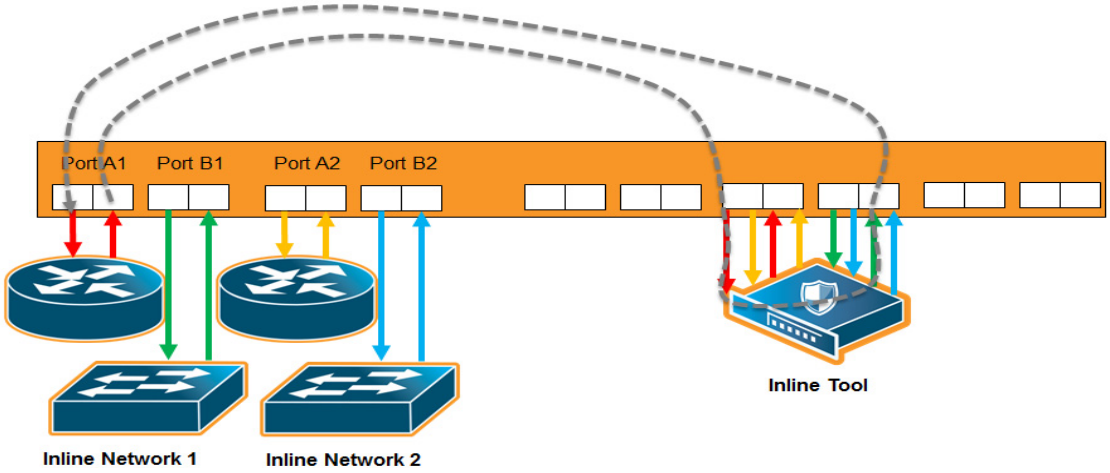

Figure 25-7: Inline Network Group

For more information on inline network groups, refer to *Configure Inline Bypass Solutions* on page 573.

Inline Flow Mapping

Some tools are optimized for particular inline traffic. With inline flow mapping, the GigaVUE node forwards packets to different inline tools based on criteria, such as TCP/UDP port number, or any other rule that can be defined in a map. Using these map rules, selected traffic can be sent to specific tools.

When inline flow mapping is combined with distribution to multiple tools, one application, such as Web traffic, can be sent to one tool or group of tools, while another application, such as email traffic, can be sent to another tool or group of tools. If there is traffic, such as encrypted traffic, that does not need to be or cannot be inspected, it can be bypassed.

An inline flow mapping based traffic distribution is shown in Figure 25-8. In Figure 25-8, traffic is only shown from A-to-B.

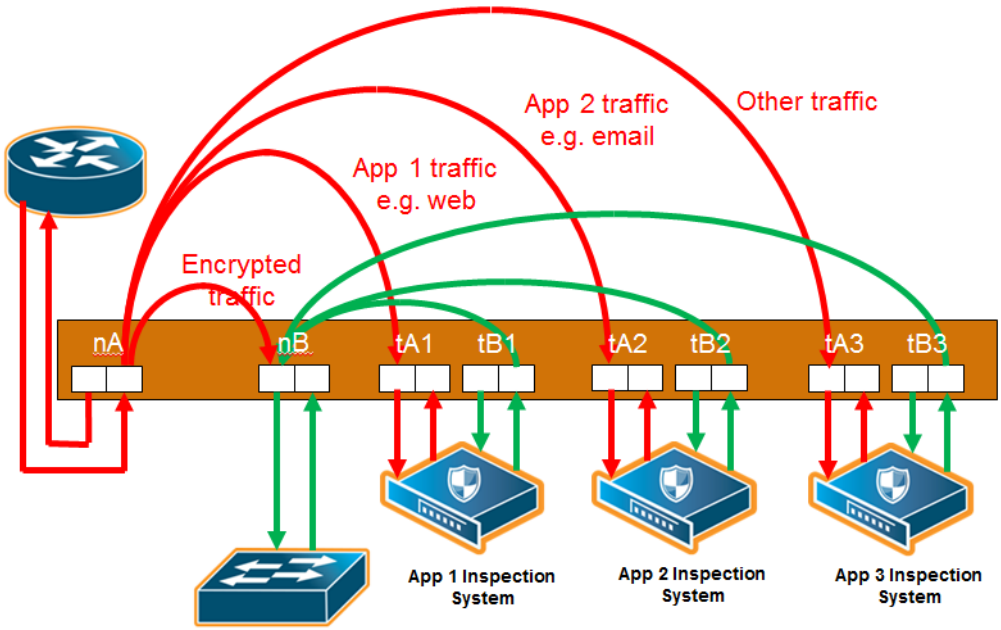

Figure 25-8: Inline Flow Mapping Based Traffic Distribution

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 25-9.

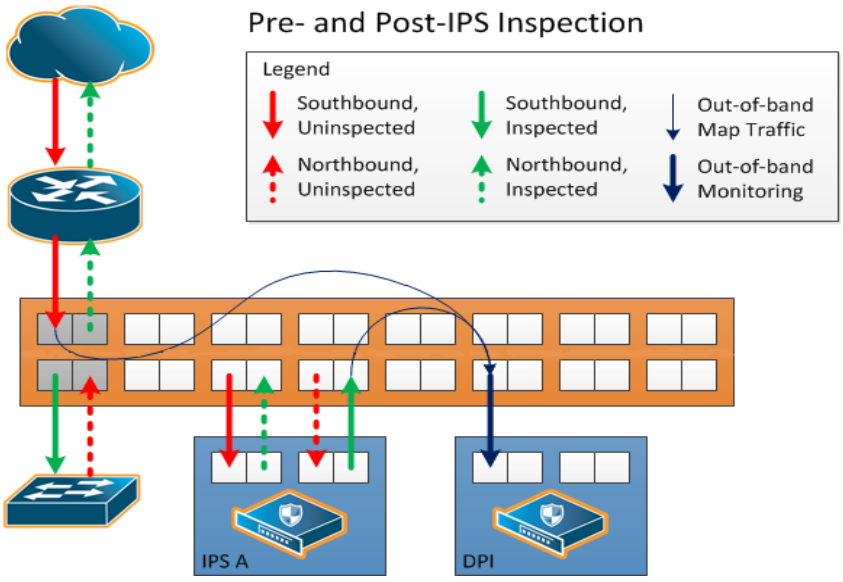

Figure 25-9: Out-of-Band Arrangement

Configure Inline Bypass Solutions

Inline bypass solutions are arrangements of inline networks that act to guide traffic through one or more inline tools attached to a GigaVUE node.

Configure inline bypass solutions using the following software constructs:

- Inline network

  For configuring inline networks, which are pairs of network ports attached to the sides of the network, refer to *About Inline Networks* on page 574.

- Inline network groups

  For configuring inline network groups, which are multiple inline networks sharing the same inline tool or tools, refer to *About Inline Network Groups* on page 583.

- Inline tool

  For configuring inline tools, which are pairs of inline tool ports plus the inline tools attached to the ports, refer to *About Inline Tools* on page 586.

- Heartbeat profile

  For configuring monitoring of inline tools using heartbeat and/or negative heart beat packets or both, refer to *About Heartbeat Profiles* on page 590.

- Inline tool group

  For configuring multiple inline tools, which are arrangements of multiple inline tools to which traffic is distributed based on hashing, refer to *About Inline Tool Groups* on page 595.

- Inline serial tools

  For configuring inline tools in a series, in which traffic is guided through the inline tools in a serial fashion. Refer to *About Inline Serial Tools* on page 601.

The steps to configure inline constructs are only applied to GigaVUE HC Series nodes. In a cluster environment, these steps are only applied to GigaVUE HC Series nodes through the cluster master.

The inline constructs must all be configured on one GigaVUE-HC3,GigaVUE-HC2, 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* on page 610.

For out-of-band maps, refer to *Out-of-Band (OOB) Map* on page 615.

For the order in which to configure the software constructs associated with inline bypass solutions, refer to *Configuration Steps* on page 617.

About Inline Networks

Currently configured in-line networks are displayed on the Inline Networks page, which is opened by selecting **Inline Bypass > Inline Networks**.

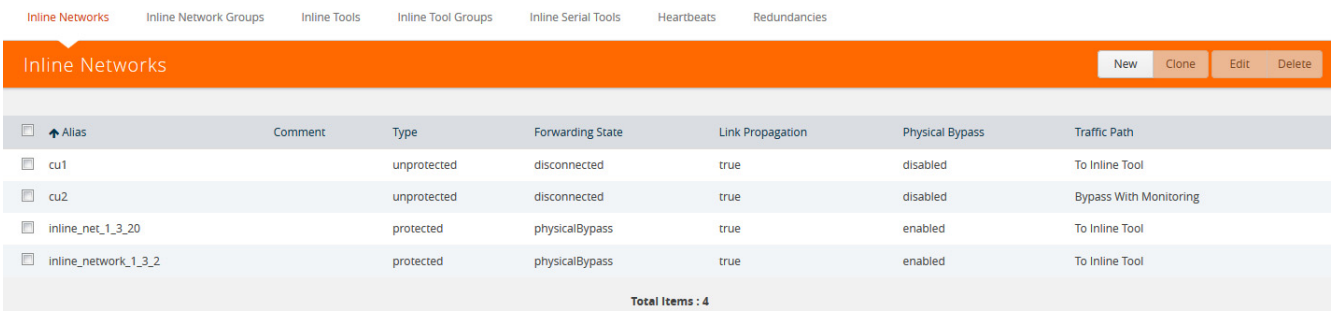

Figure 25-10: Summary Page of Inline Network Ports

Click **New** to open the **Inline Network** configuration page to configure an inline network, which is a pair of network ports arranged for inline monitoring. The arrangement facilitates access to a bidirectional link between the two networks (two far-end network devices) that need to be linked through an inline tool.

An inline network consists of inline network ports, always in pairs, running at the same speed, on the same medium (either fiber or copper). The inline network ports must be on the same GigaVUE-HC3, GigaVUE-HC2, or GigaVUE-HC1 node.

An inline network can be unprotected or protected. On the GigaVUE HC Series nodes, protected inline networks are implemented using bypass combo modules. Protected inline networks are based on the pairs of ports associated with physical protection switches on the bypass combo modules. The protected inline network ports provided by the bypass combo modules offer link speeds that can be configured at either 1Gb or 10Gb.

Starting in software version 5.0, the GigaVUE-HC2 with Control Card version 2 supports a 40Gb bypass combo module and starting in software 5.1.01, the GigaVUE-HC3 supports a 100Gb bypass combo module and starting in software version 5.2, the GigaVUE-HC3 100Gb bypass combo module supports dual speeds: 100Gb or 40Gb.

NOTE: You do not need to create the inline network ports for the protected port pairs on the bypass combo modules because they are created automatically when the bypass combo module initializes.

Other ports on the bypass combo module or on other line cards or modules can be designated as inline-network type ports. These ports will automatically be unprotected.

You also do not need to create the protected inline networks on the bypass combo modules because they are created automatically when the bypass combo module initializes. In addition, you cannot delete the protected inline networks on the bypass combo modules.

On GigaVUE-HC1, inline networks can also be configured on the copper TAP-HC1-G10040 module. On the GigaVUE-HC2, inline networks can be configured on the copper TAP-HC0-G100C0 module.

Once an inline network is created, it is fully configurable regardless of whether or not it is associated with an inline tool or an inline tool group through a map or map passall.

Use the Inline Networks page to display the configuration of inline networks as well as the current state of the inline bypass solution. The overall state of the inline bypass solution consists of the inline networks involved in the solution as well as the associated inline maps and inline tools. For details, refer to the following:

- *Display Current State of Inline Bypass Solution* on page 579
- *How to Use SNMP Polling to Obtain Inline Network State* on page 582
- *SNMP Notification of Forwarding State Change* on page 582
- *How to Use Syslog to Obtain Inline Network State* on page 582

For details on the parameters of inline networks, refer to the following:

- *Network Port Link Status Propagation Parameter* on page 575
- *Traffic Path Parameter* on page 576
- *Physical Bypass Parameter* on page 578.

Network Port Link Status Propagation Parameter

One of the parameters of inline networks is link status propagation, which controls the behavior of the link status for the inline network ports involved in a given inline network. The default is enabled.

When enabled, an inline network link failure on one side of the inline network will be propagated to the other side. For example, when the link is lost on one side of the network such that traffic cannot be sent to the inline tools, the link on the opposite side of the network is also brought down.

When the link is restored to the side that originally went down, the link will automatically be restored to the other side of the network. The GigaVUE node will not forward packets to the inline tools until the link is restored on both sides.

Link status propagation is enabled by selecting **Link Failure Propagation** when configuring an inline network port.

Traffic Path Parameter

One of the parameters of inline networks is **Traffic Path**, which specifies the path of the traffic received at an inline network port. The values are as follows:

- **Drop**—no traffic is exchanged through the inline network ports. All traffic to these ports is dropped. No traffic is forwarded to or from the inline tool or tools. No traffic is passed from inline network port A to inline network port B or from inline network port B to inline network port A. Refer to Figure 25-11 on page 577.
- **ByPass**—there are two cases for bypass as follows, depending on the inline map configuration:
  - If there are no inline maps associated with the inline network or if the set of inline maps associated with the inline network guarantees that no traffic is dropped when the traffic path is set to **To Inline Tool**, then setting the traffic path to **Bypass** leads to the following: all traffic arriving at the side A inline network port is forwarded to the side B inline network port and all traffic arriving at the side B inline network port is forwarded to the side A inline network port through a logical bypass. Refer to Figure 25-12 on page 577 for a simple scenario involving a map-passall in which there is no possibility of traffic drops.
  - If the set of inline maps associated with the inline network involves some traffic drop when the traffic path is set to **To Inline Tool**, then setting the traffic path to **Bypass** leads to the following: all traffic arriving at the side A inline network port that would not have been dropped with traffic path set to **To Inline Tool** is forwarded to the side B inline network port and all traffic arriving at the side B inline network port that would not have been dropped with traffic path set to **To Inline Tool** is forwarded to the side A inline network port through a logical bypass. In other words, packets that were dropped when the traffic path was set to **To Inline Tool** will also be dropped when set to **Bypass**.

  In either of these two bypass cases, no traffic is forwarded to the inline tool or tools.
- **ByPass with Monitoring**—there are two cases for monitoring as follows, depending on the inline map configuration:
  - If there are no inline maps associated with the inline network or if the set of inline maps associated with the inline network guarantees that no traffic is dropped when the traffic path is set to **To Inline Tool**, then setting the traffic path to **ByPass with Monitoring** leads to the following: all traffic is forwarded as for bypass, but a copy of the traffic is forwarded to the inline tool or tools according to the configured maps between the inline network and the inline tool or tools. Refer to Figure 25-13 on page 578 for a simple scenario involving a map-passall in which there is no possibility of traffic drops.
  - If the set of inline maps associated with the inline network involves some traffic drop when the traffic path is set to **To Inline Tool**, then setting the traffic path to **ByPass with Monitoring** leads to the following: all traffic that would not have been dropped with traffic path set to **To Inline Tool** is forwarded as for bypass, but a copy of the traffic is forwarded to the inline tool or tools according to the configured maps between the inline network and the inline tool or tools.

  In either of these two monitoring cases, no traffic is taken from the inline tools.
- **To Inline Tool**—traffic received at the inline network ports is forwarded according to the following factors:

- the configured maps between the inline network and the inline tools
- the failover actions of the inline tool or tools
- the health state of the inline tool or tools

Refer to Figure 25-14 on page 578.

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 25-11 to Figure 25-14 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.

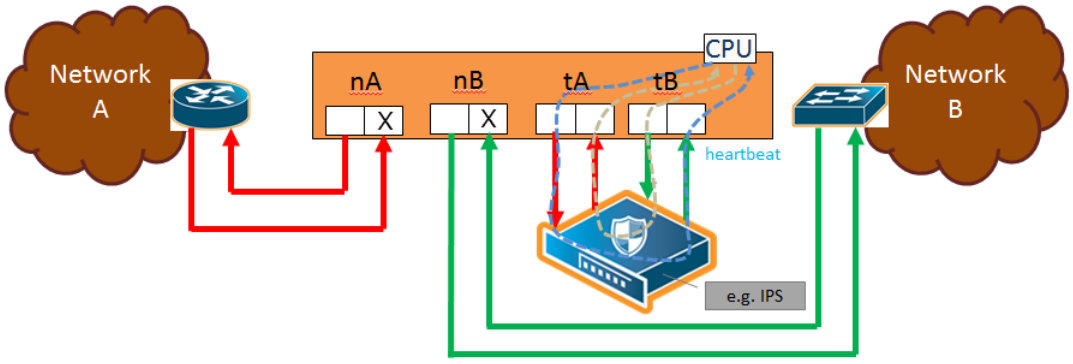

Figure 25-11: Traffic Path of Drop

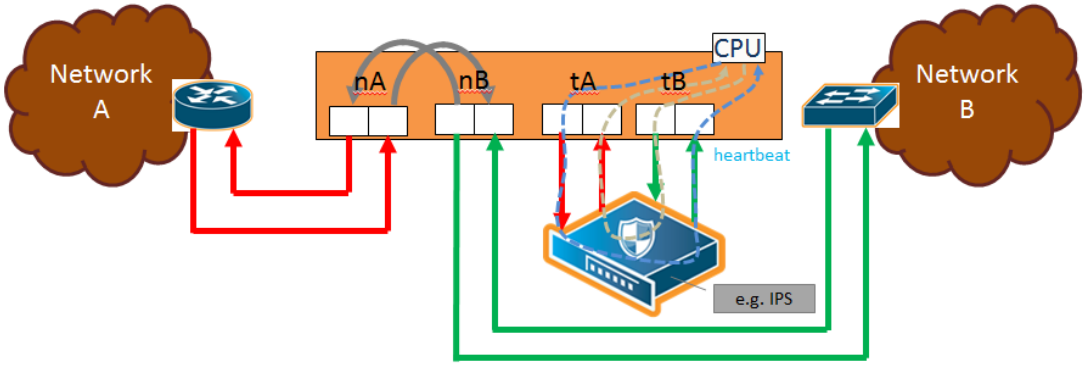

Figure 25-12: Traffic Path of Bypass

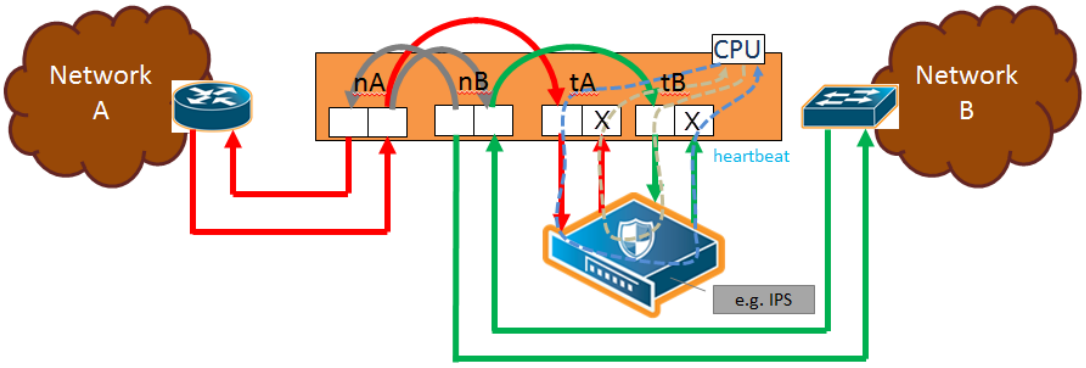

Figure 25-13: Traffic Path of Monitoring

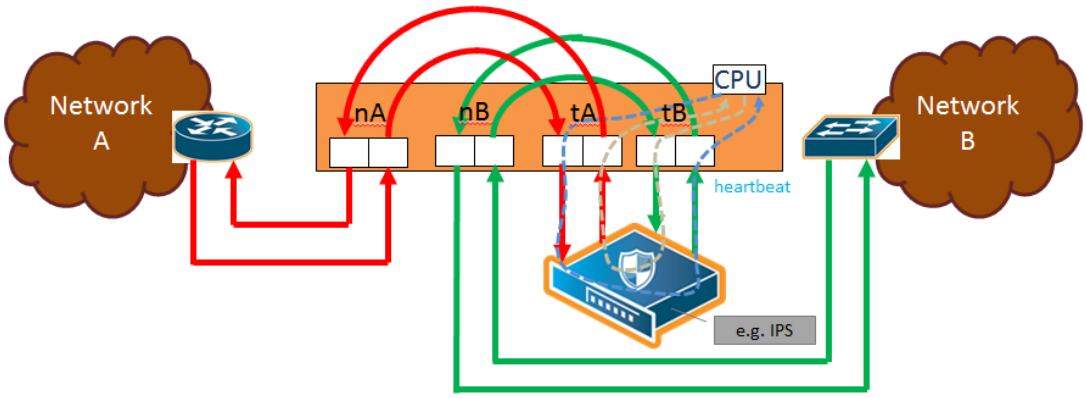

Figure 25-14: Traffic Path of To-Inline-Tool

Physical Bypass Parameter

One of the parameters of inline networks is physical bypass, which controls the state of the optical protection switch on the bypass combo module or copper TAP module when the module is powered on. The optical protection switch can have one of the following states:

- close—the fiber connected to the side A network port is passively coupled with the fiber connected to the side B port without any transceivers or switching fabric. Therefore, any traffic coming in on these fibers is exchanged between the two inline network ports without being noticed by the system.
- open—the fiber connected to the inline network ports is coupled through transceivers with the switching fabric that is under software control. Therefore, any traffic coming in on these fibers is subject to the traffic forwarding rules imposed by the current configuration as well as the current state of the inline tools.

When the bypass combo module or copper TAP modules is powered off, the optical protection switch is always in the close state.

When the bypass combo module or copper TAP modules is powered on, the state of the optical protection switch is as follows:

- the close state if **Physical Bypass** is set to enable (that is, selected on the configuration page)

- the open state if physical bypass is set to disable (that is, not selected on the configuration page)

Physical Bypass is enabled by default.

NOTE: Physical bypass only applies to protected inline networks.

Redundancy Profile Parameter

One of the parameters of inline networks is redundancy profile. A redundancy profile is used to configure Gigamon Resiliency for Inline Protection (GRIP). For more detailed information about GRIP, refer to the *Configure Gigamon Resiliency for Inline Protection* on page 627.

To created a redundancy profile, doing the following:

1. Select **Inline Bypass > Redundancies**.
2. Click New to open the Add Redundancy Profile page.
3. Type a name for the profile in the Alias field to help identify the profile.
4. Select a stack port by clicking the **Signaling Port** field and selecting an available port.

   The Signaling Port field specifies the used to signal the state of the two GigaVUE-HC2 nodes to each other. The ports provide the mechanism to detect loss of power in one of the GigaVUE-HC2 nodes.
5. Click in the **Protection Role** field and select a role.

   The protection role specifies the role of the GigaVUE-HC2, as primary, secondary, or suspended. The default is suspended. When suspended, the protection role is on hold. Changing a GigaVUE-HC2 from the primary role to the suspended role can be used to manually force one of the GigaVUE-HC2 nodes to become active. The suspended role is also used when performing maintenance.
6. Click **Save**.

To add the redundancy profile to the inline network, select the profile from the Redundancy Profile drop-down list on the Inline Network configuration page.

Display Current State of Inline Bypass Solution

Inline networks, inline tools, and inline maps work together to form an inline bypass solution. The inline bypass solution has an overall state, which can change in response to hardware conditions and user configuration.

Several factors make up the overall state of an inline bypass solution, as follows:

- The physical bypass configuration of the inline network is protected.
- The inline network configuration, in particular, if physical bypass is enabled or disabled, what traffic path is configured, and if link failure propagation is enabled or disabled.
- The inline tool configuration, in particular, if the state of the inline tool is enabled or disabled, if there is a heartbeat profile configured and if the heartbeat is enabled or disabled, and what failover actions are configured.

- The inline tool group configuration, in particular, if the state of the inline tool group is enabled or disabled, what failover mode is configured, what failover action is configured, and what number of minimum healthy tools is configured.
- The status of the links attached to the inline network ports and inline tool ports.

Table 25-1 on page 580 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 Table 25-1 on page 580.

Table 25-1: Forwarding States and Determining Factors

physical-bypass	traffic-path	Status of far-end ports connected to inline network ports	Status of inline tool side	Description
enable	drop, bypass, monitoring, or to-inline-tool	any status	any status	Forwarding state—PHYSICAL BYPASS If physical bypass is enabled, all traffic is exchanged directly between side A network and side B network without any monitoring by the GigaVUE node. Only applies to protected inline networks.
disable	drop	any status	any status	Forwarding state—DISABLED If the inline network is configured with a traffic path of drop, no traffic is exchanged between side A network and side B network because all packets coming to the inline network ports are dropped. **NOTE:** If the inline network is configured with link failure propagation set to true (enabled), the status of the inline network ports will be determined by the status of the far-end ports connected to the inline network ports. If both far-end ports are up, then both inline network ports will be up. If any far-end ports are down, then both inline network ports will be down.
disable	bypass, monitoring, or to-inline-tool	at least one far-end port is down	any status	Forwarding state—DISCONNECTED When one of the inline network ports is down due to a link down caused by a far-end device, no traffic is exchanged between side A network and side B network. **NOTE:** If the inline network is configured with link failure propagation enable as true, the status of both inline network ports will be down. That is, if only one inline network port is down, the other will be brought down.

Table 25-1: 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
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.
disable	to-inline-tool	both far-end ports are up	at least one inline tool involved in the inline bypass solution is disabled or associated with any inline tool port that is down or associated with the heartbeat protocol in a down state	Forwarding state—FAILURE-INTRODUCED BYPASS All traffic is exchanged directly between side A network and side B network through the switching fabric as a result of an inline tool port failure or heartbeat failure or inline tool being disabled that led to the network-level bypass due to the configured failover actions.
disable	to-inline-tool	both far-end ports are up	at least one inline tool involved in the inline bypass solution is disabled or associated with any inline tool port that is down or associated with the heartbeat protocol in a down state	Forwarding state—FAILURE-INTRODUCED DROP No traffic is exchanged between side A network and side B network. All packets coming to the inline network ports are dropped as a result of an inline tool port failure or heartbeat failure or inline tool being disabled that led to the network-level drop due to the configured failover actions.

Table 25-1: 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
disable	to-inline-tool	both far-end ports are up	at least one inline tool involved in the inline bypass solution is disabled or associated with any inline tool port that is down or associated with the heartbeat protocol in a down state	Forwarding state—NETWORK PORTS FORCED DOWN No traffic is exchanged between side A network and side B network. The inline network ports are forced down as a result of an inline tool port failure or heartbeat failure or inline tool being disabled that led to the network ports being forced down due to the configured failover actions.
disable	to-inline-tool	both far-end ports are up	at least one inline tool involved in the inline bypass solution is disabled or associated with any inline tool port that is down or associated with the heartbeat protocol in a down state	Forwarding state—ABNORMAL Any condition of the inline bypass solution that does not meet the criteria of the forwarding states listed in this table. At least one inline tool passes traffic as desired. There are many scenarios that can lead to the abnormal forwarding state, for example, when one inline tool member of an inline tool group has failed, but the number of remaining healthy tools is still above the minimum required number of healthy tools.

How to Use SNMP Polling to Obtain Inline Network State

The overall inline network state can also be obtained through SNMP polling using an SNMP-compliant network management application or a MIB browser. The names of the MIB files that need to be loaded in order to poll the inline network state are: GIGAMON-COMMON-SMI and GIGAMONINLINEBYPASS.

The inline network forwarding states are described in Table 25-1 on page 580. Some of the states have slightly different names in SNMP.

SNMP Notification of Forwarding State Change

Use the following steps to configure a notification that will be sent when the inline bypass forwarding state changes:

1. Select **Settings > SNMP Traps**.
2. Click **Trap Settings**. The Edit SNMP Trap Settings page displays.
3. On the Edit SNMP Trap Settings page, select **Enable** for **Inline Bypass Forwarding State Change**.
4. Click **Save**.

How to Use Syslog to Obtain Inline Network State

The overall inline network state can also be obtained through syslog.

When the inline network state change is triggered by a configuration change, a sample syslog message is as follows:

```
IBFE STATE CHANGE inline network default_inline_net_1_1_1 status
NORMAL changed to PHYSICAL BYPASS triggered by config
```

When the inline network state change is triggered by a link status change or heartbeat status change, a sample syslog message is as follows:

```
IBFE STATE CHANGE inline network inNet1 status NORMAL changed to
FAILURE INTRODUCED BYPASS triggered by 1/x10
```

In the syslog messages, IBFE indicates the Inline Bypass Failover Engine.

The inline network forwarding states are described in Table 25-1 on page 580.

About 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, GigaVUE-HC2, or GigaVUE-HC1 node. The inline networks participating in the inline network group can be different speeds and different mediums.

An inline network group can share an inline tool, or tools in an inline tool group or inline tool series. Many-to-one means from an inline network group to an inline tool and is shown in Figure 25-15 on page 583. Many-to-many means from an inline network group to an inline tool group or inline tool series.

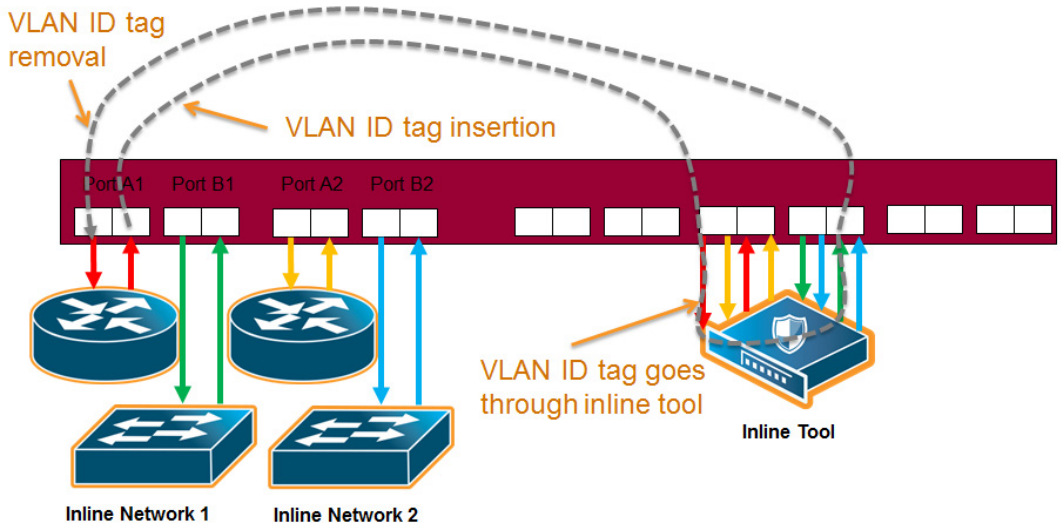

Figure 25-15: Inline Network Group Many-to-One

When an inline tool or inline tool group is configured between several pairs of networks, each pair operates independently of the others. The traffic coming from the inline tool or tools must be segregated into individual substreams according to the traffic source.

For example, the traffic that came into the GigaVUE node on inline network port net_A_3 must be sent through inline network port net_B_3 (with net_A_3 and net_B_3 belonging to the same inline tool that is the third member of the inline tool group). Refer to Figure 25-16.

Figure 25-16: Inline Tool Inserted in Multiple NetA-NetB Links

To accomplish the segregation, packets are tagged coming from an individual side A inline network port before they are sent to the inline tool or tools. When tagged packets come back from the inline tool or tools, the tag identifies the respective side B inline network port through which the packets should be sent. The tags are removed before the packets are sent through the side B inline network ports.

Traffic is guided to a particular inline network through internal VLAN ID tagging. This VLAN tagging affects packets only on their way from inline tool ports to the attached inline tool and back from the attached inline tool to the inline tool ports. The packets sent out from inline network ports remain untagged. Refer to Figure 25-15 on page 583.

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* on page 585.

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* on page 585.

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* on page 424.

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. Select **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. Select **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.

About Inline Tools

There are two meanings to the term inline tool. The inline tool software construct consists of a pair of inline tool ports plus the inline tool attached to the ports. The software construct has attributes that are configured on the GigaVUE HC Series node.

The term inline tool also refers to the pass-through device itself that performs packet inspection and selective forwarding, such as an Intrusion Protection System (IPS). This is a physical device, external to the GigaVUE node.

Use the Inline Tool page to configure the inline tool software construct. An inline tool consists of inline tool ports, always in pairs, running at the same speed, on the same medium (fiber or copper). The inline tool ports must be on the same GigaVUE-HC3, GigaVUE-HC2, or GigaVUE-HC1 node. The inline tool ports must also be on the same GigaVUE-HC3, GigaVUE-HC2, or GigaVUE-HC1 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-HC2 or 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 25-17 on page 587.

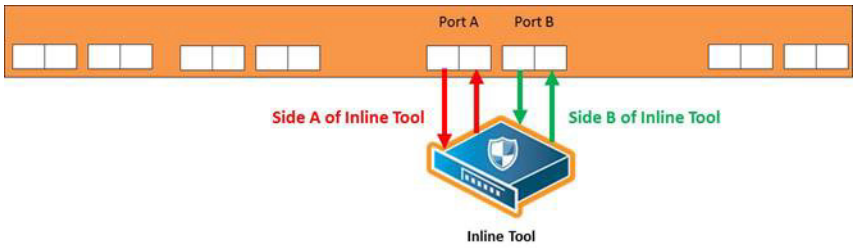

Figure 25-17: Inline Tool and Inline Tool Ports

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.

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 25-18 on page 588 for automatic and manual inline tool recovery from failover.

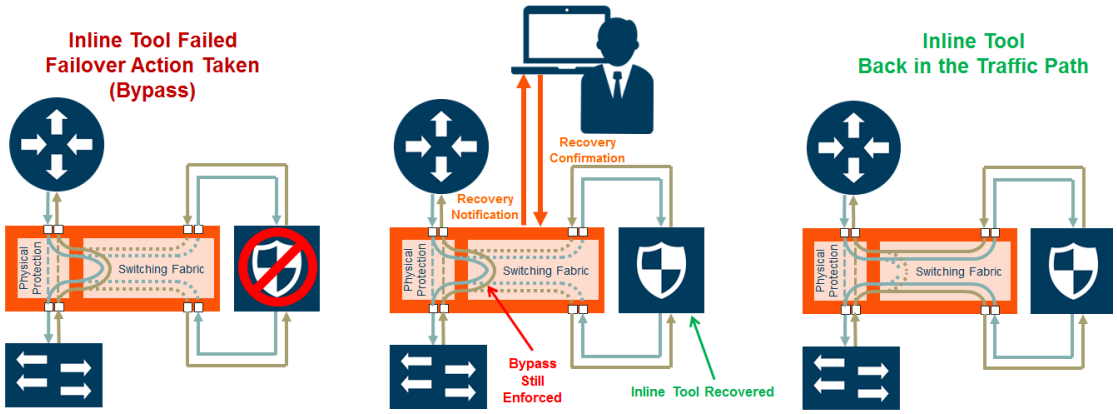

Figure 25-18: Automatic and Manual Inline Tool Recovery from Failover

The left side of Figure 25-18 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 25-18. When the inline tool recovers, traffic is automatically directed back to it.

Manual recovery is shown in the center of Figure 25-18. When the recovery mode is configured as manual, an SNMP notification, if enabled, will send a notification when the inline tool is ready to be put back into service. The failover action, in this case, bypass, will be enforced until you manually put the inline tool back into service.

When the recovery mode is configured as manual, an SNMP notification, when enabled, will notify you when the inline tool is ready to be put back into service.

Use the following steps to configure notification that will be sent when the inline tool is ready:

1. Select **Settings > SNMP Traps**.
2. Click **Trap Settings**. The Edit SNMP Trap Settings page displays.
3. On the Edit SNMP Trap Settings page, select **Enable** for **Inlinetool Recovery**.
4. Click **Save**.

The default for **Inlinetool Recovery** is disabled.

Use the following steps to put an inline back in service when the recovery mode is manual and the inline tool has an operational state of ready.

1. Select Inline **Bypass > Inline Tools**.
2. On the Inline Tools page, select the inline tool.
3. Click **Recover**.

NOTE: Also use **Recover** after the GigaVUE HC Series node is reloaded or rebooted (even though the inline tool has not failed). Issue the **Recover** on all the inline tools that are configured with manual recovery after a reload or **Actions > Shut Down** from the Chassis page on a selected card in Chassis Table View followed by **Actions > Start Up**. Refer to “Chassis” in the *GigaVUE-OS H-VUE Administration Guide* for more details.

The **Inline Tool** page displays the operational state of each inline tool as up, down, disable, or ready. Refer to Table 25-2 on page 590 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 25-2. They are an aggregate of the inline tool port statuses and the heartbeat status.

Table 25-2: 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.

About Heartbeat Profiles

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 through 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.

Heartbeat Profiles

A heartbeat profile supports health monitoring of inline tools. A heartbeat profile is a group of attributes applied to an inline tool to configure its heartbeat operation. Multiple inline tools can share a heartbeat profile.

To display the configured heartbeat profiles, select Inline Bypass > Heartbeats to open the Heartbeats page. An example is shown in Figure 25-19.

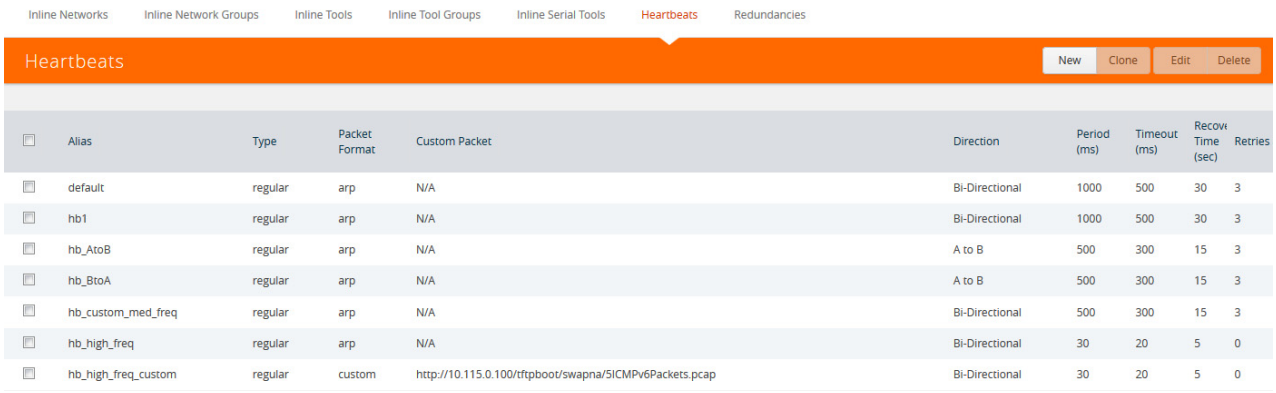

Alias	Type	Packet Format	Custom Packet	Direction	Period (ms)	Timeout (ms)	Recove Time (sec)	Retries
default	regular	arp	N/A	Bi-Directional	1000	500	30	3
hb1	regular	arp	N/A	Bi-Directional	1000	500	30	3
hb_AtoB	regular	arp	N/A	A to B	500	300	15	3
hb_BtoA	regular	arp	N/A	A to B	500	300	15	3
hb_custom_med_freq	regular	arp	N/A	Bi-Directional	500	300	15	3
hb_high_freq	regular	arp	N/A	Bi-Directional	30	20	5	0
hb_high_freq_custom	regular	custom	http://10.115.0.100/tftpboot/swapna/5ICMPv6Packets.pcap	Bi-Directional	30	20	5	0

Figure 25-19: 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
- 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 and GigaVUE-HC2, and 16 on the GigaVUE-HC1.

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.

Use the **Add Heartbeats** page shown in Figure 25-20 to configure a heartbeat or negative heartbeat profile by selecting **Inline Bypass > Heartbeats > Heartbeats**, and then clicking **New**.

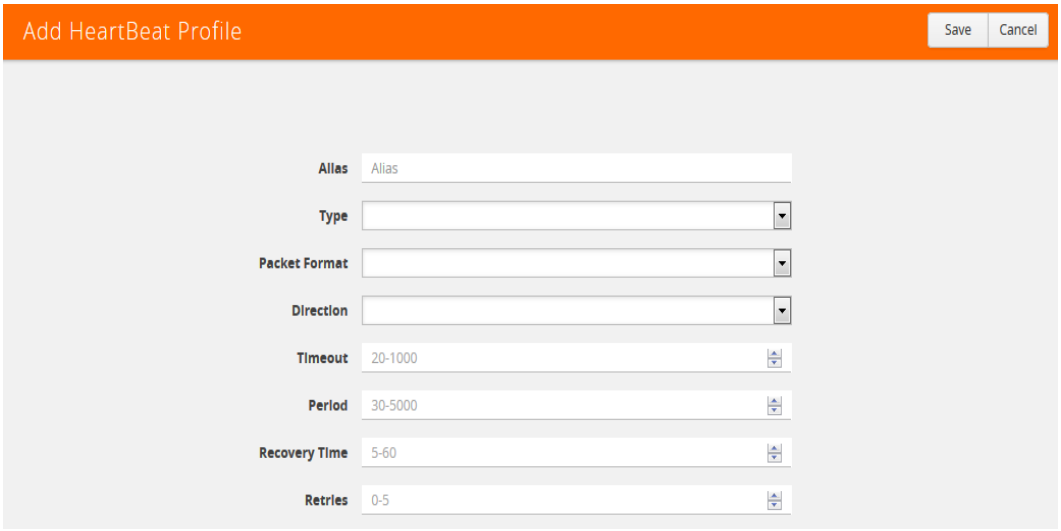

Figure 25-20: Configured Heartbeat Profile for Inline Tool

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
Heartbeat Profile	The alias of the heartbeat profile
Inline Tool Alias	The alias of the inline tool to which the heartbeat profile is related.
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.

You can clear the statistics for a specific heartbeat profile or all heartbeat profile. To clear a single heartbeat profile,

1. select the profile in the Heartbeat profile column.
2. Click **Clear**.

To clear the statistics for all heartbeat profiles, do the following:

1. Select **Clear All**.
2. Choose one of the following from the menu:
   - Select **Clear All Heartbeat Stats** to clear all statistics for all heartbeat profiles.

- Select **Clear All Negative Heartbeat Stats** to clear only the statistics for negative heart profiles. The statistics for regular heartbeat profiles will remain.

For details on the parameters of heartbeat profiles, refer to the following:

- *Standard Heartbeat* on page 593
- *Standard or Custom Heartbeat Packet* on page 593
- *Detect Inline Tool Failure* on page 593.

Standard Heartbeat

The standard heartbeat is a packet sent by the GigaVUE node that passes through the inline tool to verify that it is passing traffic, even if the link is *up*. If the packet is not passed through the tool, the tool is considered to have failed and a bypass action is triggered.

Even when the tool is considered down, heartbeat packets continue to be sent so that the bypass action can be reversed when the tool is healthy again.

Heartbeats are sent bidirectionally to the inline tool.

Standard or Custom Heartbeat Packet

The format of the heartbeat packet can be the standard ARP packet or a custom packet. For a custom packet, you must provide a URL from which a PCAP file can be imported. If the PCAP file contains several packets, the first packet present in the file is taken as the heartbeat packet. The size of a custom heartbeat packet must be less than 128 bytes.

NOTE: The system will overwrite the MAC address portion of the custom heartbeat packet.

If the inline tool through which the heartbeat packets are passed is expecting IPv6 traffic exclusively, you must select a custom heartbeat packet.

Custom heartbeat packets are needed in situations in which inline tools do not reliably pass standard ARP packets. For example, if an inline tool is configured to pass only IPv6 traffic, an ICMPv6 ARP packet might be appropriate.

If a custom heartbeat packet is specified, the **Heartbeats** page displays the name of the PCAP file from which it was imported.

Detect Inline Tool Failure

The health of the inline tool is critical to the proper handling of traffic. An inline tool is determined to have failed if:

- link is lost to the tool
- inline heartbeat fails

When the tool fails in one direction, it is considered to have failed in both directions. For example, if the heartbeat stops flowing in the northbound direction, neither northbound or southbound packets are sent to the tool.

Negative Heartbeat Profiles

A negative heartbeat profile is a group of attributes applied to an inline tool to configure its negative heartbeat operation. Multiple inline tools can share a negative heartbeat profile. The content of a negative heartbeat is configurable using the same PCAP file mechanism as for a custom heartbeat packet.

Use the Add Heartbeats page shown in Figure 25-20 on page 592 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:

- **Packet Format** 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 and GigaVUE-HC2, and 16 on the GigaVUE-HC1.

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 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.

About 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, GigaVUE-HC2, or GigaVUE-HC1 node, but can be on different modules on the node. The inline tool ports must also be on the same GigaVUE-HC3, GigaVUE-HC2, or GigaVUE-HC1 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 25-21 on page 596.
- 1+1 redundancy—single inline tool with a spare inline tool. Refer to Figure 25-22 on page 596.
- 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 25-23 on page 597.

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 25-21 on page 596 for a non-redundant inline tool group.

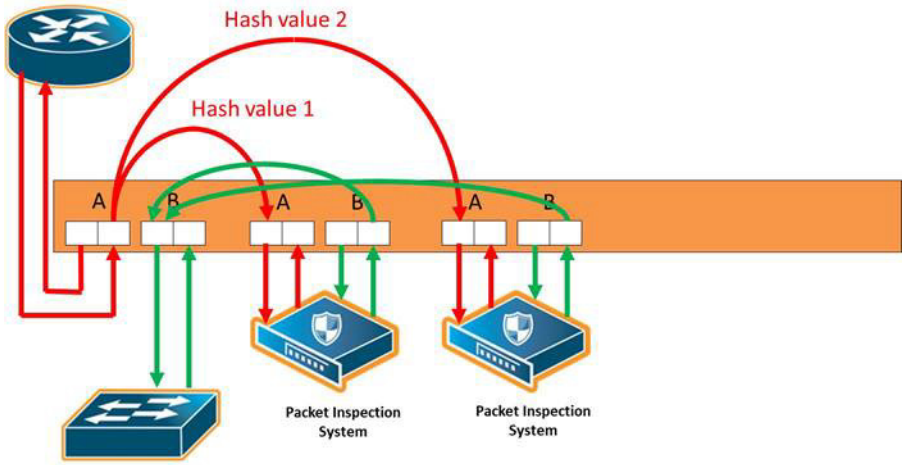

Figure 25-21: Inline Tool Group With No Spare, Non-Redundant

Refer to Figure 25-22 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 25-22, traffic is only shown from A-to-B.

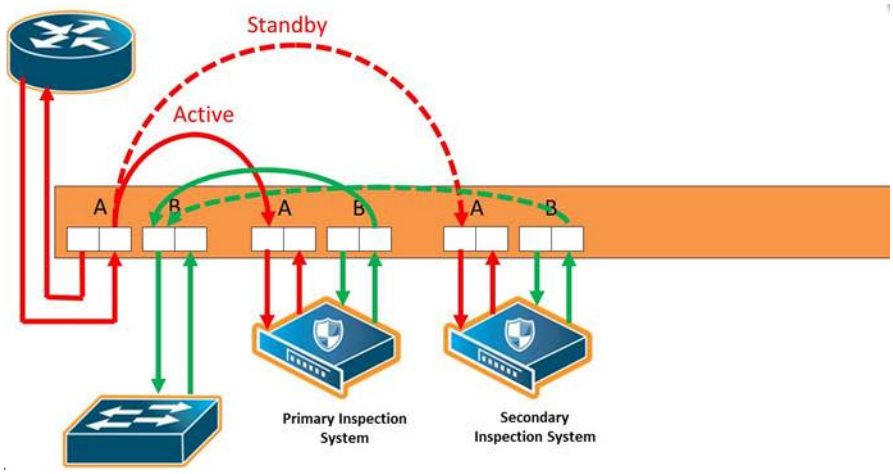

Figure 25-22: Inline Tool Group With Spare, Redundant 1+1 Scenario

Refer to Figure 25-23 on page 597 for an inline tool group in an N+1 redundant scenario, in which N is greater than 1. In Figure 25-23, traffic is only shown from A-to-B.

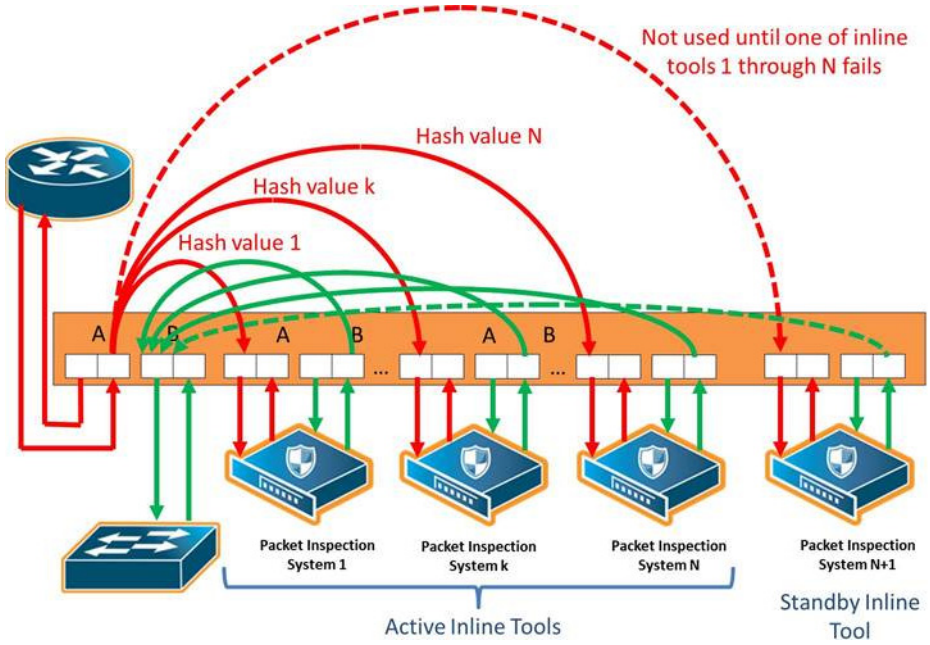

Figure 25-23: Inline Tool Group With Spare, Redundant N+1 Scenario

For details on the parameters of inline tool groups, refer to the following:

- *Inline Tool Group Failover Action* on page 597
- *Inline Tool Group Spare Inline Tool* on page 598
- *Symmetrical and Asymmetrical Hashing* on page 599.
- *Resilient Weighted Hashing* on page 601

Inline Tool Group Failover Action

One of the parameters of inline tool groups is the failover action, taken in response to a failure of an inline tool group, when the number of healthy inline tools in the inline tool group (including the spare inline tool, if configured) falls below the configured minimum. You can configure one of the following failover actions:

- **ToolBypass**—when the inline tool group fails, the traffic that normally was directed to the inline tool group is redirected to the bypass path. Use this failover action for configurations involving multiple inline tools or inline tool groups associated with an inline network or inline network group using rule-based maps. For configurations using map passalls, tool-bypass is the same as network-bypass.
- **NetworkBypass**—when the inline tool group fails, all traffic that would not have been dropped when the inline network or networks had a NORMAL forwarding state is directed to the bypass path. That is, all such traffic arriving at the side A inline network port or ports is forwarded to the side B inline network port or ports and all traffic arriving at the side B inline network port or ports is forwarded to the side A inline network port or ports.
- **NetworkPortForcedDown**—when the inline tool group fails, the inline network ports of the respective inline network (or inline network group) are forced down.

- **ToolDrop**—when the inline tool group fails, the traffic that normally was directed to the inline tool group is dropped. Use this failover action for configurations involving multiple inline tools or inline tool groups associated with an inline network or inline network group using rule-based maps. For configurations using map passalls, tool-drop is the same as network-drop.
- **NetworkDrop**—when the inline tool group fails, all traffic coming to the respective inline network (or inline network group) is dropped.

The default is tool-bypass.

The bypass path is between side A and side B of the inline network ports.

The failover action of all the inline tools specified by the inline tool list is overwritten by the failover mechanism of the inline tool group. This means that when a given inline tool specified by the inline tool list fails, the traffic originally directed to this inline tool is redirected to the spare inline tool (if one is configured and available) or handled according to the failover mode of the active tools, so long as the total number of healthy inline tools in the inline tool group is not smaller than the minimum required number of healthy inline tools.

When the total number of healthy inline tools in the inline tool group drops below the minimum required number of healthy inline tools, the failover action of the inline tool group determines the action to be taken.

Inline Tool Group Spare Inline Tool

One of the parameters of inline tool groups is a spare inline tool. If a spare is configured, the inline tool group becomes a redundant arrangement of inline tools. When the first failure occurs in a set of active inline tools, traffic will be forwarded to the spare with no loss, thus the spare will replace the failed tool in the active set.

The inline tools in the inline tool list are considered to be active inline tools. The traffic is hash-distributed over the active inline tools as long as all the inline tools are healthy, When one of the active inline tools fails, the spare inline tool takes the place of the failed inline tool and the new set operates as a new active set. If another inline tool fails, the traffic is redistributed according to the failover mode, as if there was no spare.

When the number of failed inline tools is such that the number of healthy inline tools is less than the minimum-group-healthy-size, the group heals itself by re-spreading the traffic over the healthy tools. When the number of healthy tools falls below the minimum-group-healthy-size, the failover action of the inline tool group takes place, while the failover action of the member inline tools is ignored.

The spare inline tool works with another parameter called release-spare-if-possible. When the inline tool that had been replaced with the spare inline tool recovers, the release-spare-if-possible parameter determines if the recovering inline tool is included in the active set of inline tools or if it becomes the new spare inline tool.

The default of the release-spare-if-possible parameter is disabled. Disabled means that even if the original inline tool recovers, the spare that replaced it will remain in the active set of inline tools. Enabled means that after the original inline tool recovers, the

spare that replaced it will be released, if possible, from the active set of tools to become the spare again.

Configure the minimum-group-healthy-size and release-spare-if-possible parameters at the same time you configure the spare inline tool.

Symmetrical and Asymmetrical Hashing

One of the parameters of inline tool groups is hashing, which is used for distributing packets across a group of inline tools belonging to the inline tool group. The values for the hash parameter are as follows:

- **advanced**—Specifies symmetrical hashing, which is derived from the combination of packet fields based on the criteria selected for the advanced-hash algorithm configured under the **Ports > Port Groups > GigaStreams > Advanced Hash Settings** page.

  For inline bypass applications, the most common choice of criteria for the advanced-hash algorithm is the combination of source IP and destination IP addresses. This produces a hash value that sends all traffic associated with the same session to the same inline tool in the inline tool group.
- **a-srcip-b-dstip**—Specifies asymmetrical hashing, which is derived from the source IP address for side A of the network and the destination IP address for side B of the network. This produces a hash value that sends all traffic associated with the same source address residing on side A to the same inline tool in the inline tool group, regardless of destination or session.
- **b-srcip-a-dstip**—Specifies asymmetrical hashing, which is derived from the destination IP address for side A of the network and the source IP address for side B of the network. This produces a hash value that sends all traffic associated with the same source address residing on side B to the same inline tool in the inline tool group, regardless of destination or session.

The default is **advanced**.

Use asymmetrical hashing if all traffic exchanged between a particular node on one side of the network and any nodes on the other side of the network that communicate with that node need to go to the same inline tool. The asymmetrical hashing options involve only source IP address (srcip) in one direction and only destination IP address (dstip) in the opposite direction. Bi-directional traffic, such as between a given user and all the Internet sites visited by the user, will be sent to the same inline tool in the group.

NOTE: When asymmetric hashing is configured, the portsrc and portdst packet fields are not included in the advanced-hash calculation for any GigaStream and inline tool groups across the GigaVUE node.

With symmetrical hashing, the inline network traffic path parameter can be configured to different values on the inline networks. With asymmetrical hashing, there is a restriction. Refer to *Asymmetrical Hashing Restrictions* on page 600.

Refer to Figure 25-24 for asymmetrical hashing.

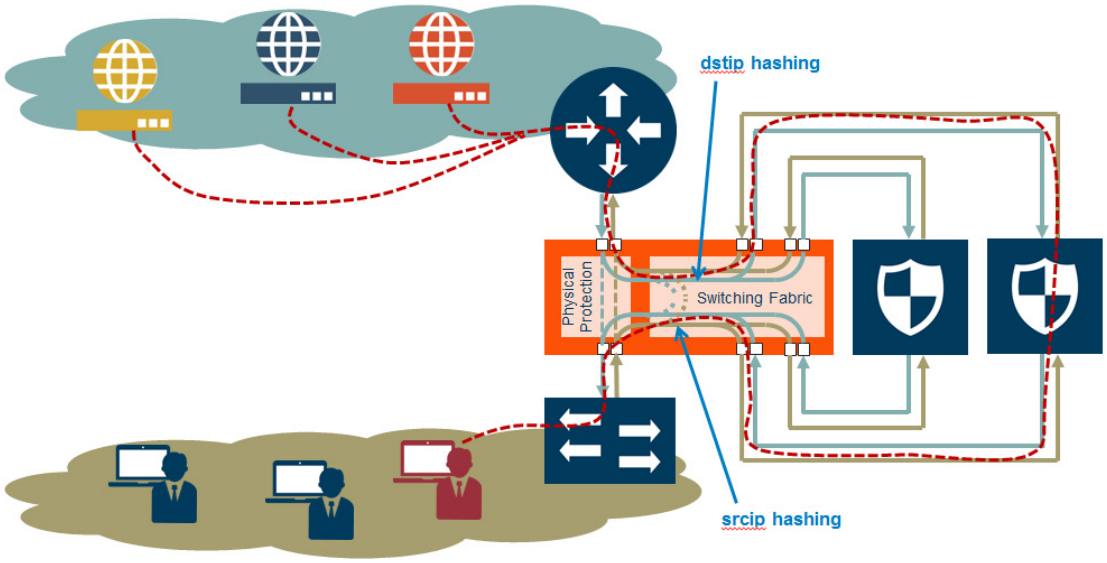

Figure 25-24: Asymmetrical Hashing

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.

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.

About 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, GigaVUE-HC2, or GigaVUE-HC1 node. The inline tool ports and inline tool groups must also be on the same GigaVUE-HC3, GigaVUE-HC2, or GigaVUE-HC1 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 25-25 for an illustration of an inline tool series. In Figure 25-25, traffic is only shown from A-to-B.

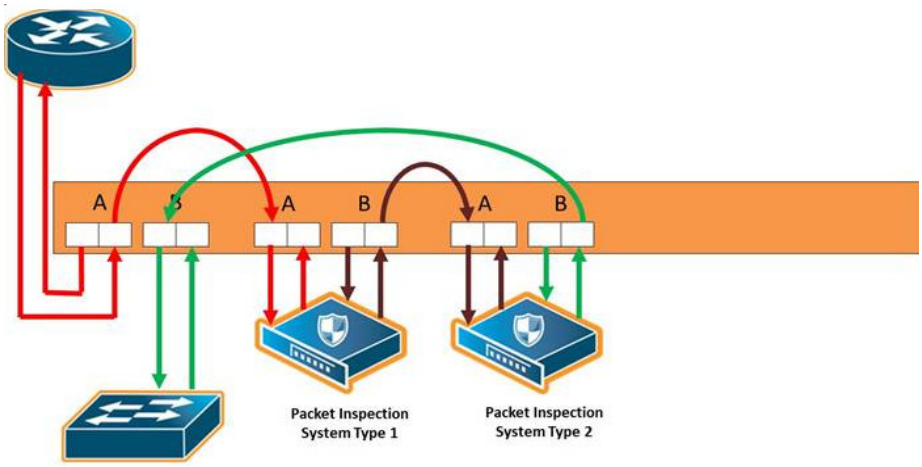

Figure 25-25: Inline Serial Tools

Refer to Figure 25-26 on page 602 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 25-26, 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.

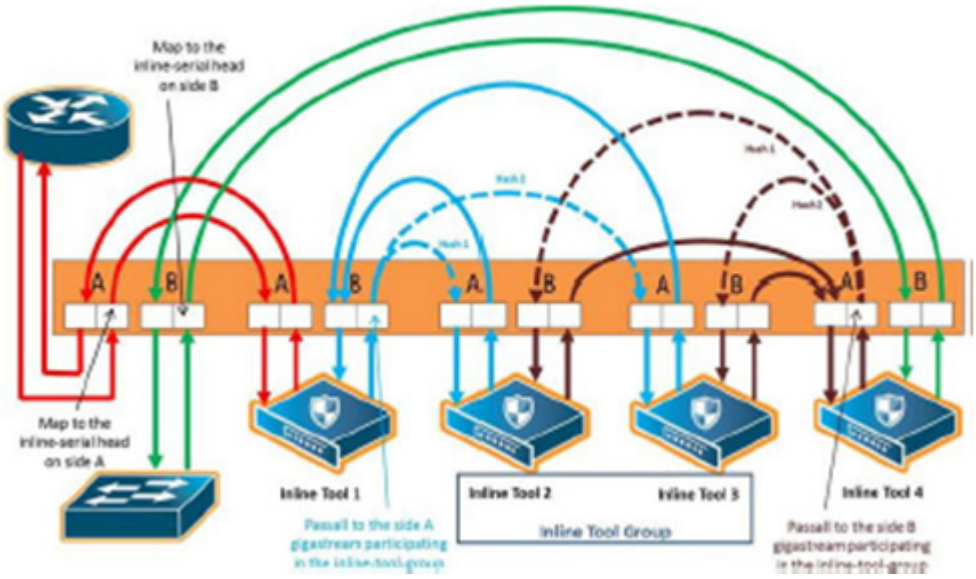

Figure 25-26: Inline Tool Series, Including Inline Tool Group

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 **Comment** 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* on page 605 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 25-27 on page 604.
- **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 25-27 on page 604.

- **NetworkPortForcedDown**—when the inline serial tools fails, the inline network ports of the respective inline network are forced down. Refer to Figure 25-28 on page 605.
- **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 25-29 on page 605.
- **NetworkDrop**—when the inline serial tools fails, all traffic coming to the respective inline network (or inline network group) is dropped. Refer to Figure 25-29 on page 605.

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 25-27 to Figure 25-29 show the global failover actions for the inline series when any individual inline tool in the series fails.

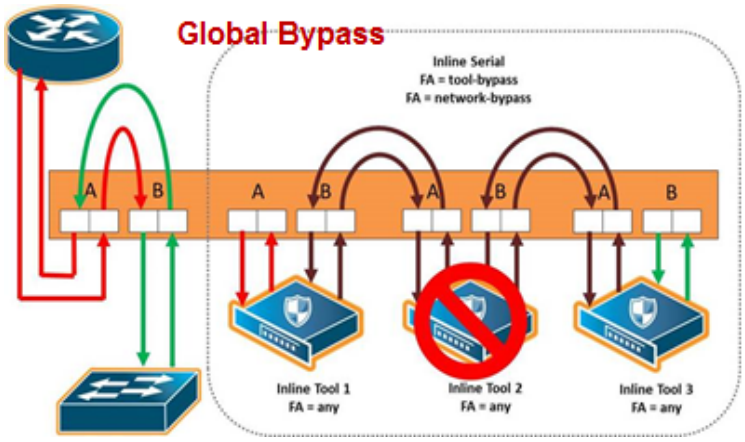

Figure 25-27: Inline Tool Series Global Bypass Failover Action

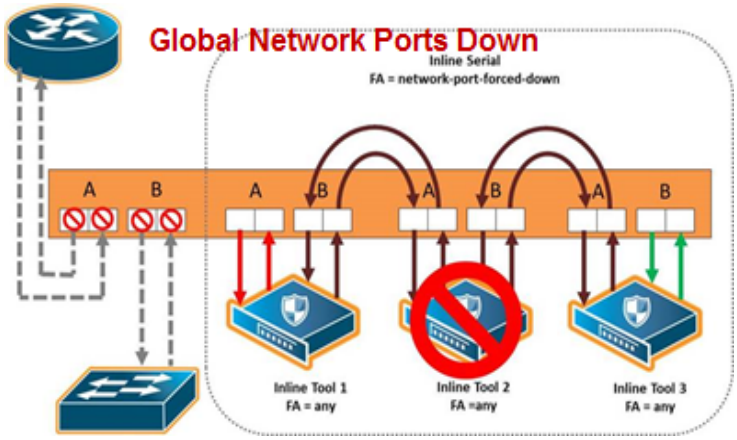

Figure 25-28: Inline Tool Series Global Network Ports Forced Down Failover Action

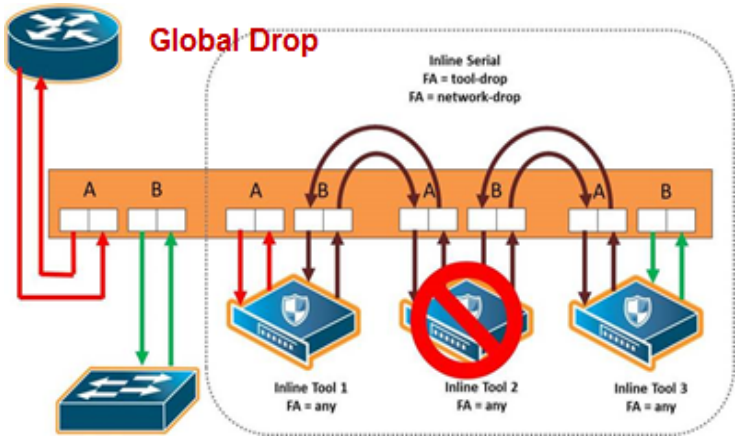

Figure 25-29: Inline Tool Series Global Drop Failover Action

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* on page 587.

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 25-30 on page 606.

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 25-31 on page 606.
- **NetworkBypass**—when the inline tool fails, a bypass is established between the inline network ports. Refer to Figure 25-32 on page 607.
- **NetworkDrop**—when the inline tool fails, traffic is dropped at the inline network ports. Refer to Figure 25-33 on page 607.
- **NetworkPortForcedDown**—when the inline tool fails, the links for the inline network ports are brought down. Refer to Figure 25-34 on page 607.

Figure 25-30 to Figure 25-34 show the local failover actions when an individual inline tool in a series fails

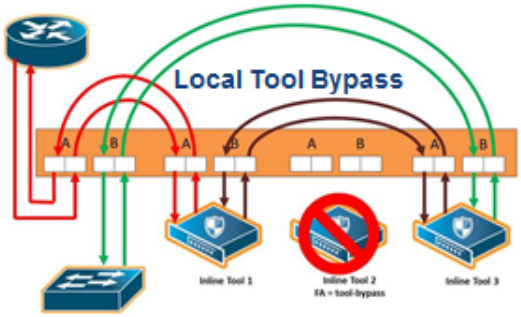

Figure 25-30: Inline Tool Series Local Tool Bypass Failover Action

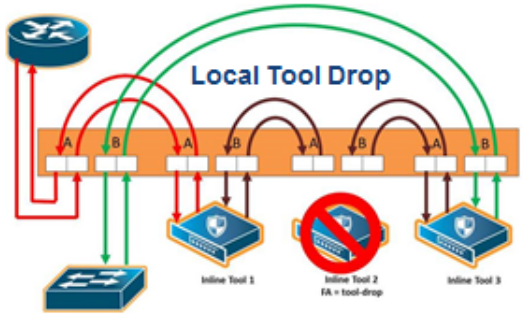

Figure 25-31: Inline Tool Series Local Tool Drop Failover Action

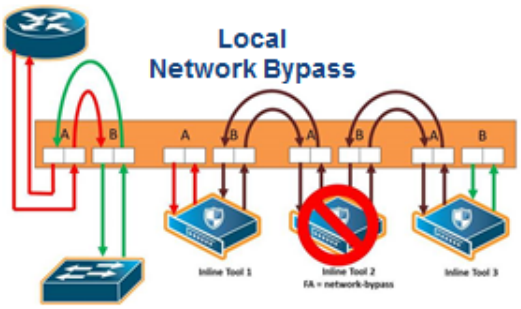

Figure 25-32: Inline Tool Series Local Network Bypass Failover Action

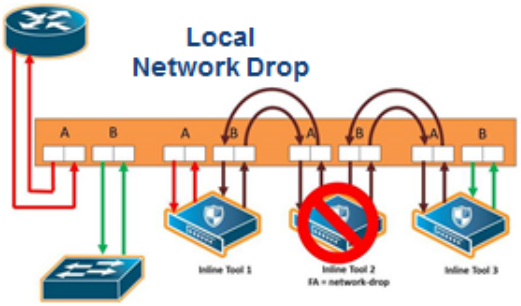

Figure 25-33: Inline Tool Series Local Network Drop Failover Action

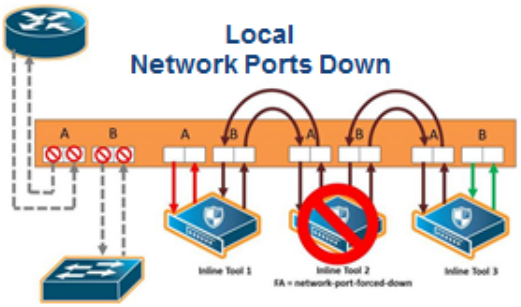

Figure 25-34: Inline Tool Series Local Network Ports Forced Down Failover Action

Figure 25-35 shows the failure of two individual inline tools in a series with different configured failover actions.

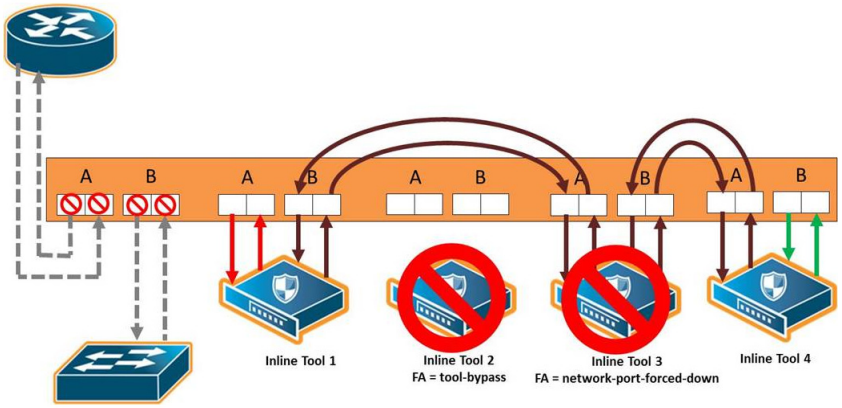

Figure 25-35: Inline Tool Series Local Failover of Two Tools

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 25-36 shows a simplified view of the flow through the tools:

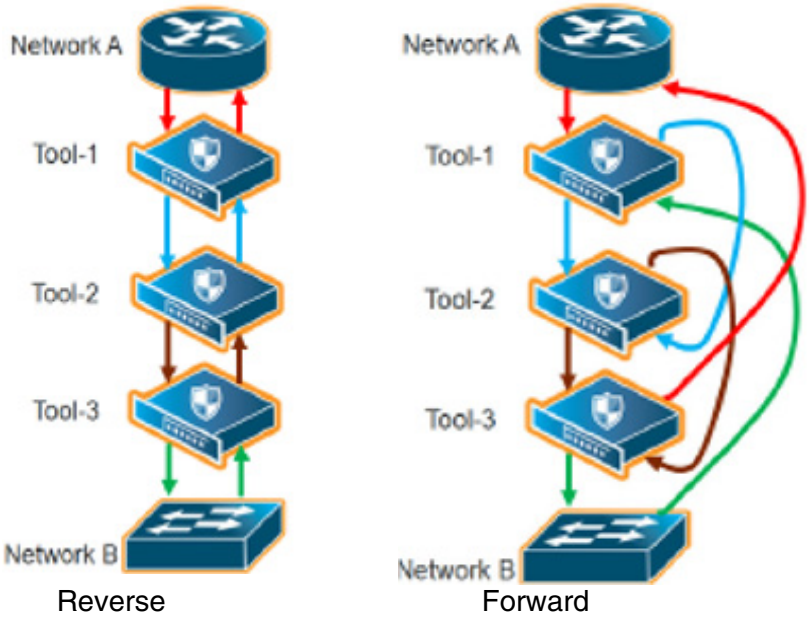

Figure 25-36: Inline Tool Series Per-Direction Order: Simplified Reverse and Forward

In Figure 25-36, **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 25-37 on page 609 shows the reverse direction with the tools connected to the GigaVUE node.

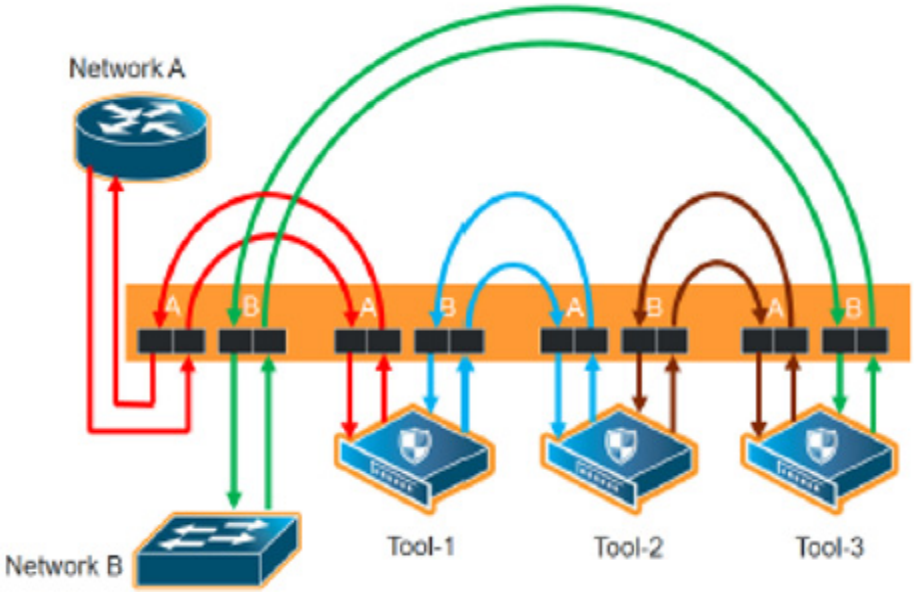

Figure 25-37: Inline Tool Series Per-Direction: Reversed

Figure 25-37 shows the forward direction with the tools connected to the GigaVUE node.

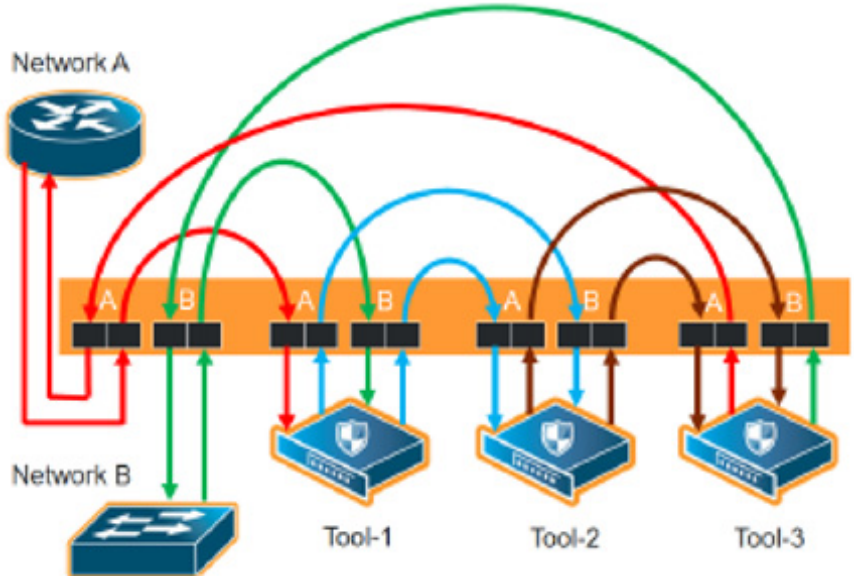

Figure 25-38: Inline Tool Series Per-Direction: Forward

Since the per-direction default value in software version 4.7 is different than the default value in older software versions, refer to *Recommendation When Upgrading from Older Software Versions* on page 610.

Recommendation When Upgrading from Older Software Versions

In software versions prior to 4.4, the only direction was **Forward**, so it was the default. Starting in software version 4.7, there are options for both **Forward** and **Reverse**, with **Reverse** as the default.

When upgrading from an older software version such as 4.3, the recommendation is to remove the existing inline serial configuration before the upgrade, then reconfigure the inline series after the upgrade to 4.7.

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, GigaVUE-HC2, or GigaVUE-HC1 node. Even if nodes are in a cluster, the inline ports cannot be on different nodes.

Refer to the following sections:

- *Inline Map Passall* on page 611
- *Inline Map* on page 612
- *Inline Map Shared Collector* on page 612
- *Inline Maps to Individual Members of an Inline Tool Group* on page 613
- *Out-of-Band (OOB) Map* on page 615
- *Symmetric and Asymmetric Maps* on page 617

Inline Map Passall

Use the **Pass All** subtype to configure a type of inline map that passes all traffic. Map-passalls facilitate the sending of traffic between inline network ports (through inline tools or bypass). All of the inline network ports and inline tool ports involved in an inline map must be located on the same Gigamon node.

Use the **Source** of the **Pass All** to specify a single inline network.

Use the **Destination** of the **Pass All** to specify an inline tool alias, an inline tool group alias, an inline tool series alias, or an inline bypass. An inline bypass is a special construct that is used as a pseudo inline tool to allow a portion of traffic to bypass any inline tools.

An inline bypass, using the physical bypass option, is only valid if the **Source** is an inline network port. This applies only to the **Pass All** subtype for asymmetrical scenarios.

Inline Map

Use the **New Map** page to configure a type of inline map that uses rules to direct traffic. These inline maps are referred to as rule-based maps. All of the inline network ports and inline tool ports involved in an inline map must be located on the same GigaVUE node.

Use the **Source** field to specify an inline network alias or an inline network group alias.

Support for rule-based maps is limited to symmetric scenarios, which means that the **Destination** of rule-based inline maps can only be aliases of inline networks or inline network groups (not individual ports).

Use the **Destination** to specify an inline tool alias, an inline tool group alias, an inline tool series alias, or an inline bypass. An inline bypass is a special construct that is used as a pseudo inline tool to allow a portion of traffic to bypass any inline tools.

For rule-based maps, the **Destination** can be configured to inline bypass with no restrictions, so long as the **Source** specifies either an inline network or an inline tool.

Use the **Priority** to order inline maps by priority. For example, you can specify the highest priority map to be for encrypted traffic and lowest priority map to be for a shared collector. You can also place inline maps before or after one another using the **Priority**.

Use the **Type** to specify any map rule.

Inline Map Shared Collector

Use the **Collector** subtype to configure a shared collector to which to send any packets that do not match the map rules in the inline maps. Use a shared collector with one or more rule-based inline maps. All of the inline network ports and inline tool ports involved in an inline shared collector map must be located on the same GigaVUE node.

Use the Source argument to specify an inline network alias or an inline network group alias.

Use the Destination argument to specify an inline tool alias, an inline tool group alias, an inline tool series alias, or an inline bypass. An inline bypass is a special construct that is used as a pseudo inline tool to allow a portion of traffic to bypass any inline tools.

For shared collector maps, the **Source** field can be configured to inline bypass with no restrictions, so long as the **Destination** field specifies either an inline network or an inline tool.

Support for shared collector inline maps is limited to symmetric scenarios, which means that the **Destination** field of rule-based inline maps can only be aliases of inline networks or inline network groups (not individual ports).

Inline Maps to Individual Members of an Inline Tool Group

Prior to software version 4.4, the **Map** page was used to configure an inline map that directed traffic to the inline tool group as a whole. The map could be rule-based or passall. There was only one type of hashing available, which distributed the traffic across the tools in the inline tool group.

Starting in software version 4.4, the **Map** page can also be used to configure inline maps to the individual members of an inline tool group. The maps must be rule-based. There are also more hashing options that can be specified for traffic that does not match any of the map rules. The hashing options are described in *Symmetrical and Asymmetrical Hashing* on page 599.

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 25-39 for the rule-based maps to the individual members of the inline tool group. In Figure 25-39, traffic is only shown from A-to-B.

.

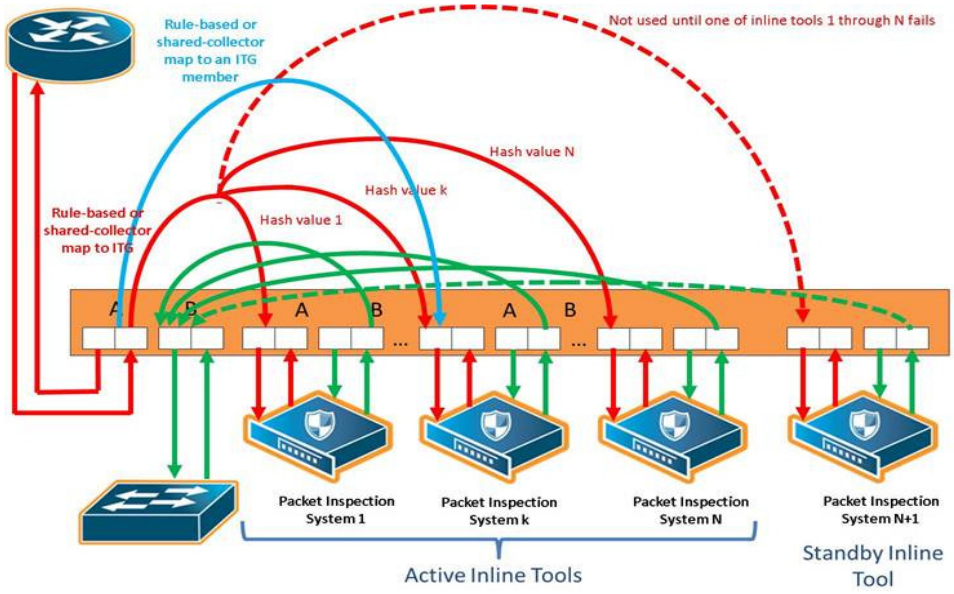

Figure 25-39: Rule-Based Maps to Individual Tools in an Inline Tool Group

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.

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 Table 25-3 on page 615 for the failover actions when an individual inline tool in an inline tool group fails.

Table 25-3: 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 25-40 on page 616.
- If the inline bypass solution use rule-based maps or map shared collectors, the OOB arrangements can use only map passalls. Refer to Figure 25-41 on page 617.
- 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

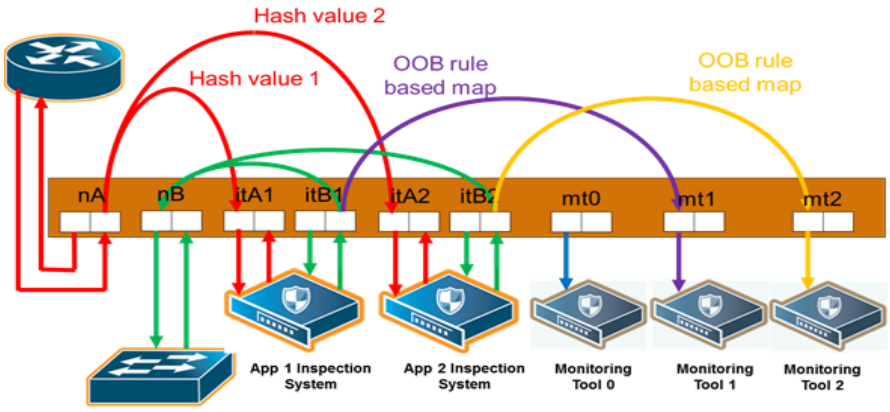

Figure 25-40: Out-of-Band Rule-Based Maps

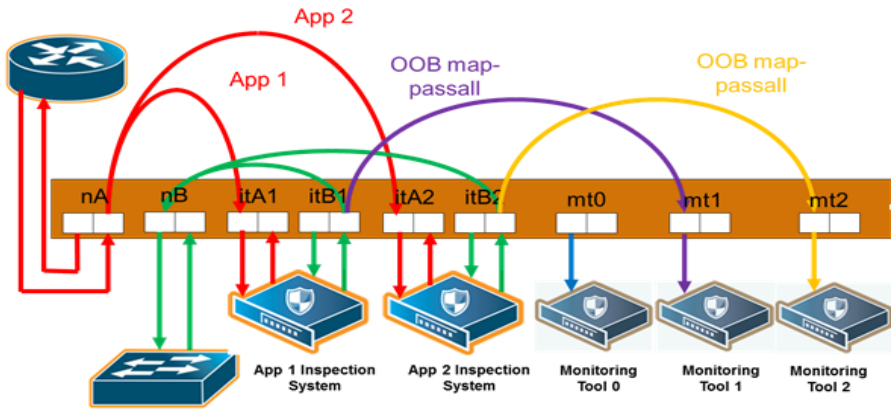

Figure 25-41: Out-of-Band Map Passalls

Symmetric and Asymmetric Maps

In a symmetric map configuration, the southbound and northbound forwarding to the tools is the same. For example, traffic from A to B goes through an inline tool, as does traffic from B to A. Rule-based maps are limited to symmetric configurations.

In an asymmetric map configuration, the southbound traffic is distributed to the inline tools, but northbound traffic can be sent through uninspected. For example, traffic from A to B goes through an inline tool, but traffic from B to A bypasses it.

Asymmetric configurations are only supported with map passalls. Traffic can come from individual inline networks, be sent to individual inline tools or inline tool groups, or to bypass.

Symmetrical combinations of two asymmetrical arrangements (that is, with both side A and side B pointing to the inline tool or to bypass) are not allowed.

Some IPSs do not need to inspect northbound traffic. For example, those focused on preventing Denial of Service (DoS) attacks, may not need to keep track of session flows and may only be concerned about southbound traffic.

Conversely, data loss prevention systems, those that are more concerned about what sensitive data is leaving the protected network than what is coming in, may focus solely on northbound traffic.

Configuration Steps

The configuration steps in summary for an inline bypass solution are as follows:

1. Configure inline network ports. (Optional for protected inline network.)
2. Configure inline network. (Optional for protected inline network.)
3. (Optional) Configure inline network group.
4. (Optional) Configure heartbeat or negative heartbeat profile.
5. Configure inline tool ports.
6. Configure inline tool.

7. (Optional) Configure inline tool group.
8. (Optional) Configure inline tool series.
9. Configure inline maps, either passall, map (rule-based), map shared collector, or bypass.
10. Configure non-default values for parameters of the inline networks or inline tools.

The summary steps are shown in Figure 25-42.

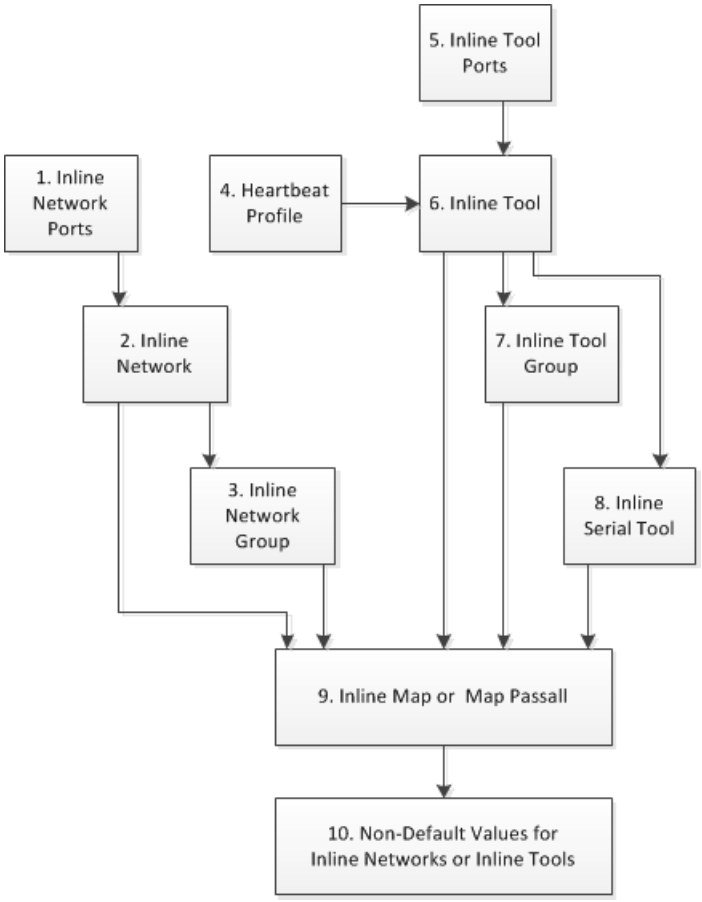

Figure 25-42: Configuration Steps for Inline Bypass Solutions

The configuration details for an inline bypass solution are as follows:

1. Configure inline network ports. (Optional for protected inline network.)

   The configuration begins with defining the inline network ports that will participate in the inline network. Use Quick Port Editor to configure a port type of inline-network.

   For an unprotected inline network, you configure the inline network ports.

   For a protected inline network, the ports are created automatically when the bypass combo modules are recognized by the GigaVUE HC Series node.

   For detailed steps, refer to *Configure Inline Network Ports* on page 620.

2. Configure inline network. (Optional for protected inline network.)

   Next configure the inline network or inline networks using the **inline Network** configuration page and the port pairs defined in step 1.

For an unprotected inline network, you configure the inline network.

For a protected inline network, the inline network is created automatically when the bypass combo modules are recognized by the GigaVUE HC Series node.

In either case, the inline network will have parameters set to default values, such as, the **Traffic Path** will be set to **Bypass** and the **Physical Bypass** will enable.

The initial forwarding state of the unprotected inline network will be DISABLED. The initial forwarding state of the protected inline network will be PHYSICAL BYPASS.

For detailed steps, refer to *Configure Inline Network (Unprotected)* on page 621.

3. (Optional) Configure inline network group.

   If the inline bypass solution involves an inline network group, first configure the participating inline networks before configuring the inline network group. Use the **Inline Network Group** configuration page to configure the inline network group and see the list of inline networks defined in step 2 by selecting **Inline Bypass > Inline Network Groups.**

   For detailed steps, refer to *Configure Inline Network Group* on page 621

4. (Optional) Configure heartbeat or negative heartbeat profile.

   If any of the inline tools will be using a heartbeat profile, a default heartbeat profile is provided, so no configuration is needed except a name. However, if any of the inline tools will be using a heartbeat profile with non-default settings, first configure the heartbeat profile using the **Heartbeats** configuration page, before configuring the inline tools that will use that profile.

   If any of the inline tools will be using a negative heartbeat profile, configure the negative heartbeat profile by providing an alias and a PCAP file using the **Heartbeats** configuration page, before configuring the inline tools that will use that profile

   For detailed steps, refer to *Create Heartbeat Profile* on page 622 and *Create Negative Heartbeat Profile* on page 623.

5. Configure inline tool ports.

   Next configure inline tool ports. Use Quick Port Editor to configure the ports a port type of **Inline Tool**

   For detailed steps, refer to *Configure Inline Tool Ports* on page 623.

6. Configure inline tool.

   Next configure the inline tool or inline tools using the **inline Tool** configuration page and the port pairs defined in step 5.

   For detailed steps, refer to *Configure Inline Tool* on page 624.

7. (Optional) Configure inline tool group.

   If the inline bypass solution involves an inline tool group, first configure the participating inline tools, before configuring the inline tool group. Use the **inline Tool Group** configuration page to configure the inline tool groups and see the list the inline tools defined in step 6 by selecting **Inline Bypass > Inline Tool Groups**.

   For detailed steps, refer to *Create Inline Tool Group* on page 625.

8. (Optional) Configure inline tool series.

   If the inline bypass solution involves an inline tool series, first configure the participating inline tools, before configuring the inline tool series. Use the **inline Serial Tools Group** configuration page and list the inline tools defined in step 6 by selecting **Inline Bypass > Inline Serial Tools**.

9. Configure inline maps, either map passall, map (rule-based), or map shared collector.

   The next configuration step is to configure inline maps that specify how to direct the traffic from the configured inline networks and inline network groups to the configured inline tools, inline tool groups, and inline tool series. You can configure either a map passall, a map (rule-based), or a map shared collector. Create a map with type **Inline** and a subtype of **ByRule**, **Pass All**, or **Collector**.

   For details about configuring maps, refer to *Manage Maps* on page 518.

10. Configure non-default values for parameters of the inline networks or inline tools.

    Now configure non-default values for inline network parameters. For example, for an unprotected inline network, when you change the **Traffic Path** to **To Inline Tool**, traffic will start flowing through the inline tools from the unprotected inline network. For a protected inline network, when you uncheck **Physical Bypass**, traffic will start flowing through the inline tools from the protected inline network.

For protected inline networks, to start the traffic flowing, perform the following steps under Configuration on the Inline Network configuration page:

1. Change the **Traffic Path** to **To Inline Tool**.
2. Uncheck the **Physical Bypass**.

Configuration Step Details

This section provides detailed steps for configuring inline bypass through the GigaVUE-OS H-VUE UI.

Configure Inline Network Ports

Use the following procedure to create inline network ports:

1. Select **Ports > All Ports.**
2. Open Quick Port Editor by clicking the **Port Editor** button.
3. Use the Quick search field to Find the ports to configure.
4. In th an Alias field enter a name to help identify the inline port.
5. For Type, select **Inline Network**.

   An Inline Network (unprotected) is a software arrangement of two network-type ports allocated to facilitate access to a bidirectional link between two networks (far end network devices) that are linked to inline tool ports.

6. Click **OK**.

NOTE: Any available network-type ports on a Gigamon node can be used to form an unprotected inline network.

Configure Inline Network (Unprotected)

An Inline Network Group is an arrangement of multiple inline network ports to which traffic is distributed based on calculated hash values used by a Gigamon node.

Perform the steps to configure an inline network:

1. After configuring the Inline Network ports, select **Inline Bypass > Inline Networks**.
2. Click **New**.
3. In the **Alias** field, enter an alias for the inline network to help identify the inline network.
4. From the **Port A** drop-down list select an inline network port.
5. An inline network port is automatically selected for **Port B**. To select a different port, select one from the **Port B** drop-down list if there is more than one inline network port.
6. Select a traffic path from the **Traffic** Path drop-down list. The types of traffic paths are:
   - **Bypass** — All traffic arriving at the Port A inline network port is directly forwarded to the Port B inline network port and all traffic arriving at the Port B inline network port is directly forwarded to the Port A inline network port.
   - **Drop** — No traffic is exchanged through the inline network ports (all traffic coming to these ports is dropped).
   - **ByPass with Monitoring** — All traffic is forwarded as a forced bypass value and a copy of the traffic is also forwarded to the inline tools. A traffic map must first be configured between the inline network and the inline tool to have the traffic forwarded with no traffic taken from the inline tools.
   - **To Inline Tool** — The traffic received at the inline network ports is forwarded based on:
     1. The traffic map between the inline network and the respective inline tools.
     2. The failover action attributes of the inline tools
     3. The health state of the inline tools.
7. Select the **Link Propagation** check box to enable whether the inline network link on one side of the inline network gets propagated to the other side.
8. Click **Save**.

Configure Inline Network Group

An Inline Network Group is an arrangement of multiple inline network ports to which traffic is distributed based on calculated hash values used by Gigamon node.

Perform the following steps to configure an Inline Network Group:

1. After configuring the Inline Network Ports, select **Inline Bypass > Inline Network Groups**.
2. Click **New**.
3. In the **Alias** field, enter a name for the network group.

4. From the **Inline Network** drop-down list, select the Inline Network ports.
5. Click **Save**.

Create Heartbeat Profile

The Heartbeat Profile is a data structure that contains the heartbeat attributes that are applied to an Inline Tool for configuring its heartbeat. The Create Heartbeat Profile wizard allows you to apply attribute values for the heartbeat profile. Use the following procedure:

1. Select **Inline Bypass > Heartbeats**.
2. Click **New**.
3. In the **Alias** field, enter a name for your heartbeat profile.
4. In the **Type** field, select **Regular**.

   For details about regular heartbeats, refer to *Standard Heartbeat* on page 593.
5. Use **Packet Format** drop-down list to select a packet type. The formats are:
   - **ARP**—This protocol (Address Resolution Protocol) is used for resolution of network layer address into link layer addresses, which is critical for multiple-access network operation. ARP is the default.
   - **Custom**—This format is a binary packet content associated with a packet capture (pcap) file. For details about custom packet format, refer to *Standard or Custom Heartbeat Packet* on page 593.

     When you select **Custom**, a **Custom Format** field displays with a **Browse** button. Use the **Browse** button to upload the pcap file.
6. Use the **Direction** drop-down menu to select the direction for sending heartbeat. The directions are:
   - **A to B**—From Port A to Port B of the inline tool.
   - **B to A**—From Port B to Port A of the inline tool.
   - **Bi-directional**—Both directions.
7. In the **Timeout** field, enter a number in milliseconds to indicate a timeout period for heartbeat packets between sending and receiving. The acceptable range is 20 to 1000 milliseconds. The default is 500 milliseconds.
8. In the **Period** field, enter a number in milliseconds for sending subsequent heartbeat packets. The acceptable range is 30 to 5000 milliseconds. The default is 1000 milliseconds.
9. In the **Recovery Time** field, enter a number in seconds to indicate that the inline tool is declared up with successfully receiving packets. The acceptable range is 5 to 60 seconds. The default is 30 seconds.
10. In the **Retries** field, enter the number for consecutive timed-out heartbeat packets at which the system will trigger (retry) a fail over condition.
11. Click **Save**.

    The heartbeat profile appears in the Heartbeat Profile table.

    NOTE: Highlight the heartbeat profile and click **Edit** to modify the parameters, if needed.

Create Negative Heartbeat Profile

The Negative Heartbeat Profile is a data structure that contains the negative heartbeat attributes that are applied to an Inline Tool for configuring its negative heartbeat. The Create Heartbeat Profile wizard allows you to apply attribute values for the negative heartbeat profile. Use the following procedure:

1. Select **Inline Bypass > Heartbeats**.
2. Click **New**.
3. In the **Alias** field, enter a name for your heartbeat profile.
4. In the **Type** field, select **Negative**

   This is a negative heartbeat profile. For details about negative heartbeats, refer to *Negative Heartbeat Profiles* on page 594.
5. Use **Browse** button in the **Custom Format** field to upload binary packet content associated with a packet capture (pcap) file. For details about the custom format, refer to *Standard or Custom Heartbeat Packet* on page 593.
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. Select **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* on page 605
7. Select the **Per Direction Order**.

   For details about per-direction order, refer to *Inline Tool Series Per-Direction Order* on page 608.

Configure When GigaVUE HC Series Modules are Operationally Up

Ensure that the GigaVUE HC Series modules are in the operationally *up* state before configuring them. Configuration changes done when a module is operationally *down* are not supported.

Also, when an inline tool or inline tool group is in the operationally *down* state, do not modify the current failover action of that inline tool or inline tool group until the tool has recovered from the failover state.

Avoid Oversubscription

In general, traffic received at inline network ports is delivered to the destination ports according to the inline maps and the out-of-band maps regardless of whether the destination ports have the capacity to absorb all the traffic or not.

NOTE: When an inline network is involved in an inline map or an out-of-band map to a destination port (tool port or inline tool port), when there is temporary oversubscription, some packets arriving at the inline network port will be dropped. This can happen when the traffic path is set to bypass or monitoring.

Ensure that destination ports of maps originating from inline network ports have enough capacity to absorb the amount of traffic coming to the inline network ports.

Configure Gigamon Resiliency for Inline Protection

Gigamon Resiliency for Inline Protection (GRIP)™ is an inline bypass solution that connects two GigaVUE nodes together so that one node provides high availability to the other node when there is a loss of power. This redundant arrangement of two GigaVUE nodes maintains traffic monitoring by inline tools when one of the nodes is down.

GRIP makes use of the bypass protection switch relays for protected inline networks on GigaVUE-HC3, GigaVUE-HC2, and GigaVUE-HC1 nodes. The following modules are required to provide physical protection:

- bypass combo modules (BPS), for a protected pair of optical inline network ports on GigaVUE-HC3, GigaVUE-HC2, or GigaVUE-HC1
- TAP-HC0-G100C0 module, for a protected pair of copper inline network ports on GigaVUE-HC2
- TAP-HC1-G10040 module, for a protected pair of copper inline network ports on GigaVUE-HC1

NOTE: GRIP is supported on GigaVUE-HC3 only when there are other modules installed in the node that can provide the stack link. The GRIP solution synchronizes the nodes through a signaling link using a stack link between two stack ports. The BPS-HC3-C25F2G module does not support stack ports, so another module such as PRT-HC3-C08Q08 or SMT-HC3-C05 must be available to be used for that purpose.

In the GRIP solution, two GigaVUE nodes are cabled so that traffic is guided through one GigaVUE node, acting in the primary role, while the other GigaVUE node is on standby, acting in a secondary role. If the primary node fails, the bypass protection switch relays on the modules switch the traffic over from the primary node to the secondary node.

Using the physical protection for either copper or fiber, traffic is guided through inline tools by one of the GigaVUE nodes. The GigaVUE node with the open bypass protection switch relays is the one through which traffic flows. The traffic only flows through one GigaVUE node or the other.

To configure the GRIP solution for copper, use two TAP-HC0-G100C0 modules on GigaVUE-HC2 or two TAP-HC1-G10040 modules on GigaVUE-HC1. The capacity will be 1Gb.

To configure the GRIP solution for fiber, use the following:

- two BPS-HC0-D25A4G, BPS-HC0-D25B4G, or BPS-HC0-D35C4G modules on GigaVUE-HC2. The capacity will be 10Gb.
- two BPS-HC1-D25A24 modules on GigaVUE-HC1. The capacity will be 10Gb.
- two BPS-HC0-Q25A28 modules on GigaVUE-HC2. The capacity will be 40Gb.

- two BPS-HC3-C25F2G modules on GigaVUE-HC3. The capacity will be either 100Gb or 40Gb, depending on the configured port speed of the inline network port pairs.

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

Refer to Figure 25-43 on page 628, Figure 25-44 on page 629, and Figure 25-45 on page 629.

Figure 25-43 on page 628 shows traffic coming from a network (for example, the Internet) through an edge router at the top of the figure. Two GigaVUE nodes with an inline monitoring tool attached to each node are shown in the middle of the figure. Traffic to end devices on a private network are shown at the bottom of the figure.

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

As shown in Figure 25-43, traffic only flows through the node with the primary role. On the primary node, the bypass protection switch relays are open. Traffic is directed to the inline tool attached to the primary node. The node with the secondary role watches the state of the signaling link. If the primary node is up, the link is up, and the secondary node takes no action. The bypass protection switch relays on the secondary node are in a closed state. In Figure 25-43 on page 628, the dotted lines depict the inactive traffic path.

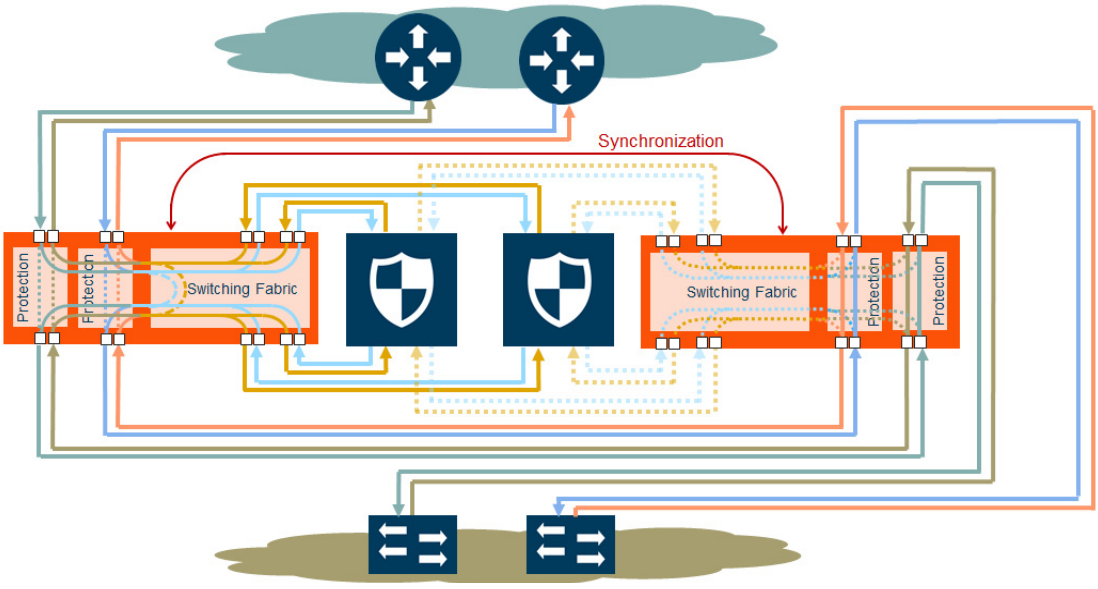

Figure 25-43: Traffic Flows Through Node with Primary Role

In Figure 25-44 on page 629, power is lost to the GigaVUE node in the primary role. The bypass protection switch relays on the primary node close automatically when the node is down. The secondary node receives a signal through the signaling link that the primary node is down. The secondary node opens its bypass protection switch relays. Now traffic flows through the secondary node and traffic is directed to the inline tool attached to the secondary node.

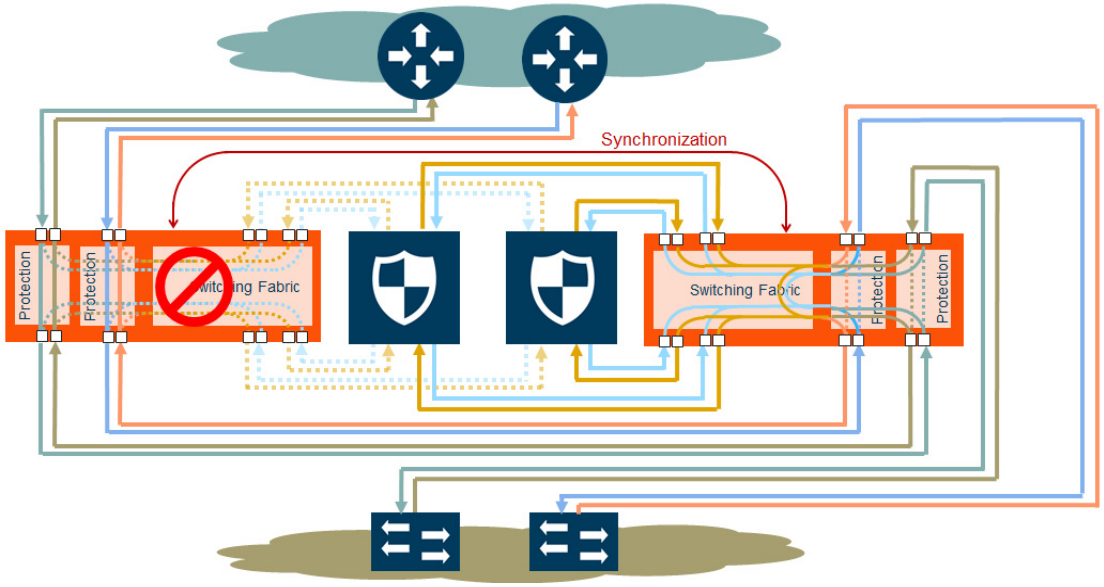

Figure 25-44: Traffic Flows Through Node with Secondary Role after Primary is Lost

In Figure 25-45 on page 629, both nodes have lost power. The bypass protection switch relays are closed on both GigaVUE nodes. Traffic flows between the networks, but without going through the inline tools, which are both bypassed.

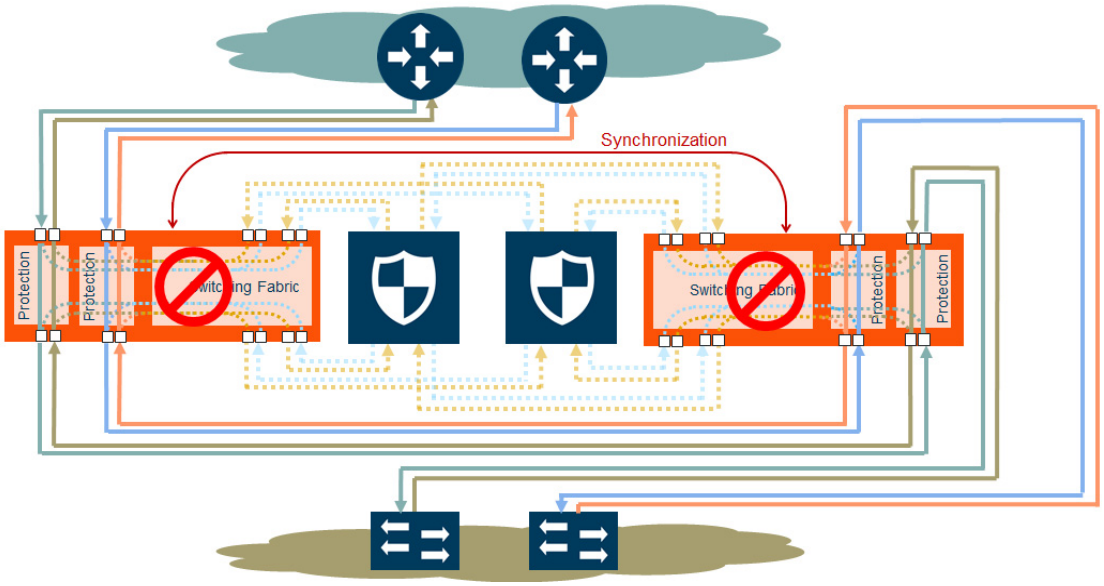

Figure 25-45: Both Nodes Fail; No Traffic Monitoring

How to Handle Recovery

In the scenario in Figure 25-44, after traffic is flowing through the secondary node, at some point the primary node will come back up. The primary node will establish the configured inline traffic paths, bring the signaling link up, and open its relays. Traffic will then flow through the primary node again.

Both Nodes Go Down and Only Secondary Comes Up

In the GRIP solution, if both primary and secondary nodes are powered down or if there is a power outage causing both primary and secondary nodes to go down, powering up the secondary alone without the primary ever coming up will cause network traffic to be bypassed instead of being sent to inline tools.

It is not recommended to power up/recover only the secondary node without the primary. The recommendation is to eventually bring the primary up also.

If the primary node is prone to failures or frequent power outages, another recommendation is to change the role of the secondary node to the primary.

How to Cable GigaVUE Nodes

To cable two GigaVUE nodes, as shown in Figure 25-43 with the primary on the left and the secondary on the right:

- Connect the network shown at the top of Figure 25-43 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 25-43.
- Connect the signaling port on the primary node to the signaling port on the secondary node.

Configure Software

To configure the GRIP solution in software, first specify a name for a redundancy profile by selecting **Inline Bypass > Redundancies**, and then clicking **New**.

The redundancy profile specifies the following:

- **Signaling Port**—specifies the ports used to signal the state of the two GigaVUE nodes to each other. The ports provide the mechanism to detect loss of power in one of the GigaVUE nodes.
- **Protection Role**—specifies the role of the GigaVUE node, as primary, secondary, or suspended. The default is suspended. When suspended, the protection role is on hold. Changing a GigaVUE node from the primary role to the suspended role can be used to manually force the primary node down so the secondary node can become active. The suspended role is also used when performing maintenance. Refer to *Limitation for Suspended Role* on page 631 and *How to Use Suspended Role for Maintenance* on page 632.

The link between the signaling ports on the two GigaVUE-HC nodes is for synchronization. When the node acting in the primary role is up, the signaling link is up, and the node acting in the secondary role sees the link as up. When the primary node loses power, the signaling link is brought down, and the secondary node sees the link as down and takes over.

The redundancy profile combines the protection role with the signaling port. The same redundancy profile is applied to the inline networks, so they have the same properties. If multiple inline networks on each GigaVUE node share the signaling link, they must be configured with the same protection role.

The primary and secondary roles on the two GigaVUE nodes do not change. That is, the role of the primary node stays the same and the role of the secondary node stays the same. The secondary always watches the state of the signaling port for whether the link is up or down.

For example, in Figure 25-44 on page 629, after the primary node recovers, it will open its bypass protection switch relays. Through the signaling port, the primary node will indicate that it is ready to receive traffic by setting the link state to up. The secondary node will notice that the link is up and will close its bypass protection switch relays. After recovery, the primary node automatically goes back into service.

Limitation for Suspended Role

Though GRIP is supported in a cluster, there is a limitation when the suspended protection role is used on the standby node in the cluster. The recommendation is to either switch the standby to the master or apply the suspended role in the redundancy profile to the master node.

Configure Synchronization

You must synchronize the configuration of the two GigaVUE nodes involved in the GRIP solution. The configuration items that must be synchronized are as follows:

- the signaling ports, as dictated by the signaling link cabling
- the inline networks, as dictated by the network path cabling between the two GigaVUE nodes
- the redundancy profiles. The redundancy profile of each GigaVUE nodes needs to have the same signaling port as well as a redundancy role that is compatible with the redundancy role on the other GigaVUE node. For example, one is configured with the primary role and one is configured with the secondary role.
- the inline tools
- the inline maps

For a configuration example of two GigaVUE-HC2 nodes, refer to *Example 4: Gigamon Resiliency for Inline Protection* on page 639. In the example, the configuration is the same on both nodes, except for the protection role (primary or secondary).

Display Redundancy Control State

To display the Redundancy Control State, go to the Inline Networks page and click on the alias of the Inline Network for which you want to display the redundancy control state. The state is displayed on the Quick View under Configuration. Figure 25-46 shows an example where the inline network cu2 has a neutral redundancy control state.

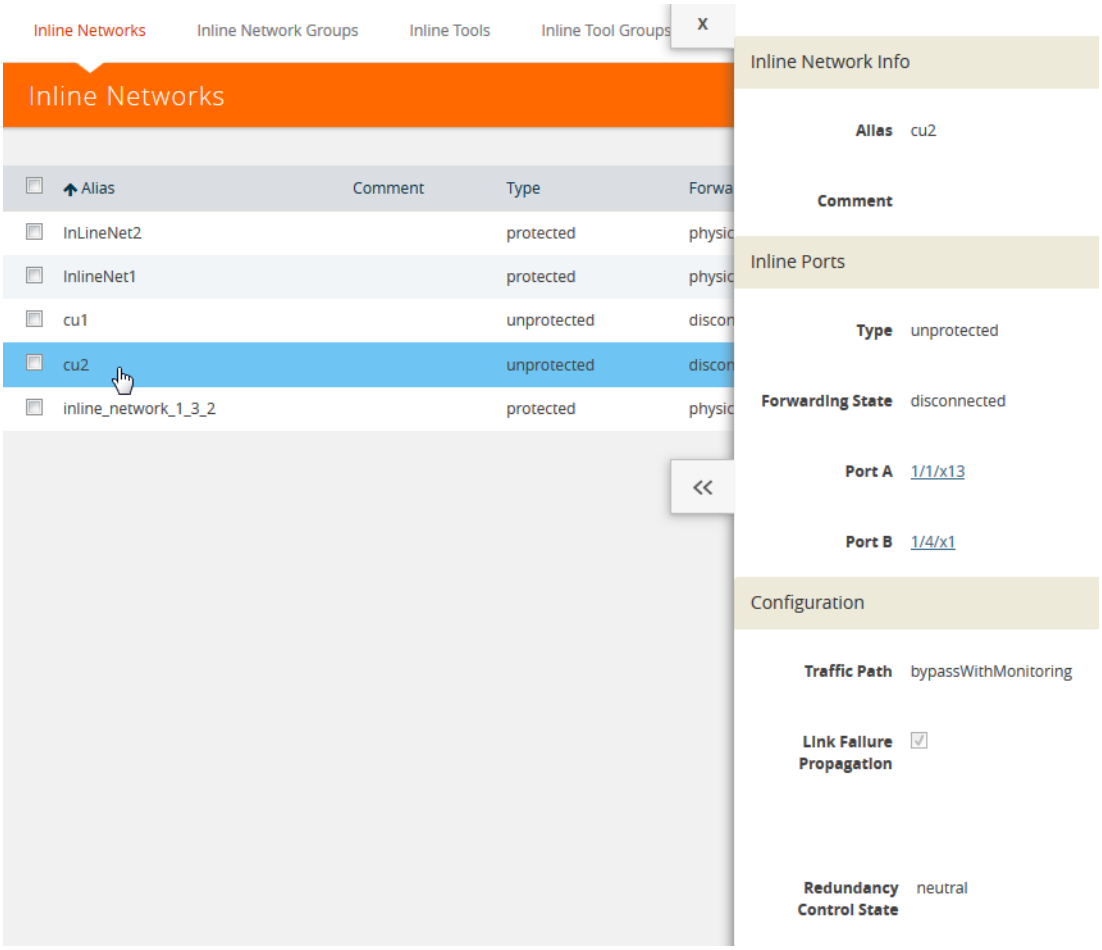

Figure 25-46: Redundancy Control State for Inline Network

Table 25-4 describes each redundancy control state.

Table 25-4: 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-HC2 nodes, use the following steps on that module:

1. Select **Inline Bypass > Redundancies**.
2. On the Redundancies page, select the redundancy profile, and then click **Edit**.
3. For **Protection Role**, select **suspended**, and then click **Save**.

GRIP for Mixed Topologies

GRIP supports mixed topologies. The two nodes in the GRIP configuration do not both have to be the same GigaVUE HC Series nodes. For example, one node can be a GigaVUE-HC1 and the other node can be a GigaVUE-HC2. However the port speed on both nodes (10Gb) must match. In the current software version, GigaVUE-HC3 is not supported in a mixed topology with either GigaVUE-HC1 or GigaVUE-HC2.

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* on page 633
- *Example 2: Unprotected Inline Bypass with Default Heartbeat* on page 635
- *Example 3: Protected Inline Bypass Using Combo Modules* on page 636
- *Example 4: Gigamon Resiliency for Inline Protection* on page 639

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-HC2, 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-HC2 node. Refer to Figure 25-47 on page 634 which shows a GigaVUE-HC2.

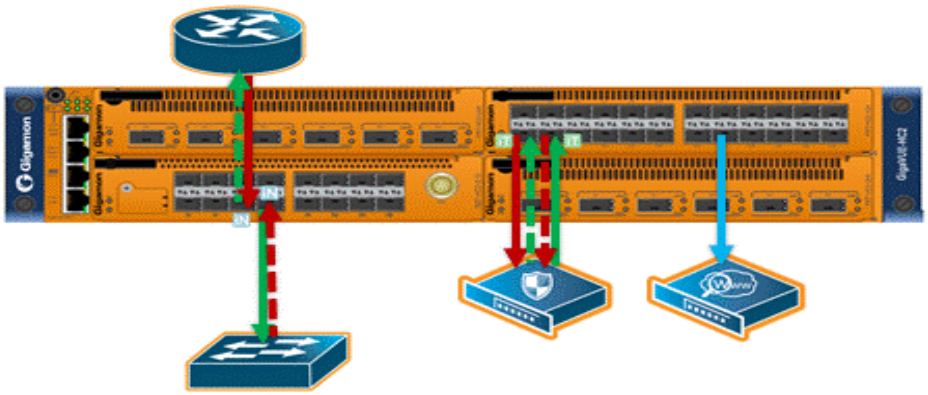

Figure 25-47: Logical Bypass Without Bypass Combo Module

On GigaVUE-HC1, an unprotected inline bypass solution can be configured on the base module, with the inline networks and iniine 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.** Select 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.** Select 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.	**1.** Select 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.

Task	Description	UI Steps
3.	Configure inline tool ports, port type (inline-tool), and administratively enable inline tool ports.	**1.** Select Ports > All Ports. **2.** Click Quick Port Editor. **3.** Use Quick search to find the ports to configure. In this example, the ports are 3/1/x3 and 3/1/x4 **4.** Set port 3/1/x4 to Type Inline-Tool and select Enable. Enter iT1 for the port alias. **5.** Set port 3/1/x4 to Type Inline-Tool and select Enable. Enter iT2 for the port alias. **6.** Make sure Enable is selected for Admin on the ports. **7.** Click OK.
4.	Configure default heartbeat profile.	In GigaVUE-HVUE, the default heartbeat profile is already configured. To view the profile, select Inline Bypass > Heartbeats.
5.	Configure inline tools, and enable it. Also enable the default heartbeat profile.	**1.** Select Inline Bypass > Inline Tools. **2.** Click New. **3.** In the Alias field, type InTool1. **4.** For Port A, select iT1. **5.** For Port B, select iT2. **6.** Under Configuration: **a.** Select Enabled **b.** Select Enabled Heartbeat and set Profile to default. **7.** Click Save.
6.	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.
7.	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 3: Protected Inline Bypass Using Combo Modules

Example 3 is a protected inline bypass solution using bypass combo modules on GigaVUE-HC2. 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-HC2, the port pairs are numbered for example: 2/2/x17 and 2/2/x18, 2/2/x19and 2/2/x20, 2/2/x21 and 2/2/x22, 2/2/x23 and 2/2/x24.

You do not need to configure inline networks because they are also created automatically on bypass combo modules. The aliases of the default inline networks

are: default_inline_net_2_2_1, default_inline_net_2_2_2, default_inline_net_2_2_3, default_inline_net_2_2_4.

On GigaVUE-HC3, protected inline bypass can be configured on the bypass combo module on ports c1..c4.

On GigaVUE-HC1, protected inline bypass can be configured on the bypass combo module. It can also be configured on the TAP-HC1-G10040 module placed in either bay 2 or bay 3, so the ports will be 1/2/g1..g8 or 1/3/g1..g8.

NOTE: The default value of the physical-bypass attribute of protected inline networks is set to enable, which means that the fibers attached to ports net-a and net-b of the inline network are optically coupled and the traffic is exchanged between end nodes without coming to the switching fabric. of the GigaVUE node. As shown in Example 4, after configuring the inline tool and the map passall, the physical-bypass attribute is set to disable in order to activate the inline-bypass solution.

Task	Description	UI Steps
1.	Configure inline too aliases, port type (inline-tool), and administratively enable inline network ports.	**1.** Select Ports > All Ports. **2.** Click Quick Port Editor. **3.** Use Quick search to find the ports to configure. In this example, the ports are 2/2/x11 and 2/2/x12 **4.** Set port 2/2/x11 to Type Inline Tool and select Enable. Enter iT1 for the port alias. **5.** Set port 2/2/x12 to Type Inline Tool and select Enable. Enter iT2 for the port alias. **6.** Click OK.
2.	Configure a heartbeat profile.	**1.** Select Inline Bypass > Heartbeats. **2.** Click New. **3.** In the Alias field, type hb2. **4.** For Type, select Regular. **5.** Click Save.
3.	Configure negative heartbeat profile alias and PCAP file	**1.** Select Inline Bypass > Heartbeats. **2.** Click New. **3.** In the Alias field, type nhb1. **4.** For Type, select Negative. **5.** Click the Browse button for Custom Format and upload the pcap file; for example, hnb.pcap. **6.** Click Save.

Task	Description	UI Steps
4.	Configure inline tools, and enable them. Also specify the heartbeat profile.	**1.** Select Inline Bypass > Inline Tools. **2.** Click New. **3.** In the Alias field, type InTool1. **4.** For Port A, select iT1. **5.** For Port B, select iT2. **6.** Under Configuration: **a.** Select Enabled **b.** Select Enabled Regular Heartbeat and set Profile to hb2. **c.** Select Enabled Negative Heartbeat and set Negative Heartbeat Profile to nhb1. **7.** Click Save. Configure the second inline tool. **1.** Select Inline Bypass > Inline Tools. **2.** Click New. **3.** In the Alias field, type InTool2. **4.** For Port A, select iT3. **5.** For Port B, select iT4. **6.** Under Configuration: **a.** Select Enabled **b.** Select Enabled Heartbeat and set Profile to hb_custom. **7.** Click Save.
5.	Configure the inline tool group and enable it.	**1.** Select Inline Bypass > Inline Tool Group **2.** Click New. **3.** In the Alias field, type inToolGroup. **4.** For Inline Tools, select InTool1 and InTool2. **5.** Under Configuration, select Enabled. **6.** Click Save.
6.	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 default_inline_net_2_2_1 **6.** For Destination, select InTool1 **7.** Click Save.
7.	Configure the path of the traffic to inline tool and disable the physical bypass	**1.** Select Inline ByPass > Inline Networks. **2.** Select the Inline Network default_inline_net_2_2_1 and click Edit **3.** Under Configuration, set the Traffic Path field to To Inline Tool. **4.** Make sure Physical Bypass is not selected. **5.** Click Save.

Example 4: Gigamon Resiliency for Inline Protection

Example 5 is an inline bypass solution for GRIP using TAP-HC0-G100C0 modules on the GigaVUE-HC2, with copper ports.

First, configure the GigaVUE-HC2 with the primary role, then configure the GigaVUE-HC2 with the secondary role. The configuration is the same (is synchronized) on both nodes, except for step 3, in which the protection role (primary or secondary) is specified.

Note that in this example, link fail propagation (LFP) is disabled to reduce inline network recovery time after failover. When a primary to secondary failover occurs and LFP is enabled for copper inline bypass links, network service recovery may take several seconds because of Ethernet link renegotiation. Optical links failover faster and typically recover service much faster. For inline networks where only one path is available, this is a consideration. When GRIP is deployed with high availability networks where a second path is present, it is a best practice to leave LFP enabled.

You can use the Chassis page to view the chassis and modules.

Figure 25-48: Chassis view of GigaVUE-HC2 with TAP-HC0-G100C0 Module

Configure Primary Role GigaVUE-HC2

Task	Description	UI Steps
1.	Configure ports on the TAP-HC0-G100C0 module as passive (in passive mode, relays are closed). Also configure ports, port type (inline-network).	**1.** Select Ports > Ports> All Ports. **2.** Select a port of the TAP-HC0-G100C0 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-HC0-G100C0 module **7.** Configure Inline Network ports **a.** Select the port. **b.** Click Quick Port Editor. **d.** 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 Inline Bypass > Redundancies **2.** Click New. **3.** Enter a name for the profile in the Alias field. For example, RP_001. **4.** Click in the Signaling Port field and select the stack port configured in Task 2. **5.** Select Primary for Protection Role. **6.** Click Save.
4.	Configure inline network. This step associates the redundancy profile to the inline network and also disables link fail propagation on the inline network.	**1.** Select Inline Bypass > Inline Tools. **2.** Click New. **3.** Enter a name for the inline network in the Alias field. For example, IN_001. **4.** Click in the Port A field and select inline network port for network A. **5.** Click in the Port B field and select an inline network port f. **6.** Select Bypass for Traffic Path. (This is the default setting.) **7.** Make sure Link Failure Propagation option is NOT checked. (It is enabled by default.) **8.** Click in the Redundancy Profile field and select the profile created in step 3. For example RP_001.

Task	Description	UI Steps
5.	Configure inline tool ports, port type (inline-tool), and administratively enable them.	**1.** Select Ports > Ports > All Ports. **2.** Select the first port (for example, 1/4/x1) to configure as an inline-tool port. **3.** Click Quick Port Editor. **4.** Select Inline Network for Type and select Enable for Admin. **5.** Click OK. **6.** Select the second port (for example, 1/4/x2) and repeat steps 3 through 5.
6.	Configure inline tool and failover action. then enable inline tool.	**1.** Select Inline Bypass > Inline Tools **2.** Enter a name for the inline tool in the Alias field. For example, IT_001 **3.** Click in the Port A field and select the first inline tool port configured in Task 5. **4.** Click in the Port B field and select the second inline tool port configured in Task 5. **5.** Make sure Enable is selected. (It is enabled by default **6.** Select NetworkBypass for Failtover action. **7.** Click Save.
7.	Configure map passall, from inline network to inline tool. **NOTE:** When you delete a map on the primary node, irrespective of the inline-network traffic-path, the traffic is switched to the secondary node. The port utilization must be 0% on the primary node and active on the secondary node.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Enter a name for the map in the Alias field. For example, INtoIT. **4.** Select Regular for Type. **5.** Select Pass All for Subtype. **6.** Select the inline network created in Task 4 for Source. For example, IN_001. **7.** Select the inline tool created in Task 6. For example, IT_001.

Configure Secondary Role GigaVUE-HC2

Task	Description	UI Steps
1.	Configure ports on the TAP-HC0-G100C0 module as passive (in passive mode, relays are closed). Also configure ports, port type (inline-network).	**1.** Select Ports > Ports> All Ports. **2.** Select a port of the TAP-HC0-G100C0 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-HC0-G100C0 module **7.** Configure Inline Network ports **a.** Select the port. **b.** Click Quick Port Editor. **e.** Select Inline Network for Type. **NOTE:** You can use the Chassis page to locate the position of the module of the module in the chassis and identify port IDs

Task	Description	UI Steps
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 Inline Bypass > Redundancies **2.** Click New. **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 Save.
4.	Configure inline network. This step associates the redundancy profile to the inline network and also disables link fail propagation on the inline network.	**1.** Select Inline Bypass > Inline Tools. **2.** Click New. **3.** Enter a name for the inline network in the Alias field. For example, IN_001. **4.** Click in the Port A field and select inline network port for network A. **5.** Click in the Port B field and select an inline network port f. **6.** Select Bypass for Traffic Path. (This is the default setting.) **7.** Make sure Link Failure Propagation option is not checked. (It is enabled by default.) **8.** Click in the Redundancy Profile field and select the profile created in step 3. For example RP_001.
5.	Configure inline tool ports, port type (inline-tool), and administratively enable them.	**1.** Select Ports > Ports > All Ports. **2.** Select the first port (for example, 1/4/x1) to configure as an inline-tool port. **3.** Click Quick Port Editor. **4.** Select Inline Network for Type and select Enable for Admin. **5.** Click OK. **6.** Select the second port (for example, 1/4/x2) and repeat steps 3 through 5.
6.	Configure inline tool and failover action. then enable inline tool.	**1.** Select Inline Bypass > Inline Tools **2.** Enter a name for the inline tool in the Alias field. For example, IT_001 **3.** Click in the Port A field and select the first inline tool port configured in Task 5. **4.** Click in the Port B field and select the second inline tool port configured in Task 5. **5.** Make sure Enable is selected. (It is enabled by default **6.** Select NetworkBypass for Failtover action. **7.** Click Save.

Task	Description	UI Steps
7.	Configure map passall, from inline network to inline tool.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Enter a name for the map in the Alias field. For example, INtoIT. **4.** Select Regular for Type. **5.** Select Pass All for Subtype. **6.** Select the inline network created in Task 4 for Source. For example, IN_001. **7.** Select the inline tool created in Task 6. For example, IT_001.
8.	Configure the path of the traffic to the inline tool, disabling physical bypass on the inline network to open the relay on the node with the primary role.	**1.** Select Inline Bypas > Inline Networks. **2.** Select the Inline Network configured in Task 4 and click Edit. **3.** Select To Inline Tool for Traffic Path **4.** Make sure Physical Bypass is NOT checked. **5.** Click Save.

26 Work With Inline SSL Decryption

This chapter describes SSL decryption for inline and out-of-band tools, referred to as inline SSL decryption. It provides introductory material as well as configuration examples for GigaVUE-FM. Refer to the following sections for details:

About Inline SSL Decryption

This section introduces inline SSL decryption. Refer to the following sections for details:

SSL Decryption for Inline Tools

SSL decryption for inline tools provides visibility into encrypted traffic. Inline SSL decryption delivers decrypted packets to tools that can be placed inline or out-of-band. The tools look into decrypted packets for threats, such as viruses or other malware.

The amount of Internet traffic that is encrypted is increasing, and much of it is encrypted with Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols.

Malware increasingly uses encrypted SSL traffic, thus a significant percentage of attacks hide in SSL. Inline SSL decryption offers visibility into encrypted applications and hidden threats in your organization.

Many applications, such as email, also use SSL. Encryption protects data from being viewed in transit over the Internet such as in an exchange of emails. Encryption also keeps the data private. But when data is encrypted, packets are not inspected, which can create blind spots in your network.

Providing visibility into encrypted traffic eliminates this blind spot. SSL/TLS blind spots in your network can be eliminated across any port or application, for example, port 443, or email, Web, or VoIP applications.

Inline SSL decryption differs from the existing GigaSMART SSL decryption application, which is passive. Out-of-band SSL decryption delivers decrypted traffic to out-of-band tools that can then detect threats entering the network. When a threat is detected, the tools can send a notification.

Inline SSL decryption offloads the decryption task so that tools can inspect traffic easily and effectively. The advantage of operating inline is that tools can act when a threat is detected.

Also, the inline SSL decryption solution is able to decrypt Perfect Forward Secrecy (PFS) ciphers, for example, ECDHE-RSA-AES256-SHA384 and DHE-RSA-AES128-SHA256. Out-of-band SSL decryption does not decrypt PFS ciphers.

What Inline SSL Decryption Provides

Inline SSL decryption provides the following:

- Identifies/detects encrypted traffic flows (SSL traffic) in a network across any port.
- Intercepts encrypted traffic flows between a client and a server.
- Filters encrypted traffic flows based on policy. For example, if the encrypted traffic flows contain health care or financial information, let those flows bypass decryption.
- Decrypts packets. Inline SSL decryption decrypts packets once at a single decryption point.
- Delivers decrypted traffic flows to multiple security tools. The tools can be inline or out-of-band. The tools can detect threats such as malware in the decrypted traffic flows.

- Re-encrypts traffic flows after receiving them back from the inline tools.
- If a tool acts on traffic flows based on the threats it finds, when malware is found in the decrypted traffic flows, the tool can:
  - modify the traffic flows
  - terminate the connection
- If the tool modifies the packets, GigaSMART will re-encrypt them. If the tool terminates the connection, GigaSMART will terminate the connection between the client and the server.

When SSL traffic is decrypted, sensitive data will be exposed in the connected tools. For example, if email traffic is decrypted, user passwords might be exposed or if financial data is decrypted, social security numbers might be exposed in the decrypted traffic.

Because SSL connections might carry sensitive data, not all connections should be inspected. Some of the SSL connections carrying user data such as financial or medical information should be bypassed without inspection, based on a configured policy.

Inline SSL decryption addresses acceptable use policies and adheres to privacy and compliance requirements. It offers advanced controls to select the traffic to decrypt.

NOTE: Throughout this document, the terms Secure Sockets Layer (SSL) and Transport Layer Security (TLS) are used interchangeably.

Example Inline SSL Decryption

Refer to Figure 26-1 for an example of inline SSL decryption.

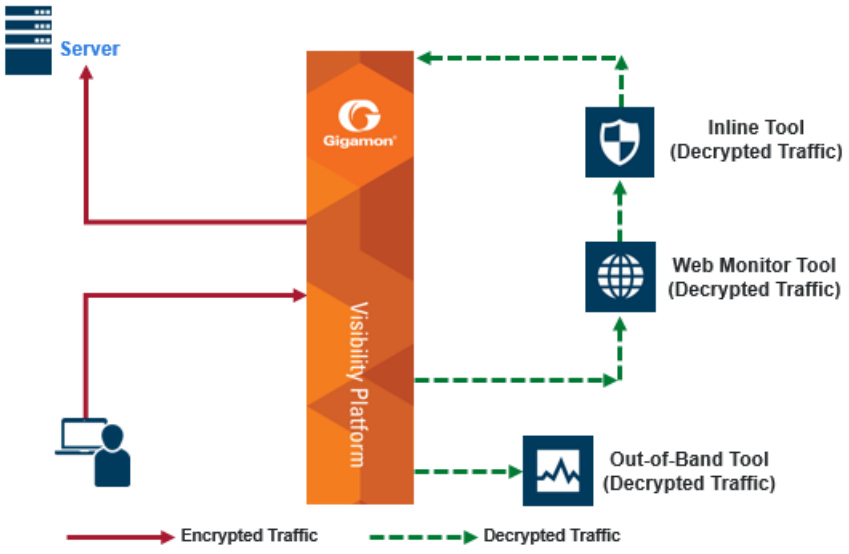

Figure 26-1: Inline SSL Decryption Example, Outbound

Use case for inline SSL decryption:

- Clients in internal network

- Servers on the Internet
- Organization does not have the private key of the server
- Diffie-Hellman and Perfect Forward Secrecy is being used

Figure 26-1 shows the inline SSL decryption solution with the client and GigaVUE node within an enterprise. The server, in the top left of the figure, resides on the Internet. This is an example of an outbound deployment.

The client is on the lower left of the figure. The client is a user who is, for example, using a browser to go to a website on the Internet, such as a bank or a search engine. The traffic from the user could be encrypted or it might not be encrypted. For example, the user might be going to a bank website using the HTTPS protocol or going to a search engine website using the HTTP protocol. The solid line from the user to the GigaVUE node represents encrypted traffic, but there might also be traffic from the user that is not encrypted. Traffic that is not encrypted can either be bypassed or it can go to tools for inspection.

In Figure 26-1, instead of the user interacting directly with the server at the top left, the GigaVUE node is placed in the middle. Thus, the GigaVUE node intercepts the client/server session.

In the GigaVUE node, encrypted packets are identified, then filtered. Selected packets are decrypted and sent to tools for inspection. The dotted lines represent decrypted packets. Packets are decrypted once, then the same decrypted packets can be sent to inline tools and/or out-of-band tools connected to the GigaVUE node.

The traffic from the inline tools is returned to the GigaVUE node to be re-encrypted and then sent to the destination on the Internet, for example, the website that the user is visiting. The solid line to the server represents traffic that has been re-encrypted in the GigaVUE node.

Deploy Inline SSL Decryption

There are two ways to deploy inline SSL decryption as follows:

- sessions are inbound
- sessions are outbound

Refer to Figure 26-2 for an example of an inbound deployment. The client is on the Internet. The server and the GigaVUE node are located within the same enterprise network, with the GigaVUE node deployed on the server side. The GigaVUE node needs access to the private keys of the server to perform Man-in-the-Middle (MitM) decryption.

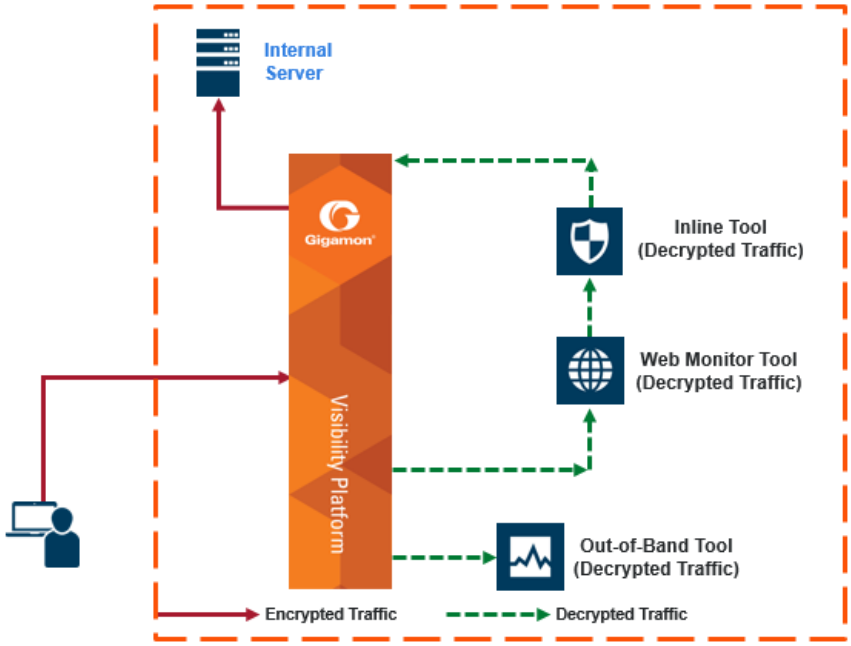

Figure 26-2: Inbound Deployment of Inline SSL Decryption

Use case for inline SSL decryption:

- Clients on the Internet
- Servers in internal network
- Organization has the private key of the server
- Diffie-Hellman and Perfect Forward Secrecy is being used

Refer to Figure 26-3 for an example of an outbound deployment. The client and the GigaVUE node are located within the same enterprise network, with the GigaVUE node deployed on the client side. The server is located in another network on the Internet. In this deployment, the role of the GigaVUE node is that of a Man-in-the-Middle (MitM). In this deployment, the GigaVUE node does not have access to the private keys of the server, but as a trusted MitM, the GigaVUE node can look at SSL traffic.

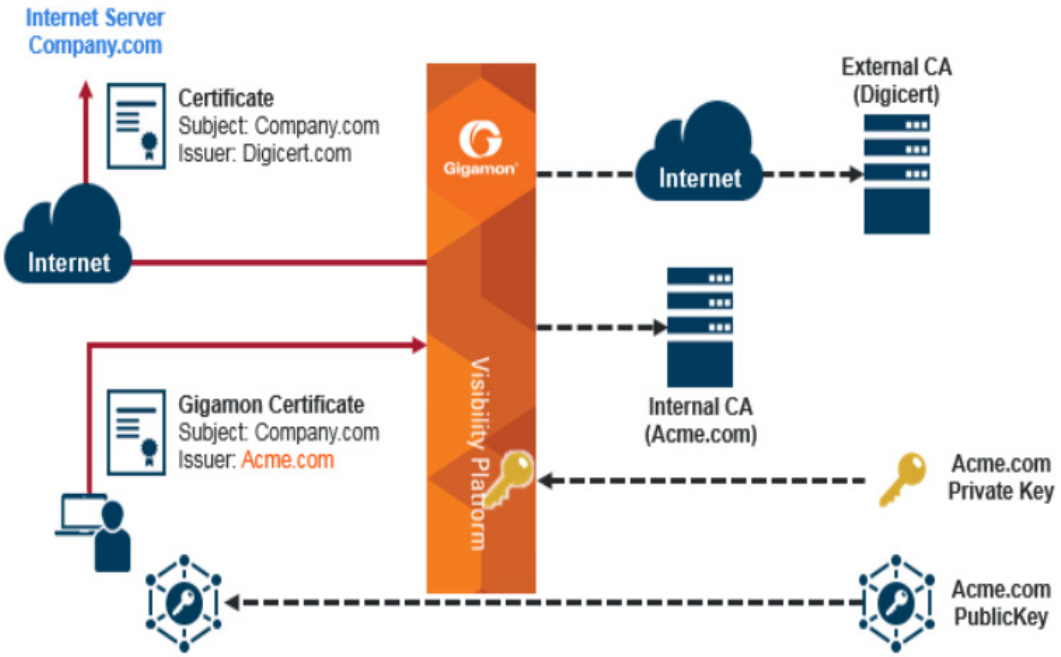

Figure 26-3: Outbound Deployment of Inline SSL Decryption

GigaVUE Modules for Inline SSL Decryption

Inline SSL decryption is supported on GigaVUE-HC1, GigaVUE-HC2, and GigaVUE-HC3 nodes, with the GigaSMART and inline bypass modules installed on the same node.

For physical inline bypass, install a fiber bypass (BPS) combo module. On GigaVUE-HC1 and GigaVUE-HC2, a copper TAP also supports physical bypass.

Table 26-1 lists the inline bypass modules.

Table 26-1: Inline Bypass Modules

GigaVUE Node	Description
GigaVUE-HC1	Bypass Combo Module 10Gb SX/SR (50/125µm multi-mode)
	Copper TAP module

Table 26-1: Inline Bypass Modules

GigaVUE Node	Description
GigaVUE-HC2	Bypass Combo Module 10Gb SX/SR (50/125µm multi-mode)
	Bypass Combo Module 10Gb SX/SR (62.5/125µm multi-mode)
	Bypass Combo Module 10Gb LX/LR (single mode)
	Bypass Combo Module 40Gb SR4 (multi-mode)
	Copper TAP module
GigaVUE-HC3	Bypass Combo Module 100Gb/40Gb SR4 MPO

Figure 26-4 shows a GigaVUE-HC2 with the GigaSMART module and the inline bypass (BPS) module. The GigaSMART module contains the SSL decryption software. The inline network ports are on the inline bypass module. The inline and out-of-band tool ports are on the same GigaVUE node.

For inline traffic, the inline network ports and the inline tool ports each need two links, called port pairs, for bidirectional traffic. For out-of-band (offline) traffic, only one link is needed, because the traffic is not bidirectional.

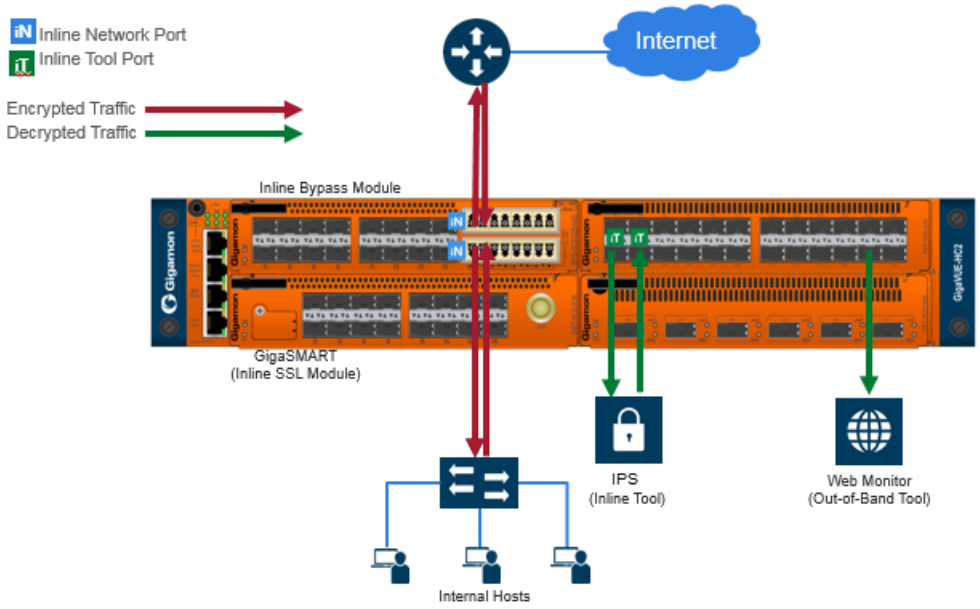

Figure 26-4: GigaVUE Modules: GigaSMART and Inline Bypass

Packet Flows

Normally, a client and server talk directly to each other, such as when you are using a browser to go to a bank website, a health care provider, or a search engine.

As shown in Figure 26-3, the GigaVUE node is placed in the middle between the client and the server. All traffic from the Internet goes through the GigaVUE node.

Incoming traffic arrives on an ingress inline network port on the inline bypass module.

SSL traffic, which is TCP traffic, is directed to the GigaSMART module. Until the traffic is processed by the GigaSMART module, it is not known if it is SSL traffic or not.

The GigaSMART module decides what traffic is bypassed, what traffic is sent to tools without decryption, and what traffic is decrypted and then sent to tools. So, there are three types of decisions as follows:

- to bypass or not
- to decrypt or not
- to send to tools or not

Figure 26-5 shows the flow for a configuration consisting of a single inline network, a single inline tool, and a single out-of-band tool.

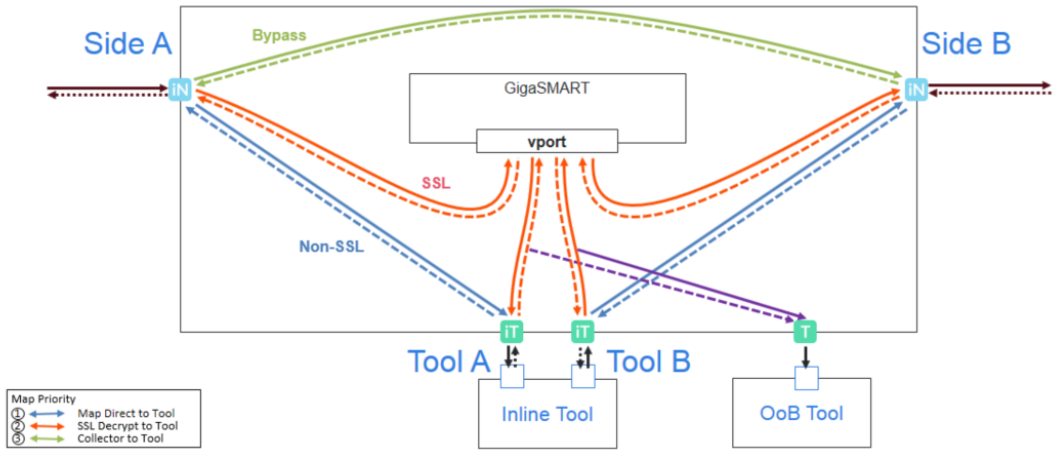

Figure 26-5: Packet Flow for Inline SSL Decryption

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 26-5. 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 26-5.

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 26-5. 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 26-5. 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 26-5.

The out-of-band tool can also receive the traffic as shown by the purple line on the right in Figure 26-5.

Starting in software version 5.2, an out-of-band map from a virtual port to a single tool port is supported. Starting in software version 5.3.01, an out-of-band map from a virtual port to multiple tool ports is supported. The ports can be tool, hybrid, or GigaStream. Out-of-band maps from a vport to port groups are also supported when the ports in the group are tool or hybrid.

Modules Matrix

For traffic that is encrypted, there is a question as to whether or not it needs to be decrypted. For example, is there a policy for or against decrypting that traffic? There might be a policy, such as do not decrypt financial or health care traffic, or there might be a blacklist that states that traffic from a particular site should always be decrypted, or there might be a whitelist that states that traffic from a particular site should always be bypassed. So encrypted packets need to be filtered because some packets will not be decrypted, while others will be decrypted.

Table 26-2 is a matrix of the GigaVUE module used for different types of traffic.

Table 26-2: GigaVUE Modules Used for Type of Traffic

Type of Traffic	GigaSMART	Inline Bypass	GigaSMART
Bypassing GigaSMART	No	Yes - to network	No
Not encrypted (non-SSL)	Yes	Yes - to tools then to network	Yes, depending on the configuration
Encrypted (SSL), and to be decrypted	Yes	Yes - to tools	Yes - to be re-encrypted, then to network

The GigaSMART module does the decryption as well as handling policies, whitelists, and blacklists. The decision to decrypt or not is made in the GigaSMART module.

Filter Traffic in GigaSMART

Because SSL connections can carry sensitive data, not all connections should be inspected. Some of the SSL connections carrying user data such as financial or health care information should be bypassed without inspection, based on a configured policy.

Based on the decryption policies, some connections will not be decrypted and will be passed on as is. The inline SSL decryption solution respects data privacy and supports compliance.

Inline SSL decryption provides different ways to filter traffic, as follows:

- Whitelists specify traffic to always pass through. A whitelist policy states that traffic from certain sites should always skip decryption. Refer to *Whitelisting Policy* on page 654.

- Blacklists specify traffic to always decrypt. A blacklist policy states that traffic from certain sites should always be decrypted. Refer to *Blacklisting Policy* on page 654.
- URL categorization categorizes 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 is not decrypted. Refer to *URL Categorization* on page 654.
- Policy rules based on network attributes, such as source or destination IP address.

Whitelisting Policy

The whitelisting policy allows certain classes such as sites, domains, and host-based IP address of traffic to bypass decryption. By default, traffic that is not to be decrypted is forwarded to the tools unless otherwise configured.

A whitelist file can contain a maximum of 10,000 domain or hostname entries.

Blacklisting Policy

The blacklisting policy allows traffic from certain sites, domains, and host-based IP address to always be decrypted. Blacklisted domains and host names will always be decrypted.

A blacklist file can contain a maximum of 10,000 entries.

URL Categorization

URL categories make it convenient to apply policies on thousands of possible URLs by simplifying the number of policy rules. Categorization is based on the hostname in the TLS Server Name Indication (SNI) or the hostname from the server certificate if there is no SNI.

The URL categorization service is provided by Webroot. 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. For cloud lookups, the stack port interface on GigaSMART must be configured to provide Internet access. Refer to *Configure Stack Port Interface* on page 680 for more information.

SSL Sessions

Secure Sockets Layer (SSL) is a protocol that allows the transmission of secure data between a server and client. Transport Layer Security (TLS) is a cryptographic protocol that adds security to TCP/IP communication.

Inline SSL decryption supports SSL version 3.0 and TLS versions 1.0, 1.1, and 1.2.

TLS and SSL are used in communications such as Web browsing, email, instant messaging, and voice over IP (VoIP). TLS and SSL encrypt these communications.

The client initiates the SSL session. The GigaVUE node intercepts the connection and negotiates an SSL session with the client.

The GigaVUE node monitors all TCP connections, then intercepts the SSL session. Non-TCP traffic is passed transparently without any changes.

All the incoming SSL traffic terminates on the GigaVUE node. The SSL connections are decrypted in inbound or outbound deployments, passed to the inline tools, and eventually to the server.

The session to the client is terminated on the GigaVUE node, but information about the session, such as the initiator's IP address is maintained, so that the GigaVUE node can "reconnect" the client and server.

The GigaVUE node performs SSL decryption and feeds tools, either inline or out-of-band.

The session to the server is from the GigaVUE node to the server. The GigaVUE node negotiates a new SSL session with the server.

SSL Handshake

SSL encryption secures traffic between a client and a server, such as a Web server. 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 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 26-6 shows the basic SSL handshake between a client and server to establish a session.

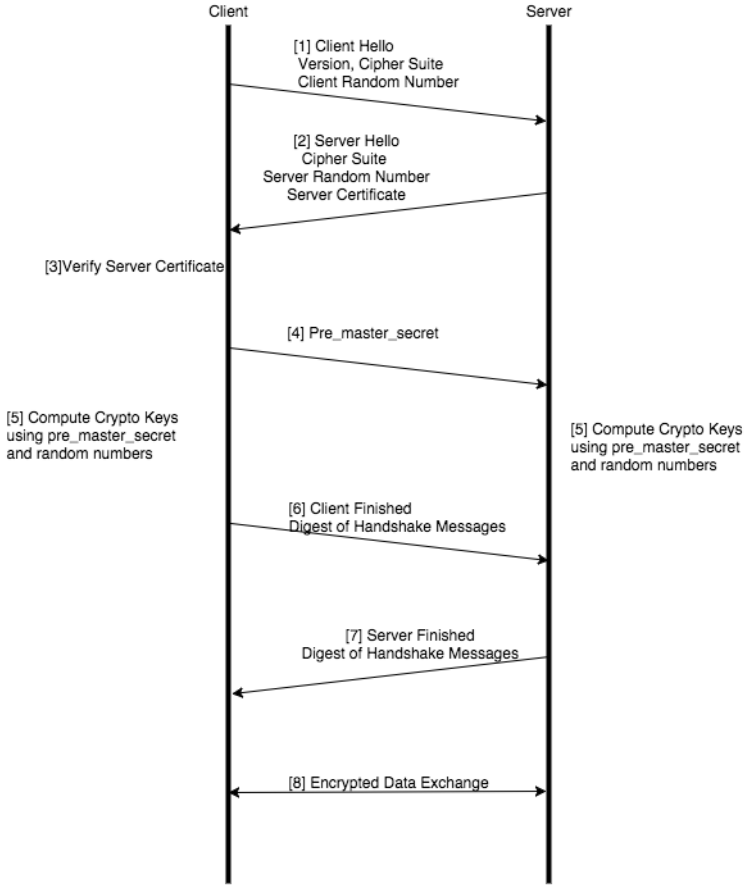

Figure 26-6: Basic SSL Handshake

SSL Handshake Steps

The SSL handshake steps in Figure 26-6 are as follows:

1. The SSL or TLS client sends a Client Hello message that contains information such as the SSL or TLS version and the list of cipher suites supported by the client. The message also contains a client random number.
2. The SSL or TLS server responds with a Server Hello message that contains the SSL or TLS version it supports and the cipher suite chosen by the server from the list provided by the client, and a server random number. The server also sends its digital certificate.
3. The SSL or TLS client verifies the server's digital certificate.
4. Using the random numbers from the Hello messages, the SSL or TLS client computes the pre_master_secret that enables both the client and the server to compute the secret key to be used for encrypting subsequent message data. The pre_master_secret is encrypted with the server's public key and sent.
5. The SSL or TLS server verifies the client's certificate.

6. The SSL or TLS client sends the server a Client Finished message containing a digest (MAC) of the messages in the handshake, which is encrypted with the secret key, indicating that the client part of the handshake is complete.
7. The SSL or TLS server sends the client a Server Finished message containing a digest (MAC) of the messages in the handshake, which is encrypted with the secret key, indicating that the server part of the handshake is complete.
8. For the duration of the SSL or TLS session, the server and client can now exchange messages that are symmetrically encrypted with the shared secret key.

Figure 26-7 shows the SSL handshake when a Man-in-the-Middle sits between the client and server.

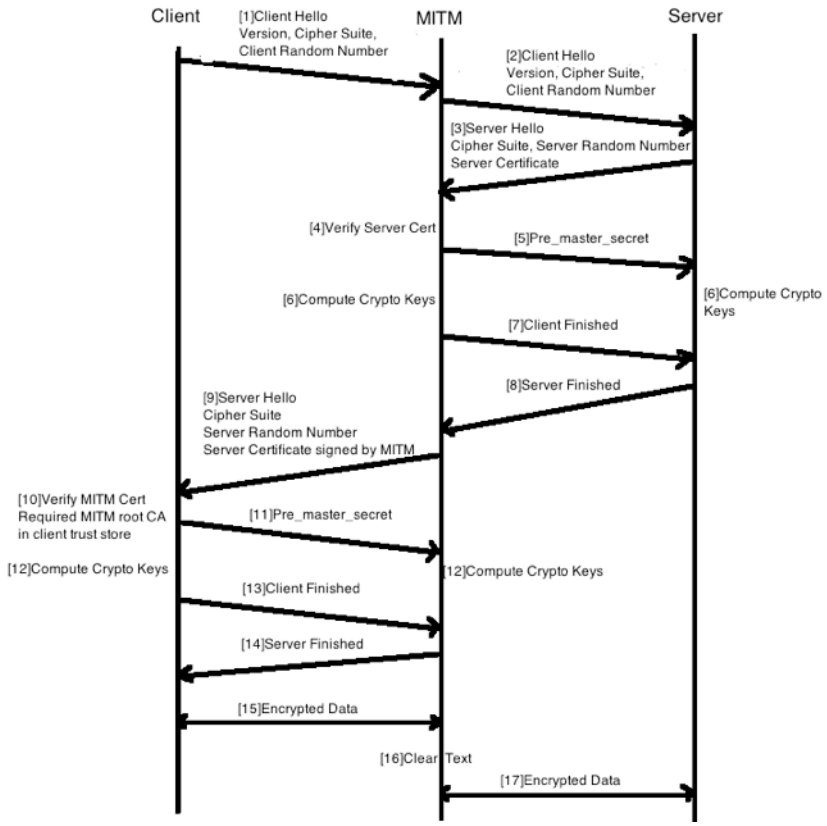

Figure 26-7: Man-in-the-Middle SSL Handshake

SSL Session, Inbound Deployment

Figure 26-2 on page 649 shows an inline SSL decryption inbound deployment. The client is on the Internet. The server and the GigaVUE node are located within the same enterprise network, with the GigaVUE node deployed on the server side. The GigaVUE node needs access to the private keys of the server.

The SSL session is created as follows:

1. Client traffic, such as from the Internet, arrives on an inline-network port on the GigaVUE node and establishes a TCP connection.

2. The GigaVUE node initiates a TCP connection to the server.
3. Gigamon establishes the SSL handshake with the server parameters from the SSL client's Hello request.
4. Gigamon establishes the SSL handshake with the client using the server's certificate and key, and using the appropriate parameters.

SSL Session, Outbound Deployment

Figure 26-3 on page 650 shows an inline SSL decryption outbound deployment. The client is in your network, and it connects to a server outside your network on the Internet. Traffic is destined to servers where you not have access to the private keys.

The SSL session is created as follows:

1. Client traffic, such as from an employee on the enterprise's Intranet is destined to a server on the Internet. The traffic arrives on a network port on the GigaVUE node.
2. Gigamon initiates a new connection to the server as the client.
3. The server responds with its server certificate and presents it to Gigamon.
4. Gigamon spoofs the server certificate and presents it to the client, but now the certificate is signed by Gigamon.
5. The end client verifies that the certificate is valid as it belongs to the server and has been signed by Gigamon, which has been listed as one of the valid CAs.
6. Gigamon now maintains two TCP connections, one with the client and one with the server.

SSL Session Resumption

SSL sessions can be resumed to improve performance. SSL session resumption speeds up the SSL handshake.

Once a session has been established, the keys are saved so a session can be resumed efficiently later. The resumed SSL handshake has fewer steps.

Session identifier-based resumption is supported. The GigaVUE node maintains the session identifier data in the cache. Session ticket-based resumption is not supported.

By default, resumption is enabled.

SSL Session Search

Starting in software version 5.2, you can search an existing session based on a hostname. The input is matched against the Server Name Indication (SNI) or the certificate subject name of the current sessions.

StartTLS

StartTLS provides a mechanism for protocols such as SMTP, IMAP, and POP3 to start in plaintext mode and upgrade to encrypted mode on the existing port instead of using another port.

The startTLS command is initiated from the client or the server. Upon receiving a response from either end, TLS handshake messages are exchanged between the client and the server.

To support startTLS, the GigaVUE node monitors the TCP connections for protocols, such as SMTP, IMAP, and POP3. Upon detecting the startTLS message on a port, the processing of packets for decryption can start.

Starting in software version 5.2, the specific ports to monitor startTLS traffic must be specified. Up to 20 ports can be monitored.

Inline SSL Decryption Behavior with StartTLS

For connections that use StartTLS to upgrade from non-TLS mode to secure mode, the inline SSL decryption solution decrypts correctly if the decision to decrypt or not is made in the certificate phase.

If the CLIENT HELLO packet does not have SNI information, the inline SSL decryption solution will apply policy rules in certificate phase of the policy evaluation.

For explicit proxy connections policy, rules are applied in certificate phase of policy evaluation. For information on the certificate phase of policy evaluation, refer to *Policy Evaluation* on page 668.

SSL Terminology and Acronyms

Table 26-3 provides definitions of SSL terminology:

Table 26-3: SSL Terminology

Term	Definition
Plaintext	The original, unencrypted data.
Ciphertext	The encrypted data.
Cryptography	The practice of secure communications.
Encryption	The process of turning plaintext into ciphertext.
Decryption	The process of turning ciphertext into plaintext.
Encryption algorithm	The algorithm used to perform encryption and decryption. It is also called the cipher.
Encryption key	The key used for encryption.
Decryption key	The key used for decryption.
Symmetrical encryption algorithm	The algorithm used for encryption in which the encryption key and the decryption key are identical.
Asymmetrical encryption algorithm	The algorithm used for encryption in which the encryption key and the decryption key are different.
Public key	The key used for encryption.
Private key	The key used for decryption.

Table 26-4 lists SSL acronyms:

Table 26-4: SSL Acronyms

Acronym	Definition
AES	Advanced Encryption Standard
CA	Certificate Authority
CBC	Cipher Block Chaining
CN	Common Name
CRL	Certificate Revocation List
DES	Data Encryption Standard
DH, D-H	Diffie-Hellman
DSA	Digital Signature Algorithm
ECDH	Elliptic Curve Diffie-Hellman
ECDSA	Elliptic Curve Digital Signature Algorithm
FQDN	Fully Qualified Domain Name
IDS	Intrusion Detection System
IPS	Intrusion Prevention System
MAC	Message Authentication Code
MD	Message Digest
MitM	Man-in-the-Middle
OCSP	Online Certificate Status Protocol
OoB	Out-of-Band
PEM	Privacy Enhanced Mail
PFS	Perfect Forward Secrecy
PKCS12	Public Key Cryptography Standard #12
PKI	Public Key Infrastructure
RSA	Rivest-Shamir-Adleman
SHA	Secure Hash Algorithm
SSL	Secure Sockets Layer
TCP	Transmission Control Protocol
TLS	Transport Layer Security
URL	Uniform Resource Locator

Keys and Certificates

The SSL protocol allows the transmission of secure data between a server and client who both have the keys to decode the transmission and the certificates to verify trust between them.

An SSL certificate is a digital document containing a public key, host information, and a digital signature from the certificate issuer, known as a Certificate Authorities (CAs). The certificate allows trust to be established between two communicating endpoints.

The inline SSL decryption solution has a trust store, which is a collection of certificates of CAs. Gigamon only trusts server certificates that have a trust anchor in the configured trust store, in other words, if the certificate chain can be built with one of the root CAs in the trust store. Gigamon ships with a default trust store, which you can replace.

The inline SSL decryption solution acts as a Man-in-the-Middle (MitM). In the outbound deployment case, the MitM generates server certificates on-the-fly signed by the installed CA. In the inbound deployment case, server certificate generation is not needed but the server's private key and certificate chain need to be made available to the MitM.

The inline SSL decryption solution also has a key store, which is a collection of SSL private keys (for inbound deployments) and SSL certificates and corresponding private keys that are used to digitally sign the emulated server certificates (for outbound deployments).

Figure 26-8 shows a sample certificate and its relevant parts.

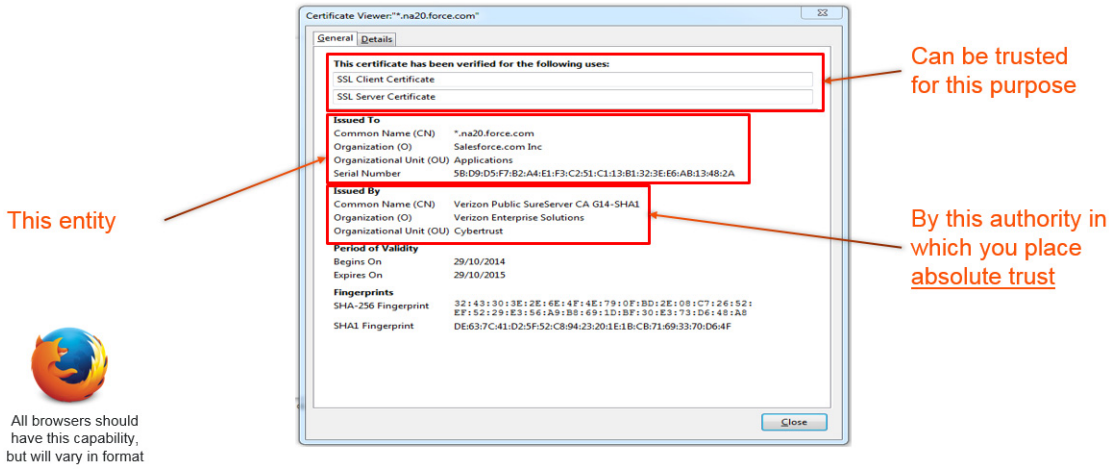

Figure 26-8: Sample Certificate

Key Store

The key store contains keys and certificate-key pairs. The key store can contain a maximum of 1000 key pairs.

A particular key in the key store can be selected only for decryption (inbound deployment) or for re-signing and decryption (outbound deployment).

Starting in software version 5.2, encrypted or password protected PEM is supported for fetching or downloading a private key.

Starting in software version 5.6.00, ECDSA keys are supported for both inbound and outbound deployments.

Generate and Add a Certificate to Key Store

To generate and add an inline-SSL signing certificate to key store for outbound deployment, perform the following steps:

1. Create an internal root CA (for example, using Microsoft AD, or OpenSSL, or any other CA implementation).
2. Push this CA root certificate to all devices.
3. Issue a sub-CA certificate for Gigamon.
4. Upload the sub-CA certificate to Gigamon with the private key of only the sub-CA.

Trust Store

The trust store contains a trusted certificate authority (CA) for server validation. A default trust store from Mozilla is included with this GigaVUE-OS software version. The trust store is updated periodically. You can also fetch a trust store file containing certificates in PEM format if you want to replace the default trust store.

Starting in software version 5.2, CA certificates can be appended to the trust store or specific certificates can be deleted from the trust store. Also, the trust store can be queried for a specific certificate by its fingerprint. The query includes the hex representation of the first four octets of the certificate's SHA1 fingerprint.

Certificate Validation

Gigamon needs to validate the server certificate so that an incoming untrusted certificate is not made legal by Gigamon re-signing the certificate.

The certificate validation process includes several steps as follows:

- Certificate expiration date and validity period: The GigaVUE node compares the current date to the validity period listed in the certificate. If the expiration date has not passed and the current date is within the period, the certificate is good.
- 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 GigaVUE node validates that the server certificate is valid for the hostname mentioned in the SNI. This validation is not performed if the client does not send SNI.

- Certificate revocation check: The GigaVUE node validates the server certificate status using OCSP and CRL lists downloaded from the concerned CAs. Internet connectivity is required for this functionality. The certificate revocation check determines the revocation status of the server certificate.

Certificates that pass the validation are accepted. The primary MitM CA signs the forged certificate. Certificates that fail validation can be accepted if security exceptions are configured and the secondary MitM CA signs the corresponding forged certificate. For self-signed certificates, the forged certificate will also be self-signed.

Client applications will typically add the primary MitM CA to their trust store and not to the secondary. This will act as a mechanism to bubble up certificate validation errors on the client applications and provide the end users an opportunity to reject the connection.

If there are certificate validation errors, the SSL connection is dropped unless explicitly permitted by the security exceptions in the policy profile. The certificate validation errors are grouped into the following four categories:

- **Expired**: The validity period of the certificate is in the past.
- **Self-signed**: The certificate is self-signed, meaning that the subject and the issuer are the same. The validity period is current and the certificate signature is valid.
- **Unknown CA**: The CA is not valid. The issuer certificate cannot be obtained from the certificate chain or is not in the trust store.
- **Invalid**: The certificate is not valid for the given SNI (for the Client Hello containing the SNI). There might have been a failure to decrypt or decode the certificate signature or the fields, not and before/not, were not read.

For security exceptions for expired, unknown CA, and invalid certificates, the resulting certificate is signed (re-signed) by the secondary MitM CA, so the user can accept or reject the connection.

Client Authentication

A server can challenge the client by requesting for a client certificate after the server responds with its Server Hello message. The client then respond with its certificate and the server validates the client certificate.

Gigamon does not support client authentication for outbound and inbound connections. If client authentication is detected during server handshake, then the connection is bypassed.

Re-Signed Certificates

As a MitM, Gigamon re-signs certificates. The following fields are copied from the original certificate:

- subject name
- issuer name
- certificate validity

The following fields are set in the re-signed certificate:

- certificate type—v3
- serial number—randomly generated
- v3 extensions:
  - basicConstraints CA—True

    NOTE: As this is a CA certificate, basicConstraints is set to True. For leaf certificates, basicConstraints is set to False.
  - keyUsage—digitalSignature and keyEncipherment
  - extendedkeyUsage—serverAuth
  - subjectKeyIdentifier—hash
  - authorityKeyIdentifier—keyid,issuer:always

Checking Certificate Revocation Status

All server certificates for decrypted outbound connections are issued by the GigaVUE node. The issuing CA is imported into the client's browsers as a trusted CA. Thus, the clients will trust all certificates signed by the GigaVUE node. This interferes with the ability of the clients to check for the revocation status of the certificates. Thus the burden is upon the node to perform the revocation checks on the original server certificates before regenerating the certificates to the clients.

If revocation check is enabled with soft fail, decryption will continue even if the revocation status is not already known, whereas with hard fail, traffic will not be decrypted unless the revocation status is determined for certain.

If revocation check is enabled, once the GigaVUE node determines that the server certificate is revoked, further SSL connections to the server are dropped.

By default, revocation check is disabled.

There are two methods to check the revocation status of the certificates as follows:

- using a Certificate Revocation List (CRL) from the issuing Certificate Authorities (CAs). Refer to *Certificate Revocation List (CRL)* on page 664.
- using the Online Certificate Status Protocol (OCSP). Refer to *Online Certificate Status Protocol (OCSP)* on page 665.

Certificate Revocation List (CRL)

A Certificate Revocation List (CRL) is an online database of certificates that have been revoked.

Each issuing CA in the PKI infrastructure maintains a list of revoked certificates that they had issued earlier. The list contains the serial number of the revoked certificates and the reasons for the revocation. Any revoked certificate should not be trusted even if the signatures are valid. CAs publish the revocation list periodically.

Each server certificate will contain the CRL location in the “CRL distribution points” X.509 extension.

Online Certificate Status Protocol (OCSP)

The Online Certificate Status Protocol (OCSP) is an Internet protocol used for obtaining the revocation status of an X.509 certificate.

Certificate status can be verified near real-time by querying the OCSP link in the certificate. The link will be present in the “Authority Information Access” X.509 extension.

CRL and OCSP

If both CRL and OCSP are enabled, OCSP is performed first, followed by CRL.

Both CRL and OCSP require Internet connectivity. Refer to *Configure Stack Port Interface* on page 680 for more information.

Supported Ciphers

Inline SSL decryption supports modern cryptographic algorithms. It supports the commonly-supported ciphers.

Combining the following ciphers, MACs, and key exchange algorithms results in many ciphersuites:

- Ciphers: RC4_128, DES_CBC, 3DES_EDE_CBC, AES_128_CBC, AES_128_GCM, AES_256_GCM, AES_256_CBC, Camellia, Chacha20
- MAC: MD5, SHA, SHA256, SHA384, Poly1305
- Key exchange algorithms: RSA, DHE_RSA, ECDHE_RSA, ECDHE_ECDSA

Table 26-5: Supported Ciphers for Inline SSL Decryption

Cipher Name	Kx	Au	Enc	Mac
TLS_RSA_EXPORT_WITH_RC4_40_MD5	RSA_EXPORT	RSA_EXPORT	RC4_40	MD5
TLS_RSA_WITH_RC4_128_MD5	RSA	RSA	RC4_128	MD5
TLS_RSA_WITH_RC4_128_SHA	RSA	RSA	RC4_128	SHA
TLS_RSA_EXPORT_WITH_DES40_CBC_SHA	RSA_EXPORT	RSA_EXPORT	DES40_CBC	SHA
TLS_RSA_WITH_DES_CBC_SHA	RSA	RSA	DES_CBC_	SHA
TLS_RSA_WITH_3DES_EDE_CBC_SHA	RSA	RSA	3DES_EDE_CBC	SHA
TLS_DHE_RSA_EXPORT_WITH_DES40_CBC_SHA	RSA_EXPORT	RSA_EXPORT	DES40_CBC	SHA
TLS_DHE_RSA_WITH_DES_CBC_SHA	RSA	RSA	DES_CBC	SHA
TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	RSA	RSA	3DES_EDE_CBC	SHA
TLS_RSA_WITH_AES_128_CBC_SHA	RSA	RSA	AES_128_CBC	SHA
TLS_DHE_RSA_WITH_AES_128_CBC_SHA	RSA	RSA	AES_128_CBC	SHA

Cipher Name	Kx	Au	Enc	Mac
TLS_RSA_WITH_AES_256_CBC_SHA	RSA	RSA	AES_256_CBC	SHA
TLS_DHE_RSA_WITH_AES_256_CBC_SHA	RSA	RSA	AES_256_CBC	SHA
TLS_RSA_WITH_AES_128_CBC_SHA256	RSA	RSA	AES_128_CBC	SHA256
TLS_RSA_WITH_AES_256_CBC_SHA256	RSA	RSA	AES_256_CBC	SHA256
TLS_RSA_WITH_CAMELLIA_128_CBC_SHA	RSA	RSA	CAMELLIA_128_CBC	SHA
TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA	RSA	RSA	CAMELLIA_128_CBC	SHA
TLS_RSA_EXPORT1024_WITH_RC4_MD5	RSA_EXPORT	RSA_EXPORT	RC4	MD5
TLS_RSA_EXPORT1024_WITH_DES_CBC_SHA	RSA_EXPORT	RSA_EXPORT	DES_CBC	SHA
TLS_RSA_EXPORT1024_WITH_RC4_SHA	RSA_EXPORT	RSA_EXPORT	RC4_	SHA
TLS_DHE_RSA_WITH_AES_128_CBC_SHA256	RSA	RSA	AES_128_CBC	SHA256
TLS_DHE_RSA_WITH_AES_256_CBC_SHA256	RSA	RSA	AES_256_CBC	SHA256
TLS_RSA_WITH_CAMELLIA_256_CBC_SHA	RSA	RSA	CAMELLIA_256_CBC	SHA
TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA	RSA	RSA	CAMELLIA_256_CBC	SHA
TLS_RSA_WITH_SEED_CBC_SHA	RSA	RSA	SEED_CBC	SHA
TLS_DHE_RSA_WITH_SEED_CBC_SHA	RSA	RSA	SEED_CBC	SHA
TLS_RSA_WITH_AES_128_GCM_SHA256	RSA	RSA	AES_128_GCM	SHA256
TLS_RSA_WITH_AES_256_GCM_SHA384	RSA	RSA	AES_256_GCM	SHA384
TLS_DHE_RSA_WITH_AES_128_GCM_SHA256	RSA	RSA	AES_128_GCM	SHA256
TLS_DHE_RSA_WITH_AES_256_GCM_SHA384	RSA	RSA	AES_256_GCM	SHA384
TLS_ECDHE_ECDSA_WITH_RC4_SHA	ECDSA	ECDSA	RC4	SHA
TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA	ECDSA	ECDSA	3DES_EDE_CBC	SHA
TLS_ECDHE_ECDSA_WITH_AES128_CBC_SHA	ECDSA	ECDSA	AES128_CBC	SHA
TLS_ECDHE_ECDSA_WITH_AES256_CBC_SHA	ECDSA	ECDSA	AES256_CBC	SHA
TLS_ECDHE_RSA_WITH_RC4_SHA	RSA	RSA	RC4	SHA
TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA	RSA	RSA	3DES_EDE_CBC	SHA
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	RSA	RSA	AES_128_CBC	SHA
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	RSA	RSA	AES_256_CBC	SHA
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	ECDSA	ECDSA	AES_128_CBC	SHA256
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384	ECDSA	ECDSA	AES_256_CBC	SHA384
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	RSA	RSA	AES_128_CBC	SHA256
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	RSA	RSA	AES_256_CBC	SHA384
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	ECDSA	ECDSA	AES_128_GCM	SHA256
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	ECDSA	ECDSA	AES_256_GCM	SHA384
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	RSA	RSA	AES_128_GCM	SHA256

Cipher Name	Kx	Au	Enc	Mac
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	RSA	RSA	AES_256_GCM	SHA384
TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305	RSA	RSA	CHACHA20	POLY1305
TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305	ECDSA	ECDSA	CHACHA20	POLY1305
TLS_DHE_RSA_WITH_CHACHA20_POLY1305	RSA	RSA	CHACHA20	POLY1305

Diffie Hellman Ephemeral (DHE) is a key exchange protocol.

Inline SSL decryption supports key cipher suites and exchanges without downgrading cryptography levels of the organization.

Ciphersuites are a standard combination of the following:

- bulk encryption algorithm—Specifies how to encrypt communications, including the algorithm, key size, and the cryptographic mode used. For example, AES_128_CBC is AES with 128-bit keys in Cipher Block Chaining mode.
- key exchange algorithm—Specifies how both sides authenticate each other during the 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.

SSL transactions with unsupported ciphers will be bypassed/TCP proxied.

The following key sizes are supported:

- RSA—512, 1024, 2048, 3072, 4096, 8192
- DH—1024, 2048, 4096
- ECC—prime256v1, ecsecp256r1, ecsecp384r1, ecsecp521r1

The following TLS extension is supported:

- RFC7301—Application-Layer Protocol Negotiation (ALPN)

Policy Profile

The policy profile consists of multiple rules, with each rule having a decrypt or no-decrypt action for the match condition. For example, there might be a policy to decrypt all but financial-related traffic.

In addition to the rules, the profile also consists of various configuration options that affect the decryption decision as follows:

- The default action to take if none of the rules match.
- The URL cache miss action to take if the URL category-based rules are configured, but GigaSMART does not have the category information.

- For decrypted traffic, options to override expired, invalid, self-signed, and unknown CA certificates and to enable or disable the certificate revocation check.
- Whether or not to send decrypted/non-decrypted traffic through the tools.

Each policy rule consists of a match condition and the decrypt or no-decrypt action for the match. The following rule types are available:

- URL category
- hostname/domain name
- server certificate issuer
- source and destination IP address
- source and destination port numbers
- VLAN identifier

NOTE: You can configure up to 1024 policy rules under a policy profile.

Policy Evaluation

Policies are evaluated by GigaSMART at various phases as shown in Figure 26-9.

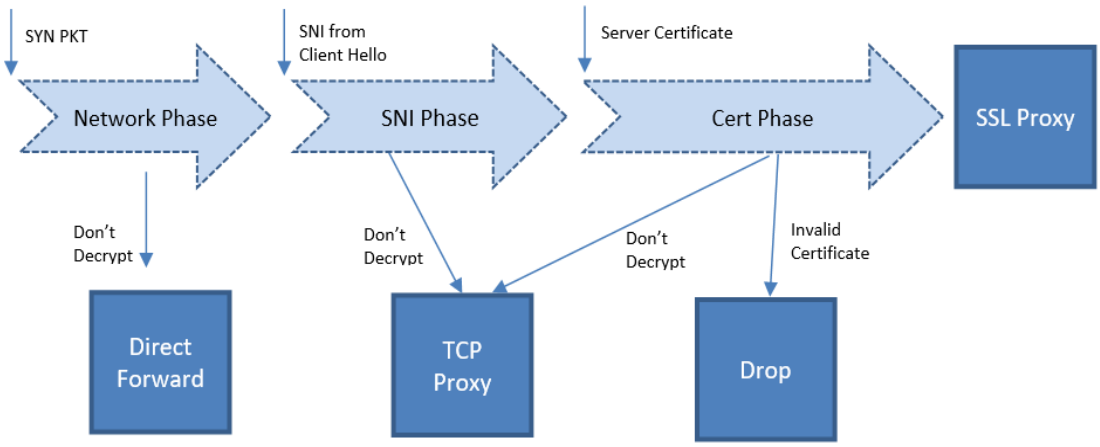

Figure 26-9: Policy Validation Flow

Network Phase

The Network Phase of the policy is done on the SYN packet. In this phase:

- The policy engine inputs are source and destination IP addresses, source and destination ports, and VLAN identifiers.
- The no-decrypt rules are evaluated before the decrypt rules, which are ordered by the following match conditions: source IP, destination IP, source port, destination port and VLAN.
- If traffic is not to be decrypted, packets are processed as non-proxy traffic. The bypass VLAN is used. Refer to Direct Forward in Figure 26-9.

  NOTE: The no-decrypt verdict is always final for the lifetime of the given TCP connection.

- Traffic that is to be decrypted continues to the next phase. The decrypt verdict from the Network Phase can be overridden by more specific rules in the next phase.

Following conditions are considered for traffic that is to be decrypted:

- If the Client Hello packet contains SNI, policy is evaluated in the SNI phase and then in the Certificate phase.
- If the Client Hello packet does not contain SNI, policy is evaluated only in the Certificate phase.
- For HTTPS proxy connections, policy rules are evaluated based on the hostname from the HTTPS proxy CONNECT request in addition to the SNI.

SNI Phase

The SNI Phase of the policy is done on the Server Network Indication (SNI) from the Client Hello. In this phase:

- The policy engine input is the hostname from SNI.
- Rules are evaluated in the following order:
  - whitelist/no-decrypt domain
  - blacklist/decrypt domain
  - no-decrypt category
  - decrypt category

  NOTE: Domain/whitelist rules apply for all subdomains as well. For example the following are recognized as being part of gigamon.com: ftp.gigamon.com, www.gigamon.com, www.gigamon.com/default.htm.
- If traffic is not to be decrypted, packets are processed as TCP proxy. Refer to TCP Proxy in Figure 26-9.
- For traffic that is to be decrypted, the plaintext version is sent through the tools based on the decrypt tool-bypass configuration.

If URL category-based rules are configured and a URL cache miss occurs, the policy verdict is based on the settings of the URL cache miss action. A cache miss action of decrypt or no-decrypt will cause the traffic to be decrypted or not decrypted immediately. The defer action will cause delays.

For compliance reasons, the cache miss action of no-decrypt is recommended.

Certificate Validation

In the case of a decrypt decision from the SNI phase or if the Client Hello does not contain SNI, GigaSMART will verify the server certificate using the configured trust store. Additional checks will be performed for certificate expiry, hostname mismatch, and self-signed certificate. If configured, revocation check will also be done on otherwise valid server certificates.

For valid server certificates, GigaSMART will issue the corresponding server certificate using the primary MitM CA. If the validation fails, the connection will be dropped unless

a security exception is configured. The secondary MitM CA, if configured, will be used to issue server certificates in that case.

Starting in software version 5.2, the primary MitM CA is not mandatory for an inbound deployment.

Cert Phase

The certificate phase of policy evaluation is done for all connections after the certificate validation is completed. In this phase:

- The policy engine inputs are the certificate issuer and the certificate subject name.
- The Common Name (CN) attribute is extracted from the server certificate subject name. Policy evaluation in this phase is performed on the value of the CN attribute. This is similar to the SNI phase. If no matching rules are found, the certificate issuer-based rules are evaluated. For issuer-based rules, the CN attribute and the Domain Name (DN) attribute of the issuer are considered.
- If traffic is not to be decrypted, packets are processed as TCP proxy. Refer to TCP Proxy in Figure 26-9. The server SSL session is reset and the Client Hello is resent to the server.
- GigaSMART supports certificate based policy evaluation for HTTPS proxied connections. It applies policy rules based on the hostname from HTTPS proxy CONNECT request for HTTPS proxy connections. Also, if required it applies no-decrypt in the certificate phase for HTTPS proxy connections.
- Certificate phase policies are also evaluated for HTTPS proxy connections.

Policy Profile Options

This section describes a few of the options for the policy profile. Refer to the following:

- *Inline SSL Decryption Port Map* on page 670
- *Enable or Disable Tool Bypass* on page 671
- *High Availability Active Standby* on page 671
- *Inline Network Group Multiple Entry* on page 671

Inline SSL Decryption Port Map

The TCP destination port for decrypted traffic sent to inline tools can be configured as part of the profile.

Following are the two priorities that GigaSMART uses to decide on the TCP port number used for decrypted traffic:

- Priority 1—This is a port map, which is user configurable. You can specify both the In Port and the Out Port. The In Port is the TCP destination port from a client. The Out Port is the TCP port used to send traffic to inline tools.
- Priority 2—This is a default Out Port. This TCP port will be used if the incoming port does not match those specified in Priority 1.

Enable or Disable Tool Bypass

Tool bypass can be enabled or disabled for the following types of traffic:

- SSL decrypted traffic
- non-decrypted SSL traffic (non-SSL TCP)
- non-SSL traffic (non-TCP)

By default, tool bypass is disabled on these traffic types, meaning that all decrypted SSL, non-decrypted SSL, and non-SSL traffic is sent to the tools. When tool bypass is enabled on a specified traffic type, that traffic is not sent to the tools.

High Availability Active Standby

Starting in software version 5.2, inline network high availability active standby is supported. When enabled, link switchover by an upstream device in active/standby scenario is detected.

For example, when there is an inline SSL network group topology with two network port pairs (Na1, Nb1 and Na2, Nb2), the incoming traffic from one network (for example, Na1) may change to another network (for example, Na2) due to upstream devices, such as firewalls performing high availability active standby failover. If an upstream device fails over, GigaSMART will forward traffic to the correct inline network.

The default is disabled.

NOTE: Do not enable this option if the inline SSL network group links are in an active/active scenario.

Inline Network Group Multiple Entry

An inline network group topology can have multiple network port pairs (for example, Na1, Nb1 and Na2, Nb2). With multiple network port pairs, traffic from a network interface might traverse GigaSMART multiple times. Intercepted traffic from GigaSMART might reenter GigaSMART through a different network interface within the same network group as shown in Figure 26-10.

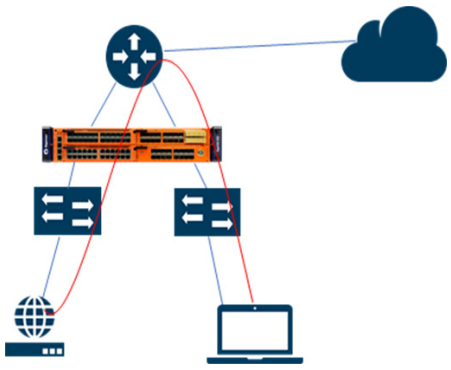

Figure 26-10: Inline SSL Inline Network Group Configuration

When the inline SSL GigaSMART sits between internal devices and the upstream router, traffic from the devices to the Internet will be intercepted by GigaSMART. When

internal devices belonging to different network port pairs within the same inline network group communicate with each other, traffic initiated from a device will be intercepted by GigaSMART and sent to the upstream router. This traffic will be routed back to GigaSMART from a different network port pair to reach the destination device.

Starting in software version 5.3, the same traffic sent from GigaSMART can reenter GigaSMART.

GigaSMART remembers the inline incoming inline network interface (for example, Na1) for each connection. When traffic from the same connection reaches GigaSMART with a different inline network interface within the same network group (for example, Na2), GigaSMART will forward the traffic to the corresponding opposite network interface (for example, Nb2), without further processing. This allows traffic from the same connection to reenter GigaSMART. GigaSMART will detect it and start forwarding traffic to the new network port pair.

However, the same traffic sent by GigaSMART reentering through the same network port pair (for example, Nb2, Na2) is not supported.

Other than the use case described above, any connection with traffic passing through GigaSMART involving more than the original network pair is not supported. If the first packet of a connection comes in through Na1, all traffic has to enter GigaSMART through the network port pair, Na1, Nb1.

You can enable or disable the inline network group multiple entry for the profile. The default is disabled.

Caches

There are four in-memory caches as follows. They are not configurable.

- re-signed certificate cache
- URL category cache
- revocation certificate cache
- session resumption cache

Cache Persistence

Caches are maintained for Internet lookups such as URL categorization and certificate revocation checks using OCSP or CRL for faster subsequent lookups. The cache persistence feature allows the information to be saved on the GigaVUE node in the control card's persistent storage so that it can be retrieved in case of reboots. This allows the GigaSMART card to start with the information learned earlier. This feature is enabled by default and can be disabled if needed.

You can search for specific entries in the caches, clear them, and display a summary of the records.

GigaSMART Overload Bypass

Packet buffers, CPU, and concurrent connections are monitored for overloaded conditions. GigaSMART goes to bypass when resource usage exceeds thresholds. Existing connections will continue to be processed by GigaSMART, but any new connections will be bypassed. Refer to Table 26-6 on page 673 for information on connections and thresholds.

Table 26-6: Overload Bypass Connections and Thresholds

Criteria	GigaVUE-HC2 (per module)	GigaVUE-HC3 (per module)	GigaVUE-HC1
Maximum connections per second	2000	5000	1500
Maximum connections	100000	200000	100000
FPA overload	Default 80% (configurable)		
Heap exhaust	Default 80% (configurable)		
CPU	Default 90% (configurable)		

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 SSL Monitor Mode

Use the inline 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 SSL decryption is configured and monitor mode is enabled, the inline 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-SSL TCP bypass action.

Monitor mode is disabled by default. To enable the monitor mode, refer to *Enable the Inline SSL Monitor Mode* on page 674.

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.

Monitor mode only captures TCP information, not SSL information. Also, it is supported for standalone nodes, but not for nodes in a cluster.

Enable the Inline SSL Monitor Mode

You can enable the monitor mode using either CLI command or GigaVUE-FM.

To enable the monitor mode using CLI, run the following CLI command:

(config) # apps inline-ssl monitor enable

To enable the monitor mode using GigaVUE-FM:

1. From the device view, go to **GigaSMART > Inline SSL > Global Defaults**.
2. Click **Edit**.
3. Select the **SSL Session Monitor** check box.
4. Click **OK**.

Inline Tool Configurations

Inline tools connect to the GigaVUE node through inline bypass (BPS) modules, available on the GigaVUE-HC1, GigaVUE-HC2, and GigaVUE-HC3.

Inline bypass is described in detail in the *GigaVUE-OS CLI User's Guide*.

The inline bypass arrangements supported by inline 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
- inline flow mapping in which traffic is classified into types to send to specific tools. Types of traffic might be: encrypted Web, email, or voice. Unencrypted traffic can be bypassed.

In summary, inline 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.

Inline tool ports can be configured in shared mode. When an inline tool is shared (true), the decrypted traffic will be VLAN tagged. The connected inline device is expected to receive VLAN tagged packets instead of untagged packets. There is an extra outer VLAN tag added to the packet, which the connected inline device needs to see. When an inline tool is not shared (false), the extra VLAN tag is not added. This allows untagged traffic to be sent to the tool ports. Use false for inline tools that are not able to handle more than one VLAN tag, such as Q-in-Q tagged packets. For tagless mode, if an inline tool is involved in an inline SSL map, the inline tool cannot be used in any other classic inline map.

Refer to Figure 26-11 to Figure 26-13 for inline tool arrangements. Encrypted traffic is shown in solid lines, decrypted traffic is shown in dotted lines.

Figure 26-11 shows a simple inline tool arrangement with one inline tool connected to the GigaVUE node. Traffic is decrypted and sent to the inline tool.

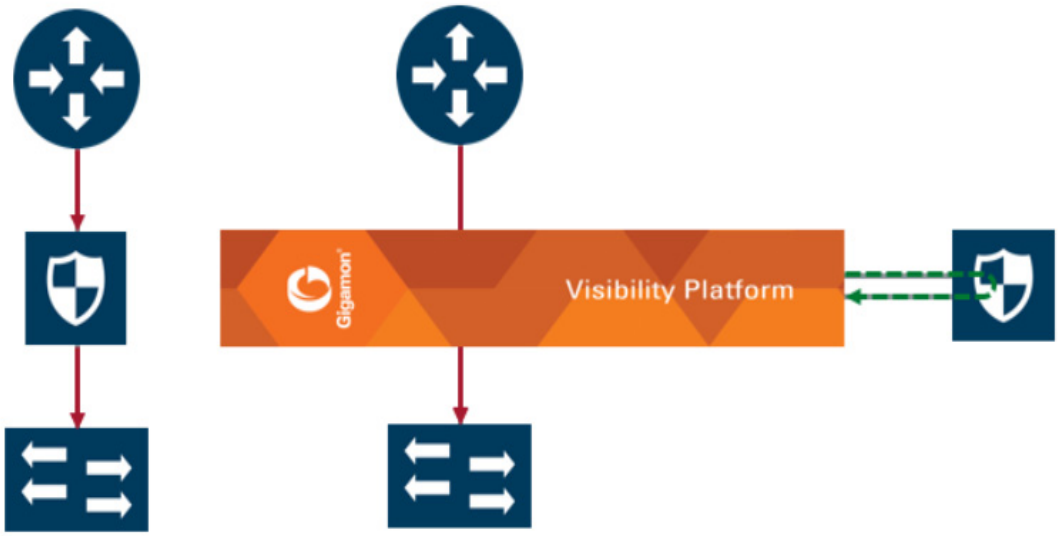

Figure 26-11: Simple Inline Tool Arrangement

Figure 26-12 shows a multiple inline tool arrangement with three inline tools connected to the GigaVUE node. Decrypted traffic is distributed across the tools.

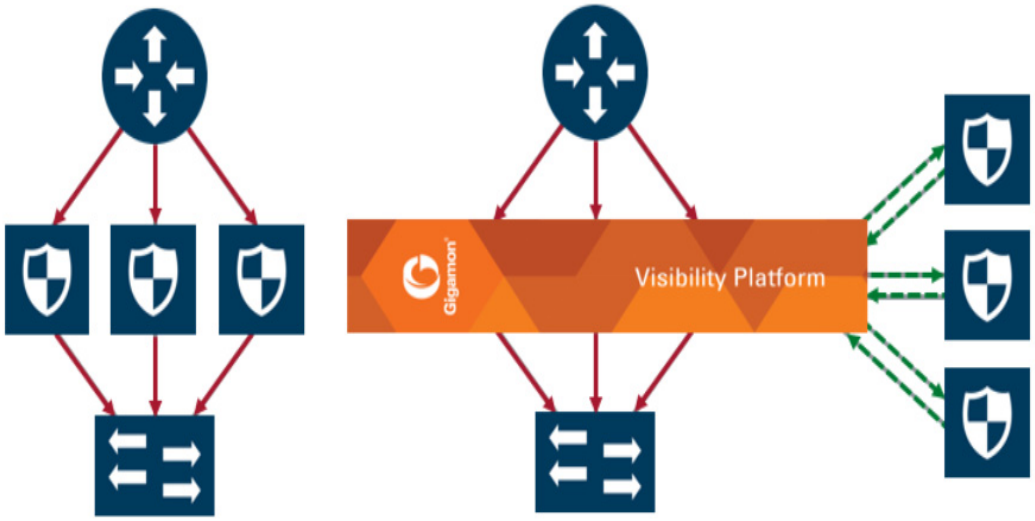

Figure 26-12: Multiple Inline Tool Arrangement

Figure 26-13 shows a serial inline tool arrangement in which traffic is decrypted on the GigaVUE node, sent serially through the inline tool, and then re-encrypted on the GigaVUE node.

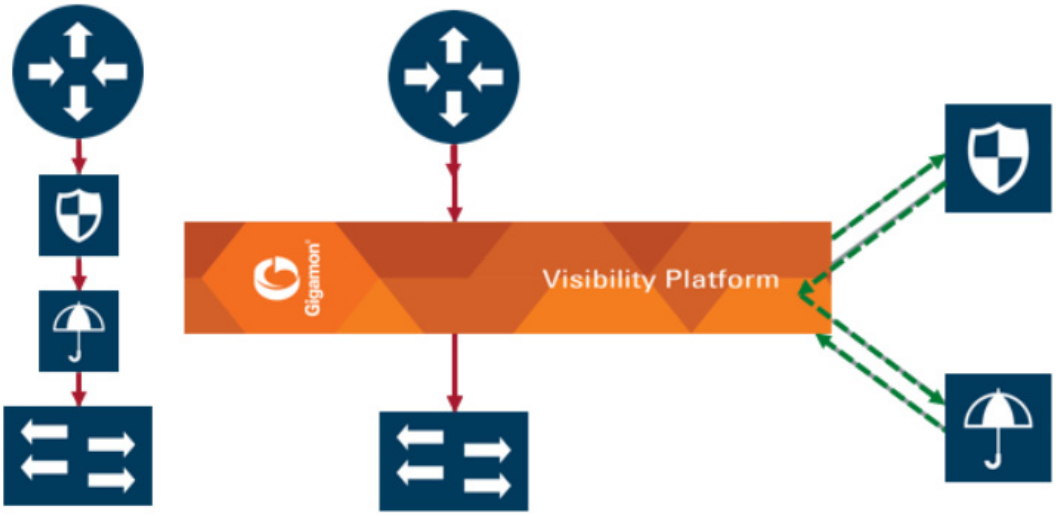

Figure 26-13: Serial Inline Tool Arrangement

Inline Bypass Restriction

The inline bypass arrangements supported by inline SSL decryption have the following restriction:

- inline network group does not support ingress VLAN tagging on the member links

Forwarding

Inline 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.
- 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.
- 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 User's 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 User's Guide* for the actions and default value of the vport command.

The GigaSMART module might also have events such as port down or card down. These failover actions are not configurable but are triggered by events. These events should not impact traffic.

Out-of-Band and Inline Tools

After decryption, traffic can be sent to multiple tools. The tools can be either inline or out-of-band.

Out-of-band tools process the decrypted packets offline. The tools are connected to the GigaVUE node through tool or hybrid ports, GigaStream, or port groups with tool or hybrid ports. The out-of-band tools receive a copy of the decrypted packets from the GigaSMART module. This is referred to as GigaSMART out-of-band forwarding.

Inline tools process the decrypted packets inline. Inline tools are connected to the GigaVUE node through inline bypass (BPS) modules.

An out-of-band tool might be an Intrusion Detection System (IDS) examining decrypted packets:

- If it detects a threat, the IDS will send a notification back, but does not have the ability to act.

An inline tool might be an Intrusion Prevention System (IPS) examining decrypted packets:

- If it does not detect a threat in the decrypted packets, the traffic comes out of the inline tool and goes back to the GigaSMART module to be re-encrypted and sent to the server.
- If it detects a threat, the IPS can act. The action depends on the tools' behavior. It can either terminate the connection or modify packets
- If the IPS terminates the connection, then GigaVUE node will terminate the connection between the client and the server.
- If the IPS modifies packets, then the modified packets will come out of the inline tool, go to the GigaSMART module to be re-encrypted, and sent to the server.

When an inline tool is shared, you must:

- configure the inline second level out-of-band map to forward proxy traffic from GigaSMART to the out-of-band tool port.

When an inline tool is not shared, you must configure only the inline second level out-of-band map to forward proxy and non-proxy traffic from GigaSMART to the out-of-band tool port.

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 SSL decryption.

Get Started with Inline SSL Decryption

This section describes the pre-requisites needed before you begin configuring inline SSL decryption. Refer to the following sections for details:

Supported Platforms

Inline SSL decryption is supported on GigaVUE-HC1, GigaVUE-HC2, and GigaVUE-HC3 nodes.

GigaSMART Licensing

A GigaSMART license is required. The required GigaSMART license is: SSL Decryption for Inline and Out-of-Band Tools.

GigaSMART Compatibility

Inline SSL decryption is not compatible with any other GigaSMART operations, including out-of-band SSL decryption. Configure inline SSL decryption on a GigaSMART engine that is not shared with any other GigaSMART operation. Moreover, for GigaVUE-HC2 nodes, it is recommended that you create separate GigaSMART operations for front and rear GigaSMART engines.

Install GigaVUE Modules

Install the GigaSMART and inline bypass module or copper TAP on the same GigaVUE-HC1 or GigaVUE-HC2 node or install the GigaSMART and inline bypass module on the same GigaVUE-HC3 node.

Install Software Version

Install software version 5.2.xx or higher for the GigaVUE-OS CLI, GigaVUE-OS H-VUE, and GigaVUE-FM.

Install U-Boot Version on GigaVUE-HC2

The U-Boot version on GigaVUE-HC2 nodes must be upgraded to version 2011.06.9 or higher. The upgrade can only be done from the CLI.

To check the U-Boot version, use the following command:

(config) # show version

For example on a GigaVUE-HC2 node, the following output is displayed:

```
U-Boot version:    2011.06.8
```

If you do not have version 2011.06.9 or higher, you will have to do a U-Boot upgrade, after the image installation. Refer to the *GigaVUE H Series Upgrade Guide* for details on installing an image.

After the image installation of the software, use the following command to upgrade the U-Boot version:

(config) # uboot install

The binary bootloader code included with the installed image is installed.

NOTE: The newer U-Boot version only goes into effect after a reload.

Install MitM Certificates in Client Trust Store

For an outbound deployment, the Man-in-the-Middle (MitM) certificates must be installed in the client trust store. Install the certificates as a Trusted Root Authority in web browsers on the client PC.

Refer to your browser's documentation for installing CA certificates to the trust store.

Configure Stack Port Interface

Internet connectivity is required for CRL, OCSP, and URL categorization. The stack port interface must be configured on the GigaSMART engine. You can configure one stack port for all the GigaSMART engines in a GigaVUE node. However, you must configure all the engines with unique DHCP or static IP address.

Refer to Figure 26-14 for the location of the stack port on the front of the GigaVUE-HC2. It is the second port from the top on the left side of the chassis.

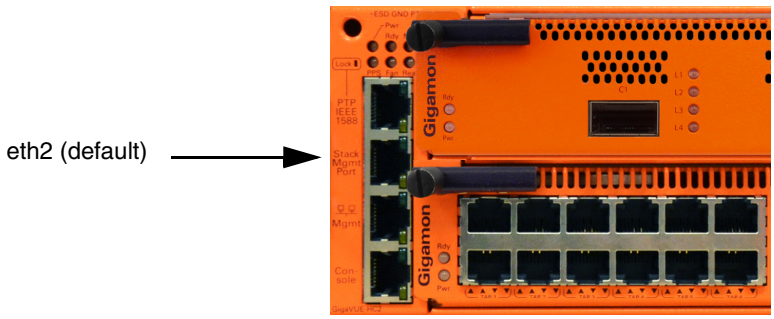

Figure 26-14: Stack Port Location on GigaVUE-HC2 Front

Refer to Figure 26-15 for the location of the two stack ports, eth2 (default) and eth3 on the control card in the front of GigaVUE-HC3. You can use either one or both the stack ports for GigaSMART connectivity.

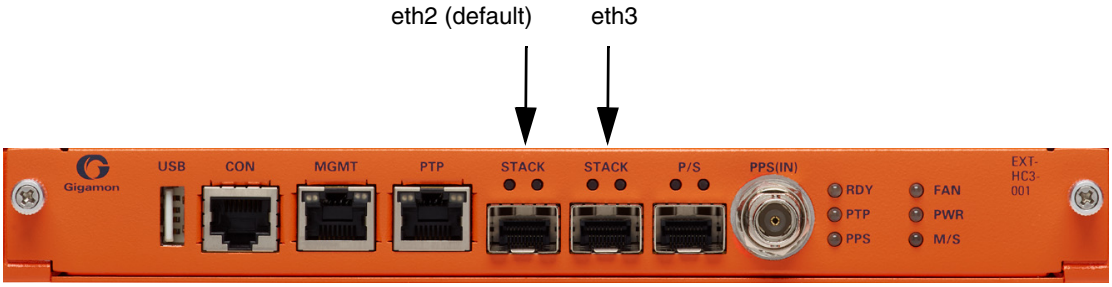

Figure 26-15: Stack Port Location on GigaVUE-HC3 Front

Refer to Figure 26-16 for the location of the stack port on the front of the GigaVUE-HC1. It is the top port on the right.

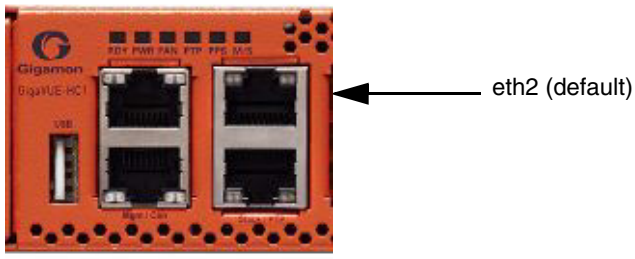

Figure 26-16: Stack Port Location on GigaVUE-HC1 Front

Configure Stack Port Interface

To configure the stack port interface:

1. Go to **Ports > All Ports**. Select a GigaSMART engine port and click **Edit**.
2. Specify either an IP Address, Netmask, Gateway, DNS IP, and optional MTU or select DHCP. Specify a VLAN ID in the range from 20 to 4094. Select the stack port interface, Eth2 or Eth3. The default is Eth2.
3. Click **OK**. The stack port interface is added.

Configure Keychain Password

For Inbound and Outbound Inline-SSL deployments, the keychain password must be configured before installing the certificates and private keys into the keystore.

Refer to *Configure Inline SSL Decryption Using GigaVUE-FM* on page 690 for the configuration steps.

Configure Primary Certificate and Key

For an outbound deployment, at least one of the CAs must be configured (primary or secondary). For an inbound deployment, a CA is not necessary.

The primary CA re-signs certificates for servers that present a valid certificate. The secondary CA re-signs certificates for servers that are invalid or that fail validation. If the secondary CA is not configured, the primary CA will be used for all certificates.

Refer to *Configure Inline SSL Decryption Using GigaVUE-FM* on page 690 for the configuration steps.

Configure Inline SSL Decryption

This section describes the workflows for configuration inline SSL decryption using GigaVUE-FM. It also provides the details of the workflows for inline SSL map. Refer to the following sections for details:

- *Introduction to Inline SSL Map Workflows* on page 682
- *Configure Inline SSL Decryption Using GigaVUE-FM* on page 690
- *View Statistics* on page 694
- *Configure Inline SSL Session Logging Server* on page 696

Introduction to Inline SSL Map Workflows

In GigaVUE-FM, workflows guide you through configuration steps. For the Inline SSL Map configuration, there are seven flows, Flow A to Flow G based on which you can perform different configurations.

Go to **Workflows** and select **Inline SSL Map** from the Inline GigaSMART Operations section as shown in Figure 26-17.

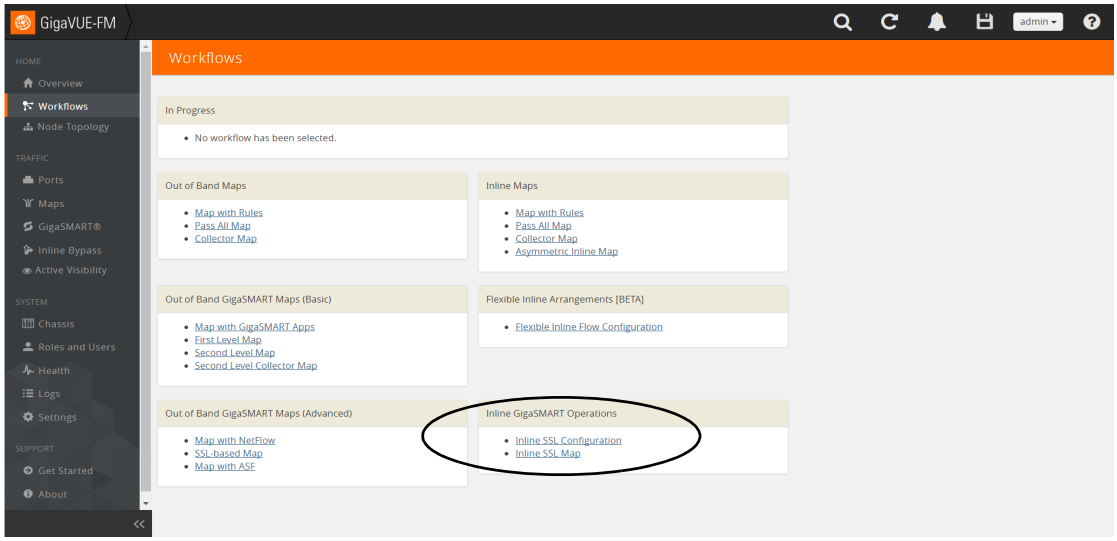

Figure 26-17: Select Inline SSL Map Configuration

Flow A to Flow G are displayed as shown in Figure 26-18.

Figure 26-18: Flow A to Flow G

The following sections describe each flow:

- *Install GigaVUE Modules* on page 679
- *Install Software Version* on page 679
- *Install U-Boot Version on GigaVUE-HC2* on page 680
- *Flow D* on page 686
- *Flow E* on page 687
- *Flow F* on page 688
- *Flow G* on page 689

Flow A

Flow A is for the following use case:

- filter HTTP traffic and direct it to the tool(s)
- filter remaining TCP traffic and direct it to the GigaSMART engine for inspection before forwarding it to the tool(s)
- direct all other traffic to the tool(s)

Refer to Figure 26-19 for a larger view of Flow A on the left and a pictorial view on the right.

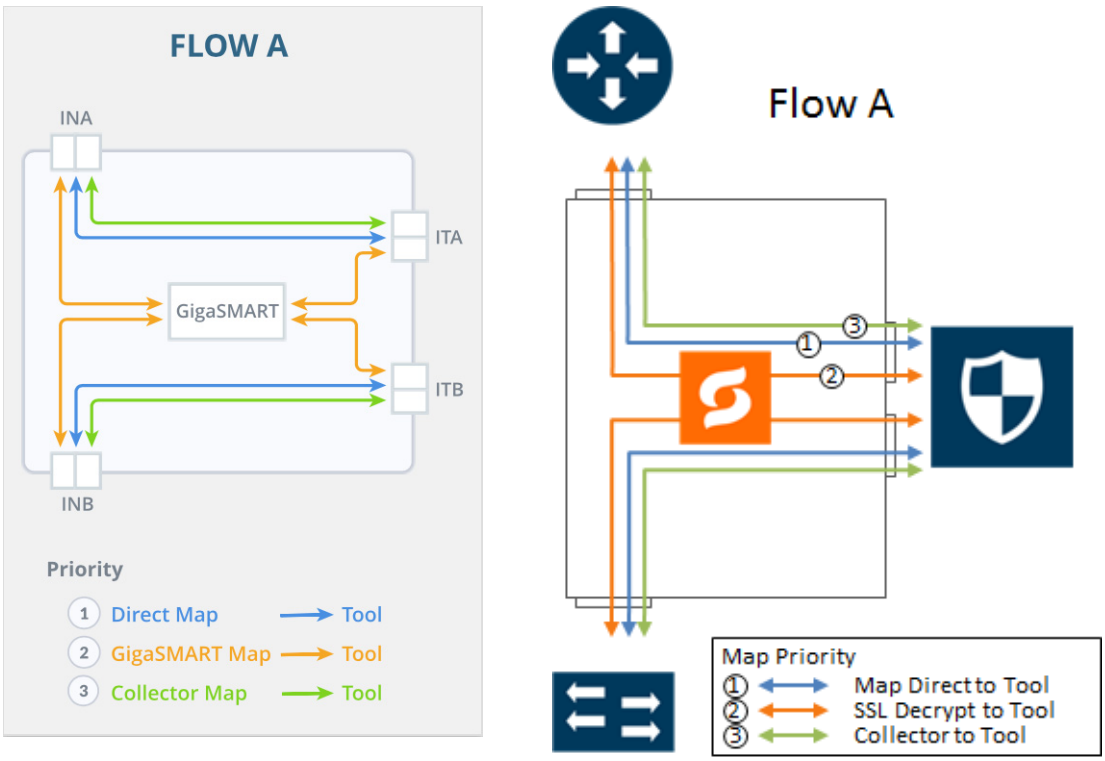

Figure 26-19: FLOW A Views

The map priorities of Flow A are as follows:

1. limit traffic going to decryption
2. selectively forward traffic for decryption
3. direct unselected traffic to a collector, which sends traffic to tool

Flow B

Flow B is for the following use case:

- filter HTTP traffic and direct it to the tool(s)
- filter remaining TCP traffic and direct it to the GigaSMART engine for inspection before forwarding it to the tool(s)
- bypass all other traffic

Refer to Figure 26-20 for a larger view of Flow B on the left and a pictorial view on the right.

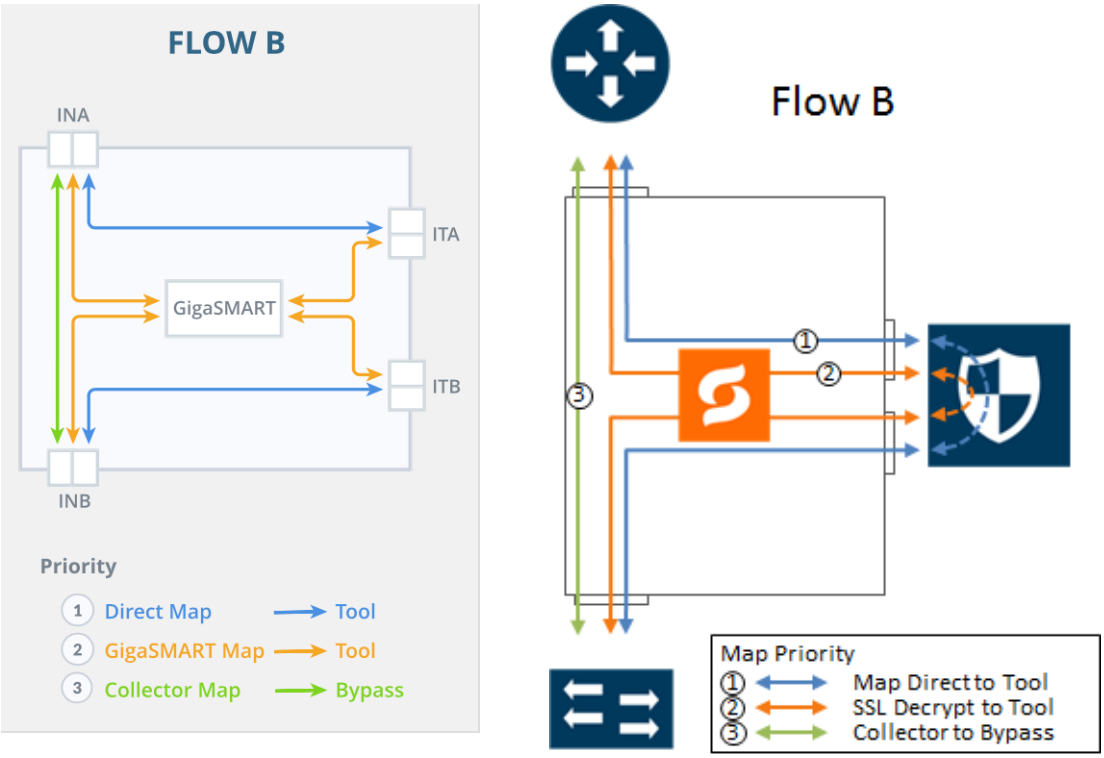

Figure 26-20: FLOW B Views

The map priorities of Flow B are as follows:

1. limit traffic going to decryption
2. selectively forward traffic for decryption
3. direct unselected traffic to a collector, which sends traffic to bypass

Flow C

Flow C is for the following use case:

- filter traffic from or to certain VLANs (for example, a guest WiFi VLAN) and direct it to bypass
- filter remaining TCP traffic and direct it to the GigaSMART engine for inspection before forwarding it to the tool(s)
- direct all other traffic to the tool(s)

Refer to Figure 26-21 for a larger view of Flow C on the left and a pictorial view on the right.

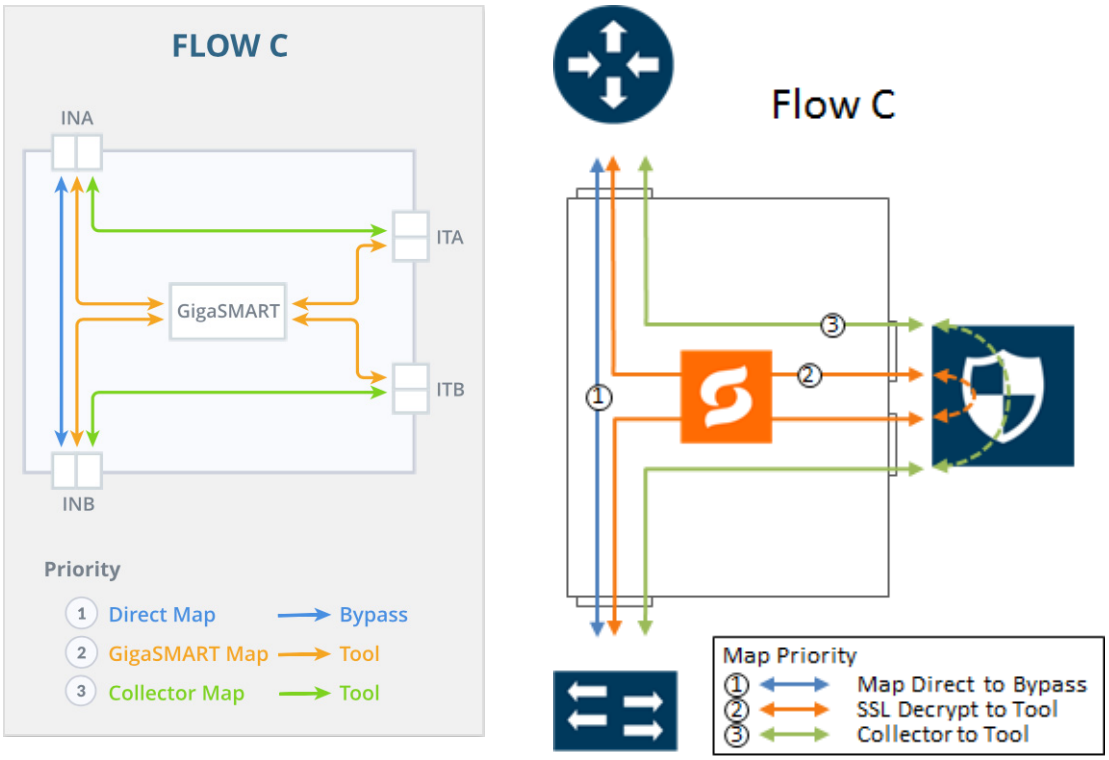

Figure 26-21: FLOW C Views

The map priorities of Flow C are as follows:

1. send trusted traffic to bypass
2. selectively forward traffic for decryption
3. direct unselected traffic to a collector, which sends traffic to tool

Flow D

Flow D is for the following use case:

- filter traffic from certain VLANs (for example, an employee WiFi VLAN) and direct it to the GigaSMART engine for inspection before forwarding it to the tool(s)
- filter traffic from or to selected internal servers and direct it to bypass
- direct all other traffic to tool(s)

Refer to Figure 26-22 for a larger view of Flow D on the left and a pictorial view on the right.

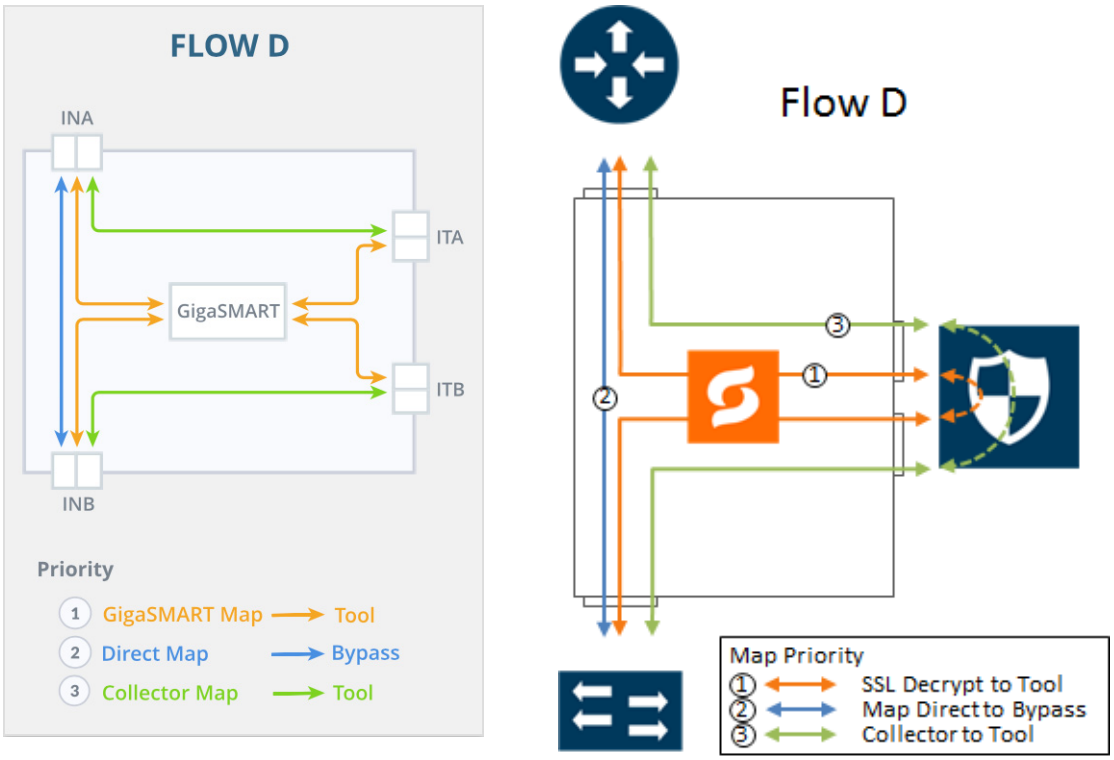

Figure 26-22: FLOW D Views

The map priorities of Flow D are as follows:

1. selectively forward traffic for decryption
2. send trusted traffic to bypass
3. direct unselected traffic to a collector, which sends traffic to tool

Flow E

Flow E is for the following use case:

- filter traffic from certain VLANs (for example, an employee WiFi VLAN) and direct it to the GigaSMART engine for inspection before forwarding it to the tool(s)
- filter traffic from or to selected internal servers and direct it to the tool(s)
- bypass all other traffic

Refer to Figure 26-23 for a larger view of Flow E on the left and a pictorial view on the right.

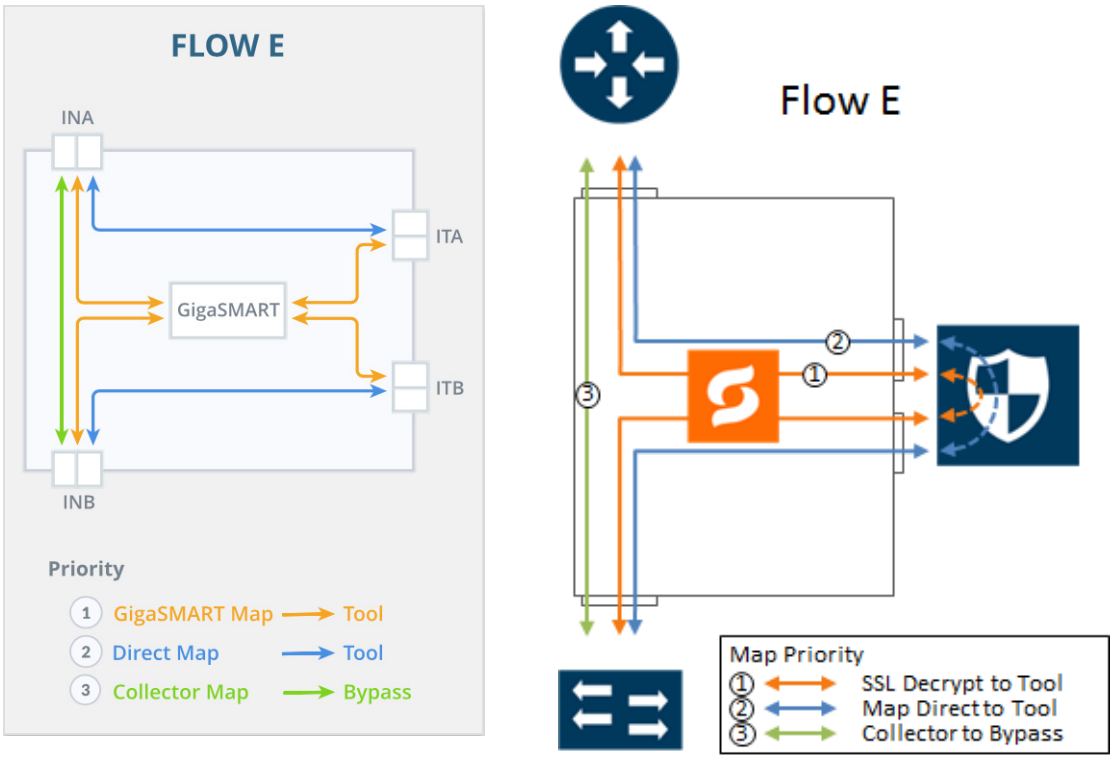

Figure 26-23: FLOW E Views

The map priorities of Flow E are as follows:

1. selectively forward traffic for decryption
2. send remaining IP traffic to tool(s)
3. direct non-IP traffic to a collector, which sends traffic to bypass

Flow F

Flow F is for the following use case:

- filter TCP traffic and direct it to the GigaSMART engine for inspection before forwarding it to the tool(s)
- direct all other traffic to the tool(s)

Refer to Figure 26-24 for a larger view of Flow F on the left and a pictorial view on the right.

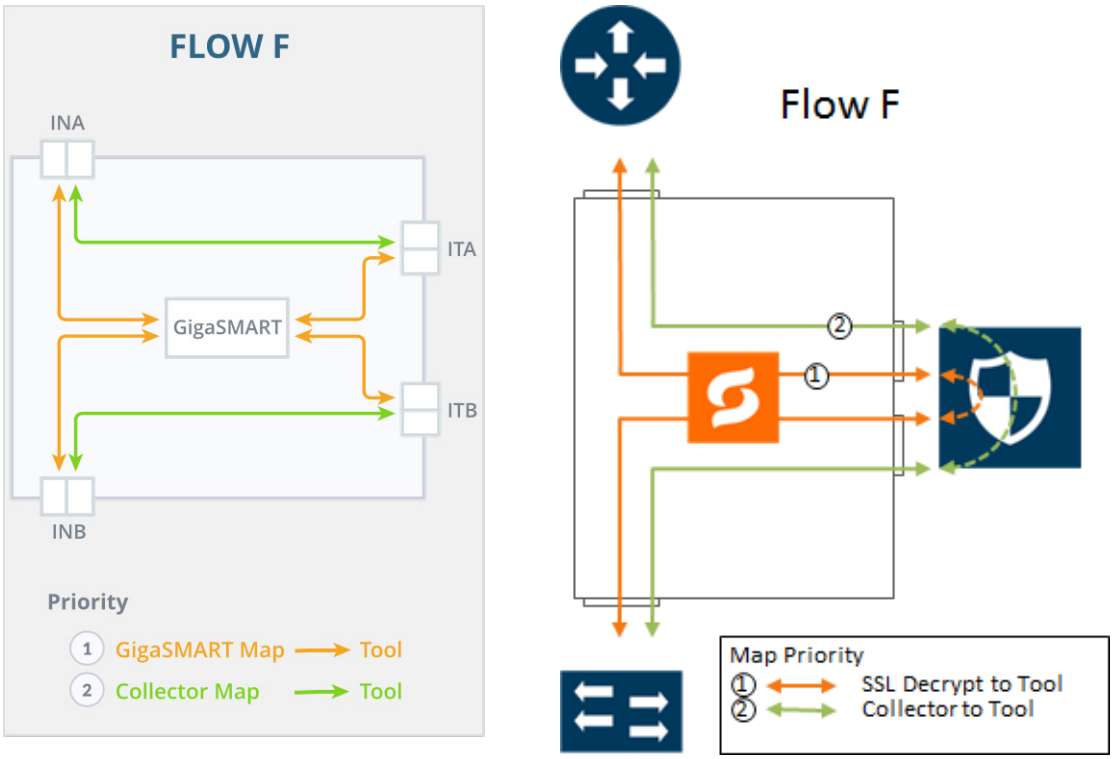

Figure 26-24: FLOW F Views

The map priorities of Flow F are as follows:

1. selectively forward traffic for decryption
2. direct unselected traffic to a collector, which sends traffic to tool

Flow G

Flow G is for the following use case:

- filter TCP traffic and direct it to the GigaSMART engine for inspection before forwarding it to the tool(s)
- bypass all other traffic

Refer to Figure 26-25 for a larger view of Flow G on the left and a pictorial view on the right.

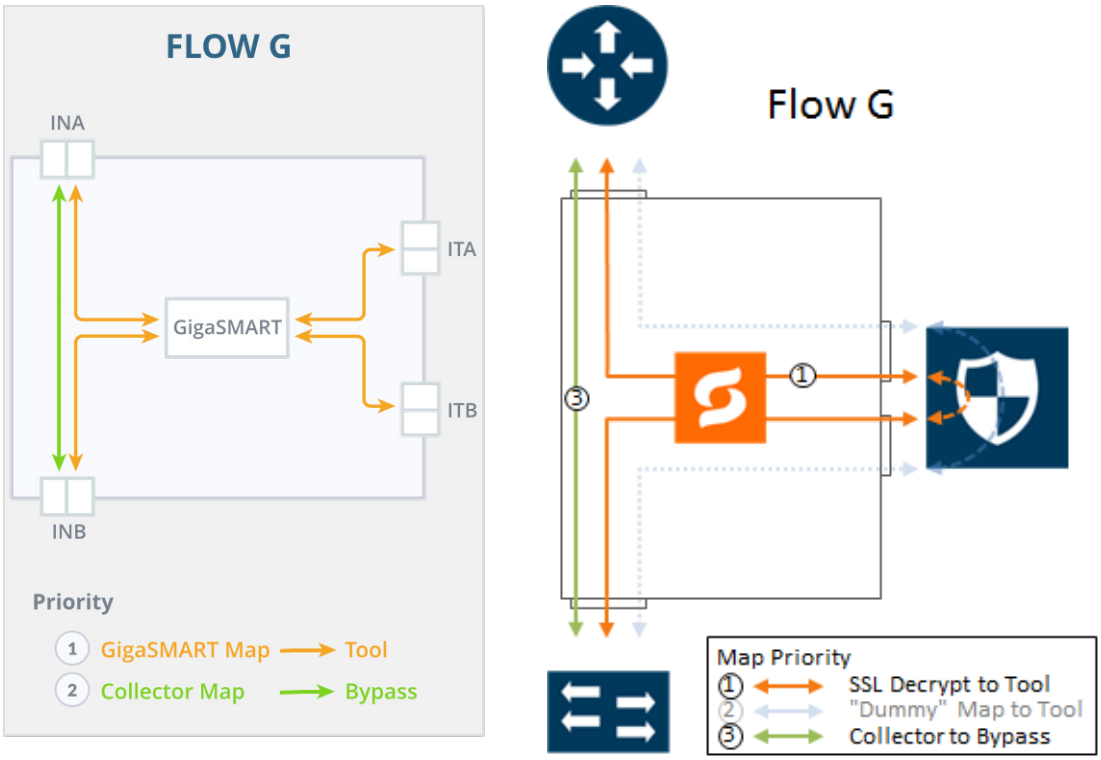

Figure 26-25: FLOW G Views

The map priorities of Flow G are as follows:

1. selectively forward traffic for decryption
2. direct unselected traffic to a collector, which sends traffic to bypass

Configure Inline SSL Decryption Using GigaVUE-FM

This section describes how to configure inline SSL decryption using GigaVUE-FM.

Before configuring, review *Get Started with Inline SSL Decryption* on page 679 for pre-requisites and review *Introduction to Inline SSL Map Workflows* on page 682.

The Inline SSL Configuration workflow is as follows:

- Keychain Password
- Key Store
- Signing CA
- Trust Store
- Policy Profile
- Network Access

The Inline SSL Map workflow (for Flow B) is as follows:

- Inline Network(s)
- Inline Tool(s)

- GS Group
- Virtual Port
- GS Operation
- Inline Rule Based Map
- Inline First Level Map
- Inline Second Level Map
- Collector Map (bypass)

To configure inline SSL decryption:

1. Go to **Physical Nodes** and select a GigaVUE-HC1, GigaVUE-HC2, or GigaVUE-HC3.
2. Go to **Workflows** and select **Inline SSL Configuration** from the Inline GigaSMART Operations section. Refer to Figure 26-17.

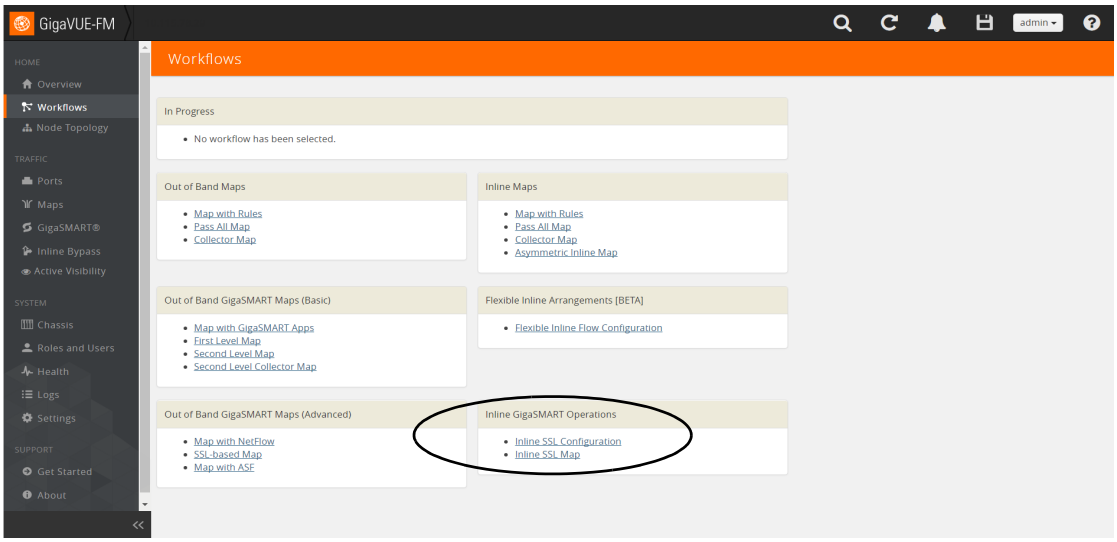

Figure 26-26: Select Inline SSL Configuration Workflow

3. The first step in the Workflow is **Keychain Password**. Click **Setup Keychain Password** to configure a keychain password.

   NOTE: The keychain password must be configured before installing certificates and keys. If the key has a passphrase, in order to install it, the keychain password and the passphrase must match.

4. Enter the Password. Hover the mouse over the **?** to review the password requirements. Confirm the password.
5. Click **Submit**. The keychain password is setup.
6. Click **Next**. The next step in the Workflow is **Key Store**.
7. Click **Add Key Pair** to configure the primary signing certificate and private key. The primary CA re-signs certificates for servers that present a valid certificate.
8. Enter an alias for the Key Pair. Click **RSA** or **ECDSA** for Key Type. Click **PEM** or **PKCS12** for Type. Click one of **Copy and Paste**, **Install from URL**, or **Install from Local Directory**.

9. Click **OK**. The Primary Key Pair is added and can be selected from the Key Pair drop-down menu.
10. Still under **Key Store**, repeat Steps 7 to 9 to configure the secondary signing certificate and private key.
11. Click **Next**. The next step in the Workflow is **Signing CA**. Click **Configure Signing CA**.
12. Map each of the key pairs installed to the Primary Root CA and Secondary Root CA.
13. Click **OK**. The signing CA is configured.
14. Click **Next**. The next step in the Workflow is **Trust Store**. No configuration is required if you use the default Trust Store.
15. Click **Next**. The next step in the Workflow is **Policy Profile**. Click **Create** to configure an inline SSL profile. The profile specifies policy configuration, such as certificate handling and actions to take for the profile.
16. Enter an alias for the profile. Under Policy Configuration, select a Default Action of Decrypt. Under Security Exceptions, select Decrypt or Drop.
17. Under Whitelist/Blacklist, select Whitelist and Blacklist. Enter the paths to the files.
18. Under Policy Rules, click **Add a Rule**. Select Category from the Rule drop-down menu. Select financial_services from the Category drop-down menu. Add another rule. Select Category from the Rule drop-down menu. Select health_and_medicine from the Category drop-down menu. Add another rule. Select Domain from the Rule drop-down menu. Enter youtube.com in **Value** text box for the Domain. Click **OK**. The inline SSL profile is added.
19. Click **Next**. The next step in the Workflow is **Network Access**. Click **Configure Network Access**.
20. Select **DHCP** and **DHCP Enabled** for a specified GigaSMART module.
21. Click **OK**. The network access is configured.
22. Go to **Workflows** and select **Inline SSL Map** from the Inline GigaSMART Operations section.
23. Select **FLOW B**.
24. The first step in the Workflow is **Inline Network(s)**. Select a default inline network from the Inline Network(s) drop-down menu. This is a protected inline network.
25. Click **Next**. The next step in the Workflow is **Inline Tool(s)**.
26. Click **Create Inline Tool**. Then click **Port Editor**. In the Quick Port Editor, locate ports and select **Type** of Inline Tool from the drop-down menu. Click **Enable** for those ports.
27. Click **OK**. The inline tool port is added. Click **Close** to exit the Quick Port Editor.
28. Still under **Inline Tool(s)**, enter an alias for the inline tool. Select Port A and Port B from the drop-down menus for the inline tool port pair. Under Configuration, ensure that **Inline tool sharing mode** is selected. Under Heartbeats, select **Enable Regular Heartbeat**.
29. Click **OK**. The inline tool is configured.

30. Click **Next**. The next step in the Workflow is **GS Group**. Click **Create** to configure a GigaSMART group and associate it with a GigaSMART engine port. Enter an alias for the GigaSMART group and select a GigaSMART engine port from the Port List.
31. Click **OK**. The GigaSMART group is added.
32. Click **Next**. The next step in the Workflow is **Virtual Port**. Click **Create** to configure a virtual port.

NOTE: You cannot add multiple vports on the same gsgroup.

33. Enter an alias for the virtual port, select the GigaSMART group configured in Step 30, then select an Inline Failover Action.
34. Click **OK**. The virtual port is added.
35. Click **Next**. The next step in the Workflow is **GS Operation**. Click **Create** to configure a GigaSMART operation.
36. Enter an alias for the GigaSMART operation, select the previously configured GigaSMART group, select Inline SSL as the GigaSMART Operation (GSOP), then select the previously configured Inline SSL profile.
37. Click **OK**. The GigaSMART operation is added.
38. Click **Next**. The next step in the Workflow is **Inline Rule Based Map**.
39. Configure the inline rule-based map. This map directs traffic from the inline network to the inline tool, using a specified rule. It has the same source port as the inline first level map and the same destination port as the inline second level map. Enter an alias for the map, and select the map Type (Inline) and Subtype (By Rule), select the source inline network and the destination inline tool.
40. Click **Add a Rule** to specify a map rule. Click **Bi-directional**, select IPv4 Protocol from the Rule drop-down menu, and select TCP from the Protocol drop-down menu. Select Port Destination from the Rule drop-down menu and enter 80 in the text box for **Min**.
41. Click **OK**. The map is added.
42. Click **Next**. The next step in the Workflow is **Inline First Level Map**.
43. Configure the inline first level map. This map directs TCP traffic from the inline network to a virtual port (and to GigaSMART). Enter an alias for the map, and select the map Type (Inline First Level) and Subtype (Ingress to Virtual Port). Under Map Source and Destination, select the inline network as the source and the virtual port as the destination. Under Map Rules, click **Add a Rule**. Select IPv4 Protocol from the Rule drop-down menu, and select TCP from the Protocol drop-down menu.
44. Click **OK**. The map is added.
45. Click **Next**. The next step in the Workflow is **Inline Second Level Map**.
46. Configure the next map, which is the inline second level map. This map directs traffic from the virtual port, uses the inline SSL GigaSMART operation, and sends traffic to the inline tool. Enter an alias for the map and select the map Type (Inline Second Level) and Subtype (Egress from Virtual Port). Under Map Source and Destination, select the virtual port as the source and the inline tool as the destination, then select the inline SSL GigaSMART operation.
47. Click **OK**. The map is added.

48. Click **Next**. The next step in the Workflow is **Collector Map (bypass)**.
49. Configure a collector map for any unmatched traffic including non-TCP traffic, which is directed to bypass. Enter an alias for the map, and select the map Type (Inline) and Subtype (Collector), then select a Traffic Path of ByPass. Under Map Source and Destination, select the inline network as the source.
50. Click **OK**. The map is added.
51. Click **To Maps** when the workflow is complete.
52. Verify the maps created by the workflow.
53. For inline network ports, go to **Inline Bypass > Inline Networks**. Select the inline network port and click **Edit**. Under Configuration, select a traffic path of To Inline Tool. If using protected inline networks, disable Physical Bypass.

View Statistics

You can view the following inline SSL decryption statistics:

- *Inline SSL Session Statistics* on page 694
- *Monitor Statistics* on page 695

Inline SSL Session Statistics

To display inline SSL session statistics, go to **GigaSMART > Inline SSL > Session Statistics**. Refer to Figure 26-27.

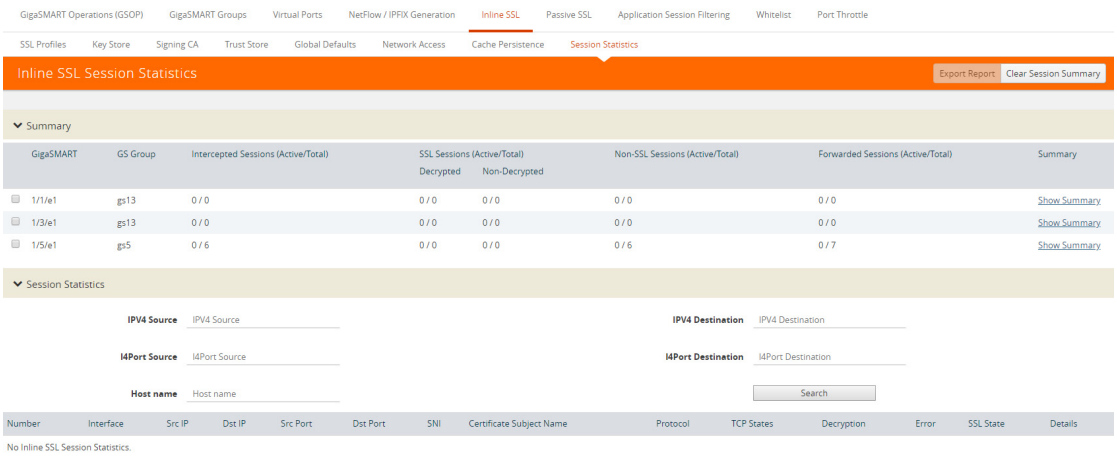

Figure 26-27: Inline SSL Session Statistics in FM

To display the inline SSL summary details, go to **GigaSMART > Inline SSL > Session Statistics**, view the Summary section, then click **Show Summary**. There are four sections: Session Statistics, Performance Statistics, Policy Statistics, and Certificate Statistics.

To search the inline SSL session statistics, go to **GigaSMART > Inline SSL > Session Statistics**. Enter an IPv4 source or destination, an L4 port source or destination, or a host name on which to search the session statistics.

Monitor Statistics

To display monitor statistics, go to **GigaSMART > Inline SSL > Monitor Statistics**.

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 26-28.

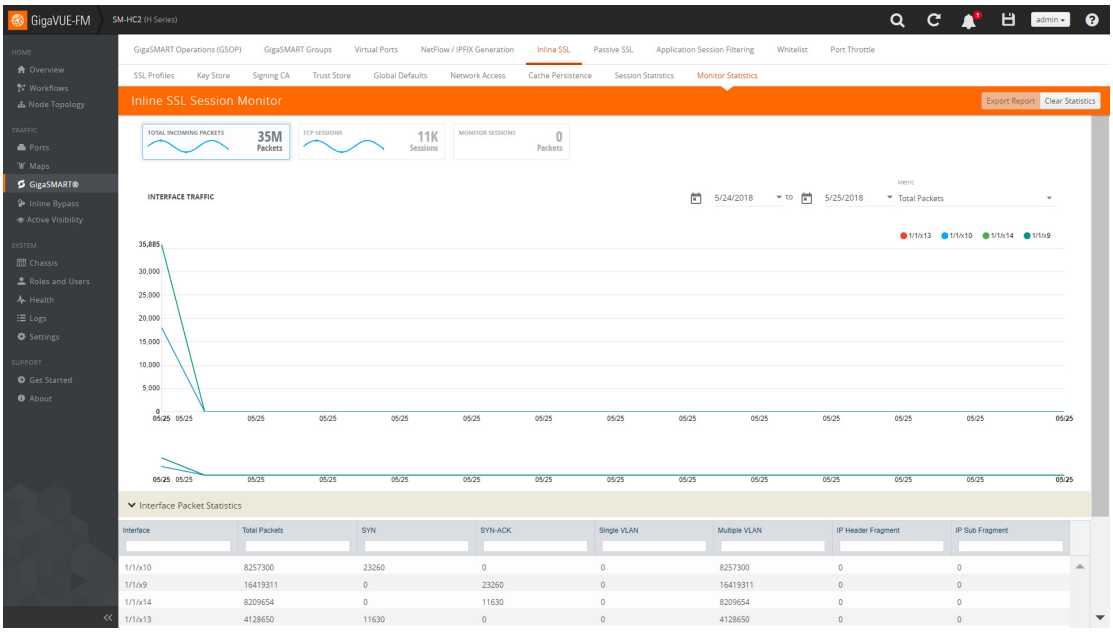

Figure 26-28: Inline SSL Session Monitor—Interface Packet Statistics

To display the graph and statistics for TCP Sessions, click the small graph, TCP SESSIONS. Refer to Figure 26-29.

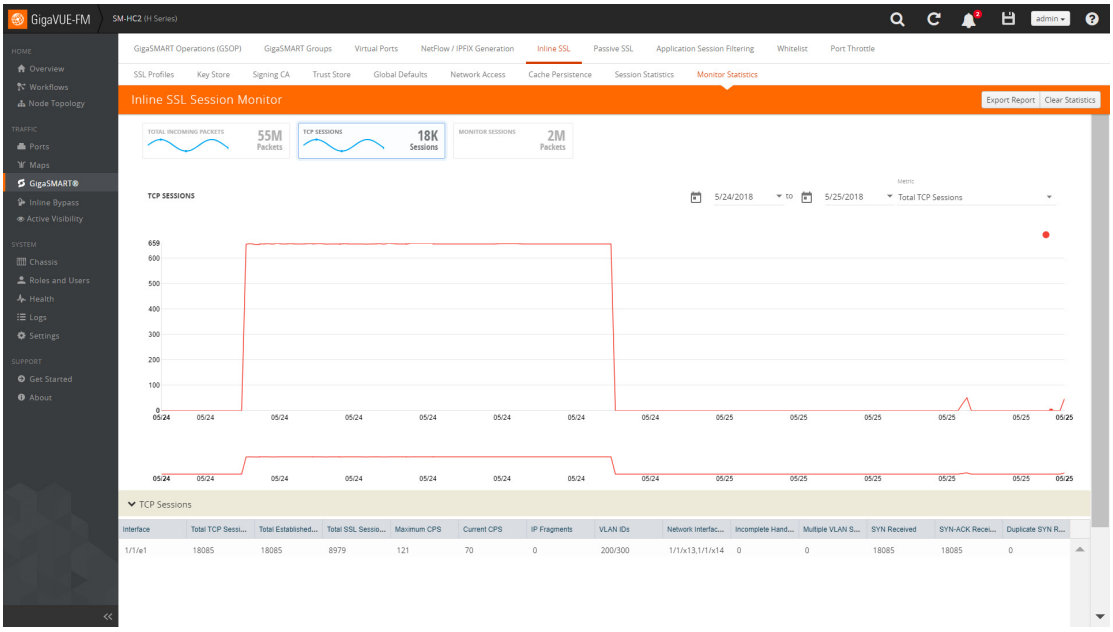

Figure 26-29: Inline SSL Session Monitor—TCP Sessions

To display Monitor Sessions, click the small graph, MONITOR SESSIONS. Refer to Figure 26-30.

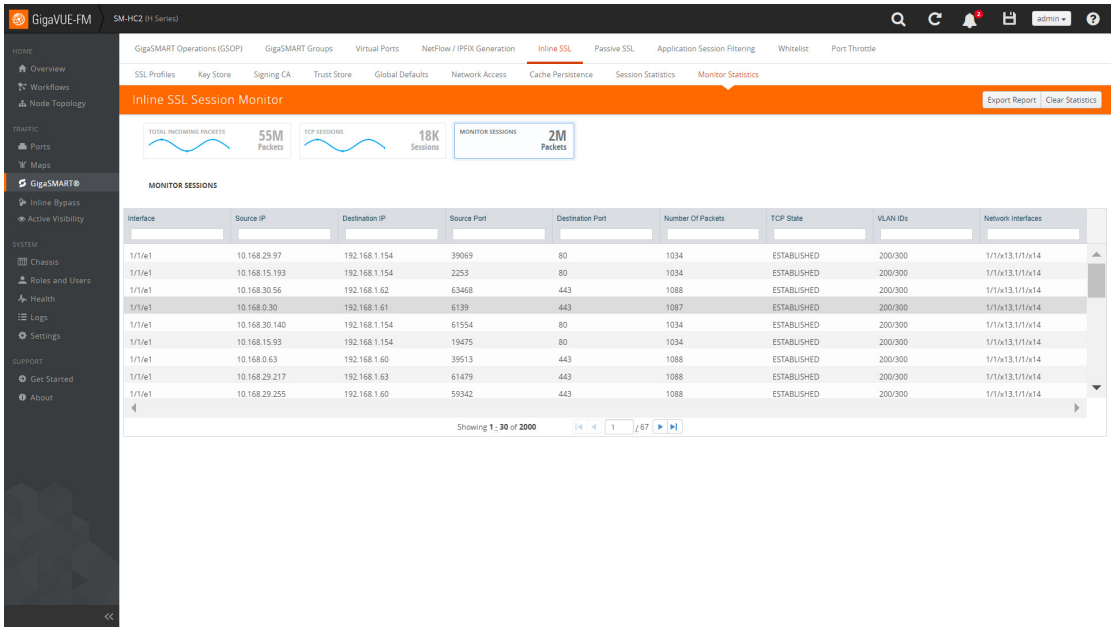

Figure 26-30: Inline SSL Session Monitor—Monitor Sessions

Configure Inline SSL Session Logging Server

You can configure an inline SSL session logging server to store the logged events that are generated when there are any changes made to the devices. You can specify the type of events that must be logged in to the server.

The following table provides a mapping of the severity, log level and its description:

Severity	Log Level	Description
0	Emergency	System is unusable
1	Alert	Action must be taken immediately
2	Critical	Critical condition
3	Error	Error condition
4	Warning	Warning condition
5	Notice	Normal but significant condition
6	Informational	Informational message
7	Debug	Debug message

The logged events are stored in the Common Event Format (CEF) as follows:

```
<SYSLOG_HEADER> <Timestamp> <hostname:engine> CEF:0|Gigamon|<Device
Model>|<GigaVUE OS Version>|<Event ID>|<Event
name>|<Severity>|[Extension]
```

Here is an example of a logged event:

```
Thu Jun 14 15:50:16 2018 hostname:hc2_test:1/1/e1
CEF:0|Gigamon|HC2|5.5.0|102|SESSION_DECRYPT|6|src=126.1.0.20
dst=126.1.0.10 spt=34267 dpt=443 dhost=example.com
cs1Label=Certificate Subject cs1=C\=US, ST\=CA, L\=Santa Clara,
CN=*.example.com cs2Label=Cipher Suite cs2=DHE-RSA-AES128-GCM-SHA256
```

You can view and track these logs to troubleshoot system issues, maintain audit trails, and for compliance purpose.

To configure an inline SSL session logging server using GigaVUE-FM:

Task	Description	UI Steps
1.	Configure a tool port.	**a.** From the device view, go to **Ports > All Ports**. **b.** Click **Quick Port Editor**. **c.** Use Quick search to find the port to configure. **d.** Set the type as Tool for the required port, and then select **Enable**. **e.** Click **OK**.
2.	Configure an IP interface with an IP address, subnet mask, default gateway, and MTU setting. Assign it to the GigaSMART group.	**a.** From the device view, go to **Ports > Ports > IP Interfaces**. **b.** Click **New**. The IP Interface page opens. **c.** In the **Alias** and **Comment** fields, enter the name and description of the IP interface. **d.** From the **Ports** drop-down list, select the tool port that you configured in step 1. **e.** Select the **Type** of the IP interface as IPv4 or IPv6. **f.** Enter the IP address, subnet mask, gateway, and MTU settings in the respective fields. **g.** From the **GS Groups** drop-down list, select a GigaSMART group under which you want to configure the inline SSL session logging server. **h.** Click **OK**.

Task	Description	UI Steps
3.	Configure the inline SSL session logging server under the GigaSMART group to which you assigned the IP interface in the task 2.	**a.** From the device view, go to **GigaSMART > GigaSMART Groups**. **b.** Choose the GigaSMART group to which you assigned the IP interface that you configured in task 2. **c.** Click **Edit**. **d.** Under **GigaSMART Parameters > Inline SSL Session Logging**, click **Add Remote Syslog Server**. **e.** In the **Remote Syslog IP** and **Remote Syslog Port Number** fields, enter the IP address and port number of the remote syslog server. **f.** From the **Associated IP Interface** drop-down list, select the IP interface that you assigned to the GigaSMART group in task 2. **g.** From the **Log Level** drop-down list, select the severity log level of the events that you want to send to the inline SSL session logging server. **h.** Click **OK**.

27 Flexible Inline Arrangements

This chapter provides an overview about the flexible inline arrangements, the supported platforms and software versions, the supported and unsupported functionalities, and limitations. It also provides details about how to configure the flexible inline maps and how to visualize the flexible inline arrangements canvas.

Refer to the following sections for details:

About Flexible Inline Arrangements

Flexible inline arrangements is 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 offers 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 “Configuring Inline Bypass Solutions” in the *GigaVUE-OS CLI User’s Guide*.

Using flexible inline maps, traffic from the same inline network can traverse different sequences of inline tools and share tools across traffic flows or with other inline networks.

You can identify specific flows of traffic using Layer 2 to Layer 4 rules, then designate the tools that will inspect the traffic, and specify the order of the

tools. For example, you can send Web traffic (defined by L4 port, 80 and/or 443) through a Web Application Firewall (WAF) and an Intrusion Prevention System (IPS), have backup traffic that might bypass inspection, and send all other traffic through the same IPS.

Figure 27-1 on page 700 illustrates a complex inspection scenario that is enabled by flexible inline arrangements.

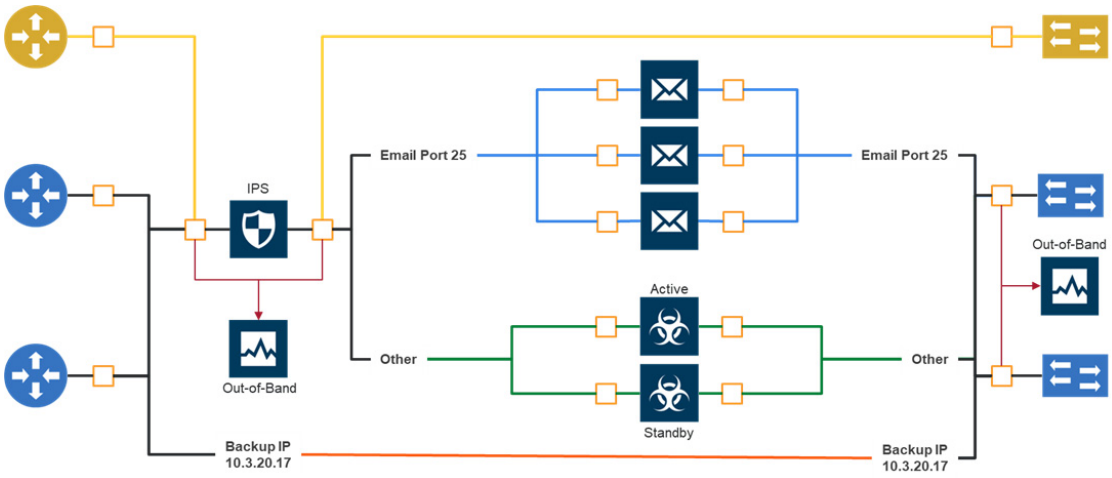

Figure 27-1: Flexible Inline Arrangements Scenario

In this example, on the left, there are three network links, all of which share the IPS. The details are as follows:

- At the top of the figure, the yellow line represents a flow of network traffic that only needs IPS inspection.
- In the middle of the figure, the Email and Other traffic represents network traffic flows that go through IPS inspection first, 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 27-1, the traffic in the reverse direction can have a different order of tools than the west-east traffic.

Supported Platforms

Flexible inline arrangements are supported on all GigaVUE HC Series nodes:

- GigaVUE-HC1
- GigaVUE-HC2 (CCv1 and CCv2)
- GigaVUE-HC3

The inline networks, inline tools, and inline tool groups involved in the flexible inline maps must be on the same node.

Software Version

GigaVUE-FM and GigaVUE-OS running on software version 5.3.xx or higher support the flexible inline arrangement functionality.

GRIP Supported by Flexible Inline Arrangements

Gigamon Resiliency for Inline Protection (GRIP™) is an inline bypass solution that connects two GigaVUE nodes together so that one node provides high availability to the other node when there is a loss of power. This redundant arrangement of two GigaVUE nodes maintains traffic monitoring by inline tools when one of the nodes is down. Flexible inline arrangements support GRIP.

Functionalities Not Supported by Flexible Inline Arrangement

Certain functionality is not supported by flexible inline arrangements. In the following cases, use classic inline bypass instead:

- for inline SSL decryption
- for inline tools that cannot tolerate the addition of VLAN tags to the traffic
- for asymmetrical hashing to an inline tool group in which one direction to hash is based on source IP address and the other direction to hash is based on destination IP address

Benefits of Flexible Inline Arrangements

Flexible inline arrangements offer flexibility in how traffic is guided through inline tools. It has the following benefits compared to classic inline bypass:

- Guides traffic through any arbitrary sequence of inspection tools.
- Shares inline tools across multiple inline network links and across multiple inline maps.
- Distributes traffic across multiple tools to meet bandwidth and throughput demands.

Flexible inline arrangements use the same software constructs as the classic inline bypass solution, such as inline network, inline tool, and inline tool group. However, inline network group and inline serial constructs are not needed.

Inline network groups have changed with flexible inline arrangements. Now every inline network is independent and can share any combination of tools in any order. The concept of inline network group is supported by creating multiple flexible inline maps. Also, multiple inline networks can be grouped into an inline network bundle. You can configure one inline map for the network bundle with the inline network bundle as the source.

Inline serial is not needed because the flexibility offered with flexible inline arrangements allows for the same configuration without the inline serial construct.

Figure 27-1 on page 700 illustrates the benefits of flexible inline arrangements by showing the kinds of deployment scenarios that can be enabled with this approach.

How to Use 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.

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.

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.

Configure 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.

Limitations of Flexible Inline Arrangements

If an inline tool is associated with a flexible inline map, it cannot be used in a classic inline map or in an inline SSL map. All inline networks and inline tools must participate exclusively in either flexible inline maps or classic inline maps.

Flexible Inline Arrangements Canvas

The GigaVUE-FM user interface provides clear visualization of inline maps. The user interface makes it easy to visualize and configure inline maps and tools. The drag-and-drop capability lets you define and add tools to maps, in any order. Refer to Figure 27-2 on page 704.

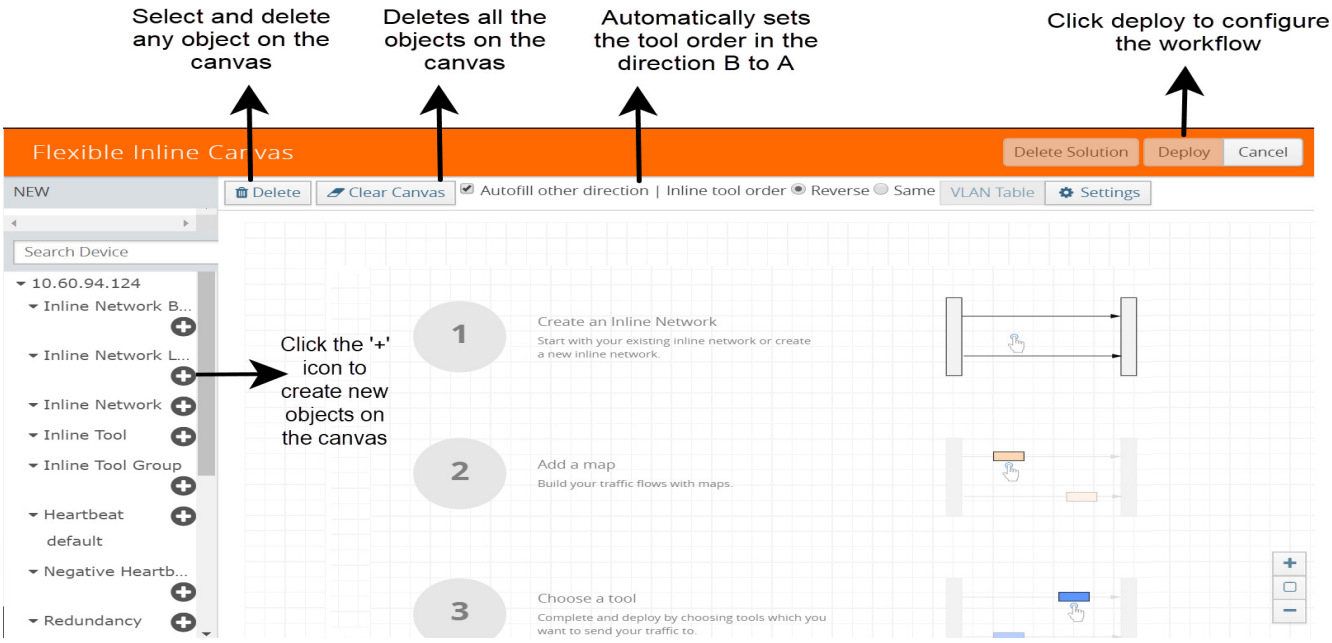

Figure 27-2: Flexible Inline Flow - Canvas

Configure Flexible Inline Flows

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

For detailed use cases of the different inline flows, refer to *"Flexible Inline Arrangements Deployment Guide for GigaVUE-OS 5.5"*.

Refer to the following sections for details:

- *Configure Inline Network Ports and an Inline Network* on page 705
- *Configure Inline Network Link Aggregation Group (LAG)* on page 706
- *Configure Inline Network Bundle* on page 709
- *Configure Inline Tool Ports and Inline Tools* on page 711
- *Configure Inline Tool Group* on page 713
- *Configure Inline Single Tag* on page 715
- *Configure Resilient Inline Arrangement* on page 718
- *Visualize Forwarding States of Inline Networks* on page 723

Configure Inline Network Ports and an Inline Network

An inline network consists of inline network ports, always in pairs, running at the same speed, on the same medium (either fiber or copper). The inline network ports must be on the same GigaVUE-HC series node.

Following are the two types of inline network:

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

To configure inline network ports and an inline network:

1. Go to **Physical** > **Inline Flows**, and then click **New** to create a new Flexible Inline Canvas.
2. In the Flexible Inline Canvas that is displayed, select the required device for which you want to configure the inline network.
3. Click the '+' icon next to the **Inline Network** option to create a new inline network. Refer to Figure 27-3 on page 705.

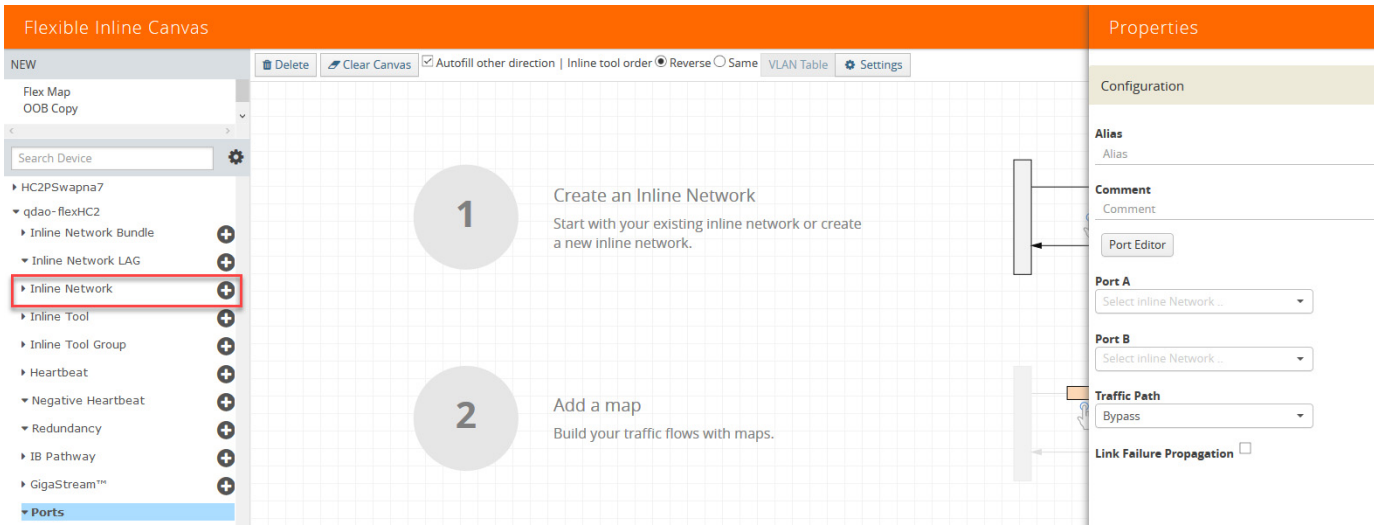

Figure 27-3: Inline Network Configuration

4. In the **Alias** and **Comment** fields, enter a name and description for the inline network, and then click **Port Editor**.
5. In the **Quick Port Editor**, scroll down to the inline network ports that you wish to configure. Select **Enable** to administratively enable inline network ports, and then click **OK**.

6. From the **Port A** and **Port B** drop-down lists, select the ports that you want to configure as the inline network pair.
7. From the **Traffic Path** drop-down list, select one of the following options:
   - **Bypass**—all traffic that originates from the inline network bypasses the sequence of inline tools and inline tool groups and is redirected to the opposite-side inline network port.
   - **Drop**—all traffic originating from the inline network is dropped.
   - **Bypass with Monitoring**—a copy of the traffic originating from the inline network bypasses the sequence of inline tools and inline tool groups and is redirected to the opposite-side inline network port. Another copy of the traffic is directed to the sequence of inline tools and inline tool groups, except that no traffic of the second copy is sent to the exit port.
   - **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.
8. Click **OK** to save the configuration.
9. Drag the **Inline Network** object to the canvas and click **Deploy**. Refer Figure 27-4 on page 706

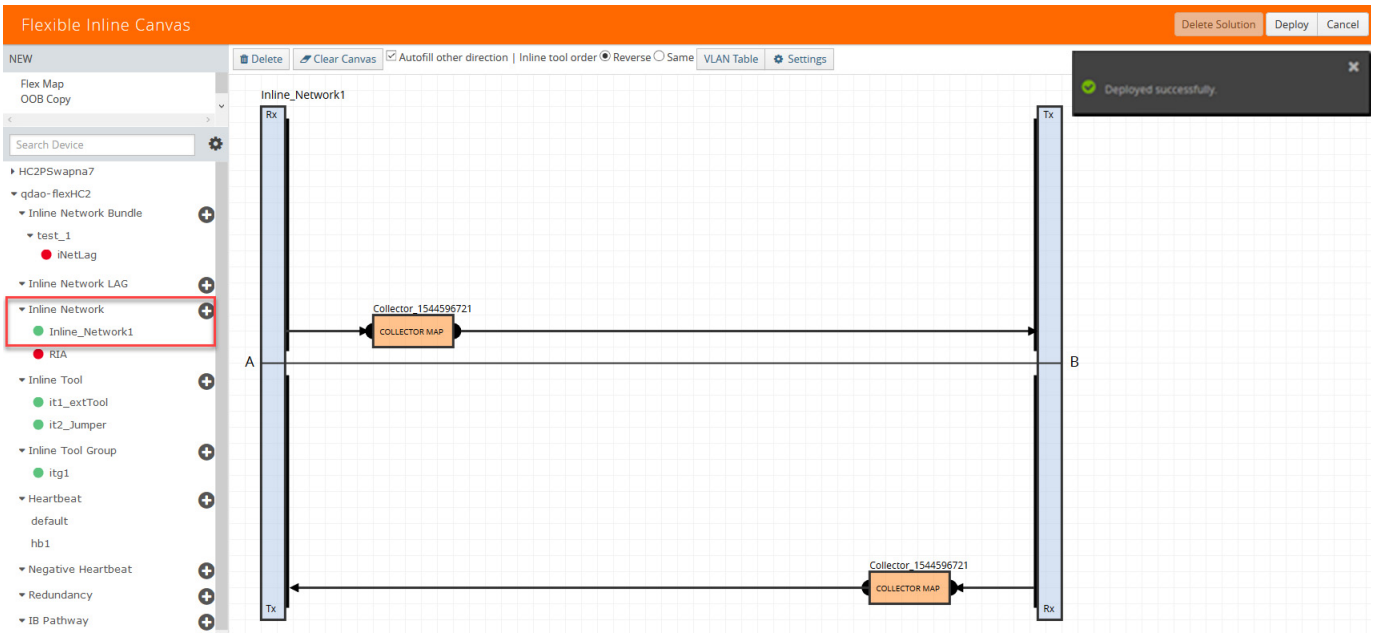

Figure 27-4: Inline Network Deployed

Configure Inline Network Link Aggregation Group (LAG)

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

- *About Inline Network LAG* on page 707
- *Limitations of Inline Network LAG* on page 707

- *Configure an Inline Network LAG* on page 707

About Inline Network LAG

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

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

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

Limitations of Inline Network LAG

Following are the limitations of the flexible inline network LAG feature:

- You cannot combine protected and unprotected inline networks, or inline networks with different speed in an inline network LAG.
- Link Aggregation Control Protocol (LACP) is not supported.

Configure an Inline Network LAG

Before you configure an inline network LAG, ensure that you configure the required inline network ports and inline networks. Refer to *Configure Inline Network Ports and an Inline Network* on page 705.

To configure an inline network LAG:

1. Go to **Physical** > **Inline Flows**, and then click **New** to create a new Flexible Inline Canvas.
2. In the **Flexible Inline Canvas** that is displayed, select the required device for which you want to configure the inline network LAG.
3. Click the '+' icon next to the **Inline Network LAG** option to create a new inline network LAG. Refer Figure 27-5 on page 708.

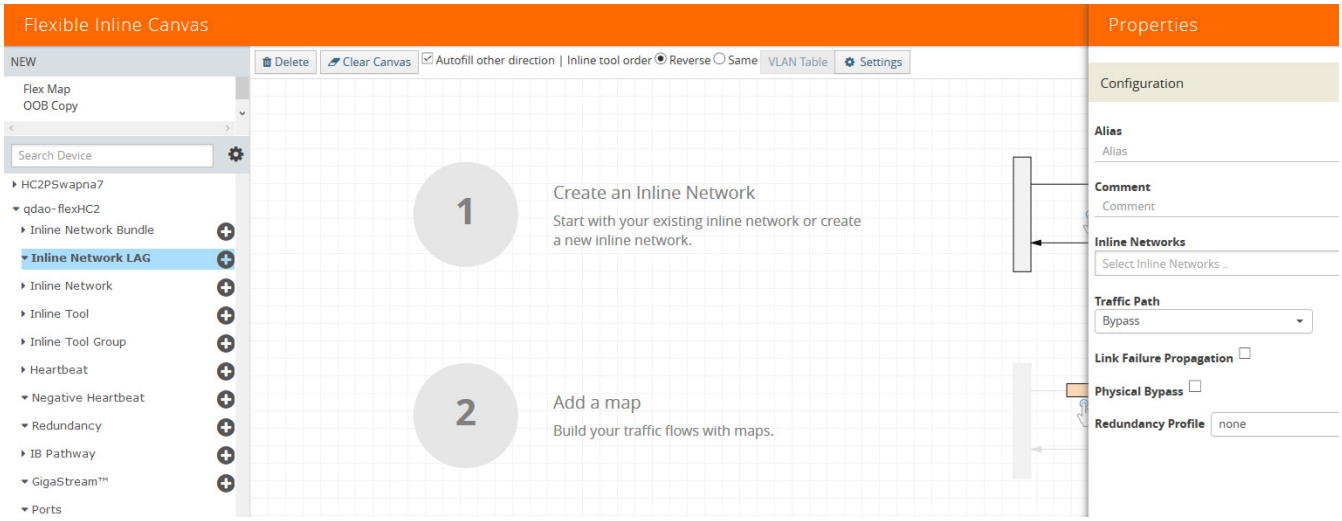

Figure 27-5: Inline Network LAG Canvas

4. In the Properties pane that appears on the right, enter the following information:
   a. In the **Alias** and **Comment** field, enter the name and description of the inline network LAG.
   b. From the **Inline Networks** drop-down list, select the required inline networks that need to be part of the inline network LAG.
   c. From the **Traffic Path** drop-down list, select the required traffic path for the inline network LAG.
   d. Enable the **Link Failure Propagation** option, if required.
   e. Enable the **Physical Bypass** option, if required.
5. Click **OK** to save the configuration.
6. Drag the **Inline Network LAG** object to the canvas.
7. Configure the required flexible inline maps and then, click **Deploy**. Refer Figure 27-6 on page 709.

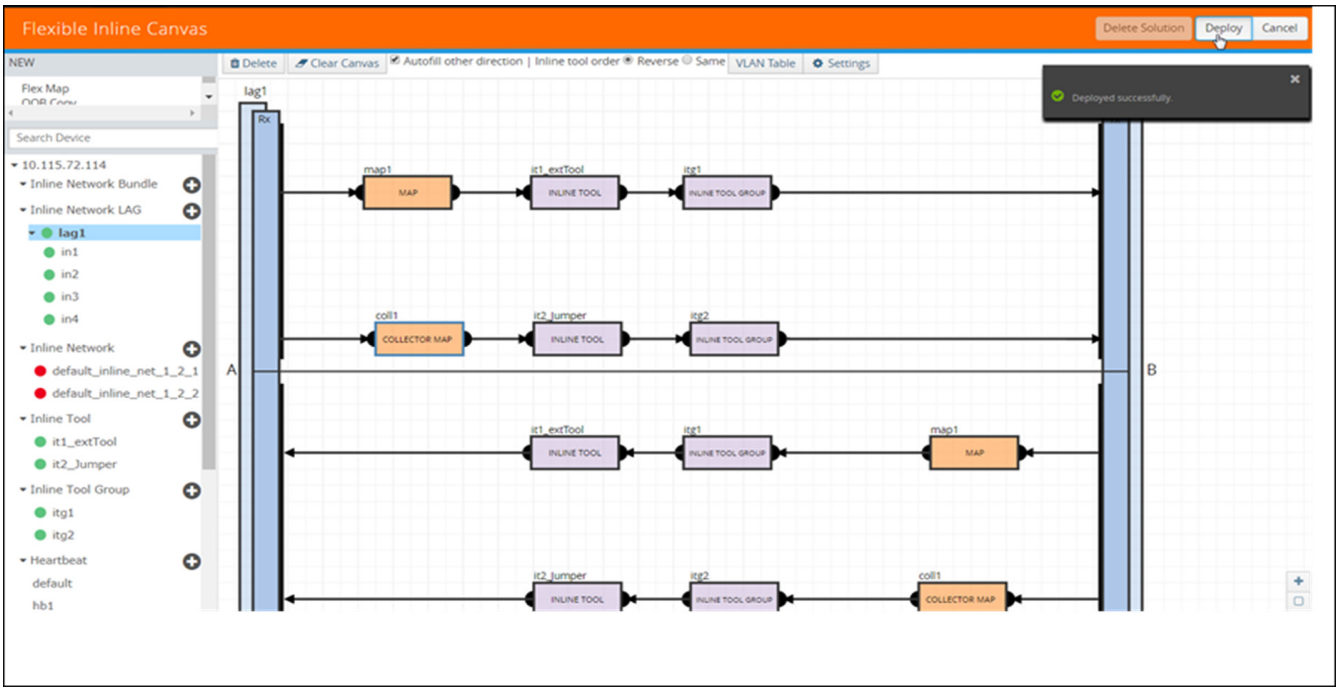

Figure 27-6: Inline Network LAG Deployed

Configure Inline Network Bundle

The Flexible inline network bundle feature allows you to group multiple inline networks into an inline network bundle. You can configure flexible inline maps with the inline network bundle as the source. GigaVUE-FM configures separate inline maps for each inline networks that are grouped in the inline network bundle. The inline maps are configured based on the Tool Side VLAN tags for the multiple inline networks and the rules that you specified when configuring the inline map for the network bundle.

Before you configure an inline network bundle, ensure that you configure the required inline network ports and inline networks. Refer to *Configure Inline Network Ports and an Inline Network* on page 705.

To configure an inline network bundle:

1. Go to **Physical** > **Inline Flows**, and then click **New** to create a new Flexible Inline Canvas.
2. In the **Flexible Inline Canvas** that is displayed, select the required device for which you want to configure the inline network bundle.
3. Click the '+' icon next to the **Inline Network Bundle** option to create a new inline network bundle.
4. In the **Alias** field, enter the name of the inline network bundle.
5. From the **Inline Networks** drop-down list, select the required inline networks that you want to add to the inline network bundle.
6. Click **OK** to save the configuration.

7. Drag and drop the inline network bundle into the canvas, and then configure the required inline map. Refer Figure 27-7 on page 710.

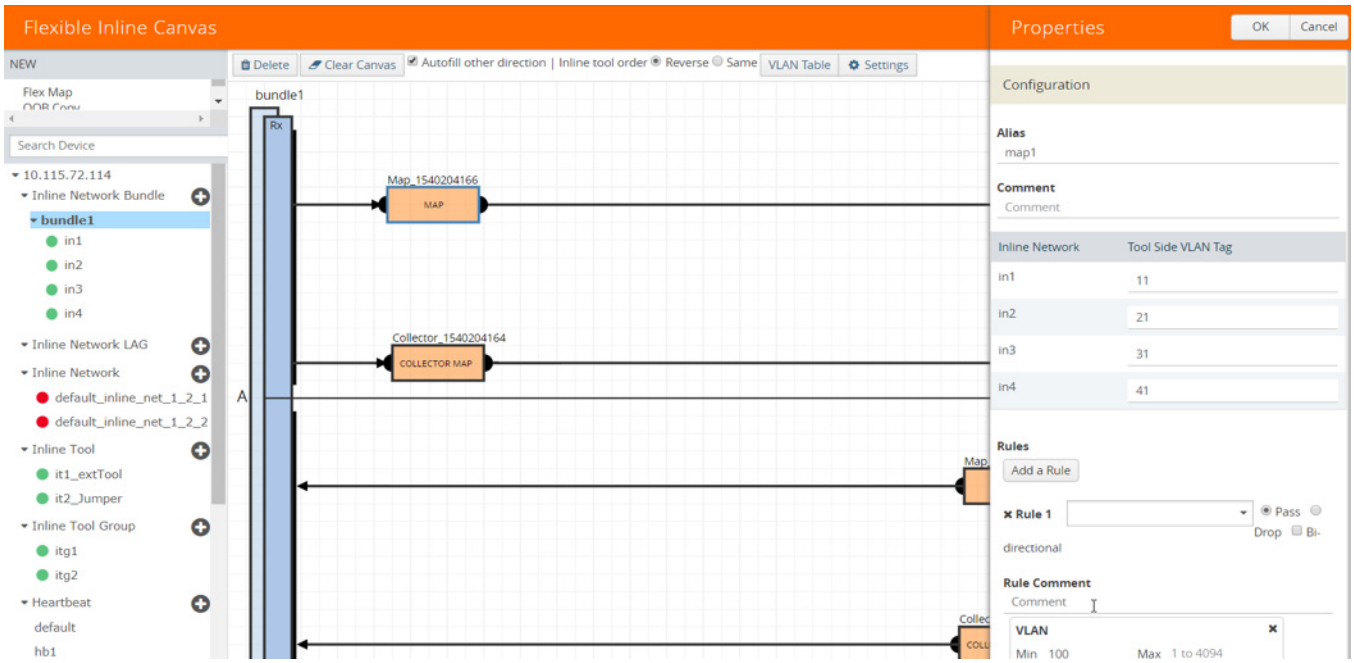

Figure 27-7: Inline Network Bundle—Map Configuration

8. Enter the **Tool Side VLAN Tag** for each inline network added in the inline network bundle.
9. Add the required rules for the inline map, and then click **OK** to save the configuration.

10. Click **Deploy**. GigaVUE-FM configures separate inline maps for each inline networks that are grouped in the inline network bundle. Refer Figure 27-8 on page 711.

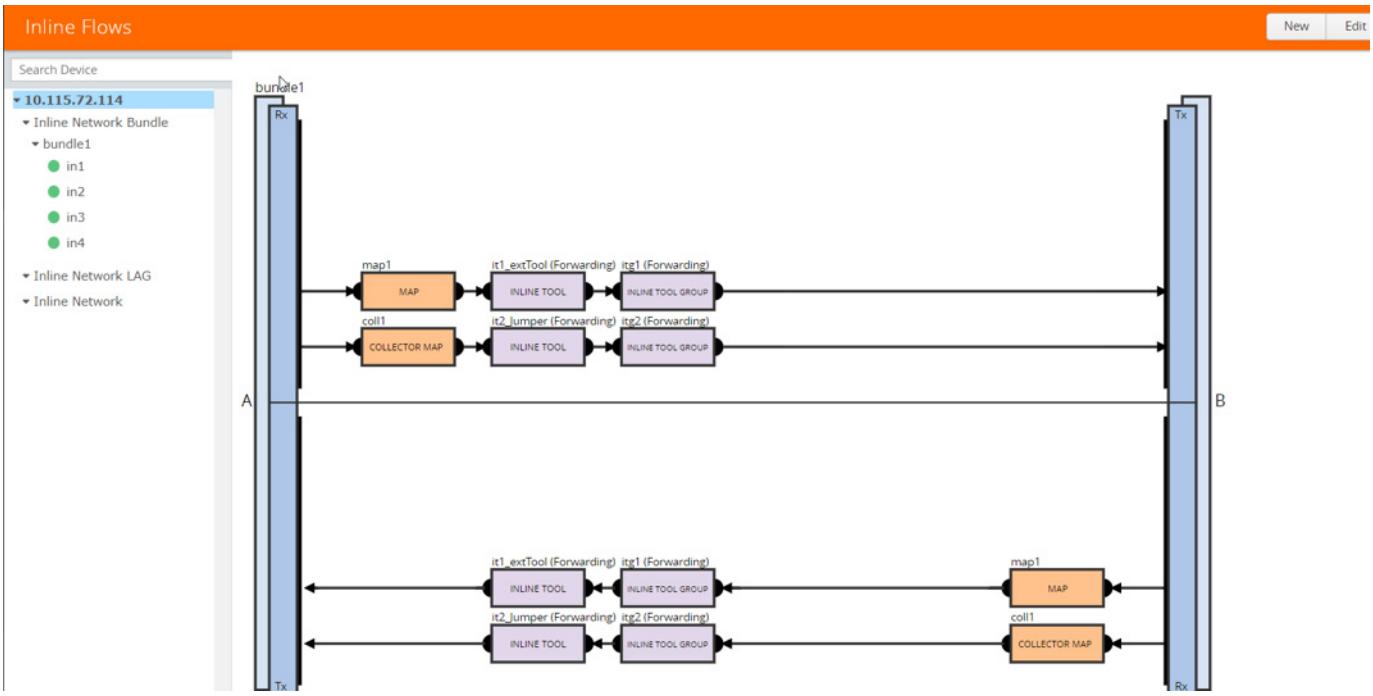

Figure 27-8: Inline Network Bundle—Deployed

Configure Inline Tool Ports and Inline Tools

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

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

To configure the inline tool ports and the inline tools:

1. Go to **Physical** > **Inline Flows**, and then click **New** to create a new Flexible Inline Canvas.
2. In the Flexible Inline Canvas that is displayed, select the required device for which you want to configure the inline tool.
3. Click the '+' icon next to the **Inline Tool** option to create a new inline tool.
4. In the **Properties** pane, in the **Alias** and **Comment** fields, enter a name and description for the inline tool, and then click **Port Editor**.
5. In the **Quick Port Editor**, scroll down to the inline tool ports that you wish to configure. Select **Enable** to administratively enable the inline tool ports, and then click **OK**.

6. 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.
7. Verify that the **Enabled** check box is selected.
8. From the **Failover action** drop-down list, select one of the following options:
   - **Tool Bypass**—For every map involving the inline tool or inline tool group that triggered this failover action, the traffic coming to such an inline tool or inline tool group is redirected to the next inline tool or inline tool group in the ordered list defined in Port A and Port B or to the respective inline network port.
   - **Network Bypass**—All inline networks configured as the source of any map involving the inline tool or inline tool group that triggered this failover action, will be put in the bypass mode, that is, all traffic coming to side A will be directed to side B and vice versa.
   - **Tool Drop**—For every map involving the inline tool or inline tool group that triggered this failover action, the traffic coming to such an inline tool or inline tool group is dropped and the traffic is redirected to a dummy VLAN with no members.
   - **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.
9. Select the **Recovery Mode** as **manual** or **automatic**.
10. Select the **Enable** check box for the **Inline tool Sharing mode** if you want to define additional tags on the tool side.

    NOTE: If you choose to disable the **Inline tool Sharing mode,** the inline tool can be used only in one flexible inline map.
11. From the **Flex Traffic Path** drop-down list, select one of the following options:
    - **Drop**—Traffic is dropped at the inline tool.
    - **Bypass**—Traffic bypasses the inline tool. Use this option for performing maintenance on an inline tool.
    - **Monitoring**—Traffic is fed to the inline tool and absorbed, while a copy of the traffic is sent to the next inline tool in the sequence. Traffic returned from side B of the network is also absorbed at the inline tool in the monitoring mode.
    - **To Inline Tool**—Traffic is forwarded to the inline tool.
12. Select the **Enable** check box below the **Regular Hearbeat**, if required, and then from the **Regular Hearbeat Profile** drop-down list, select a suitable profile.
13. In the HB IP Address A and HB IP Address B fields, enter the IP address of side A and side B defined in the Heartbeat profile.
14. Select the **Enable** check box below the **Negative Hearbeat**, if required, and then from the **Negative Hearbeat Profile** drop-down list, select a suitable profile.
15. Click **OK** to save the configuration.

16. Drag the **Inline Tool** object to the canvas.
17. Configure the required flexible inline maps and then, click **Deploy**.

Configure Inline Tool Group

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

The inline tool ports that make up the inline tools participating in the inline tool group are always in pairs, running at the same speed, on the same medium (fiber or copper). All inline tool ports of the inline tool group must be on the same GigaVUE-HC3 or GigaVUE-HC2 node, but can be on different modules on the node. On the GigaVUE-HC1, all the inline tool ports of the inline group must be on either the base module or the bypass combo module. The inline tool ports must also be on the same GigaVUE-HC3, GigaVUE-HC2, or GigaVUE-HC1 node as the inline network ports.

In a cluster environment, you can configure the inline tool group on GigaVUE HC Series nodes through the cluster master. The inline constructs must all be configured on one GigaVUE-HC3, GigaVUE-HC2, or GigaVUE-HC1 node, not across nodes, even if the nodes are in a cluster.

Resilient weighted hashing provides you the ability to distribute traffic to the inline tools by assigning either an equal weight or a custom weight to the inline tools. You can assign custom weight in percentage or ratio. If an inline tool in a group goes down and the group maintains the **Minimum Healthy Group Size** that is defined for the group, the traffic is redistributed to the remaining tools based on the equal weight or the custom weight assigned to the tools. If the inline tool group does not meet the **Minimum Healthy Group Size** defined for the group, the traffic is redistributed based on the **Failover Action** defined for the group.

NOTE: Resilient hashing is not supported for classic inline maps.

Before you configure an inline tool group, ensure that you configure the required inline tools. Refer to *Configure Inline Tool Ports and Inline Tools* on page 711.

To configure an inline tool group:

1. Go to **Physical** > **Inline Flows**, and then click **New** to create a new Flexible Inline Canvas.
2. In the Flexible Inline Canvas that is displayed, select the required device for which you want to configure the inline tool group.
3. Click the '+' icon next to the **Inline Tool Group** option to create a new inline tool group.
4. In the **Properties** pane, in the **Alias** and **Comment** fields, enter a name and description for the inline tool group.
5. From the **Inline Tools** drop-down list, select the required inline tools.
6. From the **Weighting** drop-down list, select one of the following options:

- **Equal**—Traffic is distributed equally to all the inline tools in the inline tool group.
- **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.

7. From the **Inline Spare Tool** drop-down list, select the inline tool to which the traffic will be forwarded when the first failure occurs in the set of primary inline tools.

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

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

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

12. From the **Minimum Healthy Group Size** drop-down list, select the minimum number of inline tools that must be up so that the entire inline tool group is

considered to be up. The minimum number must include the inline spare tool as well.

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

Configure Inline Single Tag

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

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

To configure an inline single tag:

1. Go to **Physical** > **Inline Flows**, and then click **New** to create a new Flexible Inline Canvas.
2. In the Flexible Inline Canvas that is displayed, select the required inline network from the list of devices.
3. Drag and drop the inline network into the canvas.
4. Click **Settings** to open the **Settings** pane. Refer Figure 27-9 on page 716.

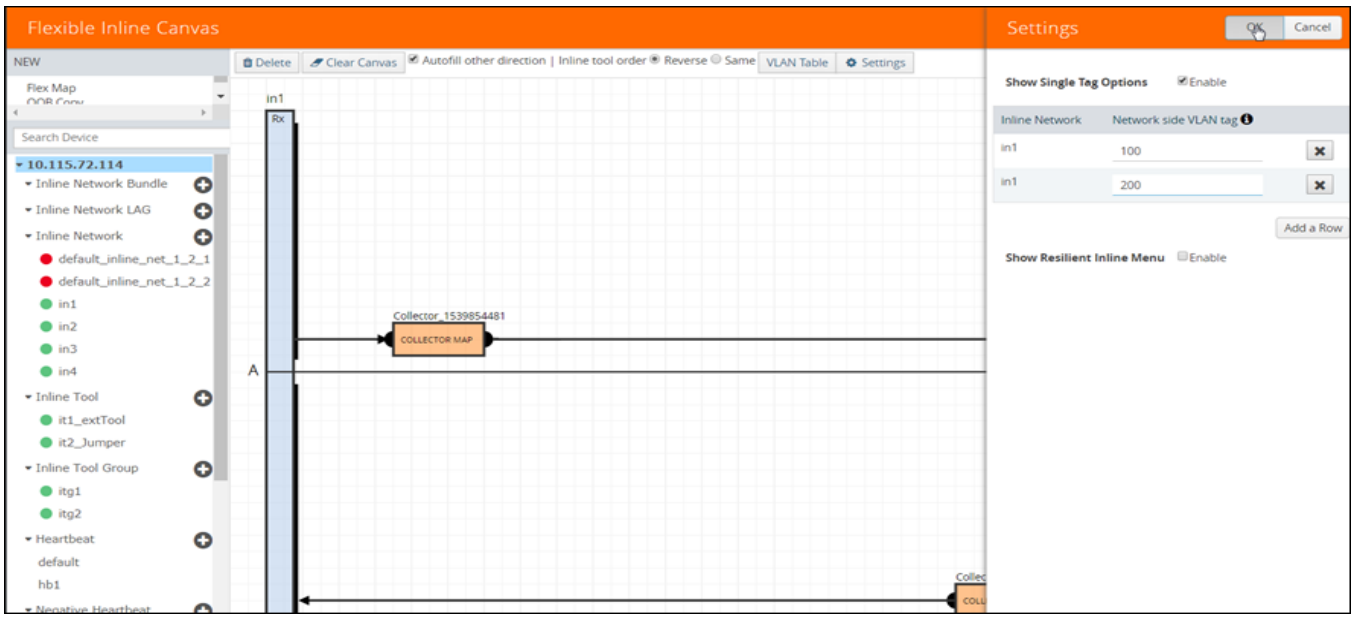

Figure 27-9: Flexible Inline - Enable Show Single Tag Options

5. Select the **Enable** check box for the **Show Single Tag Options**, and then enter the VLANS expected on the inline network.
6. Drag and drop a flexible inline map object into the canvas, and then click the map to open the **Properties** pane.
7. Select the **Enable** check box for the **Single Tag Mode**, and then enter the tool side VLAN tags. The VLAN qualifier is added to the rules by GigaVUE-FM. If you do not specify any rules, GigaVUE-FM adds a rule with the VLAN qualifier to the map. Refer Figure 27-10 on page 717.

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

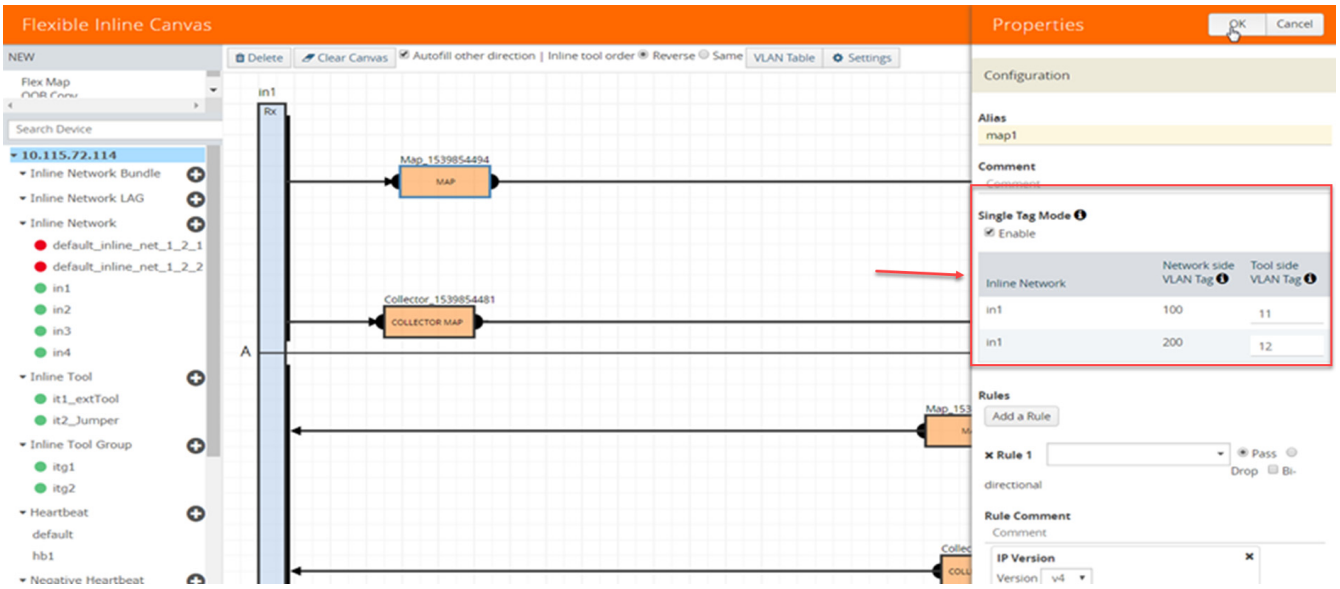

Figure 27-10: Flexible Inline - Single Tag -VLANS

8. Drag and drop the required inline tools into the canvas.
9. Drag and drop the OOB Copy into the canvas, if required. Refer Figure 27-11 on page 717.

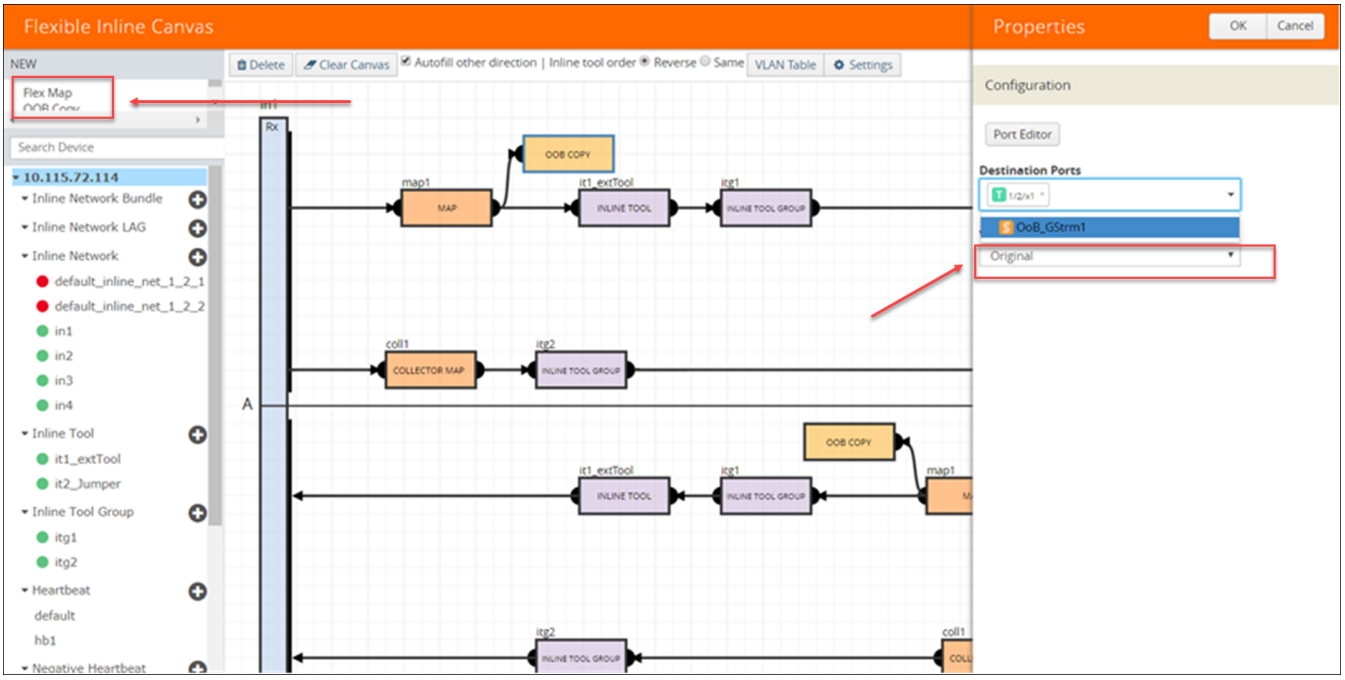

Figure 27-11: Flexible Inline Single Tag - OOB-Copy

10. Click **Port Editor**, and then in the **Quick Port Editor**, scroll down to the hybrid or tool ports that you wish to configure. Select **Enable** to administratively enable the ports, and then click **OK**.

11. From the **Destination Ports** drop-down list, select the required hybrid or tool ports that you want to configure as destination ports. You can also select a hybrid or tool GigaStream. For information about GigaStream, refer to *"Using GigaStream"* in the *GigaVUE-FM User's Guide*.
12. From the **VLAN Tag** drop-down list, select the required tags.
13. Click **OK** to save the configuration.
14. Click **Deploy**. Refer Figure 27-12 on page 718.

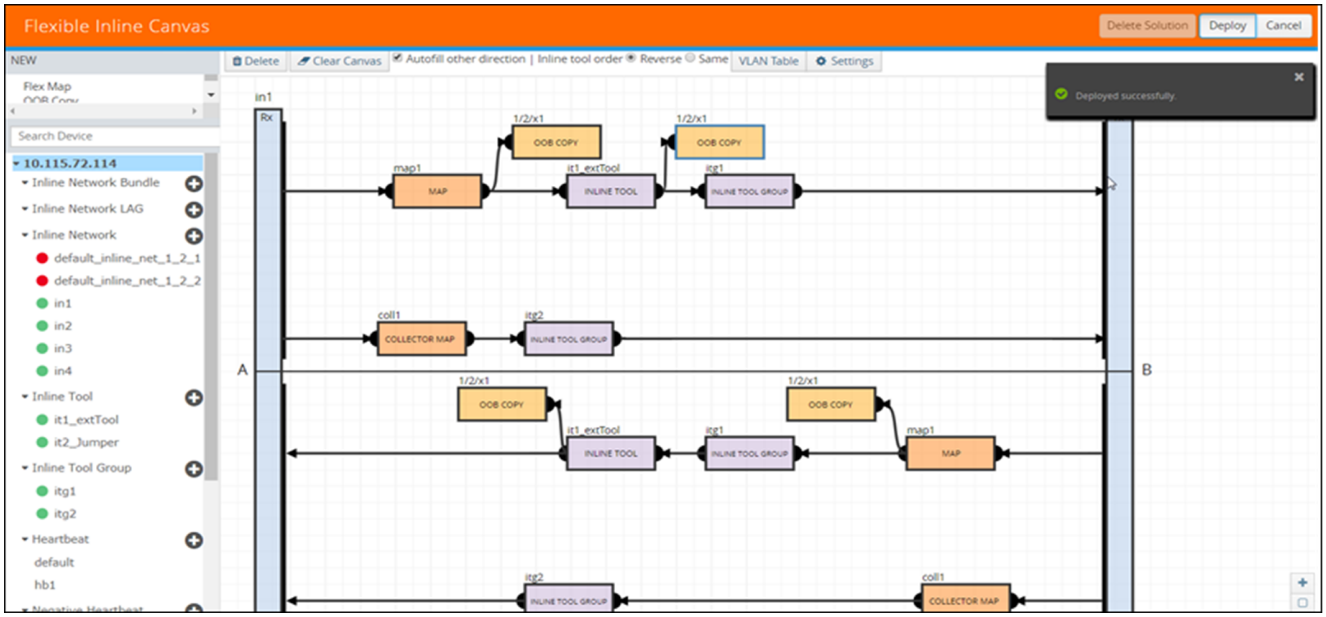

Figure 27-12: Flexible Inline - Single Tag Deploy

Configure Resilient Inline Arrangement

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

- *About Resilient Inline Architecture* on page 718
- *About Inter-broker Pathway (IB-P)* on page 720
- *Resilient Inline Arrangement—Rules and Notes* on page 720
- *Configure Resilient Inline Arrangement* on page 721

About Resilient Inline Architecture

Resilient inline arrangements 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.

The following figure illustrates the resilient inline architecture.

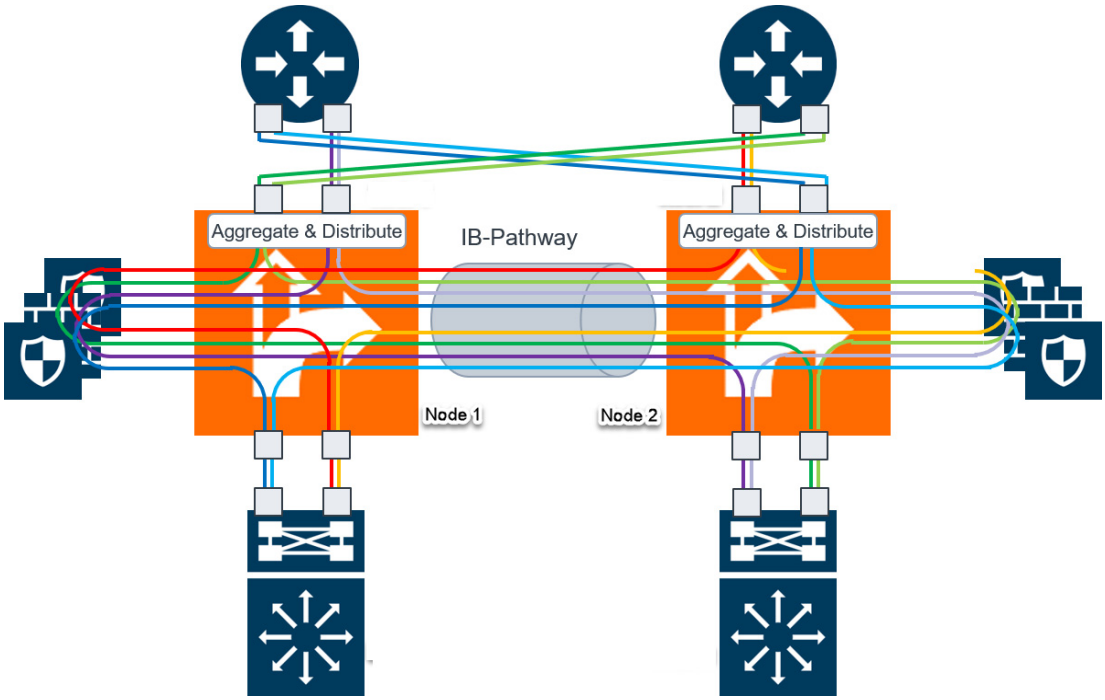

Figure 27-13: Resilient Inline Architecture

The resilient inline architecture 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.

About 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 opposite-side inline network port.
  - 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 (Refer Figure 27-13 on page 719):

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

Resilient Inline Arrangement—Rules and Notes

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

- Ensure that the names on both GigaVUE devices are identical, that is, the inline networks, inline tools, inline tool groups, and out-of-band tools must all have the same alias names on both the devices.

- If you choose to use the inline network bundle, the alias of the inline network bundle on both the devices must be identical. However, the inline networks that are grouped into the bundle can have different aliases.

Configure Resilient Inline Arrangement

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

- Configure the required inline networks. Refer to *Configure Inline Network Ports and an Inline Network* on page 705.
- Configure the required inline network LAG. Refer to *Configure Inline Network Link Aggregation Group (LAG)* on page 706.
- Configure the required inline tools. Refer to *Configure Inline Tool Ports and Inline Tools* on page 711.
- Configure the required inline tool group. Refer to *Configure Inline Tool Group* on page 713.

To configure resilient inline arrangement:

1. Go to **Physical** > **Inline Flows**, and then click **New** to create a new Flexible Inline Canvas.
2. In the Flexible Inline Canvas that is displayed, select the required device for which you want to configure the resilient inline arrangement.
3. Click the '+' icon next to the **IB Pathway** option to create a new inter-broker pathway.
4. In the **Properties** pane, in the **Alias** and **Comment** fields, enter a name and description for the inter-broker pathway, and then click **Port Editor**.
5. In the **Quick Port Editor**, scroll down to the inline tool ports that you wish to configure. Select **Enable** to administratively enable the required network ports, and then click **OK**.
6. From the **Ports** drop-down lists, select the network ports that you have configured.
7. In the **Minimum Ports Up** field, enter the minimum number of network ports that must be operationally up so that the status of the inter-broker pathway will be up.
8. 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 that you have configured.
9. Click **OK** to save the configurations.
10. Drag and drop the required inline network or inline network LAG in to the flexible inline canvas, and then click **Settings**.
11. In the **Settings** pane, select the **Show Resilient Inline Menu** check box.

12. Select **Node 1**, **Node 2**, **IB Pathway1**, **IB Pathway2**, **Hashing Source**, **Hashing LSB Node** that you want to configure for resilient inline arrangement.
13. Click **OK** to save the settings.
14. Drag and drop the flexible inline map into the canvas.
15. In the **Properties** pane, in the **Alias** and **Comment** fields, enter the name and description of the inline map.
16. Enter the **Tool Side VLAN Tag** for the inline network for which you are configuring the map.
17. From the **FlexInline Failover** drop-down list, select one of the following options:
    - **Bypass**—the traffic is passed directly between the respective inline network ports.
    - **Original Map**—the traffic is passed through the path that is defined in this flexible inline map.
18. Add the required rules for the inline map, and then click **OK** to save the configuration.
19. Drag and drop the required inline tools or inline tool group into the canvas.
20. Drag and drop the **OOB Copy** into the canvas, if required.
21. Click **Deploy**. Refer Figure 27-14 on page 722.

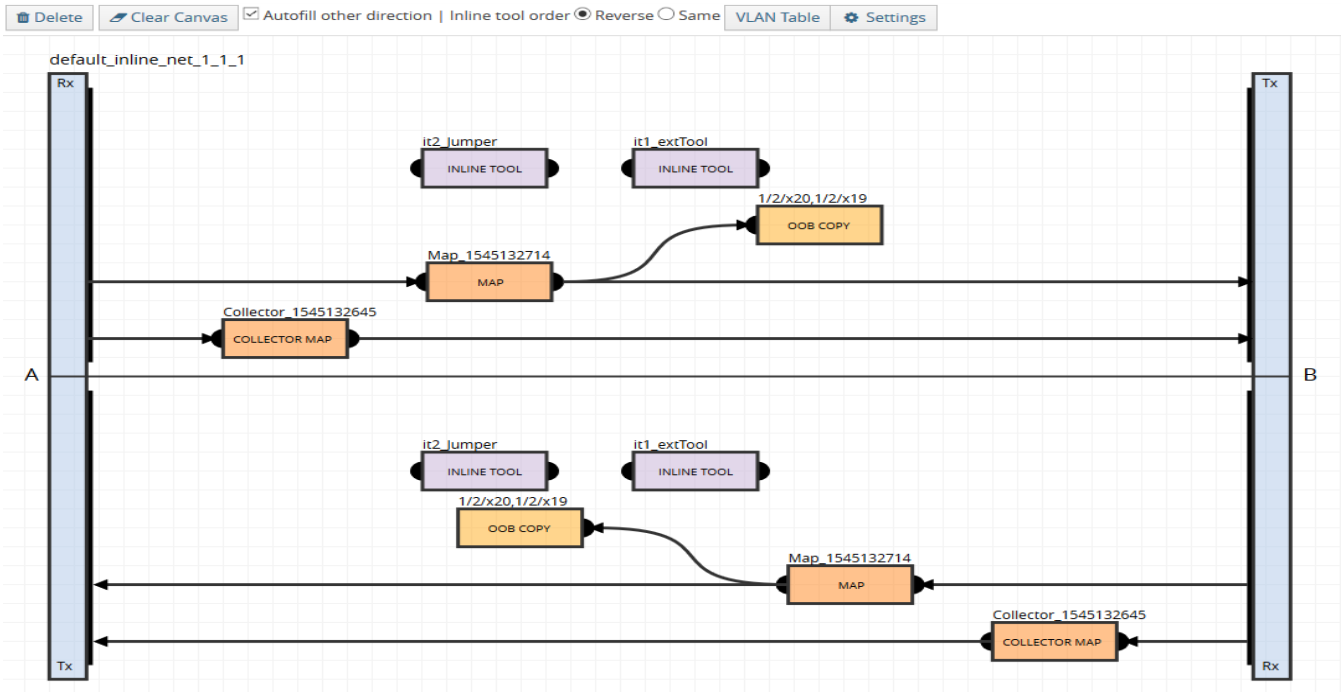

Figure 27-14: Resilient Inline Arrangement–Deployed

Visualize Forwarding States of Inline Networks

You can view the forwarding states of the inline networks in the flexible inline canvas. Refer to Figure 27-15 on page 723.

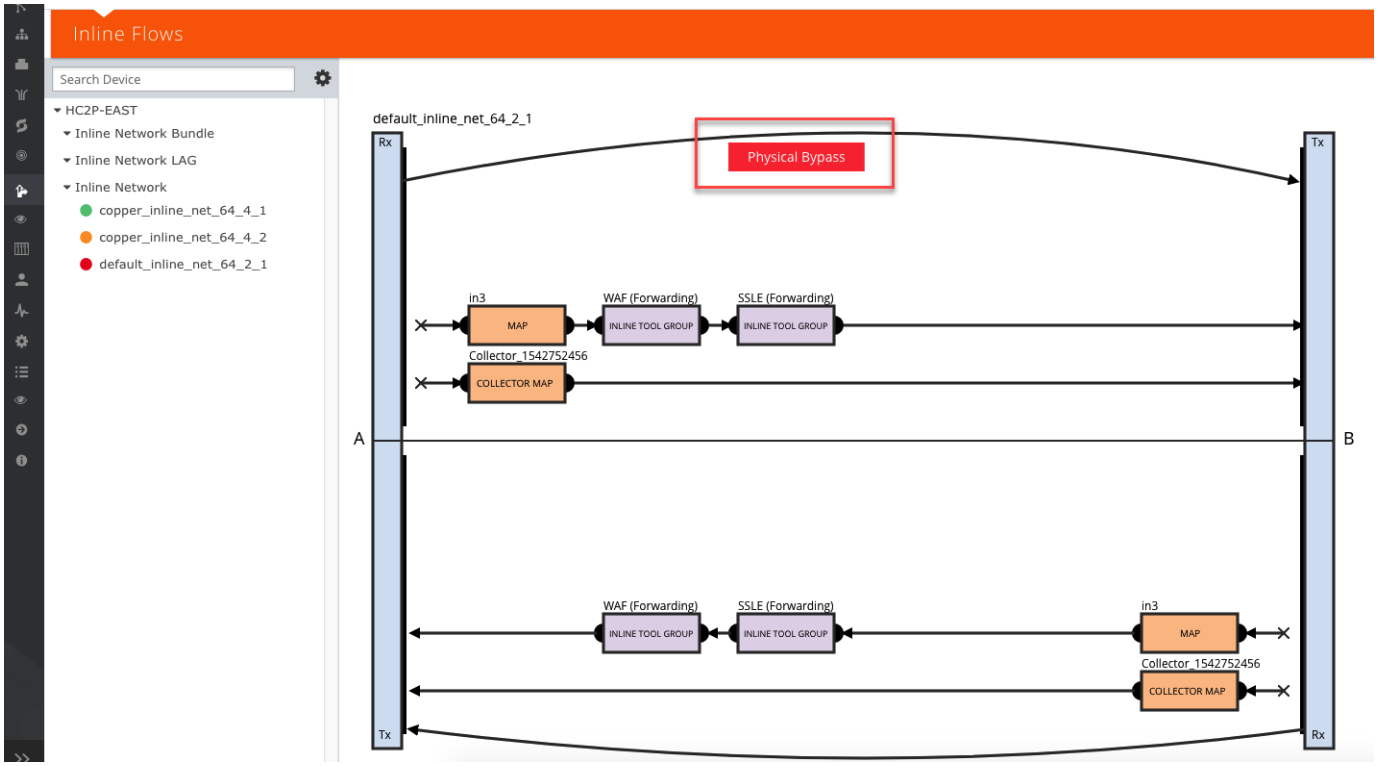

Figure 27-15: Inline Network Forwarding States

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.

Table 27-1 provides the list of forwarding states of inline network and their description.

Table 27-1: Forwarding State of Inline Networks

Inline Network Physical Bypass	Inline Network Traffic Path	Far-End Status of Links Connected to Inline Network Ports	Operational State of *Not* Forced Inline Tools and Inline Tool Groups Involved in Maps from the Inline Network	Forwarding State	Description
enable	any inline network traffic path configuration	any combination of far-end ports status	any combination of operational state of the inline tools or inline tool groups involved in the maps originating from the inline network	PHYSICAL BYPASS	all traffic exchanged directly between the end nodes without being noticed by the switching fabric (GigaVUE node acting as a wire or fiber)

Table 27-1: Forwarding State 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
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
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.

Table 27-1: Forwarding State 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
disable	traffic path set to to-inline-tool	both far-end ports are up	all inline tools involved (directly or indirectly as members of inline tool groups) in the maps originating from the inline network are in the *up* operational state	NORMAL	The traffic is guided between the source inline network port and the destination inline network port according to the status of the inline tools and inline tool groups **NOTE:** The state of all inline tools must be *up*, including inline tools configured as spare in an inline tool group, inline tools or inline tool group members in the **a-to-b** and **b-to-a** lists configured with any traffic path other than to-inline-tool.
disable	traffic path set to to-inline-tool	both far-end ports are up	at least one of the inline tools or 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	**a.** none of the inline tools or inline tool groups involved in the maps originating from the inline network configured with to-inline-tool and failover action network-port-forced-down is in the *down* operational state **b.** at least one of the inline tools or inline tool groups involved in the maps originating from the inline network configured with to-inline-tool and failover-action of network-drop is in the *down* operational state	FAILURE INTRODUCED DROP	All traffic arriving at the inline network ports is dropped

Table 27-1: Forwarding State 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
disable	traffic path set to to-inline-tool	both far-end ports are up	**a.** none of the inline tools or inline tool groups involved in the maps originating from the inline network configured with to-inline-tool and failover action of network-port-forced-down or network-drop is in the *down* operational state **b.** at least one of the inline tools or inline tool groups involved in the maps originating from the inline network configured with to-inline-tool and failover action of network-bypass is in the *down* operational state	FAILURE INTRODUCED 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 to-inline-tool	both far-end ports are up	any combination of conditions not listed for the forwarding state definitions of PHYSICAL BYPASS, DISABLED, DISCONNECTED, FORCED BYPASS, FORCED BYPASS WITH MONITORING, NORMAL, NETWORK PORTS FORCED DOWN, FAILURE-INTRODUCED DROP, or FAILURE-INTRODUCED BYPASS	ABNORMAL	The traffic is guided between the source inline network port according to the status of the inline tools and inline tool groups **NOTE:** If any inline tool involved in flexible inline maps (directly or indirectly as a member of an inline tool group) is in the *down* operational state and there is no network-level failover action in effect, the inline network is in the ABNORMAL state.

28 Application Intelligence

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

Required License: Application Filtering Intelligence. The Application Monitoring capability is included as part of this license.

Licensing and Limited Availability

Application Filtering Intelligence includes Application Monitoring functionality.

- Application Filtering Intelligence is generally available as of GigaVUE-FM 5.6.00.
- Application Metadata Intelligence is available for preview purposes only through Gigamon's Beta Program. Application Filtering Intelligence is required to run Application Metadata Intelligence.

  If you have interest in exploring Application Metadata Intelligence, please reach out to your Gigamon sales representative and ask to be included in the Beta Program. Refer to *GigaVUE-FM Application Metadata Intelligence User's Guide 5.6.00 (Beta)* for additional information about this feature.

About Application Intelligence

GigaVUE-FM Application Intelligence provides a comprehensive solution that:

- identifies the applications contributing to the network traffic.
- isolates preferred application-specific traffic and directs it to the appropriate tools.

- exports relevant application metadata for further analytics and analysis.

Application Intelligence provides the following capabilities:

- **Application Monitoring** - Identifies and monitors all applications contributing to the network traffic, and reports on the total applications and the total bandwidth they consume over a select period. Able to identify more than 3,000 applications. It displays the traffic statistics in bytes, packet and flows.
- **Application Filtering Intelligence** - Enables traffic filtering by layer 7 applications, which means you can filter out high-volume, low-risk traffic from reaching the tools and distribute high-risk network traffic of interest to the right tool at the right time.
- **Application Metadata Intelligence** - Supports exporting over 5200 attributes of metadata that provide relevant usage context on over 3,000 applications.

  NOTE: Application Metadata Intelligence is available for preview purposes only through Gigamon's Beta Program. If you have interest in exploring Application Metadata Intelligence, please reach out to your Gigamon sales representative and ask to be included in the Beta Program. Refer to *GigaVUE-FM Application Metadata Intelligence User's Guide 5.6.00* for additional information about this feature.

How Application Intelligence Works

Figure 28-1 illustrates how the Application Intelligence solution works.

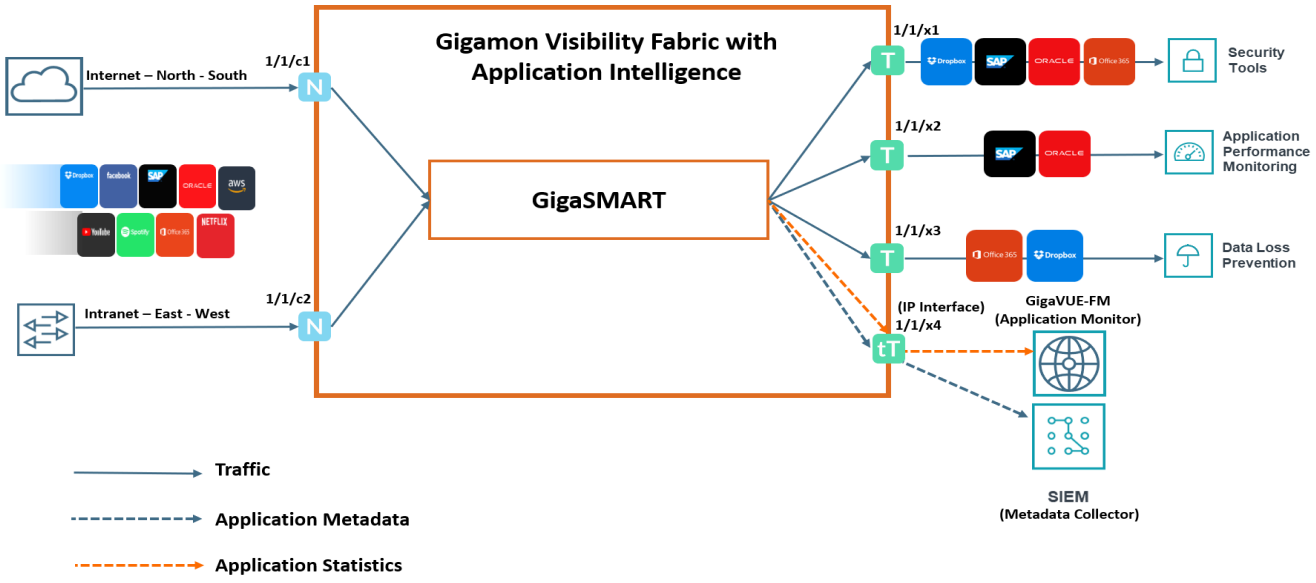

Figure 28-1: Application Intelligence Solution

The Gigamon device that is configured with the Application Intelligence capability receives the traffic through the network ports. Application Monitoring forwards the network traffic to the GigaSMART engine that exports applications related information to GigaVUE-FM, which renders the information on the Application Intelligence Dashboard. Based on the pass or drop rules configured in the maps, Application

Filtering Intelligence lets relevant applications be guided to the tools. In this example, you can see that except some of the audio or video streaming applications such as Spotify®, YouTube, and Netflix, other applications are filtered and sent to the security tool. This is because the audio or video streaming applications are high-volume, low-risk traffic. The threat detection tool need not inspect such traffic. Hence, these applications are dropped based on the configured drop rules.

Application Metadata intelligence lets exporting metadata from applications that are passed by the Application Filtering Intelligence. The records can be exported to a collector either in IPFIX or CEF format via the IP interface. You can also use the application metadata attributes for purposes other than security, such as to determine the network or application health, to track the long-lived sessions seen in the network, and so on.

Application Intelligence—Rules and Notes

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

- The Application Intelligence solution is supported on all the GigaSMART Modules of GigaVUE-HC Series devices.
- Whenever you perform a backup and restore operation, you must create a backup of both the device and GigaVUE-FM and then restore the backed-up data on both as well.

Create Application Intelligence Session

Application Monitoring gathers the application statistics, and sends this information to GigaVUE-FM, which acts as an application monitor. The application statistics appear as an array of monitoring reports that provide application-usage data in an easy-to-read graphical interface. This provides you with greater insight and control over how your network is being used and what applications are utilizing the most resources. To perform Application Monitoring, you must create the required application intelligence sessions for a device.

To create an Application Intelligence Session:

1. Go to **Physical** > **Physical Nodes**.
2. Click the Cluster ID of the node for which you want to create the Application Intelligence Session.
3. From the navigation pane, click **App Intelligence**.
4. Click **Create**. The **Create Application Intelligence Session** page appears.
5. In the Basic Info section, enter the name and description for the session.
6. In the **Configurations** section, complete the following:
    a. Select an export interval during which you want the Application Intelligence session to generate the reports for application monitoring. The valid range is 60–900.

b. Select a GigaSMART Group. You can also choose to create a new GigaSMART Group.

- Provide a name in the **Alias** field.
- Select a port or multiple ports from the **Port List**.

  Click **Save**.

  For more information about creating a new GigaSMART group, refer to the *Using GigaSMART Operations - Example*.

c. Select an IP interface that is used to create a dedicated channel to communicate all application statistics to GigaVUE-FM. You can also choose to create a new IP interface.

   If you are unable to view the required port in the **Port** field, perform these steps:

- Click **Port Editor.** Select the **Type** as **Tool** from the drop-down list for the required **Port Id**. Select **OK**.

  The selected Port appears in the list.
- Provide the **IP Address**, **IP Mask, Gateway,** and **MTU**.

  Click **Save**.

  For more information, refer to the *Create an IP Interface section*.

NOTE: You can view the **IP Interface** field only after you select a GigaSMART Group.

7. In the **Destination Settings,** enter the destination IP address. By default, the IP address of the GigaVUE-FM interface is displayed.
8. In the **Source Traffic** section, select a source port that require application monitoring in the **Source ports** field. Source port can be a single port, multiple ports, and port groups.
9. Configure the rules for filtering the required traffic in the **L2-L3 Rules** fields. To configure a rule:

   a. Click **Select Conditions**. Select the required parameters from the drop-down list.

   b. Select the value for the parameters from the drop-down.

   c. Select the required options:

   - Pass or Drop - Based on the parameter selected in the Conditions fields, the traffic that matches the conditions will either be passed or dropped.
   - Bidirectional - Allows the traffic in both directions of the flow.

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

10. Click **Save**. The session created is added in the list view.

The total applications participating in the network traffic are displayed in the Application Intelligence Dashboard. For more information about the dashboard, refer to the *View Application Intelligence Dashboard* on page 735.

If the session configuration is unsuccessful, troubleshoot the error notified (refer to *Health Status of a Solution* on page 737). Click the **Reapply all pending solutions** button in the dashboard to redeploy the configuration.

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

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

View Details of Application Intelligence Session

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

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

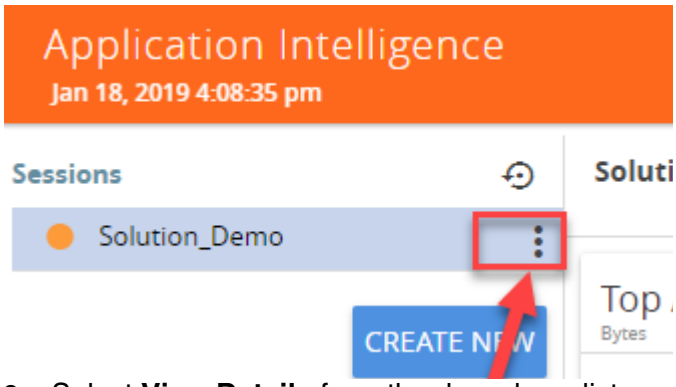

2. Select **View Details** from the drop-down list.

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

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

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

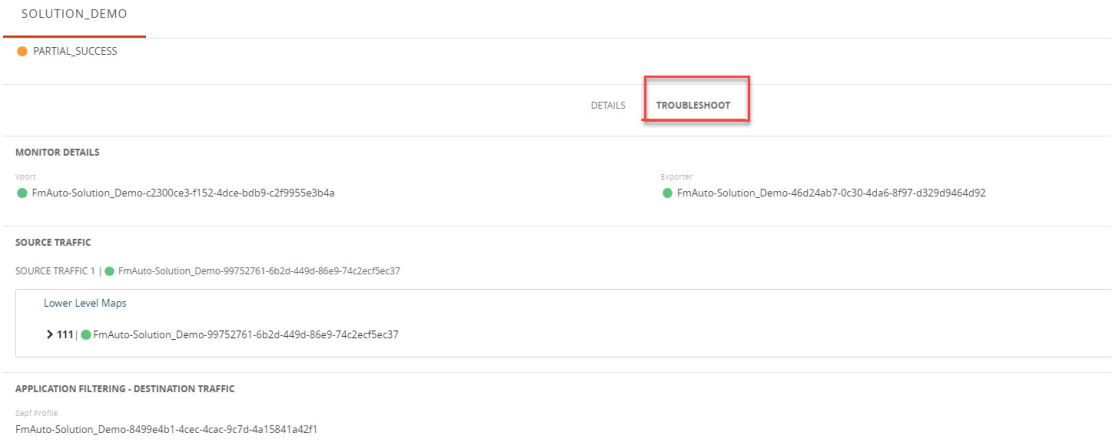

The details of the components that can be viewed for a solution are shown in the following table:

Components	Details of Components
Source Traffic	• Network Ports • First Level Maps • L2-L4 Rules
GigaSMART	• GigaSMART Port • GigaSMART Group • Virtual Port
Application Monitoring	• IP interface • Exporter
Application Filtering -Destination Traffic	• Second Level Map • Application and Advanced Rules • Tool Ports • GigaSMART Operation (GSOP)
Application Filtering	• Application Session Filtering

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

To learn more about the color indication and the health status of a solution refer to *Health Status of a Solution* on page 737.

Create Application Filtering Intelligence

Application Filtering Intelligence functionality on GigaSMART allows filtering of traffic based on the application (such as MySQL, ORACLE RAC, Sophos, or Facebook) or application family (such as antivirus, web, erp, or instant-messaging).

Application Filtering Intelligence supports filtering over 3000 applications.

You can create Application Filtering Intelligence in GigaVUE-FM by following either of the two ways:

- *Create Application Filtering Intelligence by Selecting Applications from Dashboard*
- *Create Application Filtering Intelligence by Editing Monitoring Session from Dashboard*

Create Application Filtering Intelligence by Selecting Applications from Dashboard

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

1. Go to **Physical** > **Physical Nodes**.
2. Click the Cluster ID to open a device.
3. Click **App Intelligence**.

NOTE: If you are creating Application Filtering Intelligence immediately after creating Application Monitoring, then you can proceed from Step 4.

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

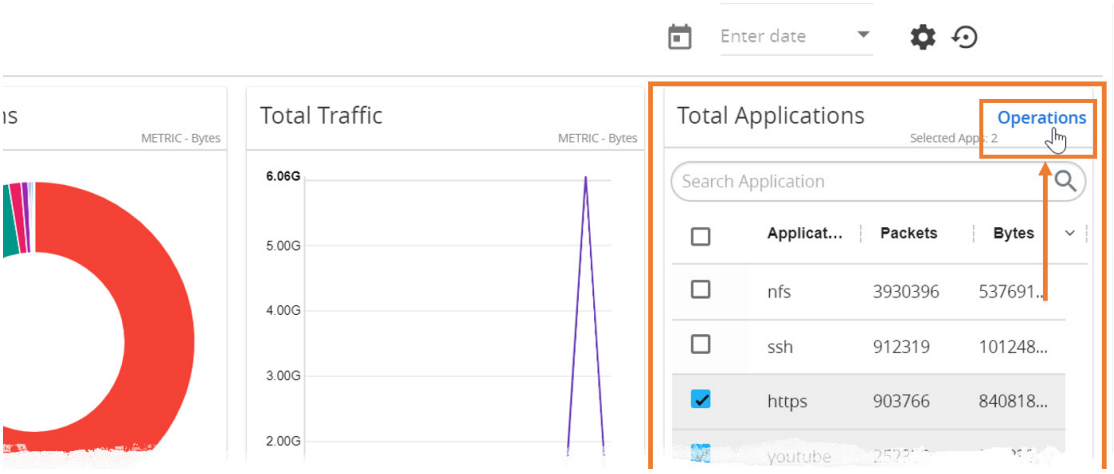

You can view the list of applications selected in the **Selected Applications** section.

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

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

In the **Priority** section, click **Advanced Rules > Add a rule** to add new rules to perform advanced filtering on the application.

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

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

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

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

- Bidirectional (Default option).
- Buffer/Buffer count - The option is enabled default.
- Protocol- The default value is TCP-UDP.
- Packet Count
- Timeout in seconds - The default value is 15secs.
- Sessions Field

10. Click **Save**.

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

Create Application Filtering Intelligence by Editing Monitoring Session from Dashboard

You can select the required Applications from the list of applications in the Application Editor by following these steps:

1. Go to **Physical** > **Physical Nodes**.
2. Click the Cluster ID to open a device.
3. From the navigation pane, click **App Intelligence**.
4. In the left pane of **Application Intelligence** dashboard, hover over the monitor sessions for which you need to apply the **Application Filtering**, and click the vertical ellipsis from the **Application Intelligence Solutions**.
5. Click **Edit** from the Actions drop-down list.
6. Click the **App Editor** to select the applications to be passed or dropped.

The Application Editor screen appears as shown:

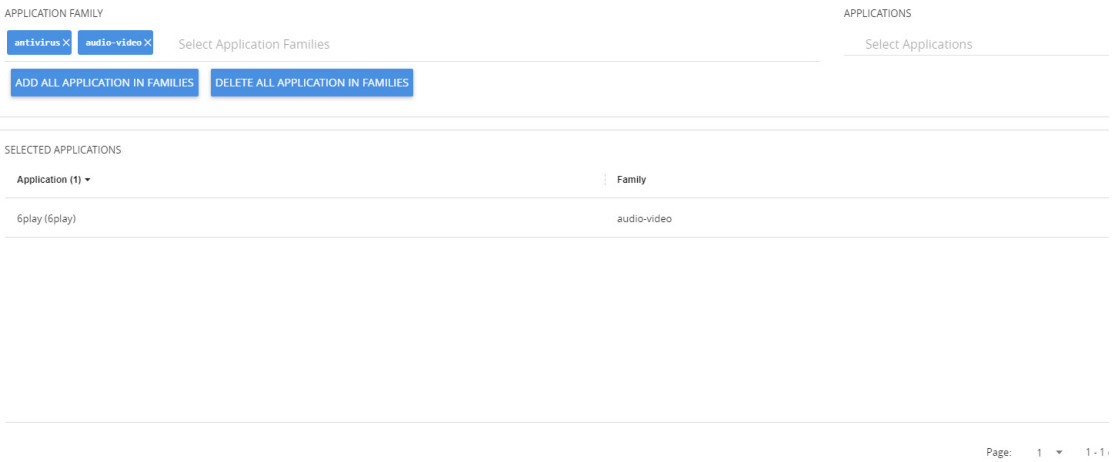

7. Click the **Application Family** field and select a Application Family such as antivirus, webmail that needs to be filtered from the traffic. You can also select multiple application families.

   If you choose to add or delete all the applications in a family, click **Add All Application in Families** or **Delete All Application in Families**.

8. Click the **Application** field and select an application or multiple applications that needs to be filtered from the traffic. You can also select multiple applications.

NOTE: GigaVUE-FM allows you to select the required protocols even without selecting the Protocol family. You must select the non-TLS version for the below TLS version protocol:

- SMTPS
- POP3S
- IMAPS
- FTPS
- LDAPS

The selected protocols and their families appear in the **Selected Applications** field.

9. Click **Add.** The selected applications appear in the **Application Filtering** tab.
10. Perform the steps 7 and 8 in the *Create Application Filtering Intelligence by Selecting Applications from Dashboard*.

After creating the Application Filtering Intelligence Session, you can view the **Application Filtering Traffic** statistics in the Application Intelligence Dashboard.

View Application Intelligence Dashboard

After creating the Application Intelligence Session, you can monitor the applications in the network by the reports displayed in the Dashboard as shown in Figure 28-2:

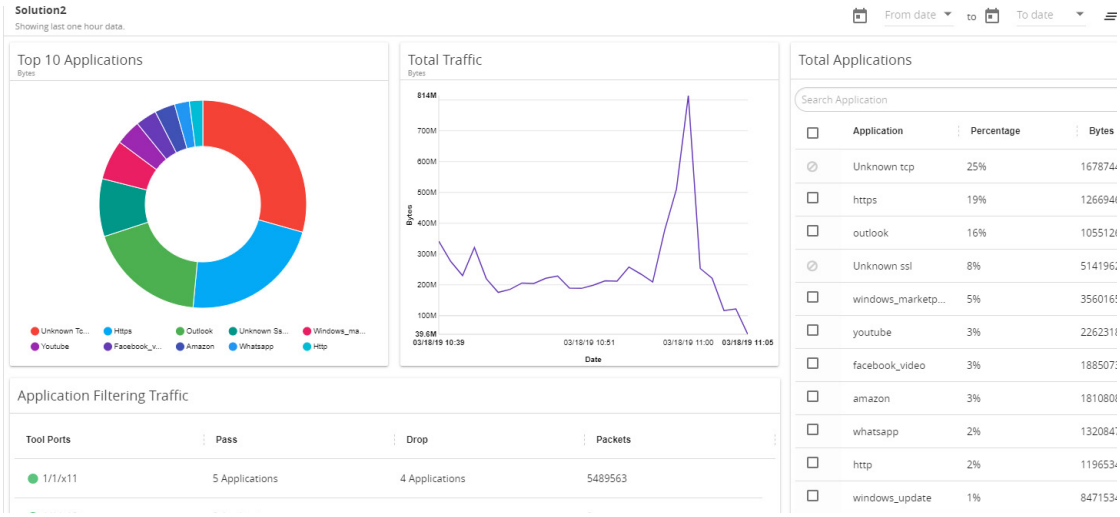

Figure 28-2: Application Intelligence Dashboard

Application Intelligence Dashboard displays the following metrics:

- **Top 10 Applications**: You can view a graphical representation of top 10 applications running in the network based on the metrics. When you hover over the Pie-chart, GigaVUE-FM shows the application name in the network. The legend for the Pie-chart appears at the bottom. When you select a pie, you can view the corresponding data highlighted in the Total Applications table.
- **Total Traffic**: You can view the total traffic of the network represented in the linear form of a graph.
- **Total Applications**: You can view the applications and their bandwidth in the network. You can also select the required application for filtering and exporting metadata by using the Operations field.
- **Application Filtering Traffic**: You can view the statistics of the applications that are filtered in the tool ports in the dashboard after creating an application filtering intelligence session for a device.

GigaVUE-FM enables you to view the above metrics for a particular period by selecting the dates from the dashboard. You can also choose to view the graphs in the dashboard for the metrics in bytes, packets or flows. To view the metrics in bytes, packets or flows, click the ⚙ button in the right corner of the Application Intelligence dashboard.

About De-duplication

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

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

NOTE: You must configure Application Filtering Intelligence for configuring De-duplication and also install licenses for De-duplication on your GigaSMART Modules for enabling the features in GigaVUE-FM.

Health Status of a Solution

The health of an Application Intelligence solution is determined by the health status of the following components, and the configuration status during deployment of the solution in a device:

- IP interface
- Source Port
- Destination Port
- GigaSMART Engine

You can view the health status of the solution by a color indication next to the name of the solution as shown in the following figure:

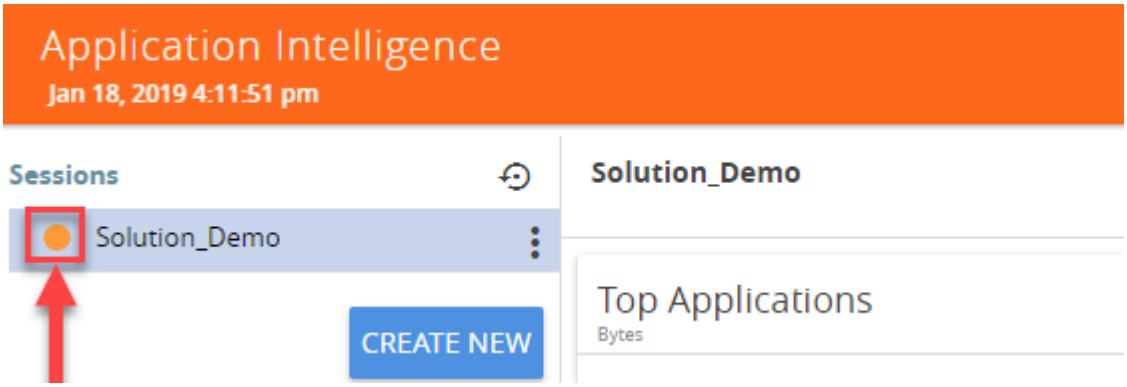

The health status of a solution is indicated by the following colors:

Color	Health Status of a Solution
Green	Healthy - All the components in a solution are functioning properly.
Red	Unhealthy - Any of the components in a solution is not functioning properly.
Amber	Partially Healthy

You can also view the reason for a failover when you hover your mouse over the color indicator next to the name of the solution. To avoid this scenario:

1. Navigate to Administration > System > Thresholds.
2. Set the threshold value of GigaSMART engine port packet drops to zero.

The following table provides the health state of the Application Intelligence solution corresponding to the health state of its associated components:

Table 28-1: Health status of Solution Vs Health status of Associated Components

Health Status of a Solution	GSGroup	IP Interface	Network Port and Tool Port	Metadata Exporter	Configuration Deployment Status
Red	Unhealthy	Healthy	Healthy	Healthy	Success
Amber	Partially Healthy	Healthy	Healthy	Healthy	Success
Red	Healthy	Unhealthy	Healthy	Healthy	Success
Amber	Healthy	Partially Healthy	Healthy	Healthy	Success
Amber	Healthy	Healthy	Some Maps are Unhealthy/ Partially Healthy	Healthy	Success
Red	Healthy	Healthy	All Maps are Unhealthy	Healthy	Success
Red	Healthy	Healthy	Healthy	All Metadata Exporters are Unhealthy	Success
Amber	Healthy	Healthy	Partially Unhealthy	Healthy	Success
Amber	Healthy	Healthy	Healthy	Some metadata exporters are Unhealthy	Success
Red	Healthy	Healthy	Healthy	Healthy	Failed
Amber	Healthy	Healthy	Healthy	Healthy	Partial Success
Red	Healthy	Healthy	Healthy	Healthy	Failed
Green	Healthy	Healthy	Healthy	Healthy	Success

Work with Application Intelligence Using GigaSMART

You can perform some expert actions and troubleshoot the existing configuration for Application Filtering Intelligence through the Application Identification page in GigaSMART.

Application Identification provides you greater visibility into the details of the ports and the maps associated with Application Filtering Intelligence configurations. This expert option helps you to configure Application Filtering Intelligence without configuring Application Monitoring.

NOTE: The Application Identification in GigaSMART interface is deprecated functionality and is focused more towards troubleshooting/debugging—it should only be used if you know what you are doing. The preferred and recommended way to use the Application Intelligence solution is through the Application Intelligence Dashboard described throughout this guide.

Part 6: GigaSMART

This section provides information about working with GigaSMART operations:

Refer to:

- *Work with GigaSMART Operations* on page 741 for rules, tips, and general guidance about working with GigaSMART Operations.
- *About GigaSMART Applications* on page 742 for devices that support GigaSMART.
- *Create GigaSMART Operations – A Summary* on page 756 to get started with GigaSMART.
- *How to Use GigaSMART Operations* on page 809 for detailed instructions about each GigaSMART operation.
- *GigaSMART Logs* on page 1203 to learn about GigaSMART application logs.

29 Work with GigaSMART Operations

This chapter describes how to use GigaSMART operations – advanced processing features available for use on GigaVUE-HB1 and GigaVUE-HC1 nodes, GigaVUE-HC2 nodes with front or rear GigaSMART modules, and GigaSMART-HC3 nodes with SMT-HC3-C05 modules (Figure 29-1).

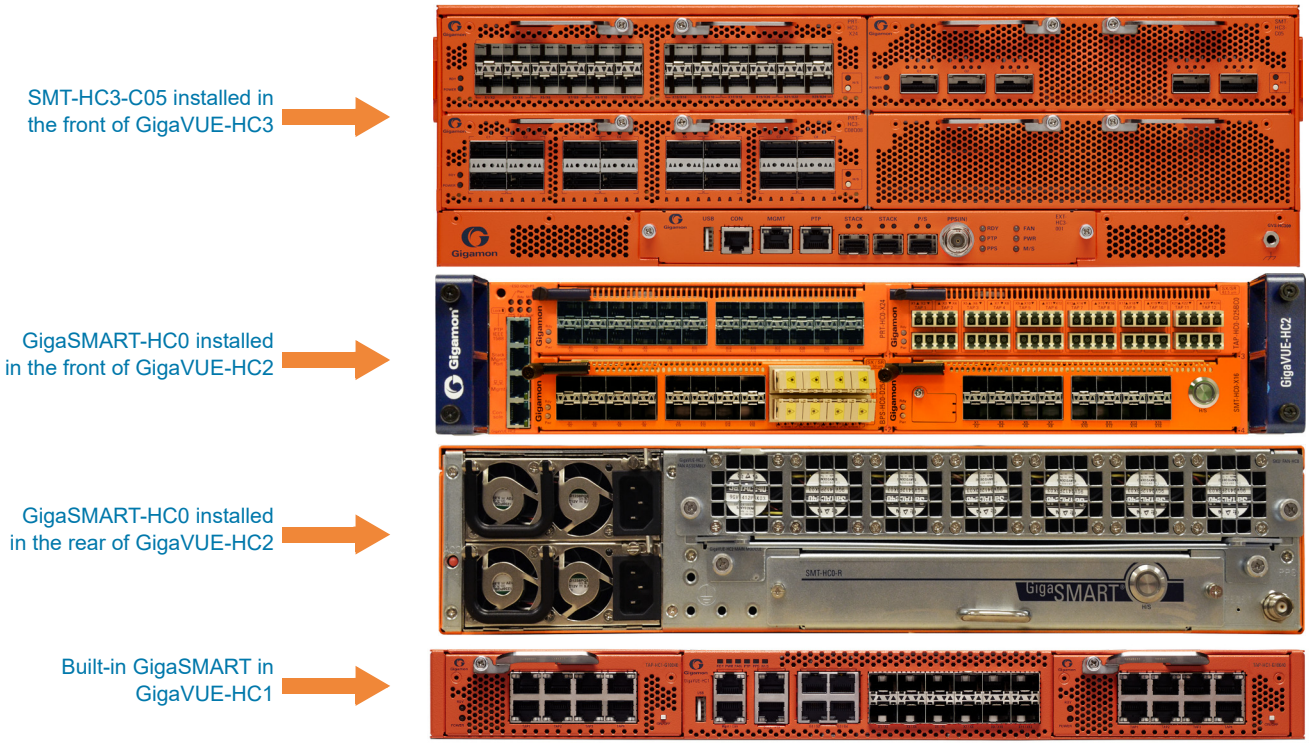

Figure 29-1: GigaSMART in GigaVUE Nodes

Refer to the following sections for details:

**Refer to the *Access GigaSMART from GigaVUE-FM* on page 754 for an overview of how to access GigaSMART from GigaVUE-FM.

***Refer to *How to Use GigaSMART Operations* on page 809 for comprehensive HowTo's on using all GigaSMART operations.

About GigaSMART Applications

GigaSMART applications are packet modification features available on the following GigaVUE H Series nodes:

- Standalone GigaVUE-HC3 with SMT-HC3-C05 module installed.
- Standalone GigaVUE-HC2 with front or rear GigaSMART-HC0 module installed.
- Standalone GigaVUE-HC1 nodes.
- Standalone GigaVUE-HB1 nodes.
- Any GigaVUE H Series node operating in a cluster with one of these node types.

NOTE: This section refers to any of these nodes as *GigaSMART-enabled* – they are all capable of using GigaSMART operations.

You can use both H-VUE and the CLI to create GigaSMART operations combining the GigaSMART applications, and then use them with other map rule criteria, and apply them in map rules on any network port in the node or cluster.

GigaSMART Perpetual Licenses

Table 29-1 lists GigaSMART perpetual licenses available on GigaVUE H Series nodes.

Table 29-1: GigaSMART Applications by License Type

Base GigaSMART Applications	
GigaVUE-HC0 Module – The base applications include Packet Slicing, Masking, Trailer, and IP and L2GRE Tunnel Decap. **GigaVUE-HC3 SMT-HC3-C05 Module** – The base applications include Packet Slicing, Masking, Trailer, and IP and L2GRE Tunnel Decap. **GigaVUE-HC1 Node** – The base applications include Packet Slicing, Masking, and Trailer. **GigaVUE-HB1 Node** – The base applications include Packet Slicing, Masking, and Trailer.	
GigaSMART Packet Slicing	Packet slicing lets you truncate packets after a specified header and offset or simply an offset, preserving the portion of the packet required for network analysis and adding a recalculated CRC to match the new packet length. GigaSMART packet slicing can parse variable-header packets, starting slicing after a named header or tunnel type (VLAN, MPLS, GTP, and so on). This way, you can slice packets without having to rely on a fixed offset. Because they are smaller, sliced packets are analyzed more efficiently and require less disk space to store. Your tools can process fewer bits and have more room to store the vital portions of each packet, enhancing storage and analysis performance.
GigaSMART Masking	Masking lets you overwrite specific packet fields with a specified pattern so that sensitive information is protected during network analysis. Privacy compliance is crucial for financial, insurance, and healthcare enterprises. GigaSMART masking lets you give network engineers the data they need while still keeping private data private. As with the slicing feature, GigaSMART masking can automatically compensate for variable length headers, allowing you to specify a mask target in terms of a particular packet header.
GigaSMART Trailers	When trailers are enabled in a GigaSMART operation, each packet is tagged with a trailer field containing metadata about the packet and how it was processed. You can configure the trailer to include the original packet's CRC as well as a Source ID field identifying the port on the GigaVUE H Series node where the packet entered the system. The GigaVUE node type, box ID, and port ID are all included in the Source ID field, making it easy to identify the source of each packet entering the Visibility Platform.

Table 29-1: GigaSMART Applications by License Type

GigaSMART IP Encapsulation/ Decapsulation (GigaSMART Tunnel)	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. IP fragmentation and reassembly is supported. Refer to *GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation* on page 828. **GigaSMART IP encapsulation/decapsulation is licensed differently for the GigaVUE line cards, modules, and nodes as follows:** • GigaVUE-HC2 and GigaVUE-HC3 – GigaSMART IP decapsulation is included with base license. • GigaVUE-HC2 and GigaVUE-HC3 – GigaSMART IP encapsulation requires Advanced Tunneling license. • GigaVUE-HC1 Node – GigaSMART IP encapsulation/decapsulation requires Tunneling license. • GigaVUE-HB1 Node – GigaSMART IP encapsulation/decapsulation requires Tunneling license. **NOTE:** GigaSMART operations with a tunnel component can only be assigned to GigaSMART groups consisting of a single GigaSMART engine port.
GigaSMART Layer 2 GRE Tunnel Encapsulation/ Decapsulation	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. Then use GigaSMART L2GRE tunnel decapsulation at the second GigaSMART node to decapsulate the traffic before sending it to local tool ports. IP fragmentation and reassembly is supported. Refer to *IP Fragmentation and Reassembly on L2GRE and GMIP Tunnels* on page 843. **GigaSMART L2GRE encapsulation/decapsulation is licensed differently for the GigaVUE line cards, modules, and nodes as follows:** • GigaVUE-HC2 and GigaVUE-HC3 – GigaSMART L2GRE decapsulation is included with base license. • GigaVUE-HC2 and GigaVUE-HC3 – GigaSMART L2GRE encapsulation requires Advanced Tunneling license. • GigaVUE-HC1 Node – GigaSMART L2GRE encapsulation/decapsulation requires Tunneling license. • GigaVUE-HB1 Node – GigaSMART L2GRE encapsulation/decapsulation requires Tunneling license. **NOTE:** GigaSMART operations with a tunnel component can only be assigned to GigaSMART groups consisting of a single GigaSMART engine port.

Table 29-1: GigaSMART Applications by License Type

GigaSMART VxLAN Tunnel Decapsulation	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 solution or a customer-specific device, will be received on an IP interface associated with a network port, then the packets 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. **GigaSMART VxLAN tunnel decapsulation is licensed differently for the GigaVUE line cards, modules, and nodes as follows:** • GigaVUE-HC2 and GigaVUE-HC3 – GigaSMART VxLAN tunnel decapsulation is included with base license. • GigaVUE-HC1 Node – GigaSMART VxLAN tunnel decapsulation requires Tunneling license. • GigaVUE-HB1 Node – GigaSMART VxLAN tunnel decapsulation requires Tunneling license. **NOTE:** GigaSMART operations with a tunnel component can only be assigned to GigaSMART groups consisting of a single GigaSMART engine port.
GigaSMART Custom Tunnel Decapsulation	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 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. GigaSMART leverages the existing generic header stripping operation to remove the tunnel header. The reassembly of fragmented IP packets is also supported. **GigaSMART custom tunnel decapsulation is licensed differently for the GigaVUE line cards, modules, and nodes as follows:** • GigaVUE-HC2 and GigaVUE-HC3 – GigaSMART custom tunnel decapsulation is included with base license. • GigaVUE-HC1 Node – GigaSMART custom tunnel decapsulation requires Tunneling license. • GigaVUE-HB1 Node – GigaSMART custom tunnel decapsulation requires Tunneling license. **NOTE:** GigaSMART operations with a tunnel component can only be assigned to GigaSMART groups consisting of a single GigaSMART engine port.
Advanced Tunneling License/Tunneling License The **Advanced Tunneling License/Tunneling License** enables the following GigaSMART applications:	
GigaSMART IP Encapsulation/ Decapsulation (GigaSMART Tunnel)	As described above, GigaSMART IP encapsulation requires Advanced Tunneling license on GigaVUE-HC2, and GigaVUE-HC3. On GigaVUE-HB1 and GigaVUE-HC1, GigaSMART IP encapsulation/decapsulation requires Tunneling license.
GigaSMART Layer 2 GRE Tunnel Encapsulation/ Decapsulation	As described above, GigaSMART L2GRE encapsulation requires Advanced Tunneling license on GigaVUE-HC2, and GigaVUE-HC3. On GigaVUE-HB1 and GigaVUE-HC1, GigaSMART L2GRE encapsulation/decapsulation requires Tunneling license.
GigaSMART VxLAN Tunnel Decapsulation	As described above, GigaSMART VxLAN tunnel decapsulation requires Tunneling license on GigaVUE-HB1 and GigaVUE-HC1.
GigaSMART Custom Tunnel Decapsulation	As described above, GigaSMART Custom tunnel decapsulation requires Tunneling license on GigaVUE-HB1 and GigaVUE-HC1.

Table 29-1: GigaSMART Applications by License Type

GigaSMART ERSPAN Tunnel Decapsulation	Some Cisco equipment provides the ability to mirror specific traffic to a remote destination through an ERSPAN tunnel. A GigaSMART-enabled GigaVUE H Series node with the Advanced Tunneling license installed can act as the receiving end of an ERSPAN tunnel, providing GigaVUE packet distribution for packets sent from remote Cisco equipment. Both ERSPAN Type II and Type III header decapsulation are supported. **GigaSMART ERSPAN decapsulation is licensed differently for the GigaVUE line cards, modules, and nodes as follows:** • GigaVUE-HC2, and GigaVUE-HC3 – GigaSMART ERSPAN decapsulation requires Advanced Tunneling license. • GigaVUE-HB1 and GigaVUE-HC1 nodes – GigaSMART ERSPAN decapsulation requires Tunneling license. **NOTE:** GigaSMART operations with a tunnel component can only be assigned to GigaSMART groups consisting of a single GigaSMART engine port.
De-Duplication License The **De-Duplication License** enables the following GigaSMART applications:	
GigaSMART De-Duplication	Packet de-duplication lets you remove duplicate IPv4 and IPv6 packets (and also non-IP packets) before they are unnecessarily processed or stored by tools. Duplicate packets can occur when 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.
Header Stripping License The **Header Stripping License** enables the following applications:	
GigaSMART Header Addition	Use this feature to add VLAN tags to packets. VLAN tag addition is handy when differentiating stripped packets from non-stripped packets on common IP ranges (for example,10.x.x.x; 192.168.x.x) or when 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.
GigaSMART Header Stripping	Stripping operations let you remove headers from tagged packets or headers and trailers from tunneled (encapsulated) packets: • **Header Stripping** – Remove headers from MPLS, MPLS+VLAN, VLAN, VN-Tag, Cisco FabricPath Header, GRE, or VXLAN tagged packets before they are sent to tool ports. Remove headers from any protocol by specifying the offset of the fragment and the length of the packet. 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. • **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.

Table 29-1: GigaSMART Applications by License Type

Adaptive Packet Filtering (APF) License	
The **Adaptive Packet Filtering License** enables the following GigaSMART applications:	
GigaSMART Adaptive Packet Filtering (APF)	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. 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. APF also provides a powerful filtering engine that identifies content (based on patterns) across any part of the data packet, including the data packet payload. APF filters packet-by-packet, but does not have the concept of sessions. For application session filtering (ASF) and packet buffering on application session filtering, refer to *GigaSMART Application Session Filtering (ASF) and Buffer ASF*.
Application Session Filtering (ASF) License	
The **Application Session Filtering License** enables the following GigaSMART applications:	
GigaSMART Application Session Filtering (ASF) and Buffer ASF	Application Session Filtering (ASF) provides additional filtering on top of Adaptive Packet Filtering (APF). With APF, you can filter on any data patterns within a packet. With ASF, you apply the pattern matching and then send all the packet flows associated with the matched packet to one or more monitoring tools. Use ASF to create a flow session and send the packets associated with the flow session to one or more tools. A flow session consists of one or more fields that you select to define the session. Either the packets for the whole session can be captured or only the packets following a pattern match. ASF captures packets of a session after an APF rule match. When the APF match occurs in the middle of a session, packets in the session prior to the match are not captured. With some tools needing all the packets of a flow session in order to perform data analysis, GigaSMART uses buffering to ensure that all packets belonging to a flow session are captured and forwarded to the tools. This is referred to as Application Session Filtering with buffering, or buffer ASF. **NOTE:** ASF and buffer ASF also require the Adaptive Packet Filtering (APF) license. **NOTE:** Stateful load balancing for the ASF application is included with the Application Session Filtering (ASF) license.

Table 29-1: GigaSMART Applications by License Type

GTP Filtering & Correlation License The **GTP Filtering & Correlation License** enables the following GigaSMART applications:	
GigaSMART GTP Correlation	The GPRS Tunneling Protocol (GTP) carries mobile data across service provider networks. GTP includes both the control plane (GTP-c) and a user-data plane (GTP-u) network traffic. Visibility into a subscriber's session requires the ability to understand the stateful nature of GTP (v1 and/or v2). To gain an accurate view into the subscriber's session, GTP tunnels are used to correlate subscriber-specific control plane and user-data plane traffic. With GTP correlation, you can gain access to the subscriber's data in these GTP tunnels by reliably correlating and passing all of the identified subscriber's control and data plane traffic to the analytics/monitoring tools and billing subsystems. Using GTP correlation, you can filter, replicate, and forward specific subscriber sessions to specific tools by correlating the subscriber IDs that are exchanged as part of the control sessions to the corresponding tunnel IDs (TEID) that are part of the user-data plane traffic. GTP correlation supports a maximum of 6 million GTP subscriber sessions for GigaVUE-HC2 nodes, whereas, it supports 12 million GTP subscriber sessions for GigaVUE-HC3 nodes. **NOTE:** Tiered License model for 250k/500k/Max Subscribers applies only to GigaVUE-HC2. GigaVUE-HC3 has Max license only. **NOTE:** Stateful load balancing for the GTP application is included with the GTP Filtering & Correlation license. **NOTE:** The Adaptive Packet Filtering (APF) license is included with GTP for filtering inside GTP headers.
GigaSMART GTP Whitelisting and GTP Flow Sampling	Starting in software version 4.3, use GTP whitelisting and GTP flow sampling to provide subsets of GTP correlated flows to tools. GTP whitelisting selects specific subscribers based on IMSI, while GTP flow sampling uses map rules to select subscribers. GTP whitelist-based sampling and GTP flow sampling (rule-based flow sampling) are performed prior to GTP filtering. **NOTE:** In addition to the GTP Filtering & Correlation License, GTP whitelisting and GTP flow sampling also require the FlowVUE license.
GTP Scaling	Starting in software version 4.5, GTP can be scaled as follows: • A GigaSMART group (gsgroup) associated with GTP applications can have multiple GigaSMART engine port members (**e** ports), up to four, forming a GTP engine group. Refer to *GTP Engine Grouping* on page 959.
GTP Stateful Session Recovery	Starting in software version 4.4, use GTP stateful session recovery to back up GTP sessions periodically so they can then be recovered faster after a GigaSMART line card or module reboot or a node reboot. GTP stateful session recovery provides session persistence for GigaSMART GTP applications, including GTP flow filtering, GTP whitelisting, and GTP flow sampling.

Table 29-1: GigaSMART Applications by License Type

SIP/RTP Correlation License	
The **SIP/RTP Correlation License** enables the following GigaSMART application:	
GigaSMART SIP/RTP Correlation	Session Initiation Protocol (SIP) is the dominant method to initiate, maintain, modify, and terminate voice calls in service provider and enterprise networks. Real-time Transport Protocol (RTP) is used to manage the real-time transmission of voice payload across the same networks. Visibility into a subscriber's voice traffic requires the ability to understand the subscriber attributes and stateful information contained within SIP to correlate subscriber-specific RTP traffic so that monitoring tools can achieve an accurate view of the subscriber's traffic on the network. The GigaSMART SIP/RTP correlation application correlates the subscriber-specific attributes and the endpoint identifiers of the RTP streams where the session is carried, as well as other SIP-related attributes that are exchanged as part of the control sessions. Use SIP/RTP correlation to leverage a subscriber-aware monitoring policy on Gigamon's Visibility platform and to optimize current tool infrastructure investments by providing only relevant data to tools while increasing visibility into subscriber traffic. This helps improve QoE and performance. Carriers gain access to the subscriber's traffic by reliably correlating and passing all the identified subscriber's control and data sessions to the analytics/monitoring probes and/or billing subsystems for an accurate view of the subscriber's sessions. **NOTE:** The FlowVUE license is needed for session-aware load balancing for RTP.
FlowVUE License	
The **FlowVUE License** enables the following GigaSMART applications:	
GigaSMART FlowVUE	FlowVUE allows for the active sampling of a subscriber's device (also known as a "User Endpoint IP" or UE IP) across IP networks or GTP-u tunnels. The integrity of subscriber flows is preserved by forwarding all flows associated with the sampled UE IP to all probes and analysis tools. Intelligent sampling executed by FlowVUE is deployed for understanding usage patterns. Operators can also gain visibility in to the subscribers QoE by forwarding all GTP control sessions to the monitoring tools. By combining FlowVUE with other GigaSMART functions such as APF, network traffic can be further reduced by filtering on specific Layer-4 application ports that the operator is interested in monitoring. Overall, this helps service providers address rising tool costs by enabling them to preserve or increase tool utility and offset ARPU reduction by monetizing Big Data with tools seen in Customer Experience Management offerings. **NOTE:** The FlowVUE license also enables GTP whitelisting and GTP flow sampling.

Table 29-1: GigaSMART Applications by License Type

NetFlow Generation License	
The **NetFlow Generation License** enables the following GigaSMART applications:	
GigaSMART NetFlow Generation	NetFlow Generation is a simple and effective way to increase visibility into traffic types and usage patterns across systems. Data can be used to build relationships and usage patterns between nodes on the network (traditionally, routers and switches collected IP traffic statistics and exported them as NetFlow Generation Records). The advanced capabilities of GigaSMART® technology can be leveraged to summarize and generate unsampled NetFlow Generation statistics from incoming traffic streams. Offloading NetFlow Generation to an out-of-band solution like the Gigamon Visibility Platform completely eliminates the risk of expending expensive production network resources in generating these analytics. Combined with the flexibility offered by Gigamon's patented Flow Mapping® technology, operators can pick and choose from the incoming flows to generate NetFlow Generation statistics, without losing critical information. NetFlow Generation supports NetFlow version 5, 9, IPFIX. The port used to export NetFlow records is configured as an IP interface. When NetFlow collects SSL metadata, it makes use of the GigaSMART SSL application, however, only the NetFlow Generation license is needed for NetFlow to collect SSL metadata. NetFlow supports second level maps that are used for configuring filtering rules enabled through Adaptive Packet Filtering (APF). After the APF rules are applied, second level maps send traffic to NetFlow and then to IP interface with tool ports. **NOTE:** NetFlow with second level maps requires the Adaptive Packet Filtering license. **NOTE:** GigaSMART operations with a NetFlow component can be assigned to multiple GigaSMART groups or GigaSMART groups consisting of multiple GigaSMART engine ports.
SSL Decryption Licenses	
The **SSL Decryption Licenses** enable the following GigaSMART applications:	
GigaSMART Out-of-Band SSL Decryption	Secure Sockets Layer (SSL) Decryption is a cryptographic protocol that adds security to 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. Out-of-band SSL decryption delivers decrypted traffic to out-of-band tools that can then detect threats entering the network. On GigaVUE H Series nodes, GigaSMART line cards or modules perform the decryption of SSL traffic. Using GigaSMART for decryption provides a single decryption point, enabling decrypted data to be sent to tools for inspection. Using GigaSMART removes the decryption function from tools and offers improved performance. Before 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. Use out-of-band SSL decryption on the GigaSMART line card or module with passive or offline traffic. Tap the traffic to and from a server and pass it to the GigaVUE H Series node with the GigaSMART line card or module. For secure storage of private keys, Thales Hardware Security Module (HSM) is integrated with out-of-band SSL decryption. Refer to *Thales HSM for SSL Decryption for Out-of-Band Tools* on page 1181.

Table 29-1: GigaSMART Applications by License Type

GigaSMART SSL Decryption for Inline and Out-of-Band Tools	SSL decryption for inline tools provides visibility into encrypted traffic. Inline SSL decryption delivers decrypted packets to tools that can be placed inline or out-of-band. The tools look into decrypted packets for threats, such as viruses or other malware. The amount of Internet traffic that is encrypted is increasing, and much of it is encrypted with Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols. Malware increasingly uses encrypted SSL traffic, thus a significant percentage of attacks hide in SSL. Inline SSL decryption delivers a complete view of encrypted applications and hidden threats in your organization. Many applications, such as email, also use SSL. Encryption protects data from being viewed in transit over the Internet such as in an exchange of emails. Encryption also keeps the data private. But when data is encrypted, packets are not inspected, which can create blind spots in your network. Providing visibility into encrypted traffic eliminates this blind spot. SSL/TLS blind spots in your network can be eliminated across any port or application, for example, port 443, or email, Web, or VoIP applications. **NOTE:** Inline SSL decryption is supported on GigaVUE-HC1, GigaVUE-HC2, and GigaVUE-HC3. Refer to the "*Inline SSL Decryption"* chapter in the *GigaVUE-FM User's Guide*.

NOTES:

- GigaSMART load balancing does not require a separate license. Stateless load balancing is included with base licenses. Stateful load balancing for GTP and ASF are included with the GTP Filtering & Correlation and Application Session Filtering (ASF) licenses. Stateful load balancing for tunnel is included with the tunneling licenses. Refer to *GigaSMART Load Balancing* on page 1147.
- GigaSMART MPLS traffic performance enhancement does not require a separate license. Refer to *GigaSMART MPLS Traffic Performance Enhancement* on page 1166.

Licensing GigaSMART Applications

GigaSMART applications are enabled using license keys. Table 29-1 on page 743 summarizes the GigaSMART applications included with the base license, as well as those included with each of the separately available licenses.

Contact your Sales Representative for information on obtaining a license key to enable additional GigaSMART applications. Refer to the **license** command in the GigaVUE-OS CLI Configuration Guide for details.

For perpetual licenses, the Expiration Date column has the word Never to indicate that there is no expiration date. For evaluation licenses, the Expiration Date column has a specific date on which the license expires. For more information on evaluation licenses, refer to *GigaSMART Evaluation Licenses* on page 752.

GigaSMART Evaluation Licenses

Use an evaluation license to evaluate GigaSMART applications. During the evaluation period of 45 days, you will have access to the full functionality of the GigaSMART applications under evaluation. You can obtain an evaluation license for any GigaSMART application, for either a single or for a number of GigaSMART applications combined in a bundle.

To obtain an evaluation license, contact your Sales Representative. A license key will be generated by Gigamon and sent to you. You then install the license, which enables the GigaSMART application for evaluation purposes.

Install Evaluation Licenses

You install an evaluation license the same way you install a perpetual license, using the **license** command.

The key consists of a long string beginning with LK2, which is a protocol, followed by the line card or module (SMT_HC0_R), followed by the content of the license key.

Notify Evaluation License Expiry

After installation, the evaluation license will expire after 45 days, on a specific date.

To notify you as the evaluation license approaches the expiry date, you can enable a notification. When enabled, the notification will be sent when there are 30, 15, 10, 5, 4, 3, 2, and 1 days remaining before the license expires.

Use the following CLI command to enable the evaluation license reminder:

```
(config) # snmp-server enable notify evallicensereminder
```

You can also use the **show license** command to display the expiration date of an evaluation license.

How to Combine Evaluation and Perpetual Licenses

An evaluation license can be for a number of GigaSMART applications combined in a bundle. If you have a perpetual license for one GigaSMART application, for example, de-duplication, and you want to evaluate a bundle that contains 10 GigaSMART applications, including de-duplication, the 45-day evaluation period will apply to the other 9 GigaSMART applications, while the perpetual license will apply to de-duplication.

If you obtain a perpetual license after an evaluation license, the perpetual license will overwrite the evaluation license.

GigaSMART Application after Expiry

Once an evaluation license expires, access to the GigaSMART application is disabled. If maps were configured using GigaSMART applications on the evaluation license, traffic will be dropped when the evaluation license expires.

NOTE: Traffic will flow through maps with perpetually licensed GigaSMART applications.

In addition, the **gsop** command will not be available once the evaluation license has expired.

However, if a new evaluation license for the same GigaSMART application is installed, a new 45-day evaluation period will begin.

Move Evaluation and Perpetual Licenses

Evaluation and perpetual license keys are saved on the GigaSMART line card or module, while license information is stored in the configuration database. The license key on the line card or module has to match the license information stored in the database, otherwise a license mismatch will result.

Line cards or modules may sometimes need to be moved or swapped. For the procedure to move a license, refer to *Move Licensed GigaSMART Line Card to a New Slot* on page 753. This procedure will clear a license mismatch under certain circumstances. Moving a license depends on the license type, as well as the expiry date, as follows:

License Key Saved on GigaSMART Line Card/Module	License Information Stored on Configuration Database	Can be Moved?
Perpetual License	Evaluation License	Yes
Evaluation License	Perpetual License	No
Evaluation License with an earlier expiry date than the one stored on the configuration database	Evaluation License	No
Evaluation License with a later expiry date than the one stored on the configuration database	Evaluation License	Yes

Move Licensed GigaSMART Line Card to a New Slot

On the GigaVUE HD Series, you can move a GigaSMART line card from one slot to another. On the GigaVUE-HC2 or GigaVUE-HC3, you can move a GigaSMART front module from one bay to another. However on the GigaVUE-HC2, you cannot move the GigaSMART rear module from the rear to the front.

If there are no GigaSMART operations (gsops) configured on the line card or module to be moved, you can move the line card or module to the new slot or bay.

If there are GigaSMART operations (gsop), GigaSMART groups (gsgroup), and maps configured on the line card or module to be moved, the system will report a license mismatch if you try to move it without first removing the related configuration.

To clear the settings related to the GigaSMART line card or module from its previous slot, allowing you to create new GigaSMART operations, GigaSMART groups, and maps using the GigaSMART line card or module in its new slot, use the following procedure:

1. Issue the following CLI command:
   (config) # show running-config
2. Copy and paste the output to a file such as Notepad, for reference.
3. Remove the map that uses the gsop defined on the gsgroup of the GigaSMART line card or module to be moved, using the following CLI command:
   (config) # no map alias <alias>
4. Remove the gsop that was defined on the gsgroup of the GigaSMART line card or module to be moved, using the following CLI command:
   (config) # no gsop alias <alias>
5. Remove the gsgroup that was defined on GigaSMART line card or module to be moved, using the following CLI command:
   (config) # no gsgroup alias <alias>
6. Issue the following CLI command on GigaSMART line card or module to be moved:
   (config) # no card slot <slot ID>
7. Assuming that the new slot does not have a GigaSMART line card or module inserted, issue the following CLI command on the new slot:
   (config) # no card slot <slot ID>
8. Issue the following CLI command on the new slot:
   (config) # card slot <slot ID>

On the new slot, configure gsgroup, gsop, and reapply the map that uses the gsop on the GigaSMART line card or module.

Access GigaSMART from GigaVUE-FM

You can access GigaSMART operations from within GigaVUE-FM, by accessing a device that has been added to FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation pane of the device view on supported devices.

To access GigaSMART from the GigaVUE-FM interface:

1. Select **Physical** from the top navigation menu.
2. Select **Physical Nodes** from the side panel. This displays the list of Devices/ Cluster Nodes managed by this instance of FM.
3. Click the Cluster ID of any node to open the node.

   Once you are in the node, this part of the interface should behave just like HVUE for devices that support GigaSMART. Refer to *About GigaSMART Applications* on page 742.
4. Click **GigaSMART** from the side navigation pane.

Refer to:

- *About GigaSMART Applications* on page 742 for devices that support GigaSMART.
- *Create GigaSMART Operations – A Summary* on page 756 to get started with GigaSMART.

- *How to Use GigaSMART Operations* on page 809 to learn how to use GigaSMART operations.

Create GigaSMART Operations – A Summary

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to *Access GigaSMART from GigaVUE-FM* on page 754 for details.

GigaSMART operations require the following steps:

1. **Create GigaSMART groups on the GigaSMART Groups page (select GigaSMART > GigaSMART Groups > GigaSMART Groups)**

   Use GigaSMART groups to manage and budget GigaSMART processing power. Use the ***New*** button on the GigaSMART Group page to create groups of the available GigaSMART ports in a given chassis.

   NOTE: The GigaSMART engine ports in a GigaSMART group can be on different line cards or modules in the same GigaVUE-HD or GigaVUE-HC2 chassis. However, all GigaSMART engine ports must be on the same chassis.

   The number of GigaSMART engine ports are as follows:

   - The GigaSMART-HC0 module includes one GigaSMART engine port.
   - The GigaVUE-HB1 node includes one GigaSMART engine port.

   The number of GigaSMART engine ports available in a chassis depends on the number of GigaSMART line cards in the chassis—up to five in the GigaVUE-HC2 (four front GigaSMART modules with one GigaSMART engine port each, and one rear GigaSMART module with one GigaSMART engine).

   The processing power of the GigaSMART engine ports is as follows:

   - Each GigaSMART port on the GigaSMART-HC0 module can process packets at **40Gb**.
   - GigaSMART port on the GigaVUE-HB1 node can process packets at **10Gb**.

   GigaSMART engine ports are numbered with an **e** prefix using **<bid/sid/e1..e2>** nomenclature – **1/1/e1**, for example.

   NOTE: The ports in a GigaSMART group can be on different line cards or modules in the same GigaVUE-HC2 chassis. However, they must all be on the same chassis.

   NOTE: The slot ID for a GigaVUE-HB1 chassis is fixed at **1**.

   NOTE: The bay ID for a GigaVUE-HC2 with a rear GigaSMART module is fixed at **5**. The bay ID for a GigaVUE-HC2 with GigaSMART front modules, will be 1 to 4, depending on where the module or modules are installed.

   Each GigaSMART operation you create in the next step must be assigned to one of the GigaSMART groups you create in this step.

2. **Create GigaSMART operations using the GigaSMART *Operations page (select* GigaSMART > GigaSMART *Operations (GSOP))***

   Give your GigaSMART operation a name, include a valid combination of GigaSMART operations, and assign it to one of the GigaSMART groups created in the previous step.

Refer to *How to Combine GigaSMART Operations* on page 778 for details on supported combinations of GigaSMART operations.

You can also configure how (or, in some cases, whether) a GigaSMART operation attaches a trailer that indicates where a packet arrived in the Gigamon Visibility Platform and how it was modified. This trailer can be interpreted using a recent version of the Wireshark® Protocol Analyzer. Refer to *GigaSMART Trailer Reference* on page 1196 for details on the GigaSMART Trailer.

3. **Apply GigaSMART operations to network ports in maps**

   The New Map and Edit Map pages contain a **GigaSMART Operations (GSOP)** field that lets you select a GigaSMART operation to be used in a map. Refer to *Manage Maps* on page 518 for details. Keep in mind, however, that GigaSMART operations *must* be selected before destination tool ports or collector destinations.

Groups of GigaSMART Engine Ports

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to *Access GigaSMART from GigaVUE-FM* on page 754 for details.

Use the GigaSMART Group page to create groups of GigaSMART engine ports in a given GigaVUE chassis. (To go to the GigaSMART Group page, select **GigaSMART > GigaSMART Groups > GigaSMART Groups > New**.)

The GigaSMART engine ports in a GigaSMART group can be on different line cards or modules in the same GigaVUE HD Series, GigaVUE-HC2, or GigaVUE-HC3 chassis. However, all GigaSMART engine ports must be on the same chassis.

Use groups of GigaSMART engine ports to increase the processing power of GigaSMART. The processing power of the GigaSMART engine ports is as follows:

- Each of the two GigaSMART engine ports in a GigaSMART-HD0 line card on GigaVUE HD Series can process packets at up to **40Gb**.
- Each of the two GigaSMART engine ports in an SMT-HC3-C05 module on GigaVUE-HC3 can process packets at up to **100Gb**.
- The GigaSMART engine port in a GigaSMART-HC0 module on GigaVUE-HC2 can process packets at up to **40Gb**.
- The GigaSMART engine port in the GigaVUE-HC1 node can process packets at up to **20Gb**.
- The GigaSMART engine port in the GigaVUE-HB1 node can process packets at up to **10Gb**.

The number of GigaSMART engine ports are as follows:

- Each GigaSMART-HD0 line card includes two GigaSMART engine ports.
- Each GigaSMART-HC0 module includes one GigaSMART engine port.
- Each SMT-HC3-C05 module on GigaVUE-HC3 includes two GigaSMART engine ports.
- The GigaVUE-HC1 node includes one GigaSMART engine port.

- The GigaVUE-HB1 node includes one GigaSMART engine port.

The number of GigaSMART engine ports available in a chassis depends on the number of GigaSMART line cards or modules in the chassis as follows:

- Up to five modules in the GigaVUE-HC2 (four front GigaSMART modules and one rear GigaSMART module with one GigaSMART engine each, for a total of five).
- Up to four modules in the GigaVUE-HC3 (eight GigaSMART engine ports).

The following table provides a summary:

GigaVUE Node	Maximum GigaSMART Line Cards or Modules per Node	Number of GigaSMART engine ports per Line Card or Module	Maximum Number of GigaSMART engine ports in a GigaSMART group (gsgroup)
GigaVUE-HC3	4	2	8
GigaVUE-HC2	4 front, 1 rear	1	5
GigaVUE-HC1	1	1	1
GigaVUE-HB1	1	1	1

Engine grouping for GTP is a special case, which is described in *GTP Engine Grouping* on page 959.

GigaSMART engine ports are numbered with an **e** prefix using **<bid/sid/e1..e2>** nomenclature, such as 1/1/e1. For example:

s

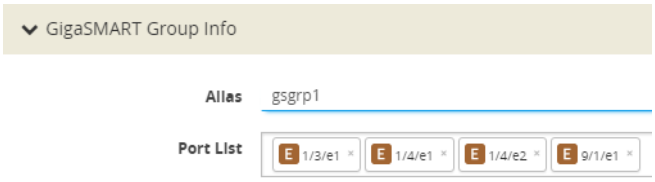

How to Use GigaSMART Operations – Example

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to *Access GigaSMART from GigaVUE-FM* on page 754 for details.

The following procedure summarizes the major steps in creating and using a GigaSMART operation.

1. From the device view, start by selecting **GigaSMART > GigaSMART Groups > GigaSMART Groups** and clicking **New** to create a *GigaSMART group* – a collection of one or more internal GigaSMART ports available in a given chassis. GigaSMART groups are used to process GigaSMART operations – each GSOP you create is assigned to a GigaSMART group.

In this example, a GigaSMART group called **GS1** is created that uses virtual port **e1** on the GigaSMART-HD0 line card in slot 2 of box 16 (**16/2/e1**):

On the GigaSMART Group configuration page:

a. Enter GS1 in the **Alias** field.

b. Click in the **Port List** field to select the port.

   On the GigaSMART Group configuration page, you can also set parameters for specific types of GigaSMART operations.

c. Click **Save**.

2. Next, you can create a *GigaSMART operation* – a combination of packet modification actions that can be used in a map – and assign it to a GigaSMART group for processing.

In this example, a GigaSMART operation called **tcpmask** is created that will overwrite 16 bytes of packet data starting 64 bytes after the end of the TCP header using a hexadecimal **ee** pattern. The GigaSMART operation is assigned to the **GS1** GigaSMART group created in the first step.

To create a GigaSMART operation called **tcpmask**:

a. From the device view, select **GigaSMART > GigaSMART Operations**, and then click **New**.

b. In the **Alias** field, enter **tcpmask**.

c. Click in the **GigaSMART Groups** field and select **GS1** from the list of GigaSMART groups.

d. Click in the **GigaSMART Operations (GSOP)** field and select **Masking** from the list. The configuration dialog for Masking displays.

e. Configure Masking as follows:

   - **Protocol**: TCP
   - **Offset**: 64
   - **Pattern**: ee
   - **Length**: 16

3. Once you have set up a GigaSMART operation, you can include it as part of a map with the **GigaSMART Operations (GSOP)** field in the Map configuration page. In this example, the **tcpmask** GigaSMART operation is combined with an IP Version pass rule so that all IPv4 traffic processed is masked according to the GSOP created in the previous step.

If you are not sure which GigaSMART operation you want to use, click in the GSOP field to display a list of the operations you have already configured.

To configure the map:

a. Select **Maps > Maps > Maps**, and then click **New** to open the New Map page.

b. On the New Map page, configure the map as follows:

   - **Alias**: gsmap

- **Type**: Regular
- **Subtype**: By Rule
- **Source**: select the network ports (for example: 16/3/x7, 16/3/x8, 16/3/x9, 16/3/x10, 16/3/x11, 16/3/x12)
- **Destination**: select the tool port (for example: 16/3/x1)
- **GigaSMART Operations (GSOP)**: tcpmasking (GS1)

c. Click **Add Rule** and specify the following for the rule:

- Select **Pass**
- Click in the **Rule** field and select **IP Version**
- Select **v4** for **Version**.

d. Click **Save**.

Here, a map named **gsmap** is created that forwards IPv4 traffic from network ports 16/3/x7..x12 to tool port 16/3/x1. The traffic will be masked using the **tcpmask** GigaSMART operation created in Step 2.

Engine Watchdog Timer in GigaSMART

In rare scenarios, a packet processing core in the CPU of a GigaSMART engine can enter a deadlocked state. The engine watchdog timer detects the issue and reloads the GigaSMART engine after a specified number of seconds. The engine watchdog timer is enabled by default with a value of 60 seconds. The maximum number of seconds is 600 seconds.

NOTE: If a core is in a deadlocked state, all packets are dropped.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to *Access GigaSMART from GigaVUE-FM* on page 754 for details.

To configure the engine watchdog timer, do the following:

1. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.
2. If you are modifying an existing GigaSMART Group, select the GS Group and click **Edit**. Otherwise, click **New**.
3. In the **Alias** field, enter an alias for this GS Group.
4. In **Port List** field, select the engine port for this GS Group.
5. Under **Engine Timer**, do the following:
   a. Select **Enable** to enable the time or clear the checkbox to disable the timer.
   b. In the **Engine Watchdog Time field**, set the number of seconds to wait before reloading the engine. The minimum is 60. The maximum is 600.

In the following example, the timer is enabled and set to 100 seconds.

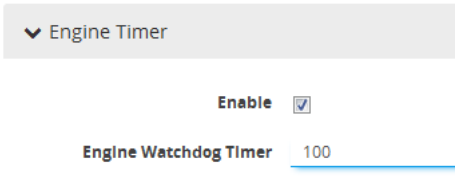

Figure 29-2: Engine Watchdog Timer

Tunnel Health Checks

Starting in software version 5.3, there are tunnel health checks. The reachability of IP destinations is checked and, if the destinations are not reachable, packets will not be sent or will stop being sent.

The tunnel health check on the GigaSMART card defines destinations as follows:

- IP destinations used for sending packets from a single IP interface with tool port to a single IP destination
- tunnel endpoints used for load balancing from a single IP interface with tool port to multiple IP destinations

An SNMP notification can be sent when the status of a IP destination or tunnel endpoint changes, either from Up to Down or from Down to Up.

To enable the SNMP notification, refer to *Enable or Disable Events for SNMP Notifications* on page 192.

To configure ICMP health check parameters for the GigaSMART group, refer to Figure 29-3 on page 762.

Figure 29-3: Configure ICMP Health Check Parameters

To configure UDP health check parameters for the GigaSMART group, refer to Figure 29-4 on page 762.

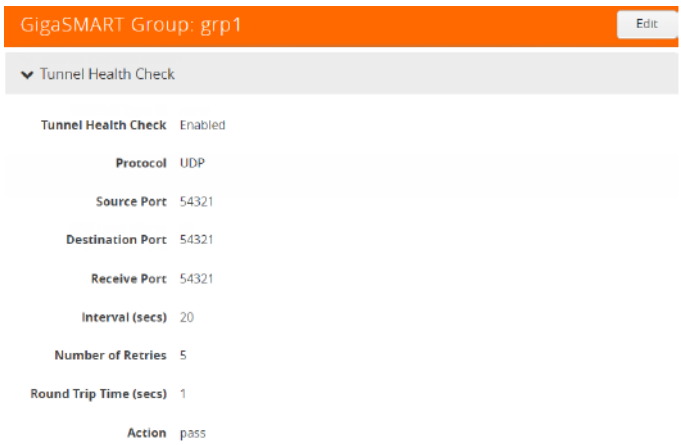

Figure 29-4: Configure UDP Health Check Parameters

To view IP interface status, refer to Figure 29-5 on page 763.

Figure 29-5: View IP Destination Status

To view tunnel endpoint status, refer to Figure 29-6 on page 763.

Figure 29-6: View Tunnel Endpoint Status

To view IP interface statistics, refer to Figure 29-7 on page 763.

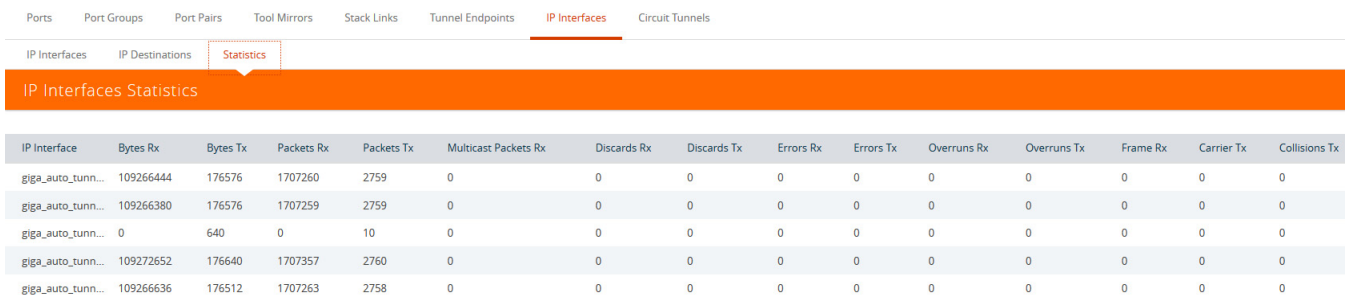

Figure 29-7: View IP DestinationStatistics

To view GigaSMART operation statistics, refer to Figure 29-8 on page 764.

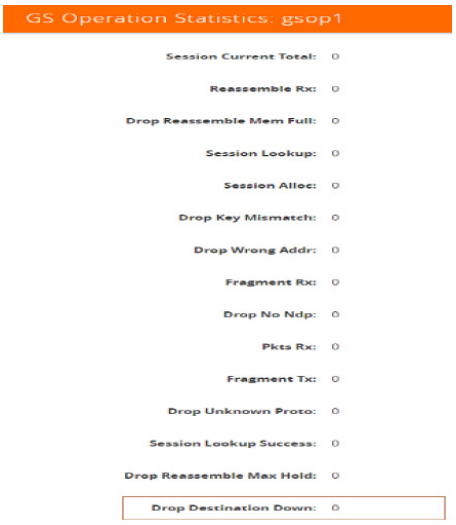

Figure 29-8: View GigaSMART Operation Statistics

Configure Hashing

The **Hash** option in the GigaSMART Group page allows you to select the required hashing option for the GigaSMART Groups. The following options are available:

- **Advanced**: Advanced hashing. *Refer to Fabric Advanced Hashing section in the GigaVUE-HVUE Admin User's Guide.*
- **IPSrcIPDst**: Fixed two tuple hashing based on outer IP, Src and Dst.

GigaSMART Rules and Tips

When using GigaSMART operations, keep in mind the following rules and tips:

Note	Description
Use on Any GigaSMART-Enabled Node	Maps including GigaSMART operations can be bound to *any network port* on a GigaSMART-enabled node: • Standalone GigaVUE-HC3 with SMT-HC3-C05 module. • Standalone GigaVUE-HC2 with front or rear GigaSMART-HC0 module. • Standalone GigaVUE-HC1 nodes. • Standalone GigaVUE-HB1 nodes. • Any GigaVUE H Series node operating in a cluster with one of these node types. Clustered nodes can use maps with GigaSMART operations so long as there is at least one node available in the cluster with GigaSMART capabilities. The GigaSMART group used to power the GigaSMART operation *does not* need to reside on the same physical chassis as the network or tool ports for the map. Refer to *GigaSMART Operations in Clusters* on page 776 for some illustrations of how this works.

Note	Description
Use in Maps with Standard Rule Criteria	Combine GSOPs with map rules carefully to ensure selective application. For example, headers in an IPv4 packet end at a different offset than those in an IPv6 traffic. You can create maps for the following: • Identify IPv4 traffic, slice it at a 64-byte offset and forward the results to tool port 5. • Identify IPv6 traffic, slice it at an 82-byte offset and forward the results to tool port 6. GSOPs are applied to **all** packets matching **any** rule in the map in which the GSOP is included.
Editing Maps and GSOPs	Maps containing GigaSMART operations (GSOPs) should not be edited. Also GSOPs should not be edited once they are associated with maps.
Rules for Tunneling Operations	The rules for tunneling operations are as follows: • A map including a GigaSMART operation with a tunnel decapsulation component (**tunnel-decap**) cannot also include a collector rule. The system prevents situations that would violate this rule. • GigaSMART operations with a tunnel decapsulation component can only be assigned to GigaSMART groups consisting of a single GigaSMART port. • IP interfaces cannot be shared with map-passalls, tool-mirrors, port-pairs, or other regular maps. • For devices involved in tunneling using GMIP. the recommendation is that they run the same software version on both sides of the tunnel (encapsulating/decapsulating).
Combine Multiple Components in a Single Operation	You can combine multiple GigaSMART components into a single operation. For example, you could set up a single GigaSMART operation that masks a packet, strips its VLAN header, and applies a trailer. **NOTE:** With the exception of slicing and masking, most GigaSMART components can be combined in a single operation. Slicing and masking can be combined with other components **but not with each other**. Refer to *How to Combine GigaSMART Operations* for details on the combinations of GigaSMART operations. **NOTE:** The **[trailer <remove>]** argument cannot combined with others – it must be used by itself. Refer to *Remove GigaSMART Trailers* for details. **NOTE:** NetFlow can only be combined with de-duplication and with APF (using second level maps).
GigaVUE-HC3 Nodes with Multiple GigaSMART Modules	The GigaVUE-HC3 chassis supports up to four GigaSMART SMT-HC3-C05 modules. Each module has two GigaSMART engine ports. Each GigaSMART engine port can process packets at up to 100Gb.

Virtual Ports

Virtual ports (vports) are used in flow maps to aggregate and redirect traffic to the GigaSMART ports. It is an aggregation point for traffic to be directed to the GigaSMART second level maps. Second level maps are used for configuring filtering rules enabled through GTP correlation and Adaptive Packet Filtering (APF).

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to *Access GigaSMART from GigaVUE-FM* on page 754 for details.

Create Virtual Port

To create a virtual port, use the following procedure:

1. From the device view, select **GigaSMART > Virtual Ports**.
2. Click **New**.
3. On the Virtual Ports page, do the following:
   a. In the **Alias** field, enter a name for the virtual port. For example, gsTraffic.
   b. From the GigaSMART Group drop-down list, select the GigaSMART group to associate with the virtual port. For example, gsgrp1.
   c. Select the GTP Overlap check box to enable multiple GTP flow sampling maps to receive traffic from the same virtual port.
   d. In the Inline Failover Action drop-down list, select one of the following options:
   - **Tool Bypass** — When the inline tool fails, all traffic coming to the respective inline tool is directed via the bypass path.
   - **Network Bypass** — When the inline tool fails, the traffic is directed to multiple inline tools associated with an inline network or inline network group using rule-based inline maps.
   - **Tool Drop** — When the inline tool fails, all traffic coming to the respective inline tool is dropped.
   - **Network Drop**—When the inline tool fails, all traffic coming to the respective inline tool is dropped.
   - **Network Port Forced Down**—When the inline tool fails, the inline network ports of the respective inline network are forced as "down".

   e. Click **OK**.
4. Create a GigaSMART operation with an Adaptive Packet Filtering (APF) component.
   a. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.
   b. Click **New**.
   c. In the **Alias** field, enter g1 for the GigaSMART Operation.
   d. From the **GigaSMART Groups** list, select a GigaSMART group.

e. From the GigaSMART Operations (GSOP) list, select **Adaptive Packet Filtering**.

5. Create a first-level map name map1 to direct traffic to the virtual port gsTraffic.
   a. Go to the Maps page and click **New**. The New Map page opens.
   b. Enter map1 in the **Alias** field.
   c. For **Type**, select **First Level** and for **Subtype** select **By Rule** map
   d. For **Source**, select a network port. For example, a network port with the alias N1.
   e. For **Destination**, select the virtual port (gsTraffic).
   f. Click Add a Rule, and create rule with the following conditions:
      - Pass
      - VLAN 20
      - IPv4 Protocol UDP
      - Port Destination 2152
   g. Click **Save**.
6. Create second-level maps named map2 and map3 to direct traffic from the virtual port to the GigaSMART operation.

   Configure the map map2 with the **Source** as virtual port gsTraffic, the **Destination** as port T1, and select g1 for **GigaSMART Operation (GSOP)**. For map2, add the following two rules.

   Rule 1:

   - Pass
   - IPv4 Destination 65.128.7.21 Cidr 32
   - IPv4 Protocol TCP
   - Port Destination 80

   Rule 2:

   - Pass
   - IPv4 Destination 98.43.132.70 Cidr 32
   - IPv4 Protocol TCP
   - Port Destination 80

   Configure the map map3 with the **Source** as virtual port gsTraffic, the destination **Destination** as port T2, and g1 for **GigaSMART Operation (GSOP)** For map3, add the following two rules.

   Rule 1:

   - Pass
   - IPv4 Destination 65.128.7.21 Cidr 32
   - IPv4 Protocol TCP
   - Port Destination 443

   Rule 2:

- Pass
- IPv4 Destination 98.43.132.70 Cidr 32
- IPv4 Protocol TCP
- Port Destination 443

If there are other GigaSMART applications defined in the GigaSMART operation, filtering will be done on the packets before sending the GigaSMART applications.

7. Create a shared collector map name mapC1 to direct traffic not matching the second-level maps to the tool port with the alias T3. The collector for the first-level map named map1 uses the standard collector available in prior releases.
   a. Go the Maps page and click **New**. The New Page page opens.
   b. In the **Alias** field, enter mapC1.
   c. For **Type**, select **Regular**.
   d. For **Subtype**, select **Collector**.
   e. For **Source**, select the virtual port gsTraffic.
   f. For **Destination**, select the tool port with the alias T3, and the click **Save**.

Virtual Port Rules

The following rules apply to single virtual ports:

1. A given vport can only belong to one GigaSMART group.
2. Different first level maps with the same network ports can use the same vport. However, you must keep in mind the limitation described in Rule 1.
3. Different first level maps with different network ports can use the same vport. However, you must keep in mind the case described in Rule 2.
4. Different vports can be configured on the same GigaSMART group, but must be used in different maps.
5. A GigaSMART operation can only belong to one GigaSMART engine group.
6. In a first level map, you can specify a vport but not a GigaSMART operation.
7. The vport and the GigaSMART operation used in a second level map must be defined on the same GigaSMART group.
8. In a second level map, a maximum of 5 maps are allowed to be attached to a vport. The maximum number of gsrules in each map is 5. The maximum number of flowrules in each map is 32.

   NOTE: Starting in software version 4.7, this limit is increased for GTP only. Refer to *GigaSMART GTP Whitelisting and GTP Flow Sampling* on page 913. Starting in software version 5.2, this limit is increased for APF and ASF. Refer to *GigaSMART Adaptive Packet Filtering (APF)* on page 1003 and *GigaSMART Application Session Filtering (ASF) and Buffer ASF* on page 1054.
9. Multiple GTP flow sampling maps can receive traffic from the same virtual port when **GTP Overlap** is enabled.

Multiple Virtual Ports for First Level Map

A first level map can have multiple virtual ports (vports). When multiple vports are configured on a first level map, data is sent from network ports to the multiple vports destined to specific GigaSMART groups.

The tool ports of a first level map can be a combination of vports and tool ports. Each vport is bound to a GigaSMART group (gsgroup). Multiple vports are bound to multiple gsgroups. However, in a single first level map, all the vports must be bound to different gsgroups.

NOTE: For a second level map, only a single vport can be configured.

When a second level map is configured using a vport, the data that is sent to the gsgroup is forwarded to the tool ports according to the gsrules or flow rules configured on the map.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to *Access GigaSMART from GigaVUE-FM* on page 754 for details.

To configure multiple vports for a first level map:

1. Create gsgroups on the GigaSMART engine, using the following steps to create a gsgroup with the alias gsgrp1 and one with the alias gsgroup2:
   a. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.
   b. Click **New**.
   c. Enter an alias for the GigaSMART Group in the **Alias** field.
   d. Click in the **Port List** field to select an engine port.
   e. Enable parameters on the GigaSMART Group as needed. For example, to enable GTP correlation, enter the Timeout value under **GTP Flow**.
   f. Click **Save**.
   g. Repeat steps b through f to add another GigaSMART Group.
2. Enable GTP correlation on the GigaSMART groups.
   a. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.
   b. Click **New**.
   c. Enter a name for the GigaSMART Operation in the **Alias** field (for example gs1).
   d. Select the GigaSMART groups from the **GigaSMART Groups** list.
   e. Select **Flow Filtering** from the GigaSMART Operations (GSOP) list.
   f. Click **Save**.
   g. Repeat steps a through f to create the second GigaSMART Operation.

3. Create vports (for example, vp1 and vp2) and assign them to the gsgroups. For the steps to create a virtual ports, refer to *Virtual Ports* on page 766.
4. Create a **First Level By Rule** map and direct traffic to both vports and a tool port by navigating to the Maps page and clicking **New**.
   a. Enter a name for the map in the **Alias** field. For example, map1.
   b. Select **First Level** for **Type** and **By Rule** for **Subtype**.
   c. Specify networks ports in the **Source** field.
   d. Select the virtual ports and a tool port in the **Destination** field.
   e. Click **Add Rule**.
   f. Select **Pass** and **Port Destination** for ports 251 to 252.
   g. Click **Save**.
5. Create a second level map named to direct traffic from the first vport (vp1) to the GTP correlation GigaSMART operation by navigate to the Maps page and clicking **New**.
   a. Enter a name for the map in the **Alias** field. For example, map2.
   b. Select **Second Level** for **Type** and **Flow Filtering** for **Subtype**.
   c. Specify the first virtual port (vp1) in the **Source** field.
   d. Select a tool port in the **Destination** field.
   e. Click in the **GigaSMART Operation (GSOP)** field and select GSOP from the list. For example, gs1.
   f. Click **Add Rule**.
   g. Select **Pass** and **GTP IMSI**.
   h. Enter 302720* in the IMSI field and select **Version V1**.
   i. Click **Save**.
6. Create a shared collector map to direct traffic not matching the second level map to a tool port by navigating to the Maps page and clicking **New**.
   a. Enter a name for the map in the **Alias** field. For example, sc1.
   b. Select **Second Level** for **Type** and **Collector** for **Subtype**.
   c. Select the first virtual port (vp1) in the **Source** field.
   d. Select a tool port in the **Destination** field.
7. Create a second level map named **map3** to direct traffic from the second vport (vp2) to the GTP correlation GigaSMART operation by navigating to the Maps page and clicking **New**.
   a. Enter a name for the map in the **Alias** field. For example, map3.
   b. Select **Second Level** for **Type** and **Flow Filtering** for **Subtype**.
   c. Specify the second virtual port (vp2) in the **Source** field.
   d. Select a tool port in the **Destination** field.
   e. Click in the **GigaSMART Operation (GSOP)** field and select GSOP from the list. For example, gs2.
   f. Click **Add Rule**.

g. Select **Pass** and **GTP IMSI**.

h. Enter * in the IMSI field.

i. Click **Save**.

Multiple Virtual Port Rules

The following rules apply to multiple virtual ports for first level maps:

1. Multiple vports with different GigaSMART groups (gsgroups) can be used on a first level map.
2. Different vports can be configured on the same GigaSMART group, but must be used in different maps.
3. A gsgroup can have multiple GigaSMART operations (gsops).
4. Only one vport is allowed on egress (second level) maps.
5. A vport on a first level map cannot be edited. To make a map change, delete the vport from the **Destination** field first and then add a new vport by clicking in the **Destination** field and selecting another vport.
6. In a standalone system, the number of vports for a first level map is limited by the number of GigaSMART engines (eports) in the chassis.
7. In a cluster environment, the number of vports for a first level map is limited by the number of eports in the cluster.

Multiple Virtual Port with Other GigaSMART Applications

The example in *Multiple Virtual Ports for First Level Map* on page 769 uses the GTP correlation GigaSMART operation. You can also use multiple vports with other GigaSMART operations, such as Adaptive Packet Filtering (APF). You can also chain multiple GigaSMART applications. This allows you to perform different functions and filtering with the same packets.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to *Access GigaSMART from GigaVUE-FM* on page 754 for details.

To configure multiple vports for a first level map:

1. Create GigaSMART Groups on the GigaSMART engines, using the following steps to create three gsgroups with the aliases gg2, gg3, and gg5 and associate with gsops gs2, gs3, and gs5, respectively.
   a. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.
   b. Click **New**.
   c. Enter an alias for the GigaSMART Group in the **Alias** field. (In this example, gg2, gg3, or gg5.)
   d. Click in the **Port List** field to select an engine port.

e. Enable parameters on the GigaSMART Group as needed. For example, to enable GTP correlation, enter the Timeout value under **GTP Flow**.

f. Click **Save**.

g. Repeat steps b through f to add another GigaSMART Group.

2. Enable APF on the GigaSMART groups, using the following steps:

   a. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.

   b. Click **New**.

   c. Enter an alias for the GigaSMART Operation in the **Alias** field. (In this example, gs2, gs3, or gs5.)

   d. Select the GigaSMART group from the **GigaSMART Groups** list. (For gs2, select gg2. For gs3, select gg3. For gs5, select gg3.)

   e. Select **APF** from the GigaSMART Operations (GSOP) list and select **Enable**.

   f. Select **Slicing** from the GigaSMART Operations (GSOP) list, and then set **Protocol** to none. For gs2 set the **offset** to80. For gs3, set the **offset** to 90. For gs5, set the **offset** to 100.

   g. Click **Save**.

   h. Repeat steps a through f to create another GigaSMART Operation.

3. Create vports and assign them to the GigaSMART Groups, using the following steps to create three virtual ports named vp2, vp3, and vp4, assigning them to gsgroup gg2, gg3, and gg5 respectively.

   a. From the device view, select **GigaSMART > Virtual Ports**.

   b. Click **New**.

   c. Enter an alias for the virtual in the **Alias** field. (In this example, vp2, vp3, or vp5.)

   d. Select the gsgroup from the **GigaSMART Group** field (gg2 for vp2, gg3 for vp2, and gg5 for vp5).

   e. Click **Save**.

   f. Repeat step b through step e to create another virtual port.

4. Create a first level map, and direct traffic to two vports and a tool port, using the following steps:

   a. Select **Maps > Maps > Maps**.

   b. Click **New**.

   c. Enter m1 in the **Alias** field.

   d. Select **First Level** for **Type** and **By Rule** for **Subtype**.

   e. Select a network port in the **Source** field.

   f. Select virtual ports vp2 and vp3 plus a tool port in the **Destination** field.

   g. Use the **Add a Rule** button to add two rules to the map.

      For the first rule, select **pass** and select **VLAN** for the condition. Set the VLAN value to 100.

For the second rule, select **pass** and select **VLAN** for the condition. Set the VLAN value to 200.

h. Click **Save**.

5. Create a second level map named to direct traffic from the vports to the GigaSMART operation, using the following steps:
   a. Select **Maps > Maps > Maps**.
   b. Click **New**.
   c. Enter m2 in the **Alias** field.
   d. Select **Second Level** for **Type** and **By Rule** for **Subtype**.
   e. Select a virtual port vp2 the **Source** field.
   f. Select a tool port in the **Destination** field.
   g. Select gs2 for **GigaSMART Operation (GSOP)**.
   h. Use the **Add a Rule** button to add a rules to the map.

      For the rule, select **pass** and select **VLAN** for the condition. Set the VLAN value to 100 with **Subset** set to none and **Position** set to 0.

   i. Click **Save**.
6. Create another second level map named to direct traffic from the vports to the GigaSMART operation, using the following steps:
   a. Select **Maps > Maps > Maps**.
   b. Click **New**.
   c. Enter m4 in the **Alias** field.
   d. Select **Second Level** for **Type** and **By Rule** for **Subtype**.
   e. Select a virtual port vp3 in the **Source** field.
   f. Select a tool port in the **Destination** field.
   g. Select gs3 for **GigaSMART Operation (GSOP)**.
   h. Use the **Add a Rule** button to add a rules to the map.

      For the rule, select **pass** and select **VLAN** for the condition. Set the VLAN value to 200 with **Subset** set to none and **Position** set to 0.

   i. Click **Save**.
7. Create another second level map to direct traffic from the vports to the GigaSMART operation.
   a. Select **Maps > Maps > Maps**.
   b. Click **New**.
   c. Enter m6 in the **Alias** field.
   d. Select **Second Level** for **Type** and **By Rule** for **Subtype**.
   e. Select a virtual port vp3 in the **Source** field.
   f. Select a tool port in the **Destination** field.
   g. Select gs3 for **GigaSMART Operation (GSOP)**.
   h. Use the **Add a Rule** button to add a rules to the map.

For the rule, select **pass** and select **VLAN** for the condition. Set the VLAN value to 200 with **Subset** set to none and **Position** set to 0.

i. Click **Save**.

8. Create a shared collector map named to direct traffic not matching the maps to a tool port.
   a. Select **Maps > Maps > Maps**.
   b. Click **New**.
   c. Enter sc1 in the **Alias** field.
   d. Select **Second Level** for **Type** and **Collector** for **Subtype**.
   e. Select a virtual port vp1 in the **Source** field.
   f. Select a tool port in the **Destination** field.
   g. Click **Save**.

Virtual Port Statistics

The Virtual Ports page displays statistics about the configured virtual ports. To view the statistics select **GigaSMART > Virtual Ports > Statistics**.

The following table describes virtual port statistics:

Statistic	Description
Rx Packets	The number of packets received into the virtual port.
Tx Packets	The number of packets transmitted out of the virtual port.
Rx Octets	The number of bytes received into the virtual port.
Tx Octets	The number of bytes transmitted out of the virtual port.
Packets Drops	The number of packets dropped at the virtual port.
Packet Drops No Init	For internal debugging.

Differences in GigaSMART Nomenclature Between the CLI and H-VUE

The CLI and the Web-based H-VUE interface occasionally use different names to refer to the same GigaSMART-related functionality. The following table summarizes the differences:

Documentation Term	CLI Term	H-VUE Term	Description
GigaSMART Port	e1, e2	GigaSMART Engine	Internal ports on GigaSMART line card, module, or GigaVUE-HB1 node used to power GigaSMART features.
GigaSMART Group	gsgroup	GigaSMART Engine Group	Group of internal ports on GigaSMART line card, module, or GigaVUE-HB1 node used to power GigaSMART features. You configure these in the CLI with the **gsgroup** command.
GigaSMART Component	gsop	GigaSMART Operation	One of the available GigaSMART packet processing features (de-duplication, slicing, and so on)
GigaSMART Operation	gsop	GigaSMART Operation Group	A combination of GigaSMART packet-processing features into a single entity used in a map. You configure these in the CLI with the **gsop** command.

GigaSMART Operations in Clusters

Clustered environments with at least one GigaSMART-enabled chassis can take advantage of GigaSMART operations in maps on network ports elsewhere in the cluster. As shown in Figure 29-9 and Figure 29-10, the GigaSMART group providing the packet processing power for a GigaSMART operation does not have to be on the same chassis as the network or tool ports for a map.

- A map's network ports, tool ports, and the GigaSMART group ports can all be on different nodes in a clustered environment.
- GigaSMART operations do not require any specific role for the chassis with the GigaSMART group – it can be a master, standby, or normal node.
- The main requirements to keep in mind is that the GigaSMART ports in a GigaSMART group must all be on the same chassis. So, for example, if you had two GigaSMART-HD0 line cards in a single chassis, you could create a single GigaSMART group out of the four GigaSMART engine ports available across the two line cards. However, if the two line cards were in two different chassis, they could not be combined into a single GigaSMART group.

Figure 29-9 and Figure 29-10 illustrate examples of network, tool, and GigaSMART group ports on different nodes in a clustered environment. For example, Figure 29-9 shows a map accepting ingress packets on GV1, sending them to the GigaSMART group on GV3 for GigaSMART processing (for example, de-duplication), and sending the results to a tool port on GV4.

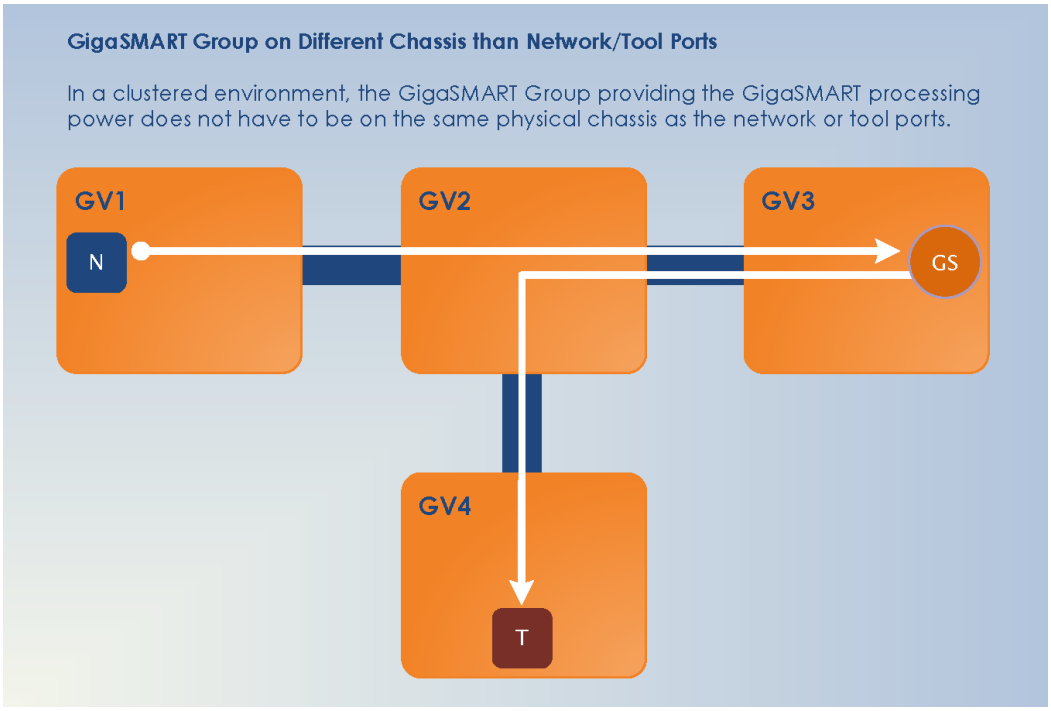

Figure 29-9: GigaSMART Group on Different Chassis than Network/Tool Ports

You could also have the GigaSMART group on the same chassis as either the network or tool ports for the map. An example of this is shown in Figure 29-10. The GigaSMART group on GV3 is on the same chassis as the egress port.

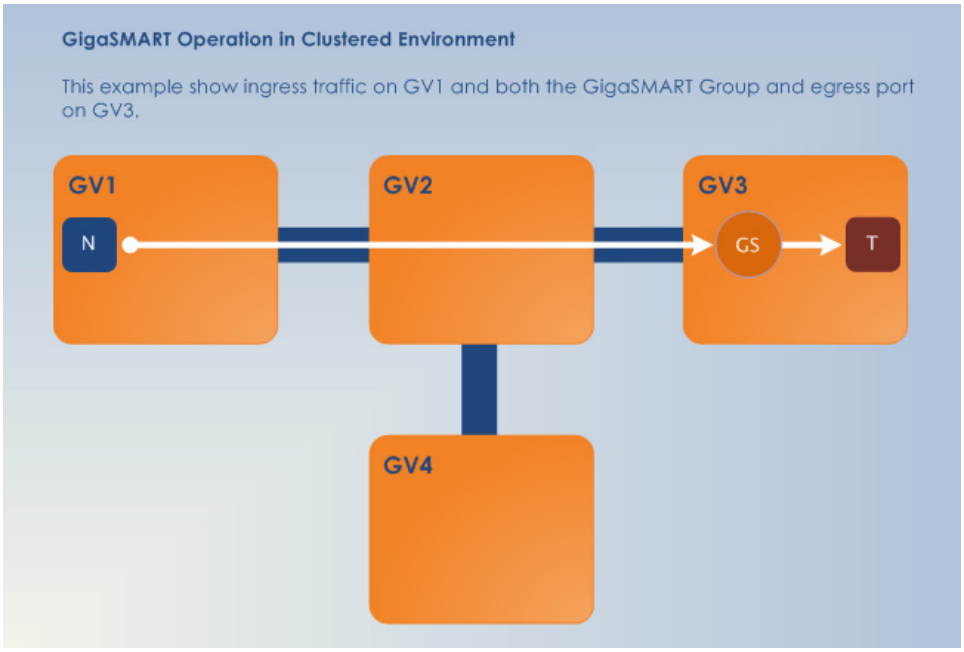

Figure 29-10: GigaSMART Operations in Clusters

How to Combine GigaSMART Operations

You can combine multiple GigaSMART components into a single operation. For example, you can set up a single GigaSMART operation that masks a packet, strips its VLAN header, and inserts a field in the GigaSMART Trailer.

The following table summarizes the valid combinations of GigaSMART operations (an X in the table indicates an invalid combination).

	Slicing	Masking	Source Id	Header/Trailer Remove	Dedup	Tunnel Encap	Tunnel Decap	Strip Header	Add Header	Flow-Ops (FlowVUE)	Flow-Ops (Flow-Filter GTP)	Flow-Ops (GTP Whitelist)	Flow-Ops (GTP Flow Sampling)	APF	ASF	NetFlow (1st Level Maps)	NetFlow (2nd Level Maps)	Load Balance	SSL Decryption	Inline-SSL
Slicing																				
Masking																				
Source Id (Add Trailer)																				
Header/Trailer Remove	X																			
Dedup																				
Tunnel Encap																				
Tunnel Decap				X		X														
Strip Header																				
Add Header						X														
Flow-Ops (FlowVUE)						X	X													
Flow-Ops (Flow-Filter GTP)		X	X	X	X	X	X	X	X	X										
Flow-Ops (GTP Whitelist)		X	X	X	X	X	X	X	X	X										
Flow-Ops (GTP Flow Sampling)		X	X	X	X	X	X	X	X	X	X	X								
APF				X			X				X	X	X							
ASF				X			X				X	X	X							
NetFlow (1st Level Maps)	X	X	X	X		X	X	X	X	X	X	X	X	X	X					
NetFlow (2nd Level Maps)	X	X	X	X		X	X	X	X	X	X	X	X		X					
Load Balance																X	X		X	X
SSL Decryption	X	X	X	X		X	X	X	X	X	X	X	X	X	X	X	X	X		
Inline-SSL	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	

NOTE: Masking operations cannot be combined with slicing unless the offset of the slicing is after the offset of the masking.

How to Read GigaSMART Operations Table

The GigaSMART operations table is read both across and down. The following is an example of how to read the GigaSMART operations table:

1. Begin in the left-most column and select a GigaSMART operation, for example: Add Header.
2. Move to the right along the Add Header row. Add Header can be combined with Slicing or Masking, Source Id, Header/Trailer Remove, Dedup, Tunnel Decap, and Strip Header. It cannot be combined with Tunnel Encap.
3. Move to the right another square to the gray square at end of the Add Header row. This is the Add Header column.
4. Move down the Add Header column, below the gray square at the end of the Add Header row. Add Header can be combined with Flow-Ops (FlowVUE), APF, ASF, and Load Balance. It cannot be combined with Flow-Ops (Flow-Filter GTP), Flow-Ops (GTP Whitelist), Flow-Ops (GTP Flow Sampling), NetFlow (1st Level Maps), NetFlow (2nd Level Maps), SSL Decryption (for Out-of-Band Tools), Inline-SSL (SSL Decryption for Inline Tools), Flow-Ops (SIP Flow Sampling), or Flow-Ops (SIP Flow Whitelist).

Work with GigaSMART Operation Combinations in H-VUE

When combining GigaSMART operations in H-VUE, the drop down option in the GSOP screen will gray out if a set of combinations is invalid.

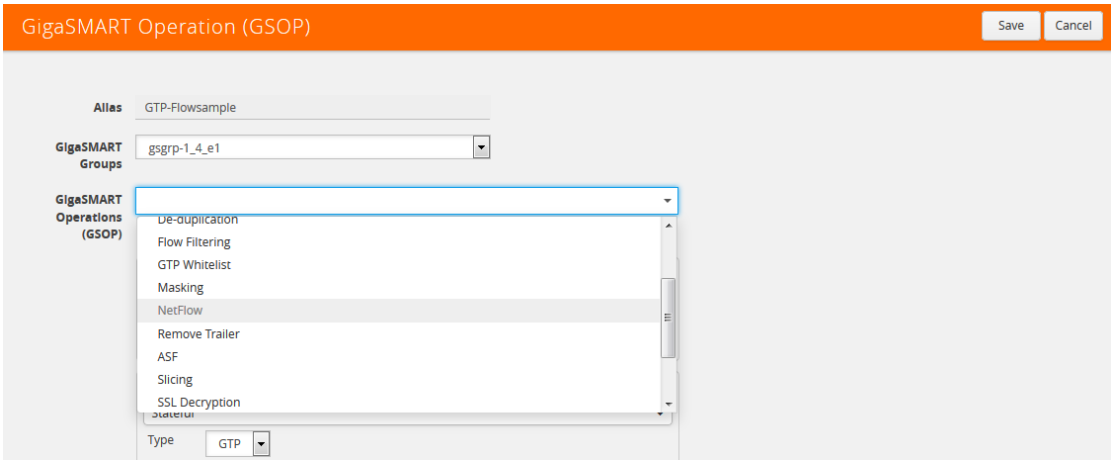

Figure 29-11: GigaSMART Operation Combinations Not Available

However, in the following cases, you may not see the combinations grayed out but when trying to save the combination a pop-up is displayed stating that the combination is invalid.

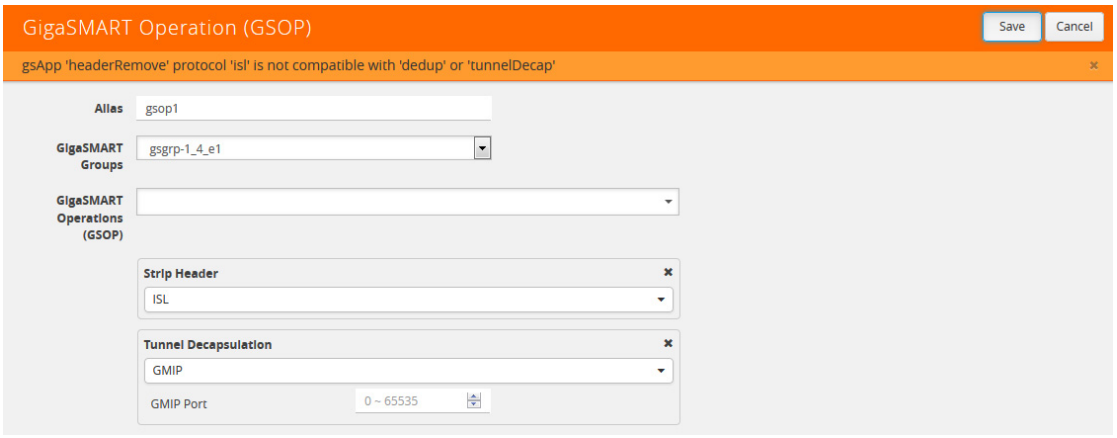

Figure 29-12: GigaSMART Operation Combinations Unable to Save

To get the valid combinations, refer to the table at the beginning of this section that summarizes the combinations.

Supported GigaSMART Operations

The following table lists the supported GigaSMART operations by GigaVUE node.

GigaSMART Operation	GigaVUE-HD4/ GigaVUE-HD8	GigaVUE-HC1	GigaVUE-HC2	GigaVUE-HC3	GigaVUE-HB1
Masking	✓	✓	✓	✓	✓
Slicing	✓	✓	✓	✓	✓
Source ID (add Trailer)	✓	✓	✓	✓	✓
Header/Trailer Remove	✓	✓	✓	✓	✓
De-duplication	✓	✓	✓	✓	✓
Tunnel Encapsulation	✓	✓	✓	✓	✓
Tunnel Decapsulation	✓	✓	✓	✓	✓
Strip Header	✓	✓	✓	✓	✓
Add Header	✓	✓	✓	✓	✓
FlowVUE (IP-based)	✓	✓	✓	✓	✓
GTP Flow Filtering	✓	×	✓	✓	×
GTP Whitelisting	✓	×	✓	✓	×
GTP Flow Sampling	✓	×	✓	✓	×
APF	✓	✓	✓	✓	✓
ASF	✓	✓	✓	✓	✓
NetFlow (1st Level Maps)	✓	✓	✓	✓	✓
NetFlow (2nd Level Maps)	✓	✓	✓	✓	✓
Load Balancing (Stateless)	✓	✓	✓	✓	✓
Load Balancing (Stateful)	✓	✓	✓	✓	✓
SSL Decryption for Out-of-Band Tools	✓	✓	✓	✓	✓
SSL Decryption for Inline Tools	×	×	✓	✓	×
SIP Flow Sampling	✓	×	✓	✓	×
SIP Flow Whitelist	✓	×	✓	✓	×

NOTE: For GTP on GigaVUE-HB1, the following are not supported:

- GTP whitelisting in a cluster
- GTP engine grouping
- Access Point Name (APN)

Order of GigaSMART Operations

When combining multiple GigaSMART components into a single operation, the components are applied in the following order:

1. [apf/asf/buffer asf] OR [flow-ops <flow-filtering>] or [flow-ops <gtp-whitelist>] or [flow-ops <gtp-flowsampling>] AND [lb] (Stateful load balancing with GTP or ASF)
2. [tunnel-decap]
3. [dedup]
4. [ssl-decrypt]
5. [flow-ops <flow-sampling>]
6. [strip-header]
7. [slicing]
8. [masking]
9. [trailer]
10. [tunnel-encap]
11. [add-header]
12. [flow-ops <netflow>]
13. [lb] (Stateless load balancing)

View GigaSMART Statistics

To view the GigaSMART statistics from the UI, select **GigaSMART Operations (GSOP) > Statistics**. The Statistics page displays as shown in <HyperText>Figure 29-13.

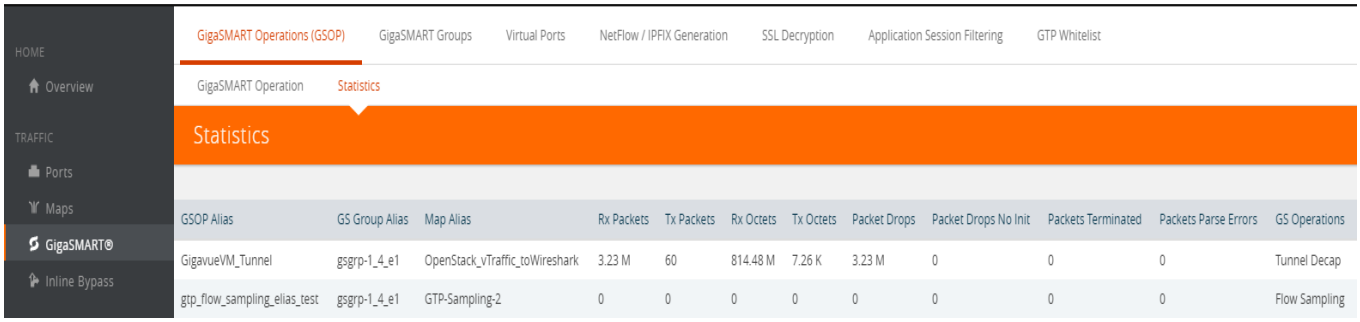

Figure 29-13: GigaSMART Operations Statistics Page

The Statistics page shows the aliases for GS Operation, GS Group, the alias of the map that is using GS Operation, and the GS Operations being used. In <HyperText>Figure 29-13, Flow Sampling and Load Balance operations are assigned to the alias GTP-Flowsample in the gsgrp-1_4_e1 GS Group. The GS operations is used by the map with the alias map-gtpFS. For a description of the other columns, refer to *GigaSMART Operations Statistics Definitions* on page 794.

Other statistics available for viewing from the UI are the following:

- GigaSMART Groups Statistics (select **GigaSMART > Statistics**)
- NetFlow / IPF Generation Statistics for Exporters and Monitors (select **NetFlow / IPFIX > Exporter Statistics or NetFlow / IPFIX Generation > Monitor Statistics**)
- IP Interface Statistics (select **Ports > IP Interfaces > Statistics**)

Definitions of GigaSMART Statistics

The following sections provide definitions for the statistics displayed. Refer to the following:

NetFlow Monitor Statistics Definitions

To view NetFlow Monitor statistics, select **GigaSMART > NetFlow / IPFIX Generation > Monitor Statistics** to open the Monitor Statistics page shown in Figure 29-14

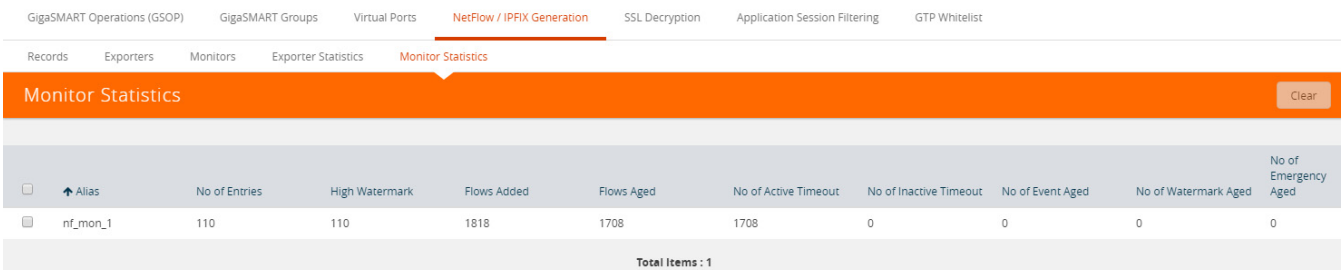

Figure 29-14: NetFlow Monitor Statistics

The following table describes NetFlow Monitor statistics displayed on the Monitor Statistics page:

Statistic	Description
No Entries	The current number of flows in the monitor cache.
High Watermark	The maximum number of flows that have ever been in the monitor cache.
Flows Added	The sum of all flows added to the monitor cache.
Flows Aged	The sum of the flows that have aged due to the following: • Active Timeout—The configured active timeout for the monitor cache was exceeded. • Inactive Timeout—The configured inactive timeout for the monitor cache was exceeded. • Event Aged—The number of entries aged from the cache because a TCP FIN/RST flag was received. • Watermark Aged—The number of entries aged from the cache because the CPU utilization of the cache exceeded 75%. • Emergency Aged—The number of entries aged from the cache because a user requested a forced flush through the CLI.
No of Active Timeout	The number of times the configured active timeout for the monitor cache was exceeded.
No of Inactive Timeout	The number of times the configured inactive timeout for the monitor cache was exceeded.
No of Event Aged	The number of entries aged from the cache because a TCP FIN/RST flag was received.
No of Watermarked Aged	The number of entries aged from the cache because the CPU utilization of the cache exceeded 75%.
No of Emergency Aged	The number of entries aged from the cache because a user requested a forced flush through the CLI.

NetFlow Exporter Statistics Definitions

To view NetFlow Exporter statistics, select **GigaSMART > NetFlow / IPFIX Generation > Monitor Statistics** to open the Exporter Statistics page. Figure 29-14 shows an example.

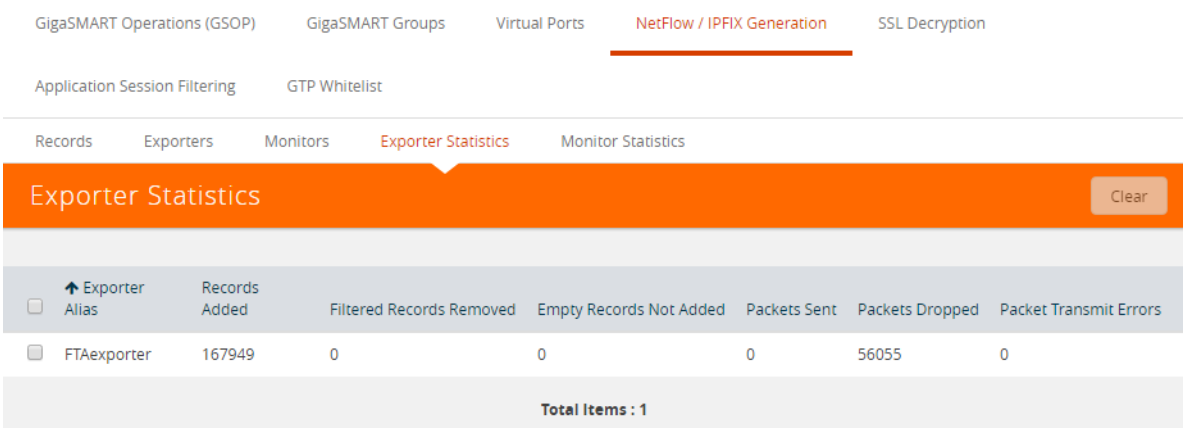

Figure 29-15: NetFlow Exporter Statistics

The following table describes NetFlow Exporter statistics:

Statistic	Description
Templates Added	The number of data templates added to the exporter.
Records Added	The number of data records added to the exporter.
Filtered Records Removed	The number of records filtered by exporter filters
Empty Records Not Added	The number of records not added to NetFlow because they were empty or blank. In the NetFlow record, if all the collect fields contain only enterprise extensions such as URL, HTTP, or DNS, and if during run-time, the records are blank or empty, they will be counted as Empty Records Not Added.
Packets Sent	The number of packets sent from the exporter to the collector.
Packet Dropped	The number of packets dropped, which could be due to the inability to send packets, such as there is no network connection or the port is down.
Packet Transmit Errors	The number of packets dropped, which could be due to the inability to send packets, such as there is no network connection or the port is down.

IP Interfaces Statistics Definitions

To view IP Interfaces statistics, select **Ports > IP Interfaces > Statistics** to open the IP Interfaces Statistics page. The statistics of the control traffic such as ARP, ICMP, and ICMPv6 for the physical node will be displayed. Figure 29-16 shows an example.

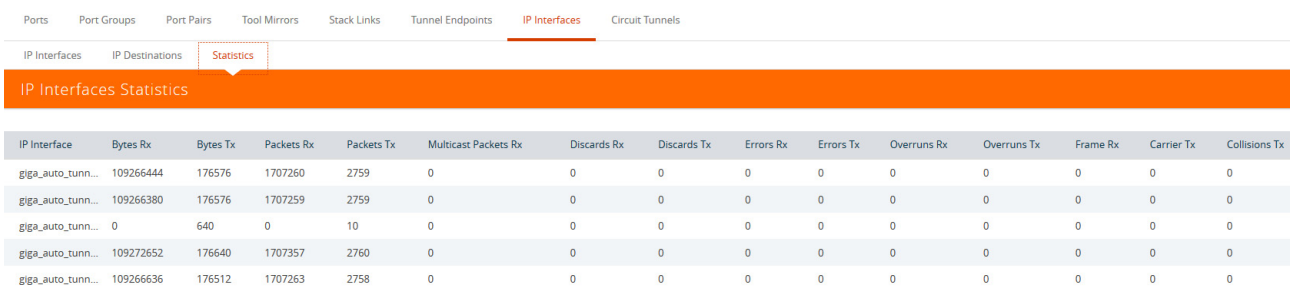

Figure 29-16: IP Interfaces Statistics

The following table describes IP interfaces statistics:

Statistic	Description
IP Interface	The IP interface for which the statistics is displayed.
Bytes Rx	The number of bytes received into the IP interface.
Bytes Tx	The number of bytes transmitted out of the IP interface.
Packets Rx	The number of packets received into the IP interface.
Packets Tx	The number of packets transmitted out of the IP interface.
Multicast Packets Rx	The number of multicast packets received into the IP interface.
Discards Rx	The number of packets that were received and discarded by the IP interface.
Discards Tx	The number of packets that were transmitted out of the IP interface and were discarded.
Errors Rx	The number of packets with errors that were received into the IP interface.
Errors Tx	The number of packets with errors that were transmitted out of the IP interface.
Overruns Rx	The number of first-in-first-out (FIFO) buffer errors that were received into the IP interface.
Overruns Tx	The number of first-in-first-out (FIFO) buffer errors that were transmitted out of the IP interface.
Frame Rx	The number of frames received into the IP interface.
Carrier Tx	The number of packets in which carrier losses were detected when transmitted out of the IP interface.
Collisions Tx	The number of packets that were colliding when transmitted out of the IP interface.

GigaSMART Group Statistics Definitions

To view GigaSMART Group statistics, select **GigaSMART > GigaSMART Groups > Statistics** to open the Statistics page. Figure 29-17 shows an example.

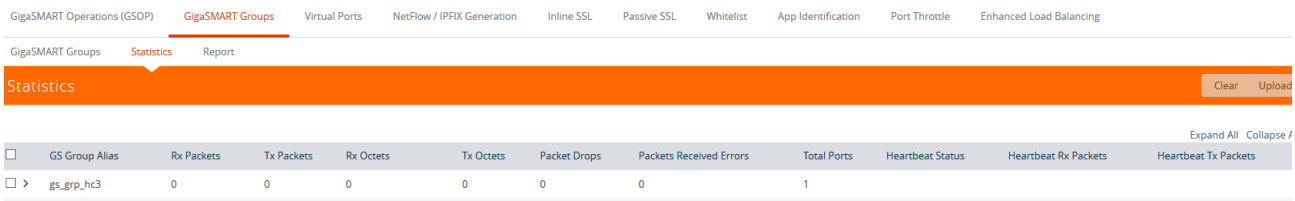

Figure 29-17: GigaSMART Group Statistics

The following table describes GigaSMART group statistics:

Statistic	Description
Rx Packets	The cumulative number of packets coming into a GigaSMART group.
Tx Packets	The cumulative number of packets going out of a GigaSMART group.
Rx Octets	The cumulative number of bytes coming into a GigaSMART group.
Tx Octets	The cumulative number of bytes going out of a GigaSMART group.
Packet Drops	The cumulative number of packets dropped at a GigaSMART group.
Packet Received Errors	The cumulative number of received packets with errors at a GigaSMART group.
Heartbeat Status of eport	The status of heartbeat. The valid values are up, down, or NA. This heartbeat status is for a GigaSMART engine port on a GigaSMART group for inline SSL decryption On a supported platform, such as GigaVUE-HC2, the status will be either up or down, followed by number of packets sent and received. On an unsupported platform, the status will be NA and the number of packets will not be displayed. The heartbeat status for the GigaSMART cards is fully operational only when it is configured with the inline SSL application. In all other cases, the default value of the heartbeat status is Up and the status is generally ignored.
Heartbeat Rx Packets	The heartbeat receive packet count.
Heartbeat Tx Packets	The heartbeat transmit packet count.

GigaSMART Group Flow Ops Report Statistics Definitions

The following sections provide definitions for the statistics that are reported for the GigaSMART group. Refer to the following:

Flow Ops Report Statistics Definitions for FlowVUE

The following table describes Flow Ops report statistics for FlowVUE:

Statistic	Description
Device IP	The IP address of the flow.
In Sample	The sample selected for pass (1) or drop (0).
Num Packets	The number of packets seen for the flow.
Num Octets	The sum of the packet lengths of all packets seen for the flow.
num_devices	The number of devices.
num_devices_in_sample	The number of devices in the sample.

Flow Ops Report Statistics Definitions for GTP

The following table describes Flow Ops report statistics for GTP, including GTP flow filtering, GTP whitelisting, and GTP flow sampling:

Statistic	Description
Tunnel[Ver]	The tunnel type (CTRL for control plane, USER for user data plane) and version (1 or 2).
Interface EBI:LBI[QCI]	The interface and Evolved Packet System (EPS) Bearer Identifier (EBI), Linked Bearer Identity (LBI), and QoS Class Identifier.
IP:Tunnel-ID ==> IP:Tunnel-ID	The IP addresses and tunnel identifiers of both sides of the tunnel. **NOTE:** For LTE nodes only, the IP addresses can be both IPv4 and IPv6.
IMSI	The International Mobile Subscriber Identity (IMSI) value.
APN	The Access Point Name (APN) value.
WL	The IMSI had a match in the whitelist (WL) or did not. The values are N for no and Y for yes.
FS	The IMSI was flow sampled or not. The values are A for accepted, R for rejected, and N for no match.
ID	The rule ID of the flow sample rule.

Statistic	Description
LB port	The load balancing port number.
Pkts	The number of packets.
Timestamp	The internal clock time of when the session was created.
GTP Resource Summary	
Num Sessions In Use	The number of correlated session in use.
Num Tunnels In Use	The number of used tunnels in the tunnel resource pool.
Tunnels Available	The number of available tunnels in the tunnel resource pool.
UPN CTunnels Available	The number of available tunnels in the UPN C tunnels resource pool.
Tunnels Pending Free	The number of tunnels marked free, ready to be returned to the tunnel resource pool.
Tunnels Marked Free	The number of times used tunnels are marked free.
Tunnels Returned	The number of times used tunnels are returned to the tunnel resource pool.
Current Time	The current time (in CPU cycles). Used for debugging.
Flow Filtering Report Summary	
Control Tunnels	The total number control tunnels.
Control & User Tunnels	The total number of control and user tunnels.
GTP Session Stats	The interface type: Gn/Gp, S1U/S11, S5/S8, S3/S4, or Other.
Sessions	The number of sessions by interface. Note the following: • The counter for S3/S4 will always be 0 (not supported). • The S10 interface is counted under Other.
Tunnels	The number of tunnels by interface.
Pending Session	The number of sessions waiting for control message response.
Control Only Session	The number of sessions without user bearers.
Flow Sampling Report Summary	
Total Devices	The total number devices in the flow sample
Number of Device in Sample	The number of devices used in each sample session
Flow SIP Report Summary	
SIP Sessions	This graph displays the total sessions, total sessions in use and the number of parse errors.
RTP Sessions	This graph displays the total number of sessions and the total data pool in use.
GTP Interface Stats For overlap maps, refer to the notes below the table for *Flow Ops Report Statistics Definitions for GTP Overlap* on page 791	

Statistic	Description
Rx Pkts	The received (Rx) packets for GTP correlation statistics for GTP traffic by interface type: S11, S1u, S5S8 (control and user), Gn (control and user), total (control and user), collector (control and user). **NOTE:** If traffic does not match any map rules, it will be sent to the collector.
Rx Bytes	The received (Rx) bytes for GTP correlation statistics for GTP traffic by interface type.
Sample/WL/Filter (Tx) Pkts	The transmitted (Tx) packets sampled in for flow sampling, whitelisting, and flow filtering by interface type. For example, if sampling is 60%, then 60% is sampled in.
Sample/WL/Filter (Tx) Bytes	The transmitted (Tx) bytes sampled in for flow sampling, whitelisting, and flow filtering by interface type.
Sample Out (Dropped) Pkts	The packets sampled out (dropped) by interface type. For example, if sampling is 60%, then 40% is sampled out.
Sample Out (Dropped) Bytes	The bytes sampled out (dropped) by interface type.
Xaui Drop	The total traffic (Rx bytes and packets) dropped due to oversubscription on the interface into the GigaSMART.
Statistics for Control Message (GTP-c)	
GTPV1	The GTPv1 (version 1) message type.
Cre PDP Req	Create Packet Data Protocol (PDP) context request.
Cre PDP Rsp	Create PDP context response.
Upd PDP Req	Update PDP context request.
Upd PDP Rsp	Update PDP context response.
Del PDP Req	Delete PDP context request.
Del PDP Rsp	Delete PDP context response.
GTPV2	The GTPv2 (version 2) message type.
Cre Ssn Req	Create session request.
Cre Ssn Rsp	Create session response.
Mod Bea Req	Modify bearer request.
Mod Bea Rsp	Modify bearer response.
Del Ssn Req	Delete session request.
Del Ssn Rsp	Delete session response.
Cre Bea Req	Create bearer request.
Cre Bea Rsp	Create bearer response.
Upd Bea Req	Update bearer request.
Upd Bea Rsp	Update bearer response.

Statistic	Description
Del Bea Req	Delete bearer request.
Del Bea Rsp	Delete bearer response.
Mod Bea Cmd	Modify bearer command.
Mod Bea Fai	Modify bearer failure indication.
Bea Rsr Cmd	Bearer resource command.
Bea Rsr Fai	Bearer resource failure indication.
Message Counters	
Tool Pass	The number of control messages passed to the tools.
Col NoSess	The number of control messages sent to the collector without matching sessions.
Col NoTnlx	The number of “out of tunnels” requests sent to the collector for the control message.
Col ParseEr	The number of control messages sent to the collector with unsupported options.
Col NoRule	The number of control messages sent to the collector without matching rules.
Col Other	The number of control messages sent to the collector with other conditions, for example, the message matched a drop rule.
Statistics for User Data Message (GTP-u)	
Tool Pass	The number of user data messages passed to the tools.
Collector	The number of user data messages sent to the collector without matching sessions.
Drop	The number of user data messages dropped.

Flow Ops Report Statistics Definitions for GTP Overlap

The following table describes Flow Ops report statistics for GTP overlap flow sampling. Since most fields are the same as *Flow Ops Report Statistics Definitions for GTP* on page 719, they are not repeated. Refer to the note below the table for GTP Interface Stats for overlap maps.

Statistic	Description
Tunnel[Ver]	The tunnel type (CTRL for control plane, USER for user data plane) and version (1 or 2).
Interface EBI:LBI[QCI]	The interface and Evolved Packet System (EPS) Bearer Identifier (EBI), Linked Bearer Identity (LBI), and QoS Class Identifier.
IP:Tunnel-ID ==> IP:Tunnel-ID	The IP addresses and tunnel identifiers of both sides of the tunnel.
IMSI	The International Mobile Subscriber Identity (IMSI) value.
APN	The Access Point Name (APN) value.

Statistic	Description
Overlap Result	The overlap result. A map group for GTP overlap flow sampling maps contains six maps. Each character in the result represents one of the maps. The result is in alphabetical order for each map within the map group. (There is no map priority.) Refer to the following legend: • A—Flow Sample Accept • R—Flow Sample Reject • W—Whitelist Accept • N—No Match
Pkts	The number of packets.
Timestamp	The internal clock time of when the session was created.

The following notes are for the combination of overlap and non-overlap maps:

- If at least one flow sample map accepts the packets among all the pairs of overlap maps, the Sample (Tx) counters in the GTP Interface Stats will be incremented. If more than one pair of maps accepts the packets the Sample (Tx) counters in the GTP Interface Stats will still be incremented only once.
- If at least one whitelist map matches the packets among all the pairs of overlap maps, the Sample (Tx) counters in the GTP Interface Stats will be incremented. If more than one pair of maps matches the packets the Sample (Tx) counters in the GTP Interface Stats will still be incremented only once.
- If there are no whitelist maps and all flow sample maps are no-rule-match, the Sample (Tx) and Sample Out counters in the GTP Interface Stats will not be incremented.

Flow Ops Report Definitions for SIP/RTP Correlation

The following table describes Flow Ops report statistics for SIP/RTP correlation.

Statistic	Description
PROTO	The protocol, such as SIP, RTP, or RTCP. **NOTE:** The RTCP port number is assumed to be the corresponding RTP port number plus 1. This results in even numbered RTP ports and odd numbered RTCP ports.
TRANSPORT	The transport layer, such as UDP or TCP.
METHOD	The SIP method (or SIP message).
CALLER: IP	The corresponding IP address of the caller.
CALLEE: IP	The corresponding IP address of the callee.
PDU	The number of packets seen for this session.
CALL-ID	The call identifier.

Statistic	Description
WL	The caller/callee ID had a match in the whitelist (WL) or did not. The values are N for no and Y for yes.
FS	The caller ID was flow sampled or not. The values are A for accepted, R for rejected, and N for no match.
ID	The rule ID
LB port	The load balancing port number, which is the port over which the session has been load balanced. All SIP and corresponding RTP packets will be sent to this port.
Timestamp	The internal clock time of when the session was created.
Message counters	
SIP messages	All the SIP messages that have been correlated, such as, ACK, BYE 200, CANCEL, INFO, and so on. These counters are cumulative
Tool Pass	The number of SIP messages passed to the tools
NoSess	The number of SIP messages sent to the collector without matching sessions
NoRule	The number of SIP messages sent to the collector without matching rules.
NoMatch	Reserved
Other	The number of SIP messages sent to the collector with other conditions.
SIP Resource Summary	
Num Sessions In Use	The number of SIP sessions.
Sessions Available	The number of remaining SIP sessions available. (This varies by GigaVUE node.)
RTP Resource Summary	
Num Data Pools In Use	The number of RTP sessions.
Data Pools Available	The number of remaining RTP sessions available. (This varies by GigaVUE node.)

Flow Ops Report Statistics for Out-of-Band SSL Decryption

The following table describes Flow Ops report statistics for out-of-band SSL decryption:

Statistic	Description
Server IP	The IP address of the server.
Server Port	The port number of the server.
Client IP	The IP address of the client.

Statistic	Description
Client Port	The port number of the client.
Version	The version.
First Error	The first error encountered on a session. After the first error, subsequent packets are dropped. If a session encounters errors, the packets for that session are ignored. The session will be cleared after the session times out.
In Pkts	The number of packets going into the session.
Out Pkts	The number of packets transmitted out of the session.
num_total_sessions	The total number of out-of-band SSL decryption sessions.
num_ssl30_sessions	The cumulative total number of SSL 3.0 sessions.
num_tls10_sessions	The cumulative total number of TLS 1.0 sessions.
num_tls11_sessions	The cumulative total number of TLS 1.1 sessions.
num_tls12_sessions	The cumulative total number of TLS 1.2 sessions.
num_session_ids	The number of current session IDs.
num_tickets	The number of current TLS tickets.

GigaSMART Operations Statistics Definitions

To view GigaSMART Operations statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics** to open the Statistics page. Figure 29-14 shows an example.

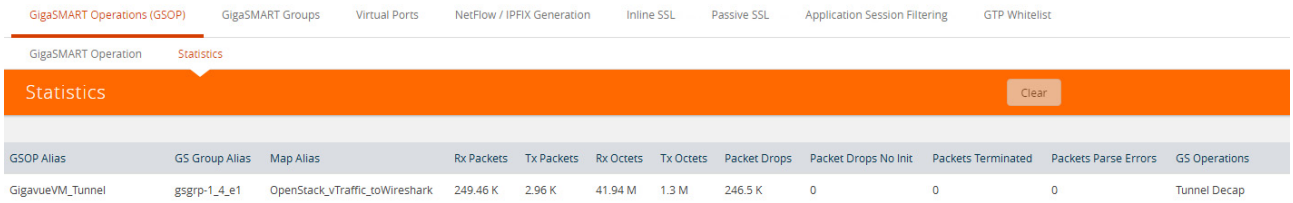

The following table describes the statistics that are common to all GigaSMART operations:

Statistic	Description
Pkts Drop	The cumulative number of packets dropped at a GigaSMART operation for a map.
Pkts Rx	The cumulative number of packets coming into a GigaSMART operation for a map.
Octets Rx	The cumulative number of bytes coming into a GigaSMART operation for a map.
Pkts Term	The cumulative number of packets of a terminated session of a GigaSMART operation for a map.

Statistic	Description
Octets Tx	The cumulative number of bytes going out of a GigaSMART operation for a map.
Pkts Tx	The cumulative number of packets going out of a GigaSMART operation for a map.
Pkts Drop No Init	For internal debugging.
Pkts Parse Err	The cumulative number of packets with invalid or unsupported header types of a GigaSMART operation for a map.

The following sections provide definitions for the statistics that are specific to a particular GigaSMART operation. Refer to the following:

Out-of-Band SSL Decryption Statistics Definitions

The following table describes GigaSMART operations statistics for out-of-band SSL decryption:

Statistic	Description
Sessions Total	The cumulative total number of sessions.
Sessions Active	The number of currently active sessions.

Inline SSL Decryption Statistics Definitions

The statistics for inline SSL decryption are described in *Inline SSL Decryption Guide*.

De-duplication Statistics Definitions

The following table describes GigaSMART operations statistics for de-duplication (including the statistics that are displayed in a cluster environment):

Statistic	Description
Ip 6 Missed Op Busy	For IPv6 packets, the number of non-duplicate packets passed through without storing because the GigaSMART engine was too busy processing other packets. **NOTE:** Applies only to non-duplicate packets.
Non Ip Dupl	The number of non-IPv4 and non-IPv6 duplicate packets detected.
Non Ip	The number of non-IPv4 and non-IPv6 packets received for de-duplication.
Ip 6 Missed Op Space	For IPv6 packets, the number of non-duplicate packets passed through without storing because the GigaSMART engine did not have enough storage space. **NOTE:** Applies only to non-duplicate packets.
Non Ip Missed Op Busy	For non-IPv4 and non-IPv6 packets, the number of non-duplicate packets passed through without storing because the GigaSMART engine was too busy processing other packets. **NOTE:** Applies only to non-duplicate packets.
Non Ip Missed Op Space	For non-IPv4 and non-IPv6 packets, the number of non-duplicate packets passed through without storing because the GigaSMART engine did not have enough storage space. **NOTE:** Applies only to non-duplicate packets.
Ip 4 Missed Op Busy	For IPv4 packets, the number of non-duplicate packets passed through without storing because the GigaSMART engine was too busy processing other packets. **NOTE:** Applies only to non-duplicate packets.
Ip 4 Missed Op Space	For IPv4 packets, the number of non-duplicate packets passed through without storing because the GigaSMART engine did not have enough storage space. **NOTE:** Applies only to non-duplicate packets.
Ip 4 Ipid Tcp Rst	The number of TCP RESETs for IPv4 plus TCP packets. (IPv4 plus TCP RESET packets are not de-duplicated.)
Ip 6 Dupl	The number of IPv6 duplicate packets detected.
Ip 4 Dupl	The number of IPv4 duplicate packets detected.

ERSPAN Statistics Definitions

The following table describes GigaSMART operations statistics for ERSPAN tunnel:

Statistic	Description
Drop Type 3 Marker Bad Sig	The number of ERSPAN Type III packets that have a bad marker packet signature. The expected marker packet signature is 0xa5a5a5a5.
Pkts Too Big	The number of ERSPAN packets that are larger than 9600 bytes after the timestamp trailer is added. These packets will be dropped.
Pkts Rx Type 3	The total number of ERSPAN Type III packets received.
Pkts Tx Type 3 Marker	The total number of ERSPAN Type III marker packets received. (These packets are not forwarded.)
Pkts Rx Type 2	The total number of ERSPAN Type II packets received.
Type 3 Marker Overdue	The number of ERSPAN Type III marker packets that are overdue. Based on granularity, if a marker packet does not arrive by the time specified, it is considered overdue. Refer to *ERSPAN Granularity* on page 846.
pkts Rx	The total number of ERSPAN packets received into the IP interface.
Drop Unknown Proto	The number of packets dropped because they were not recognized as ERSPAN packets or ERSPAN marker packets.
Drop Id No Hit	The number of ERSPAN packets with a wrong flow ID/ERSPAN ID.

Tunnel Decapsulation Statistics Definitions

The following table describes GigaSMART operations statistics for tunnel decapsulation:

Statistic	Description
GMIP Tunnel	
Drop Wrong Addr	The number of GMIP tunneled packets dropped whose destination UDP port does not match the configured value.
Drop Other	The number of GMIP tunneled packets dropped because of fragmented packets.
Pkts Rx	The number of packets received into the IP interface.
Reassemble Rx Success	The number of incoming packets successfully reassembled.
Drop Unknown Proto	The number of packets through the network IP interface dropped if they are neither IPv4 or UDP packets.
Reassemble Rx	The number of incoming packets to be reassembled.
Drop Reassemble Mem Full	The number of packets not successfully reassembled and dropped due to lack of memory for reassembly.
Sliced Mtu	Not valid for tunnel decapsulation.

Statistic	Description
Drop Reassemble Max Hold	The number of packets not successfully reassembled and dropped due to the arrival of too many fragments not yet reassembled. The maximum fragment sessions are already held for reassembly.
Drop No Arp	The number of packets dropped because ARP was not resolved on the IP interface (in particular, on a tool IP interface).
Drop Reassemble Overlap Frag	The number of packets not successfully reassembled and dropped due to overlapping reassembly fragments.
L2GRE Tunnel	
Drop Key Mismatch	The number of packets dropped due to key mismatch.
Pkts Rx	The number of packets received into the IP interface.
Reassemble Rx Success	The number of incoming packets successfully reassembled.
Drop Unknown Proto	The number of packets dropped due to an unknown protocol.
Reassemble Rx	The number of incoming packets to be reassembled.
Drop Reassemble Mem Full	The number of packets not successfully reassembled and dropped due to lack of memory for reassembly.
Drop Reassemble Max Hold	The number of packets not successfully reassembled and dropped due to the arrival of too many fragments not yet reassembled. The maximum fragment sessions are already held for reassembly.
Drop Other	The number of packets dropped due to other reasons.
Drop Reassemble Overlap Frag	The number of packets not successfully reassembled and dropped due to overlapping reassembly fragments.

Tunnel Encapsulation Statistics Definitions

The following table describes GigaSMART operations statistics for tunnel encapsulation:

Statistic	Description
GMIP Encap Tunnel	
Fragment Rx	The number of incoming packets to be fragmented.
Drop Other	Not valid for tunnel encapsulation.
Pkts Rx	The number of packets received into the IP interface.
Fragment Tx	The number of outgoing packets sent to the tunnel after fragmentation.
Sliced Mtu	The number of packets that are sliced to the MTU size of the tool IP interface.
Drop No Arp	The number of packets dropped because ARP was not resolved on the IP interface (in particular, on a tool IP interface).

Statistic	Description
L2GRE Encap Tunnel	
Fragment Rx	The number of incoming packets to be fragmented.
Pkts Rx	The number of packets received into the IP interface.
Fragment Tx	The number of outgoing packets sent to the tunnel after fragmentation.
Session Current Total	The number of currently active sessions.
Session Alloc Fail	The number of session allocations that failed.
Session Lookup	The number of lookups in a session.
Session Lookup Success	The number of lookups in a session that were a success.
Session Alloc	The number of sessions allocated.
Session Timedout	The number of sessions that timed out after a configured timer value.
Sliced Mtu	The number of packets that are sliced to the MTU size of the tool IP interface.
Drop No Arp	The number of packets dropped because ARP was not resolved on the IP interface (in particular, on a tool IP interface).
CUSTOM Decap Tunnel	
rx_packets	The number of packets received into the IP interface.
pkts_drop_unknown_protocol	The number of packets dropped due to an unknown protocol.
pkts_drop_portsrc_mismatch	The number of packets dropped due to source port mismatch.
pkts_drop_portdst_mismatch	The number of packets dropped due to destination port mismatch.
pkts_in_reassemble	The number of incoming packets to be reassembled.
pkts_in_reassemble_success	The number of incoming packets successfully reassembled. For example, if four (4) packets are reassembled into one (1), 4 is displayed in this field.
pkts_out_reassembled	The actual number of packets sent out the tool port. For example, if four (4) packets are reassembled into one (1), 1 is displayed in this field.
pkts_drop_reassemble_overlap_frag	The number of packets not successfully reassembled and dropped due to overlapping reassembly fragments.
pkts_drop_reassemble_max_hold	The number of packets not successfully reassembled and dropped due to the arrival of too many fragments not yet reassembled. The maximum fragment sessions are already held for reassembly.
pkts_drop_reassemble_mem_full	The number of packets not successfully reassembled and dropped due to lack of memory for reassembly.
pkts_drop_reassemble_timed_out	The number of packets dropped due to timeout in the reassembly queue.
VXLAN Decap Tunnel	

Statistic	Description
rx_packets	The number of packets received into the IP interface.
pkts_drop_unknown_protocol	The number of packets dropped due to an unknown protocol.
pkts_drop_portsrc_mismatch	The number of packets dropped due to source port mismatch.
pkts_drop_portdst_mismatch	The number of packets dropped due to destination port mismatch.
pkts_in_reassemble	The number of incoming packets to be reassembled.
pkts_in_reassemble_success	The number of incoming packets successfully reassembled. For example, if four (4) packets are reassembled into one (1), 4 is displayed in this field.
pkts_out_reassembled	The actual number of packets sent out the tool port. For example, if four (4) packets are reassembled into one (1), 1 is displayed in this field.
pkts_drop_reassemble_overlap_frag	The number of packets not successfully reassembled and dropped due to overlapping reassembly fragments.
pkts_drop_reassemble_max_hold	The number of packets not successfully reassembled and dropped due to the arrival of too many fragments not yet reassembled. The maximum fragment sessions are already held for reassembly.
pkts_drop_reassemble_mem_full	The number of packets not successfully reassembled and dropped due to lack of memory for reassembly.
pkts_drop_reassemble_timed_out	The number of packets dropped due to timeout in the reassembly queue.

APF Statistics Definitions

The following table describes GigaSMART operations statistics for APF:

Statistic	Description
Apf Drop	The number of packets matching the GigaSMART drop rules.
Apf Pass	The number of packets matching the GigaSMART pass rules.
Rule Not Match	The number of packets not matching any GigaSMART rules in the map.
Masking Err	The number of masking errors in the map. This number is usually zero (0), which indicates no issues with the masking operation. If the number is non-zero, it indicates there is some issue with the masking operation.

ASF Statistics Definitions

The following table describes GigaSMART operations statistics for ASF:

Statistic	Description
Session Created	The number of sessions created.

Statistic	Description
Session Deleted	The number of sessions deleted.
Session Timeout	The number of sessions that were deleted due to inactivity (expiry of the session timer).
Session Matched (pkt)	The number of incoming packets matching ASF sessions. This count does not include the packets that triggered the creation of the session.
Exceed Count Drop	The number of packets dropped, even if they matched a flow session, because a packet-count was configured and exceeded.

Masking Statistics Definitions

The following table describes GigaSMART operations statistics for masking:

Statistic	Description
No Header	The number of packets with no configured masking protocol.
Too Short	The number of packets with a length less than the masking offset.

Slicing Statistics Definitions

The following table describes GigaSMART operations statistics for slicing:

Statistic	Description
No Header	The number of packets with no configured header for slicing.
Too Short	The number of packets with a length less than the slicing length.
Min Len	The number of packets that are sliced to less than 64 bytes.

Header Stripping Statistics Definitions

The following table describes GigaSMART operations statistics for header stripping:

Statistic	Description
Id No Hit	The number of packets that do not match the configured VXLAN ID to be stripped.
Fm 6000 Pkts Ts	The number of packets received with the FM6000 timestamp.
Unknown Next	The number of packets not stripped of their configured header type as the packets will have an unknown header after the header is stripped.
No Header	The number of packets with no configured header type to be stripped.
Fm 6000 Data Pkt Too Big	The number of FM6000 packets that are larger than 9600 bytes after the timestamp trailer is added. These packets will be dropped.

Statistic	Description
Fm 6000 Keyframe Overdue	The number of FM6000 key frames that are overdue. By default, the key frame rate is 1 packet per second. If a key frame is not received in 1 second, this statistic is incremented.
Fm 6000 Keyframe	The number of key frames received from the FM6000 device. Key frames are marker packets that carry information to convert the FM6000 timestamp to UTC time.

Generic Header Stripping Statistics Definitions

The following table describes GigaSMART operations statistics for generic header stripping:

Statistic	Description
first_anchor_header_not_present	The number of packets with no configured first anchor header.
offset_beyond_anchor_header_size	The number of packets that are not stripped as the configured offset size exceeds the size of the first anchor header.
second_anchor_header_not_present	The number of packets with no configured second anchor header.
strip_length_beyond_pkt_len	The number of packets that are not stripped as the configured strip length exceeds the packet length.
strip_success	The number of packets that are successfully stripped.
incompatible_anchor_headers	The number of packets that are dropped as the first anchor header is incompatible with the second anchor header after the stripping operation is complete.

Trailer Statistics Definitions

The following table describes GigaSMART operations statistics for trailers:

Statistic	Description
pkts_too_big	The number of packets for which the size of the packet has become greater than the maximum supported size (9600) when adding the trailer. This count is incremented with one trailer added to the packet.

FlowVUE Statistics Definitions

The following table describes GigaSMART operations statistics for FlowVUE:

Statistic	Description
Dev Ip Src Match	The number of packets with a source IP matching the range defined in the GigaSMART parameters (gsparams).
Exceed License Warn	The number flows that exceed the installed license.
Exceed License Err	The number of flows that exceed the installed license error limit. (The limit is the number of sessions in the license plus 5%.)

Statistic	Description
Dev Ip No Match	The number of packets with a source and destination IP that does not match any of the ranges defined in the GigaSMART parameters (gsparams).
Dev Ip Non Ip	The number of packets with no IP headers (and hence, not sampled).
Dev Ip Dst Match	The number of packets with a destination IP matching the range defined in the GigaSMART parameters (gsparams).
Out of Resource	The number of times there was a failure to allocate resources for recording a new flow. **NOTE:** The maximum supported flows for each engine port is 2 million.
Dev Ip Drop Not In Sample	The number of packets dropped by the sampling application, not because of errors, but because the flow was sampled to be dropped.

NetFlow Statistics Definitions

The following table describes GigaSMART operations statistics for NetFlow:

Statistic	Description
Out of Resource	The number of times there was a failure to allocate resources for recording a new flow. **NOTE:** The maximum supported flows for each engine port is 2 million.
Non Ip	The number of packets received that are not IPv4 or IPv6.
Non Configured	Packets received when NetFlow is not enabled in the GigaSMART group.
Ssl Active Sessions	The number of currently active SSL sessions monitored by NetFlow.
Total Ssl Sessions	The cumulative total number of SSL sessions monitored by NetFlow.

Display GigaSMART Application Resource Usage

GigaSMART applications, such as De-duplication and inline and out-of-band SSL decryption, use memory resources on the GigaSMART line card or module. As new GigaSMART applications are configured, the total resources on the GigaSMART line card or module can become fully used.

Starting in software version 4.4, you can display the GigaSMART application resource usage, which provides information about the applications that use resources. With this

information, you can choose to free up resources on one application to use them on another.

Table 29-2: GigaSMART Application Resource Information

Name	Format
GSgroup	The alias of the GigaSMART group associated with the GigaSMART application.
Application	The list of licensed GigaSMART applications that use resources, including Adaptive Packet Filtering (APF), de-duplication, GPRS Tunneling Protocol (GTP), NetFlow Generation, out-of-band SSL decryption, Adaptive Session Filtering (ASF) with buffering, and inline SSL decryption. Other GigaSMART applications (such as flow sampling, header addition, header stripping, ERSPAN tunnel decapsulation, slicing, masking, and others) do not have databases that store data, therefore they do not use resources and are not displayed with the **show gsgroup gsapp-resource** command.
% of Total	The percentage of the total amount of memory used by each GigaSMART application.
Configured Resource	The amount of resources configured for each GigaSMART application. The valid values are as follows: • app-max—Indicates the maximum amount of memory configured for the application. It is a pre-allocated amount. • integer, such as 2—Indicates the number of sessions configured, in the units specified. M indicates millions.
Installed Resource	The amount of resources installed for each GigaSMART application.
Licensed Quantity	The amount of resources licensed for each GigaSMART application.
Units	The units, such as sessions, or millions (M) of sessions.

Overview of GigaSMART Application Resources

GigaSMART application resources are managed per GigaSMART group (gsgroup). A gsgroup can be configured with one or more GigaSMART engine ports on one or more GigaSMART line cards or modules.

For most GigaSMART applications, resources are allocated automatically, based on configuration. For some GigaSMART applications, resources are allocated when they are configured in the gsgroup. For other applications, resources are allocated when a GigaSMART operation (gsop) using the application is created. For buffer ASF however, you can explicitly configure resources, in millions of sessions.

The allocation of resources for a new application will be successful if the application is licensed and if there is sufficient space for the new application.

If there is insufficient space, then the resources need to be managed to free up memory for the new application. Managing resources includes deleting applications that are no longer used.

However, deleting a GigaSMART application does not result in the immediate deletion of application resources. Once a resource has been allocated, it remains allocated. To delete resources for APF, out-of-band or inline SSL decryption, de-duplication, and

GTP, remove the configuration related to the application, then reload the GigaSMART line card or module.

To remove the configuration related to an application, delete the gsop first, then delete the gsgroup. If the gsop or gsgroup is bound to a map, you will also have to delete the map.

Resources for Buffer ASF

The resources for buffer ASF depend on the number of sessions and the type of node. For example, on GigaVUE-HC2, 5 million sessions uses 62% of total resources, but on GigaVUE-HC3, 5 million sessions uses 26%.

For GigaVUE-HC3, the resources for buffer ASF for the number of sessions is as follows:

- 2 million—11% of total resources
- 3 million—16% of total resources
- 4 million—21% of total resources
- 5 million—26% of total resources

Reload GigaSMART Line Card or Module

Occasionally, the GigaSMART line card or module will need to be reloaded for changes to take effect and to allocate resources accordingly. Reloading also provides applications with contiguous memory.

The following message displays at the bottom of the output of the **show gsgroup gsapp-resource** command when the GigaSMART line card or module needs to be reloaded:

```
*Resource allocation changes have been made that require GigaSMART card 2/1/1 to be
reloaded in order for them to take effect.
```

When this message is displayed, you cannot change the configuration relating to that application until after the reload. For example, you cannot use the gsop, associated with the gsgroup, in a map.

Use the following command to reload a GigaSMART line card or module:

(config) # card slot <slot ID> down

Use the following command to bring the GigaSMART line card or module back up:

(config) # no card slot <slot ID> down

GigaSMART CPU Utilization Statistics

You can display CPU utilization statistics for GigaSMART. The statistics indicate the performance of GigaSMART, improve visibility, and help identify high load conditions.

Show commands display instantaneous CPU utilization as well as historical, providing trends for CPU utilization.

You can also configure a rising threshold, as a percentage, to indicate when high CPU utilization occurs. When the aggregate CPU utilization percentage exceeds the rising the threshold, an SNMP notification can be triggered.

This feature is supported on all products that support GigaSMART: GigaVUE-HB1, GigaVUE-HC1, GigaVUE-HC2, GigaVUE-HC3, and GigaVUE-HD8/GigaVUE-HD4.

The GigaSMART engine port (**e** port) numbers are e1 on nodes with one GigaSMART engine and e1 and e2 on nodes with two GigaSMART engines, for example: 10/1/e1 or 8/1/e1 and 8/1/e2.

Refer to the following sections for viewing statistics, configuring the threshold, and configuring a notification that can be sent when the threshold is exceeded:

- *Display GigaSMART CPU Utilization* on page 806
- *Configure Threshold* on page 807
- *Configure Threshold Crossing Notification* on page 807

Display GigaSMART CPU Utilization

Use the **show gsgroup stats all** command to display the statistics on all GigaSMART groups on the node.

The statistics are displayed for 1 second, 1 minute, 5 minute, 10 minute, and 15 minute intervals. The 1 second interval displays the statistics for the previous second. The 1 minute, 5 minute, 10 minute, and 15 minute intervals display statistics containing history.

Statistics are displayed in an aggregate form. For example, if there are two GigaSMART **e** ports: e1 and e2, there will be two aggregates. One aggregate will be for e1, the other will be for e2. The term aggregate refers to aggregation across all packet processing cores (up to 31) in the CPU. It does not refer to an aggregate across CPUs.

The statistics are as follows:

- Useful Time—Amount of time during which the CPU is processing packets, in milliseconds (ms) or seconds (s).
- Idle Time—Amount of time during which the CPU is not processing packets, for example, when it is busy looping, in milliseconds (ms) or seconds (s).
- In Packets (pkts/s)—Number of packets per second coming into the CPU. For the 1 second interval, In Packets is the actual number of incoming packets for that second. For the 1 minute, 5 minute, 10 minute, and 15 minute intervals, In Packets is an average number of incoming packets per second.

- Packets Drop (pkts/s)—Number of packets per second dropped by the CPU. For the 1 second interval, Packets Drop is the actual number of dropped packets for that second. For the 1 minute, 5 minute, 10 minute, and 15 minute intervals, Packets Drop is an average number of dropped packets per second.
- Packets Recv Error—Number of received packets per second with errors. For the 1 second interval, Packets Recv Error is the actual number of errored packets for that second. For the 1 minute, 5 minute, 10 minute, and 15 minute intervals, Packets Recv Error is an average number of errored packets per second.
- CPU Usage %—Percentage of time during which the CPU is processing packets. CPU Usage % plus CPU Idle % equals 100.
- CPU Idle %—Percentage of time during which the CPU is not processing packets. CPU Idle % plus CPU Usage % equals 100.

NOTE: When the node is restarted, the 1 minute, 5 minute, 10 minute, and 15 minute statistics will not be exactly for 1 minute, 5 minute, 10 minute, and 15 minute intervals, until the full interval has elapsed and the history is available.

Configure Threshold

An upper threshold (rising) can be configured. When the aggregate value of the CPU utilization on the GigaSMART engine exceeds the threshold, an SNMP notification can be triggered. Refer to *Configure Threshold Crossing Notification* on page 807.

Configure Threshold Crossing Notification

When the aggregate value of the CPU utilization exceeds the upper (rising) threshold, a message is logged, and optionally, an SNMP notification is sent to all configured destinations.

When enabled, the SNMP notification is sent when the rising threshold is exceeded in any 5-second interval over which the CPU utilization is averaged.

NOTE: Once the rising threshold is exceeded for 5 seconds, the SNMP notification is generated. However, if the CPU utilization falls below the upper threshold but does not remain below that threshold continuously for 5 seconds, a new notification is not generated when the upper threshold is exceeded again. A new notification is generated only when the CPU utilization falls below the threshold, stays below the threshold continuously for 5 seconds, exceeds the threshold again, and stays above it for 5 seconds.

30 How to Use GigaSMART Operations

Use the **GigaSMART Operations** page to create GigaSMART operations. GigaSMART operations consist of a name and a supported combination of the available GigaSMART applications that you have licensed.

- Refer to *How to Combine GigaSMART Operations* on page 778 for details on supported combinations of GigaSMART operations.
- Refer to *Order of GigaSMART Operations* on page 781 for information on the order in which GigaSMART components are applied in a single operation.

The details of each GigaSMART operation are described in the following sections:

GigaSMART Masking

Required License: Base

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 30-1.

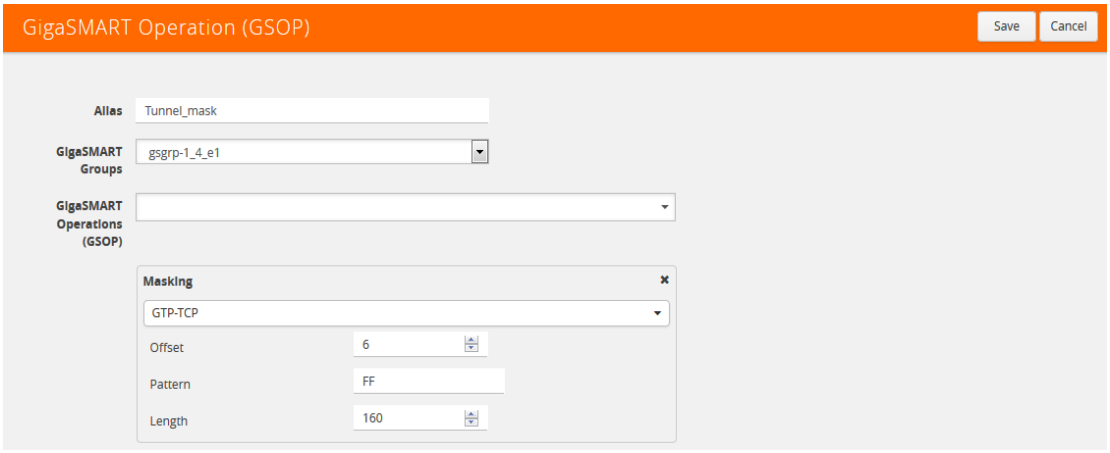

Figure 30-1: GigaSMART Operations Page with Masking Selected

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: • **IPV4** – Mask starting a specified number of bytes after the IPv4 header. • **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.
GigaSMART Trailer	Masking operations can optionally include the GigaSMART Trailer. If you do elect to include the GigaSMART Trailer, it will include the original packet length. Refer to *GigaSMART Trailer Reference* on page 1196 for details.
In Combination with Slicing	Masking operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to *Groups of GigaSMART Engine Ports* on page 757 for details.

Examples – GigaSMART Masking

The example shown in Figure 30-2 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.

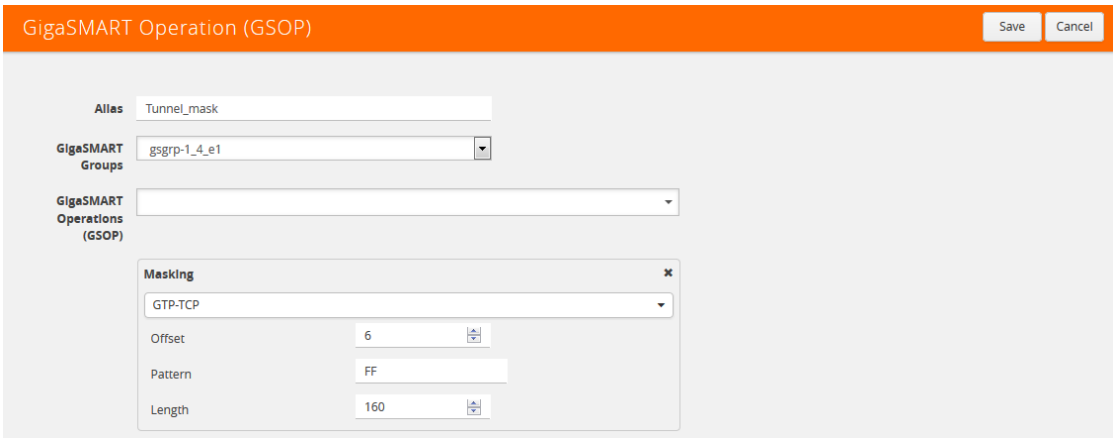

Figure 30-2: GigaSMART Masking Operation

This example shown in Figure 30-3 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.

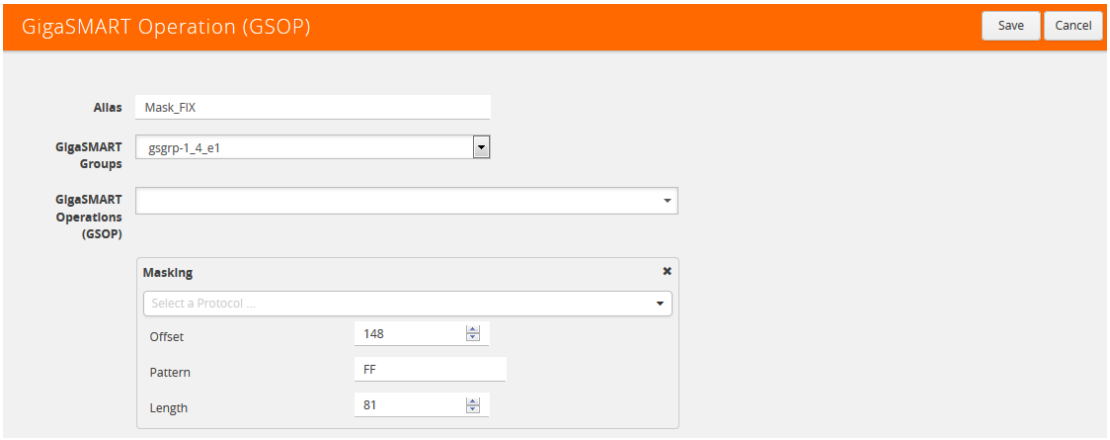

Figure 30-3: GigaSMART Operation with a Static Offset

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* on page 801 for descriptions of the masking statistics as well as to *GigaSMART Operations Statistics Definitions* on page 794.

GigaSMART Packet Slicing

Required License: Base

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: • **IPV4** – Slice starting a specified number of bytes after the IPv4 header. • **IPV6** – Slice starting a specified number of bytes after the IPv6 header. • **UDP** – Slice starting a specified number of bytes after the UDP header. • **TCP** – Slice starting a specified number of bytes after the TCP header. • **FTP** – Identify using TCP port 20 and slice payloads using offset from the TCP header. • **HTTPS** – Identify using TCP port 443. Slice encrypted payloads using offset from the TCP header. • **SSH** – Identify using TCP port 22. Slice encrypted payloads using offset from the TCP header. The GigaSMART can provide slicing for GTP tunnels, provided the user payloads are 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. • **GTP** – Slice starting a specified number of bytes after the outer GTP header. • **GTP-IPV4** – Slice starting a specified number of bytes after the IPv4 header inside the encapsulating GTP packet. • **GTP-UDP** – Slice starting a specified number of bytes after the UDP header inside the encapsulating GTP packet. • **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:

Feature	Description
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.
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 *GigaSMART Trailers* on page 1191 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* on page 757 for details.

Example – GigaSMART Packet Slicing

This example shown in 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.

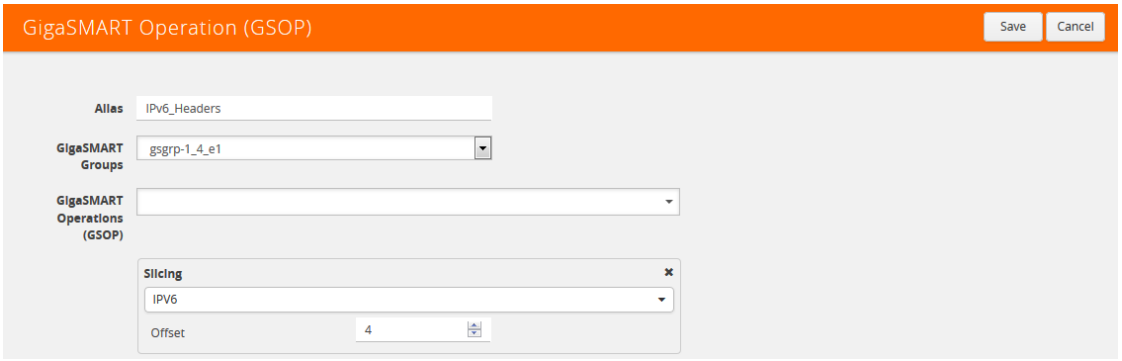

Figure 30-4: GigaSMART Operations (GSOP) Page with Slicing Selected

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* on page 801 for descriptions of these statistics as well as to *GigaSMART Operations Statistics Definitions* on page 794.

GigaSMART IP Encapsulation/Decapsulation (GigaSMART Tunnel)

Required License for IP Decapsulation: Base (GigaVUE-HC2, and GigaVUE-HC3), Tunneling (GigaVUE-HC1 and GigaVUE-HB1)
Required License for IP Encapsulation: Advanced Tunneling (GigaVUE-HC2, and GigaVUE-HC3), Tunneling (GigaVUE-HC1 and GigaVUE-HB1)

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.

IP fragmentation and reassembly are supported. Refer to *IP Fragmentation and Reassembly on L2GRE and GMIP Tunnels* on page 843.

The source of the GigaSMART tunnel can be any of the following:

- **GigaSMART-Enabled GigaVUE H Series Node**
  - Standalone GigaVUE-HC3 node with SMT-HC3-C05 modules installed.
  - Standalone GigaVUE-HC2 node with GigaSMART-HC0 front and/or rear modules installed.
  - Standalone GigaVUE-HC1 nodes.
  - Any GigaVUE H Series node operating in a cluster with the previous node types.
- **GigaVUE-2404 G Series node with a GigaSMART-6X line card installed**
- **GigaVUE V Series node or a GigaVUE-VM**

NOTE: You can also create GigaSMART operations that allow a GigaVUE H Series node to act as the receiving end of an ERSPAN tunnel for data mirrored over the Internet from Cisco equipment. However, this feature requires the Advanced Tunneling license; refer to *GigaSMART ERSPAN Tunnel Decapsulation* on page 845.

Configure Both Ends of GigaSMART Tunnel

Creating a GigaSMART tunnel requires configuration on both the sending and receiving ends:

Sending End of Tunnel	Receiving End of Tunnel
The sending end of a GigaSMART tunnel can be either a GigaVUE-VM deployment or a GigaSMART-enabled GigaVUE H Series or G Series node. **Sending Data from a GigaSMART-Enabled GigaVUE H Series Node** • Configure an IP interface with an IP address, subnet mask, default gateway, MTU setting and assign it to a GigaSMART group. • Create a GigaSMART operation with a tunnel-encap component. The encapsulation settings include the IP address and listening UDP port of the P interface that is associated with a network port on the destination GigaVUE H Series. • 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: • **UDP IP** – The IP address of the P interface that is associated with a network port on the receiving end of the tunnel. • **UDP Source Port** – The source port from which traffic will be sent to the receiving end of the GigaSMART tunnel. • **UDP Destination Port** – The listening UDP port at the destination end of the GigaSMART tunnel. **Sending Data from GigaVUE-2404/GigaSMART-6X** • Configure an IP interface with an IP address, subnet mask, default gateway, and MTU setting. Associate the IP interface with a tool port. • Create a GigaSMART operation with an encapsulation component. The encapsulation settings include the IP address and listening UDP port of the IP interface thst is associated with network port on the destination GigaVUE G Series. • 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.	• 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. • 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. • 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.

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* on page 782 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 H 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 only be assigned to GigaSMART groups consisting of a single GigaSMART engine port. Refer to *Groups of GigaSMART Engine Ports* on page 757 for more information.

Example: GigaSMART Encapsulation/Decapsulation (GigaVUE-HB1 Node)

The following figures demonstrate how to create a sample IP tunnel between a sending GigaVUE-HB1 in Reno and a receiving GigaVUE H Series cluster in San Francisco. First, the overall tunnel is summarized, followed by configuration descriptions for the sending and receiving ends.

GigaVUE-HB1 in Reno

The GigaVUE-HB1 in Reno takes IPv4 packets arriving at network port 15/1/g1, encapsulates them in UDP packets, and forwards them across the Internet from IP interface 1/1/g10 to the receiving GigaVUE-HD4 in San Francisco.

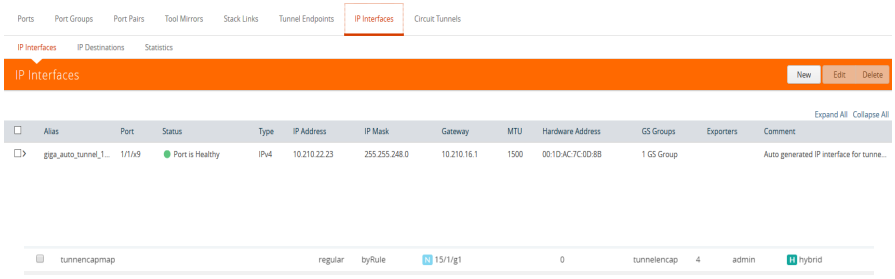

Set up the IP interface in the **Port > IP Interfaces** page. For example, 1/1/x9.

The **tunnelencap** GigaSMART operation bound as part of a map on network port 15/1/g1 specifies the destination IP address and UDP port (the address of the IP interface in San Francisco). Note that this map sends IPv4 packets to the 15/1/g1 IP interface configured in the **Port > IP Interfaces** page.

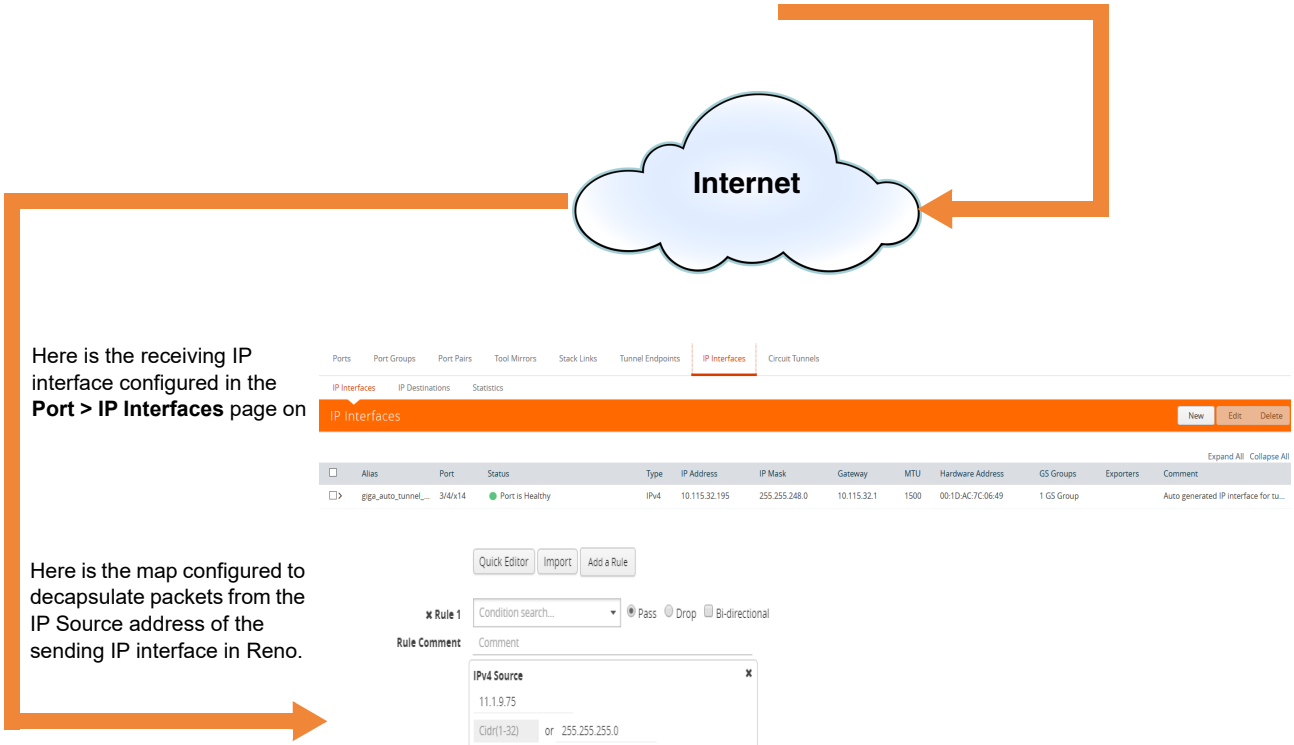

GigaVUE-HD4 with GigaSMART in San Francisco

IP interface 5/1/g2 in San Francisco decapsulates traffic arriving from Reno and makes it available for standard GigaVUE packet distribution to local tool ports. The **decapper1** GigaSMART operation bound to IP interface with network port 5/1/g2 with a map rule specifies the listening UDP port matching the one configured in Reno. It is sending traffic from IP address 11.1.9.75 to local tool port 1/1/g5 (a port on another node in the cluster).

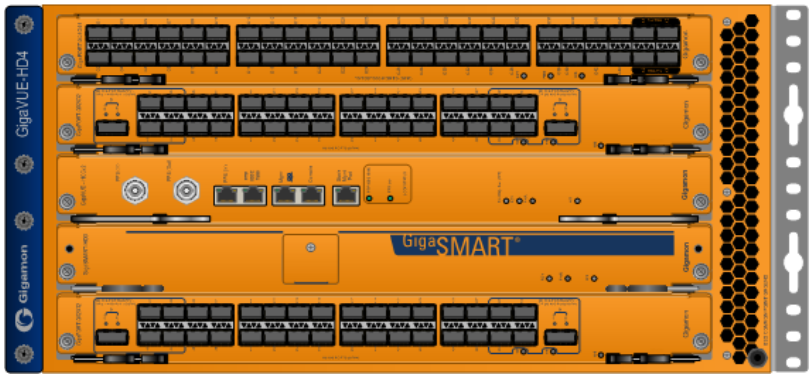

Configure Sending End of Tunnel: GigaVUE-HB1 in Reno

The GigaVUE-HB1 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 CLI:

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.
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: IPv6 • Port Source: 5000 • Port Destination: 10000 • Destination IP: 21.2.9.75 • DSCP: 0 • Precision: 1 • TTL: 1 **5.** Click Save.

Task	UI Steps
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/g1 • GigaSMART Operations (GSOP): tunnelencap (gsport1) **4.** Click Add Rule. **5.** Select Pass. **6.** Select IP Version for Rule 1. **7.** Select v4 or v6 for Version. **8.** Save.

Configure Receiving End of Tunnel: GigaVUE-HD4 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-HD4 in this location will have an IP interface associated with network port configured on network port 5/1/g2 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/g2 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/g2. **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/g2. This command sets the IP address, subnet mask, default gateway, and MTU for the IP interface associated with network port on port 5/1/g2. 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_g2 • Port: 1/1/g2 • 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/g2 from IP address 21.2.9.25 (the start of the tunnel) and sends it to port 1/1/g5. 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/g2 • Destination: 1/1/g5 • 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.

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* on page 797 and *GigaSMART Operations Statistics Definitions* on page 794 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* on page 798 and *GigaSMART Operations Statistics Definitions* on page 794 for descriptions of these statistics.

Example: GigaSMART Encapsulation/Decapsulation (GigaVUE-VM)

The following figures demonstrate how to create a sample IP tunnel between a sending GigaVUE-VM node in Sydney and a receiving GigaVUE H Series in Melbourne. First, the overall tunnel is summarized, followed by configuration descriptions for the sending and receiving ends.

GigaVUE-VM vMap in Sydney

The GigaVUE-VM vMap in Sydney takes virtual packets matching the criteria specified by the vMap, encapsulates them in UDP packets, and forwards them across the Internet to the receiving GigaVUE H Series in Melbourne.

The vMap includes the destination IP address and UDP port (the address of the IP interface associated with network port in Melbourne, shown in highlighting).

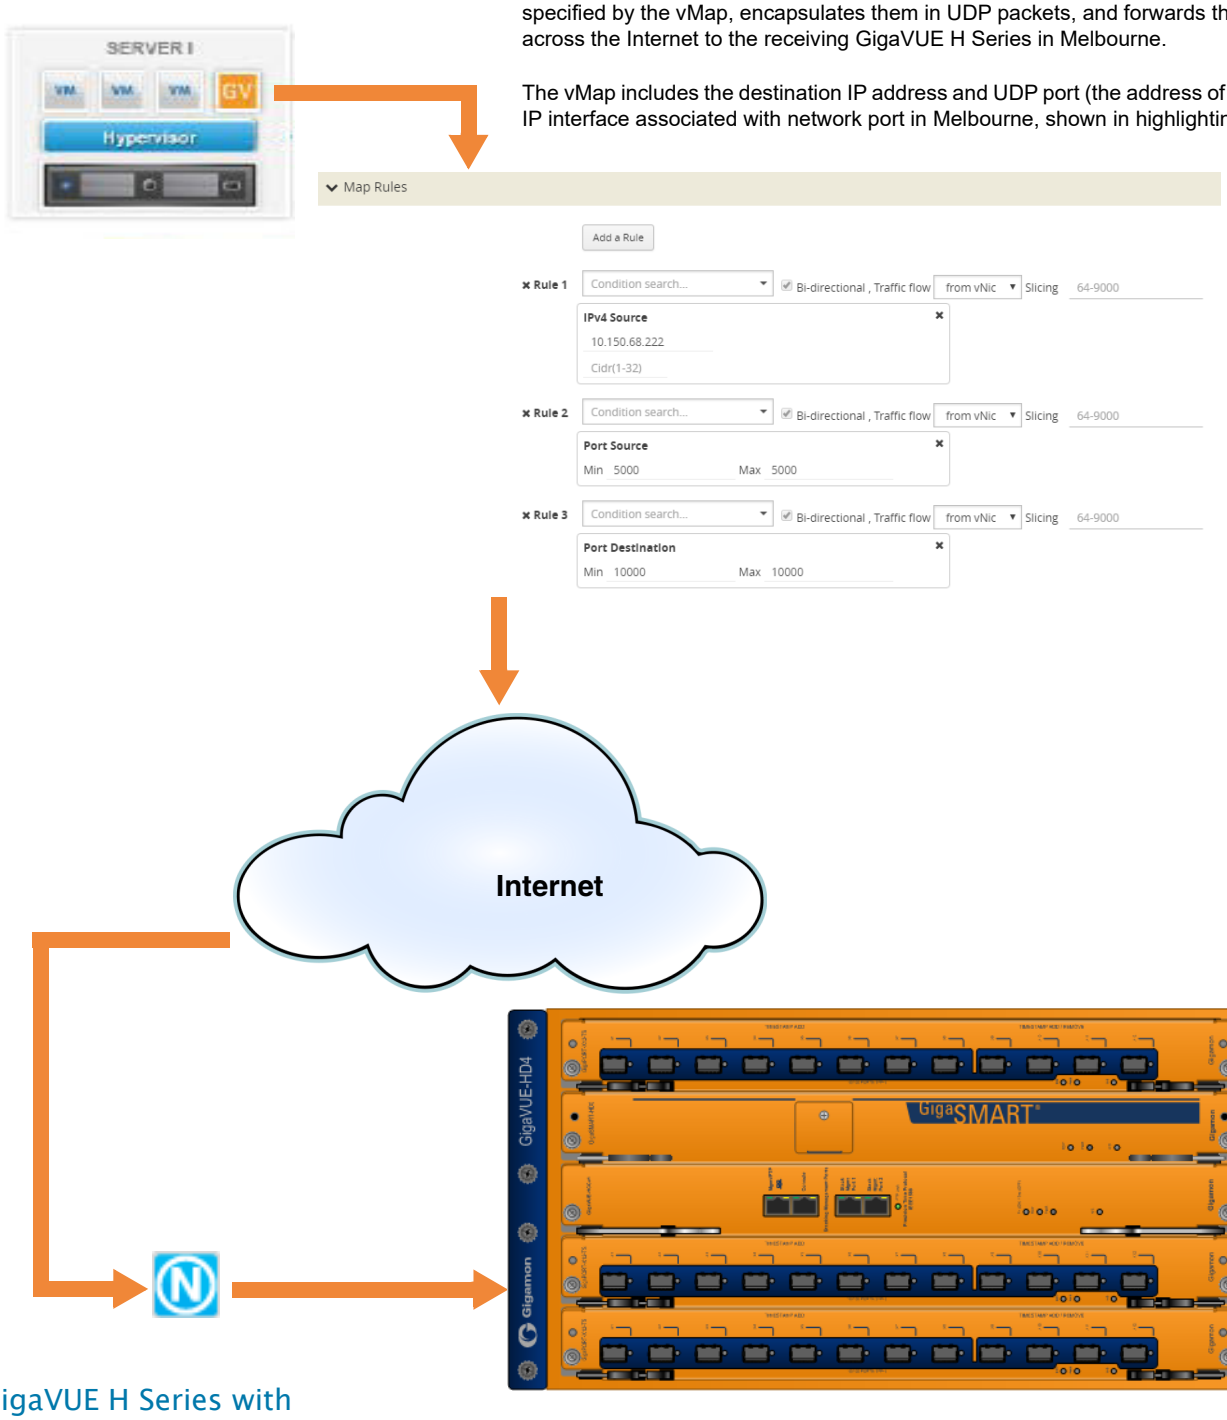

GigaVUE H Series with GigaSMART in Melbourne

IP interface associated with network port 1/1/x3 in Melbourne decapsulates traffic arriving from Sydney and makes it available for standard GigaVUE packet distribution to local tool ports. The decapsulation operation bound to IP interface associated with network port 3 as part of a map specifies the listening UDP port matching the one configured in Sydney.

The GigaSMART decapsulation operation is based on the GigaVUE-HD0 line card in slot 3.

Configure Sending End of Tunnel: GigaVUE-VM vMap in Sydney

A GigaVUE-VM node in this location is configured with a vMap that will send data over the Internet to the IP interface associated with a network port on a GigaVUE H Series with a GigaSMART-HD0 line card installed.

VMaps are created in the GigaVUE-FM user interface – Step 2 in the Create Map wizard includes **Tunnel Traffic To** settings that specify where matching traffic should be sent:

Create vMap "Tunnel Traffic To" Option	Setting
UDP IP	This is the destination IP address for the IP interface associated with network port on the GigaVUE H Series in Melbourne. We will set it to 10.150.68.222
UDP Source Port	This is the UDP source port from which tunneled packets will be sent. We will set this to 5000.
UDP Destination Port	This is the listening port on the receiving GigaVUE H Series IP interface associated with network port. We will set this to 10000.

Configure Receiving End of Tunnel: GigaVUE H Series with GigaSMART in Melbourne

Now we need to configure the receiving end of the tunnel with an IP interface associated with network port. The GigaVUE H Series in this location will have an IP interface associated with network port configured on network port 1/1/3 with an IP address of 10.150.68.222 and a GigaSMART decapsulation operation that listens on UDP port 10000.

The following table summarizes the steps necessary to configure the receiving end of the tunnel using the UI:

Task	UI Steps
Start by designating port 1/1/x3 as an IP interface with network port, configuring its IP profile, and assigning its GigaSMART operations to a GigaSMART group. This command sets the IP address, subnet mask, default gateway, and MTU for the IP interface associated with a tool port on port 1/1/x3.	**1.** Select Ports > IP Interfaces. **2.** Click New. **3.** Configure the IP Interface: • Alias: 1_1_x3 • Port: 1/1/x3 • IP Address: **10.150.68.222** • IP Mask: 255.255.255.255 • Gateway: **10.150.68.1** • MTU: 9400 • GigaSMART Group: GS2 **4.** Save

Task	UI Steps
Now, create an IP decapsulation GigaSMART operation (**gmipdecap**) that will decapsulate traffic received on UDP port 10000. Recall that we configured the sending end of the tunnel to send to that UDP port. The operation has the alias **gv_ipdecap**. Note that this operation uses the same GigaSMART group (GS2) as the IP interface associated with network port we set up in the first step.	**1.** From the device view, select GigaSMART > GigaSMART Operations > GigaSMART Operation. **2.** Click New. **3.** Configure the GigaSMART Operation: • Alias: **gv_ipdecap** • GigaSMART Groups: GS2 • GigaSMART Operations (GSOP): Tunnel Decapsulation **4.** Configure the Tunnel Encapsulation: • GMIP • GMIP Port: 10000 **5.** Save.
Once we have our IP decapsulation operation, we can include it as part of a map. • Open the map configuration page to create a map named **decapper**. • The **Source** field specifies the ingress ports for this map. • The **GSOP** field applies the gv_ipdecap GigaSMART operation to all packets matching the rules in the map, decapsulating them from the tunnel. • The **Destination** field specifies where matching packets will be sent (tool port 1/1/x11). • The rule with **Pass** selected specifies that packets arriving on this port with an IP Source address of 10.10.10.10 /32 will be processed by the gv_ipdecap GSOP and sent to tool port 1/1/x11.	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. • Alias: decapper • Type: Regular • Subtype: By Rule • Source: 1/1/x3 • Destination: 1/1/x11 • GigaSMART Operation (GSOP): gv_ipdecap (GS2) **4.** Click Add Rule. • Select Pass. • Select IPv4 Source for Rule 1. • Set the IPv4 Address to 10.10.10.10 • Set the Net Mask to 255.255.255.255 **5.** Click Save.

GigaSMART IP Encapsulation (GigaSMART Tunnel)

Required License for IP Encapsulation: Advanced Tunneling (GigaVUE-HC2, and GigaVUE-HC3), Tunneling (GigaVUE-HC1 and GigaVUE-HB1)

GigaSMART-enabled nodes with the Advanced Tunneling license installed can encapsulate traffic and send it through a GigaSMART tunnel to a destination GigaSMART-enabled node.

1. Configure an IP interface with an IP address, subnet mask, default gateway, and MTU setting and assign it to a GigaSMART group.
2. Create a GigaSMART operation with a **Tunnel Encapsulation** component. The encapsulation settings include the IP address and listening UDP port of the IP interface associated with network port on the destination GigaVUE H Series.
3. Bind the GigaSMART operation to one or more network ports as part of a map. The network ports must be mapped to the IP interface associated with a tool port.

Refer to the sections beginning with *Configure Both Ends of GigaSMART Tunnel* on page 817 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-HC2, and GigaVUE-HC3), Tunneling (GigaVUE-HC1 and GigaVUE-HB1)
Required License for L2GRE Encapsulation: Advanced Tunneling (GigaVUE-HC2, and GigaVUE-HC3), Tunneling (GigaVUE-HC1 and GigaVUE-HB1)

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* on page 843.

Figure 30-5 on page 829 shows the remote site encapsulating the filtered traffic and routing it to the main office from the remote end.

The encapsulated packet is sent out of the tool port, which is connected to the public network (the Internet). This packet is routed in the public network to reach the main office site. It ingresses at the routed network port of the GigaVUE node at the main office.

The ingress encapsulated packet is then sent to the GigaSMART at the main office, where the packet is decapsulated and sent to the tool port. The received packet's destination IP is checked against the source IP/IPv6 configured for the network port. If they match, decapsulation is applied. The Ethernet + IP + GRE or Ethernet + IPv6 + GRE header is stripped and the remaining packet is sent to the tool port.

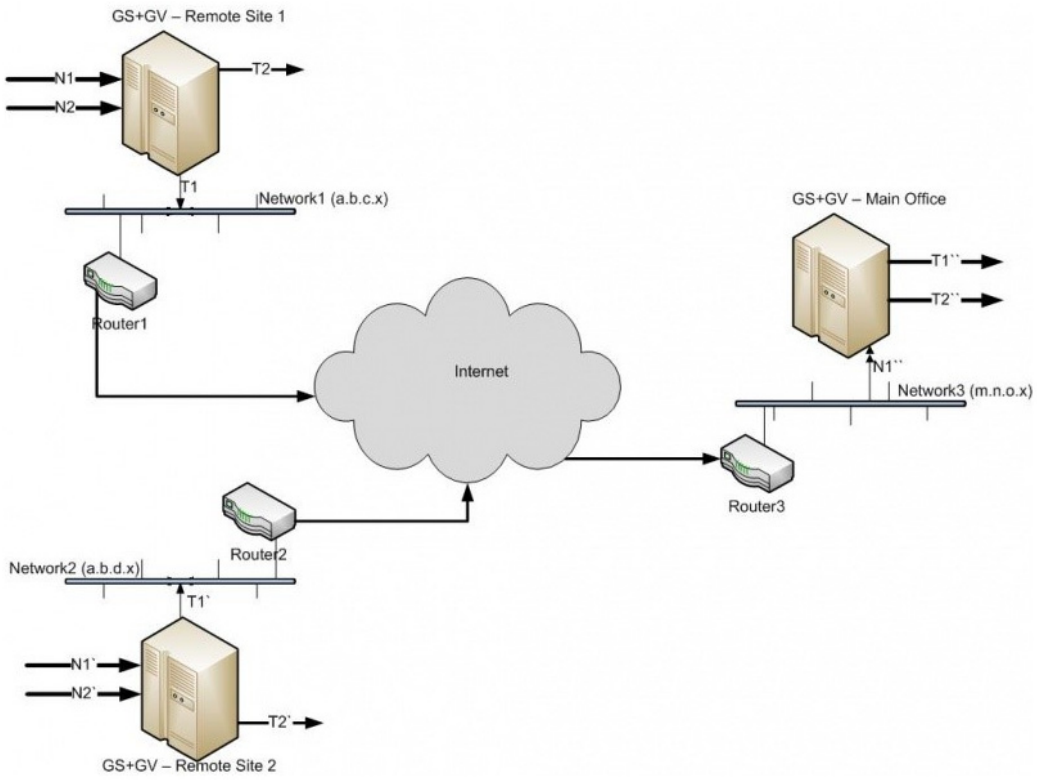

Figure 30-5: L2GRE Tunnel Encapsulation/Decapsulation

For L2GRE tunnel encapsulation/decapsulation configuration examples, refer to *Example 1 – GigaSMART L2GRE Tunnel Encapsulation* on page 832 and *Example 4 – GigaSMART L2GRE Tunnel Decapsulation* on page 836.

For statistics for encapsulated packets, refer to *Display L2GRE Tunnel Encapsulation Statistics* on page 842. For statistics for decapsulated packets, refer to *Display L2GRE Tunnel Decapsulation Statistics* on page 842.

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* on page 1147.

For examples of load balancing on L2GRE encapsulation, refer to *Example 2 – GigaSMART L2GRE Tunnel Encap Stateful LB* on page 833 and *Example 3 – GigaSMART L2GRE Tunnel Encap Stateless LB* on page 836.

L2GRE IPv6 Encapsulation/Decapsulation

Gigasmart L2GRE IPv6 lets you route filtered traffic to the remote end using IPv6-based L2GRE tunneling. At the receiving end, filtered traffic is sent to GigaSMART, which adds an L2GRE header and a IPv6 header to make it routable. The remote end decapsulates the packet and sends it to the tool port.

GigaVUE nodes act as L2GRE encapsulation and decapsulation devices. The IPv6 protocol is used to deliver all packets received in the encap tunnel to the termination node using the configured source and destination IPv6 address. The tunnel termination (decap) 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 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 *Example 5 – GigaSMART L2GRE IPv6 Tunnel Encamp/Decap with Load-Balancing* on page 838

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 H Series device at a remote site traverses through the L2GRE tunnel and reaches the GigaVUE-H 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 H 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 H Series device. After tunnel termination, the packet is presented to the flow mapping module to filter based on map rule parameters.

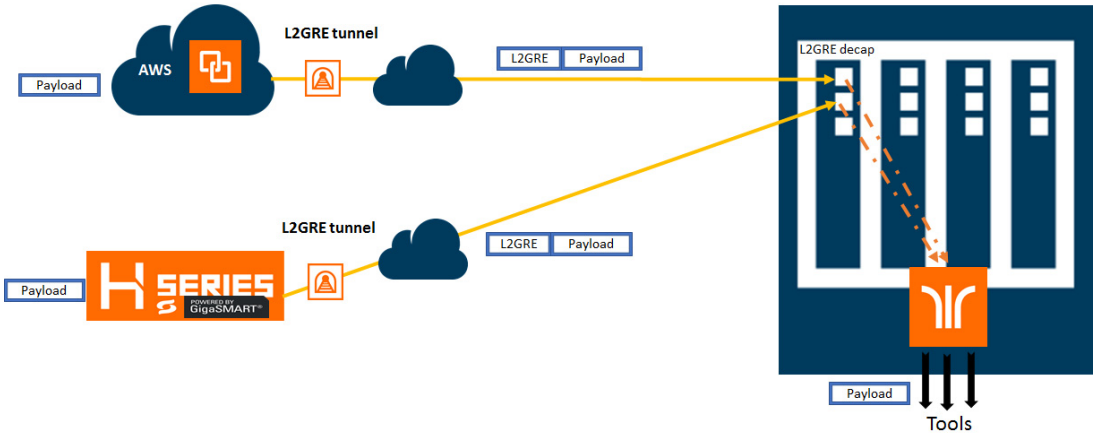

Configure L2GRE Tunnel Encapsulation and Decapsulation

Refer to the following configuration examples:

- *Example 1 – GigaSMART L2GRE Tunnel Encapsulation* on page 832
- *Example 2 – GigaSMART L2GRE Tunnel Encap Stateful LB* on page 833
- *Example 3 – GigaSMART L2GRE Tunnel Encap Stateless LB* on page 836
- *Example 4 – GigaSMART L2GRE Tunnel Decapsulation* on page 836
- *Example 5 – GigaSMART L2GRE IPv6 Tunnel Encamp/Decap with Load-Balancing* on page 838

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* on page 754 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 Comment 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.

Task	Description	UI Steps
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.
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* on page 754 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 Comment 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 30-6 on page 834.

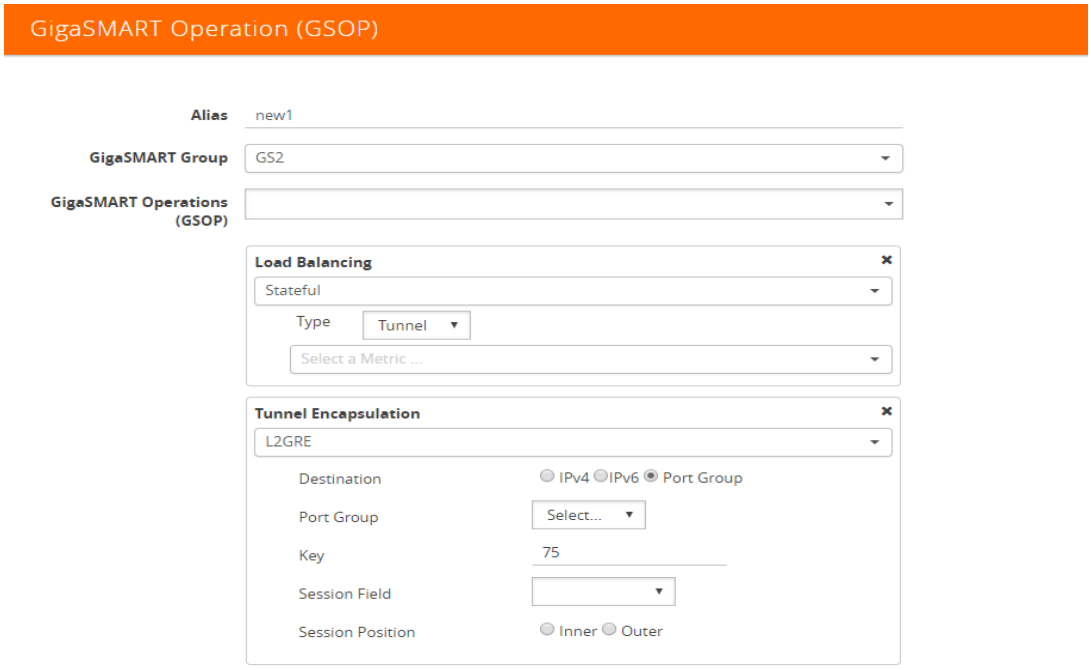

Figure 30-6: New GigaSMART Operation for Stateful Load Balancing

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 30-7 on page 835.

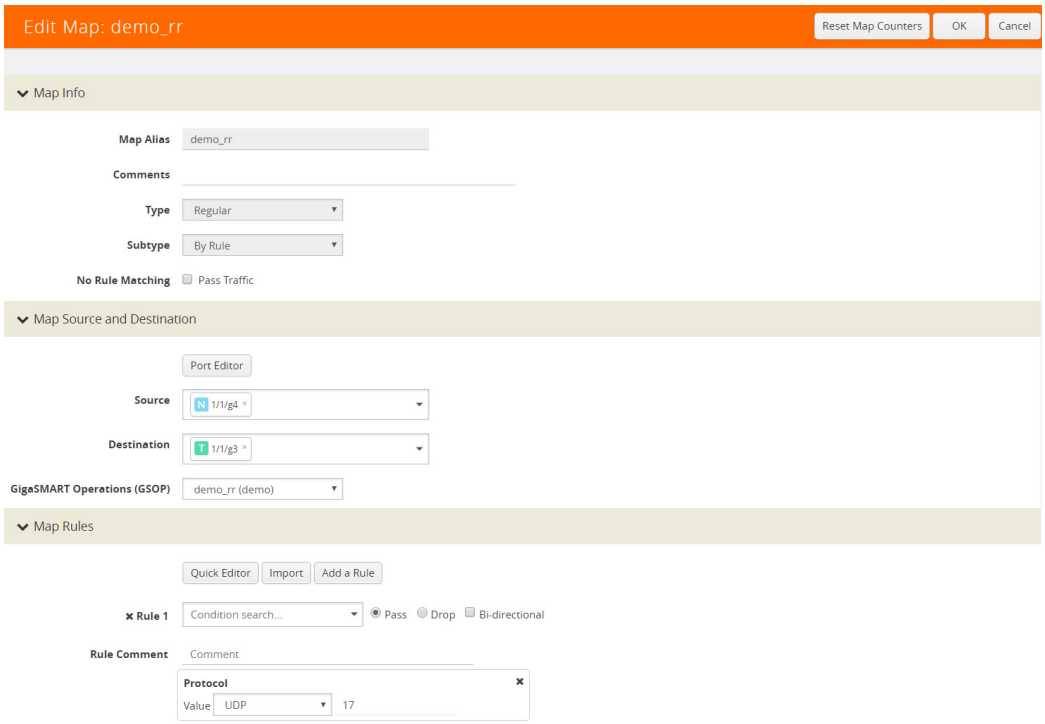

Figure 30-7: New Map Configuration

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 30-8 on page 836.

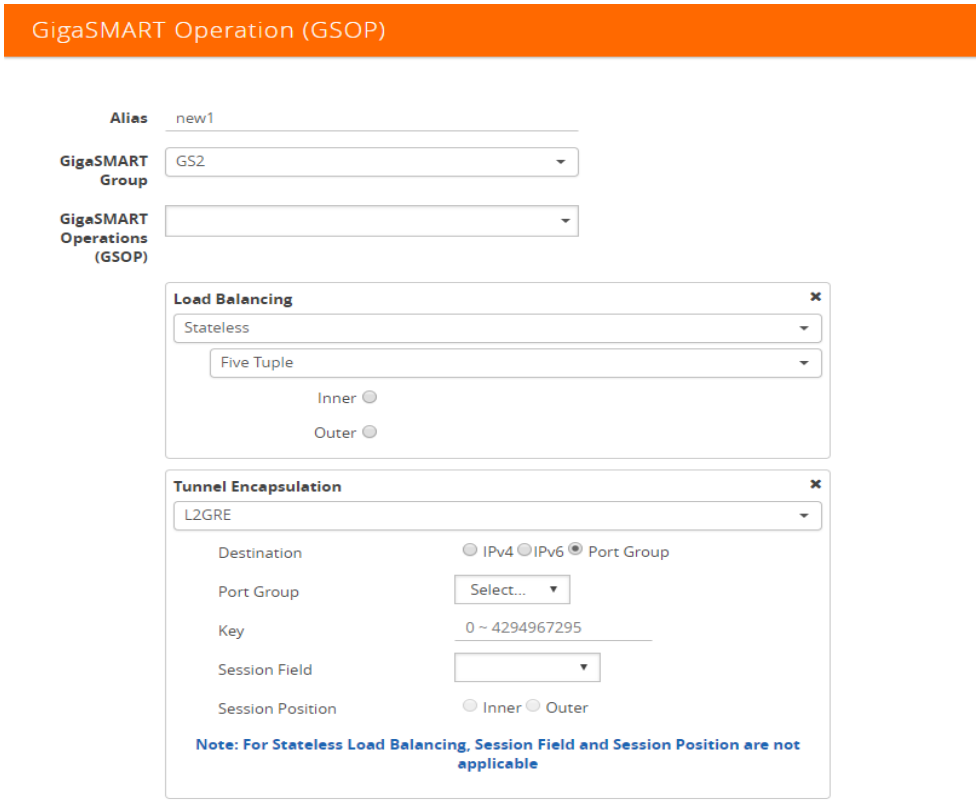

Figure 30-8: New GigaSMART Operation for Stateless

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.
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 Comment 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 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* on page 754 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 **Comment** 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**.

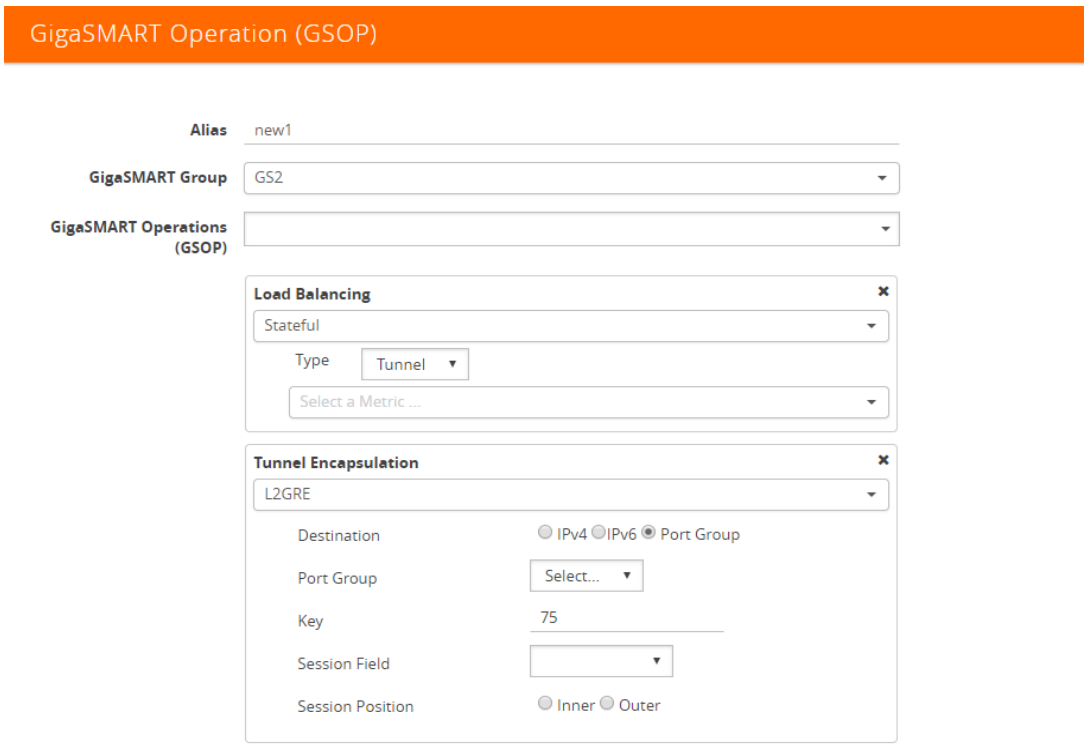

Figure 30-9: GigaSMART Operation (GSOP)

3. In the **Alias** field, enter a name for the GigaSMART operation.
4. From the **GigaSMART Group** drop-down list, select the **GigaSMART Group** 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**.

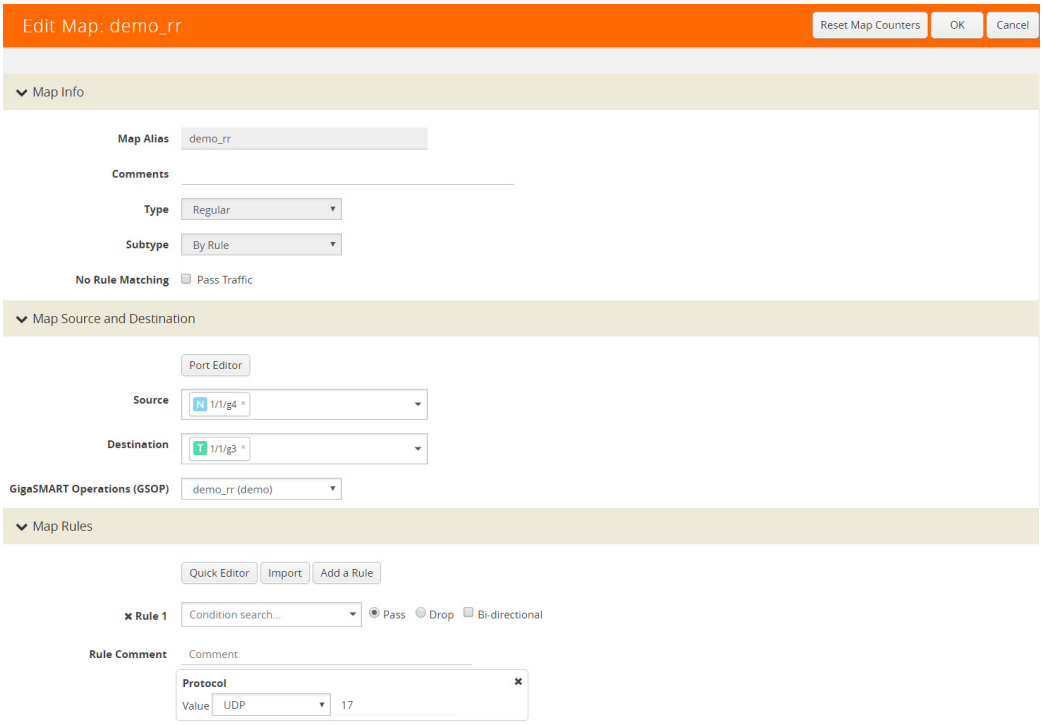

Figure 30-10: Map Configuration

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 **Comment** 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**.

Display L2GRE IPv6 Tunnel Statistics

To view IP Interfaces statistics, **select Ports > IP Interfaces > Statistics** to open the IP Interfaces Statistics page.

The IPv6 tunnel statistics pane displays the gateway status as **Reachable** if neighbor discovery is completed with gateway or Unreachable if neighbor discovery failed. Neighbor discovery is done only on the encapsulation node. On the decapsulation node, the gateway 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* on page 798 for descriptions of these statistics as well as to *GigaSMART Operations Statistics Definitions* on page 794.

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* on page 797 for descriptions of these statistics as well as to *GigaSMART Operations Statistics Definitions* on page 794.

IP Fragmentation and Reassembly on L2GRE and GMIP Tunnels

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: Each fragmented packet will contain the tunnel header.

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.

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.

Feature	Description
Tunnel MTU	The tunnel MTU is configured using the MTU field on the **IP Interfaces** configuration page. (Select **Ports > IP Interfaces**, and then click **New** to open the page.) The MTU is fixed at 9400 for all network/tool ports on the following platforms except the following • GigaVUE-HB1 • GigaVUE-TA1, GigaVUE-TA10, and GigaVUE-TA40 • Certified Traffic Aggregation White Box The MTU is fixed at 9400 for all network/tool ports on the following platforms: • GigaVUE-HC2 and GigaVUE-HC2 equipped with Control Card version 2 (HC2 CCv2) • GigaVUE-HC1 • GigaVUE-HC3 • GigaVUE-TA100
Encapsulation Statistics	The encapsulation statistics count the number of fragmented packets. Refer to *Display GMIP Tunnel Encapsulation Statistics* on page 822 and *Display L2GRE Tunnel Encapsulation Statistics* on page 842. For definitions, refer to *Tunnel Encapsulation Statistics Definitions* on page 798.
Decapsulation Statistics	The decapsulation statistics count the number of reassembled packets. Refer to *Display GMIP Tunnel Decapsulation Statistics* on page 822 and *Display L2GRE Tunnel Decapsulation Statistics* on page 842. For definitions, refer to *Tunnel Decapsulation Statistics Definitions* on page 797.
GigaSMART Engine Ports	GigaSMART operations with a tunnel component can only be assigned to GigaSMART groups consisting of a single GigaSMART engine port. Refer to *Groups of GigaSMART Engine Ports* on page 757 for details.

GigaSMART ERSPAN Tunnel Decapsulation

Required License for ERSPAN Decapsulation: Advanced Tunneling (GigaVUE-HC2 and GigaVUE-HC3), Tunneling (GigaVUE-HC1 and GigaVUE-HB1)

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* on page 845.

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 (refer to the following figure). 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* on page 848 and *ERSPAN Type III Tunnel Header Removal* on page 849.

For an example of APF and ERSPAN tunneling, refer to *ERSPAN Tunneling* on page 1034.

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 30-11 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.

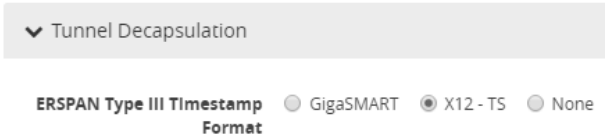

Figure 30-11: ERSPAN Type III Timestamp Formats on GigaSMART Groups Page

The GigaSMART timestamp is added to the Gigamon trailer. For the format of the GigaSMART trailer, refer to *GigaSMART Trailer Reference* on page 1196. The x12-ts timestamp is added to the PRT-H00-X12-S trailer. For the format of the PRT-H00X12TS trailer, refer to the *GigaVUE-OS CLI User's 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
- 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* on page 851 for how to display the output and to *ERSPAN Statistics Definitions* on page 797 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* on page 754 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 30-12.

NOTE: A flow ID of zero is a wildcard value that matches all flow IDs.

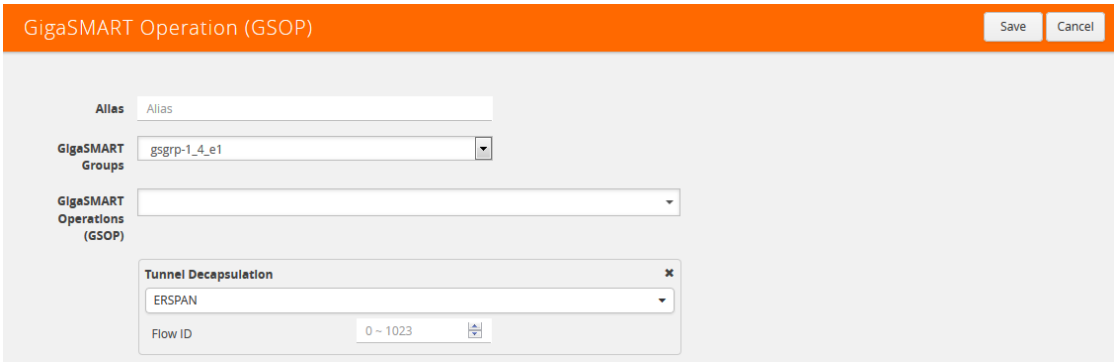

Figure 30-12: Decapsulation Flow ID Set to Zero.

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 Comment 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 30-12 on page 848. **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.

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 > 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 Comment fields, enter the name and description of the IP interface. **d.** Click the Ports field and select the tool port from the drop-down list. **e.** Enter the IP address, subnet mask, gateway, and MTU settings in the respective fields. For example, port 1/1/g2, IP address 10.10.10.10, mask 255.255.225.0, gateway 0.10.10.1, and MTU 1500. **f.** Click on the GigaSMART Group field to select the GigaSMART group. **g.** Click Save.
4.	Configure the GigaSMART operation and assign it to the GigaSMART group. **NOTE:** A flow ID of zero is a wildcard value that matches all flow IDs.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Type an alias in the Alias field. **d.** From the GigaSMART Groups drop-down list, select the GigaSMART Group that you created in the second task. **e.** From the GigaSMART Operations (GSOP) drop-down list, select Tunnel Decapsulation. **f.** Select ERSPAN for the decapsulation type. **g.** Enter a value of 0 in the Flow ID field. The configuration should look like the example shown in Figure 30-12 on page 848. **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 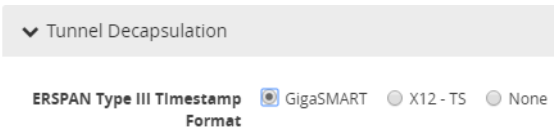 **e.** Click Save.

Task	Description	UI Steps
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* on page 797 for descriptions of these statistics as well as to *GigaSMART Operations Statistics Definitions* on page 794.

GigaSMART VxLAN Tunnel Decapsulation

Required License for VxLAN Decapsulation: Base (GigaVUE-HC2 and GigaVUE-HC3), Tunneling (GigaVUE-HC1 and GigaVUE-HB1)

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* on page 855

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* on page 852

NOTE: GigaSMART operations with a tunnel component can only be assigned to GigaSMART groups consisting of a single GigaSMART engine port.

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* on page 754 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 alias in 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 **Comment** 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 **GigaSMART Group** 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).
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.

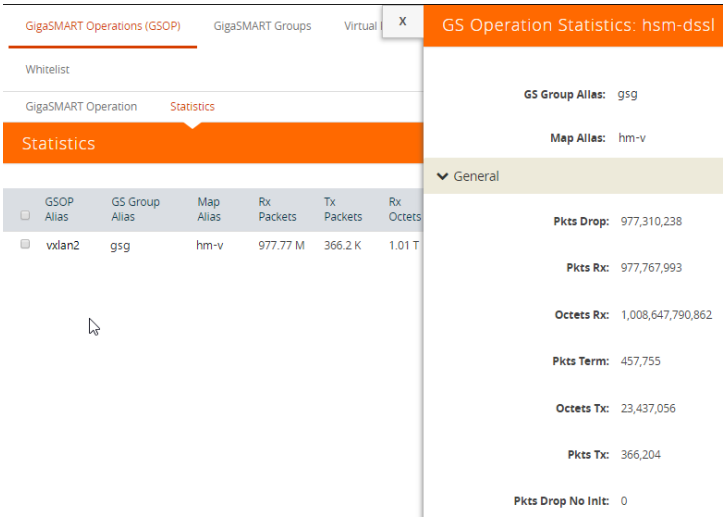

Figure 30-13: GSOP VxLAN Statistics Quick View

GigaSMART Custom Tunnel Decapsulation

Required Licenses for Custom Decapsulation: Base (GigaVUE-HC2 and GigaVUE-HC3), Tunneling (GigaVUE-HC1 and GigaVUE-HB1), 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 *Custom Tunnel Decapsulation Configuration Example* on page 856

NOTE: GigaSMART operations with a tunnel component can only be assigned to GigaSMART groups consisting of a single GigaSMART engine port.

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* on page 754 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**.

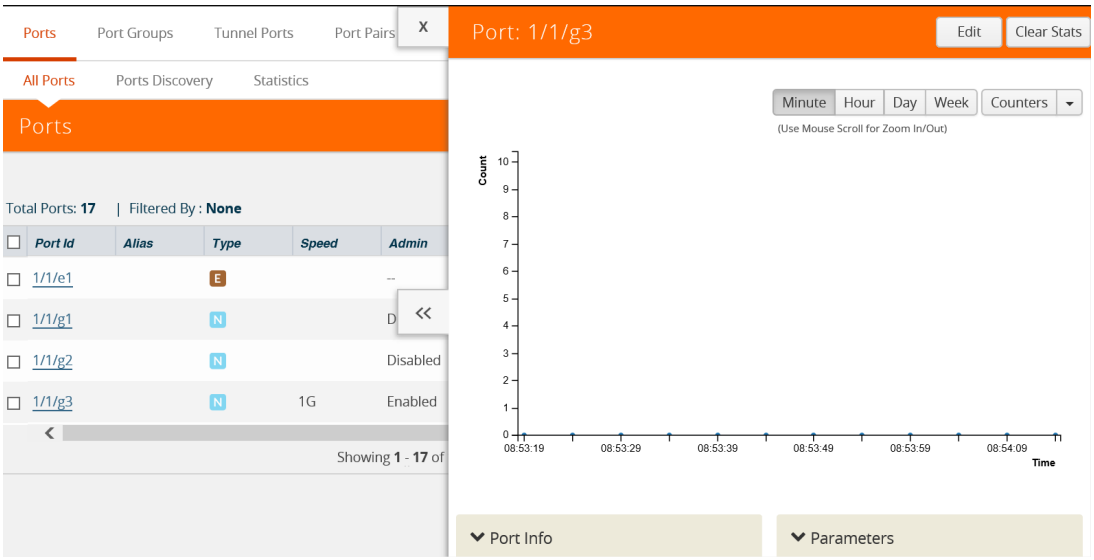

Figure 30-14:

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 **Comment** 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 **GigaSMART Group** 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 **Tunnel Decapsulation**.
   a. Select **Custom**
   b. Type **Source Port**
   c. Type **Destination Port**
6. Click **Save**

Create 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**.

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.

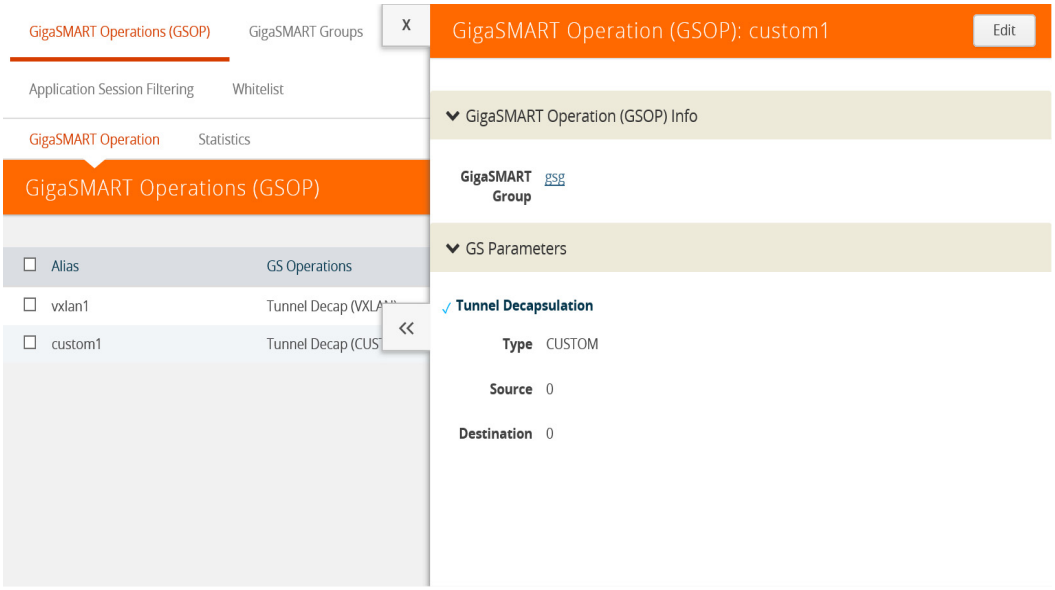

Figure 30-15: GSOP Custom Operation

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.

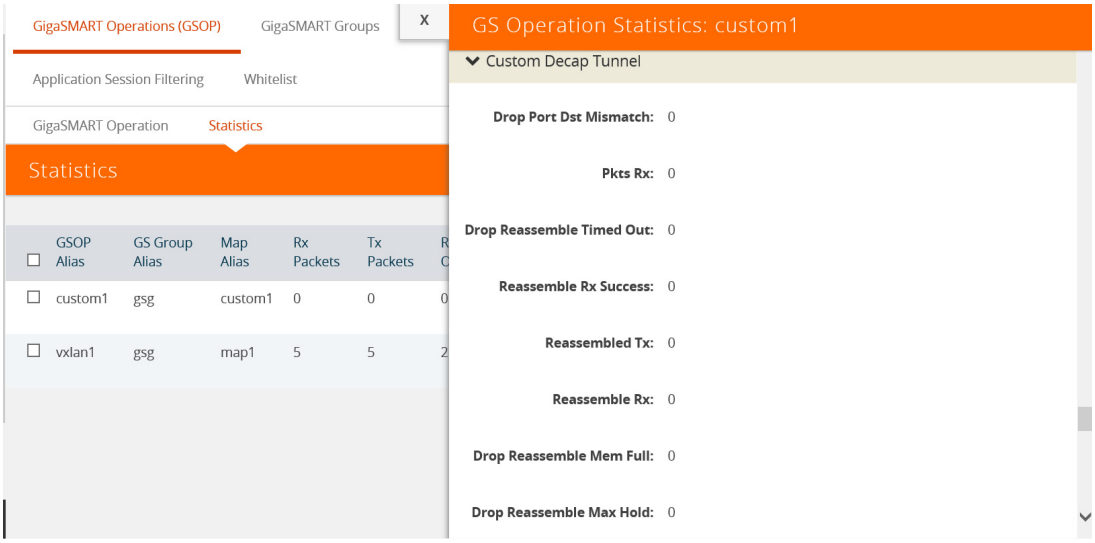

Figure 30-16:

4. **Click** the **X** to close the **GS Operation Statistics** pane.

GigaSMART Header Addition

Required License: Header Stripping

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: • **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. • **CFI** – 0 • **Priority** – 0 • **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* on page 504 for TPID.
CRCs Recalculated	The GigaVUE H Series node automatically recalculates and applies correct CRC checksums based on the new packet length after the header is stripped.
Viewing Statistics	Use **GS Operations Statistics** page to see statistics related to ongoing GigaSMART operations. Refer to *View GigaSMART Statistics* on page 782 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* on page 778 for details on the combinations of GigaSMART operations. Refer to *Order of GigaSMART Operations* on page 781 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* on page 757 for details.

GigaSMART De-Duplication

Required License: De-Duplication

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 field 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.

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 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.
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* on page 865.
GigaSMART Engine Ports	De-duplication operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to *Groups of GigaSMART Engine Ports* on page 757 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 *Example – GigaSMART De-Duplication* on page 863.

Configure GigaSMART Parameters for Packet De-Duplication

Use the **Dedup** 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: • Count – GigaSMART counts the duplicate packets, but does not drop them. • Drop – GigaSMART drops the duplicate packets. The default is drop.

Parameter	Description
IP Tclass **IP TOS** **TCP Sequence** **VLAN**	These options are useful when applying de-duplication operations to packets in a NAT environment. Different NAT implementations can change certain packet header fields (for example, the TCP sequence number). If you want to be able to detect duplicates without requiring that these fields match (ToS field, TCP sequence number, VLAN ID), you can disable the corresponding option. • **IP Tclass** – Ignore or include IPv6 traffic class. Use for IPv6. The default is include. • **IP TOS** – Ignore or include the IP ToS bits when detecting duplicates. Use for IPv4. The default is include. • **TCP Sequence** – Ignore or include the TCP Sequence number when detecting duplicates. The default is include. • **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* on page 754 for details.

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

Task	Description	UI Steps
1.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**a.** From the device view, select GigaSMART > GigaSMART Groups. **b.** Click New. **c.** Type an alias in the Alias field. For example, gs2port1. **d.** Click in Port List field and select an engine port. For example, 2/1/e1. **e.** Go to Task 2.

Task	Description	UI Steps
2.	Configure parameters on the GigaSMART group.	**a.** Under the Dedup 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 Deduplication 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 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-duplications statistics will be in the row labeled Dedup in the GS Operations column.

Refer to *De-duplication Statistics Definitions* on page 796 for descriptions of these statistics as well as to *GigaSMART Operations Statistics Definitions* on page 794.

GigaSMART Header Stripping

Required License: Header Stripping

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 30-17 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 30-18 shows an example of the GS Operations page configured for tunnel stripping.

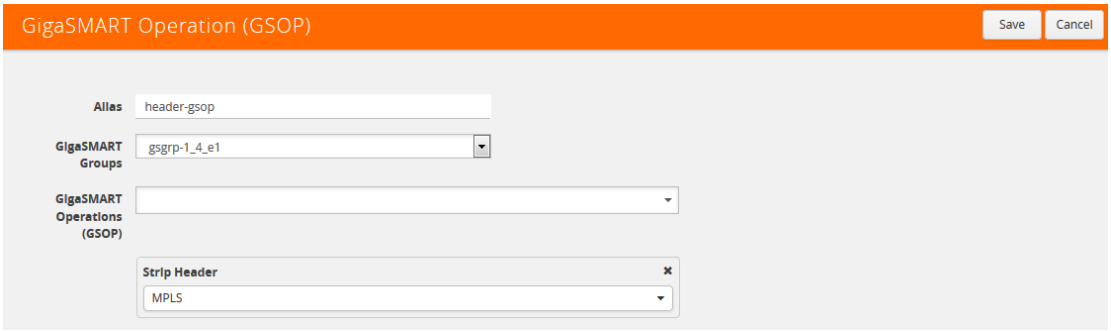

Figure 30-17: Strip Header GigaSMART Operation Configured

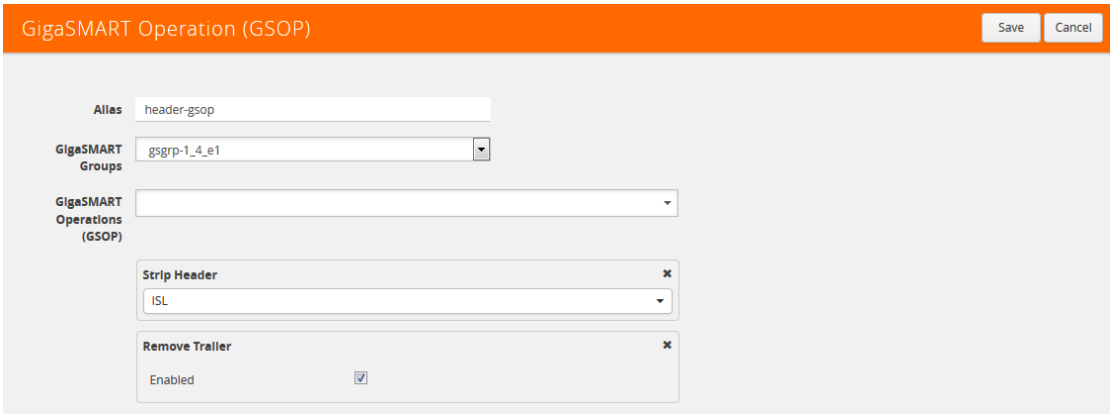

Figure 30-18: Tunnel Stripping GigaSMART Operation Configured

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.

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.

Summary	Description
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* on page 1196. 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 User's 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 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. For an FM6000 timestamping example, refer to *Example – FM6000 Timestamping* on page 872.
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 delivery header and the GRE header to expose the encapsulated packet. Only IPv4 as the delivery protocol is supported. 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* on page 870. **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* on page 871.
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* on page 504.

Summary	Description
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 will strip all the way to the end of the VXLAN layer instead of just the VXLAN tag.
ISL Tunnel Stripping	ISL tunnel stripping removes the 26-byte header and the 4-byte FCS trailer associated with Cisco ISL VLAN encapsulation. **IMPORTANT:** Make sure the packets processed by a GigaSMART operation with a **Strip Header ISL** component are all using ISL encapsulation. GigaSMART operations do not distinguish between packets using ISL and packets that do not – it strips the requisite bytes from all packets it processes.
GTP Tunnel Stripping	GTP tunnel stripping removes the header and trailer for GTP-u packets inside the GTP tunnel between the SGSN and GGSN interfaces in a 3G network, and between the eNodeB (eNb) and the SGW and between the SGW and the PGW in an LTE network. The SGSN and GGSN interfaces are also referred to as the Gn (or Gp) interface. The interface between eNb and SGW is referred to as S1U. The user plane interface between SGW and PGW is referred to as S5-U/S8-U. Both use GTPv1. Both GTPv1 and GTPv0 are supported for stripping. GTP-c control packets are not stripped. GTP[1] (also referred to as “GTP-Prime”) is not supported for stripping.
Ethertype Replaced	After the VLAN/MPLS headers are stripped, the original ethertype carried in the Layer Two header is no longer valid. The GigaSMART replaces the ethertype field differently for MPLS and VLAN packets: Ethertype Replacement for VLAN Packets VLAN-tagged packets carry the original value for the ethertype field immediately after the VLAN tag. After the four-byte VLAN header is stripped, GigaSMART simply sets the ethertype field in the Layer 2 header to the value that was originally present in the packet past the VLAN tag. Ethertype Replacement for MPLS Packets Unlike VLAN-tagged packets, the Layer 3 protocol type is not carried in the packet for an MPLS packet – instead, it is applied by an egress router. To handle this, the GigaSMART examines the byte following the MPLS header to determine whether the packet is IPv4/IPv6 and takes the following actions: • **IPv4** – The GigaSMART replaces the ethertype from the MPLS packet with the IPv4 ethertype (0x0800) • **IPv6** – The GigaSMART replaces the ethertype from the MPLS packet with the IPv6 Ethertype (0x86DD). • **Non-IPv4/IPv6** – The GigaSMART passes the packet to destination tool ports without stripping the header. MPLS header stripping is only supported for IPv4/IPv6 packets.
CRCs Recalculated	The GigaVUE H Series node automatically recalculates and applies correct CRC checksums based on the new packet length after the header is stripped.

Summary	Description
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* on page 782 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* on page 778 for details on the combinations of GigaSMART operations. Refer to *Order of GigaSMART Operations* on page 781 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* on page 757 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 *Generic Header Stripping* on page 874.

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.

Typical MPLS Packet arriving at network port:

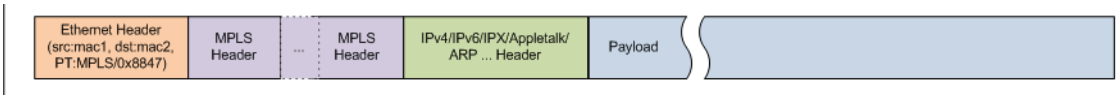

Packet after MPLS header is stripped and ethertype is replaced:

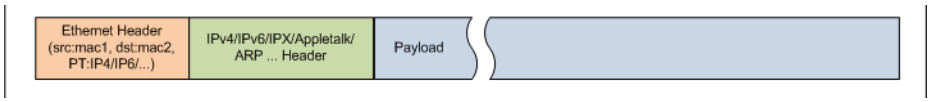

Example – Stripping MPLS Headers and Adding a VLAN ID

The example shown in Figure 30-19 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.

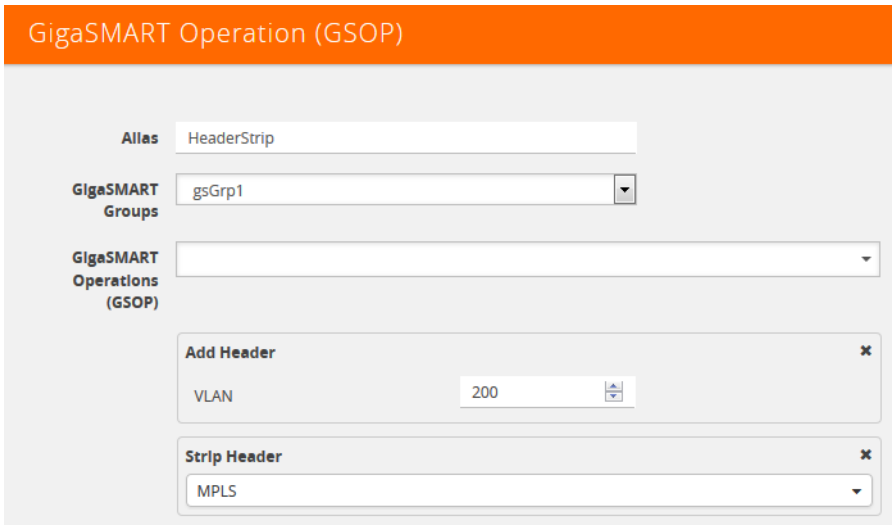

Figure 30-19: Strip Header GigaSMART Operation for Stripping MPLS Headers

The example shown in Figure 30-20illustrates 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.

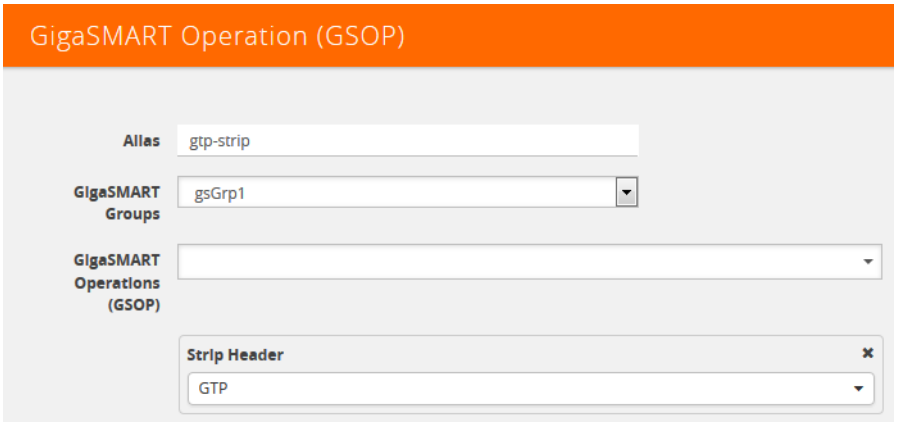

Figure 30-20: Strip Header GigaSMART Operation for Stripping GTP Tunnel Information

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.

Before L3 GRE Stripping:

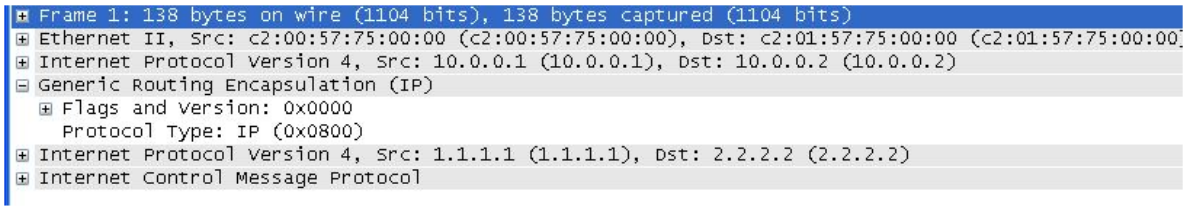

After L3 GRE Stripping:

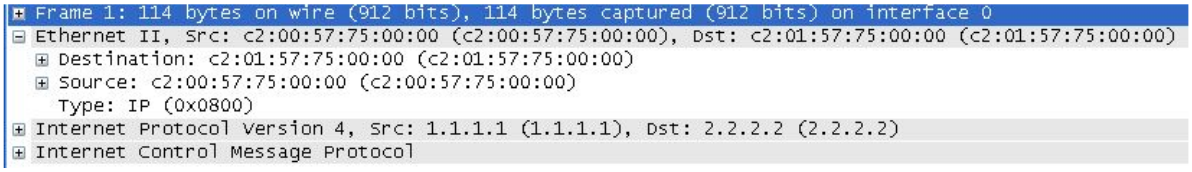

Example – Stripping Layer 2 GRE Ethernet Encapsulated Packets

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.

Before L2GRE Stripping:

```
⊞ Frame 1458: 192 bytes on wire (1536 bits), 192 bytes captured (1536 bits)
⊞ Ethernet II, Src: Vmware_65:a5:7f (00:50:56:65:a5:7f), Dst: GigamonL_17:1a:a7 (00:1d:ac:17:1a:a7)
⊞ Internet Protocol Version 4, Src: 192.168.8.17 (192.168.8.17), Dst: 192.168.8.5 (192.168.8.5)
⊞ Generic Routing Encapsulation (Transparent Ethernet bridging)
⊞ Ethernet II, Src: JuniperN_f5:98:30 (64:87:88:f5:98:30), Dst: Vmware_01:03:8f (00:50:56:01:03:8f)
⊞ Internet Protocol Version 4, Src: 10.208.74.152 (10.208.74.152), Dst: 172.29.99.28 (172.29.99.28)
⊞ Transmission Control Protocol, Src Port: diameter (3868), Dst Port: diameter (3868), Seq: 1, Ack: 1, Len: 80
⊞ Diameter Protocol
```

After L2GRE Stripping:

```
⊞ Frame 1: 146 bytes on wire (1168 bits), 146 bytes captured (1168 bits) on interface 0
⊞ Ethernet II, Src: JuniperN_f5:98:30 (64:87:88:f5:98:30), Dst: Vmware_01:03:8f (00:50:56:01:03:8f)
⊞ Internet Protocol Version 4, Src: 10.208.74.152 (10.208.74.152), Dst: 172.29.99.28 (172.29.99.28)
⊞ Transmission Control Protocol, Src Port: diameter (3868), Dst Port: diameter (3868), Seq: 1, Ack: 1, Len: 80
⊞ Diameter Protocol
```

Display Header Stripping Statistics

To display header stripping statistics, select **GigaSMART > GigaSMART Operations (GSOP) > Statistics**.

Refer to *Header Stripping Statistics Definitions* on page 801 for descriptions of these statistics as well as to *GigaSMART Operations Statistics Definitions* on page 794.

Example – FM6000 Timestamping

Figure 30-21 shows an example GigaSMART Operation configured to strip packets containing the FM6000 timestamp:

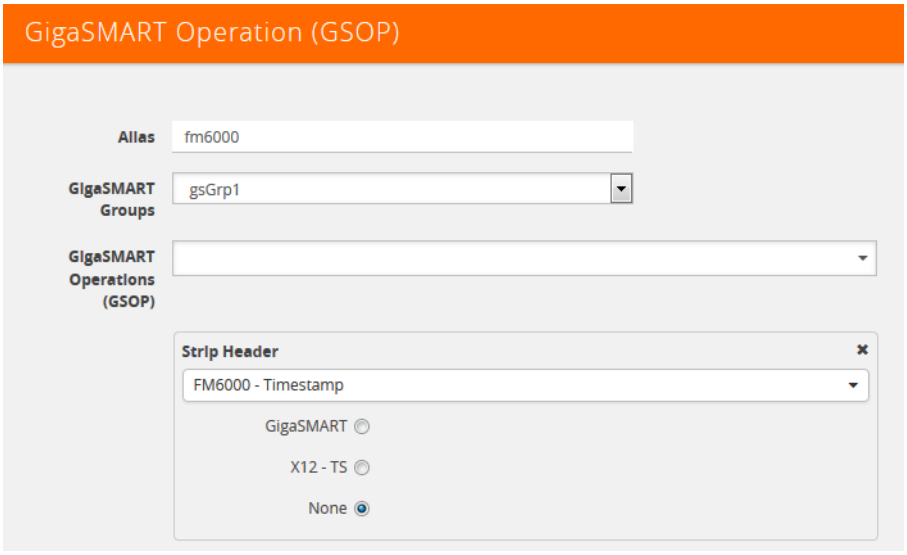

Figure 30-21: Strip Packets with FM6000 Timestamp

Figure 30-3 is an example GigaSMART Operation configured to convert packets containing the FM6000 timestamp to UTC and append the UTC timestamp to the Gigamon trailer:

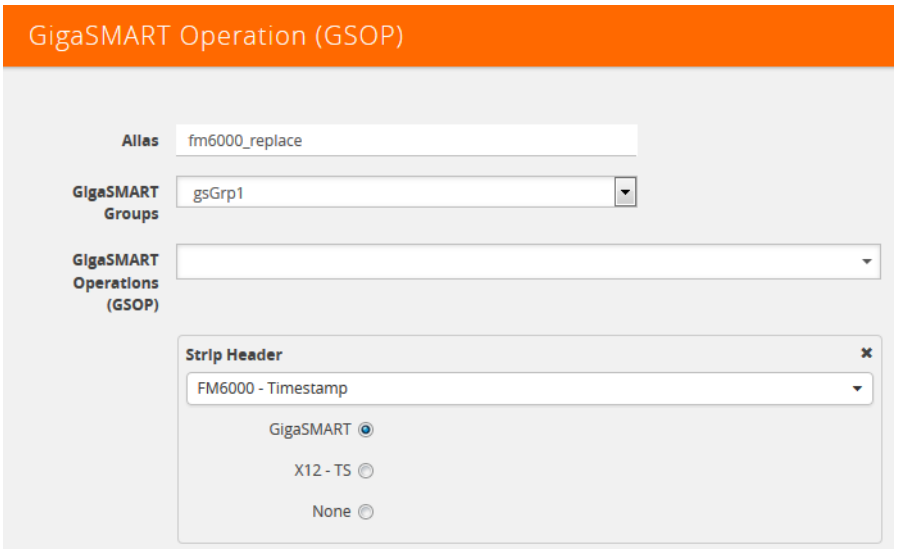

Figure 30-22: Strip Packets and Append UTC Timestamp to the Gigamon Trailer

Figure 30-23 on page 873 is an example map using the strip header GigaSMART operation in Figure 30-22.

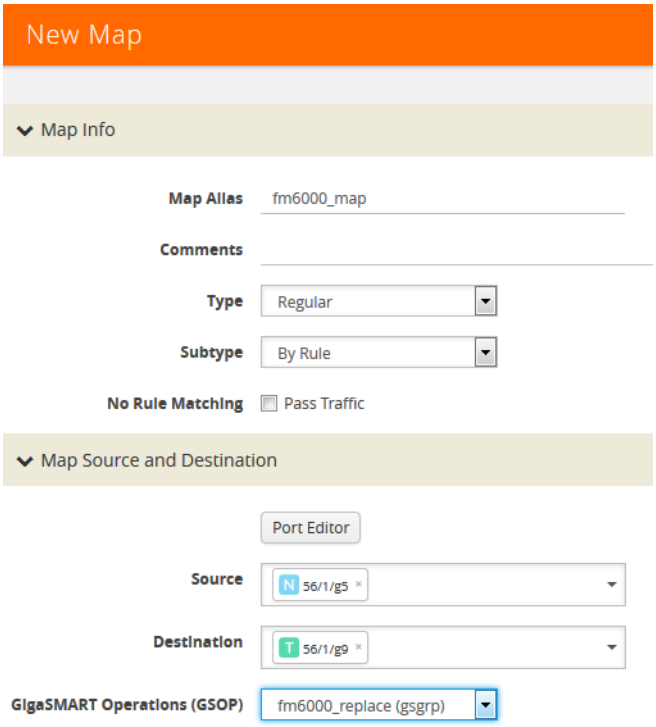

Figure 30-23: Map with Strip Header GigaSMART Operation

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* on page 754 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.

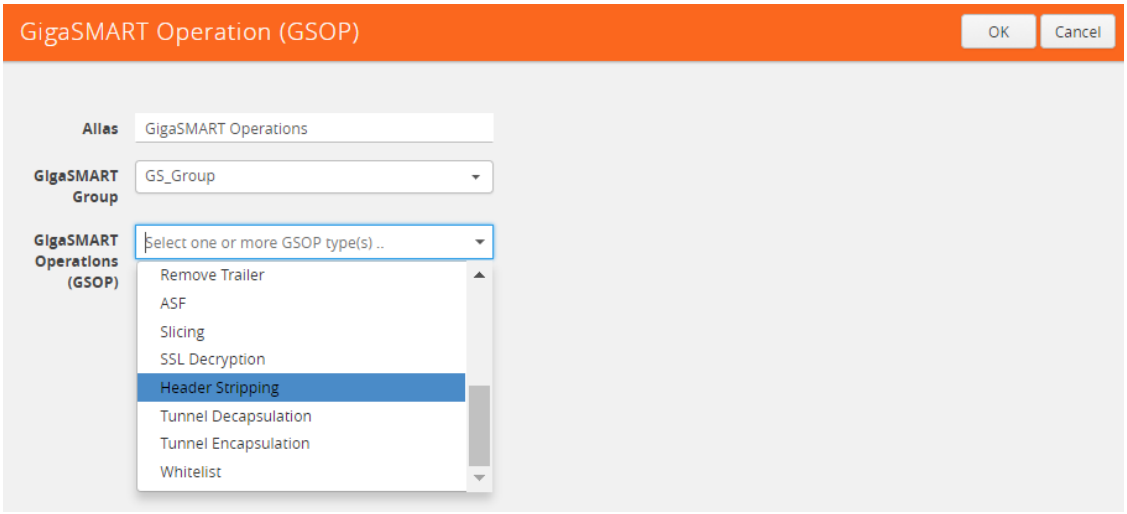

Figure 30-24: GigaSMART Header Stripping

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.

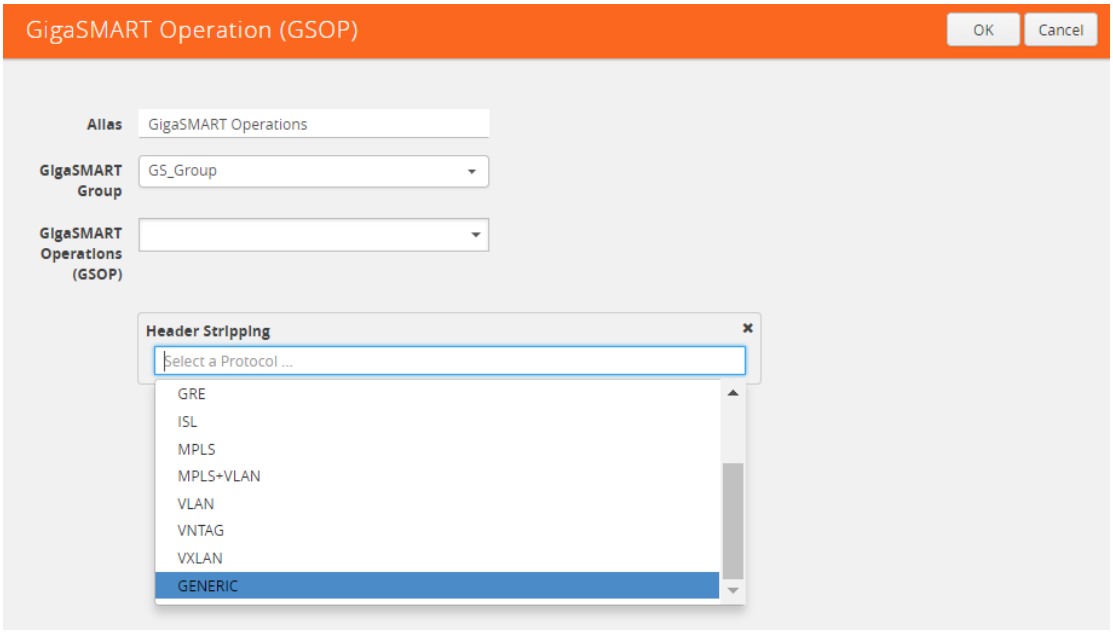

Figure 30-25: GigaSMART Generic Header Stripping

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

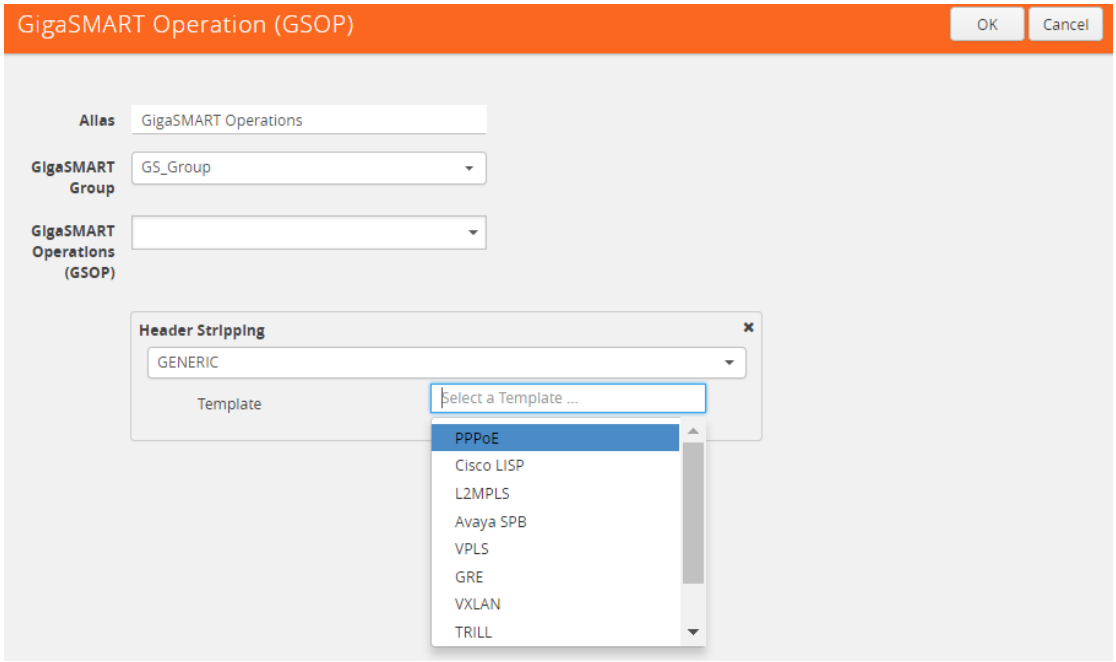

Figure 30-26: Generic Header Stripping Template

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.**

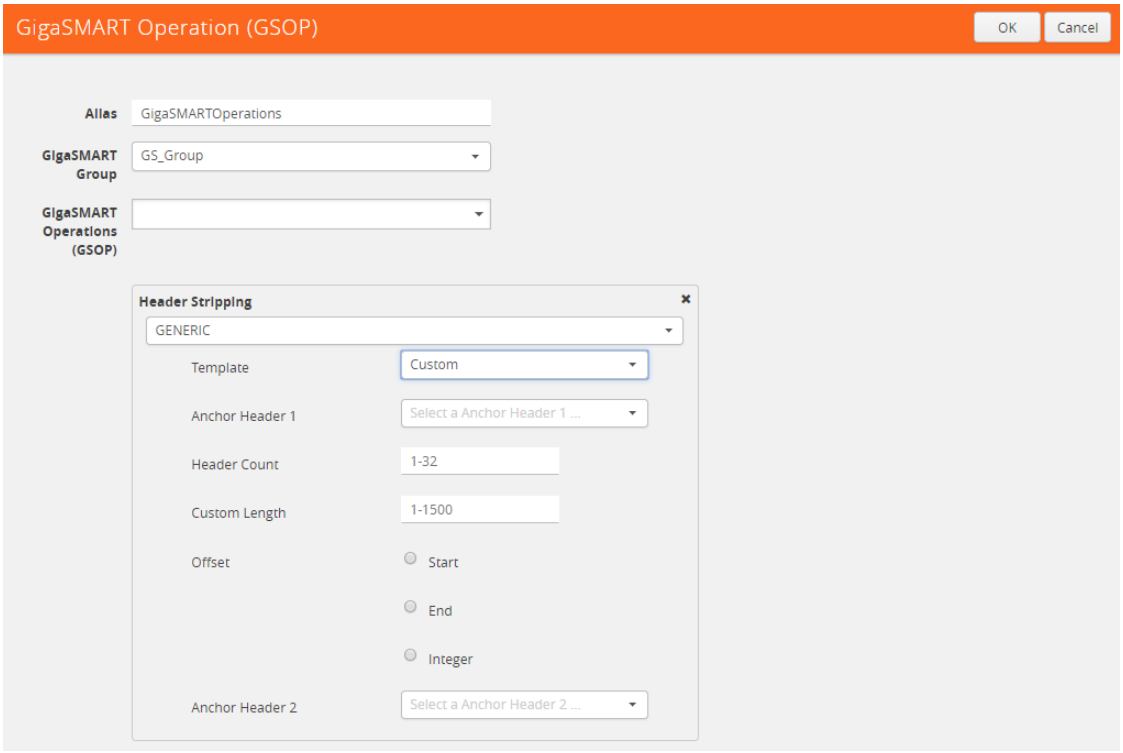

Figure 30-27: Generic Header Stripping - Custom

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: None, Eth, VLAN, MPLS, 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**.

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

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

PPPoE

The PPPoE template lets you strip the PPPoE encapsulated packets from the packet structure. Figure 30-28 illustrates a red outline around the frame that needs to be striped.

Figure 30-28: PPPoE Encapsulated Packets

To strip the PPoE encapsulated packets:

1. From the **Template** drop-down list, select **PPPoE**. Refer to Figure 30-29.

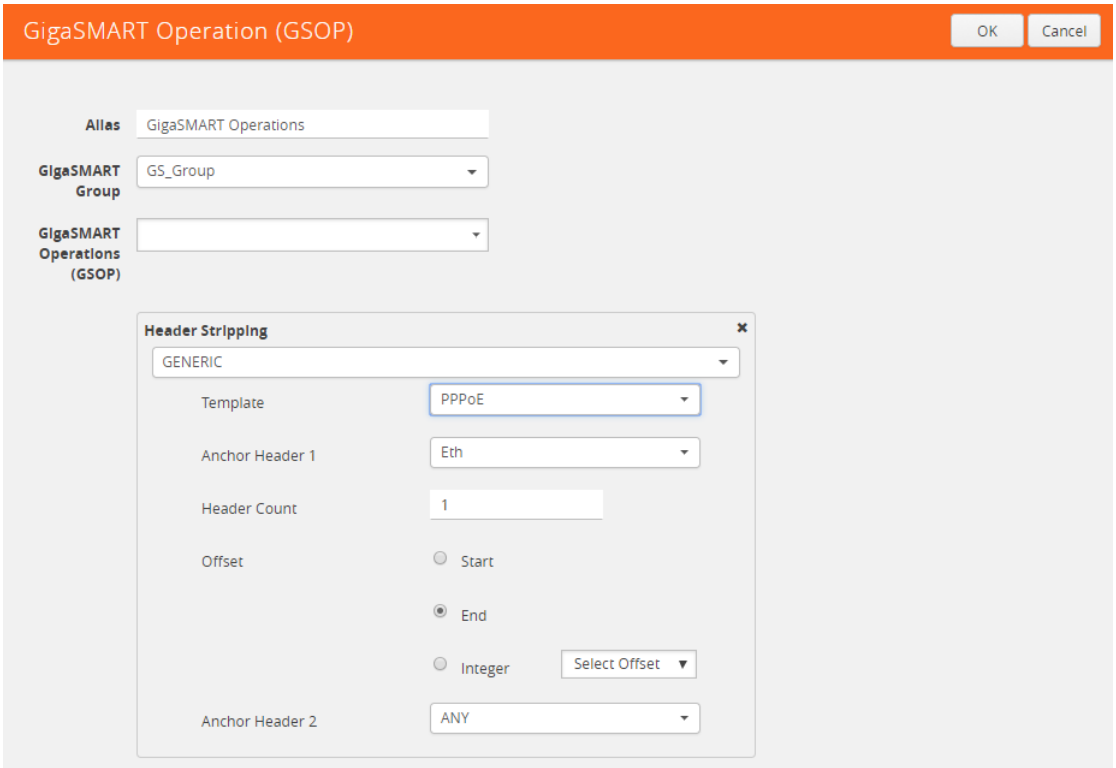

Figure 30-29: Generic Header Stripping - PPPoE

By default, the following values are selected:

Field	Value	Description
Anchor Header 1 **Offset**	Eth End	Starts the header stripping operation from the right end of the Ethernet header.
Header Count	1	Strips the header next to the Ethernet Header.
Anchor Header 2	Any	Updates a valid protocol as the Anchor Header 2 in the packet. In this case, any IPv4 or IPv6 protocol can become the Anchor Header 2.

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

Cisco LISP

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

Figure 30-30 illustrates a red outline around the frame that needs to be striped.

Figure 30-30: Cisco LISP Encapsulated Packets

To strip the Cisco LISP encapsulated packets:

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

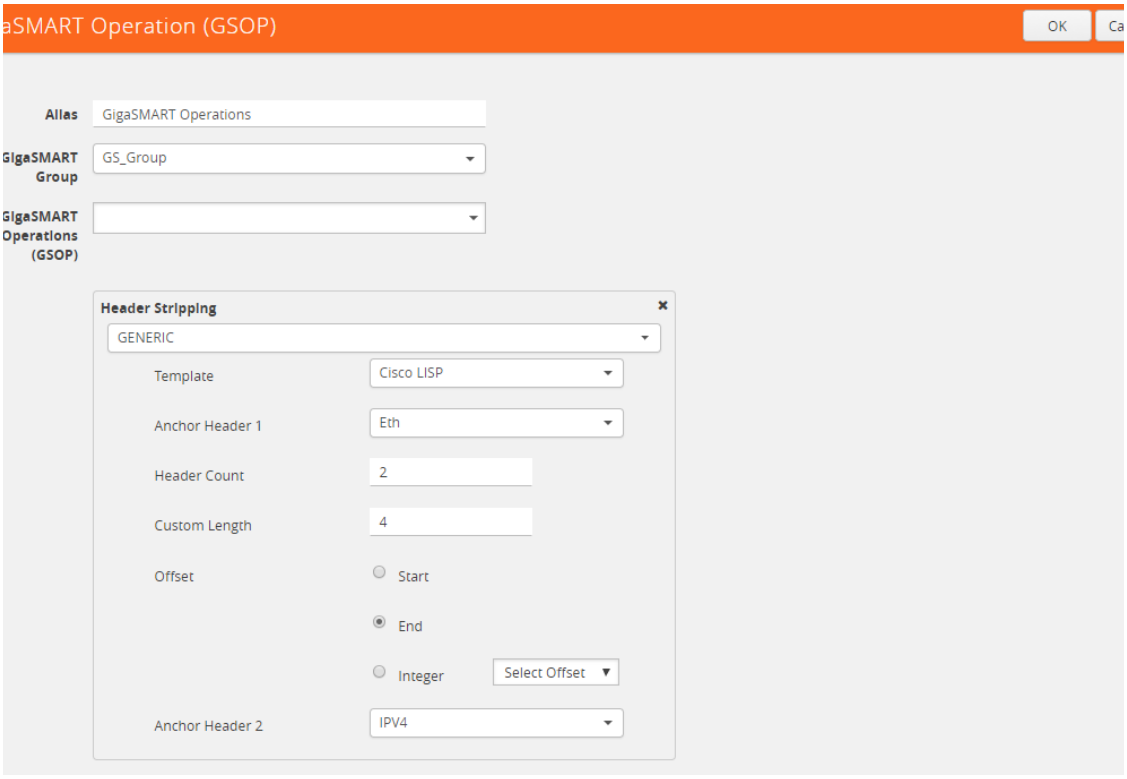

Figure 30-31: Generic Header Stripping - Cisco LISP

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 30-32 illustrates a red outline around the frame that needs to be striped.

Figure 30-32: 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**.

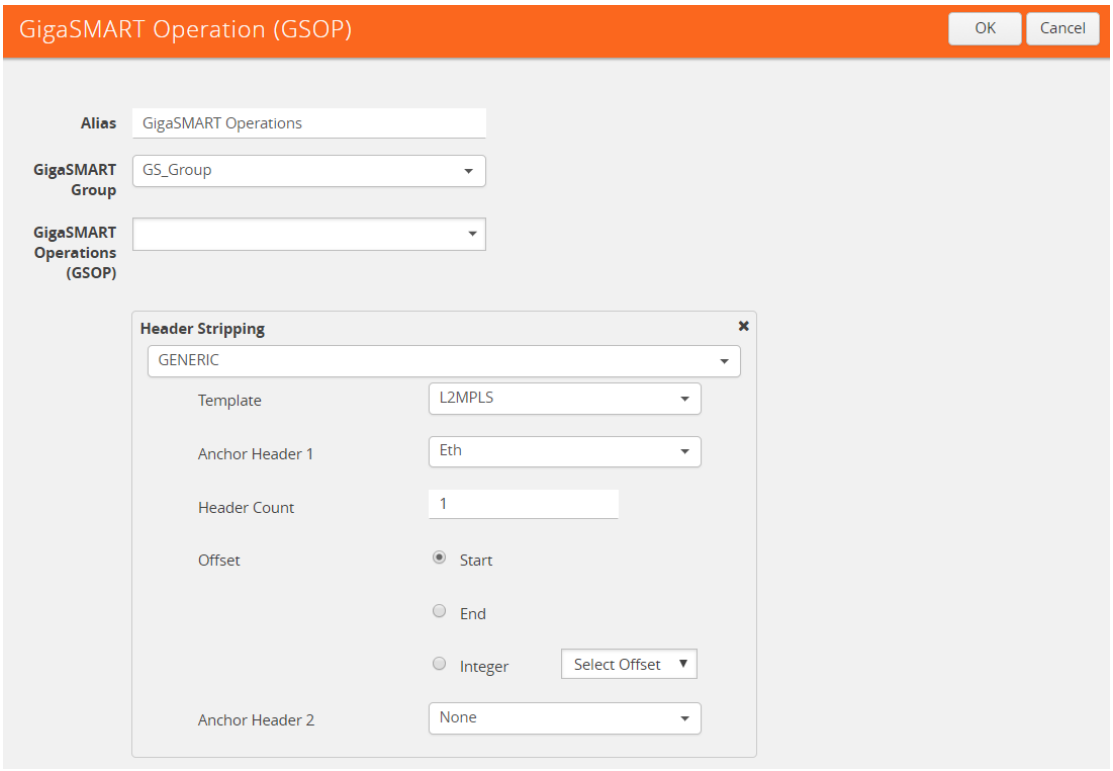

Figure 30-33: Generic Header Stripping - L2 MPLS

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.

Field	Value	Description
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 30-34 illustrates a red outline around the frame that needs to be striped.

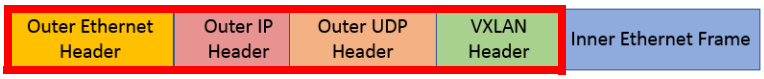

Figure 30-34: VxLAN Encapsulated Packets

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

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

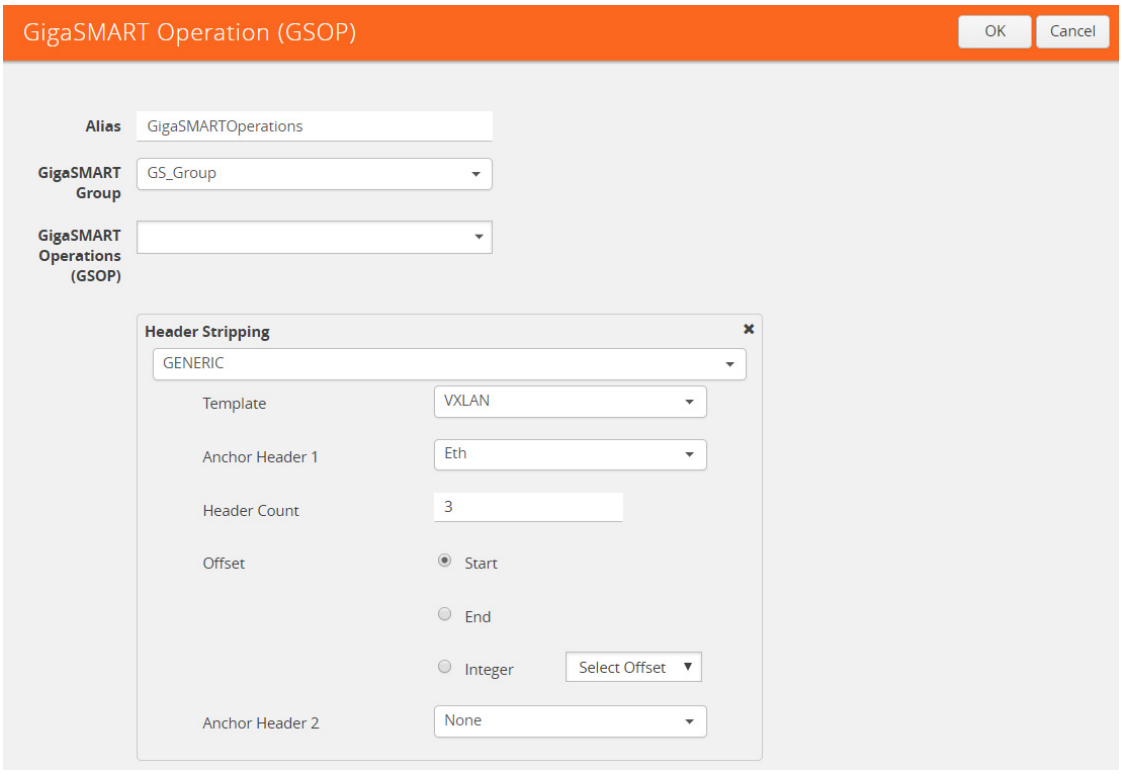

Figure 30-35: Generic Header Stripping - VXLAN

By default, the following values are selected:

Field	Value	Description
Anchor Header 1 Offset	Eth Start	Starts the header stripping operation from the start of the Ethernet header.
Header Count	3	Strips the next three headers—outer IP header, outer UDP header, and VXLAN header.
Anchor Header 2	None	Signifies that there is no need to specify the Anchor Header 2. In this case, the IPv4 protocol forms the first header of the packet. **NOTE:** When the Anchor Header 1 is set to **None**, the Anchor Header 2 must also be set to **None**.

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

TRILL

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

Figure 30-36 illustrates a red outline around the frame that needs to be striped.

Figure 30-36: TRILL Encapsulated Packets

To strip TRILL from the encapsulated packets:

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

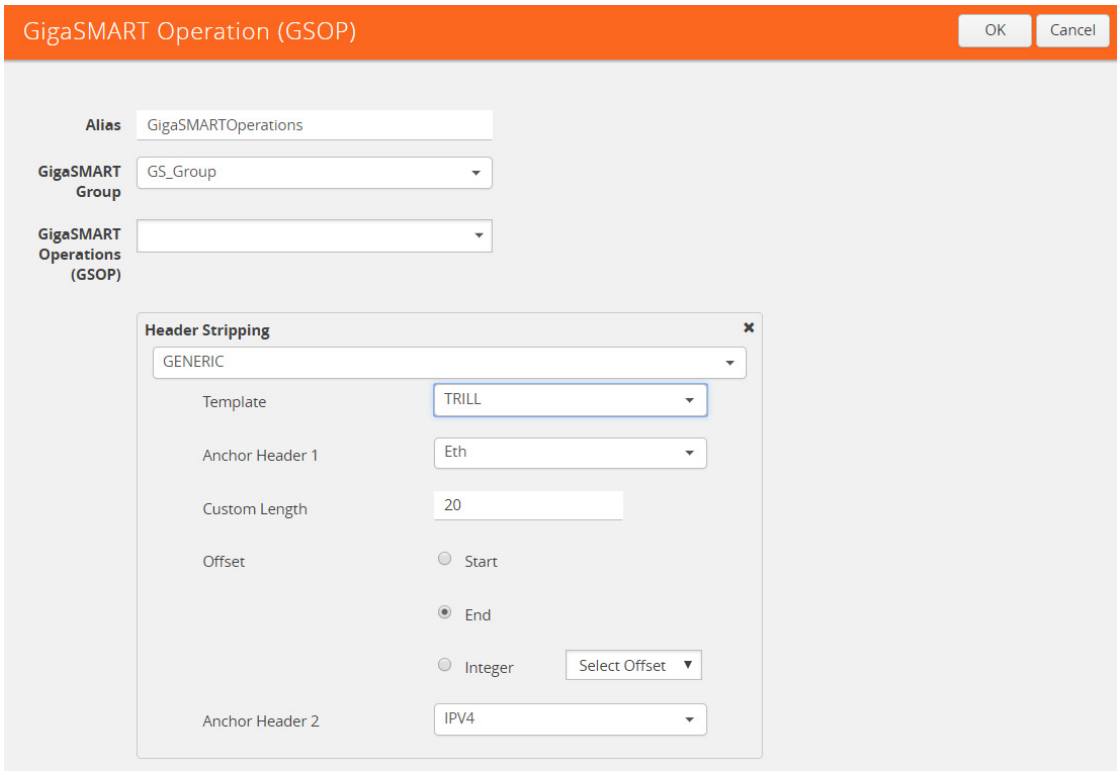

Figure 30-37: Generic Header Stripping - TRILL

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 30-38 illustrates a red outline around the frame that needs to be striped..

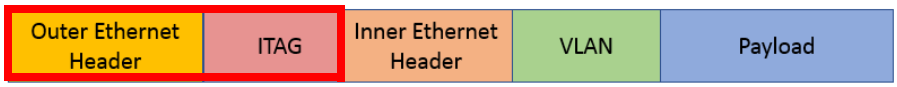

Figure 30-38: Avaya SPB Encapsulated Packets

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

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

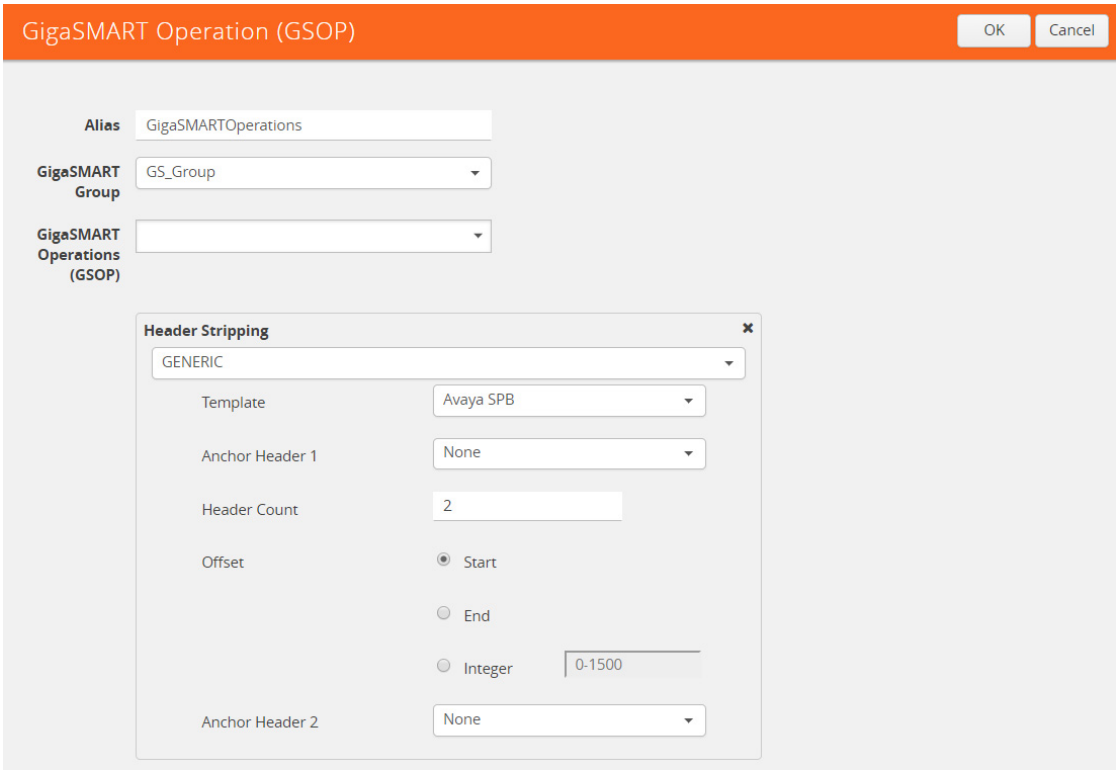

Figure 30-39: Generic Header Stripping - Avaya SPB

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 30-40 illustrates a red outline around the frame that needs to be striped.

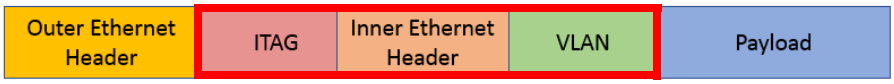

Figure 30-40: Avaya SPB Encapsulated Packets

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 GTP Correlation

Required License: GTP Filtering & Correlation

The GigaSMART GTP application correlates traffic based on mobile subscriber IDs in the packet data networks of service providers. It provides a mechanism to filter and forward session traffic for subscribers to tools. GTP correlation assists mobile carriers in debugging and analyzing GTP traffic in their 3G/4G networks.

GPRS Tunneling Protocol (GTP) is an IP/UDP-based protocol that carries mobile data across service provider networks. The protocol is used in General Packet Radio Service (GPRS) networks such as: Global System for Mobile Communications (GSM), Universal Mobile Telecommunications System (UMTS), and Long Term Evolution (LTE). The protocol encapsulates user data that passes through the core network and carries subscriber-specific signaling traffic.

GTP includes both control plane (GTP-c) and user-data plane (GTP-u) traffic. To gain an accurate view into the subscriber's session, GTP tunnels are used to correlate the subscriber-specific control plane and user-data plane traffic. A GTP session is the minimum unit of GTP correlation consisting of one control and multiple user tunnels. All GTP traffic belonging to the same session is forwarded to the same tool port.

Using GTP correlation, you can filter, replicate, and forward specific subscriber sessions to specific tools by correlating the subscriber IDs that are exchanged as part of the control sessions to the corresponding tunnel IDs (TEIDs) that are part of the user-data plane traffic.

GTP correlation provides the following:

- stateful filtering based on subscriber IDs (IMSI, IMEI, and MSISDN)
- stateful filtering based on GTP version or EPC interface
- stateful correlation of GTP-c with GTP-u traffic
- correlation of subscriber ID with corresponding tunnel ID
- forwarding of the subscriber-specific control and user-data plane traffic to a tool or group of tools
- supports a maximum of 6 million GTP subscriber sessions for GigaVUE-HC2 nodes, whereas, it supports 12 million GTP subscriber sessions for GigaVUE-HC3 nodes
- combine with GigaSMART Load Balancing to load balance GTP traffic to a set of tool ports. For information on GTP load balancing, refer to stateful load balancing in the section *GigaSMART Load Balancing* on page 1147. For examples of GTP load balancing, refer to *Configure GTP Correlation Examples* on page 897. Starting in software version 4.6, GTP load balancing in a cluster is supported for GTP flow filtering. For an example of GTP load balancing in a cluster, refer to *GTP Whitelisting and GTP Flow Sampling Examples* on page 924.

Starting in software version 4.5, a GigaSMART group (gsgroup) associated with GTP applications can have multiple GigaSMART engine port members (**e** ports), up to four, forming an engine group. Refer to *GTP Scaling* on page 959.

Filtering on Subscriber IDs and Version

GTP stateful filtering supports filtering of GTP sessions based on the following subscriber IDs:

Component	Description
imsi	The International Mobile Subscriber Identity (IMSI) is a number that identifies a subscriber of a cellular network. It is a unique identification associated with all cellular networks. An IMSI is usually a 15 digit number, associated with GSM, UMTS, and LTE network mobile phone users.
imei	The International Mobile Station Equipment Identity (IMEI) is a number, usually unique, that identifies 3rd Generation Partnership Project (3GPP), for example, GPRS, LTE, as well as Integrated Digital Enhanced Network (iDEN) mobile phones, and some satellite phones. The IMEI identifies the device, but has no permanent relationship to the subscriber. Instead, the subscriber is identified by transmission of an IMSI number, stored on a SIM card.
msisdn	The Mobile Station International Subscriber Directory Number (MSISDN) is a unique number that identifies subscribers in a GSM or UMTS mobile network. This numbering plan is defined in the ITU-T recommendation E.164. The maximum length of an MSISDN is 15 digits.

In addition to filtering on subscriber IDs, you can optionally filter on GTP version (v1 or v2) or Evolved Packet Core (EPC) interface. Filtering on the EPC interface allows traffic to be segmented for a given interface.

The supported interfaces for EPC filtering are as follows:

- Gn/Gp
- S11U
- S11/S1-U
- S5/S8
- S10
- S2B

When filtering on EPC interface, you do not also need to specify version, as the version is implied.

To create maps using GTP, specify a **Second Level Flow Sample** map and select **GTP** for the rule. When adding a map rule, you can specify the following:

- subscriber IDs (IMSI, IMEI, or MSISDN)
- number of digits. The maximum number of digits for the IMSI or MSISDN value is 15. The maximum number of digits for the IMEI value is 16. To specify the prefix for IMSI, IMEI, or MSISDN, you can use a wild card character or a digit string followed by a wild card character.
- map comment to label the purpose of a rule or the type of traffic covered by a rule

- GTP version 1 or version 2 (refer to Figure 30-41) or EPC interface Gn/GP, S5/S8, or S10, S2B, or S11/S1U (refer to Figure 30-42)

NOTE: In a map, version and EPC interface cannot be specified in the same flowrule, but they can be specified in different flowrules.

NOTE: The maximum number of GTP flowrules is 32 per map.

For examples of filtering on GTP version, refer to *Configure GTP Correlation Examples on page 897*.

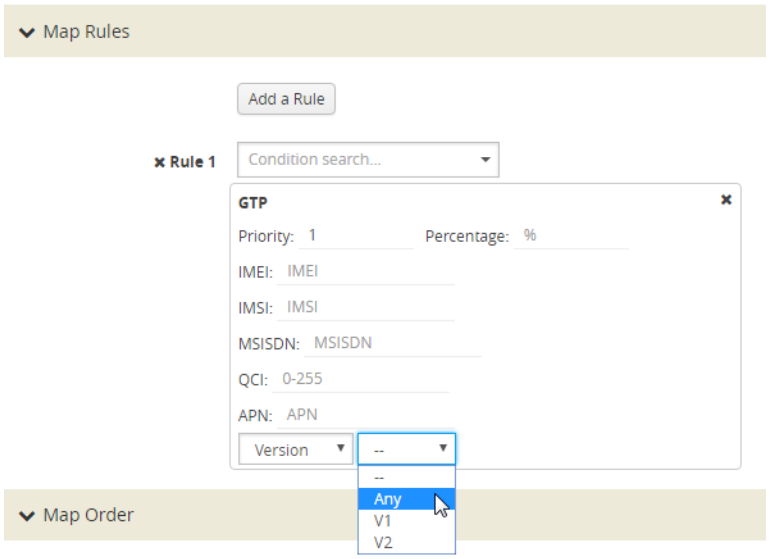

Figure 30-41: GTP Version

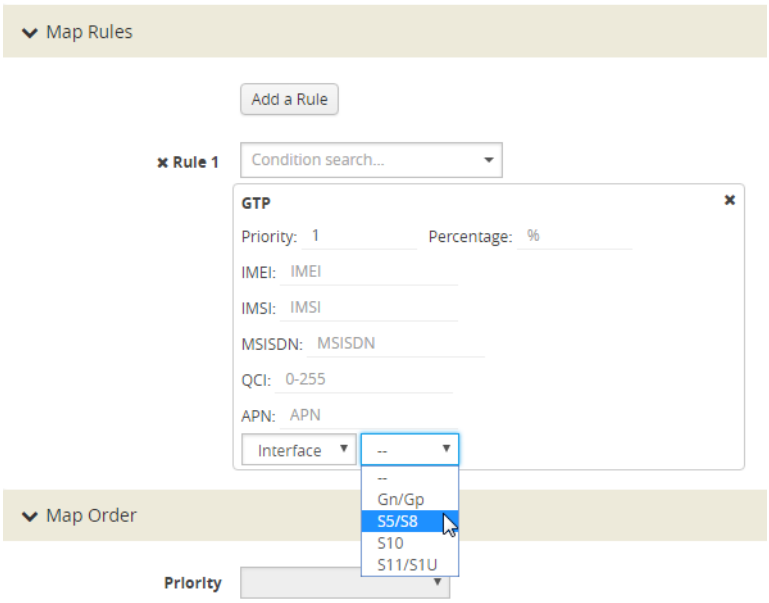

Figure 30-42: GTP EPC Interface

Session Correlation

Each GTP session has one control tunnel and one or more user tunnels. All the tunnels are correlated together into a session. Packets belonging to the same session will be forwarded to the same tool port. Refer to the following figure.

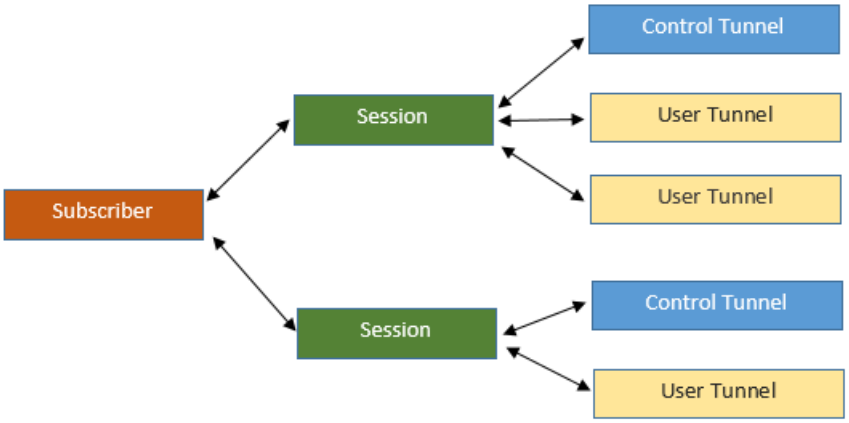

In a second level map, the following can be specified:

- one tool port—packets from one subscriber (same subscriber ID), from one or more GTP sessions, will be forwarded to the same tool port.
- multiple tool ports—packets from one subscriber (same subscriber ID), from multiple GTP sessions will be correlated and forwarded to same tool port. Using load balancing, GTP traffic that matches the same map but belongs to different subscribers can be load balanced to multiple tool ports.

Supported Interfaces

GTP is used at multiple interfaces by multiple devices in the core network. GTP stateful correlation is implemented for the following interfaces:

- Gn/Gp (for GPRS). The Gn interface is between SGSN-GGSN only.
- S5/S8 (for LTE)
- S1-U, S11U and S11 (for LTE)
- S10 (for S1-based mobility) Refer to *Conditional S10 Support* on page 891
- S2B

Support for interfaces for both GPRS and LTE networks includes the handovers between the different networks. Refer to the following figure.

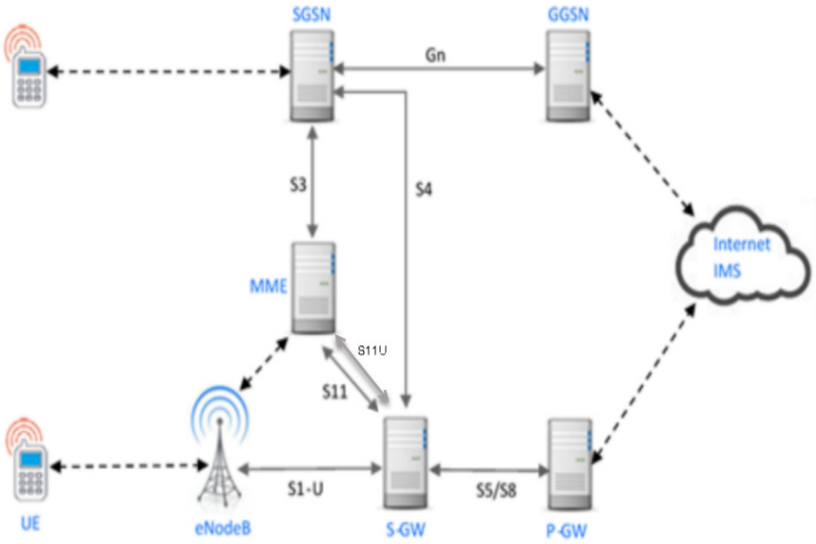

For LTE networks, the following GTP traffic will be correlated to the specific mobile subscriber and routed to the same tool port:

- GTP-c traffic on the S11 interface between MME and S-GW
- GTP-u traffic on the S11u interface between MME and S-GW
- GTP-u traffic on the S1u interface between eNodeB and S-GW
- GTP-c traffic on the S10 interface between MMEs
- GTP-c traffic on the S5/S8 interface between S-GW and P-GW
- GTP-u traffic on the S5u interface between S-GW and P-GW
- GTP-c traffic on the S2b interface between P-GW and ePDG
- GTP-u traffic on the S2b-U interface between P-GW and ePDG

In order to correlate GTP-c and GTP-u traffic running on different interfaces, you must tap into the correct interfaces, as follows:

- Gn/Gp—one interface runs both GTP-c and GTP-u
- S5/S8—one interface runs both GTP-c and GTP-u
- S1u, S11u, and S11—these three interfaces have to be tapped at the same time to get both GTP-c and GTP-u to perform the correct correlation.
- S2b-C and S2b-U interfaces needs to be tapped to get GTP-c and GTP-u traffic for correlation.

For examples of filtering on GTP interface types Gn/Gp, S5/S8, S1 and S11, refer to *Configure GTP Correlation Examples* on page 897.

Conditional S10 Support

The following table outlines support for the S10 interface. In the table, No means not supported, Conditional means there is limited support of the S10 interface.

S10	Support
Forward Relocation Request/Resp	Conditional; IMSI must be present in Forward Relocation Request
Forward Relocation Complete Notification/Ack	Conditional; IMSI must be present in Forward Relocation Request for Forward Relocation Complete Notification/Ack to be supported
Context Request/Response and Ac	Conditional; IMSI must be present
Identification Request/Response	No
Forward Access Context Notification/Ack	No
Relocation Cancel Request/Response	No
Configuration Transfer Tunnel	No

GigaSMART 3GPP CUPS Support

CUPS stands for Control and User Plane Separation of Evolved Packet Core nodes. 3GPP CUPS enhances the Evolved Packet Core nodes by introducing the following three new interfaces between the control plane and user plane functions of the S-GW, P-GW and TDF, respectively:

- SXa
- SXb
- SXc

Figure 30-43 illustrates 3GPP CUPS Network topology

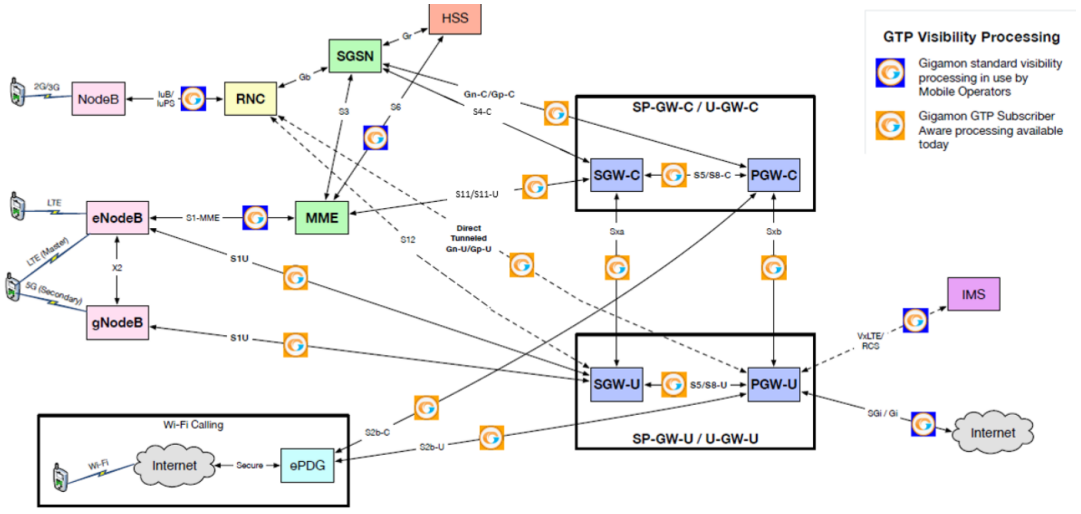

Figure 30-43: 3GPP CUPS Network topology

The advantages of 3GPP CUPS include the following:

- Increased flow of data traffic.
- Reduced latency.
- Independent User Plane and Control Plane scaling.

Starting in software version 5.6, GigaSMART GTP correlation leverages the SXa and SXb interfaces of the 3GPP CUPS architecture to receive additional traffic that is used to include the GTP session with the Packet Forwarding Control Protocol (PFCP) session. With this enhancement, control traffic and user traffic are processed at the following nodes (engines):

- Control Processing Node (CPN) - to process control traffic
- User Processing Node (UPN) - to process user traffic

The CPN and UPN communicate with each other.

The following topology diagram explains about the communication between CPN and UPN:

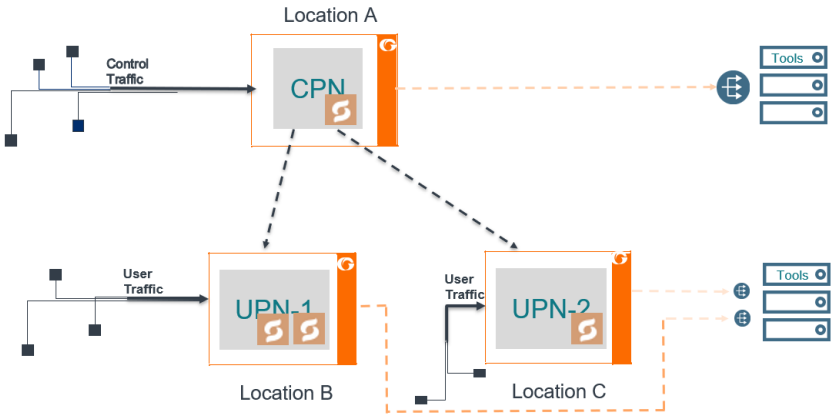

Figure 30-44: Topology diagram of CPN and UPN Connections

The user traffic is processed in UPN-1 and UPN-2 present in *Location B* and *Location C*, respectively. The control traffic is processed at CPN present in *Location A*, and few control parameters are sent to UPN-1 and UPN-2. *Location A* handles control traffic corresponding to the user traffic of *Location B* and *Location C*. The CPN and UPN can be present at the same location or at a different location.

When the GigaSMART engine in the CPN receives the control traffic, it generates a session table with information such as IMSI, MSISDN, IMEI, QCI, APN values and performs whitelisting, flow sampling and load balancing. UPN supports engine grouping, whereas CPN do not support engine grouping.

A GigaSMART Transport Agent (GTA) profile is created from CPN to transfer the appropriate control parameters to the corresponding UPN. The GTA profile instructs

the control processing node about the number of user groups it supports, and the source and destination port to route the packets.

The control parameters helps the UPN to generate a session table based on the Transport Agent (TA) packet information received at the GTA interface. GTA interface is a communication channel between the CPN and the UPN. The session table performs whitelisting, flow sampling and load balancing at the UPN. UPN correlates the user traffic by using the populated session table.

Limitations

- Engine Grouping on CPN is not suported.
- GTA communication is based on UDP.
- Supports only maximum of 16 GTA profiles.
- SXa, SXb packets are broadcasted to all the tool ports.
- A loop back IP connection with a router is required in-between the CPN and UPN if they are present within the same chassis.

The following table summarizes the required tasks for configuring the GigaSMART 3GPP CUPS:

S.No	Task	Refer to...
1.	Configure an IP interface to send and receive the Gigamon Transport Agent packets.	*IP Interfaces* on page 417
2.	Configure a GTA profile on Control Processing Plane (CPN) for routing the GTA packets.	*Configure GTA Profile* on page 893
3	Configure a first level map with port number 5000 on the UPN to receive the GTA packets. In the first level map configuration create a new rule to pass SXa, SXb packets through port number 8805.	
4.	Configure a node role for the control and user node, and attach a traffic profile to the CPN node in the GigaSMART engine.	*Associate Node Role and GTA Profiles* on page 894

Configure GTA Profile

To configure a GTA profile, do the following:

1. Select **Physical Nodes** from the Navigation pane.
2. Select the device on which you want to configure a GTA profile by clicking the Cluster ID of the device.
3. Select **GigaSMART** > **GTA Profile**.
4. Click **New**.
5. On the GTA Profile page, do the following:
   a. In the **Alias** field, enter a name for the GTA profile.

b. In the **Control Node** field, enter the IP address of the node that you want to assign as control processing node.

c. In the **User Node** field, enter the IP address of the node that you want to assign as the user processing node.

d. In the **Source Port** field, enter the port number from which the traffic is to be transferred. The value must be between 1 and 65535.

e. In the **Destination Port** field, enter the port number to which the traffic is received. The value must between 1 and 65535.

f. In the **Core Network Node**, you have the following options to enter the IP addrss value:

- List - You can enter multiple IP addresses by separating the IP addresses with a comma.
- Range- You can provide the range of IP address.
- Subnet- You can provide the IP address with subnet mask.

6. Click **OK**.

Associate Node Role and GTA Profiles

To associate the node role and GTA profiles, do the following:

1. From the device view, select **GigaSMART** > GigaSMART Groups.
2. Select a GigaSMART group alias for which you want to associate the GTA profile and click **Edit**.
3. Select anyone of the node type from the **Node Role** drop-down list in the **3GPP & CUPS** section.
4. Select the GTA profile that has to be associated with the node from the **GTA Profiles** drop-down list.
5. Click **OK**.

GTP Session Timeout

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to *Access GigaSMART from GigaVUE-FM* on page 754 for details.

In prior software versions, the complete GTP session timeout was eight hours. Starting with software version 4.2, the GTP session timeout is configurable, with eight hours as the default.

To configure the GTP session timeout, do the following:

1. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.
2. Click **New** to create a new GigaSMART Group or **Edit** to modify an existing one.
3. Under GigaSMART Parameters, go to GTP Flow.

4. Enter the timeout in the **Timeout** field. The following figure shows an example where the timeout value is set to 60.

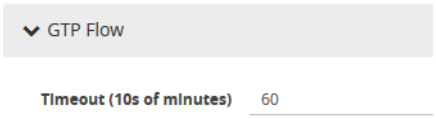

5. Click **Save**.

The GTP session timeout disconnects a GTP session if it has been inactive for the timeout value. The timeout can be configured as an integer from 1 to 6000, in increments of 10 minutes. The default value is 48, which is 480 minutes, which is 8 hours.

Priorities for Flow Rules and Maps

One virtual port can have multiple maps, and for each map, you can add multiple flow rules with different filtering attributes. The priorities for flow rules are as follows:

- a rule with a drop action has a higher priority than a rule with a pass action
- for the same pass or drop action, the priorities are IMSI, IMEI, or MSISDN in descending order

In a GTP session, if one IMSI, IMEI, or MSISDN rule is matched, the map is matched. For example, if any one of the following matches any rule shown in the following figure, map1 (which is a Second Level Flow Filter map) is matched:

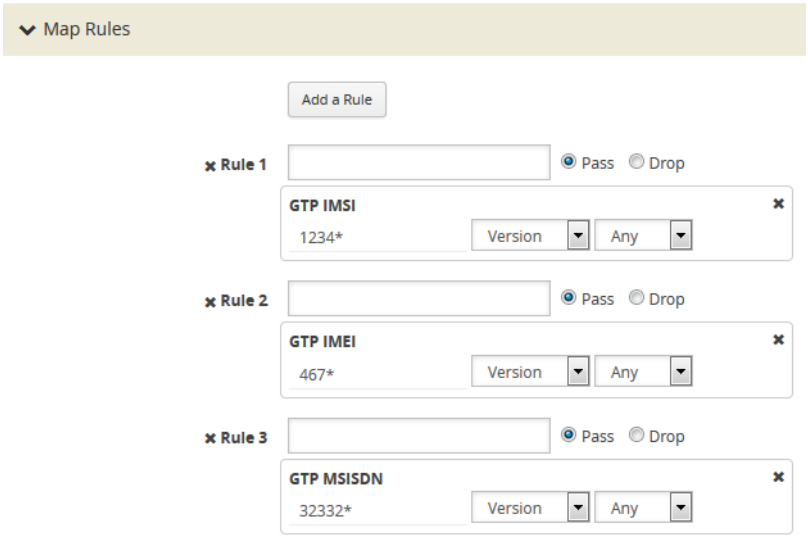

In addition, in one map, all drop rules are matched first and all pass rules are matched next.

For example, in a GTP session, if an IMSI matches the first rule in map1 and an IMEI matches the second rule in map1, because the drop rule has higher priority, the packet will be dropped:

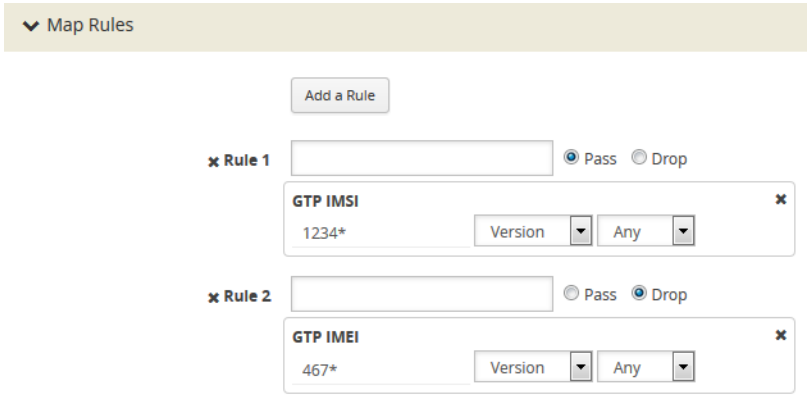

If multiple maps are matched, the map with the highest priority will be considered for further processing. For example, the Figure 30-45 shows the rule for map1 while the rule Figure 30-93 shows the rule for map2. In a GTP session, if an IMSI matches the first rule in map1 and an IMEI matches the first rule in map2, because map1 has higher priority, the packet will be passed.

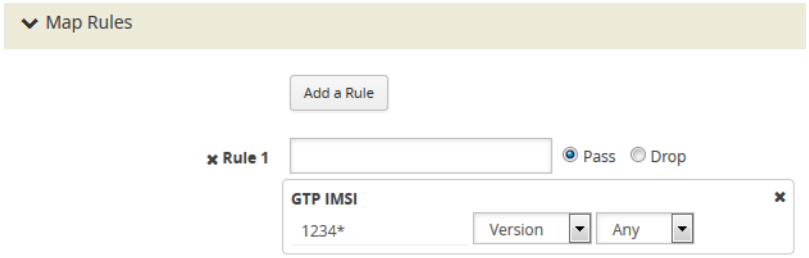

Figure 30-45: Rule in Map1

Figure 30-46: Rule in Map2

Configure GTP Correlation Examples

The following sections provide examples of GTP correlation and GTP load balancing. Refer to the following examples:

- *Example 1: Identifying High-Value and/or Roaming Subscribers Based on IMSIs* on page 897
- *Example 2: Identifying GTP Versions* on page 900
- *Example 3: Same Subscriber, Filter on Different Versions* on page 904
- *Example 4: Same Subscriber, Filter on Different Interfaces* on page 906
- *Example 5: EPC Filtering* on page 909
- *Example 6: EPC Filtering* on page 911

Example 1: Identifying High-Value and/or Roaming Subscribers Based on IMSIs

Use GTP correlation to identify high value subscribers based on an IMSI or group of IMSIs. GTP correlation keeps track of the IMSIs that you are interested in monitoring. It correlates them to the corresponding data/user-plane sessions for the subscriber and/or group of subscribers. Filtering on subscriber ID (IMSI) limits the amount of traffic that is sent to monitoring tools.

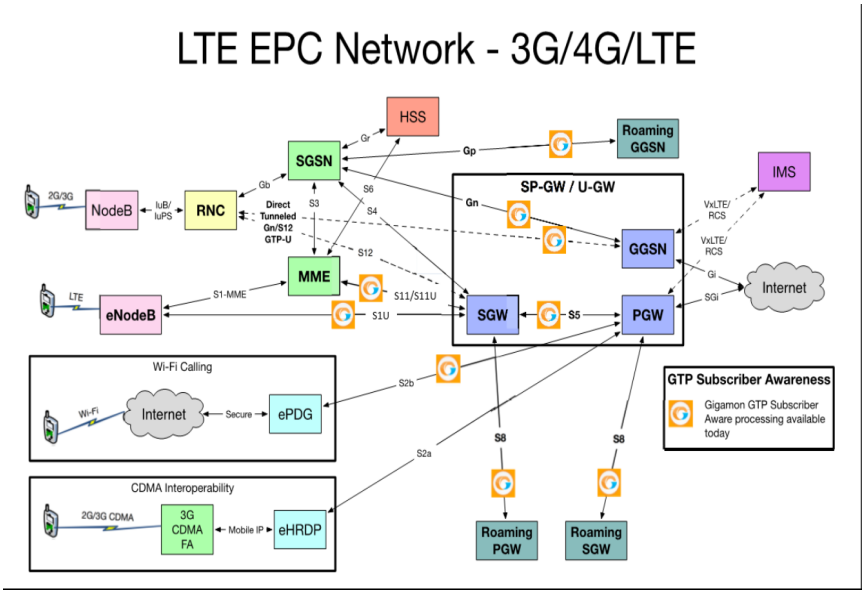

In Example 1, filter rules are configured to identify and forward all the traffic related to subscribers identified by an IMSI prefix. All traffic specific to the filtered IMSIs 22222222222223*, including GTP-c and GTP-u, is forwarded to a monitoring tool. A shared collector is configured to which traffic not matching the filters is sent.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to *Access GigaSMART from GigaVUE-FM* on page 754 for details.

Task	Description	UI Steps
1	Configure one network and two tool type of ports.	1. Select Ports > Ports > All Ports. 2. Click Quick Port Editor. 3. In the Quick View Editor set one port to Network and two ports to Tool. 4. Select Enable on each port. 5. Click OK. 6. Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type a name for the GigaSMART Group in the Alias field. 4. Click in the Port List field and select an engine port. 5. Under GTP Flow set the Timeout. The default is value is 48, which is 480 minutes. 6. Click Save.
3	Configure the GigaSMART operation and assign it to the GigaSMART group to enable GTP correlation.	1. From the device view, select GigaSMART > GigaSMART Operations (GSOP > GigaSMART Operation. 2. Type a name for the GigaSMART Group in the Alias field. 3. Click in the GigaSMART Operations (GSOP) field and select Flow Filtering. 4. Click OK.
4	Configure a virtual port and assign it to the same GigaSMART group.	1. From the device view, select GigaSMART > Virtual Ports. 2. Click New. 3. Type a name for the virtual port in the Alias field. 4. Click in the GigaSMART Groups field, and select the GigaSMART Group created in Task 2. 5. Click OK.

Task	Description	UI Steps
5	Create a first level map that directs GTP traffic from physical network port/s to the virtual port you created in the previous step. **NOTE:** In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. • Enter an alias • Select First Level for Type • Select By Rule for Subtype • Select a network port for the Source • Select the virtual port configured int Task 4 for the Destination **4.** Create Rule 1. **a.** Click Add a Rule **b.** Select Pass. **c.** Select Bi Directional. **d.** Select Port Source **e.** Set the source to 2123 **5.** Create Rule 2. **a.** Click Add Rule **b.** Select Pass. **c.** Select Bi Directional **d.** Select Port Source **e.** Set the source to 2153 **6.** Create Rule 3. **a.** Click Add Rule **b.** Select Pass. **c.** Select IPv4 Fragmentation **d.** Set Value to allFragNoFirst **7.** Click Save.
6	Create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, matches IMSIs specified by the flow rule, and sends matching traffic to physical tool ports.	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. • Enter an alias • Select Second Level for Type • Select Flow Filter for Subtype • Select a the virtual port configured in Task for the Source • Select the a tool port configured in Task 1 for the Destination • Select the GigaSMART Operation created in Task 3 from the GSOP list. **4.** Create a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select GTP IMSI. **d.** Enter 22222222222223* in the IMSI field. **5.** Click Save.

Task	Description	UI Steps
7	Add a shared collector for any unmatched data and send it to the second tool port.	1. Select Maps > Maps > Maps 2. Click New. 3. Configure the map. • Enter an alias • Select Second Level for Type • Select Collector for Subtype • Select a the virtual port configured in Task for the Source • Select the second tool port configured in Task 1 for the Destination 4. Click Save.
8	Display the configuration for Example 1.	To display the configuration for the GigaSMART Group: 1. From the device view, select GigaSMART > GigaSMART Groups. 2. Click on the alias for the GigaSMART Group to display the Quick View. To display the configuration for the GigaSMART Operation: 1. From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. 2. Click on the alias for the GigaSMART Operation to display the Quick View. To display the configuration for the maps: 1. Select Maps > Maps > Maps. 2. Click on a map alias to display the Quick View for the map.

Display GTP Correlation Flow Ops Report Statistics

To display GTP correlation statistics associated with the GigaSMART group, select **GigaSMART > GigaSMART Operations > Statistics**.

Example 2: Identifying GTP Versions

As part of GTP correlation, GigaVUE nodes also provide the flexibility to identify GTPv1 and GTPv2 messages. GTP version information is typically exchanged only as part of the control sessions. By correlating the control and user-plane sessions, GigaVUE nodes can identify, filter, and forward all sessions specific to a GTPv1 or v2 to one or more monitoring/analytic tools.

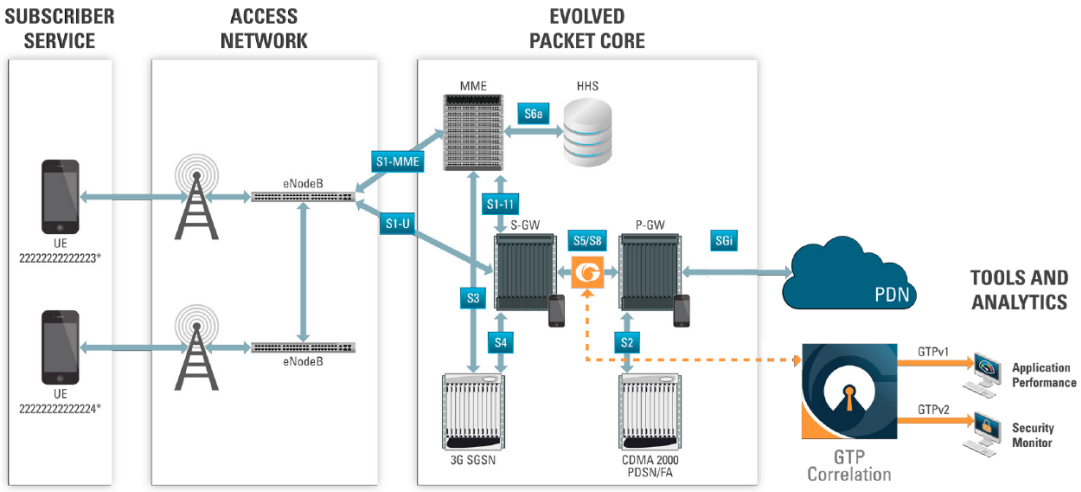

In Example 2, EMEI traffic is distributed based on GTP versions as follows:

- Filter and forward GTPv1 to a tool port
- Filter and forward GTPv2 to another tool port

Task	Description	UI Steps
1	Configure one network and two tool type of ports.	1. Select Ports > Ports > All Ports. 2. Click Quick Port Editor. 3. In the Quick View Editor set one port to Network and two ports to Tool. 4. Select Enable on each port. 5. Click OK. 6. Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type a name for the GigaSMART Group in the Alias field. 4. Click in the Port List field and select an engine port. 5. Under GTP Flow set the Timeout. The default is value is 48, which is 480 minutes. 6. Click Save.
3	Configure the GigaSMART operation and assign it to the GigaSMART group to enable GTP correlation.	1. From the device view, select GigaSMART > GigaSMART Operations > GigaSMART Operation. 2. Type a name for the GigaSMART Group in the Alias field. 3. Click in the GigaSMART Operations (GSOP) field and select Flow Filtering. 4. Click Save.

Task	Description	UI Steps
4	Configure a virtual port and assign it to the same GigaSMART group.	1. From the device view, select GigaSMART > Virtual Ports. 2. Click New. 3. Type a name for the virtual port in the Alias field. 4. Click in the GigaSMART Groups field, and select the GigaSMART Group created in Task 2. 5. Click Save.
5	Create a first level map that directs GTP traffic from physical network ports to the virtual port you created in the previous task. **NOTE:** In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	1. Select Maps > Maps > Maps 2. Click New. 3. Configure the map. • Enter an alias • Select First Level for Type • Select By Rule for Subtype • Select a network port for the Source • Select the virtual port configured int Task 4 for the Destination 4. Create Rule 1. a. Click Add a Rule b. Select Pass. c. Select Bi Directional. d. Select Port Source e. Set the source to 2123 5. Create Rule 2. a. Click Add Rule b. Select Pass. c. Select Bi Directional d. Select Port Source e. Set the source to 2153 6. Create Rule 3. a. Click Add Rule b. Select Pass. c. Select IPv4 Fragmentation d. Set Value to allFragNoFirst 7. Click Save.

Task	Description	UI Steps
6	Create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, matches IMEIs specified by the flow rule, and sends matching traffic to a tool port.	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. • Enter an alias • Select Second Level for Type • Select Flow Filter for Subtype • Select a the virtual port configured in Task for the Source • Select the a tool port configured in Task 1 for the Destination • Select the GigaSMART Operation created in Task 3 from the GSOP list. **4.** Create a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select GTP IMSI. **d.** Enter * in the IMSEI field and set Version to V1 **5.** Click Save.
7	Create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, matches IMEIs specified by the flow rule, and sends matching traffic to another tool port.	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. • Enter an alias • Select Second Level for Type • Select Flow Filter for Subtype • Select a the virtual port configured in Task for the Source • Select the second tool port configured in Task 1 for the Destination • Select the GigaSMART Operation created in Task 3 from the GSOP list. **4.** Create a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select GTP IMSEI. **d.** Enter * in the IMSI field and set Version to V2 **5.** Click Save.
8	Display the configuration for Example 2.	To display the configuration for the GigaSMART Group: **1.** From the device view, select GigaSMART > GigaSMART Groups > Maps. **2.** Click on the alias for the GigaSMART Group to display the Quick View. To display the configuration for the GigaSMART Group: **1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations **2.** Click on the alias for the GigaSMART Operation to display the Quick View. To display the configuration for the maps: **1.** Select Maps > Maps > Maps. **2.** Click on the map alias to display the Quick View for the map.

Example 3: Same Subscriber, Filter on Different Versions

In this example, traffic from the same subscriber is forwarded to two different load balancing groups based on version. GTP version 1 traffic is sent to one load balancing group and GTP version 2 traffic is sent to another load balancing group.

Task	Description	UI Steps
1	Configure one network and multiple tool type of ports.	1. Select Ports > Ports > All Ports. 2. Click Quick Port Editor. 3. In the Quick View Editor set one port to Network and multiple ports to Tool. For example, set 1/2/g1 as a Network port and ports 1/2/g5 through 1/2/g9 as Tool ports. 4. Select Enable on each port. 5. Click OK. 6. Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type a name for the GigaSMART Group in the Alias field. 4. Click in the Port List field and select an engine port. 5. Click Save.
3	Configure the GigaSMART operation and assign it to the GigaSMART group.	1. From the device view, select GigaSMART > GigaSMART Operations > GigaSMART Operations. 2. Type a name for the GigaSMART Group in the Alias field. 3. Click in the GigaSMART Operations field and select Flow Filtering. 4. Click in the GigaSMART Operations filed and select Load Balancing. 5. Configure the Load Balancing as follows:. • Select Stateful • Select GTP for Type • Selecting Hashing • Select IMSI 6. Click Save.
4	Configure a virtual port and assign it to the GigaSMART group.	1. From the device view, select GigaSMART > Virtual Ports. 2. Click New. 3. Type a name for the virtual port in the Alias field. 4. Click in the GigaSMART Groups field, and select the GigaSMART Group created in Task 2. 5. Click Save.

Task	Description	UI Steps
5	Create two port groups (one for version 1 traffic and one for version 2 traffic) and enable load balancing on the port groups.	1. Select Ports > Port Groups > All Port Groups. 2. Click New. 3. Type a name for the GigaSMART Group in the Alias field. For example, pglbv1. 4. Select SMART Load Balancing to enable load balancing 5. Click in the Ports field and select half the tool ports configured in Task 1. 6. Repeat steps 2 through 4, creating a second port group with the other ports configured in Task 1
6	Create an ingress (first level) map.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Enter an alias. For example, map1_1. • Select First Level for Type • Select By Rule for Subtype • Select a network port for the Source • Select the virtual port configured in Task 4 for the Destination 4. Click Add a Rule to create a rule. • Select pass • Select MAC Destination • Enter a MAC address. For example, **00:a0:d1:e1:02:01** • Enter a MAC mask. For example, 0000.0000.0000 5. Click Save.
7	Create a second level map.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Enter an alias. For example, map2_1. • Select Second Level for Type • Select Flow Filter for Subtype • Select the virtual port configured in Task 4 for the Source • Select the first port group for Destination • Select the GigaSMART Operation configured in Task 3 from the GSOP list. 4. Click Add a Rule to create a rule. • Select Pass • Select GTP IMSI • Enter * in the IMSI field • Select V1 for Version 5. Click Save.

Task	Description	UI Steps
8	Create another second level map.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Enter an alias. For example, map2_2. • Select Second Level for Type • Select Flow Filter for Subtype • Select the virtual port configured in Task 4 for the Source • Select the second port group for Destination • Select the GigaSMART Operation configured in Task 3 from the GSOP list. 4. Click Add a Rule to create a rule. • Select Pass • Select GTP IMSI • Enter * in the IMSI field • Select V2 for Version 5. Click Save.

Example 4: Same Subscriber, Filter on Different Interfaces

In this example, traffic from the same subscriber is forwarded to two different load balancing groups based on interface. In this example, VLANs 1601 and 1602 are from S5/S8 interface and VLANs 1611 and 1612 are from S11/S1-U interface. The first level maps split the VLAN traffic to different virtual ports. The second level maps send the traffic to different load balancing groups.

Task	Description	UI Steps
1	Configure one network and multiple tool type of ports.	1. Select Ports > Ports > All Ports. 2. Click Quick Port Editor. 3. In the Quick View Editor set one port to Network and multiple ports to Tool. For example, set 1/2/g1 as a Network port and ports 1/2/g5 through 1/2/g9 as Tool ports. 4. Select Enable on each port. 5. Click OK. 6. Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type a name for the GigaSMART Group in the Alias field. 4. Click in the Port List field and select an engine port. 5. Click Save.

Task	Description	UI Steps
3	Configure the GigaSMART operation and assign it to the GigaSMART group.	1. From the device view, select GigaSMART > GigaSMART Operations > GigaSMART Operations. 2. Type a name for the GigaSMART Group in the Alias field. 3. Click in the GigaSMART Operations field and select Flow Filtering. 4. Click in the GigaSMART Operations filed and select Load Balancing. 5. Configure Load Balancing as follows: • Select Stateful • Select GTP for Type • Selecting Hashing • Select IMSI 6. Click Save.
4	Configure virtual ports and associate them with the GigaSMART group.	1. From the device view, select GigaSMART > Virtual Ports. 2. Click New. 3. Type a name for the virtual port in the Alias field. 4. Click in the GigaSMART Groups field, and select the GigaSMART Group created in Task 2. 5. Click Save. 6. Repeat steps 1 through 5 to create a second virtual port.
5	Create two port groups (one for version 1 traffic and one for version 2 traffic) and enable load balancing on the port groups.	1. Select Ports > Port Groups > All Port Groups. 2. Click New. 3. Type a name for the GigaSMART Group in the Alias field. For example, pglbv1. 4. Select SMART Load Balancing to enable load balancing 5. Click in the Ports field and select half the tool ports configured in Task 1. 6. Repeat steps 2 through 4, creating a second port group with the other ports configured in Task 1
6	Create a first level map.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Enter an alias. For example, map1_1. • Select First Level for Type • Select By Rule for Subtype • Select a network port for the Source • Select the first virtual port configured in Task 4 for the Destination 4. Click Add a Rule to create a rule. • Select Pass • Select VLAN • Enter 1601 for the Min value • Enter 1602 for the Max value 5. Click Save.

Task	Description	UI Steps
7	Create another first level map.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Enter an alias. For example, map1_2. • Select First Level for Type • Select By Rule for Subtype • Select a network port for the Source • Select the second virtual port configured in Task 4 for the Destination 4. Click Add a Rule to create a rule. • Select Pass • Select VLAN • Enter 1611 for the Min value • Enter 1611 for the Max value 5. Click Save.
8	Create a second level map.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Enter an alias. For example, map2_1. • Select Second Level for Type • Select Flow Filter for Subtype • Select the virtual port configured in Task 4 for the Source • Select the first port group for Destination • Select the GigaSMART Operation configured in Task 3 from the GSOP list. 4. Click Add a Rule to create a rule. • Select Pass • Select GTP IMSI • Enter * in the IMSI field 5. Click Save.
9	Create another second level map.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Enter an alias. For example, map2_2. • Select Second Level for Type • Select Flow Filter for Subtype • Select the second virtual port configured in Task 4 for the Source • Select the second port group for Destination • Select the GigaSMART Operation configured in Task 3 from the GSOP list. 4. Click Add a Rule to create a rule. • Select Pass • Select GTP IMSI • Enter * in the IMSI field 5. Click Save.

Example 5: EPC Filtering

In this example, traffic for all subscribers on interfaces S11/S1-U and Gn/Gp is sent to the same load balancing group. All other traffic is dropped.

Task	Description	UI Steps
1	Configure one network and two tool type of ports.	1. Select Ports > Ports > All Ports. 2. Click Quick Port Editor. 3. In the Quick View Editor set one port to Network and two ports to Tool. For example, set 1/2/g1 as a Network port and ports 1/2/g5 and 1/2/g6 as Tool ports. 4. Select Enable on each port. 5. Click OK. 6. Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type a name for the GigaSMART Group in the Alias field. 4. Click in the Port List field and select an engine port. 5. Click Save.
3	Configure the GigaSMART operation and assign it to the GigaSMART group.	1. From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. 2. Type a name for the GigaSMART Group in the Alias field. 3. Click in the GigaSMART Operations field and select Flow Filtering. 4. Click in the GigaSMART Operations (GSOP) field and select Load Balancing. 5. Configure the Load Balancing as follows. • Select Stateful • Select GTP for Type • Selecting Hashing • Select IMSI 6. Click Save.
4	Configure a virtual port and assign it to the GigaSMART group.	1. From the device view, select GigaSMART > Virtual Ports. 2. Click New. 3. Type a name for the virtual port in the Alias field. 4. Click in the GigaSMART Groups field, and select the GigaSMART Group created in Task 2. 5. Click Save.
5	Create a port group enable load balancing on the port group.	1. Select Ports > Port Groups > All Port Groups. 2. Click New. 3. Type a name for the GigaSMART Group in the Alias field. For example, pglbv1. 4. Select SMART Load Balancing to enable load balancing 5. Click Save.

Task	Description	UI Steps
6	Create an ingress (first level) map. **NOTE:** In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. • Enter an alias • Select First Level for Type • Select By Rule for Subtype • Select a network port for the Source • Select the virtual port configured int Task 4 for the Destination **4.** Create Rule 1. **a.** Click Add a Rule **b.** Select Pass. **c.** Select Bi Directional. **d.** Select Port Source **e.** Set the source to 2123 **5.** Create Rule 2. **a.** Click Add Rule **b.** Select Pass. **c.** Select Bi Directional **d.** Select Port Source **e.** Set the source to 2152 **6.** Click Save.
7	Create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, matches IMSIs specified by the flow rules, and sends matching traffic to physical tool ports.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. • Enter an alias. • Select Second Level for Type • Select Flow Filter for Subtype • Select the virtual port configured in Task 4 for the Source • Select the port group for Destination • Select the GigaSMART Operation configured in Task 3 from the GSOP list. **4.** Create the rule 1. **a.** Click Add a Rule **b.** Select Pass **c.** Select GTP IMSI **d.** Enter * in the IMSI field **e.** Select Interface and set it to Gg/Gp **5.** Create rule 2. **a.** Click Add a Rule **b.** Select Pass **c.** Select GTP IMSI **d.** Enter * in the IMSI field **e.** Select Interface and set it to S11/S1U **6.** Click Save.

Example 6: EPC Filtering

In this example, traffic for all subscribers from all interfaces except S5/S8 is sent to the same load balancing group. Traffic from the S5/S8 interface is dropped.

Step	Description	UI Steps
1	Configure one network and two tool type of ports.	1. Select Ports > Ports > All Ports. 2. Click Quick Port Editor. 3. In the Quick View Editor set one port to Network and two ports to Tool. For example, set 1/2/g1 as a Network port and ports 1/2/g5 and 1/2/g6 as Tool ports. 4. Select Enable on each port. 5. Click OK. 6. Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with a GigaSMART engine port.	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type a name for the GigaSMART Group in the Alias field. 4. Click in the Port List field and select an engine port. 5. Click Save.
3	Configure the GigaSMART operation and assign it to the GigaSMART group.	1. From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. 2. Type a name for the GigaSMART Group in the Alias field. 3. Click in the GigaSMART Operations (GSOP) field and select Flow Filtering. 4. Click in the GigaSMART Operations (GSOP) field and select Load Balancing. 5. Configure Load Balancing as follows: • Select Stateful • Select GTP for Type • Selecting Hashing • Select IMSI 6. Click Save.
4	Configure a virtual port and assign it to the GigaSMART group.	1. From the device view, select GigaSMART > Virtual Ports. 2. Click New. 3. Type a name for the virtual port in the Alias field. 4. Click in the GigaSMART Groups field, and select the GigaSMART Group created in Task 2. 5. Click Save.
5	Create a port group and enable load balancing on the port group.	1. Select Ports > Port Groups > All Port Groups. 2. Click New. 3. Type a name for the GigaSMART Group in the Alias field. For example, pglbv1. 4. Select SMART Load Balancing to enable load balancing 5. Click Save.

Step	Description	UI Steps
6	Create an ingress (first level) map. **NOTE:** In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**1.** Select Maps > Maps > Maps **2.** Click New. **3.** Configure the map. • Enter an alias • Select First Level for Type • Select By Rule for Subtype • Select a network port for the Source • Select the virtual port configured int Task 4 for the Destination **4.** Create Rule 1. **a.** Click Add a Rule **b.** Select Pass. **c.** Select Bi Directional. **d.** Select Port Source **e.** Set the source to 2123 **5.** Create Rule 2. **a.** Click Add Rule **b.** Select Pass. **c.** Select Bi Directional **d.** Select Port Source **e.** Set the source to 2152 **6.** Click Save.
7	Create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, matches IMSIs specified by the flow rules, and sends matching traffic to physical tool ports.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. • Enter an alias. • Select Second Level for Type • Select Flow Filter for Subtype • Select the virtual port configured in Task 4 for the Source • Select the port group for Destination • Select the GigaSMART Operation configured in Task 3 from the GSOP list. **4.** Create the rule 1. **a.** Click Add a Rule **b.** Select drop **c.** Select GTP IMSI **d.** Enter * in the IMSI field **e.** Select Interface and set it to S5/S8 **5.** Create rule 2. **a.** Click Add a Rule **b.** Select Pass **c.** Select GTP IMSI **d.** Enter * in the IMSI field **6.** Click Save.

GigaSMART GTP Whitelisting and GTP Flow Sampling

Required Licenses: GTP Filtering & Correlation and FlowVUE

Use GTP whitelisting and GTP flow sampling to provide subsets of GTP correlated flows to tools. GTP whitelisting selects specific subscribers based on IMSI, while GTP flow sampling uses map rules to select subscribers. Starting in software version 4.8, GigaSMART supports GTP overlap mapping, which combines both whitelisting and flow sampling maps as part of a map group. Refer to *GTP Overlap Flow Sampling Maps* on page 950

Starting in software version 4.5, a GigaSMART group (gsgroup) associated with GTP applications can have multiple GigaSMART engine port members (**e** ports), up to four, forming an engine group. Refer to *GTP Scaling* on page 959.

Refer to the following sections:

- *GTP Whitelisting* on page 913
- *GTP Flow Sampling* on page 919
- *GTP Whitelisting and GTP Flow Sampling Examples* on page 924

GTP Whitelisting

GTP whitelisting selects specific subscribers based on IMSI. The whitelist contains up to 500,000 subscriber IMSIs. For subscribers in the whitelist, 100% of their traffic is always sent to a specified tool port.

For example, when a subscriber session comes in, GTP whitelisting checks the IMSI of the subscriber. If the incoming IMSI matches an IMSI in the whitelist, the session is sent to the tool port or load balancing group specified in the whitelist map.

Starting in software version 4.7, GTP whitelisting is supported in a cluster. Refer to *GTP Whitelisting in a Cluster* on page 919 for more information.

Create Whitelist

Subscriber IMSIs are added to a whitelist that can contain up to 500,000 subscriber IMSIs. One active whitelist per GigaSMART group is supported.

Entries in the whitelist can be added one at a time or whitelist files containing multiple IMSIs can be created and downloaded. Entries are added by using the GTP Whitelist page by selecting **GigaSMART > GTP Whitelist**. The GTP Whitelist page shows alias for the currently configured GTP Whitelists, the IMSI count for each Whitelist and the GigaSMART Group associated with the GTP Whitelist. (Refer to Figure 30-47 on page 914.) The GTP Whitelist is associated with the GigaSMART group by specifying its alias in the **GTP Whitelist Alias** field in GigaSMART Group configuration page.

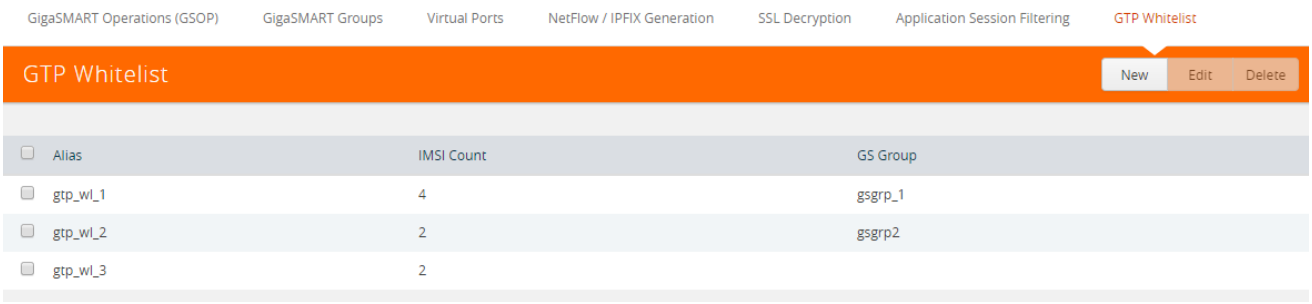

Figure 30-47: GTP Whitelist Page

and then clicking **New**. Figure 30-48 shows an example of an a whitelist with the alias MyIMSIs and ready to fetch multiple IMSIs from a whitelist file at remote location.

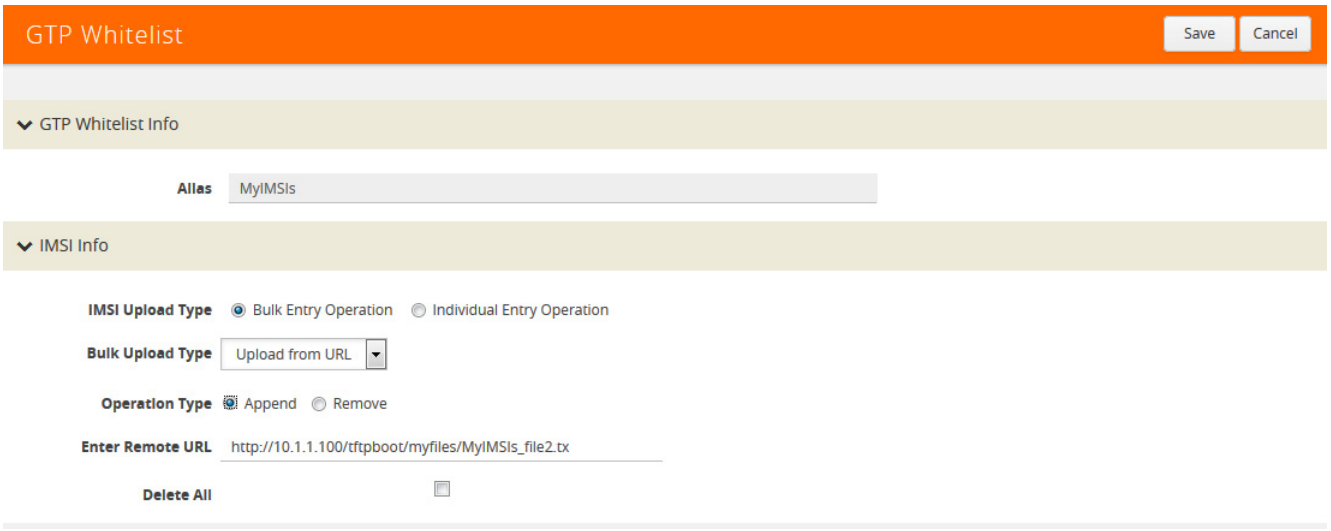

Figure 30-48: GTP Whitelist Bulk Upload

An individual IMSI is added by selecting Individual **Entry Operation** and specifying the IMSI in the **Individual IMSI Entry** field as shown in Figure 30-49.

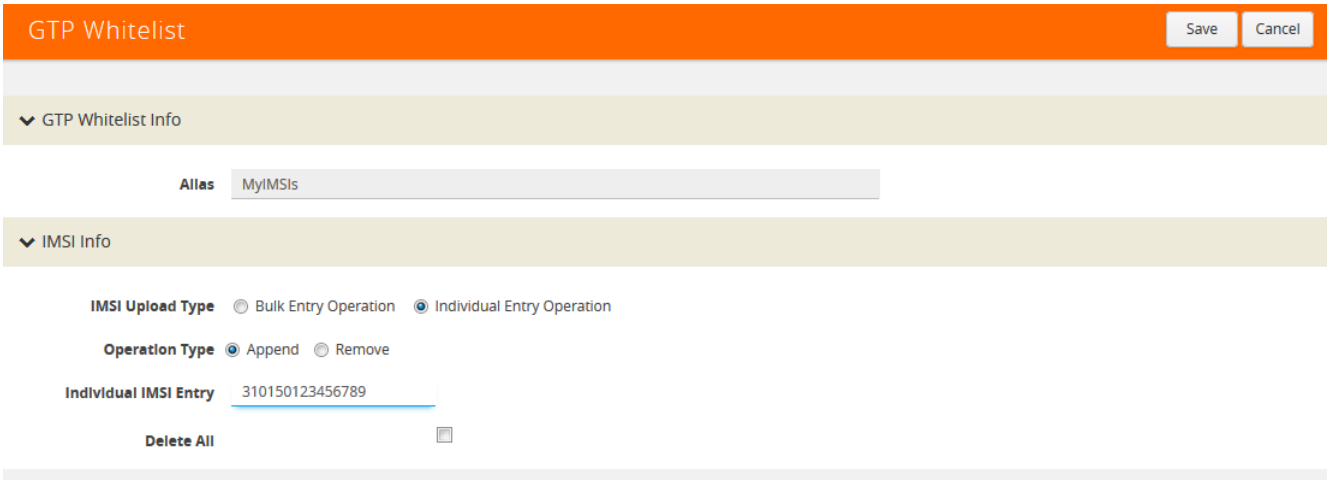

Figure 30-49: GTP Whitelist Add Individual IMSI

The IMSIs in whitelist files must be distinct entries, with one IMSI on each line of a file and a maximum of 20,000 entries in each file. This means that 25 files of 20,000 entries will be needed to populate the whitelist to its capacity. Wildcards are not supported in whitelist files.

Whitelist files must have a filename with a .txt suffix. Use the GTP Whitelist page to fetch the entries from a whitelist file at a specified location, using one of the following formats, which are specified in the **Enter Remote URL** field when **Bulk Entry Operation** is selected and the **Bulk Upload Type** is **Upload from URL**:

- http://IPaddress/path/filename.txt
- scp://username:password@IPaddress:/path/filename.txt
- tftp://IPaddress/path/filename.txt

To fetch a whitelist file from a local location, select **File Upload** for **Bulk Upload Type** and use the **Browse** button to select the file.

When a whitelist file is downloaded, the entries are compared to the whitelist on the node. There may be new entries in the file that might already be part of the existing whitelist. GigaSMART will add the new, non-duplicate entries to the whitelist, without rejecting the entire file.

If the current number of entries in the whitelist plus the new entries in the whitelist file is greater than the whitelist capacity of 500,000 IMSIs, the **Append** operation will fail and the new entry or the entries from the new whitelist file will not be added.

GTP whitelisting does not use map rules like GTP flow sampling does. The whitelist is associated with a GigaSMART group, GigaSMART operation, and second level maps, called whitelist maps.

You can create multiple whitelists, each with 500,000 IMSIs. However, even though you can create multiple whitelists, you can only have one active whitelist in use at a time in a GigaSMART group, a GigaSMART operation, and whitelist maps. To switch from one whitelist to another, you must first either delete or destroy the currently active whitelist before you can make another whitelist active. Refer to *Delete Whitelist* on page 918.

For the sequences of steps to create a whitelist with the UI, refer to the configuration example for whitelisting in *Example 1: GTP Whitelisting* on page 925.

Configure Whitelist Maps

The whitelist maps are configured per GigaSMART group. Each whitelist map, associated with the same vport, uses the same underlying whitelist.

Up to ten (10) whitelist maps are supported. Multiple whitelist maps provide a granular selection of tool ports for whitelisting. Using multiple maps, traffic can be segregated and sent to multiple destinations. Whitelist map rules allow you to select the subset of IMSIs sent to a particular tool.

Each whitelist map can contain up to four rules. The rules specify the type of traffic to be whitelisted by that map. Within any single map, the rules are evaluated in order. The

rules in the first map have a higher priority than the rules in the second, third, and subsequent maps.

The rules will specify either an Evolved Packet Core (EPC) interface type (refer to Figure 30-50) or a GTP version (refer to Figure 30-51) as the attribute to match. An Access Point Name (APN) (refer to Figure 30-50 and Figure 30-50) can also be specified in a rule of a Second Level Flow Whitelist map, either by itself, or preceding the EPC interface type or in combination with the GTP version.

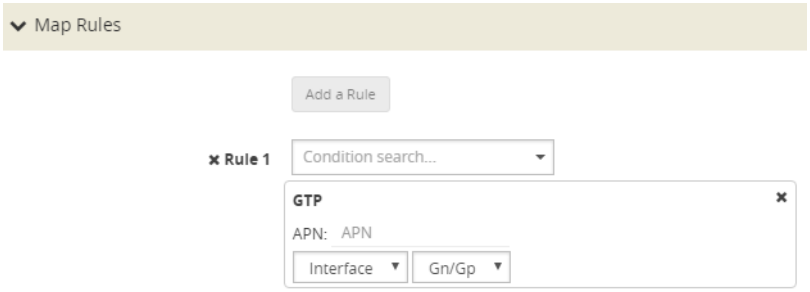

Figure 30-50: GTP Rule for EPC Interface Type

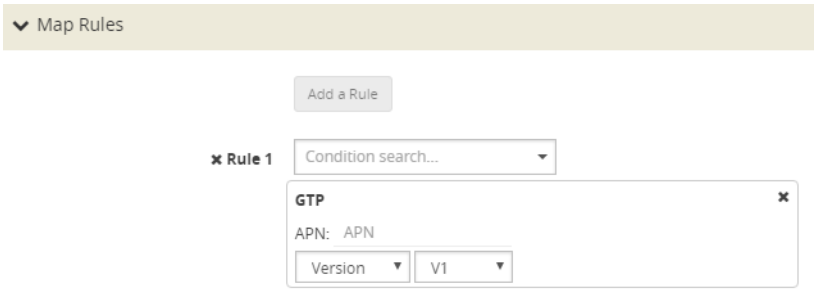

Figure 30-51: GTP Rule for GTP Version

For APN, you must specify a pattern (a name) to match. Use APN to direct the traffic that matches the pattern to a specific tool.

GTP version and EPC interface are mutually exclusive. A mix of versions and interface types across whitelist maps, associated with the same vport, is not supported. For example, you can configure two whitelist maps with one map specifying a rule for version 1 and another map specifying a rule for version 2, OR four whitelist maps with each map specifying a rule for each interface type (Gn, S11, S5, and S10). For more information on interfaces, refer to *Supported Interfaces* on page 889.

An APN pattern is for example, three.co.uk. Wildcard prefixes and suffixes are supported, for example, *mobile.com or *ims*. The pattern can be specified in up to 100 case-insensitive alphanumeric characters and can include the following special characters: period (.), hyphen (-), and wildcard (*). A standalone wildcard (*) is not allowed for APN.

NOTE: APN is not supported on GigaVUE-HB1.

Each new subscriber session will be evaluated by the whitelist maps in the order of priority, which, by default, is the order in which the maps were created.

When a subscriber session comes in, GTP whitelisting will check the IMSI of the subscriber. If the IMSI is present in the whitelist, the rules in the first whitelist map is evaluated to qualify the match further. Otherwise, the packet is evaluated against the rules in the subsequent whitelist maps for a possible match.

For example, with one whitelist map having a rule specifying GTP version 1 and another whitelist map having a rule specifying GTP version 2, when a subscriber session comes in, GTP whitelisting will check the IMSI of the subscriber. if the IMSI is present in the whitelist and if there is a match to version 1, the session (100% of subscriber packets) will be forwarded to the tool port, GigaStream, or load balancing group specified in the whitelist map. If there is not a match to version 1, the next map is evaluated. If there is a match to version 2 in the next map, the session will be forwarded to the tool port, GigaStream, or load balancing group specified in the second whitelist map.

NOTE: Both maps can specify the same destination.

Rules can be added to, or deleted from, a whitelist map. Use the **Add a Rule** button to add a new whitelist rule (a pass rule). Click **x** to delete a rule. A rule in a whitelist map cannot be edited. To edit a rule, first delete it, then recreate it.

The default map configuration in which neither GTP version, EPC interface, or APN is specified in the map, continues to be supported. If the incoming IMSI matches an IMSI in the whitelist, the session will be sent to the tool port, GigaStream, or load balancing group specified in the whitelist map.

Whitelist maps cannot contain any other rules such as GigaSMART rules (gsrule), flow filtering rules (flowrule), or flow sampling rules (flowsample).

GTP whitelist-based forwarding is performed prior to GTP flow sampling (rule-based flow sampling) and GTP flow filtering.

NOTE: For GTP second level maps, a maximum of fifteen maps can be attached to a vport. For example, for the same vport you can have five whitelist maps and ten flow sampling maps, or ten whitelist maps, four flow sampling maps, and one flow filtering map. In addition, you can have a collector map, which is not counted.

For the steps to create a whitelist map with the UI, refer to the configuration example for whitelisting in *GTP Whitelisting and GTP Flow Sampling Examples* on page 924.

Change Priority of Whitelist Maps

Use the **Priority** field in the map to change the priority of whitelist maps.

Delete Whitelist Maps

When a whitelist map is deleted, the priority of the remaining whitelist maps will be re-prioritized. For example, if the first whitelist map is deleted, the second whitelist map will increase in priority.

For the deleted whitelist map, the traffic associated with the rules in the map will be reevaluated and then passed to subsequent maps.

When a whitelist map is re-prioritized, the existing sessions will be reevaluated according to the new priority of the map. The traffic associated with the rules in the map will be reevaluated and then passed to subsequent maps.

When the last whitelist map is deleted, the traffic associated with the rules in the map will also be reevaluated before being passed to subsequent maps. But the traffic associated with the rules in maps that were not matched, will not be reevaluated because that traffic was already passed to subsequent maps.

Apply Whitelist

When a single whitelist entry is added, whitelisting is applied for new as well as existing subscribers. When a new whitelist file is fetched, whitelisting is applied only for new subscribers.

Whitelisted traffic is then sent to the port or load balancing group specified in the whitelist map.

Delete an Entry from Whitelist

Entries in the whitelist can be deleted one at a time. Each entry is a single IMSI.

When a whitelist entry is deleted, the session associated with the whitelist entry stays active and traffic is still sent to the whitelist map. The whitelist session will not be reevaluated or passed to subsequent maps.

To delete a single entry from the whitelist, select **Individual Entry Operation**, set **Remove** as the **Operation Type,** and enter the IMSI in the **Individual IMSI Entry** field.

Delete Multiple Entries from Whitelist

Multiple IMSIs can be deleted from the whitelist. Specify the IMSIs to be deleted in a whitelist file, which can contain up to 20,000 IMSIs.

Whitelist files must have a filename with a .txt suffix. To remove multiple entries from the whitelist, select **Bulk Entry Operation** and set **Remove** as the **Operation Type**.

Delete Whitelist

The entire whitelist can be deleted using one of the following two options:

- Delete the whitelist by deleting all the IMSI entries. With this option, you do not have to delete the whitelist map, GigaSMART operation, or disassociate the GigaSMART group from the whitelist. To delete all the IMSI entries, select **Delete All**.
- Destroy the whitelist. With this option, you must first delete the whitelist map, GigaSMART operation, and disassociate the GigaSMART group from the whitelist before deleting the whitelist.

Destroy Whitelist

To destroy a whitelist, use the following sequence:

Task	UI Steps
Delete the whitelist map	**1.** Select Maps > Maps > Map. **2.** Select the whitelist map. **3.** Click Delete.
Delete the GigaSMART Operation	**1.** From the device view, select GigaSMART > GigaSMART Operation(s) > GigaSMART Operation. **2.** Select the GigaSMART Operation. **3.** Click Delete.
Disassociate the GigaSMART group from the whitelist. (You do not need to delete the GigaSMART group.)	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Select the GigaSMART group. **3.** Click Edit. **4.** Under GigaSMART Parameters, go the GTP Whitelist and set GTP Whitelist Alias to None.
Destroy (delete) the whitelist	**1.** From the device view, select GigaSMART > GigaSMART GTP Whitelist. **2.** Select the GTP Whitelist. **3.** Click Delete.

GTP Whitelisting in a Cluster

The whitelist (all whitelist files) must reside on the master node of the cluster. The non-master nodes receive a copy of the whitelist from the master. Updates to the whitelist are synchronized from the master to the non-master nodes. If a non-master node leaves the cluster and rejoins, its whitelist will be resynchronized.

GigaVUE-HB1 nodes do not support GTP whitelisting in a cluster due to their limited storage. If there are GigaVUE TA Series nodes in the cluster, they will not receive a copy of the whitelist.

GTP Flow Sampling

GTP flow sampling samples a configured percentage of GTP sessions. GTP flow sampling uses map rules to select subscribers and then forward a percentage of the packets to tool ports.

Pass rules are defined in flow sampling maps. Each rule contains some combination of IMSI, IMEI, and MSISDN numbers or patterns, Evolved Packet Core (EPC) interface type, GTP version, Access Point Name (APN), or QoS Class Identifier (QCI), as well as a percentage to sample. The flow is sampled to see if it matches a rule. The percentage of the subscriber sessions matching each rule are selected.

Map rules specify the type of traffic to be flow sampled by that map. For each new session, map rules are evaluated in top-down order of decreasing priority. If there is a match, the indicated percentage of the subscriber session is either accepted or

rejected. If accepted, the traffic is sent to the tool port or load balancing group specified in the map. If rejected, the traffic is dropped. If there is not a match to a rule, the traffic is passed to subsequent maps.

Starting in software version 4.6, GTP load balancing in a cluster is supported for GTP flow sampling. For an example of GTP load balancing in a cluster, refer to *GTP Whitelisting and GTP Flow Sampling Examples* on page 924.

About Flow Sampling Rules and Maps

Flow sampling rules are configured in maps called flow sampling maps. Up to ten (10) flow sampling maps per GigaSMART group are supported. Each flow sampling map supports up to 20 flow sampling rules, for a maximum of 200 rules per GigaSMART group.

GTP flow sampling (rule-based flow sampling) is performed after GTP whitelist-based forwarding but before GTP flow filtering. So, flow sampling maps have a priority lower than whitelist maps and higher than flow filtering maps.

NOTE: For GTP second level maps, a maximum of fifteen maps can be attached to a vport. For example, for the same vport you can have one whitelist map and ten flow sampling maps, or ten whitelist map, four flow sampling maps, and one flow filtering map. In addition, you can have a collector map, which is not counted.

In the flow sampling maps, the rules in the first map have a higher priority than the rules in the second, third, and subsequent maps. Within any single map, rules are evaluated in order.

Rules can be added to, deleted from, or inserted into a flow sampling map when the subtype selected for a **Second Level** map is **Flow Sample**. Suffix wildcarding, such as IMSI 100*, is supported in the flow sampling map rules.

Use the **Add a Rule** button in the Maps page to add a new flow sampling rule (a pass rule). Specify IMSI, IMEI, or MSISDN subscriber IDs, as well as the percentage of the flow to be sampled. The percentage is a range from 1 to 100%. Use 0% to drop sampled data.

A rule can specify other packet attributes, such as an EPC interface type or GTP version. An APN pattern can also be specified in a rule, either by itself or preceding the EPC interface or GTP version. A QCI value can be specified, but only in combination with an APN pattern.

EPC interface and GTP version are mutually exclusive. They can be specified in a flow sampling rule, but not both in a single rule. The supported interface types for filtering are: Gn/Gp, S11/S1-U, S5/S8, S10, or S2B. The supported versions for filtering are 1 or 2. For example, you can send version 1 traffic to one tool port and version 2 traffic to another tool port. For more information on interfaces, refer to *Supported Interfaces* on page 889.

For APN, specify a pattern (a name) to match, for example, three.co.uk. Wildcard prefixes and suffixes are supported, for example, *mobile.com or *ims*. The pattern can be specified in up to 100 case-insensitive alphanumeric characters and can include the following special characters: period (.), hyphen (-), and wildcard (*).

NOTE: APN is not supported on GigaVUE-HB1.

QCI is a mechanism used in Long Term Evolution (LF TE) networks to ensure bearer traffic is allocated to the appropriate Quality of Service (QoS). For QCI, specify a value from 0 to 255. Wildcard prefixes and suffixes are not supported.

Use APN and QCI to send traffic that matches a certain APN pattern or that belongs to a certain bearer with a certain QCI to specified tool ports, based on the sampling percentage.

Click the **x** next to a rule to delete a specific rule. Rules are identified by a priority ID, which indicates the order of rules in a flow sampling map. For example, if a map has 12 pass flow sampling rules, there will be 12 priority IDs. Figure 30-52 shows rules in a Flow Sampling map and their priority IDs.

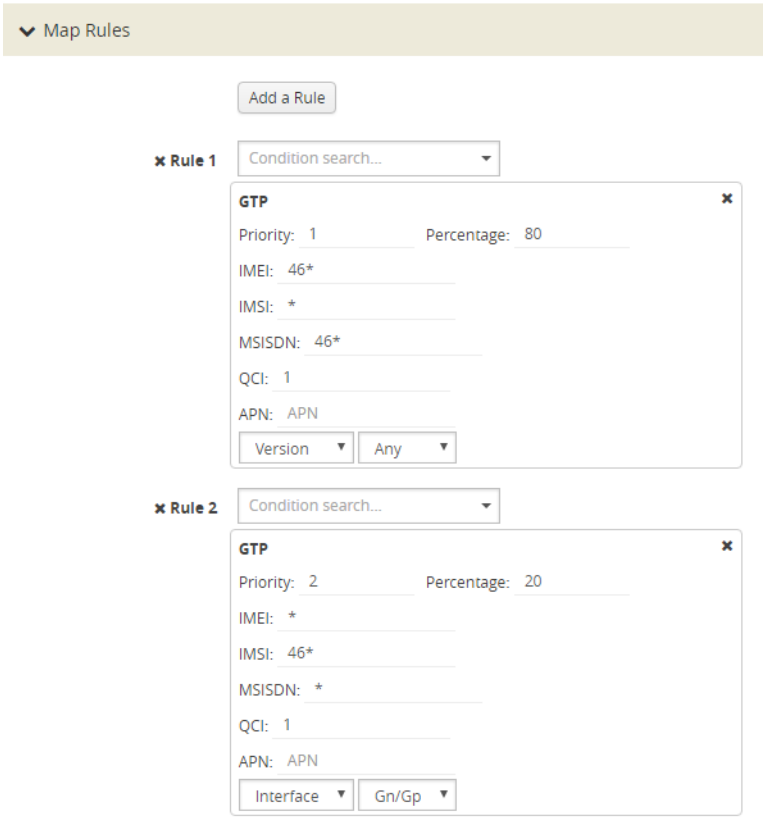

Figure 30-52: Flow Sampling Rules with Priorities

When creating Flow Sampling rules on the Maps page, the first rule created has the highest priority and the priority of subsequent rules is in the order that they are added. To change the priority of a Flow Sampling rule in a new map, do the following:

1. **Save** the rule.
2. Select the map and click **Edit**.
3. Enter a priority in the **Priority** field of each rule to order the rules in the map. (For details about map priority, refer to *Map Priority* on page 490)

NOTE: A flow sampling map can contain only flowsampling rules. A flow sampling map cannot contain other GigaSMART rules (gsrule) or flow filtering rules (flowrule).

For configuration examples for flow sampling, refer to *GTP Whitelisting and GTP Flow Sampling Examples* on page 924.

Add Rule to Flow Sampling Map

Flow sampling is applied for new subscribers. When a new rule is added to the rules in a flow sampling map, traffic will be sent to the port or load balancing group specified in the map.

Delete Rule from Flow Sampling Map

When a rule is deleted from a flow sampling map, the session associated with the rule stays active. The traffic associated with the rule will not be reevaluated by subsequent maps.

Change Priority of Flow Sampling Maps

Use the **Priority** field in the GTP map rule to set the priority of flow sampling maps.

Delete Flow Sampling Map

When a flow sampling map is deleted, the priority of the remaining flow sampling maps will be re-prioritized. For example, if the first flow sampling map is deleted, the second flow sampling map will increase in priority.

For the deleted flow sampling map, the traffic associated with the rules in the map will be reevaluated and then passed to subsequent maps.

When a flow sampling map is re-prioritized, the existing sessions will be reevaluated according to the new priority of the map. The traffic associated with the rules in the map will be reevaluated and then passed to subsequent maps.

When the last flow sampling map is deleted, the traffic associated with the rules in the map will also be reevaluated before being passed to subsequent maps. But the traffic associated with the rules in maps that were not matched, will not be reevaluated because that traffic was already passed to subsequent maps.

Flow-Ops Report Limitation for Multiple Flow Sampling Maps

The flow-ops report displays the flow sampling rule ID for sessions that have been accepted or rejected by the flow sampling map.

However, since rule IDs are not unique across maps, when there are multiple flow sampling maps, the flow-ops report is unable to identify the exact rule that the session matched. For example, with multiple flow sampling maps, each map can have a rule ID of 1. The rule ID will be identified in the flow-ops report, but not the map associated with it.

GTP Flow Sampling Percentage

The sampling Percentage field in a map for GTP flow sampling, represents the percentage of subscribers that will be sampled (not the sessions). For example, Figure 30-53, shows a GTP flow sampling rule with the percentage set to 80.

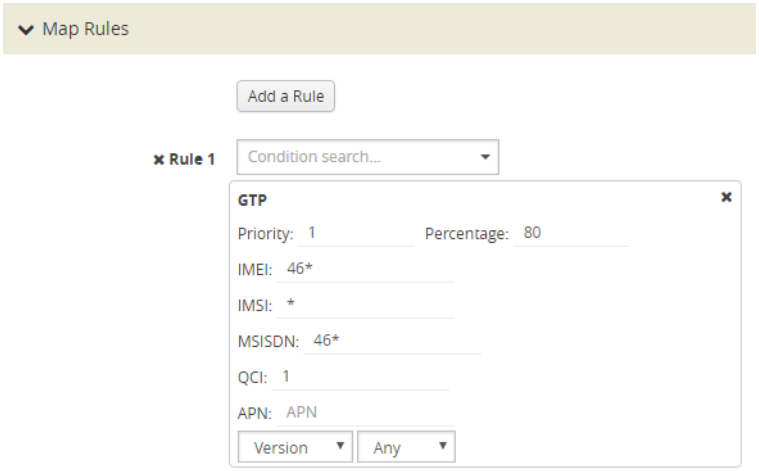

Figure 30-53: GTP Flow Sampling Percentage

The GTP correlation engine tracks all of the subscribers and all of their sessions that it sees on the network. In this example, for those subscribers with an IMSI starting with the value 46*, the GTP correlation engine keeps a list of them and randomly selects 80% of those subscribers and sets them to be in the sample, which means that a tool port (or load balanced group) will see 100% of the packets for 100% of the sessions for those randomly selected 80% of subscribers.

For the other 20% of subscribers, the GTP correlation engine continuously tracks those subscribers through the network, but does not send any packets to the tool port (or load balanced group).

Refer to the GTP flow sampling configuration examples in *GTP Whitelisting and GTP Flow Sampling Examples* on page 924.

Unmatched Traffic

When a session matches one of the configured flow sampling rules, it is either accepted for sampling or rejected.

If it is accepted, all packets belonging to that GTP session are sent to the tool port or ports specified in the flow sampling maps. If a subscriber is in the sample, then both the control plane packets and the user-data plane packets are sent to the tools.

If it is rejected, all packets belonging to the session are dropped. If the subscriber is not in the sample, then neither the control plane packets nor the user-data plane packets are sent to the tools.

Control plane (GTP-c) and user-data plane (GTP-u) traffic are treated the same. For a matching session, all the control plane and user-data plane traffic will be accepted. Otherwise, all the control plane and user-data plane traffic will be rejected and

dropped. Instead, to enable or disable GTP control plane traffic sampling, refer to *Enable or Disable GTP Control Plane Traffic Sampling* on page 924.

Enable or Disable GTP Control Plane Traffic Sampling

GTP control plane (GTP-c) traffic is typically a small percentage of total GTP traffic, but it contains useful information for analytics. Therefore, it is not always expedient to drop control plane traffic for sampled sessions.

Subscriber traffic by IMSI can be sampled such that network traffic for a subset of mobile subscribers can be selected to be sent to network monitoring tools. In some cases, network monitoring tools will want to see GTP control plane and GTP user plane traffic for a percentage of the subscribers. In other cases, network monitoring tools will want to see all of the GTP control plane traffic, but see only the GTP user plane traffic for the sampled percentage of subscribers.

Starting in software version 4.5, all control plane traffic for all subscribers will be sent to tools if GTP control plane traffic sampling is disabled. When disabled, 100% of the control traffic that matches any of the flow sampling rules will be sent to the tool ports specified in the flow sampling maps. Control traffic for both accepted and rejected sessions will be sent to the tool ports.

When GTP control plane traffic sampling is enabled, GTP-c packets will be sampled and only the indicated percentage of the control traffic that matches any of the flow sampling rules will be sent to the tool ports specified in the flow sampling maps, as described in *GTP Flow Sampling Percentage* on page 923.

The default is enable.

To disable sampling of GTP-c traffic, which enables 100% of control plane traffic, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**. Under GigaSMART Parameters, go to **GTP Sampling** and make sure that **GTP Control Sampling** is not selected.

To enable sampling of GTP-c traffic, which enables 100% of control plane traffic, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**. Under GigaSMART Parameters, go to **GTP Sampling** and make sure that **GTP Control Sampling** is selected as shown in Figure 30-54. This setting applies to all the flow sampling maps for a GigaSMART group.

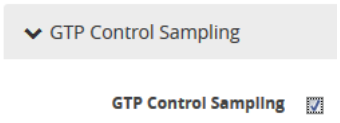

Figure 30-54: GTP Control Sampling Enabled

GTP Whitelisting and GTP Flow Sampling Examples

Refer to the following examples:

- *Example 1: GTP Whitelisting* on page 925
- *Example 2: GTP Whitelisting with Multiple Maps* on page 927

Example 1: GTP Whitelisting

Example 1 is a GTP whitelisting configuration example. Traffic from network ports go to the three first level maps (GTP-Control, GTP-User, and Fragments-Not_First) and then to the virtual port (vport1). If there is a match to an IMSI in the whitelist (MyIMSIs), it is forwarded to a port.

Task	Description	UI Steps
1.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**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.
2.	Create a virtual port.	**a.** From the device view, select GigaSMART > Virtual Ports. **b.** Click New. **c.** Type an alias in the Alias field and enter an engine port in the Port List field. **d.** From the GigaSMART Groups drop-down list, select the GigaSMART group created in task 1. **e.** Click Save.
3.	Create the GTP whitelist.	**a.** From the device view, select GigaSMART > GTP Whitelist. **b.** Click New. **c.** Type an alias in the Alias field. **d.** From the GigaSMART Groups drop-down list, select the GigaSMART group created in Task 1. **e.** Go to Task 4.
4.	Fetch whitelist files from a specified location to populate the GTP whitelist.	**a.** On the GTP Whitelist page, select Bulk Upload. **b.** Select Bulk Entry Operation for IMSI Upload Type **c.** Select Upload from URL from the Bulk Upload Type list. **d.** Enter the URL in the Enter Remote URL field. For example, http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx **e.** Click Save.

Task	Description	UI Steps
5.	Associate the GigaSMART group to the GTP whitelist.	**a.** From the device view, select GigaSMART > GigaSMART Groups **b.** Select the GigaSMART Group created in Task 1 and click Edit. **c.** Type an alias in the Alias field. **d.** From the GigaSMART Groups drop-down list, select the GigaSMART group created in task 1. **e.** Under GTP Whitelist, click on the GTP Whitelist Alias field and select the alias from Task 3. **f.** Click Save.
6.	Configure the GigaSMART operation for GTP whitelisting.	**a.** From the device view, select GigaSMART > GigaSMART Operations > GigaSMART Operation. **b.** Click New. **c.** Type an alias in the Alias field. For example, GTP-Whitelist. **d.** Select the GigaSMART group created in task 1. **e.** From the GigaSMART Operations (GSOP) drop-down list, select the following: • GTP Whitelist and select Enabled. • Load Balancing. **f.** For Load Balancing, do the following: • Choose Stateful • For Type select GTP • Choose Hashing for the metric and select IMSI **g.** Click Save.
7.	Configure three first level maps. **NOTE:** In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**a.** Configure the first map as follows: • Alias: GTP-Control • Type and subtype: First Level By Rule • Source: network port or ports • Destination: virtual port created in Task 2. • Rule: Pass, Bi Directional, Port Destination 2123 • Map Permissions: Select current user's group for Owner • Save the map **b.** Configure the second map as follows: • Alias: GTP-User • Type and subtype: First Level By Rule • Source: Same network port or ports as first map. • Destination: virtual port created in Task 2. • Rule: Pass, Bi Directional, Port Destination 2152 • Map Permissions: Select current user's group for Owner • Save the map **c.** Configure the third map as follows: • Alias: Fragments-Not-First • Type and subtype: First Level By Rule • Source: Same network port or ports as first map • Destination: virtual port created in Task 2 • Rule: Pass, IPv4 Fragmentation and select allFragNoFirst • Map Permissions: Select current user's group for Owner • Save the map

Task	Description	UI Steps
8.	Configure a second level map for GTP whitelisting, the whitelist map. If there is a match to an IMSI in the whitelist (MyIMSIs), it is forwarded to a port.	**1.** Configure the second level map as follows: • Alias: GTP-Whitelist • Type and subtype: Second Level By Rule • Source: virtual port created in Task 2 • Destination: select a tool port • GSOP: GigaSMART Operation created in Task 6 • Map Permissions: Select current user's group for Owner **2.** Click Save.

Example 2: GTP Whitelisting with Multiple Maps

Example 2 is a GTP whitelisting configuration example that includes multiple GTP whitelisting maps, which provide a more granular selection of tool ports.

Traffic from network ports go to the three first level maps (GTP-Control, GTP-User, and Fragments-Not-First) and then to the virtual port (vport1). Two whitelist maps are configured. The first map specifies a rule for version 1 traffic. The second map specifies a rule for version 2 traffic.

Task	Description	UI Steps
1.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type an gsg1 in the Alias field and enter an engine port in the Port List field, for example 10/7/e1. **d.** Click Save.
2.	Create a virtual port.	**a.** From the device view, select GigaSMART > Virtual Ports. **b.** Click New. **c.** Type vport1 in the Alias field. **d.** From the GigaSMART Groups drop-down list, select the GigaSMART group created in task 1. **e.** Click Save.
3.	Create the GTP whitelist.	**a.** From the device view, select GigaSMART > GTP Whitelist. **b.** Click New. **c.** Type an MyIMSIs in the Alias field. **d.** From the GigaSMART Groups drop-down list, select the GigaSMART group created in Task 1. **e.** Go to Task 4.
4.	Fetch whitelist files from a specified location to populate the GTP whitelist.	**a.** On the GTP Whitelist page, select Bulk Upload. **b.** Select Bulk Entry Operation for IMSI Upload Type **c.** Select Upload from URL from the Bulk Upload Type list. **d.** Enter the URL in the Enter Remote URL field. For example, http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx **e.** Click Save.

Task	Description	UI Steps
5.	Associate the GigaSMART group to the GTP whitelist.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Select the GigaSMART Group created in Task 1 and click Edit. **c.** From the GigaSMART Groups drop-down list, select the GigaSMART group created in task 1. **d.** Under GTP Whitelist, click on the GTP Whitelist Alias field and select the alias from Task 3. **e.** Click Save.
6.	Configure the GigaSMART operation for GTP whitelisting.	**a.** From the device view, select GigaSMART > GigaSMART Operations > GigaSMART Operation. **b.** Click New. **c.** Type gtp-whitelist in the Alias field. **d.** Select the GigaSMART group created in task 1. **e.** From the GigaSMART Operations (GSOP) drop-down list, select the following: • GTP Whitelist and select Enabled. • Load Balancing. **f.** For Load Balancing, do the following: • Choose Stateful • For Type select GTP • Choose Hashing for the metric and select IMSI **g.** Click Save.

Task	Description	UI Steps
7.	Configure three first level maps. **NOTE:** In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	Configure the first map. **a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Configure the map: • Alias: GTP-Control • Type: First Level, Sub Type: By Rule • Source: 8/1/x40, 8/1/x6 • Destination: vport1 **d.** Click Add a Rule. • Select Pass and Bi Directional • Select Port Destination for the rule • Set port value to 2123 **e.** Click Save. Configure the second map. **a.** Click New. **b.** Configure the map: • Alias: GTP-User • Type: First Level, Sub Type: By Rule • Source: 8/1/x40, 8/1/x6 • Destination: vport1 **c.** Click Add a Rule. • Select Pass and Bi Directional • Select Port Destination for the rule • Set port value to 2152 **d.** Click Save. Configure the second map. **a.** Click New. **b.** Configure the map: • Alias: Fragment-Not-First • Type: First Level, Sub Type: By Rule • Source: 8/1/x40, 8/1/x6 • Destination: vport1 **c.** Click Add a Rule. • Select Pass • Select Port IPv4 Fragmentation for the rule • Select allFragNoFirst for Value **d.** Click Save.

Task	Description	UI Steps
8.	Configure one second level map for GTP whitelisting, the first whitelist map. If there is a match to version 1 and if the IMSI is present in the whitelist (MyIMSIs), it is forwarded to the specified port.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Configure the map: • Alias: GTP-Whitelist_v1 • Type: Second Level, Sub Type: Flow Whitelist • Source: vport1 • Destination: 1/2x23 • Select gtp-whitelist from the GSOP list. **d.** Click Add a Rule. • Select GTP • Set Version to V1 **e.** Click Save.
9.	Configure another second level map for GTP whitelisting, the second whitelist map. If there is a match to version 2 and if the IMSI is present in the whitelist (MyIMSIs), it is forwarded to the specified port.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Configure the map: • Alias: GTP-Whitelist_v2 • Type: Second Level, Sub Type: Flow Whitelist • Source: vport1 • Destination: 1/2x24 • Select gtp-whitelist from the GSOP list. **d.** Click Add a Rule. • Select GTP • Set Version to V2 **e.** Click Save.

Example 3: GTP Flow Sampling

Example 2 is a GTP flow sampling configuration example. Traffic from network ports go to the three first level maps (GTP-Control, GTP-User, and Fragments-Not_First) and then to the virtual port (vport1). The traffic flow is sampled based on the rules in one flow sampling map (GTP-Sample-01). The flow sampling rules specify IMSI, IMEI, and MSISDN numbers, as well as the percentage to sample. Packets are then accepted or rejected. Accepted packets are forwarded to a port. Rejected packets are dropped. Packets that do not match a rule will be passed to subsequent maps.

Task	Description	UI Steps
1.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type an alias in the Alias field and enter an engine port in the Port List field. **d.** Click Save.

Task	Description	UI Steps
2.	Create a virtual port.	**a.** From the device view, select GigaSMART > Virtual Ports. **b.** In the Alias field, type an alias for this virtual port. **c.** Type an alias in the Alias field and enter an engine port in the Port List field. **d.** From the GigaSMART Groups drop-down list, select the GigaSMART group created in task 1. **e.** Click Save.
3.	Configure three first level maps. **NOTE:** In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**a.** Configure the first map as follows: • Alias: GTP-Control • Type and subtype: First Level By Rule • Source: network port or ports • Destination: virtual port created in Task 2. • Rule: Pass, Bi Directional, Port Destination 2123 • Map Permissions: Select current user's group for Owner • Save the map **b.** Configure the second map as follows: • Alias: GTP-User • Type and subtype: First Level By Rule • Source: Same network port or ports as first map. • Destination: virtual port created in Task 2. • Rule: Pass, Bi Directional, Port Destination 2152 • Map Permissions: Select current user's group for Owner • Save the map **c.** Configure the third map as follows: • Alias: Fragments-Not-First • Type and subtype: First Level By Rule • Source: Same network port or ports as first map • Destination: virtual port created in Task 2 • Rule: Pass, IPv4 Fragmentation and select allFragNoFirst • Map Permissions: Select current user's group for Owner • Save the map
4.	Configure the GigaSMART operation for GTP flow sampling.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations. **b.** Click New. **c.** Type an alias in the Alias field. For example, GTP-Whitelist. **d.** Select the GigaSMART group created in task 1. **e.** From the GigaSMART Operations (GSOP) drop-down list, select the following: • GTP Whitelist and select Enabled. • Load Balancing. **f.** For Load Balancing, do the following: **a.** Choose Stateful **b.** For Type select GTP **c.** Choose Hashing for the metric and select IMSI **d.** Click Save.

Task	Description	UI Steps
5.	Configure a second level map for GTP flow sampling, the flow sampling map. The traffic flow is sampled based on the rules in this map.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Configure the map. • Type GTP-Sample-01 in the Alias field • Select Second Level for Type • Select Flow Sample for Subtype. • Select the virtual port configured in Task 2 for the Source • Select a tool port for the Destination • Select the GigaSMART Operation configured in Task for from the GSOP list **d.** Use the Add a Rule button to create the following flow sampling rules: • Percentage to 50, IMEI 01416800* • Percentage to 80, IMSI 46* • Percentage to 25, MSISDN 1509* • Percentage to 15, IMSI 01400* • Percentage to 20, IMSI, 31*, MSISDN 1909* **e.** Click Save.

Example 4: GTP Whitelisting, GTP Flow Sampling, and Load Balancing

Example 4 combines the GTP whitelisting configuration from Example 1 with the GTP flow sampling configuration from Example 3, and adds GigaSMART load balancing.

In Example 4, traffic from network ports go to the three first level maps (GTP-Control, GTP-User, and Fragments-Not-First) and then to the virtual port (vport1). If there is a match to an IMSI in the whitelist (MyIMSIs), it is forwarded to the port group (PG-Whitelist) for load balancing.

NOTE: In Example 4, the tool ports in the port group are on the same node as the GigaSMART group and GigaSMART operation.

If there is not a match to an IMSI in the whitelist, the traffic flow is sampled based on the rules in the flow sampling map (GTP-Sample-01). The flow sampling rules specify IMSI, IMEI, and MSISDN numbers, as well as the percentage to sample. Packets are then accepted or rejected. Accepted packets are forwarded to the port group (PG-Sample) for load balancing. Rejected packets are dropped. Packets that do not match a rule will be passed to subsequent maps.

Task	Description	UI Steps
1.	Create port groups and specify the tool ports and enable load balancing.	**a.** Select Ports > Port Groups > All Port Groups. **b.** Click New. **c.** Type PG-Whitelist in the Alias field. **d.** Select SMART Load Balancing **e.** Click in the Ports field and select the tool ports for the port group. **f.** Click Save. **g.** Repeat steps 2 through 6, to create a port group with the alias PF-Sample.
2.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type an alias in the Alias field and enter an engine port in the Port List field. **d.** Click Save.
3.	Create a virtual port.	**a.** From the device view, select GigaSMART > Virtual Ports. **b.** In the Alias field, type an alias for this virtual port. **c.** Type an alias in the Alias field and enter an engine port in the Port List field. **d.** From the GigaSMART Groups drop-down list, select the GigaSMART group created in task 1. **e.** Click Save.
4.	Configure three first level maps. **NOTE:** In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**a.** Configure the first map as follows: • Alias: GTP-Control • Type and subtype: First Level By Rule • Source: network port or ports • Destination: virtual port created in Task 2. • Rule: Pass, Bi Directional, Port Destination 2123 • Map Permissions: Select current user's group for Owner • Save the map **b.** Configure the second map as follows: • Alias: GTP-User • Type and subtype: First Level By Rule • Source: Same network port or ports as first map. • Destination: virtual port created in Task 2. • Rule: Pass, Bi Directional, Port Destination 2152 • Map Permissions: Select current user's group for Owner • Save the map **c.** Configure the third map as follows: • Alias: Fragments-Not-First • Type and subtype: First Level By Rule • Source: Same network port or ports as first map • Destination: virtual port created in Task 2 • Rule: Pass, IPv4 Fragmentation and select allFragNoFirst • Map Permissions: Select current user's group for Owner • Save the map

Task	Description	UI Steps
5.	Create the GTP whitelist.	**a.** From the device view, select GigaSMART > GTP Whitelist. **b.** Click New. **c.** Type an Alias for the Whitelist in the Alias field. For example, MyIMSIs
6.	Fetch whitelist files from a specified location to populate the GTP whitelist.	**a.** On the GTP Whitelist page, select Bulk Upload. **b.** Select Bulk Entry Operation for IMSI Upload Type **c.** Select Upload from URL from the Bulk Upload Type list. **d.** Enter the URL in the Enter Remote URL field. For example, http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx **e.** Click Save.
7.	(Optional) Add a single IMSI to the GTP whitelist.	**a.** On the GTP Whitelist page, select Individual Entry Operation. **b.** Select Append for Operation Type **c.** Enter the IMSI entry in the Individual IMSI Entry field. **d.** Click Save.
8.	Associate the GigaSMART group to the GTP whitelist.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type an alias in the Alias field. **d.** Under GTP Whitelist, click on the GTP Whitelist Alias field and select the alias from Task 5. **e.** Click Save.
9.	Configure the GigaSMART operation for GTP whitelisting.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP)> GigaSMART Operation. **b.** Click New. **c.** Select the GigaSMART Group created in Task 8 from the GigaSMART Groups list. **d.** Type an alias in the Alias field. For example, gtp-whitelist. **e.** Select GTP Whitelist from the GigaSMART Operations (GSOP) list. **f.** Select Load Balancing from the GigaSMART Operations (GSOP) list. **g.** Configure Load Balancing as follows: • Select Stateful • Set Type to GTP • Select Hashing • Select IMSI **h.** Click Save.

Task	Description	UI Steps
10.	Configure a second level map for GTP whitelisting, the whitelist map. If there is a match to an IMSI in the whitelist (MyIMSIs), it is forwarded to a load balancing port group.	**a.** Select Maps > Maps > Maps. **b.** Click New **c.** Configure the map. • Type an name in the Alias field. For example GTP-Whitelist. • Select Second Level for Type • Select By Rule for Subtype • Select the GigaSMART Operation configured in Task 9 from the GigaSMART Operations (GSOP) list. • Select the virtual port configured in Task 3 for Source • Select PG-Whitelist for Destination **d.** Click Save.
11.	Configure the GigaSMART operation for GTP flow sampling.	**e.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **f.** Click New. **g.** Select the GigaSMART Group created in Task 8 from the GigaSMART Groups list. **h.** Type an alias in the Alias field. For example, gtp-flowsample. **i.** Select Flow Sampling from the GigaSMART Operations (GSOP) list. **j.** Select Flow Sampling-GTP. **k.** Select Load Balancing from the GigaSMART Operations (GSOP) list. • Select Stateful • Set Type to GTP • Select Hashing • Select IMSI **l.** Click Save.
12.	Configure a second level map for GTP flow sampling, the flow sampling map. If there is not a match to an IMSI in the whitelist, the traffic flow is sampled based on the rules in this map. Accepted packets are forwarded to a load balancing port group.	**a.** Select Maps > Maps > Maps. **b.** Click New **c.** Configure the map. • Type an name in the Alias field. For example GTP-Sample-01. • Select Second Level for Type • Select Flow Sample for Subtype • Select the GigaSMART operation for flow sampling configured in Task 11 from the GSOP list. • Select the virtual port configured in Task 3 for Source • Select PG-Sample for Destination **d.** Create the following flow sample rules: • Percentage 50, IMEI 01416800*, IMSI 31* • Percentage 80, IMSI 46* • Percentage 25, MSISDN 1509* • Percentage 15, IMEI 01400*, imsi 31* • Percentage 20, IMSI 31*, MSISDN 1909* **e.** Click Save.

Example 5: GTP Flow Sampling with Multiple Maps

Example 5 includes multiple GTP flow sampling maps, which provide a more granular selection of tool ports for flow sampling.

In Example 5, traffic from network ports go to the three first level maps (GTP-Control, GTP-User, and Fragments-Not_First) and then to the virtual port (vport1). If there is a match to an IMSI in the whitelist (VoLTE_1MM), it is forwarded to the port-group (PG-Whitelist-1) for load balancing.

NOTE: In Example 5, the tool ports in the port group are on the same node as the GigaSMART group and GigaSMART operation.

If there is not a match to an IMSI in the whitelist, the traffic flow is sampled based on the rules in the four flow sampling maps (GTP-Sample-1 to GTP-Sample-4).

The flow sampling rules in each map specify IMSI, IMEI, and MSISDN numbers, as well as the percentage to sample. Packets are then accepted or rejected. Accepted packets are forwarded to the port-group (PG-Sample-1 to PG-Sample-4) for load balancing. Rejected packets are dropped. Packets that do not match a rule will be passed to subsequent maps, in this example, to a shared collector.

Task	Description	UI Steps
1.	Create port groups, specifying the tool ports and enabling load balancing.	**a.** Select Ports > Port Groups > All Port Groups. **b.** Click New. **c.** Type PG-Sample-1 in the Alias field. **d.** Select SMART Load Balancing **e.** Click in the Ports field and select the tool ports for the port group. **f.** Click Save. **g.** Repeat steps 2 through 6, to create a port groups with the aliases
2.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type an alias in the Alias field and enter an engine port in the Port List field. **d.** Click Save.
3.	Create a virtual port.	**a.** From the device view, select GigaSMART > Virtual Ports. **b.** Type vport1 in the Alias field. **c.** Select the GigaSMART Groups created in Task 2 from the GigaSMART Group list. **d.** Click Save.

Task	Description	UI Steps
4.	Configure three first level maps. **NOTE:** In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**a.** Configure the first map as follows: • Alias: GTP-Control • Type and subtype: First Level By Rule • Source: network ports (for example, 10/1/x5, 10/3/x1,10/6/q1) • Destination: virtual port created in Task 2. • Rule: Pass, Bi Directional, Port Destination 2123 • Save the map **b.** Configure the second map as follows: • Alias: GTP-User • Type and subtype: First Level By Rule • Source: Same network ports as first map. • Destination: virtual port created in Task 2. • Rule: Pass, Bi Directional, Port Destination 2152 • Save the map **c.** Configure the third map as follows: • Alias: Fragments-Not-First • Type and subtype: First Level By Rule • Source: Same network ports as first map • Destination: virtual port created in Task 2 • Rule: Pass, IPv4 Fragmentation and select allFragNoFirst • Save the map
5.	Create the GTP whitelist.	**a.** From the device view, select GigaSMART > GTP Whitelist. **b.** Click New. **c.** Enter VoLTE_1MM in the Alias field. **d.** Go to Task 6.
6.	Fetch whitelist files from a specified location to populate the GTP whitelist.	**a.** On the GTP Whitelist page, select Bulk Upload. **b.** Select Bulk Entry Operation for IMSI Upload Type **c.** Select Upload from URL from the Bulk Upload Type list. **d.** Enter the URL in the Enter Remote URL field. For example, http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx **e.** Click Save.
7.	(Optional) Add a single IMSI to the GTP whitelist.	**a.** On the GTP Whitelist page, select Individual Entry Operation. **b.** Select Append for Operation Type **c.** Enter the IMSI entry in the Individual IMSI Entry field. **d.** Click Save.
8.	Associate the GigaSMART group to the GTP whitelist.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups **b.** Click New. **c.** Type gsg1 in the Alias field. **d.** Under GTP Whitelist, click on the GTP Whitelist Alias field and select VolTE_1MM. **e.** Click Save.

Task	Description	UI Steps
9.	Configure the GigaSMART operation for GTP whitelisting.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Select the GigaSMART Group gsg1 created in Task 8 from the GigaSMART Groups list. **d.** Enter gtp-whitelistl in the Alias field. **e.** Select GTP Whitelist from the GigaSMART Operations (GSOP) list **f.** Select Load Balancing from the GigaSMART Operations (GSOP) list. **g.** Configure Load Balancing as follows: • Select Stateful • Set Type to GTP • Select Hashing • Select IMSI **h.** Click Save.
10.	Configure a second level map for GTP whitelisting, the whitelist map. If there is a match to an IMSI in the whitelist (VoLTE_1MM), it is forwarded to a load balancing port group.	**a.** Select Maps > Maps > Maps. **b.** Click New **c.** Configure the map. • Enter GTP-Whitelist in the Alias field. • Select Second Level for Type • Select By Rule for Subtype • Select gtp-whitelist from the GSOP list. • Select the virtual port vport1 configured in Task 3 for Source • Select port group PG-Whitelist-2 for Destination **d.** Click Save.
11.	Configure the GigaSMART operation for GTP flow sampling.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Select the GigaSMART Group created in Task 8 from the GigaSMART Groups list. **d.** Enter gtp-flowsample-1 in the Alias field. **e.** Select Flow Sampling from the GigaSMART Operations (GSOP) list and then select the Flow Sampling-GTP option. **f.** Select Load Balancing from the GigaSMART Operations (GSOP) list. **g.** Configure Load Balancing as follows: • Select Stateful • Set Type to GTP • Select Hashing • Select IMSI **h.** Click Save.

Task	Description	UI Steps
12.	Configure a second level map for GTP flow sampling, the first flow sampling map. This map has 12 rules. Traffic flow is sampled based on the rules in this map. Accepted packets are forwarded to a load balancing port group.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Configure the map. • Type GTP-Sample-1 in the Alias field • Select Second Level for Type • Select Flow Sample for Subtype. • Select the virtual port vport1 configured in Task 3 for the Source • Select a port grroup PG-Sampl-1 for the Destination • Select the GigaSMART Operation configured in Task for from the GSOP list **d.** Use the Add a Rule button to create the following flow sampling rules: • Percentage 75, IMSI 3182609833*, IMEI 35609506* • Percentage 10, IMSI 3182609834*, IMEI 3560950* • Percentage 20, IMSI 31826098350*, IMEI 356095* • Percentage 20, IMSI 31826098351*, IMEI 35609* • Percentage 20, IMSI 31826098352*, IMEI 3560* • Percentage 20, IMSI 31826098353*, IMEI 356* • Percentage 20, IMSI 31826098354*, IMEI 35* • Percentage 20, IMSI 31826098355*, IMEI 31* • Percentage 20, IMSI 31826098356*, IMEI 356095* • Percentage 20, IMSI 31826098356*, IMEI 356095* • Percentage 20, IMSI 31826098357*, IMEI 3560* • Percentage 20, IMSI 31826098358*, IMEI 35* • Percentage 20, IMSI 31826098359*, IMEI 356095* **e.** Click Save.

Task	Description	UI Steps
13.	Configure a second level map for GTP flow sampling, the second flow sampling map. This map has 12 rules. Traffic flow is sampled based on the rules in this map. Accepted packets are forwarded to a load balancing port group.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Configure the map. • Type GTP-Sample-2 in the Alias field • Select Second Level for Type • Select Flow Sample for Subtype. • Select the virtual port vport1 configured in Task 2 for the Source • Select a tool port group PG-Sample-2 for the Destination • Select flow-sample-1 configured in Task 11 for from the GSOP list **d.** Use the Add a Rule button to create the following flow sampling rules: • Percentage 30, IMSI 3182609836*, IMEI 35609506* • Percentage 5, IMSI 3182609837*, IMEI 356095062* • Percentage 50, IMSI 31826098380*, IMEI 356095062* • Percentage 50, IMSI 31826098381*, IMEI 35609506* • Percentage 50, IMSI 31826098382*, IMEI 3560950* • Percentage 50, IMSI 31826098383*, IMEI 356095* • Percentage 50, IMSI 31826098384*, IMEI 35* • Percentage 50, IMSI 31826098385*, IMEI 356* • Percentage 50, IMSI 31826098386*, IMEI 3560* • Percentage 50, IMSI 31826098387*, IMEI 35609* • Percentage 50, IMSI 31826098388*, IMEI 356095* • Percentage 50, IMSI 31826098389*, IMEI 3560950* **e.** Click Save.
14.	Configure a second level map for GTP flow sampling, the third flow sampling map. This map has 5 rules. Traffic flow is sampled based on the rules in this map. Accepted packets are forwarded to a load balancing port group.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Configure the map. • Type GTP-Sample-3 in the Alias field • Select Second Level for Type • Select Flow Sample for Subtype • Select the virtual port vport1 configured in Task 3 for the Source • Select a port group PG-Sample-3 port for the Destination • Select flow-sample-1 configured in Task 11 for from the GSOP list **d.** Use the Add a Rule button to create the following flow sampling rules: • Percentage 10, IMSI 31826098390*, IMEI 35609506* • Percentage 10, IMSI 31826098391*, IMEI 35609506* • Percentage 10, IMSI 31826098392*, IMEI 35609506* • Percentage 10, IMSI 31826098393*, IMEI 35609506* • Percentage 10, IMSI 31826098394*, IMEI 35609506* **e.** Click Save.

Task	Description	UI Steps
15.	Configure a second level map for GTP flow sampling, the fourth flow sampling map. This map has one rule. Traffic flow is sampled based on the rules in this map. Accepted packets are forwarded to a load balancing port group.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Configure the map. • Type GTP-Sample-4 in the Alias field • Select Second Level for Type • Select Flow Sample for Subtype • Select the virtual port vport1 configured in Task 3 for the Source • Select a tool port for the Destination • Select flow-sample-1 configured in Task 11 for from the GSOP list **d.** Use the Add a Rule button to create the following flow sampling rule: • Percentage 10, IMSI 31826098429*, IMEI 35609506* **e.** Click Save.
16.	Configure a collector map for any packets that do not match other rules.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Configure the map. • Type GTP-Collector in the Alias field • Select Second Level for Type • Select Collector for Subtype • Select the virtual port vport1 configured in Task 3 for the Source **d.** Click Save.

Display GTP Flow Ops Report Statistics

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

Example 6: APN for GTP Whitelisting, GTP Flow Sampling

Example 7 specifies APN patterns for GTP whitelisting and GTP flow sampling.

In Example 7, traffic from network ports go to the two first level maps (gtp_to_v1_c and gtp_to_v1_u) and then to the virtual port (v1).

In the whitelist map, if there is a match to the APN pattern and if the IMSI is present in the whitelist (IMSI), packets are forwarded to a tool port.

If there is not a match to an IMSI in the whitelist, the traffic is flow sampled based on the APN pattern in the flow sampling map. Accepted packets are forwarded to the same tool port as specified in the whitelist map.

Any unmatched traffic goes to a shared collector that sends it to a different tool port.

Task	Description	UI Steps
1.	Configure a network port and two tool ports and enable them.	**a.** Select Ports > Ports > All Ports. **b.** Click Quick Port Editor. **c.** Configure a network port. Port 22/3/x3 in this example. **d.** Configure two tool ports. Port 22/4/x18 and 22/4/x19 in this example. **e.** Admin enable the ports by selecting Enable for each port. **f.** Click OK.
2.	Configure a GigaSMART group and associate it with two GigaSMART engine port.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type an gsg2 in the Alias field. **d.** In the Port List field, select the engine ports, which are 22/2/e1 and 22/2/e2 in this example **e.** Click Save.
3.	Create a virtual port.	**a.** From the device view, select GigaSMART > Virtual Ports. **b.** Type v1 in the Alias field. **c.** Select the GigaSMART Group created in Task 2 from the GigaSMART Group list. **d.** Click Save.
4.	Configure two first level maps, one for control traffic and one for user traffic.	**a.** Configure the first map as follows: • Alias: gtp_to_v1_c • Type and subtype: First Level By Rule • Source: 22/3/x3 • Destination: virtual port created in Task 2. • Rule 1: Pass, Bi Directional, Port Destination 2123 • Rule 2: Pass, Bi Directional, Port Destination 2122 • Save the map **b.** Configure the second map as follows: • Alias: gtp_to_v1_u • Type and subtype: First Level By Rule • Source: 22/3/x3. • Destination: virtual port created in Task 2. • Rule 1: Pass, Bi Directional, Port Destination 2152 • Rule 1: Pass, Bi Directional, IPv4 Fragmentation, Value: allFragNoFirst. • Save the map
5.	Create the GTP whitelist.	**a.** From the device view, select GigaSMART > GTP Whitelist. **b.** Click New. **c.** Enter gtp-whitelist in the Alias field **d.** Go to Task 6.

Task	Description	UI Steps
6.	Fetch whitelist files from a specified location to populate the GTP whitelist.	**a.** On the GTP Whitelist page, select Bulk Upload. **b.** Select Bulk Entry Operation for IMSI Upload Type **c.** Select Upload from URL from the Bulk Upload Type list. **d.** Select Append. **e.** Enter the URL in the Enter Remote URL field. For example, http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx **f.** Click Save.
7.	(Optional) Add a single IMSI to the GTP whitelist.	**a.** On the GTP Whitelist page, select Individual Entry Operation. **b.** Select Append for Operation Type **c.** Enter the IMSI entry in the Individual IMSI Entry field. **d.** Click Save.
8.	Associate the GigaSMART group to the GTP whitelist.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups **b.** Select GS Group gsg2 created in Task 2 and click Edit **c.** Under GTP Whitelist, click on the GTP Whitelist Alias field and select gtp-whitelist **d.** Click Save.
9.	Configure the GigaSMART operation for GTP whitelisting.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Select the GigaSMART Group gsg2 created in Task 2 and associated with the GTP whitelist in Step 8. **d.** Enter gtp-correlat_gsp_wl in the Alias field. **e.** Select GTP Whitelist from the GigaSMART Operations (GSOP) list **f.** Select Load Balancing from the GigaSMART Operations (GSOP) list. **g.** Configure Load Balancing as follows: • Select Stateful • Set Type to GTP • Select Hashing • Select IMSI **h.** Click Save.
10.	Configure a second level map for GTP whitelisting, the whitelist map. If there is a match to the APN pattern and if IMSI is present in the whitelist (IMSI), it is forwarded to a tool port.	**a.** Select Maps > Maps > Maps. **b.** Click New **c.** Configure the map. • Enter GTP-Whitelist in the Alias field. • Select Second Level for Type • Select Flow Whitelist for Subtype • Select gtp-correlate_gsg_wl from the GSOP list. • Select the virtual port v1 configured in Task 3 for Source • Select 22/4/x18 for Destination • Rule 1: GTP, APN: mobile.com **d.** Click Save.

Task	Description	UI Steps
11.	Configure the GigaSMART operation for GTP flow sampling.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Select the GigaSMART Group gsg2 created in Task 2 and associated with the GTP whitelist in Step 8. **d.** Enter gtp-correlat_gsp_fs in the Alias field. **e.** Select GTP Whitelist from the GigaSMART Operations (GSOP) list **f.** Select Load Balancing from the GigaSMART Operations (GSOP) list. **g.** Configure Load Balancing as follows: • Select Stateful • Set Type to GTP • Select Hashing • Select IMSI **h.** Click Save
12.	Configure a second level map for GTP flow sampling, the flow sampling map. If there is not a match to an IMSI in the whitelist, the traffic flow is sampled based on the APN pattern in this map. Accepted packets are forwarded to the same tool port as specified in the whitlelist map	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Configure the map. • Type from_vp_fs1 in the Alias field • Select Second Level for Type • Select Flow Sample for Subtype. • Select the virtual port v1 configured in Task 3 for the Source • Select a 22/4/x18 for the Destination • Select the GigaSMART Operation gtp-correlate_gsg_fs • Rule 1: GTP, Percentage: 100, APN: imsi* **d.** Click Save.
13.	Add a shared collector for any unmatched traffic from the virtual port and send it to a different tool port.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Configure the map. • Type from_vp_scoll in the Alias field • Select Second Level for Type • Select Collector for Subtype • Select the virtual port v1 configured in Task 3 for the Source **d.** Click Save.

Example 7: APN for FTP Whitelisting, APN and QCI for GTP Flow Sampling

Example 6 specified APN patterns for GTP whitelisting and GTP flow sampling. It also specifies QCI for GTP flow sampling.

In Example 7, traffic from network ports go to the two first level maps (gtp_to_v1_c and gtp_to_v1_u) and then to the virtual port (v1).

In the whitelist map, if there is a match to the APN pattern and if the IMSI is present in the whitelist (IMSI), packets are forwarded to a tool port.

If there is not a match to an IMSI in the whitelist, the traffic is flow sampled based on the APN pattern and QCI value in the flow sampling map. Accepted packets are forwarded to the same tool port as specified in the whitelist map. Only 50% of traffic with QCI 5 is sent to the tool port.

Any unmatched traffic goes to a shared collector that sends it to a different tool port.

Task	Description	UI Steps
1.	Configure a network port and two tool ports and enable them.	**a.** Select Ports > Ports > All Ports. **b.** Click Quick Port Editor. **c.** Select a port (for example, 22/2/x3) and set Type to Network. **d.** Select a port (for example, 22/2/x18) and set Type to Tool **e.** Select a second port (for example, 22/2/x19) and set Type to Tool. **f.** Select Enable for Admin on the network and two tool ports. **g.** Click OK.
2.	Configure a GigaSMART group and associate it with two GigaSMART engine ports	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type gsg2 in the Alias field. **d.** Click in the Port List field and select two engine ports. For example, 22/2/e1 and 22/2/e2 **e.** Click Save.
3.	Create a virtual port.	**a.** From the device view, select GigaSMART > Virtual Ports. **b.** Type v1 in the Alias field. **c.** Click in the GigaSMART Group field and select the GigaSMART Group created in Task 2. **d.** Click Save.
4.	Configure two first level maps, on for control traffic and one for user traffic	**a.** Configure the first map as follows: • Alias: gtp_to_v1_c • Type and Subtype: First Level By Rule • Traffic Type: select Control • Source: 22/2/3/x3 (network port created in Task 1) • Destination: v1 (virtual port created in Task 3) • Rule 1: Pass, Bi Directional, Port Destination 2123 • Rule 2: Pass, Bi Directional, Port Destination 2122 • Save the map **b.** Configure the second map as follows: • Alias: gtp_to_v1_u • Type and subtype: First Level By Rule • Source: 22/2/3/x3 (network port created in Task 1) • Destination: v1 (virtual port created in Task 3) • Rule 1: Pass, Bi Directional, Port Destination 2152 • Rule 2: Pass, Bi Directional, IPv4Fragmentation allFragNoFirst • Save the map

Task	Description	UI Steps
5.	Associate the GigaSMART group to the active GTP Whitelist	**a.** From the device view, select GigaSMART > GigaSMART Groups. **b.** Select the GigaSMART Group created in Task 1 and click Edit. **c.** Locate the GTP Whitelist param, and enter the alias of whitelist in the GTP Whitelist Alias field. For example, IMSI. **d.** Save the GigaSMART Group.
6.	Configure the GigaSMART operation for GTP whitelisting.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Select the GigaSMART Group gsg1 created in Task 2 from the GigaSMART Groups list. **d.** Enter gtp-correlate_gsp_wl in the Alias field. **e.** Select GTP Whitelist from the GigaSMART Operations (GSOP) list **f.** Select Load Balancing from the GigaSMART Operations (GSOP) list. **g.** Configure Load Balancing as follows: • Select Stateful • Set Type to GTP • Select Hashing • Select IMSI **h.** Click Save.
7.	Configure a second level map for GTP whitelisting, the whitelist map. If there is a match to the APN pattern and if the IMSI is present in the whitelist (IMSI), packets are forwarded to a tool port.	**a.** Select Maps > Maps > Maps. **b.** Click New **c.** Configure the map. • Alias: GTP-whitelist • Type an Subtype: Second Level Flow Whitelist • Source: v1 (virtual port created in Task 3) • Destination: 22/4/x18 • GSOP: gtp-corelate_gsg_wl • Select gtp-whitelist from the GSOP list. • Rule: GTP, APN: mobile.com **d.** Click Save.

Task	Description	UI Steps
8.	Configure the GigaSMART operation for GTP flow sampling.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Select the GigaSMART Group created in Task 2 from the GigaSMART Groups list. **d.** Enter gtp-corelate_gsg_fs in the Alias field. **e.** Select Flow Sampling from the GigaSMART Operations (GSOP) list and then select the Flow Sampling-GTP option. **f.** Select Load Balancing from the GigaSMART Operations (GSOP) list. **g.** Configure Load Balancing as follows: • Select Stateful • Set Type to GTP • Select Hashing • Select IMEI **h.** Click Save.
9.	Configure a second level map for GTP flow sampling, the flow sampling map. If there is not a match to an IMSI in the whitelist, the traffic flow is sampled based on the APN pattern in this map. Accepted packets are forwarded to the same tool port as specified in the whitelist map.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Configure the map. • Alias: from_vp_fs1 • Type and Subtype: Second Level Flow Sample • Source: vp1 • Destination: 22/4/x18 • GSOP: gtp-corelate_gsg_fs • Rule 1: GTP, APN: *imsi*, QCI: 5, Percentage: 50 • Rule: 2: GTP, IMSI: ims*, Percentage 100 **d.** Click Save.
10.	Add a shared collector for any unmatched traffic from the virtual port and send it to a different tool port.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Configure the map: • Alias: from_vp_scoll • Type and Subtype: Regular Collector • Source: v1 • Destination: 22/4/x19 **d.** Click Save.

Display Flow Ops Reports

GigaSMART provides support for Flow Ops reporting. The Flow Ops reports displays session table statistics for various feature.

Refer to *GigaSMART Group Statistics Definitions* on page 786 for descriptions of these statistics.

Description	UI Steps
To display Flow Filtering Reports:	1. From the device view, select **GigaSMART > GigaSMART Groups > Report**. 2. Under Report Info a. Select *Flow Sampling* as the **Report Type** b. Select a GigaSMART Group - for example: **grp1** c. Enter a specific **Device IP address/Mask** or select *Any* to search for any available device IP address. 3. Click **Generate**.

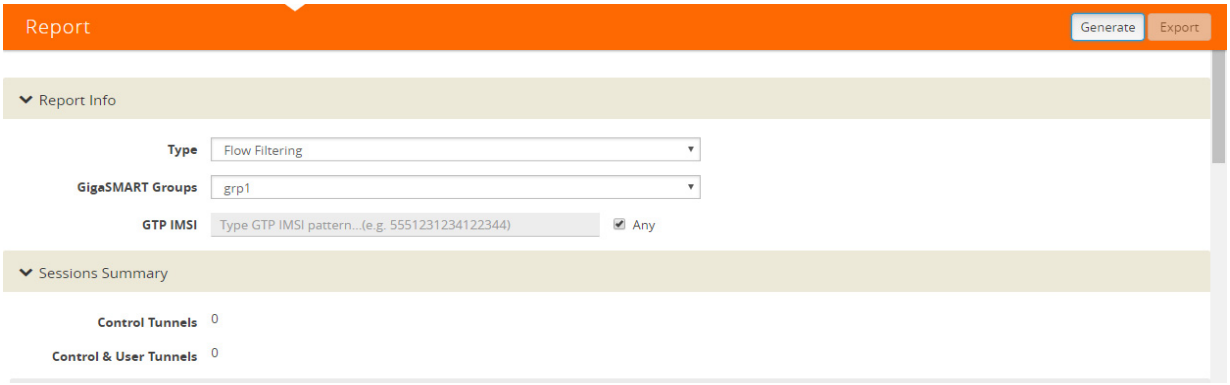

To display Flow Sampling Reports:	• Under Report Info • Select a Report Type: **Flow Sampling**

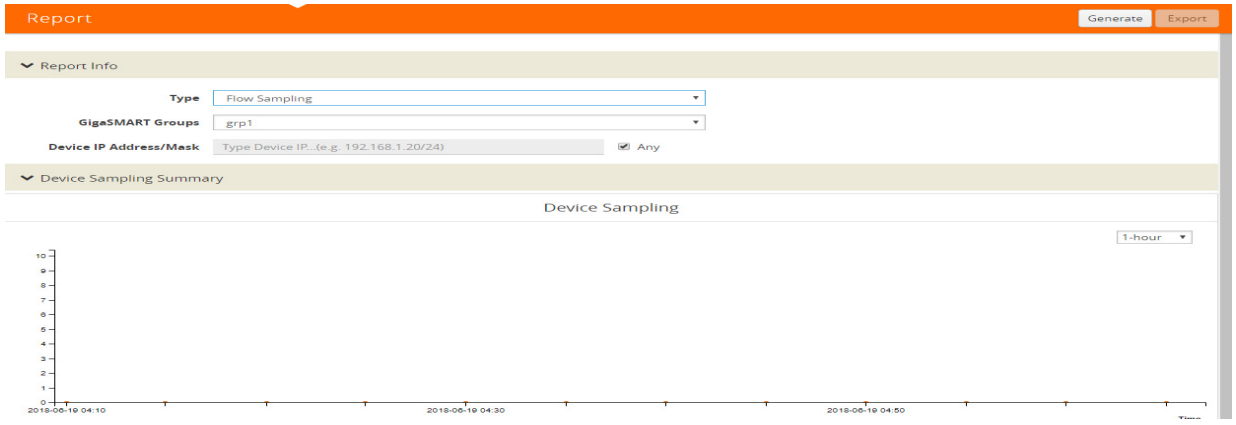

To display Flow SIP Reports:	• Under Report Info • Select a Report Type: **Flow SIP**

Description	UI Steps

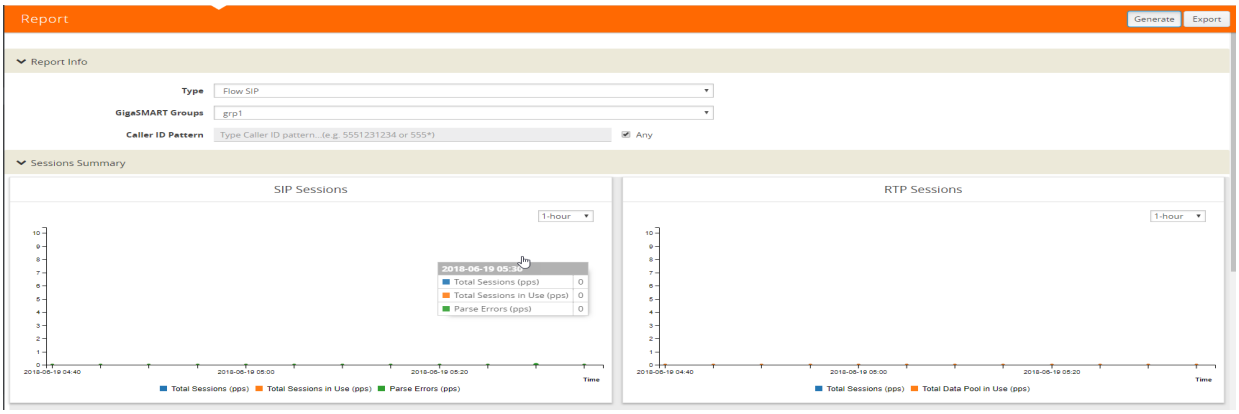

Export Flow SIP Session Reports

Starting with software version 5.4 support for exporting flow-ops session reports are available for Flow SIP.

To export Flow SIP session reports:

1. Navigate to **From the device view, select GigaSMART > GigaSMART Groups > Report**.
2. Generate a Flow SIP report.
3. Click **Export** to download the Flow SIP report you just generated. A text file of the Flow SIP report is saved to your local directory.

   NOTE: The session table displays the first 1000 sessions only.

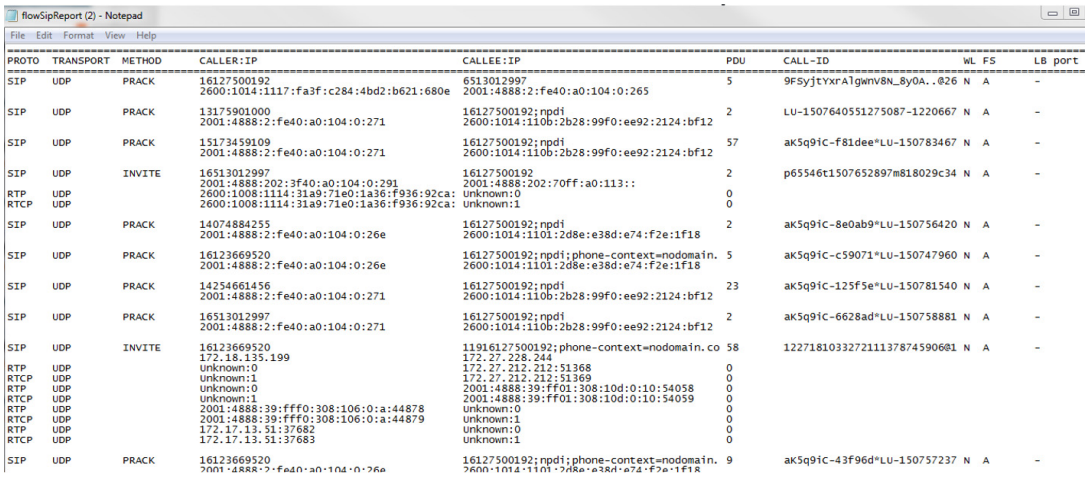

GTP Overlap Flow Sampling Maps

Starting in software version 4.8, GTP overlap flow sampling maps combines GTP whitelisting and GTP flow sampling maps into a GTP overlap flow sampling map group, which allows for selected traffic to be sent to multiple destinations simultaneously.

In this scenario, once traffic matches a map, it will be sent to the destination for that map. However, the matched traffic will also be evaluated by subsequent maps and, if a match occurs, it will be sent to each of the destinations pointed to by the subsequent maps.

Figure 30-56 on page 951 illustrates regular non-overlap mapping where, once a traffic match is achieved in one map, all other maps are ignored.

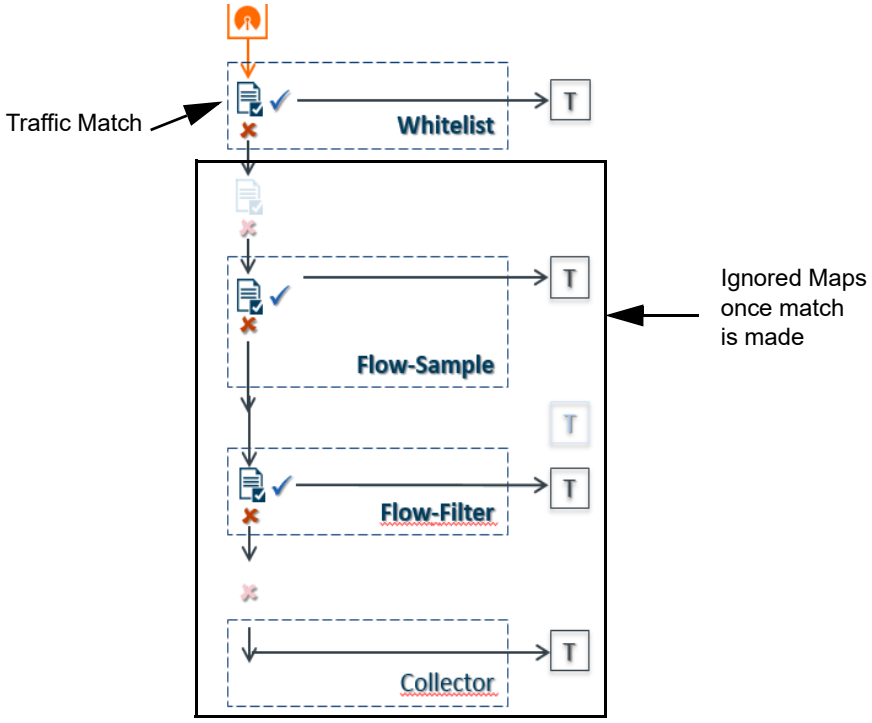

Figure 30-55: Non-Overlap GTP Mapping Mode

This contrasts with GTP overlap flow sampling maps. In Figure 30-55 on page 950 matched traffic is sent to up to six GTP whitelisting and flow sampling map pairs that in turn send accepted traffic to up to six load balanced port groups.

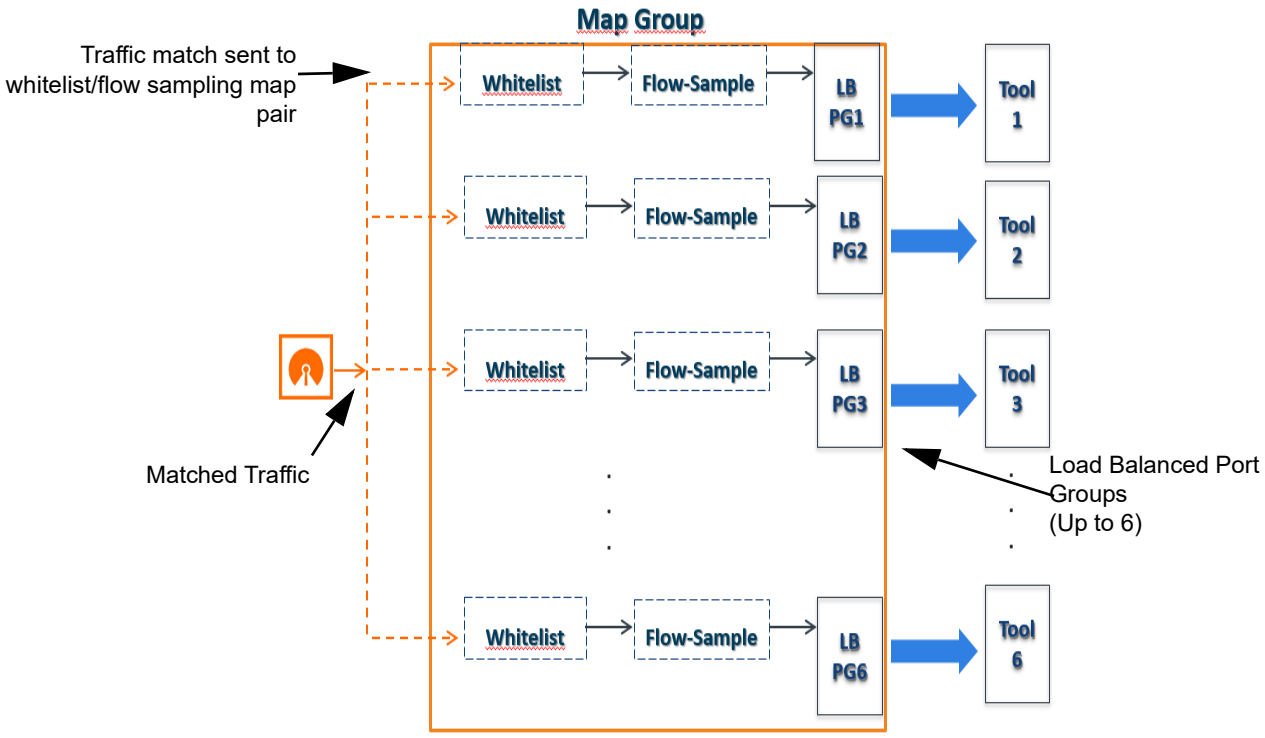

Figure 30-56: GTP Overlap Flow Sampling Map Groups

Configure GTP Overlap Mapping

The configuration of GTP whitelisting and GTP flow sampling maps that are part of the GTP overlap flow sampling map group follow the same configuration considerations discussed previously in *GigaSMART GTP Whitelisting and GTP Flow Sampling* on page 913. As is the case with regular non-overlap GTP mapping, GTP whitelisting selects specific subscribers based on IMSI, whereas GTP flow sampling uses map rules to select subscribers and then forward a percentage of the packets to tool ports.

Configuration Considerations

This section details certain configuration considerations that apply only to the configuration of GTP whitelisting and flow sampling maps for GTP overlap flow sampling maps.

About GTP Overlap Flow Sampling Map Mode and Port Groups

A second level type map specifying GTP overlap flow sampling map mode must be selected to configure GTP whitelisting and flow sampling maps.

To configure a GTP whilelisting map in overlap flow sampling map mode, select **Type Second Level** and **Subtype Flow Whitlelist Overlap** in a map as shown in the following figure.

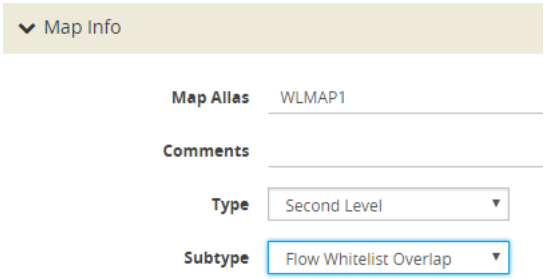

To configure a GTP flow sampling map in GTP overlap flow sampling map mode, select **Type Second Level** and **Subtype Flow Sample Overlap** in a map as shown in the following figure.

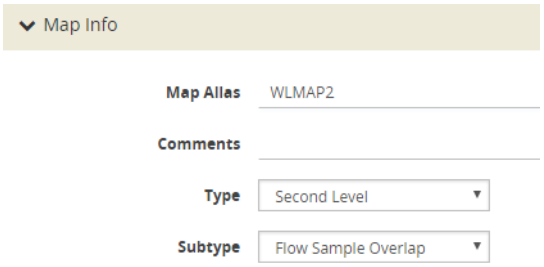

You can configure one GTP whitelisting map and one GTP flow sampling map pair that contain traffic policies corresponding to one destination port group. The load balanced port groups can contain a single port, a port range, or a GigaStream. Note that, starting in software version 4.8, port groups used in GTP overlapping maps support GigaStream.

The maximum number of port groups per single GTP overlap flow sampling map group is six.

For more information about port groups, refer to *Port Groups* on page 412.

Maximum Number of Port Group Members

Use the following sequence to help you determine the maximum number of port group members:

1. Determine the number of members per port group and add 1 to the number.
2. Multiply each port group result times each other.
3. The total multiplication should not exceed 512.

For instance, assume the following configuration in a GTP overlap mapping group:

- Port Group 1—2 load balanced Gigastreams
- Port Group 2—3 load balanced Gigastreams
- Port Group 3—1 load balanced tool port

- Port Group 4—1 load balanced Gigastream
- Port Group 5—4 load balanced tool ports

The total number becomes:

(2+1)*(3+1)*(1+1)*(1+1)*(4+1) = 240

Since this does not exceed the maximum number of multicast IDs (512), the tool configuration shown is accepted.

GTP Overlap Flow Sampling Map Priority

Since a packet matches multiple maps independently the concept of second level map priority does not apply to GTP overlap flow sampling maps. A GTP overlap flow sampling map pair consists of one GTP whitelisting map and one GTP flow sampling map having the same destination port group. Within a GTP overlap flow sampling map pair the whitelisting map rules will be applied before the flow sampling map rules.

Virtual Port Configuration in GTP Overlap Mode

In GTP Overlap map configuration, the virtual port sending traffic to all the port groups needs to be configured in GTP overlap mode.

To configure the virtual port with GTP overlap mode, select **GTP Overlap** when configuring the virtual port as shown in the following figure.

About Map Groups

To create a group of maps for GTP whitelisting and GTP flow sampling, select **Maps > Maps > Map Groups**, and then click **New**. The maps for a map group are entered in the **Maps** field. Refer to Figure 30-57 on page 954. All the maps in a map group receive traffic according to map rules, rather than map priority. Thus, multiple copies of a GTP packet can be sent to more than one tool.

The **Maps** field of the Map Group page groups the whitelisting and flow sampling maps. For example, assuming that two whitelisting maps (*WLMAP1* and *WLMAP2*) and two flow sampling maps (*FSMAP1* and *FSMAP2*) have been configured in GTP overlap

mode, the following example groups them all into the same map group called *map-group1*:

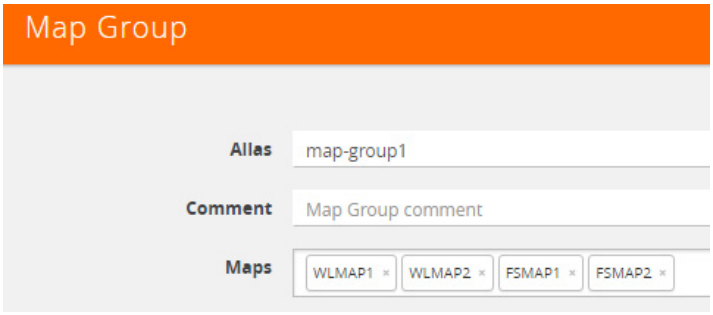

Figure 30-57: Map Group Configuration

Keep in mind the following configuration considerations for map groups:

- A map group can be associated with only one GigaSMART group (gsgroup).
- All maps within a map group must be connected to the same vport.
- A map group can consist of only one GTP whitelisting map or only one GTP flow sampling map but it cannot contains two maps of the same type.
- Once a map group is created, it cannot be edited to change the type or subtype of the map. However, you can add and edit the map rules for a map while it is configured in a map group.
- If multiple map groups are configured, the maps within each map group must point to the same port groups as the other map groups.

For more information about map groups, refer to *Create Map Groups* on page 521.

About Whitelist Maps

The GTP whitelist is an IMSI list which is common to all whitelist maps. You can configure an optional rule within a whitelist map to specify a GTP version or interface-based policy.

Other than specifying a new second level type using **Type Second Level** and **Subtype Flow Whitelist Overlap** when creating the map, the configuration of GTP whitelist maps follows the same configuration guidelines as shown in the section *GTP Whitelisting* on page 913.

A maximum of six whitelist maps sending traffic to six different port groups can be configured per GigaSMART group (gsgroup).

About Flow Sampling Maps

In GTP overlap flow sampling map mode, GTP flow sampling (rule-based flow sampling) is performed after GTP whitelist-based forwarding. Therefore, flow sampling maps have a lower priority than whitelist maps. Thus, within a GTP overlap map pair that consists of a single GTP whitelist overlap map and a GTP flow sampling overlap map, the GTP whitelist map is of higher priority.

Within the flow sampling maps, the rules in the first map have a higher priority than the rules in the second, third, and subsequent maps. Within any single map, rules are evaluated in order.

A maximum of six flow sampling maps sending traffic to six different port groups can be configured per GigaSMART group (gsgroup).

Example 1: GTP Overlap Mode

Example 1 is a GTP overlap mapping mode example.

In Example 1, traffic from a single network port goes to a single first level map (mapLevel1-GTP) which directs GTP-Control, and GTP-User traffic to a virtual port (VP31). Traffic from VP31 is replicated to two GTP whitelisting maps (WLMAP1 and WLMAP2) and two GTP flow sampling maps (FSMAP1 and FSMAP2), which then forward accepted traffic to the final port-group destinations, pg1 and pg2, for load balancing (refer to Figure 30-58 on page 955).

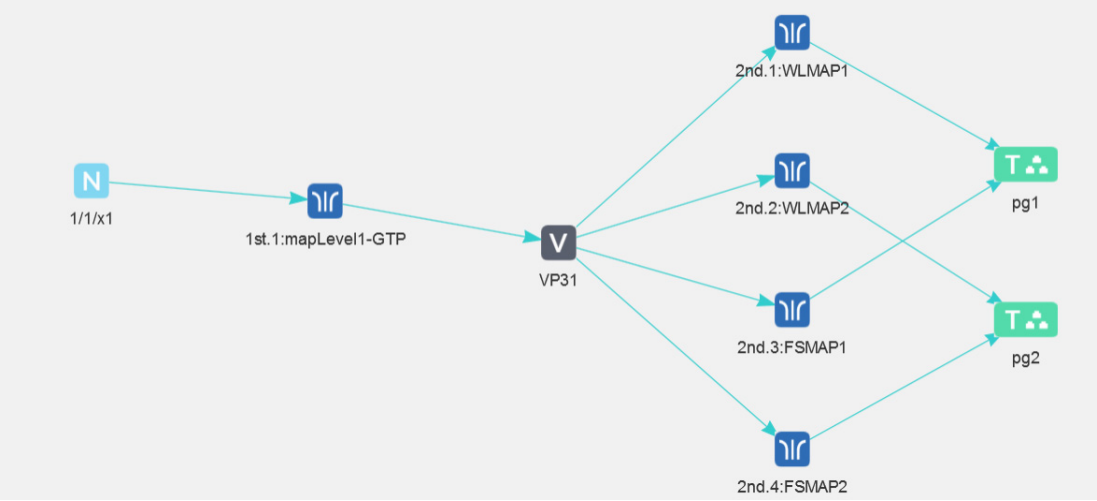

Figure 30-58: GTP Overlap Mode Example 1

NOTE: In Example 1, the tool ports and GigaStreams in the port group are on the same node as the GigaSMART group and GigaSMART operation.

Within each GTP whitelisting and flow sampling pair, if there is not a match to an IMSI in the whitelist map, the traffic flow is sampled based on the rules in the flow sampling map. The flow sampling rules specify IMSI, IMEI, and MSISDN numbers, as well as the percentage to sample.

Within each map pair, packets are then accepted or rejected. Accepted packets are forwarded to the port groups for load balancing. Rejected packets are dropped.

Use the following steps to configure example 1.

Task	Description	UI Steps
1.	Create GigaStreams that will be part of the port groups	**a.** Select Ports > Port Groups > GigaStreams **b.** Click New. **c.** Enter gs1 in the Alias field. **d.** In the Ports field, select port 1/1x16 and 1/1/x17. **e.** Click Save. **f.** Configure a second GigaStream with the alias gs2, select ports1/1/x1 and 1/1/x2 in the Ports field, and click Save.
2.	Create port groups and specify the tool ports and assign GigaStreams to the port groups. The port groups will also be load balanced.	**a.** Select Ports > Port Groups > All Port Groups. **b.** Click New. **c.** Enter pg1 in the Alias field. **d.** Select Type GigaSMART Load Balancing. **e.** In the Ports field, select ports 1/1/x6 and 1/1/x7. **f.** In the GigaStream field, select gs1 **g.** Click Save. **h.** Configure a second Port Group. with the alias pg2, select ports1/1/x18 and 1/1/x10 in the Ports field, select GigaSMART Load Balancing, select pg2 in the GigaStream field, and then click Save.
3.	Configure a GigaSMART group and associate it with a GigaSMART engine port.	**a.** From the device view, select GigaSMART > GigaSMART Groups. **b.** Click New. **c.** Enter GS31 in the Alias field. **d.** In the Port List field, select an engine port. For example 1/3/e1. **e.** Click Save.
4.	Create a virtual port. **NOTE:** You must enable GTP Overlap when configuring a virtual port for GTP overlap mapping.	**a.** From the device view, select GigaSMART > Virtual Ports > Virtual Ports. **b.** Click New **c.** Enter VP31 in the Alias field. **d.** In the GigaSMART Group field, select GS31. **e.** Select GTP Overlap, **f.** Click Save.
5.	Create the GTP Whitelist	**a.** From the device view, select GigaSMART > GTP Whitelist. **b.** Click New. **c.** Enter Whitelist in the Alias field **d.** Go to Task 6.
6.	Fetch whitelist files from a specified location to populate the GTP whitelist.	**a.** On the GTP Whitelist page, select Bulk Upload. **b.** Select Bulk Entry Operation for IMSI Upload Type **c.** Select Upload from URL from the Bulk Upload Type list. **d.** Enter the URL in the Enter Remote URL field. For example, http://10.1.1.100/tftpboot/myfiles/MyIMSIs_file2.tx **e.** Click Save.

Task	Description	UI Steps
7.	Associate the GigaSMART group to the GTP whitelist.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups **b.** Select GS Group GS31 created in Task 3 and click Edit **c.** Under GTP Whitelist, click on the GTP Whitelist Alias field and select Whitelist. **d.** Click Save.
8.	Configure the GigaSMART operation for GTP whitelisting.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Enter gtp-overlapwhitelist1 in the Alias field. **d.** Select the GigaSMART Group GS31from the GigaSMART Groups list. **e.** Select GTP Whitelist from the GigaSMART Operations (GSOP) list **f.** Select Load Balancing from the GigaSMART Operations (GSOP) list. **g.** Configure Load Balancing as follows: • Select Stateful • Set Type to GTP • Select Hashing • Select IMSI **h.** Click Save.
9.	Configure the GigaSMART operation for GTP flow sampling.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Enter gtp-overlapsampling1 in the Alias field. **d.** Select the GigaSMART Group GS31from the GigaSMART Groups list. **e.** Select Flow Sampling-GTP **f.** Select Load Balancing from the GigaSMART Operations (GSOP) list. **g.** Configure Load Balancing as follows: • Select Stateful • Set Type to GTP • Select Hashing • Select IMSI **h.** Click Save.
10.	Configure the first level maps. **NOTE:** In the rules, 2123 is GTP-c traffic and 2152 is GTP-u traffic.	**a.** Configure the first map as follows: • Alias: mapLevel1-GTP • Type and Subtype: First Level By Rule • Source: 1/1/x1 • Destination: VP31 • Rule 1: Pass, Bi Directional, Port Destination 2123 • Rule 2: Pass, Bi Directional, Port Destination 2152 • Click Save.

Task	Description	UI Steps
11.	Configure the first second level GTP overlap map for GTP whitelisting If there is a match to an IMSI in the whitelist for GTP version 1 traffic, it is then forwarded to load balancing port group *pg1*.	**a.** Select Maps > Maps > Maps. **b.** Click New **c.** Configure the map. • Alias: WLMAP1 • Type and Subtype: Second Level GTP Flow Whitelist Overlap • Source: VP3 • Destination: pg1 • GSOP: gtp-whitelist • Rule 1: GTP, APN: Version V1 **d.** Click Save.
12.	Configure a second level map for GTP flow sampling, the flow sampling map. If there is not a match to an IMSI in the whitelist, the traffic flow is sampled based on the rules in this map. Accepted packets are forwarded to load balancing port group *pg1*.	**a.** Select Maps > Maps > Maps. **b.** Click New **c.** Configure the map. • Alias: FSMAP1 • Type and Subtype: Second Level GTP Flow Sample Overlap • Source: VP3 • Destination: pg1 • GSOP: gtp-overlapsample1 • Rule 1: GTP, IMSI: 3102609834*, IMEI: 35609506*, Percentage: 20 **d.** Click Save.
13.	Configure the next second level GTP overlap map for GTP whitelisting If there is a match to an IMSI in the whitelist for GTP version 2 traffic, it is then forwarded to load balancing port group *pg2*.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Configure the map. • Alias: WLMAP2 • Type and Subtype: Second Level Flow Whitelist Overlap • Source: VP31 • Destination: pg2 • GSOP: gtp-whitelist • Rule 1: GTP, APN: Version V2 **d.** Click Save.
14.	Configure the next second level map for GTP flow sampling. If there is not a match to an IMSI in the whitelist as evaluated by the second level GTP whitelisting map *WLMAP2*, the traffic flow is sampled based on the rules in this map. Accepted packets are forwarded to load balancing port group *pg2*.	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Configure the map: • Alias: FSMAP2 • Type and Subtype: Second Level GTP Flow Sample Overlap • Source: VP31 • Destination: pg2 • Rule 1: GTP, IMSI: 3102609835*, IMEI: 35609507*, Percentage: 20 **d.** Click Save.
15.	Configure a map group. Add the GTP whitelisting and the two GTP flow sampling maps configured in previous steps.	**a.** Select Maps > Map Groups. **b.** Click New. **c.** Enter OverlapMap in the Alias field. **d.** In the Maps field, select WLMAP1,WLMAP2,FSMAP1,FSMAP2. **e.** Click Save.

Overlap Map Statistics

Starting with version 5.4 overlap and non-overlap maps are available. Overlap maps are displayed based on the following:

- If at least 1 flow-sample map accepts the packets among all the pairs of overlap maps, the Sample (Tx) counters in the GTP Interface stats will be incremented. If more than 1 pair of maps accepts the packets, the Sample (Tx) counters in the GTP Interface stats is incremented only once.
- If at least 1 Whitelist map matches the packets among all the pairs of overlap maps, the Sample (Tx) counters in the GTP Interface stats will be incremented. If more than 1 pair of maps matches the packets, the Sample (Tx) counters in the GTP Interface stats is incremented only once.
- If there are no WL maps and all flow sample maps are no rule match, then Sample(Tx) and Sample Out counters in the GTP Interface stats is not incremented.

GTP Scaling

GTP can be scaled as follows:

- *GigaSMART Cards in GigaVUE-HD8* on page 959
- *GTP Engine Grouping* on page 959

GigaSMART Cards in GigaVUE-HD8

Required License: GTP Filtering & Correlation

Starting in software version 4.5, a total of six GigaSMART-HD0 line cards are supported on a single GigaVUE-HD8 node. This provides a total of twelve GigaSMART engine ports, which increases the amount of GigaSMART processing available on the GigaVUE-HD8.

The increased number of GigaSMART line cards in the GigaVUE-HD8 can be used by the following GTP applications: GTP flow filtering, GTP flow sampling, and GTP whitelisting.

GTP Engine Grouping

Required License: GTP Filtering & Correlation

A GigaSMART group (gsgroup) associated with GTP applications can have multiple GigaSMART engine port members. Up to four engine ports can be combined to form an engine group. The engine group provides higher capacity to GTP applications by load balancing GTP user-data plane (GTP-u) traffic among the members of the group. Grouping multiple GigaSMART engine ports increases the effective throughput for GTP applications.

NOTE: Software version 5.5 supports a maximum of 6 million GTP subscriber sessions for GigaVUE-HC2 nodes, whereas, it supports 12 million GTP subscriber sessions for GigaVUE-HC3 nodes.

Starting in software version 4.5.01, it is supported on GigaVUE-HC2 nodes. GTP engine grouping can be used by the following GTP applications: GTP flow filtering, GTP flow sampling, and GTP whitelisting.

The following table lists GTP engine grouping support for GigaVUE nodes:

GigaVUE Node	Maximum Number of GigaSMART Line Cards per Node	Number of e ports per Line Card	Supported Number of e ports per GigaSMART Group	Location of e ports
GigaVUE-HC3	4	2	8	**e1** and **e2** on same module
GigaVUE-HC2	4	1	4	any front modules
GigaVUE-HC1	Not supported in this software version.			
GigaVUE-HB1	Not supported in this software version.			

Keep in mind the following recommendations and restrictions:

- Configure a GTP engine group on a single GigaVUE node.
- Configure either two or four engine ports per GigaSMART group in a GTP engine group.

  NOTE: One engine port per GigaSMART group is the default behavior supported in software version 4.5 and prior releases without GTP engine grouping.
- On the GigaVUE-HC2, use only GigaSMART front modules in the engine group (not GigaSMART rear modules).
- GTP engine grouping only supports IMSI hash-based load balancing.
- GTP engine grouping is limited to out-of-band cluster configurations in this software version.

Pass GTP Control Traffic

Selecting the **Control** option for **Traffic Type** of the a First Level By Rule map specifies an option for GTP applications to pass GTP control traffic (GTP-c) to all GigaSMART engines in a GTP engine group. GTP-c traffic is sent to all members of the engine group in order to replicate the session tables.

NOTE: In the map with the **Control** selected for **Traffic Type**, only one vport is supported.

Starting in software version 4.7, the map with the **Traffic Type Control** set can be edited. For example, **Traffic Type Control** can be added to an existing first level map, or it can be deleted from a first level map by clearing the **Control** checkbox for **Traffic Type** (refer to Figure 30-59). Other editing, such as changing the **Source** or the **Destination** in the first level map is also allowed.

Starting in software version 4.7, the map with **Traffic Type Control** set can be edited. For example, **Control** can be selected in an existing first level map, or it can be deleted from a first level map by clearing **Control**. Other editing, such as changing the **Source** or the **Destination** in the first level map is also allowed.

To set GTP Control Traffic:

1. From the device view, select **GigaSMART > Maps > Maps** and click **New**.
2. On the New Map page, select **First Level** and **By Rule** for the map type and subtype, respectively.
3. Select **Control** as shown in Figure 30-59
4. Configure the map and click **Save** when done.

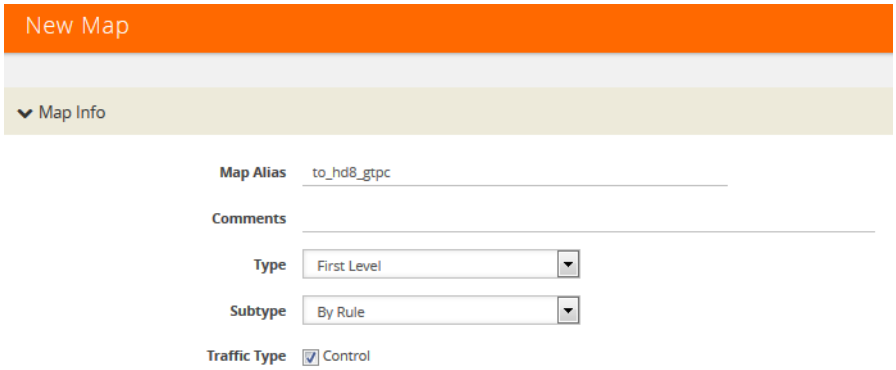

Figure 30-59: GTP Control Traffic Selected

Upgrade from an Earlier Release

When there is existing GTP configuration with one engine port per GigaSMART group in a pre-4.5 software version, an upgrade from that earlier software version to 4.5 or a higher release will succeed.

However in the 4.5 or higher release, you cannot convert that configuration to multiple engine ports per GigaSMART group. You must delete the configuration and reconfigure it, including the GigaSMART group, GigaSMART operation, virtual port, and maps. This is due to the need for separate maps for GTP control plane and GTP user plane traffic in 4.5 and higher releases.

Modify Engine Ports in GigaSMART Group

You can modify the engine ports in a GigaSMART group. For example, you can add an engine port to a GigaSMART group or remove an engine port from a GigaSMART group. After the change, reset all the GigaSMART line cards or modules that have engine ports configured in the GigaSMART group.

Modify vports in a Map

You can modify the vport relating to the first level map with Traffic Type Control set. For example, you can change the vport configured in the map. After the change, reset all the GigaSMART line cards or modules that have engine ports configured in the vport.

Configure GTP Engine Grouping

Refer to the following examples:

GTP Engine Grouping Configuration Example

This is an example of a GTP engine group consisting of two engine ports on a GigaVUE-HD4 node. This example includes a GigaSMART operation for GTP flow filtering.

Task	Description	UI Steps
1	Configure ports as follows: • one network type of port. This will be used as the **Source** attribute in two first level maps in **Task 5** and **Task 6**. • one tool type of port for the **Destination** attribute in a second level flow filtering map in **Task 7**. • one tool type of port for the **Destination** attribute in a shared collector map in **Task 8**. Then administratively enable the ports.	**1.** Select Ports > Ports > All Ports **2.** Click Quick Port Editor **3.** Configure a network port and two tool ports. For example, select Network for port 22/3/x3 and select Tool for ports 22/3/x1 and 22/1/x11. **4.** Select Enable for each port. **5.** Click OK. **6.** Close the Quick Port Editor.
2	Configure a GigaSMART group and associate it with two GigaSMART engine ports, to form the GTP engine group. The GigaSMART group will be used in **Task 5** and **Task 6**.	**1.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **2.** Type gsg2 in the Alias field. **3.** Click in the Port List field to add the two engine ports. **4.** Click Save.
3	For GTP flow filtering, configure a flow filtering GigaSMART operation and assign it to the GigaSMART group. The gsop will be used in the second level flow filtering map in **Task 7**.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **2.** Type gtp_gsg2 in the Alias field. **3.** Select gsg2 from the GigaSMART Groups list. **4.** Select Flow Filtering from the GigaSMART Operations list. (GSOP). **5.** Click Save.
4	Configure a virtual port and assign it to the same GigaSMART group. This virtual port will be used as the **Destination** in the first level maps in **Task 5** and **Task 6**, as the **Source** in the second level map in **Task 7**, and as the **Source** attribute in the shared collector map in **Task 8**.	**1.** From the device view, select GigaSMART > Virtual Ports. **2.** Click New. **3.** Type vp1 in the Alias field. **4.** Select gsg2 from the GigaSMART Groups list. **5.** Click Save.

Task	Description	UI Steps
5	Create a first level map that directs GTP control traffic from the physical network port to the virtual port created in **Task 4**. **NOTE:** In the rule, 2123 is GTP-c traffic. This map, with the **Traffic Type Control** attribute, identifies the GTP-c control traffic needed for GTP engine grouping. **NOTE:** The order of configuration is important. Set **Traffic Type Control** before any map rules.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. • Type gtp_to_vp1-c in the Alias field • Select First Level for Type • Select By Rule for Subtype • Select Control for Traffic Type • Select 22/3/x3 for Source • Select vp1 for Destination **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass and Bi Directional **c.** Select Port Destination and specify 2123 **5.** Click Save.
6	Create another first level map that directs GTP user traffic from the physical network port to the virtual port created in **Task 5**. **NOTE:** In the rule, 2152 is GTP-u traffic. GTP-u traffic corresponding to the same GTP-c traffic will be sent to the same virtual port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. • Type gtp_to_vp1 in the Alias field • Select First Level for Type • Select By Rule for Subtype • Select Control for Traffic Type • Select the network configured in Task 1 for Source • Select vp1 for Destination **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass and Bi Directional **c.** Select Port Destination and specify 2123 **5.** Click Save.
7	Create a second level map for GTP flow filtering that takes traffic from the virtual port, applies the flow filtering GigaSMART operation, matches IMEIs and version specified by the flow rule, and sends matching traffic to a tool port.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Type from_vp1 in the Alias field. **4.** Configure the map. • Select Second Level for Type • Select Flow Filter for Subtype • Select vp1 for Source • Select one of the tool ports configured in Task 1 for Destination. • Select gtp_gsg2 for from the GSOP list. **5.** Add a Rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select GTP IMEI. **d.** Specify * in the IMEI field. **e.** Select V2 for Version. **6.** Click Save.

Task	Description	UI Steps
8	Add a shared collector for any unmatched traffic from the virtual port and send it to a different tool port than in **Task 7**.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. • Type from_vp1_scoll in the Alias field. • Select Second Level for Type • Select Collector for Subtype • Select vp1 for Source. • Select the other tool port configured in Task 1 for Destination. **4.** Click Save.

GTP Engine Grouping Configuration Complex Example

This is a more complex example of GTP engine grouping than the previous example. This example has four engine ports on two GigaSMART line cards on the same GigaVUE-HD8 node. The GigaSMART line cards are in slots 1 and 3.

The GigaVUE-HD8 node is the cluster master of a two-node out-of-band cluster. A GigaVUE-HD4 is the standby node in the cluster.

This example includes GigaSMART operations for GTP flow filtering with load balancing, GTP flow sampling with load balancing, and GTP whitelisting. The whitelist must be associated with the GigaSMART group on the master node, the GigaVUE-HD8.

Task	Description	UI Steps
1	Configure ports on the GigaVUE-HD8 as follows: • One network type of port. This will be used as the **Source** attribute in two first level maps in **Task 11** and **Task 12**. • Twelve tool type of ports. There are four tool ports in each of three port groups used for load balancing. The port groups will be created in **Task 6**. • Five tool type of ports for a GigaStream that will be created in **Task 2**. • Two tool type of ports for another GigaStream that will be created in **Task 2**. Then administratively enable the ports.	**1.** On the GigaVUE-HD8, select Ports > Ports > All Ports **2.** Click Quick Port Editor. **3.** Configure one network port for **Task 11** and **Task 12**. For example, 23/7/6. **4.** Configure 12 tool ports for three port groups in **Task 6**. For example, 23/4/x1..x4 for the first port group, 23/4/x9..x12 for the second tool group, and 23/4/x13..x16 for the third port group. **5.** Configure five tool ports for a GigaStream in **Task 2**. For example, 23/4/x28..x32. **6.** Configure two tool ports for another GigaStream in **Task 2**. For example, 23/7/q1..q2 type tool. **7.** Select enable for each port. **8.** Click OK. **9.** Close the Quick Port Editor.

Task	Description	UI Steps
2	On the GigaVUE-HD8, configure one GigaStream using five tool ports. This will be used as the **Destination** attribute in the map in **Task 11**. Configure another GigaStream to be used in the stack link between the GigaVUE-HD8 and GigaVUE-HD4 that will be created in **Task 4**.	**1.** Select Ports > Port Groups > GigaStreams **2.** Configure a GigaStream with five tool ports. **a.** Click New. **b.** Type hd8-gs-1 in the Alias field. **c.** Select Tool GigaStream. **d.** Click in the Ports field and select the five tool ports configured in Task 1. **e.** Click Advanced Hash Settings and use the Default setting. **f.** Click Save. **3.** Configure another GigaStream with two tool ports. **a.** Click New. **b.** Type hd8-80g in the Alias field. **c.** Select Tool GigaStream. **d.** Click in the Ports field and select the two tool ports configured in Task 1. **e.** Click Advanced Hash Settings and use the Default setting. **f.** Click Save.
3	Configure ports on the GigaVUE-HD4 as follows: • Two tool type of ports for a GigaStream that will be created in **Task 4**. • One tool type of port that will be used as the **Destination** in a map in **Task 11**. • Four tool type of ports for a GigaStream that will be created in **Task 4**. Then administratively enable the ports.	**1.** On the GigaVUE-HD4, select Ports > Ports > All Ports **2.** Click Quick Port Editor. **3.** Configure two tool ports for a GigaStream in **Task 4**. For example, 33/2/q1..q2 **4.** Configure one tool ports for the map in **Task 11**. For example, 33/3/x11. **5.** Configure four tool ports for a GigaStream in **Task 4**. For example, 33/2/x25..x28. **6.** Select Enable for each port. **7.** Click OK. **8.** Close the Quick Port Editor.
4	On the GigaVUE-HD4, configure a GigaStream using two tool ports. This will be used in the stack link created in **Task 5**. Configure another GigaStream using four tool ports. This will be used in the shared collector in **Task 17**.	**1.** Select Ports > Port Groups > GigaStreams **2.** Configure a GigaStream with five tool ports. **a.** Click New. **b.** Type hd4-gs-4 in the Alias field. **c.** Select Tool GigaStream. **d.** Click in the Ports field and select the five tool ports configured in Task 1. For example, 33/2/x25..x28. **e.** Click Advanced Hash Settings and use the Default setting. **f.** Click Save. **3.** Configure another GigaStream with two tool ports. **a.** Click New. **b.** Type hd8-80g in the Alias field. **c.** Select Tool GigaStream. **d.** Click in the Ports field and select the two tool ports configured in Task 1. For example, 33/2/q1..q2. **e.** Click Advanced Hash Settings and use the Default setting. **f.** Click Save.

Task	Description	UI Steps
5	Configure the stack link between the GigaVUE-HD4 and GigaVUE-HD8.	1. Select Ports > Port Groups > Stack Links 2. Click New. 3. Select Stack GigaStream. 4. For First Member select hd8-80g. 5. For Second Member select hd4-80g
6	Create three port groups and specify four tool ports each, for load balancing. Also, enable load balancing on each port group. The port groups, hd8-pg-1 and hd8-pg-2, will be used as the **Destination** in two second level flow sampling maps in **Task 14** and **Task 15**. The port group, hd8-q2x32-1-4, will be used as the **Destination** in a second level flow filtering map in **Task 16**	1. Select Ports > Port Groups > All Port Groups. 2. Create the first port group. a. Click New. b. Type hd8-pg-1 in the Alias field. c. Select SMART Load Balancing. d. Click in the Ports field and select four tool ports. For example, 23/4/x9..x12. e. Click Save. 3. Create the second port group. a. Click New. b. Type hd8-pg-1 in the Alias field. c. Select SMART Load Balancing. d. Click in the Ports field and select four tool ports. For example, 23/4/x13..x16 e. Click Save. 4. Create the second port group. a. Click New. b. Type hd8-q2x32-1-4 in the Alias field. c. Select SMART Load Balancing. d. Click in the Ports field and select four tool ports. For example, 23/4/x13..x16 e. Click Save.
7	Configure a GigaSMART group and associate it with four GigaSMART engine ports, two in slot 1 and two in slot 3, to form the GTP engine group. The GigaSMART group will be used in **Task 8**, **Task 9**, and **Task 10**.	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type hd8scale-4engines-slots1and3. 4. Click in the Port List and select the engine ports. For example, 3/1/e1,23/1/e2,23/3/e1,23/3/e2. Go to **Task 8**.
8	Associate the GigaSMART group to an existing GTP whitelist. **NOTE:** The whitelist is only supported on the cluster master node, which is the GigaVUE-HD8 in this example.	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Under GTP Whitelist, select the alias of an existing whitelist. For example, gtp-whitelist. 3. Click Save.

Task	Description	UI Steps
9	For GTP flow filtering, configure a flow filtering GigaSMART operation, specify load balancing, and assign the GigaSMART operation to the GigaSMART group. The hd8-scale-ff-lb gsop will be used in the second level flow filtering map in **Task 16**. For GTP flow sampling, configure a flow sampling GigaSMART operation, specify load balancing, and assign the GigaSMART operation to the GigaSMART group. The hd8-scale-fs-lb gsop will be used in the two second level flow sampling maps in **Task 15** and **Task 16**. For GTP whitelisting, configure a whitelisting GigaSMART operation, and assign the GigaSMART operation to the GigaSMART group. (This GigaSMART operation is not load balanced.) The hd8-scale-wl gsop will be used in the second level whitelisting map in **Task 13**.	**1.** Configure a Flow Filtering operation. **a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Type hd8-scale-ff-lb in the Alias field. **d.** Select hd8scale-4engines-slots1and3 from the GigaSMART Groups list. **e.** Select Flow Filtering from the GigaSMART Operations (GSOP) list. **f.** Click Save. **2.** Configure a Flow Sampling operation. **a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Type hd8-scale-fs-lb in the Alias field. **d.** Select hd8scale-4engines-slots1and3 from the GigaSMART Groups list. **e.** Select Flow Sampling from the GigaSMART Operations (GSOP) list. **f.** Select Flow Sampling -GTP **g.** Click Save. **3.** Configure a Whitelisting operation. **a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Type hd8-scale-wl in the Alias field. **d.** Select hd8scale-4engines-slots1and3 from the GigaSMART Groups list. **e.** Select GTP Whitelist from the GigaSMART Operations (GSOP) list. **f.** Select Enabled (default). **g.** Click Save.
10	Configure a virtual port and assign it to the same GigaSMART group. This virtual port will be used as the **Destination** in the first level maps in **Task 11** and **Task 12**, as the **Source** in the second level maps in **Task 13**, **Task 14**, **Task 15**, **Task 16**, and as the **Source** in the shared collector in **Task 17**.	**1.** From the device view, select GigaSMART > Virtual Ports **2.** Click New. **3.** Type vp-hd8scale-4engines-slots1and3 in the Alias field. **4.** Select hd8scale-4engines-slots1and3 from the GigaSMART Groups list. **5.** Click Save.

Task	Description	UI Steps
11	Create a first level map that directs GTP control traffic from the physical network port to the virtual port created in **Task 10**. **NOTE:** In the rule, 2123 is GTP-c traffic. This map, with the **Traffic Type Control** enabled, identifies the GTP-c control traffic needed for GTP engine grouping. In addition to the virtual port, traffic is also sent to a GigaStream and a tool port.	**1.** Select Maps > Maps > Maps. **2.** Configure the map. • Type to_hd8_gtpc in the Alias field • Select First Level for Type. • Select By Rule for Subtype • Select Control for Traffic Type • Select a network port for the Source. For example, 23/7/q6. • Select the GigaStream port hd8-gs-1, the virtual port p-hd8scale-4engines-slots1and3, and a tool port (for example, 23/7/q6) for the Destination. **3.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass and Bi Directional. **c.** Select Port Destination for the rule. **d.** Enter 2123 for the port value. **4.** Click Save.
12	Create another first level map that directs GTP user traffic from the physical network port to the virtual port created in **Task 10**. **NOTE:** In the rule, 2152 is GTP-u traffic. GTP-u traffic corresponding to the same GTP-c traffic will be sent to the same virtual port.	**1.** Select Maps > Maps > Maps. **2.** Configure the map. • Type to_hd8_gtpu_1 in the Alias field • Select First Level for Type. • Select By Rule for Subtype • Select a network port for the Source. For example, 23/7/q6. • Select the virtual port vp-hd8scale-4engines-slots1and3 for the Destination. **3.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass and Bi Directional. **c.** Select Port Destination for the rule. **d.** Enter 2152for the port value. **4.** Add a second rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Fragmentation for the rule. **d.** Enter 2152for the port value. **e.** Select allFragNoFirst for Value. **5.** Click Save.

Task	Description	UI Steps
13	Configure a second level map for GTP whitelisting, the whitelist map, that takes traffic from the virtual port, applies the whitelisting GigaSMART operation, and sends traffic to the remote GigaVUE-HD4 node through a GigaStream.	**1.** Select Maps > Maps > Maps. **2.** Configure the map. • Type from_hd8_wl in the Alias field • Select Second Level for Type. • Select By Rule for Subtype • Select the virtual port vp-hd8scale-4engines-slots1and3 for the Source. • Select the GigaStream hd4-gs-1 for the Destination. • Select hd8-scale-wl from the GSOP list. **3.** Click Save.
14	Configure a second level map for GTP flow sampling. This is the first of two flow sampling maps. This map filters for version 2. It takes traffic from the virtual port and applies the flow sampling GigaSMART operation. Traffic flow is sampled based on the flow sampling rule in this map. Accepted packets are forwarded to load balancing port group hd8-pg-2.	**1.** Select Maps > Maps > Maps. **2.** Configure the map. • Type from_hd8_fs_v2 in the Alias field • Select Second Level for Type. • Select Flow Sample for Subtype • Select the virtual port vp-hd8scale-4engines-slots1and3 for the Source. • Select the port group hd8-pg-2 for the Destination. • Select hd8-scale-fs-lb from the GSOP list. **3.** Add a rule. **a.** Click Add a Rule. **b.** Select GTP. **c.** Enter 60 for Percentage. **d.** Enter 5* in the IMSI field. **e.** Select V2 for Version **4.** Click Save.
15	Configure a second level map for GTP flow sampling. This is the second of two flow sampling maps. This map filters for version 1. It takes traffic from the virtual port and applies the flow sampling GigaSMART operation. Traffic flow is sampled based on the flow sampling rule in this map. Accepted packets are forwarded to load balancing port group hd8-pg-1.	**1.** Select Maps > Maps > Maps. **2.** Configure the map. • Type from_hd8_fs_v1 in the Alias field • Select Second Level for Type. • Select Flow Sample for Subtype • Select the virtual port vp-hd8scale-4engines-slots1and3 for the Source. • Select the port group hd8-pg-1 for the Destination. • Select hd8-scale-fs-lb from the GSOP list. **3.** Add a rule. **a.** Click Add a Rule. **b.** Select GTP. **c.** Enter 60 for Percentage. **d.** Enter 5* in the IMSI field. **e.** Select V1 for Version **4.** Click Save.

Task	Description	UI Steps
16	Create a second level map for GTP flow filtering that takes traffic from the virtual port, applies the flow filtering GigaSMART operation, matches IMSIs specified by the flow rule, and sends matching traffic to load balancing port group hd8-q2x32-1-4.	**1.** Select Maps > Maps > Maps. **2.** Configure the map. • Type from_hd8_ff in the Alias field • Select Second Level for Type. • Select Flow Filter for Subtype • Select the virtual port vp-hd8scale-4engines-slots1and3 for the Source. • Select the port group hd8-pg-1 for the Destination. • Select port group hd8-q2x32-1-4 from the GSOP list. **3.** Add a rule. **a.** Click Add a Rule. **b.** Select GTP IMSI **c.** Enter * in the IMSI field. **d.** Select Any for Version **4.** Click Save.
17	Add a shared collector for any unmatched traffic from the virtual port and send it to a GigaStream.	**1.** Select Maps > Maps > Maps. **2.** Configure the map. • Type s_coll_hd8 in the Alias field • Select Second Level for Type. • Select Collector for Subtype • Select the virtual port vp-hd8scale-4engines-slots1and3 for the Source. • Select GigaStream hd4-gs-4 for the Destination. **3.** Click Save.

Displaying Statistics

To display the GigaSMART Group statistics, select **GigaSMART > GigaSMART Groups > Statistics**.

Refer to *GigaSMART Group Statistics Definitions* on page 786 for descriptions of the statistics.

GTP Stateful Session Recovery

Required License: GTP Filtering & Correlation

GTP sessions can be backed up periodically so they can then be recovered faster after a GigaSMART line card reboot or a node reboot. GTP stateful session recovery provides session persistence for GigaSMART GTP applications, including GTP flow filtering, GTP whitelisting, and GTP flow sampling.

GTP stateful session recovery requires additional memory for storing backups. GigaVUE-HC3 has the required memory. For GigaVUE HD Series, a memory upgrade for control card HCCv2 is available. For GigaVUE-HC2, Control Card version 2 (HC2 CCv2) is required. Contact your Sales representative or authorized partner for the required control cards for GigaVUE HD Series and GigaVUE-HC2.

Using GTP stateful session recovery, the GTP session tables in the GigaSMART line card memory will be periodically backed up to the control card memory on the node and stored.

You can configure an interval for how often the backups occur, such as every 10 minutes. If GTP stateful session recovery is enabled and the GigaSMART line card is rebooted, the GTP session tables will be restored automatically following the reboot.

The last stored backup file will be downloaded from the control card to the GigaSMART line card using FTP. The session table will be repopulated from the last stored backup file to each GigaSMART engine, up to 8 engines. Packet count statistics for sessions are saved and will also be restored.

Depending on the size of the session table, the amount of time to restore from the backup might take as much as 3 minutes. During that interval, traffic will be blocked to the virtual port on the GigaSMART line card. Once the session table is read and populated, traffic will be allowed.

Depending on the interval between backups, there could be differences between the stored state and the current state of the system, for example, map configuration could change, or sessions could be added, modified, or deleted.

Load balancing information is not persisted, so after a session table is repopulated, a session that was once sent to one load balanced port may be sent to a different load balanced port after the reboot. However, for IMSI-based load balancing, the traffic might be sent to the same port as it was before the reboot.

GTP stateful session recovery works in a cluster environment; however, the cluster master must remain the same.

To enable GTP persistence, as well as to configure timers, use to the GTP Persistence fields under GigaSMART Parameters on the GigaSMART Group configuration page shown in Figure 30-60 and select **GTP Persistence** The timers are preconfigured with default values.

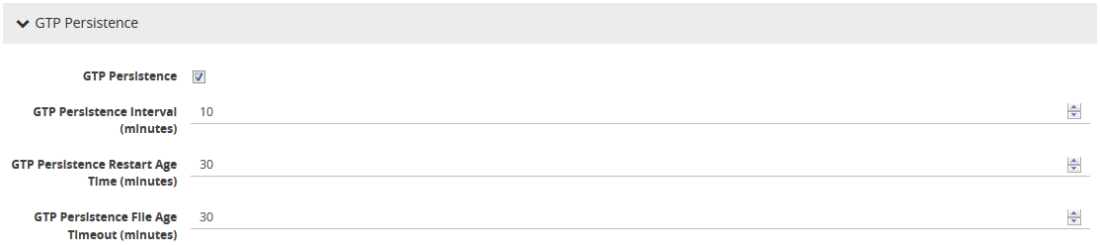

Figure 30-60: GTP Persistence GigaSMART Parameters

Use the **System** widget on the Overview page to determine the amount of memory. The size of memory will be 24Gb in an upgraded system. To view the System information, select **Overview** from the Navigation pane. The amount of free and used memory is displayed in the **Memory** field as indicated in Figure 30-61.

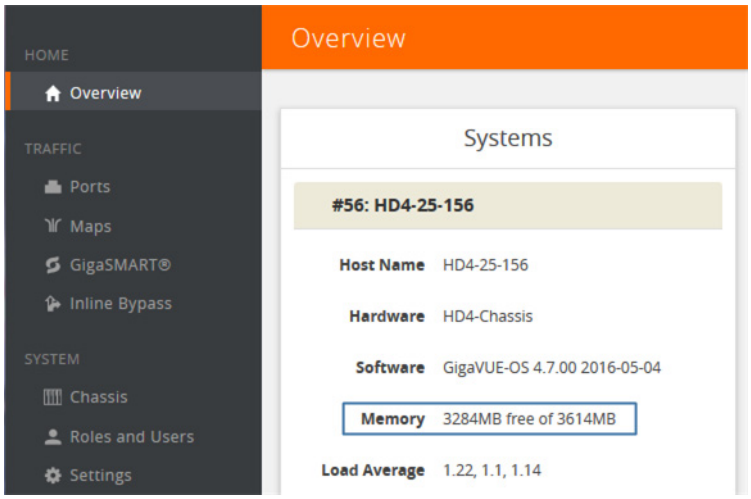

Figure 30-61: System Memory

To see backup and restore information for GTP Persistence, select **GigaSMART > GigaSMART Groups > GigaSMART Group**s, and then click on the alias of the GigaSMART group. A Quick View opens for the selected GigaSMART group. Scroll down do GTP Persistence. In Figure 30-62 on page 973, GigaSMART Group gsgrp-1_4_e1 is selected and the Quick View displayed.

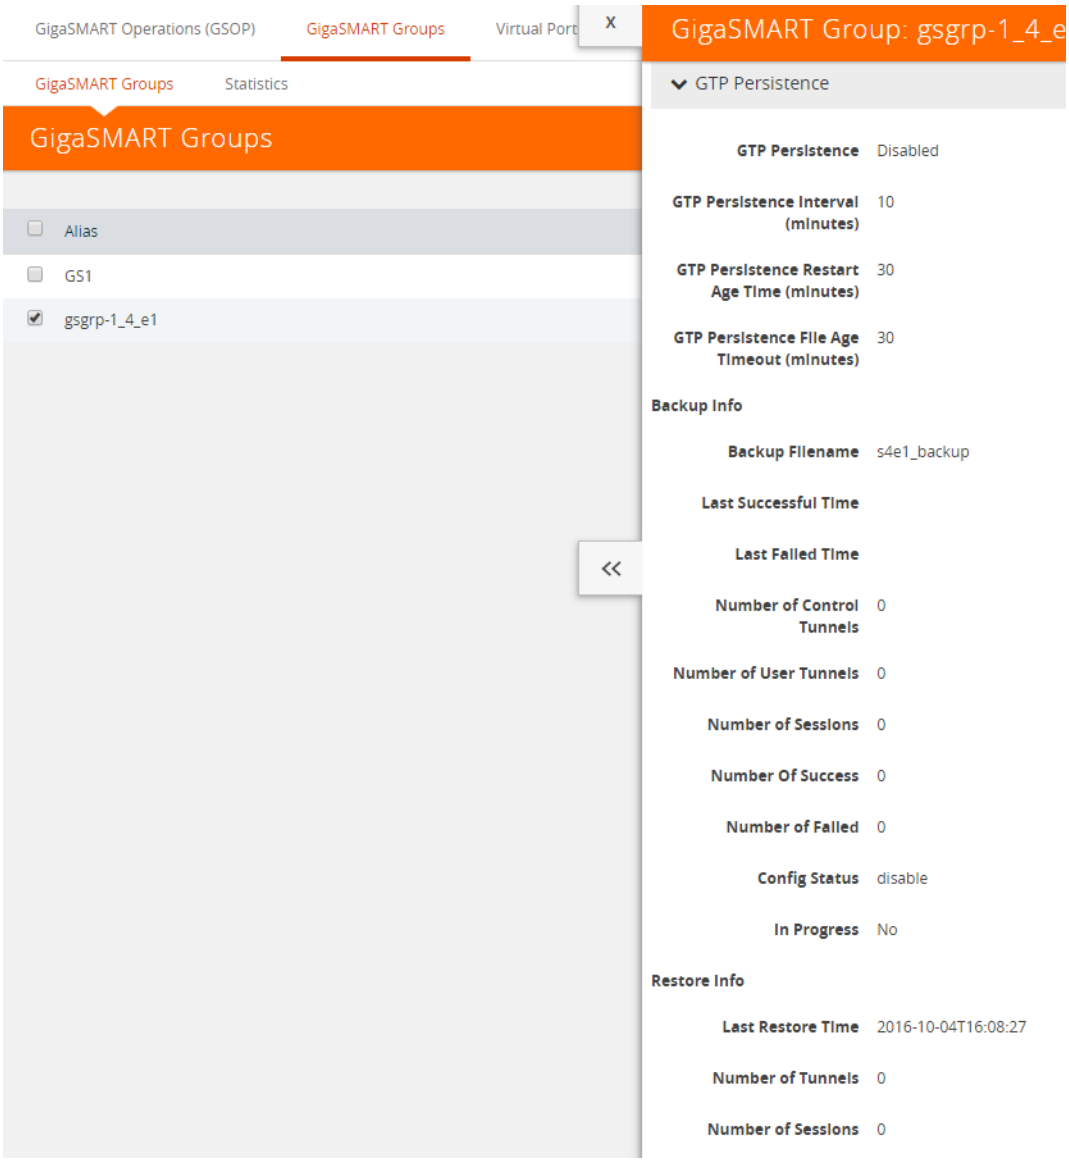

Figure 30-62: GTP Persistence Information

The following table describes persistence information.

Table 30-1: GigaSMART GTP Persistence Information

Name	Format
Backup Info	
Backup filename	The internal name of the backup file.
Last successful time	The timestamp of the last successful backup.
Last fail time	The timestamp of the last failed backup.
Number of control tunnels	The number of control tunnels backed up.
Number of user tunnels	The number of user tunnels backed up.
Number of sessions	The number of sessions backed up.

Table 30-1: GigaSMART GTP Persistence Information

Name	Format
Number of success	The number of successful backups.
Number of failed	The number of failed backups.
Config Status	The status of a backup, which will be either Enabled or Disabled.
In Progress	The progress, which will be either Yes or No.
Restore Info	
Last restore time	The timestamp of the last restore.
Number of tunnels	The number of tunnels restored.
Number of sessions	The number of sessions restored.

To delete backup files, select the alias of GigaSMART Group and click **Edit**.Scroll down to GTP Persistence (refer to Figure 30-63) and click **Delete All** under **GTP Backup Files**.

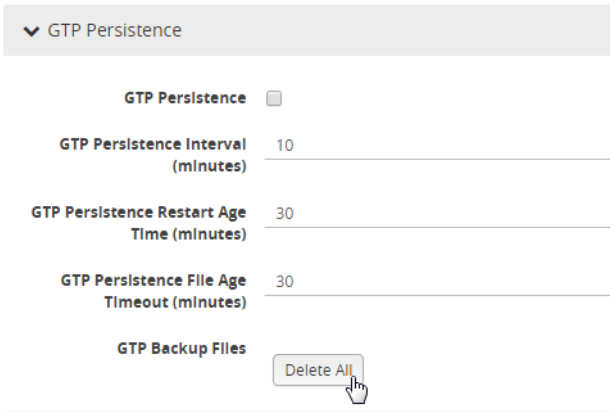

Figure 30-63: GTP Backup Files Delete

GigaSMART SIP/RTP Correlation

Required Licenses: SIP/RTP Correlation (and FlowVUE for session-aware load balancing for RTP)

Session Initiation Protocol (SIP) is the dominant method to initiate, maintain, modify, and terminate voice calls in service provider and enterprise networks. Real-time Transport Protocol (RTP) is used to manage the real-time transmission of voice payload across the same networks. Visibility into a subscriber's voice traffic requires the ability to understand the subscriber attributes and stateful information contained within SIP to correlate subscriber-specific RTP traffic so that monitoring tools can achieve an accurate view of the subscriber's traffic on the network.

The GigaSMART SIP/RTP correlation application correlates the subscriber-specific attributes and the endpoint identifiers of the RTP streams where the session is carried, as well as other SIP-related attributes that are exchanged as part of the control sessions. Use SIP/RTP correlation to leverage a subscriber-aware monitoring policy on Gigamon's Visibility platform and to optimize current tool infrastructure investments by providing only relevant data to tools while increasing visibility into subscriber traffic. This helps improve QoE and performance. Carriers gain access to the subscriber's traffic by reliably correlating and passing all the identified subscriber's control and data sessions to the analytics/monitoring probes and/or billing subsystems for an accurate view of the subscriber's sessions.

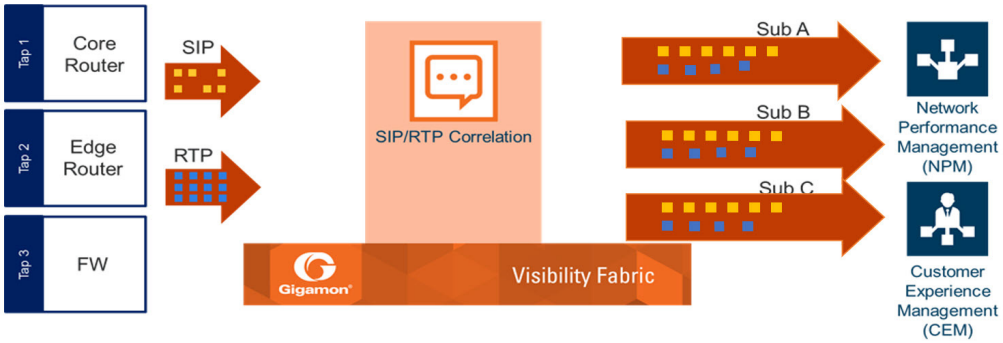

Figure 30-64: SIP/RTP Correlation

SIP is a signaling protocol for VoIP and VoLTE call initiation. It is implemented with RTP to control the payload. GigaSMART SIP/RTP correlation provides customers visibility into VoIP, VoLTE, SMSoLTE, and RCS traffic, and allows them to filter and forward traffic for a subscriber to the tools.

All SIP traffic is sent to all tool ports, as follows:

- all SIP packets sent to all ports within a load balancing port group.
- all SIP packets sent to all ports within a GigaStream.

- all SIP packets sent to all tool ports belonging to maps.

RTP traffic will be sampled and sent to the maps with the rules that match. RTP non-correlated traffic will be sent to the collector.

SIP/RTP correlation can be used for both enterprise and service providers, where ever there is SIP/RTP traffic, such as in wire-line communications, wireless communications, and packet cable networks. This includes enterprise, IP Multimedia Subsystem (IMS), or fixed network implementations of SIP, as well as any media controlled by SIP, such as voice, text, or streaming media.

In addition, SIP/RTP correlation correlates SIP signaling and RTP payload for all sessions selected by a SIP User Agent (UA), a caller ID in a whitelist, with flow sampling from 0 to 100%.

CallerID Tracking

The CallerID is tied to the Call-ID and remains constant for the duration of the session:

CallerID (A) <-----(Call-ID)-----> CalleeID (B)

CallerID tracking is based on the initial caller (A) and does not change until the Call-ID changes. Even if the callee, (B), sends a new INVITE during the SIP session, if that INVITE uses the same Call-ID, then the CallerID information in the SIP session display will still identify the initial caller, (A); it will not switch to reflect that (B) is now the caller.

Support for SIP, RTP, and RTCP

SIP/RTP correlation handles all SIP signaling and RTP/RTCP traffic, including RTCP control traffic, as well as that belonging to a call session.

The following is supported:

- Non-tunneled SIP, RTP, and RTCP correlation (all IMS interfaces)
- Tunneled SIP, RTP, and RTCP correlation (SIP/RTP/RTCP in GTP-U, through the GTP tunnel)

SIP/RTP Correlation Engine

When a packet containing SIP, RTP, or RTCP traffic is received, the SIP/RTP correlation engine looks up the session in the session table for load balancing ports and sampling maps or whitelist map. All SIP/RTP traffic with port or load balancing port group is forwarded based on the session table. The correlation engine load balancing keeps track of both the SIP session and the associated multiple RTP channels.

Each session identifies both sides of media streams (RTP) associated with the session. The SIP session has an aging timer that is configurable.

When a session matches one of the configured flow sampling rules, it is either accepted for sampling or rejected. If it is accepted, all packets belonging to that session are sent to the tool port. Otherwise, all packets belonging to the session are dropped.

All SIP traffic is sent to all tool ports, as follows:

- all SIP packets sent to all ports within a load balancing port group
- all SIP packets sent to all ports within a GigaStream
- all SIP packets sent to all tool ports belonging to maps

RTP traffic will be sampled and sent to the maps with the rules that match. RTP non-correlated traffic will be sent to the collector.

Only one SIP interface type is supported per engine, for example, S5. There is no mixing of interface types, such as S5 GTP-U with SGi.

SIP Whitelist

The SIP whitelist contains caller IDs. Each whitelist entry in a file is a SIP caller ID or callee ID. The whitelist can contain both caller IDs (the from field) and callee IDs (the to field).

Whitelist entries can be both alphabetic and numeric. For each entry, specify up to 64 alphanumeric characters. Some special characters are also supported.

You can manually add one entry at a time to a whitelist file, or you can upload files in.txt format. Each whitelist file can have up to 20,000 entries. One or more whitelist filescan be fetched from a local directory or remote URL using HTTP or SCP.

On GigaVUE-HC2 and GigaVUE HD Series nodes, the whitelist database supports 500,000 entries. On GigaVUE-HC3 nodes, the whitelist database supports 1 million entries.

Multiple whitelist databases can reside on a GigaVUE node, but only one whitelist is applied to a GigaSMART group at a time.

Only one whitelist map is supported for a GigaSMART group. The GigaSMART operation does not have any rules for whitelisting.

RTP Flow Sampling

FlowVUE is used for session-aware (stateful) load balancing and whitelisting with sampling. Only RTP traffic will be sampled. There is no sampling of SIP traffic.

Up to five flow sample maps per GigaSMART group are supported. Each flow sample map can have 20 rules. Use rules to filter on caller ID. The rules support both alphabetic and numeric characters, up to 64 characters. Some special characters are also supported, such as wildcard characters.

Sampling is based on caller ID only (the from field).

Support for Sessions

The number of supported SIP and RTP sessions are as follows:

- GigaVUE-HC3—1 million SIP sessions and 4 million concurrent RTP sessions
- GigaVUE-HC2—500,000 SIP sessions and 2 million concurrent RTP sessions

Each SIP session can handle two RTP streams in both directions (bidirectional).

The number of supported TCP sessions are as follows:

- GigaVUE-HC3—2 million sessions
- GigaVUE-HC2—1 million sessions

Support for IPv4 and IPv6

SIP is a text-based protocol, which is supported over UDP and TCP. The size of the SIP message can vary greatly, so fragmentation and segmentation are common and are supported for tunneled SIP and non-tunneled SIP (IMS).

IPv4 and IPv6 are supported as follows:

- UDP Fragmentation—in-order packets, out-of-order packets
- TCP Segmentation—in-order packets, out-of-order packets

The following is not supported:

- GTP tunneled packets where the inner IP is fragmented
- IMS packets where the outer IP is fragmented

		UDP				TCP		
		Outer Frag	Inner Frag	In order	Out of order	Segmentation	In order	Out of order
GTP	IPv4	✓	×	✓	✓	✓	✓	✓
		✓	×	✓	✓	✓	✓	✓
	IPv6	✓	×	✓	✓	✓	✓	✓
		✓	×	✓	✓	✓	✓	✓
IMS	IPv4	N/A	✓	✓	✓	✓	✓	✓
		N/A	✓	✓	✓	✓	✓	✓
	IPv6	N/A	✓	✓	✓	✓	✓	✓
		N/A	✓	✓	✓	✓	✓	✓

Figure 30-65: SIP/RTP UDP/TCP Support

Support for Content Masking

SIP Common Presence and Instant Messaging (CPIM) content masking is supported, but only when the SIP transport is UDP.

The SIP method, MESSAGE, carried over UDP, might contain user-friendly, readable text messages. Use masking to replace these messages with x's, so they cannot be read.

NOTE: SIP/RTP correlation cannot mask text messages with a content type other than message/CPIM", such as plain text.

Behaviors of Some SIP Methods

The following are behaviors for some particular SIP methods:

- The SIP method, REGISTER, might not contain a user part. When there is no user part, it will be treated as a parse error.
- The SIP method, OPTIONS, (and response messages) might not contain a user part. When there is no user part, it will be treated as a parse error.

NOTE: SIP TCP packets with parse errors are not sent to collector. SIP TCP packets will be sent to the tool and incremented as parse errors in the session table stats.

SIP Whitelisting in a Cluster

The whitelist (all whitelist files) reside on the master node of the cluster. The non-master nodes receive a copy of the whitelist from the master. Updates to the whitelist are synchronized from the master to the non-master nodes. If a non-master node leaves the cluster and rejoins, its whitelist will be resynchronized.

Use the cluster master preference command to specify the highest preference for the master node, the second highest preference for the standby node, and lower preferences for the normal nodes in the cluster.

If there are GigaVUE TA Series nodes in the cluster, they will not receive a copy of the whitelist.

Not Supported by SIP/RTP Correlation

The following list is not currently supported by SIP/RTP correlation:

- encryption
- filtering based on Codecs
- SRVCC
- roaming
- SIP-I/SIP-T
- forwarding of emergency calls to specific tools
- engine grouping. Only one engine is used for SIP/RTP correlation.

NOTE: SIP/RTP correlation and GTP correlation are not supported on the same GigaSMART engine port.

SIP/RTP Load Balancing Example

This is a load balancing configuration example of SIP/RTP.

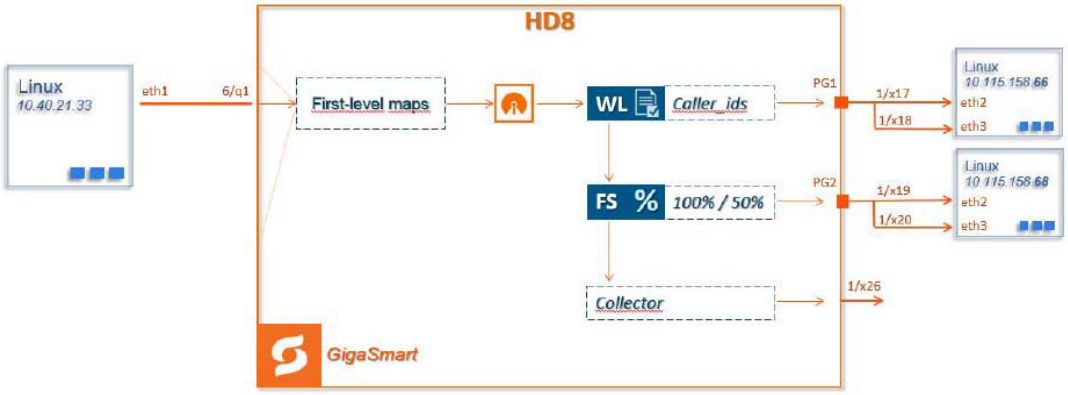

Figure 30-66: SIP/RTP UDP/TCP Support

SIP/RTP Minimum Configuration Example

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to *Access GigaSMART from GigaVUE-FM* on page 754 for details.

Configure ports

To configure a GigaSMART group and associate it with a GigaSMART engine port do the following.

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.

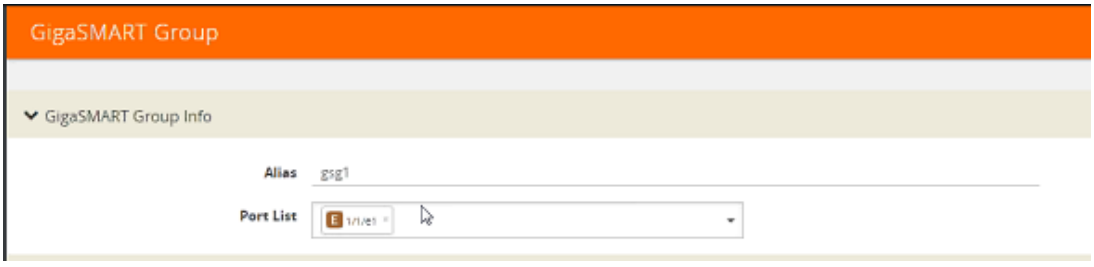

Figure 30-67: GigaSMART Group Port Info

4. Click **Save**.
5. Scroll drown the page to select **SIP Port** parameters.
6. Type parameters for **SIP Port and RTP Port.**
7. Enter parameters for **SIP Session Timeout** and the **SIP TCP Idle Timeout**.

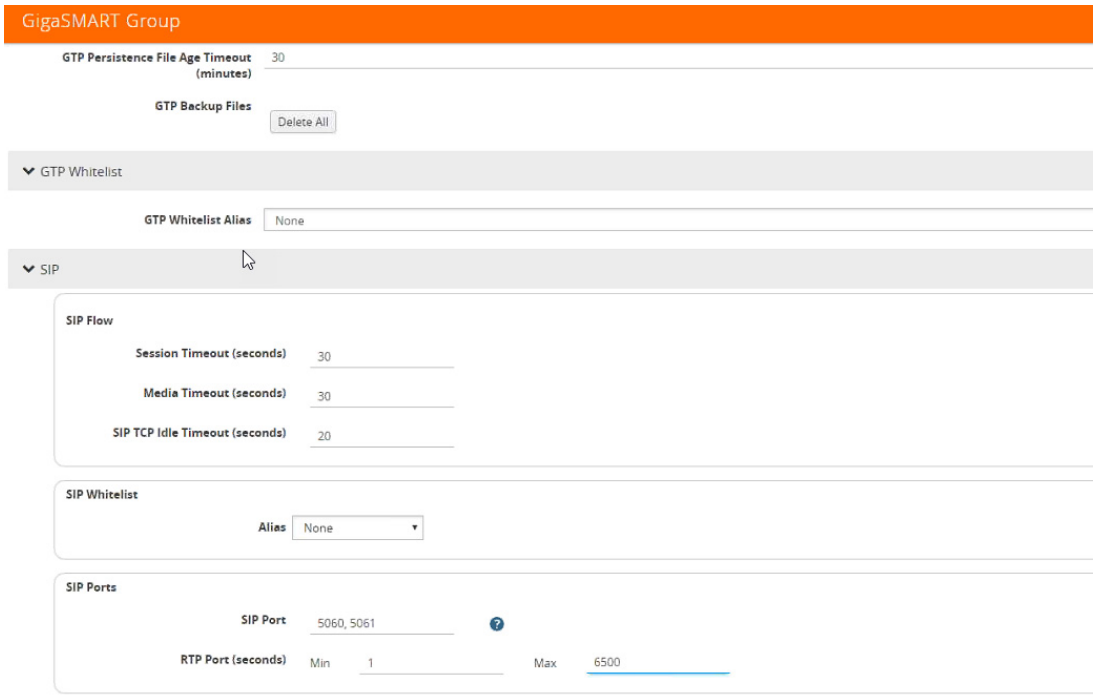

Figure 30-68: SIP Port parameters

8. Click **OK**.

Create Virtual Ports

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to *Access GigaSMART from GigaVUE-FM* on page 754 for details.

To create virtual ports, do the following:

1. From the device view, select **GigaSMART** > **Virtual Ports**.
2. Click **New**.

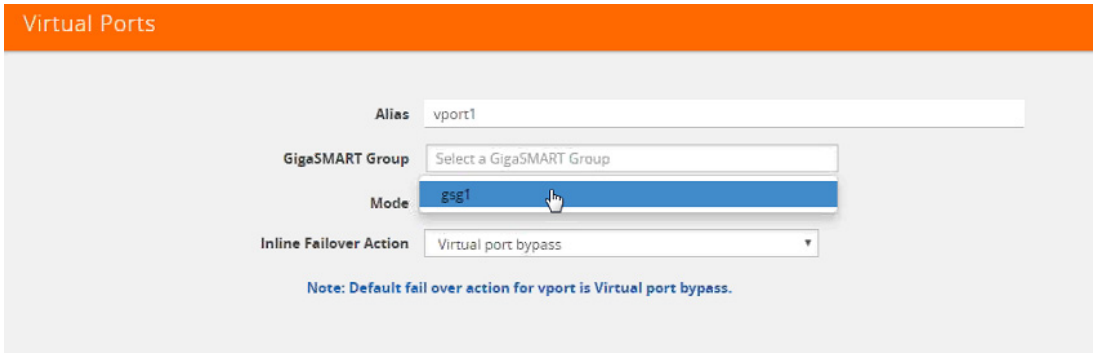

Figure 30-69: Virtual Ports

3. Enter an **alias** in the Alias field to identify the virtual port.

4. In the GigaSMART Groups field, select the GigaSMART Group configured in Step 1: of Configure a GigaSMART Group.
5. Click **Save**.

Configure GigaSMART Operation

Define a GigaSMART operation to enable SIP Flow Sampling. If combining Flow Sampling with Load Balancing GSOPs, make sure that you select both operations when creating the GigaSMART Operation.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to *Access GigaSMART from GigaVUE-FM* on page 754 for details.

To configure the GigaSMART Operation, do the following:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP).**
2. Click **New**. On the GigaSMART Operations page, do the following:

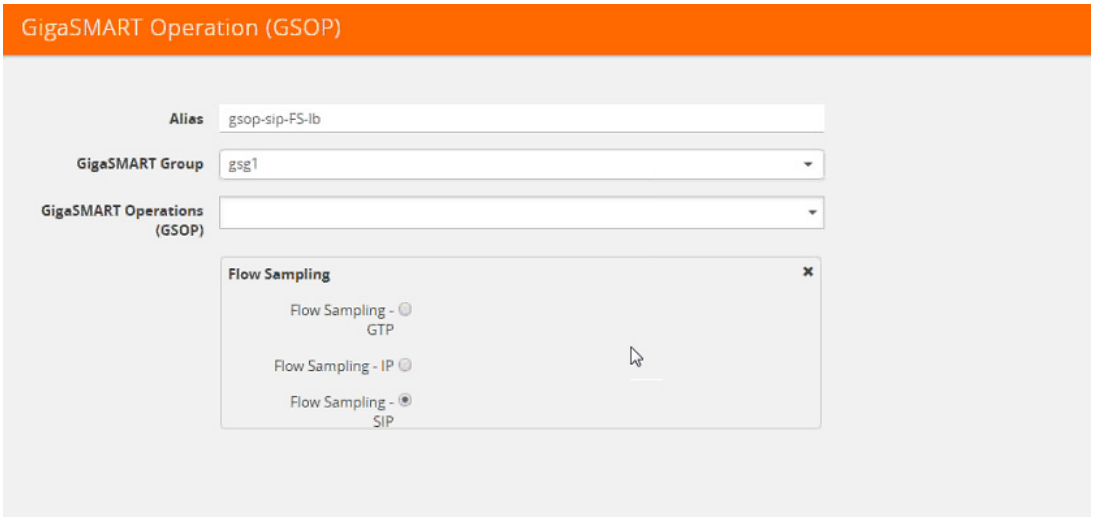

Figure 30-70: GigaSMART Operations page

3. In the Alias field, enter an **alias** to help identify this GSOP.
4. In the **GigaSMART Groups field**, select the **GigaSMART group** configured in Step 1: Configure a GigaSMART Group.
5. In the GigaSMART Operations (GSOP) field, select **Flow Sampling for SIP**.
6. Using the GSOP drop down list, select **Load Balancing** as the next GSOP operation.

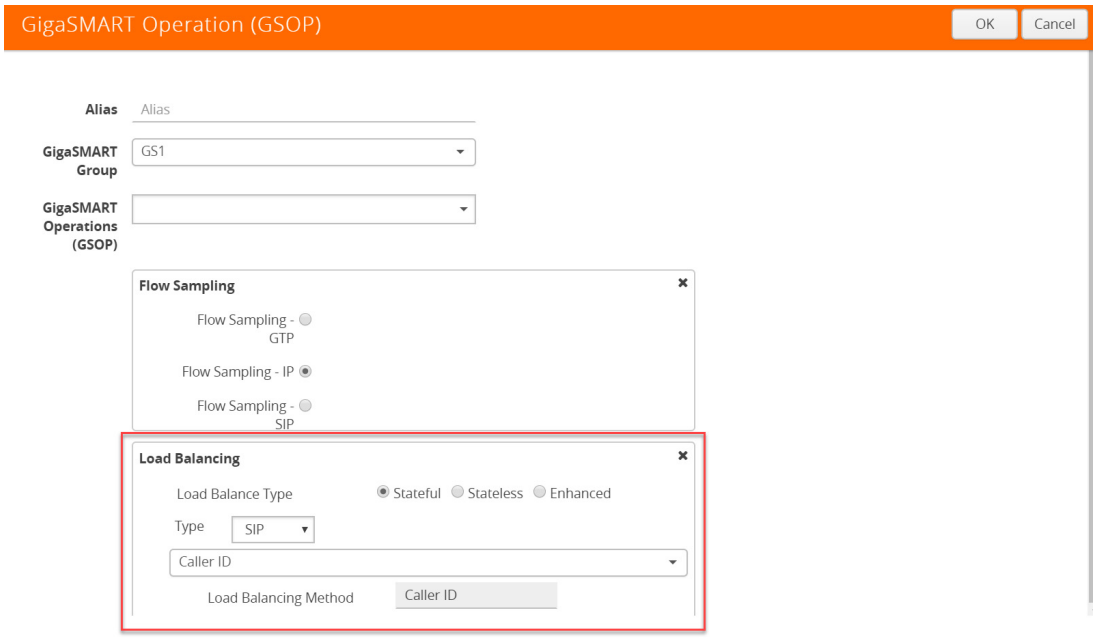

Figure 30-71: GigaSMART Operations - Load Balancing

Options:

- Stateful
- Stateless
- Enhanced

7. Select **Stateful**.
8. For Type, select **SIP** as the stateful application within a group of GigaSMART operations.
9. Select **Caller ID** as the Load Balancing Method.
10. Click **OK**.

Create first level map.

1. Select **Maps > Maps > Maps**.
2. Click **New**.
3. Type **map-level1** in the Alias field.
4. Select **First Level** for Type and **By Rule** for Subtype.
5. Select **port 1/1/x1** for the Source.

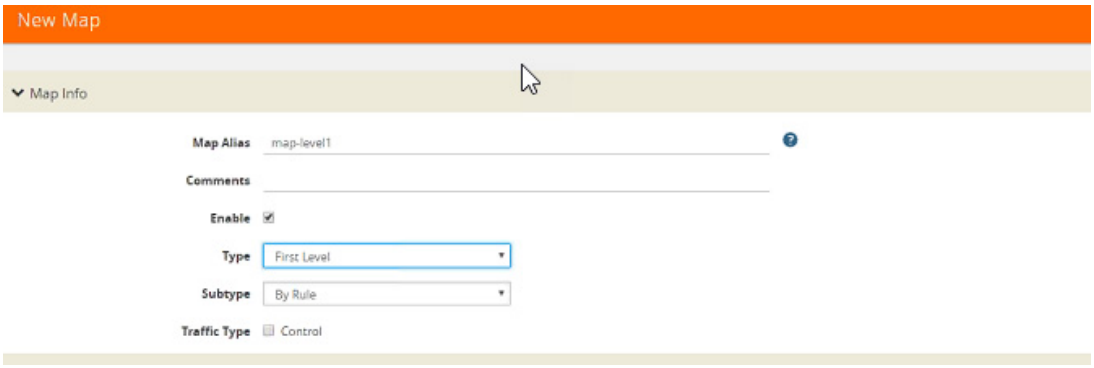

Figure 30-72: Create New Map

6. Select **virtual port vport1** for the Destination.
7. Click **Add a Rule** to add Rule 1
8. Click **Save**.

Create second level map for SIP Flow Sampling.

1. Select **Maps > Maps > Maps.**
2. Click **New**.
3. Type **an alias** in the Alias field.
4. Select Second Level for Type and **Flow Sample SIP** for Subtype.
5. Select virtual port **vport1** for the Source.
6. Select **port group** for the Destination.
7. Select **the group** from the GSOP list.
8. Click **Add a Rule**.
9. Select **SIP** for the condition.

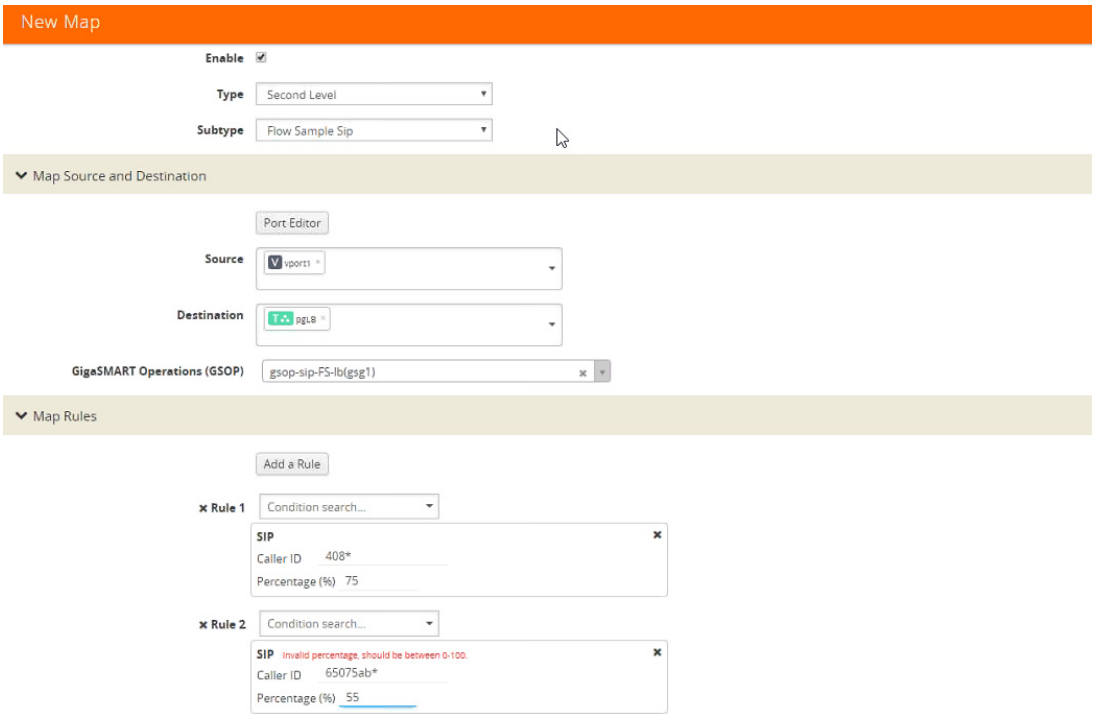

Figure 30-73: Create Second Level Map

10. Enter **408*** for Caller ID.
11. Enter a **percentage** for amount the traffic you want to be affected by SIP flow sampling.
12. Click **Add a Rule**.
13. Select **SIP** for the condition
14. Enter **6501234*** for Caller ID.
15. Enter a **percentage** for amount the traffic you want to be affected by SIP flow sampling.
16. Click **Save**.

Create the SIP whitelist

1. From the device view, select **GigaSMART > Whitelist.**
2. Click **New**.
3. Type an **alias** in the Alias field.
4. From the GigaSMART Groups drop-down list, select a .the GigaSMART group.
5. Load whitelist files from a specified location to populate the SIP whitelist.
   a. On the SIP Whitelist page, select **Bulk Upload**.
   b. Select **Bulk Entry Operation** for **Upload Type**
   c. Select Upload from URL from the Bulk Upload Type list.
   d. Enter the **URL** in the **Enter Remote URL field**.
6. Click **Save**.

Associate GigaSMART group to SIP whitelist

1. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups.**
2. Select the **GigaSMART Group** you previously created and click **Edit**.
3. From the **GigaSMART Groups drop-down list,** select a GigaSMART group previously created.
4. Under **SIP Whitelist**, click on the **SIP Whitelist Alias field** and select the **alias** previously created from the available list.
5. Click **Save**.

Configure GigaSMART operation for SIP whitelisting

1. From the device view, select **GigaSMART > GigaSMART Operations > GigaSMART Operation.**
2. Click **New**.
3. Type an **alias** in the Alias field.
4. Select the GigaSMART group created in task 1.
5. From the GigaSMART Operations (GSOP) drop-down list, select the following:
   - **SIP Whitelist** and select **Enabled.**
   - **Load Balancing**.
6. For Load Balancing, do the following:
   a. Choose: **Stateful**
   b. For Type select: **SIP**

7. Click **Save**.

Configure first level map

1. Select **Maps > Maps > Maps**.
2. Click **New**.
3. Configure the map:
   a. Enter an **Alias**
   b. Type: **First Level,** Sub Type: **By Rule**
   c. Source: **1/1/g2**
   d. Destination: **vport**
4. Click **Add a Rule.**
5. Click **Save**.

Create another second level map for SIP flow whitelist

1. Select **Maps > Maps > Maps.**
2. Click **New**.
3. Configure the map:

a. Alias: **alias name**

b. Type: **Second Level**, Sub Type: **Flow Whitelist**

c. Source: **vport1**

d. Destination: **1/2x23**

e. Select from the **GSOP** list.

4. Click **Save**.

Display SIP/RTP Reports

To display SIP report, do the following:

1. From the device view, select GigaSMART**> GigaSMART Groups> Report**.
2. Select Group Type: **Flow SIP.**
3. From the device view, select GigaSMART Group: **gsg1**
4. Specify **Caller ID Pattern.**
5. Select Any. This return any pattern specified in the Caller ID Pattern field.

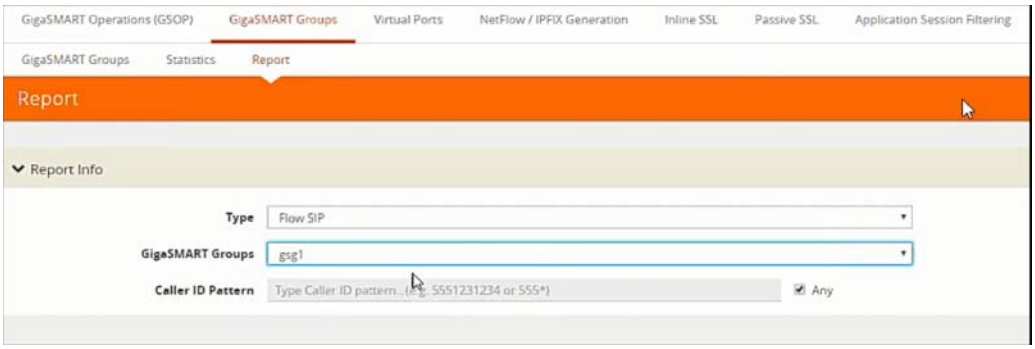

Figure 30-74: Generate SIP Report

6. Click the Generate button. The SIP Messages Report displays.

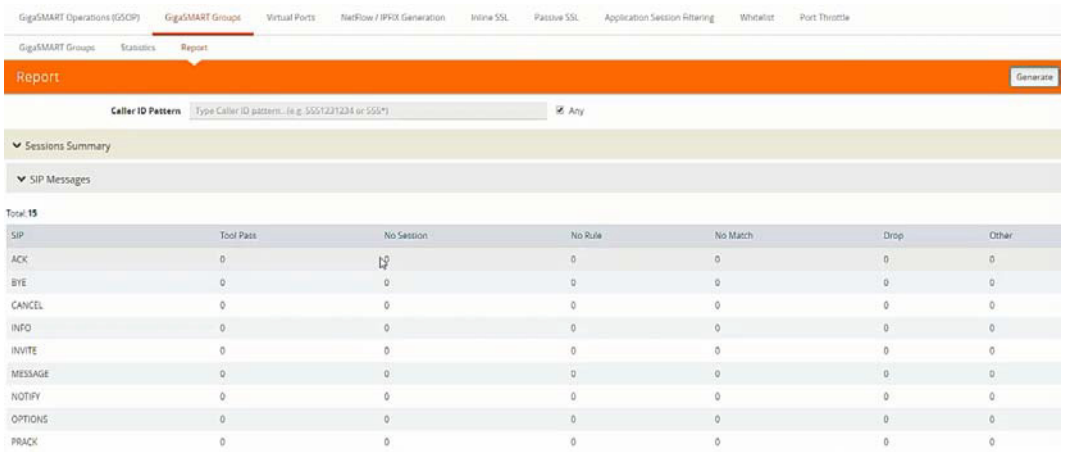

Figure 30-75: SIP Report Page

Display SIP Map Statistics

Map Statistics displays the counts sessions that matched a particular map.

For SIP Flow Sample maps, the counters show how many sessions matched the Caller-ID rule and were either accepted or rejected based on the sampling percentage.

For SIP Whitelist maps, the counter shows how many total entries are in the Whitelist and how many sessions matched those entries. .

To display SIP Map Statistics, do the following:

- Select **Maps> Statistics.** The Statistics page displays a count of the rules that actually matched in a map.

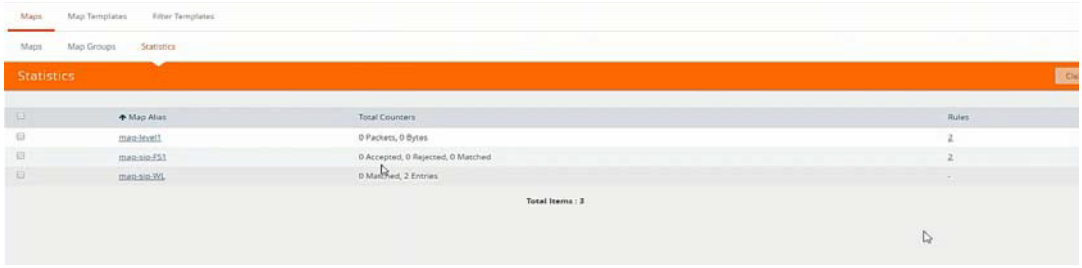

Figure 30-76: SIP Map Statistics

SIP/RTP Support for Tool Throttling

GigaSMART supports sampling/scaling based on fixed percentages, which remain in effect at all times, regardless of the tool port utilization. However, tool utilization may not be as efficient as not using fixed percentage for every use case. Staring with GigaSMART version 5.4 support for throttling sessions based on the traffic (pps) reaching a tool port is available. This feature helps avoid packets being drops during peak times by allowing users to adjust throttling start and stop levels.

The illustration below is a configuration and Intra Flow for Tool Port Throttling.

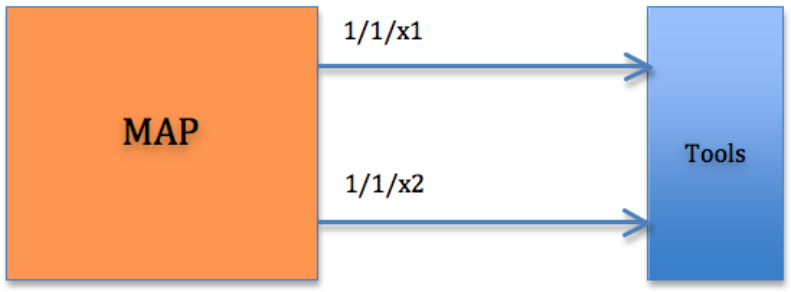

Admission Control:

Each SIP session comes with RTP streams and each of the RTP stream uses a specific codec for information transfer. We can use this codec information to our advantage and do predictive analysis on how much pps would be generated for a given SIP session.

Based on the outcome of the pps for a given RTP stream codec, admission control module will check this value against the cumulative packet throughput on the destination tool-port to decide if the session will be Accepted or Rejected.

Example: Tool port 1/1/x1 is configured with a threshold of 3k pps.

Time	Port	Session	Codec (pps)	Cumulative pps	Throttle pps	Accepted
t0	1/1/x1	0		0	3000	
t1	1/1/x1	1	500	500	3000	Yes
t2	1/1/x1	2	1500	2000	3000	Yes
t3	1/1/x1	3	800	2800	3000	Yes
t4	1/1/x1	4	500	3300	3000	No

In software version 5.4 Tool port throttling applies only to SIP sessions for audio and only load balanced ports are supported in tool port throttling. Use case where there are tapping multiple interfaces using multiple engines, one SIP session can be throttled in one engine and not in another.

GigaSMART Diameter S6a Correlation

Required Licenses: Diameter correlation license is available as a Base (100,000 users) and Max license (maximum number of users).

Diameter is an application layer protocol used in the GTP Mobile Infrastructure for authentication, authorization and accounting. The Diameter protocol uses SCTP protocol (and also TCP protocol). Thousands of diameter transactions occur every second on a relatively low network segment. Therefore, it is important to load balance the diameter transactions to the diameter processing probes (that have a finite capacity).

The GigaSMART Diameter Correlation application load balances the Diameter traffic to the LTE sessions of each of the subscriber such that the probes receive the S6a traffic for the given IMSI/MSISDN.

The application performs the following functions:

- Correlates Diameter S6a message
- Provides support for FlowVUE functionality with Stateful Load balancing (based on IMSI/MSISDN)
- Provides single engine support and allows engine grouping, as well
- Provides IPv4 support

Diameter Correlation Engine

The Diameter Correlation Engine correlates the messages exchanged during the initial attach of the subscriber (except the reset messages). Each message has a mandatory field called Session-ID and User Name, and these fields are used to correlate the messages for a given subscriber.

When a message containing Diameter traffic is received, the Diameter correlation engine looks up the session in the session table for load balancing ports and sampling maps or whitelist map.

NOTE: The GigaSMART Diameter Correlation application uses dynamic memory instead of fixed memory. It can run with GTP and other compatible applications.

Diameter Whitelisting in Clusters

The Diameter whitelist contains user names (IMSI) and command codes. Whitelist entries can only be numeric. For each entry, you can specify up to 14 to 15 numeric characters.

You can manually add one entry at a time to a whitelist file, or you can upload files in.txt format. Each whitelist file can have up to 20,000 entries. One or more whitelist files can be fetched from a local directory or remote URL using HTTP or SCP.

On GigaVUE-HC1, GigaVUE-HC2 and GigaVUE HD Series nodes, the whitelist database supports 500,000 entries. On GigaVUE-HC3 nodes, the whitelist database supports 1 million entries.

The GigaSMART operation does not have any rules for whitelisting.

Diameter Flow Sampling

FlowVUE is used for session-aware (stateful) load balancing and whitelisting with sampling.

Up to 10 flow sample maps per GigaSMART group are supported. Each flowsample map can have 20 rules. Use rules to filter based on User Name. The rules support numbers and wild card entries, up to 15 characters.

Sampling is based on **User Name** field.

Configure Diameter S6a Correlation

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* on page 754 for details.

The following table summarizes the tasks required for configuring GigaSMART Diameter S6a :

Configure GigaSMART Groups	*Configure GigaSMART Group* on page 991
Create Virtual Ports	*Create Virtual Ports* on page 992
Configure GigaSMART Operations: Diameter Flow Sampling	*Configure GigaSMART Operation: Diameter Flow Sampling* on page 992
Configure GigaSMART Operations: Diameter Whitelisting	*Configure GigaSMART Operation: Diameter Whitelisting* on page 995
Display Diameter Report	*Display Diameter Reports* on page 997

Configure GigaSMART Group

To configure a GigaSMART group and associate it with a GigaSMART engine port do the following.

1. From the device view, select **GigaSMART > GigaSMART Groups.**
2. Click **New**. Configure a GigaSMART group and associate it with a GigaSMART engine port do the following.
3. Type an alias in the Alias field and enter **an engine port** in the Port List field.

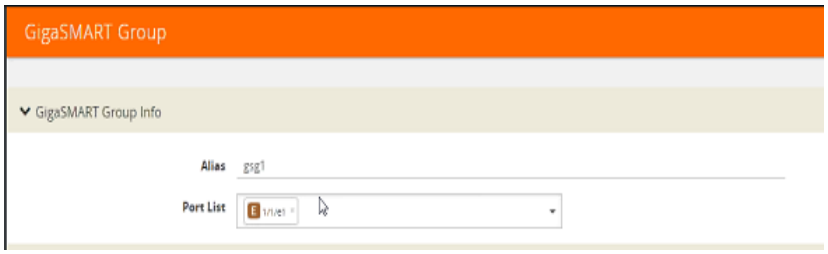

Figure 30-77: GigaSMART Group Port Info

4. Scroll drown the page to select **Diameter S6a Session** parameters.

- Session Limit (1K)
- Session Timeout (seconds)
- Diameter Packet Timeout (seconds)

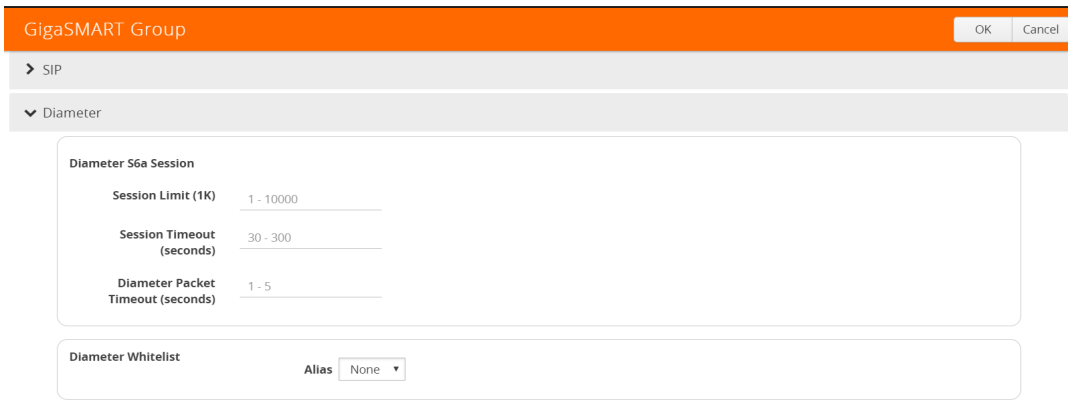

Figure 30-78: Diameter Port Parameters.

5. Click **OK**.

Create Virtual Ports

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to *Access GigaSMART from GigaVUE-FM* on page 754 for details.

To create virtual ports, do the following:

1. From the device view, select **GigaSMART > Virtual Ports**.
2. Click **New**.

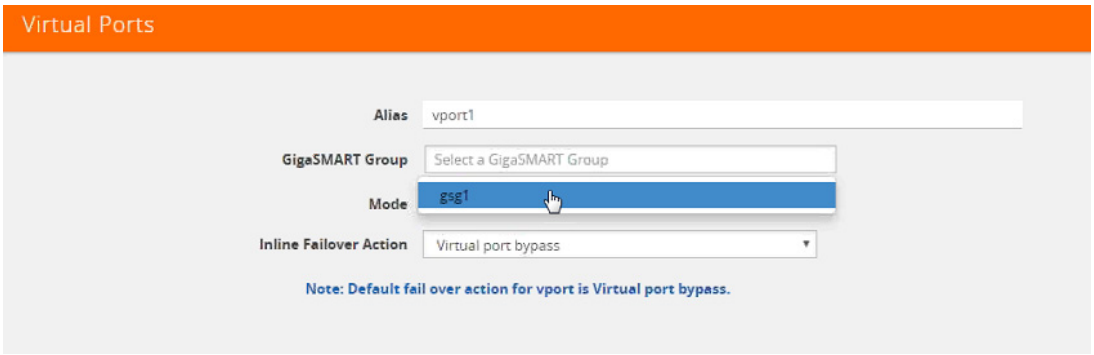

Figure 30-79: Virtual Ports

3. Enter an **alias** in the Alias field to identify the virtual port.
4. In the GigaSMART Groups field, select the GigaSMART Group configured in Step 1: of Configure a GigaSMART Group.
5. Click **Save**.

Configure GigaSMART Operation: Diameter Flow Sampling

Define a GigaSMART operation to enable Diameter Flow Sampling. If combining Flow Sampling with Load Balancing GSOPs, make sure that you select both operations when creating the GigaSMART Operation.

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. GigaSMART appears in the navigation pane of the device view on supported devices. Refer to *Access GigaSMART from GigaVUE-FM* on page 754 for details.

To configure the GigaSMART Operation, do the following:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP).**
2. Click **New**. On the GigaSMART Operations page, do the following:

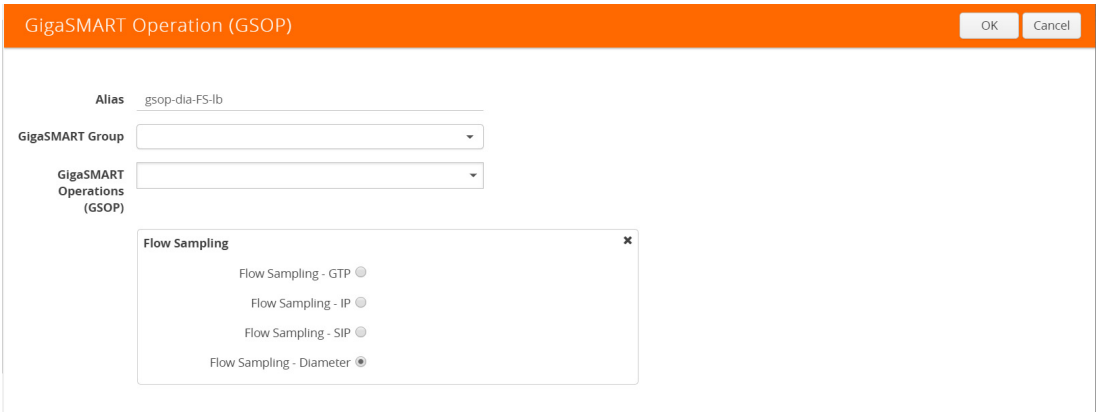

Figure 30-80: GigaSMART Operations page

3. In the Alias field, enter an **alias** to help identify this GSOP.
4. In the **GigaSMART Groups field**, select the **GigaSMART group** configured in Step 1: Configure a GigaSMART Group.
5. In the GigaSMART Operations (GSOP) field, select **Flow Sampling - Diameter**.
6. Using the GSOP drop down list, select **Load Balancing** as the next GSOP operation.

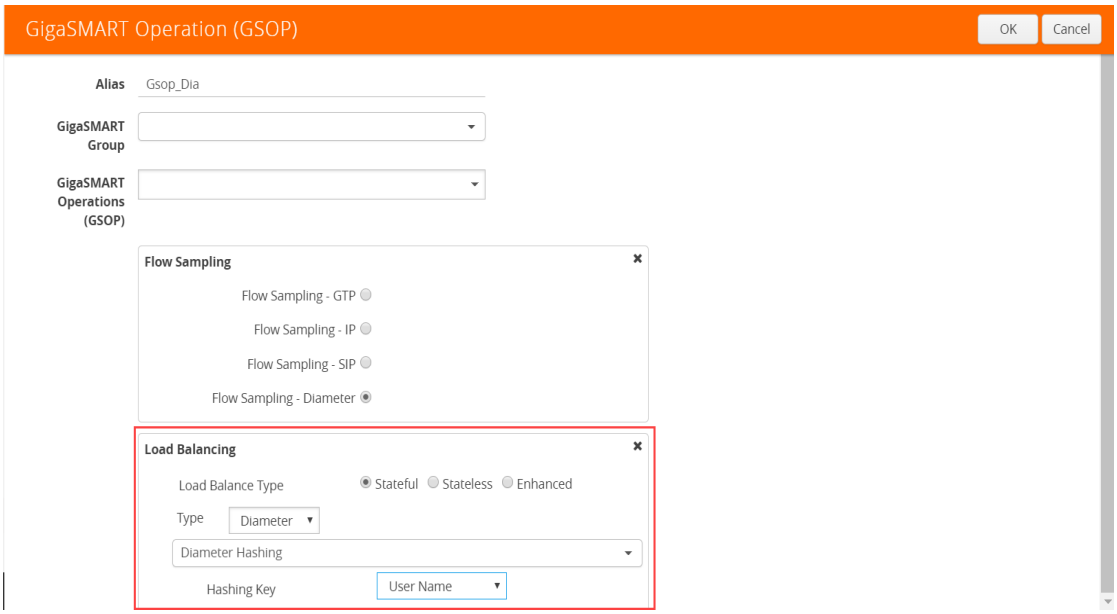

Figure 30-81: GigaSMART Operations - Load Balancing

Options:

- Stateful
- Stateless
- Enhanced

7. Select **Stateful**.

NOTE: Only Stateful load balancing option is supported for Diameter S6a Correlation.

8. For Type, select **Diameter** as the stateful application within a group of GigaSMART operations.
9. Select **Diameter Hashing** as the Load Balancing Method.
10. Select the required **Hashing Key**. Options are:

- User Name
- Command Code

11. Click **OK**.

Create first level map

1. Select **Maps > Maps > Maps**.
2. Click **New**.
3. Type **map-level1** in the Alias field.
4. Select **First Level** for Type and **By Rule** for Subtype.
5. Select **port 1/1/x1** for the Source.

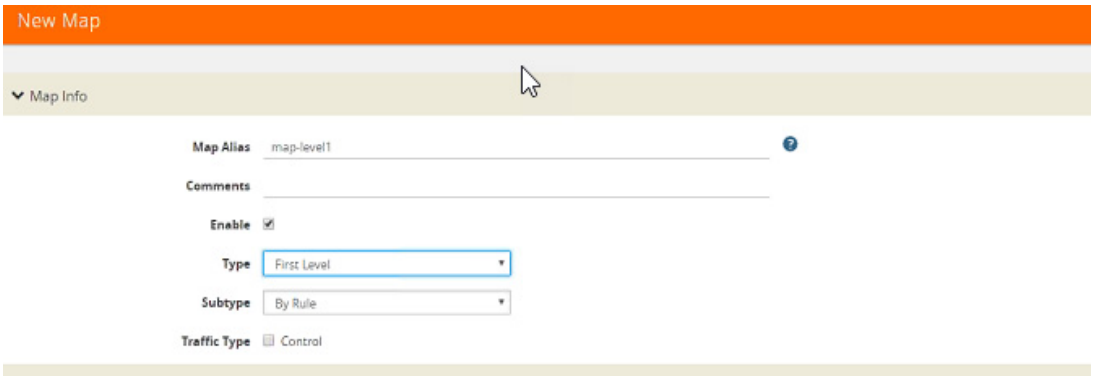

Figure 30-82: Create New Map

6. Select **virtual port vport1** for the Destination.
7. Click **Add a Rule** to add Rule 1
8. Click **Save**.

Create second level map for Diameter Flow Sampling

1. Select **Maps > Maps > Maps.**
2. Click **New**.
3. Type an **alias** in the Alias field.

4. Select Second Level for Type and **Flow Sample Diameter** for Subtype.
5. Select virtual port **vport1** for the Source.
6. Select **port group** for the Destination.
7. Select from the GSOP list.
8. Click **Add a Rule**.
9. Select **Diameter** for the condition.

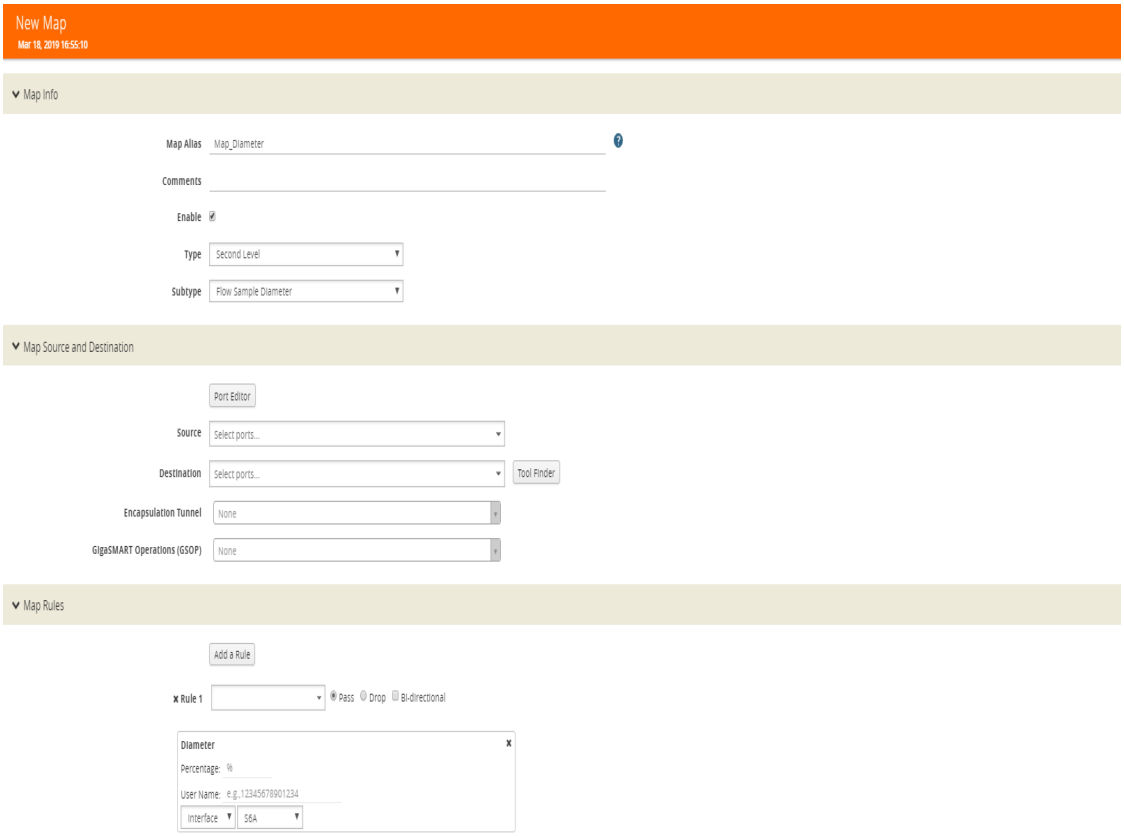

Figure 30-83: Create Second Level Map

10. Enter the following details:

- **Percentage** for the amount of traffic you want to be affected by Diameter flow sampling
- User Name
- Interface

11. Click **Save**.

Configure GigaSMART Operation: Diameter Whitelisting

1. From the device view, select **GigaSMART > Whitelist.**
2. Click **New**.
3. Type an **alias** in the Alias field.
4. Select the Whitelist type as **Diameter**.

5. You can either upload the details manually (you must enter the Diameter Entry) or load whitelist files from a specified location to populate the Diameter whitelist. You can either:

- **Upload from URL**: Enter the URL in the **Remote URL** field and the Password required to access the URL.
- **Upload a File**: Choose the location of the file name

6. Click **OK**.

Associate GigaSMART group to Diameter whitelist

1. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups.**
2. Select the **GigaSMART Group** you previously created and click **Edit**.
3. Under **Diameter Whitelist**, select the **alias** previously created from the available list.
4. Click **OK**.

Configure GigaSMART operation for Diameter whitelisting

1. From the device view, select **GigaSMART > GigaSMART Operations > GigaSMART Operation.**
2. Click **New**.
3. Type an **alias** in the Alias field.
4. Select the GigaSMART group created in task 1.
5. From the GigaSMART Operations (GSOP) drop-down list, select the following:
   - **Whitelist** and select **Diameter**.
   - **Load Balancing**.
6. For Load Balancing, do the following:
   a. Choose: **Stateful**
   b. For Type select: **Diameter**
   c. Select Diameter Hashing and enter the following Hashing Key parameters:
   - User Name
   - Command Code

7. Click **OK**.

Configure first level map

1. Select **Maps > Maps > Maps**.
2. Click **New**.
3. Configure the map:
   a. Enter an **Alias**
   b. Type: **First Level,** Sub Type: **By Rule**
   c. Source: **1/1/g2**

d. Destination: **vport**

4. Click **Add a Rule.**
5. Click **Save**.

Create another second level map for Diameter flow whitelist

1. Select **Maps > Maps > Maps.**
2. Click **New**.
3. Configure the map:
   a. Alias: **alias name**
   b. Type: **Second Level**, Sub Type: **Flow Whitelist Diameter**
   c. Source: **vport1**
   d. Destination: **Port group**
   e. Select from the **GSOP** list.
4. Click **OK**.

Display Diameter Reports

To display Diameter report, do the following:

1. From the device view, select GigaSMART**> GigaSMART Groups> Report**.
2. Select Group Type: **Flow Diameter S6a.**
3. Select the required GigaSMART Groups: **gsg1**
4. Specify the **User Name Pattern.**
5. Select Any. This returns any pattern.

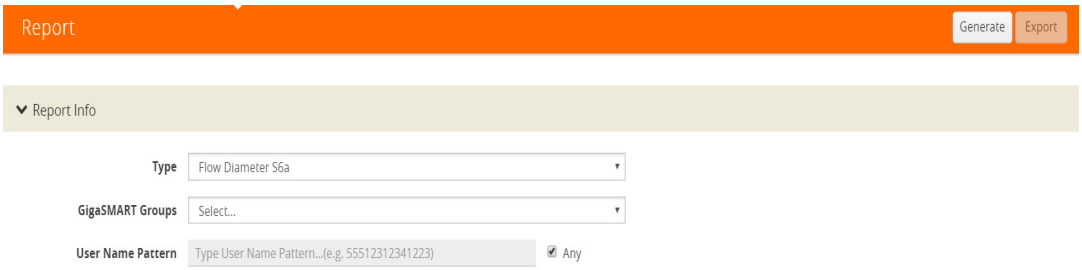

Figure 30-84: Generate Diameter Report

6. Click the **Generate** button. The Diameter Messages Report is displayed.

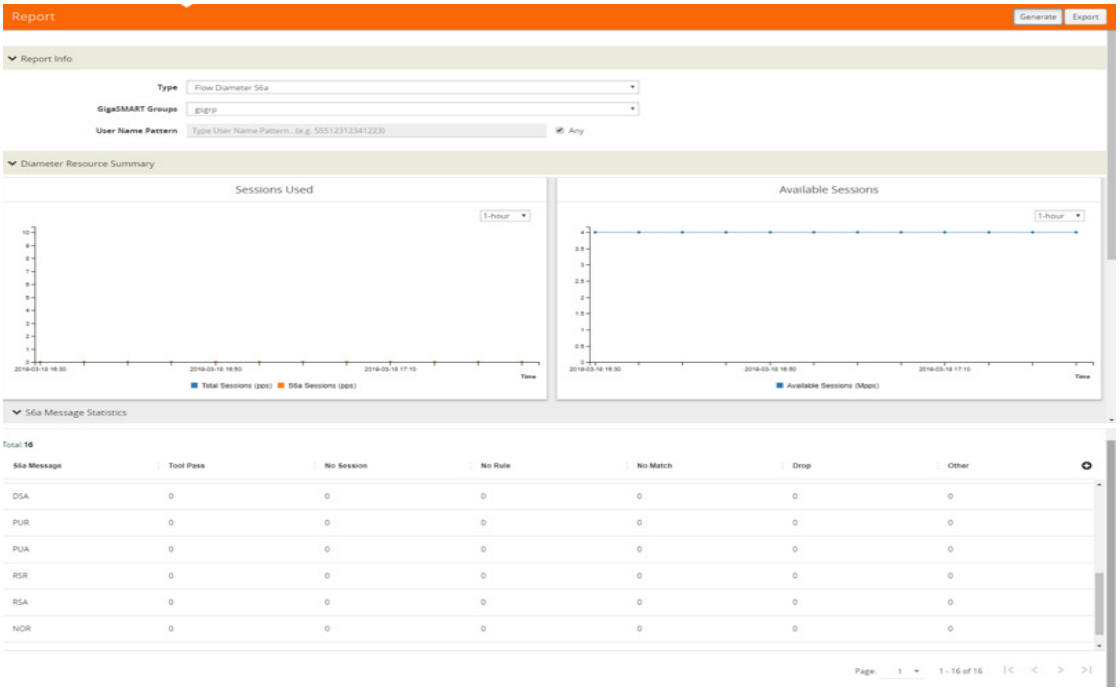

Figure 30-85: Diameter Report Page

GigaSMART FlowVUE

Required License: FlowVUE

GigaSMART FlowVUE supports the following:

- flow-aware sampling of subscriber devices to filter and forward all flows sourced from a sampled set of subscriber device IPs
- flexible sampling on subscriber IPs and IP ranges, and at specified sampling rates
- user-configurable timeouts to detect and replace inactive devices
- IP-based sampling of flows and IP-based flows encapsulated in GTP-u tunnels

FlowVUE offers subscriber IP-based flow sampling that helps carriers turn Big Data into manageable data. The GigaSMART application enables existing tools to connect to the latest high-speed pipes by providing a representative view of traffic for diagnostic coverage. GPRS Tunneling Protocol (GTP) is commonly used to carry mobile data across service provider networks and includes the control plane (GTP-c) and a user-data plane (GTP-u) traffic. FlowVUE allows for active sampling of a subscriber's device (known as a user endpoint IP or UE IP) across GTP-u tunnels. The integrity of the subscriber flows is preserved by forwarding all the flows associated with the sampled UE IP to the monitoring and analytic tools.

In contrast, traditional methods randomly sample packets without any correlation to the flows. This provides limited visibility into subscriber behavior and experience. FlowVUE intelligently reduces the amount of traffic, while keeping the integrity of the data flows intact, but at a lower speed feed within a smaller pipe.

Leveraging FlowVUE, providers can enhance quality of experience (QoE) monitoring by forwarding all of the control plane traffic to the tools infrastructure and only perform intelligent user-plane sampling (a configurable percentage of UE IPs) to get a representative view of application usage.

When combined with the advanced filtering capabilities of GigaSMART Adaptive Packet Filtering (APF), operators can further filter, replicate, and forward specific traffic flows of interest based on application ports and packet content-based payloads for all or a subset of the sampled subscribers. This further reduces the volume of traffic to the tool infrastructure. The ability to sample a subset of subscriber devices and transmit all the associated sessions of interest to the monitoring tools, reduces the amount of data while enabling Big Data throughput processing, with existing cost structures.

FlowVUE operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to *Groups of GigaSMART Engine Ports* on page 757 for details.

Configure FlowVUE

The GigaSMART parameters for configuring FlowVUE are as follows:

- flow-sampling-device-ip-ranges—Specifies the range of IP addresses that identify a valid device.
- flow-sampling-rate—Specifies how much GTP traffic from subscribers in the specified IP ranges is sampled. The values range from 5 to 95%.
- flow-sampling-timeout—Specifies after how much time a flow/device in a sampled IP range is declared idle and is no longer sampled. The values range from 1 to 60 minutes.
- flow-sampling-type—Specifies whether inner or outer IP addresses are used for FlowVUE sampling as follows:
  - device-ip—Specifies a sample subset of devices based on IP address.
  - device-ip-in-gtp—Specifies a sample subset of devices based on inner IP address in the GTP-u tunnel.

Sample of Subset of Subscribers and Sample of all Subscribers Traffic

The following example samples on GTP-u traffic where 10% of the subscribers are forwarded.

Task	Description	UI Steps
1	Configure one network and two tool type of ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** 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.

Task	Description	UI Steps
2	Configure a GigaSMART group and associate it with a GigaSMART engine port and configure sampling parameters	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type gsg1 in the Alias field. 4. Click int he Port List field and select an engine port. For example. 1/1/e1. 5. Configure the Flow Sampling parameters under GigaSMART Parameters. • Select Device IP. • Enter 1.1.1.0/255.255.255.0 in the IP Ranges field. • Enter 10 for Rate to set the flow sampling rate to 10 percent. 6. Click Save.
3	Configure the GigaSMART operation and assign it to the GigaSMART group.	1. From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations. 2. Click New. 3. Type gsfvue in the Alias field. 4. Select gsg1 from the GigaSMART Groups list. 5. Select Flow Sampling from the GigaSMART Operations (GSOP) list. 6. Select Flow Sampling - IP 7. Click Save.
4	Create an ingress (first level) map. **NOTE:** In the rule, 2152 is GTP-u traffic.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Enter to_tool in the Alias field. • Select Regular for Type. • Select By Rule for Subtype. • Select the network port for the Source. • Select a tool port for the Destination. • Select gsfvue from the GSOP list. 4. Add a rule. a. Click Add a Rule. b. Select Pass. c. Select Port Destination. d. Enter 2152 for the port value. 5. Click Save.

Sample of a Subset of Subscribers and Sample of a Subset of Traffic

FlowVUE can be used to reduce traffic to the monitoring tools. By combining FlowVUE with other GigaSMART applications such as APF, the traffic can be further reduced by filtering on specific Layer 4 application ports.

The following example samples on a subset of subscribers and forwards only the HTTP traffic related to the sampled subscriber set of devices.

Task	Description	UI Steps
1	Configure one network and two tool type of ports.	1. Select Ports > Ports > All Ports. 2. Click Quick Port Editor. 3. 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. Click int he Port List field and select an engine port. For example. 1/1/e1. 5. Configure the Flow Sampling parameters under GigaSMART Parameters. • Select Device IP. • Enter 1.1.1.0/255.255.255.0 in the IP Ranges field. • Enter 10 for Rate to set the flow sampling rate to 10 percent. 6. Click Save.
3	Configure the GigaSMART operation and assign it to the GigaSMART group. Also, configure APF.	1. From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. 2. Click New. 3. Type gsfvue_apf in the Alias field. 4. Select gsg1 from the GigaSMART Groups list. 5. Select APF from the GigaSMART Operations (GSOP) list 6. Enable APF. 7. Click Save.
4	Configure 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.

Task	Description	UI Steps
5	Create a first level map and direct 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 network port for the Source. • Select the virtual port vp1 for the Destination. • Select gsfvue from the GSOP list. **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 second level map and use the APF GigaSMART operation. APF performs filtering according to the gsrules, sending only matching traffic to the tool port.	**1.** 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 a tool port for the Destination. • Select gsfvue_apf from the GSOP list. **4.** Add a rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Port Destination. **d.** Enter 80 for the port value. **e.** Select 2 for Position. **5.** Add a rule 2. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Port Source. **d.** Enter 80 for the port value. **e.** Select 2 for Position. **6.** Click Save.

Display FlowVUE Statistics

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

Refer to *FlowVUE Statistics Definitions* on page 802 for descriptions of these statistics as well as to *GigaSMART Operations Statistics Definitions* on page 794.

GigaSMART Adaptive Packet Filtering (APF)

Required License: Adaptive Packet Filtering

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.

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. APF also provides a powerful filtering engine that identifies content (based on patterns) across any part of the data packet, including the data packet payload.

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* on page 1055.

APF operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to *Groups of GigaSMART Engine Ports* on page 757 for details.

In APF second level maps, a maximum of five (5) maps can be attached to a virtual port (vport). Each map can contain up to 25 gsrules.

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

- *Content-based Filtering* on page 1005
- *Encapsulation Awareness* on page 1009
- *Pattern Matching* on page 1010

Implement APF Through the UI

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

1. Select the **GigaSMART > GigaSMART Groups > GigaSMART Groups**, and click **New**.
2. 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.
3. From the device view, select **GigaSMART > Virtual Ports**, and then click **New**.
4. On the Virtual Ports page, enter an alias and select the GigaSMART groups created in Step1, and then click **Save**.
5. 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.

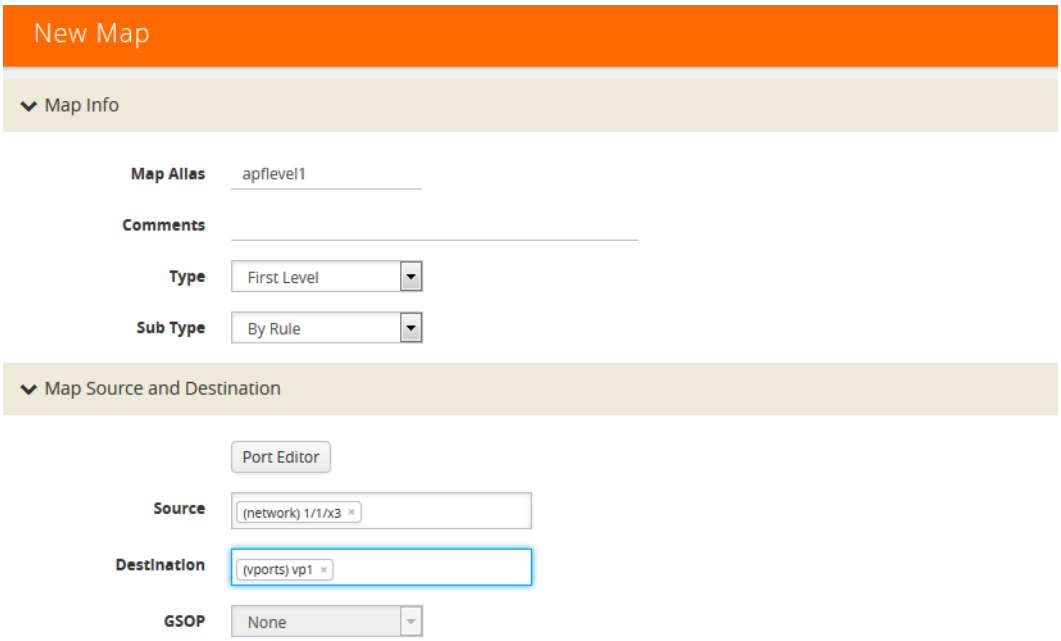

7. Create a second level map with the APF GigaSMART operation, the virtual port as the source, and a rule. The following figure shows an example.

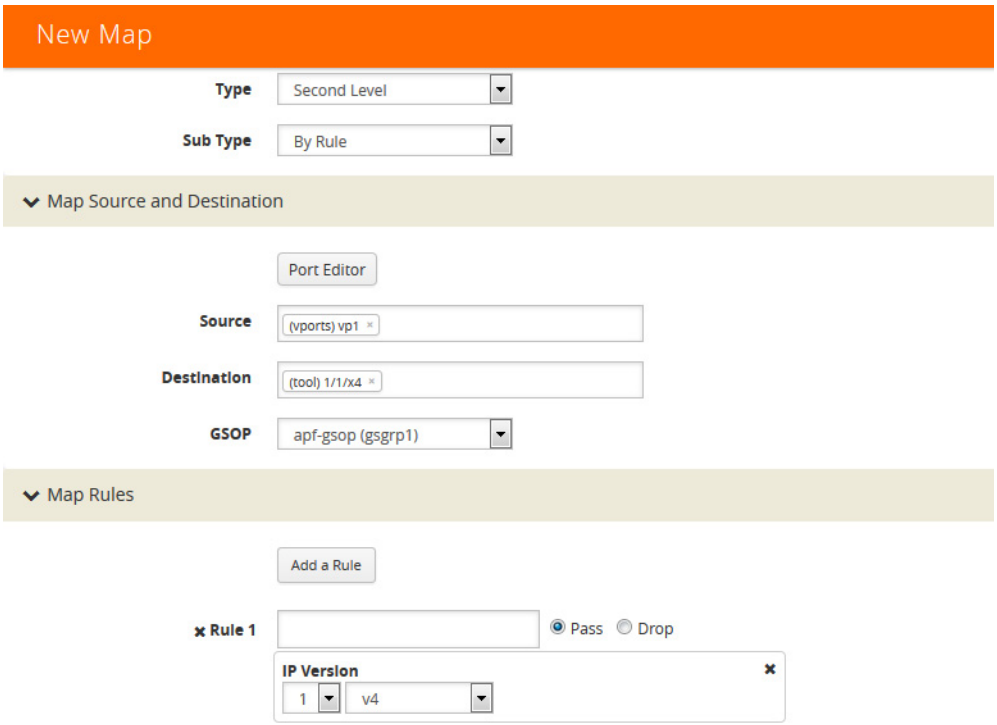

This completes the process to create an APF GigaSMART operation and corresponding rules. To learn more about the rules applicable for APF, see following sections.

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.

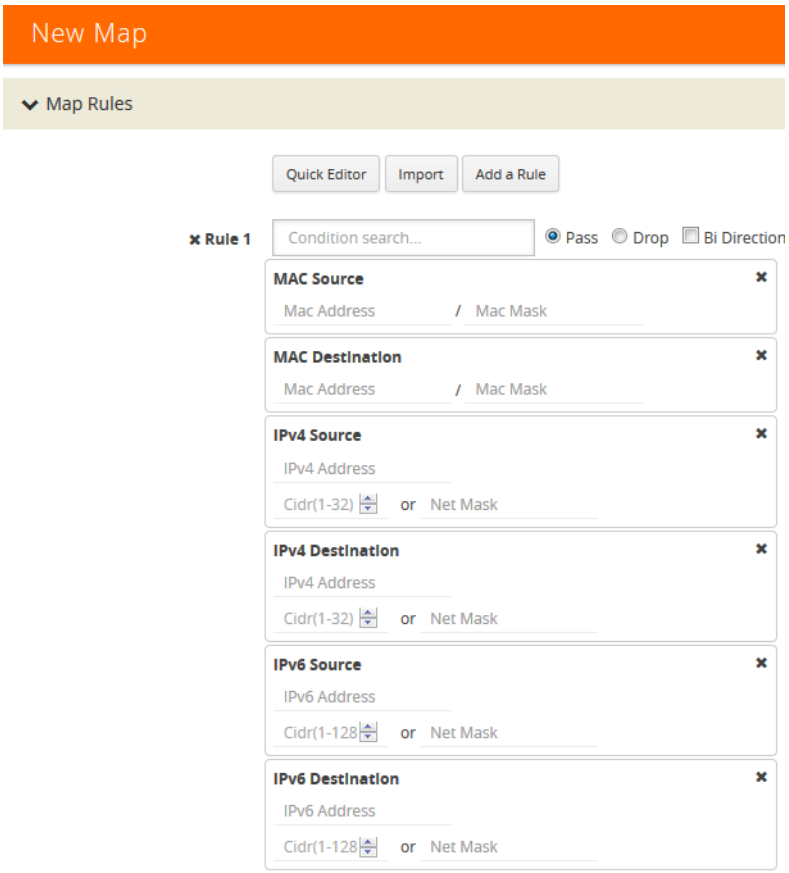

The second group of attributes has the following format:

```
<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.

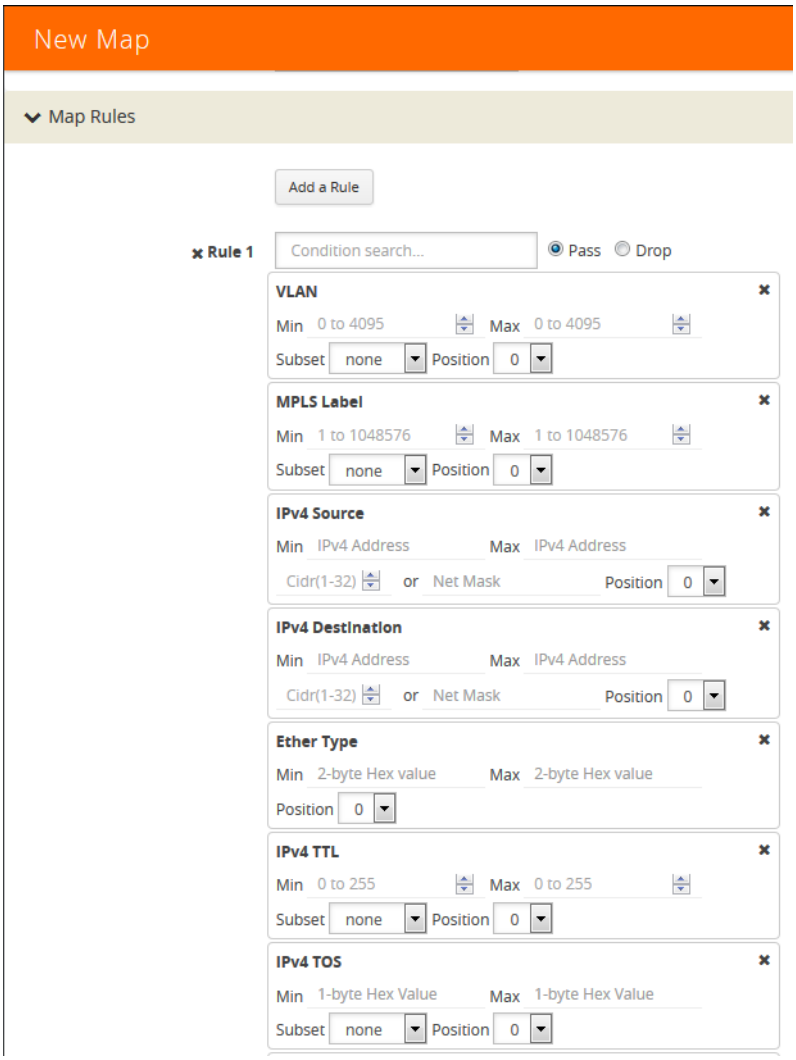

The third group of attributes has the following format: <value> <position>

```
<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.

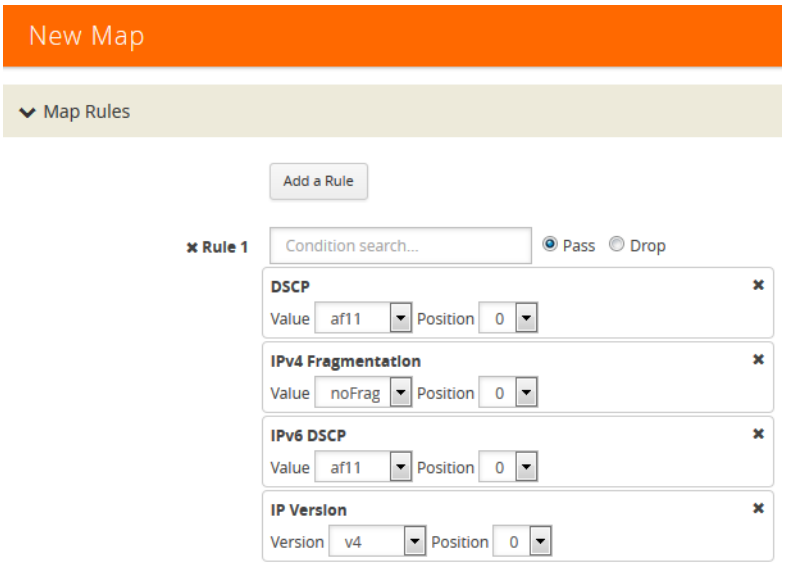

The fourth group of attributes has the following format:

```
<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.

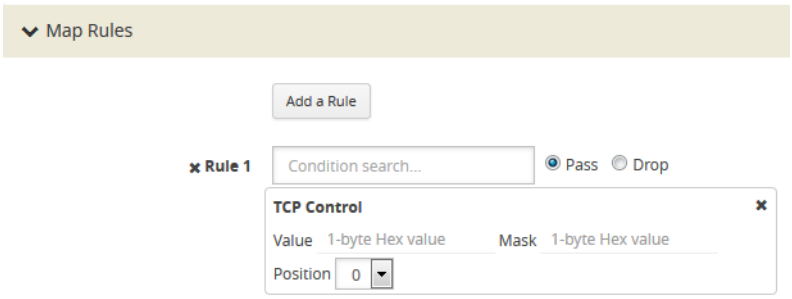

The maximum occurrences of each attribute supported are as follows:

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
Attributes in VNTag header	3

Attribute	Maximum Occurrences
Attributes in TCP header	3
IP ver	3

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.

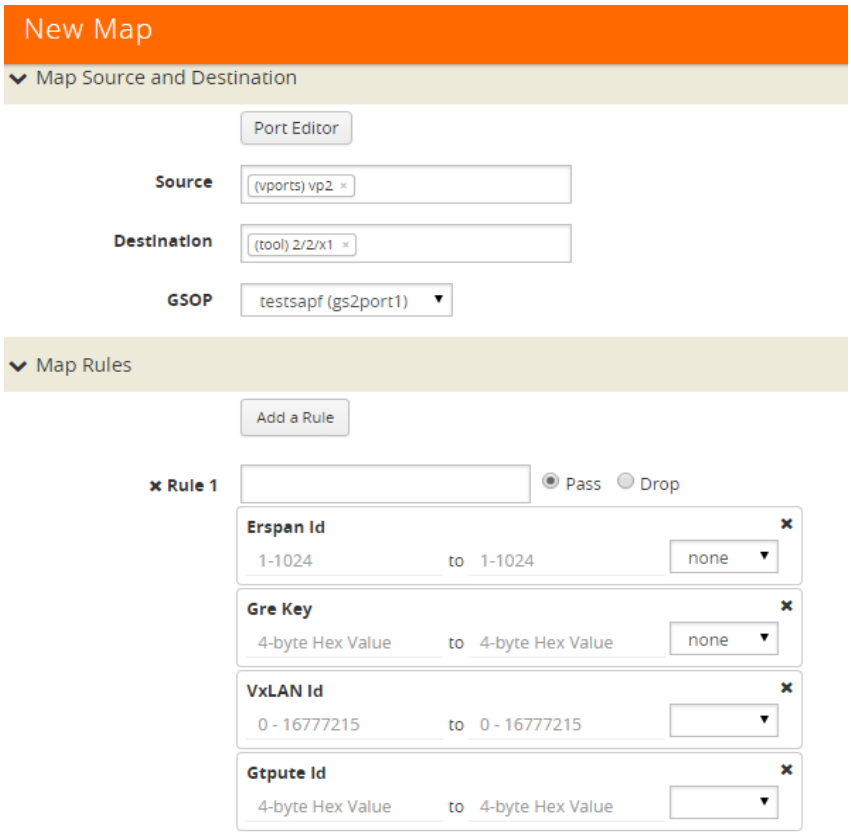

Pattern Matching

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 30-86.

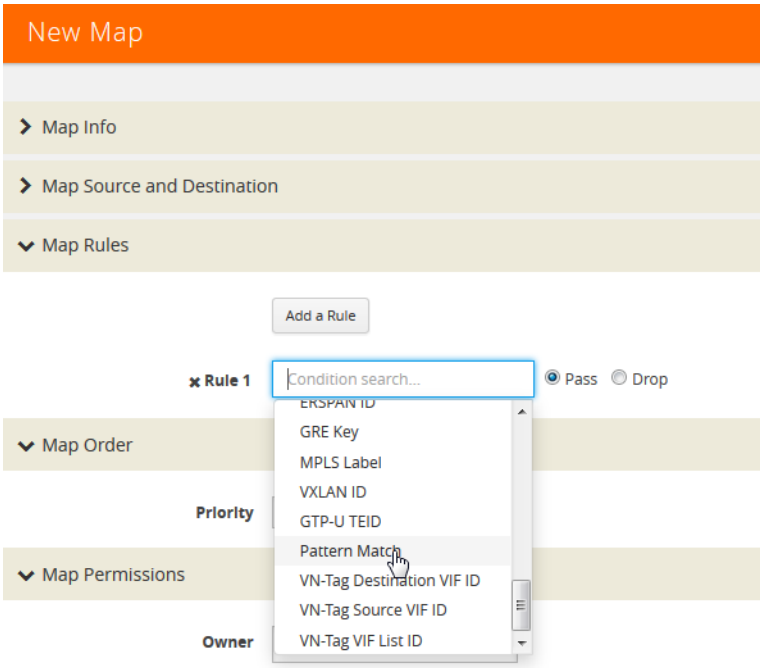

Figure 30-86: Use Pattern Match Under Maps for Pattern Matching

After selecting pattern matching for the rule, you can enter a Perl-compatible regular expression or a string to be used as a filter when pattern matching. For example to pass all packets including the string www.gigamon.com select **string** as type for the pattern match as shown in Figure 30-87.

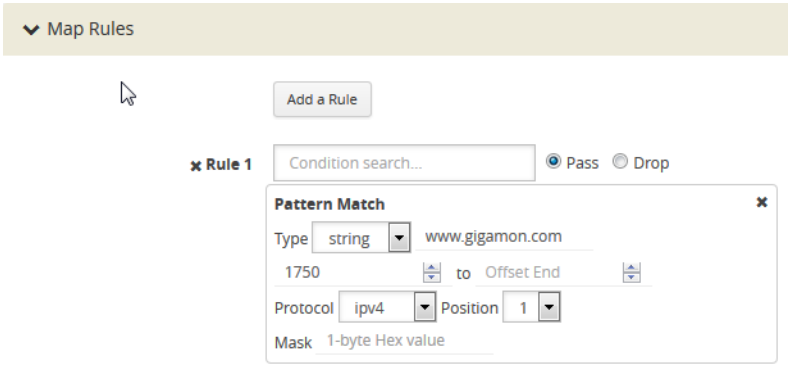

Figure 30-87: Pattern Match with Type String

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{3} as shown in Figure 30-88.

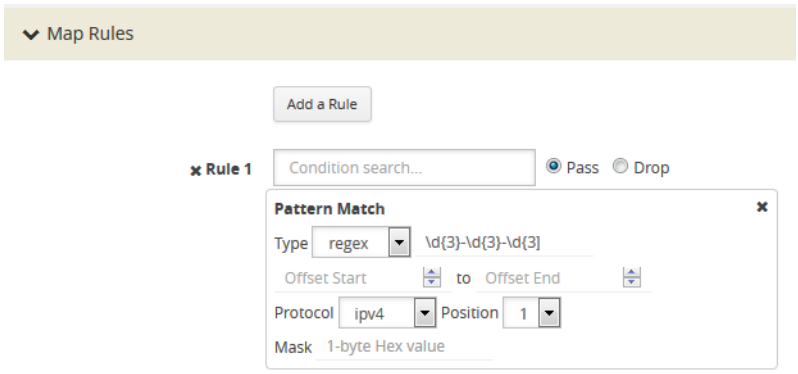

Figure 30-88: Pattern Match with Type RegEx

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 30-89.

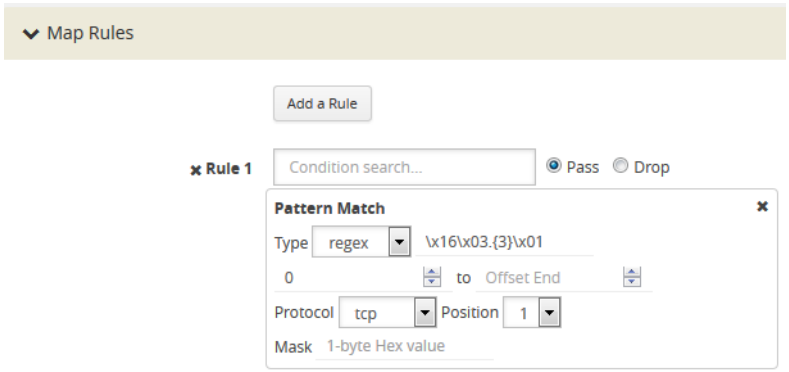

Figure 30-89: Pattern Match for SSL Client Hello Packet

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 30-90.

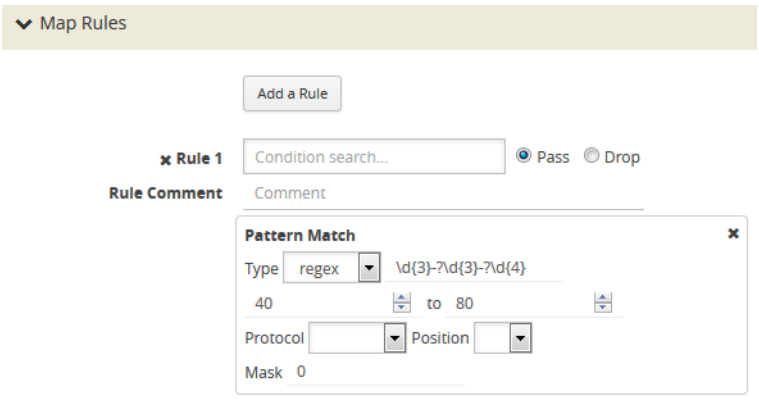

Figure 30-90: Map Rule with RegEx for Masking SSNs

Pattern Matching Hint

To optimize APF pattern matching performance in second level maps with gsrules, you can optionally use a pattern matching hint. Refer to the example in Figure 30-91.

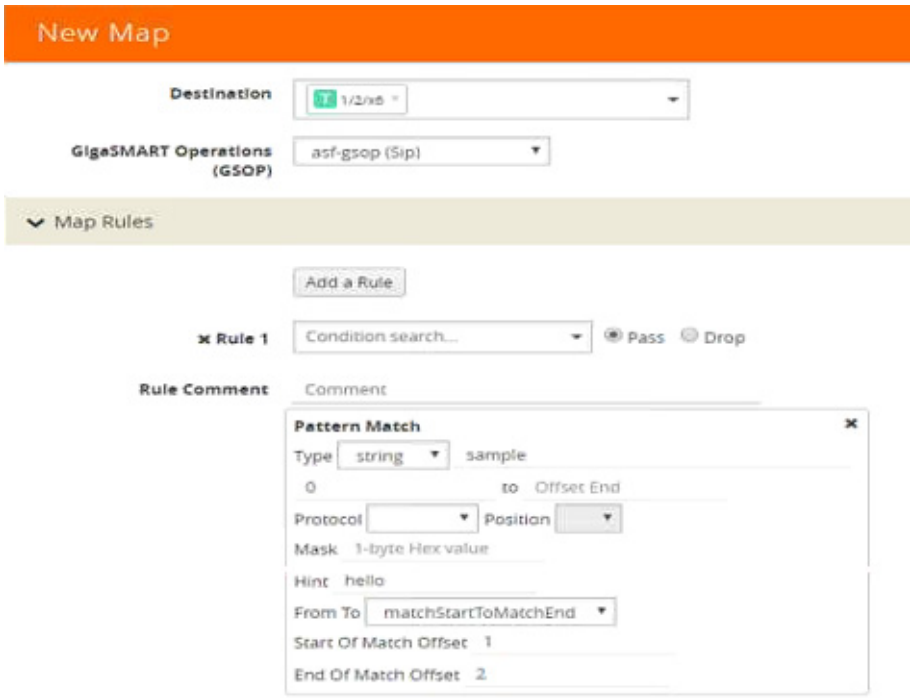

Figure 30-91: Pattern Match with Hint

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. 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.

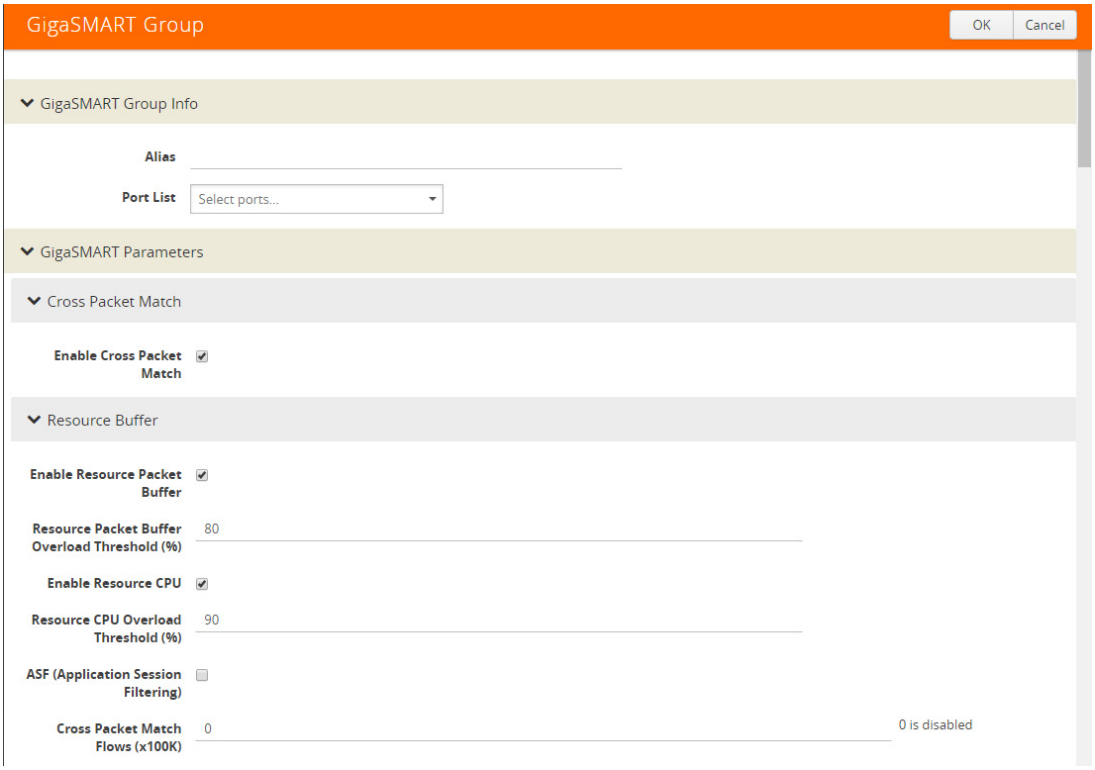

4. Click the **Enable Cross-packet Match** check box to enable.
5. Enter a range from 1 to 10 for the **Cross Packet Match Flows** parameter. Each unit is 100K bi-directional flows.
6. Click **OK**.

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.

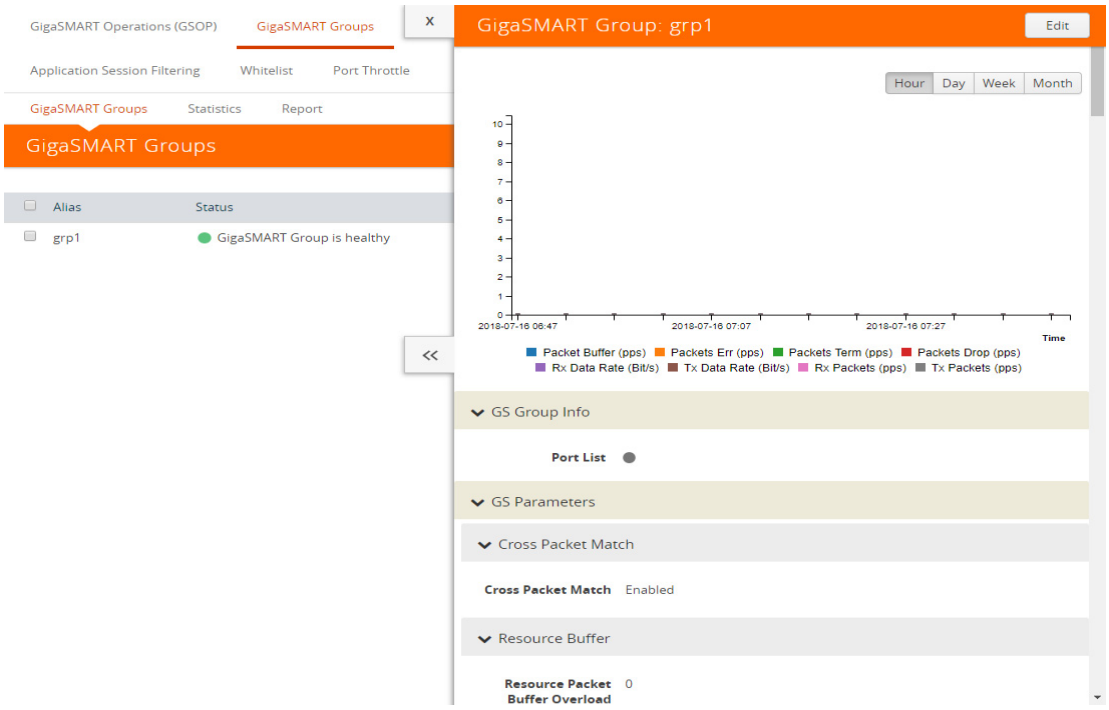

Limitations

The following constraints exist with this functionality.

- Cannot coexist with other GSOPs on same gsgroup.
- Only one second level map is allowed for each vport attached to the gsgroup.
- Disabling the feature requires a GigaSMART card reboot.

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.

APF Examples

The following are APF examples:

- *How to Identify Social Security Numbers in User-Level Transactions* on page 1017
- *How to Mask Social Security Numbers* on page 1018
- *How to Filter on Fiber Channel over Ethernet (FCOE) Traffic* on page 1020
- *Multi-Encapsulation Filtering* on page 1023
- *How to Filter on Subscriber Device IP (User-Endpoint IP or UE-IP)* on page 1025
- *How to Filter on Inner Layer 2-4 Parameters for Unrecognized Headers* on page 1028

- *GTP Tunnel ID-Based Filtering* on page 1031
- *ERSPAN Tunneling* on page 1034
- *Distribute Traffic Based on Inner IP Addresses and Inner TCP Port Values* on page 1036
- *MPLS Label Based Filtering* on page 1039
- *Combine APF with GigaSMART Operations* on page 1042
- *Conditional Header Stripping* on page 1045
- *Implement Overlapping Rules* on page 1049

How to Identify Social Security Numbers in User-Level Transactions

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

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

Task	Description	UI Steps
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. 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 5. Click Save.

How to 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.

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.

Task	Description	UI Steps
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.
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.

Task	Description	UI Steps
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 g. Enter 0 for Mask. 5. Click Save.

How to 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.

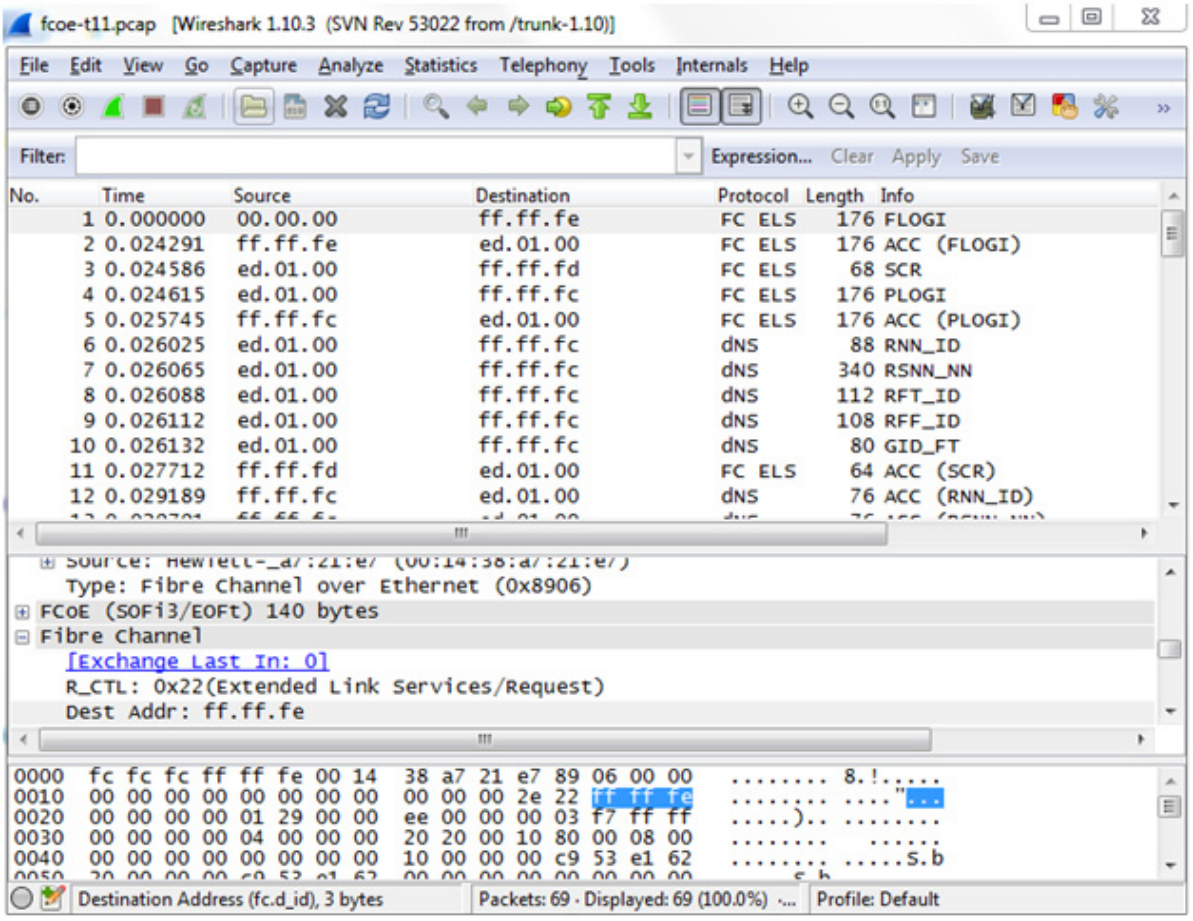

Task	Description	UI Steps
1	Configure ports.	1. Select Ports > Ports > All Ports. 2. Click Quick Port Editor. 3. Configure one network port and two tool ports. For example, select Network for port 1/1/x3. Select Tool for port 1/1/x4 and port 1/1/x1. 4. Select Enable for each port. 5. Click OK. 6. Close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

Task	Description	UI Steps
3	Configure the GigaSMART operation.	1. From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. 2. Click New. 3. Select gsfil from the GigaSMART Groups list. 4. Select Adaptive Packet Filtering from the GigaSMART Operations list. 5. Select Enable. 6. Click Save.
4	Create a virtual port and associate it with the GigaSMART group.	1. From the device view, select GigaSMART > Virtual ports. 2. Click New. 3. Enter gsTraffic in the Alias field. 4. Select gsgrp1 from the GigaSMART Groups list. 5. Click Save.
5	Create a first level map to forward FCOE traffic to the virtual port.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Enter to_vp in the Alias field. • Select First Level for Type. • Select By Rule for Subtype. • Select the port 1/1/x3 for the Source. • Select the virtual port vp1 for the Destination. 4. Add a rule. a. Click Add a Rule. b. Select Pass. c. Select Ether Type d. Enter 8906 in the Value field. 5. Click Save.

Task	Description	UI Steps
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.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Pattern Match. **d.** Select string for Type and enter txff\xff\xfe. **e.** Set the Offset Start to 0. **f.** Set the Offset End to 29 **g.** Enter 0 for Mask. **5.** 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 GigaVUE Visibility Platform 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.

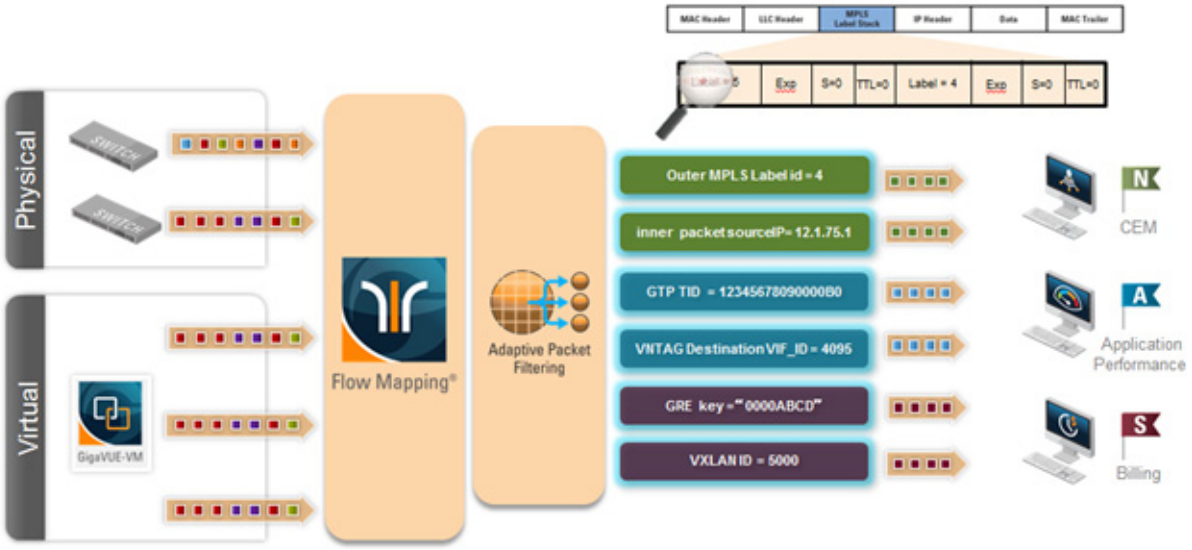

With encapsulation awareness enabled by APF, operators have multiple options to act on the packet including the flexibility to:

- 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.

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

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

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

In this example, we are:

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

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

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

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 to forward GTP-u traffic to the virtual port. **NOTE:** In the rule, 2152 is GTP-u traffic.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Enter to_vp in the Alias field. • Select First Level for Type. • Select By Rule for Subtype. • Select the port 1/1/x3 for the Source. • Select the virtual port vp1 for the Destination. 4. Add a rule. a. Click Add a Rule. b. Select Pass. c. Select Port Source d. Enter 2152 for the port value. 5. Click Save.
6	Create a first level map to forward GTP-c traffic to the tools. **NOTE:** In the rule, 2123 is GTP-c traffic.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Type to_tool in the Alias field. • Select Regular for Type. • Select By Rule for Subtype. • 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.

Task	Description	UI Steps
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. 5. 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. 6. Click Save.

Task	Description	UI Steps
8	Create another second level map to filter on source and destination IP (bi-directional).	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. • Enter map1 in the Alias field. • Select Second Level for Type. • Select By Rule for Subtype. • Select the virtual port vp1 for the Source. • Select the tool port 1/1/x4 for the Destination. • Select gsfil from the GSOP list. **4.** Add Rule 1. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select IPv4 Source. **d.** Enter 1.1.1.1 for the IPv4 Address **e.** Enter 255.255.255.255 for the Net Mask **f.** Set Position to 2. **5.** 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. **6.** Click Save.

How to 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

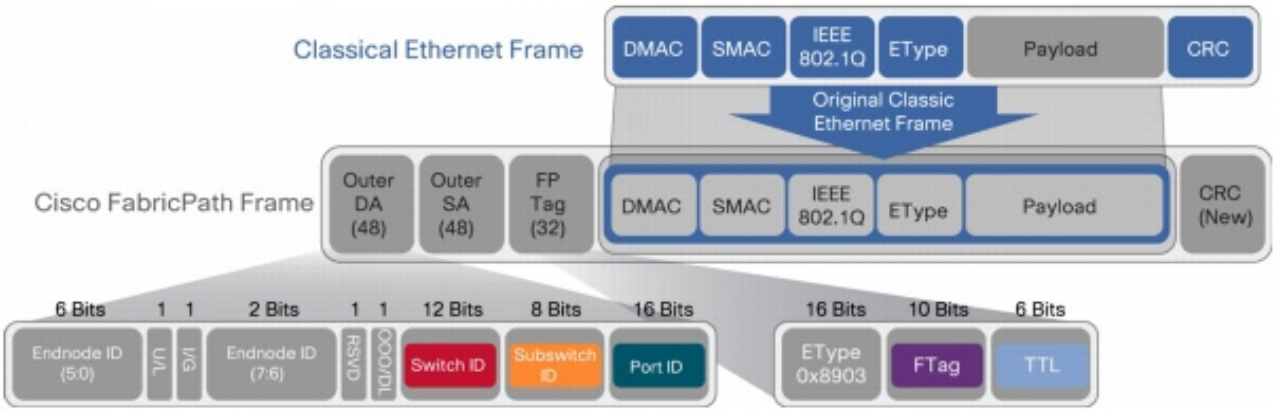

Task	Description	UI Steps
1	Configure ports.	1. Select Ports > Ports > All Ports. 2. Click Quick Port Editor. 3. Configure one network port and two tool ports. For example, select Network for port 1/1/x3. Select Tool for port 1/1/x4 and port 1/1/x1. 4. Select Enable for each port. 5. Click OK. 6. Close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

Task	Description	UI Steps
7	Create another second level map to filter on source and destination IP (bi-directional).	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Enter map1 in the Alias field. • Select Second Level for Type. • Select By Rule for Subtype. • Select the virtual port vp1 for the Source. • Select the tool port 1/1/x4 for the Destination. • Select gsfil from the GSOP list. 4. Add Rule 1. a. Click Add a Rule. b. Select Pass. c. Select IPv4 Source. d. Enter 1.1.1.2 for the IPv4 Address e. Enter 255.255.255.255 for the Net Mask f. Set Position to 1. 5. 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. 6. 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

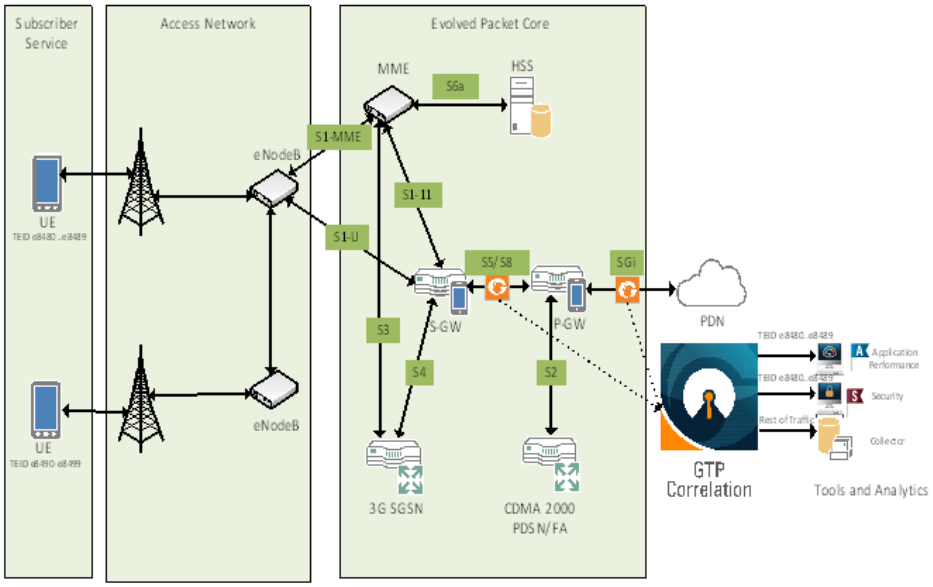

Task	Description	UI Steps
1	Configure one network and three tool type of ports.	1. Select Ports > Ports > All Ports. 2. Click Quick Port Editor. 3. Configure one network port and two tool ports. For example, select Network for port 1/1/x9. Select Tool for the port s 1/1/x13, 1/1/x14, and 1/1/x15. 4. 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.
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 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 Ether Type d. Enter 8903 in the Value field. 5. Click Save.
6	Create a first level map that directs GTP-u traffic from physical network port/s to the tool ports. **NOTE:** In the rule, 2123 is GTP-c traffic.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Enter ctrl_to_tool in the Alias field. • Select Regular for Type. • Select By Rule for Subtype. • Select the port 1/3/x9 for the Source. • Select the port 1/3/x13 and port 1/3/x15 for the Destination. 4. Add a rule. a. Click Add a Rule. b. Select Pass. c. Select Port Source. d. Enter **2123** for the port value. 5. Click Save.
7	Create a second level map that takes traffic from the virtual port, applies the GigaSMART operation, and matches tunnel IDs specified by the gsrule.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Type m1 in the Alias field. • Select Second Level for Type. • 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. 5. Click Save.

Task	Description	UI Steps
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. 5. 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.

Task	Description	UI Steps
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.
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.

Task	Description	UI Steps
6	Create a second level map.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. • Type test1b in the Alias field. • Select Second Level for Type. • Select By Rule for Subtype. • Select the network port for the Source. For example, 1/1/g3. • Select the virtual port vp for the Destination. • Select er2 from the GSOP list. **4.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select MAC Source **d.** Enter the address 0000.0000.0000 for Min and the address 0000.0000.0000 for Max. **5.** Click Save.

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

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

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

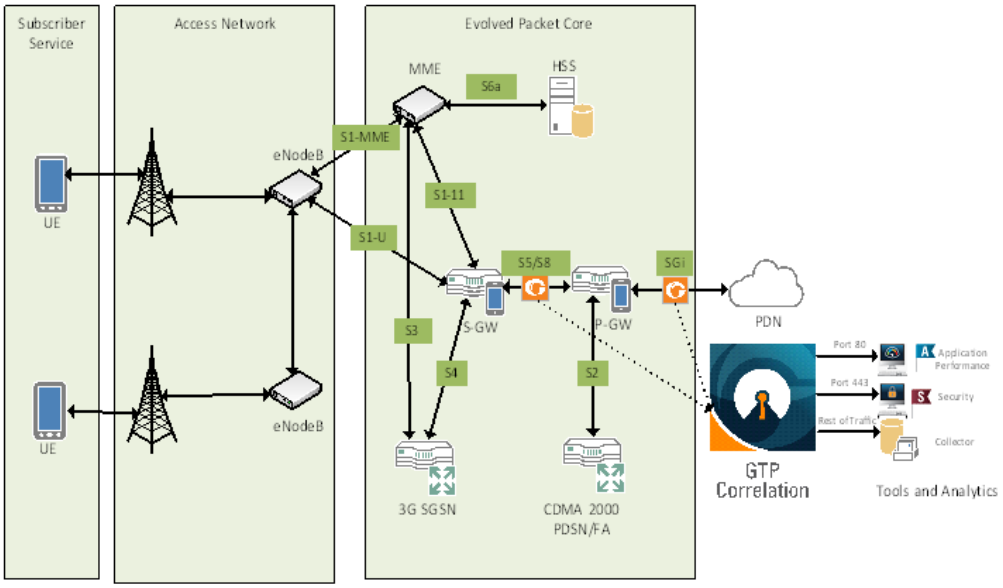

Task	Description	UI Steps
1	Configure one network and three tool type of ports.	1. Select Ports > Ports > All Ports. 2. Click Quick Port Editor. 3. Configure one network port and three tool ports. For example, select Network for port 1/1/x1. Select Tool for the port s 1/1/x10, 1/1/x11, and 1/1/x12. 4.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

Task	Description	UI Steps
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. **5.** 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.

Task	Description	UI Steps
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. • 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. **5.** 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 Visibility Platform 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

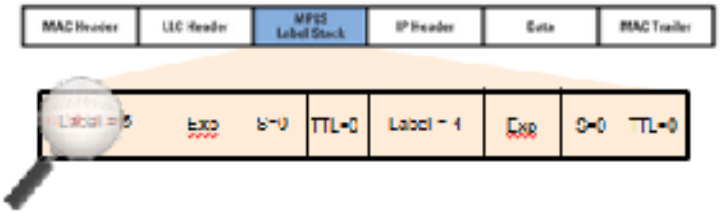

Step	Description	Command
1	Configure ports.	**1.** Select Ports > Ports > All Ports. **2.** Click Quick Port Editor. **3.** Configure one network port and two tool ports. For example, select Network for port 1/1/x3 and select Tool for the port s 1/1/x4 and 1/1/x1 **4.** Select Enable for each port. **5.** 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.

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

Step	Description	Command
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. • 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.

Step	Description	Command
5	Create a first level map to forward VXLAN traffic to the virtual port. VXLAN accepts destination UDP ports 8472 and 4789. Starting in software version 4.5.01, VXLAN also accepts destination UDP port 48879.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Type to_vp in the Alias field. • Select First Level for Type. • Select By Rule for Subtype. • Select the network port for the Source. For example, 1/1/x3 • Select the virtual port vp1 for the Destination. 4. Add a rule. a. Click Add a Rule. b. Select Pass. c. Select Port Source and set the port value to 8472. 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 2. 5. 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. 6. Click Save.

Step	Description	Command
8	Create another second level map to filter on source and destination IP (bi-directional).	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. • Enter map1 in the Alias field. • Select Second Level for Type. • Select By Rule for Subtype. • Select the virtual port vp1 for the Source. • Select the tool port 1/1/x1 for the Destination. • Select gsfil from the GSOP list. **4.** Add Rule 1. **a.** Click Add a Rule. **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. **5.** 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. **6.** 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.

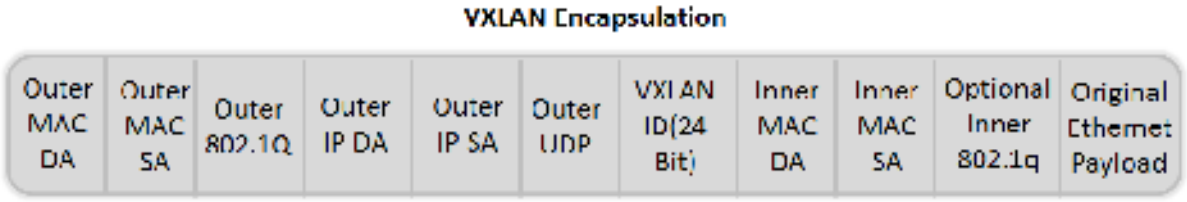

NOTE: This can be applied to any protocol that is supported through header stripping. GTP, VXLAN, ISL, MPLS, MPLS+VLAN, VLAN, VN-Tag, and fabric-path are all supported, as is Gigamon tunnel decapsulation.

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.

Step	Description	Command
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. 3. Create second first operation. a. Click New. b. Type gsfil apf in the Alias field. c. Select gsg1 from the GigaSMART Groups list. d. Select Adaptive Packet Filtering from the GS Operations (GSOP) list and Enable. e. Click Save.
4	Create a virtual port and associate it with the GigaSMART group.	1. From the device view, select GigaSMART > Virtual ports. 2. Click New. 3. Enter vp1 in the Alias field. 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. 5. Click Save.

Step	Description	Command
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. **5.** 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. **6.** Click Save.

Step	Description	Command
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. **5.** 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. **6.** Click Save.

Implement 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. 3. Create second operation. a. Click New. b. Type gsfil_slice in the Alias field. c. Select gsg1 from the GigaSMART Groups list. d. Select APF from the GigaSMART Operations (GSOP) list and Enable. e. Select Slicing from the GigaSMART Operations (GSOP) list and select None. Set Offset to 150. 4. 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. 5. 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 4. Click Save.
7	Create another second level map for the rest of the traffic.	1. Click New. 2. Configure the map. • Type map2 in the Alias field. • Select First Level for Type. • Select By Rule for Subtype. • Select the virtual port vp1 for the Source. • Select the virtual port 1/1/x1 for the Destination. • Select gsfil from the GSOP list. 3. Add a rule. a. Click Add a Rule. b. Select Pass. c. Select IP Version. d. Set Version to v4. e. Set Postilion to 4 4. 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. 5. Click Save.

Step	Description	Command
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. **5.** 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. **4.** 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. **5.** Click Save.

Display APF Statistics

Refer to *APF Statistics Definitions* on page 800 for descriptions of these statistics as well as to *GigaSMART Operations Statistics Definitions* on page 794.

GigaSMART Application Session Filtering (ASF) and Buffer ASF

Required Licenses: Adaptive Packet Filtering (APF) and Application Session Filtering (ASF)
NOTE: The ASF license requires the APF license to be installed as a prerequisite.

DEPRECATION ANNOUNCEMENT: GigaSMART Application Session Filtering is end-of-sale in 5.6.00. Refer to the *Application Intelligence* on page 727 to learn about the feature set and licensing for the improved application filtering and application metadata functionality.

Application Session Filtering (ASF) provides additional filtering on top of Adaptive Packet Filtering (APF). With APF, you can filter on any data patterns within a packet. With ASF, you apply the pattern matching and then send all the packet flows associated with the matched packet to monitoring or security tools.

ASF allows you to filter all traffic corresponding to a session. Use ASF to create a flow session and send the packets associated with the flow session to one or more tools. A flow session consists of one or more fields that you select to define the session. Either the packets for the whole session can be captured or only the packets following a pattern match.

For example, use APF to filter TCP packets to capture the SYN packet. Then use ASF with GigaSMART Load Balancing to send all subsequent packets associated with the session to multiple tool ports. This example is illustrated in *Example 1: ASF, Forward TCP Traffic* on page 1060. For information on capturing a whole session by buffering packets, refer to *Application Session Filtering with Buffering* on page 1055.

Or use APF to create pattern-matching filters in which the pattern is a sequence of data bytes at a variable or fixed offset within a packet. Then use ASF with a specified session definition to capture subsequent packets belonging to the session. When an incoming packet matches an APF rule, a flow session is created. The subsequent incoming packets that match the value of the fields in the flow session will be forwarded to the same tool port as the matching packet.

For example, use APF to pattern match the string *www.gigamon.com*. Use the 5tuple field to identify the flow session. This creates the signature of the session. All the packets associated with the session will be forwarded to a tool port, hence APF becomes flow-aware or session-aware.

ASF provides the following session capabilities:

- filter on one, two, or both MPLS labels and/or VLAN IDs
- filter on both inner and outer IP addresses, Layer 4 ports, and protocols

Pattern matching examples are illustrated in *Example 2: ASF, Forward VNC Traffic* on page 1062 to *Example 4: ASF, Forward GTP Traffic* on page 1066

For information on load balancing, refer to stateful load balancing in the section *GigaSMART Load Balancing* on page 1147.

ASF operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to *Groups of GigaSMART Engine Ports* on page 757 for details.

In ASF and buffer ASF second level maps, a maximum of five (5) maps can be attached to a virtual port (vport). Each map can contain up to 25 gsrules.

Application Session Filtering (with or without buffering) is a pillar of the GigaSECURE Security Delivery Platform. Refer to *GigaSECURE Security Delivery Platform* on page 556.

Session-Aware APF (SAPF) Renaming and Licensing Change to ASF

The GigaSMART feature named Session-Aware Adaptive Packet Filtering (SAPF) in GigaVUE-OS 4.3 was renamed to Application Session Filtering (ASF) in GigaVUE-OS 4.4. It is now encompassed within ASF as the non-buffering equivalent to Application Session Filtering with buffering.

In addition, the license has moved from the APF license to the ASF license.

Application Session Filtering with Buffering

ASF captures packets of a session after an APF rule match. When the APF match occurs in the middle of a session, packets in the session prior to the match are not captured. With some tools needing all the packets of a flow session in order to perform data analysis, GigaSMART uses buffering to ensure that all packets belonging to a flow session are captured and forwarded to the tools. This is referred to as Application Session Filtering with buffering, or buffer ASF.

Buffer ASF uses the pattern-matching and regular expression engine in APF to select packet flows based on matching criteria with one or more packets in the flow session. Buffering ensures that the entire session, from start to finish, is either dropped or forwarded to the security tools or the performance monitoring tools.

To capture all packets belonging to a flow session prior to the APF rule match, ASF needs to know the first packet of a flow session. For this, GigaSMART supports both TCP and UDP connections.

For TCP connections, the TCP SYN packet is used to indicate the start of a session. GigaSMART captures and stores (or buffers) all packets of a flow session until an APF match occurs. After that, GigaSMART will either forward or drop all stored packets belonging to that session based on the APF pass or drop rule that is configured. Subsequent packets after the APF match will be forwarded or dropped as they arrive.

For UDP connections, there is no special packet that indicates the start of a UDP flow session from a Layer 4 perspective. GigaSMART will take the first UDP packet of a session it encounters as the start of a session flow. This may result in incomplete capture at the beginning of configuration or at system boot up, but as new UDP sessions arrive, GigaSMART will capture the first packet of the flows.

ASF and Buffer ASF Session Definitions

Use the **Session Field** of the **ASF** page to define ASF and buffer ASF sessions by specifying session field attributes to add or delete. A session field is a group of one or

more fields that define a flow session. (From the device view, select **GigaSMART > ASF** and click **New** to open the ASF page.)

A flow session consists of field names and attributes. Some field names include multiple attributes, which provide a quick way to define sessions.

The field names and attributes are as follows:

- gtpu-teid
- ipv4 (ipv4-src, ipv4-dst)
- ipv4-5tuple (ipv4-src, ipv4-dst, l4port-src, l4port-dst, ipv4-protocol)
- ipv4-dst
- ipv4-l4port-dst (ipv4-src, ipv4-dst, l4port-dst)
- ipv4-protocol
- ipv4-src
- ipv4-src-l4port-dst (ipv4-src, l4port-dst)
- ipv6 (ipv6-src, ipv6-dst)
- ipv6-5tuple (ipv6-src, ipv6-dst, l4port-src, l4port-dst, ipv6-protocol)
- ipv6-dst
- ipv6-l4port-dst (ipv6-src, ipv6-dst, l4port-dst)
- ipv6-protocol
- ipv6-src
- ipv6-src-l4port-dst (ipv6-src, l4port-dst)
- l4port (l4port-src, l4port-dst)
- l4port-dst
- l4port-src
- mpls-label
- vlan-id

An ASF session definition consists of combinations of the fields and attributes in the list above. In addition, for all IP and L4 port fields in the packet, each ASF session field must specify **outer** or **inner** for location. Outer specifies the first IP or L4 port in the packet. Inner specifies the second IP or L4 port in the packet (usually inside tunneling). For VLAN ID and MPLS label fields, a position (1 or 2) must be specified. Position 1 is the first occurrence of the protocol header or field in the packet. Position 2 is the second occurrence of the protocol header or field in the packet.

A buffer ASF session definition consists of combinations of the fields and attributes in the list above. **One restriction is that ipv4-src or ipv6-src needs to be defined, as a minimum.** In addition, the following restrictions apply to buffer ASF session definitions:

- the gtpu-teid field name is not supported
- the IP and L4 port fields only support location **outer**
- the VLAN ID field only supports position 1

All packets belonging to the same source and destination IP will be considered as the same flow session. This is useful if you want to capture all packets belonging to separate TCP/UDP connections that have the same IPs, such as control or data flows.

Define ASF Session

When defining an ASF session, enter the fields, attributes, and options, then save. The changes only take effect when after you save. For example, in Figure 30-92, the settings are:

- Alias is asf2, which is name of the ASF Profile for the GigaSMART Operation
- buffer enabled
- Buffer Count before Match is 5
- ipv4-5tuple outer
- vlan-id position 1

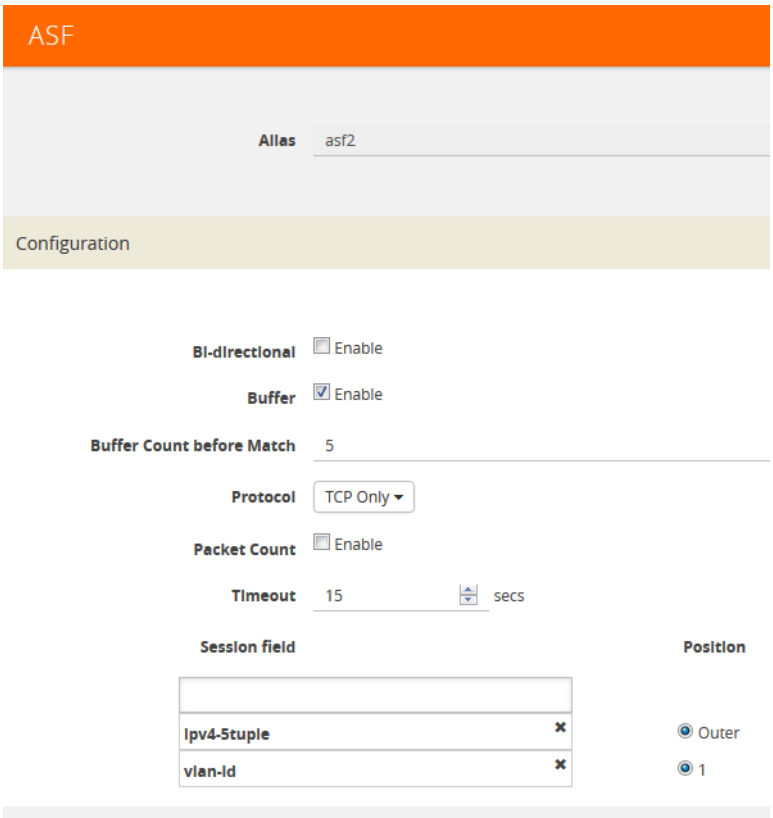

Figure 30-92: An ASF Configuration

Quick Session Delete for Buffer ASF

For a buffer ASF session defined with ipv4-5tuple or ipv6-5tuple, there is a quick session delete for TCP connections. The session is deleted 4 seconds after RST or both FIN packets are detected, signaling the end of the flow.

Specify Resources for Buffer ASF

On GigaVUE-HC3, GigaVUE-HC2, GigaVUE-HC1, and GigaVUE HD Series nodes, buffer ASF supports from 2 to 5 million sessions. On GigaVUE-HB1, buffer ASF supports 2 million sessions.

The large number of sessions can use a lot of memory resources on the GigaSMART line card or module. Occasionally, the GigaSMART line card or module will need to be reloaded for changes to take effect and to allocate resources accordingly.

The GigaSMART line card or module does not need to be reloaded the first time buffer ASF resources are allocated. But subsequent changes, such as increasing from 3 million to 4 million sessions, or decreasing from 3 million to 2 million sessions will require a reload for the changes to take effect.

You can reload the card from the UI, by doing the following:

1. Select **Chassis** in the Navigation pane.
2. Clicking the Table View button to open the Table View of the Chassis.
3. Under Cards, select the card to reload.
4. From the **Actions** menu, select **Shut Down**.
5. Now select **Start Up** from **Actions** menu.

Alternatively, you can use the following GigaVUE-FM API to reload the card:

```
PATCH /inventory/chassis/cards/{slotId}
```

For more information about the GigaVUE-FM APIs, refer to the *GigaVUE-FM REST API Getting Started Guide* and the *GigaVUE-FM API Reference*.

ASF and Buffer ASF Session Notes

The following table summarizes notes relating to ASF and buffer ASF:

Notes
A session field can only be modified or deleted if it is *not* configured in any GigaSMART operation.
A session field can only contain the same session attribute and position pair once.
A session field cannot contain overlapped session attribute and position. For example, the following is not valid: ipv4-5tuple outer and ipv4-src outer.
Up to a maximum of 25 flow session aliases are supported for ASF.
A total of 4 session tables per GigaSMART engine are supported for ASF. Each table has its own session definition.
Up to 2 million session entries are supported for ASF. The entries are shared by all session aliases.
Each session table (session alias) can only be used once within a gsgroup.
The number of buffer ASF sessions supported is configurable from 2 to 5 million.
The number of packet buffers supported for buffer ASF is from 2 to 5 million.

Buffer ASF Packet Processing Special Cases

The following are special cases of packet processing for buffer ASF:

- Non-TCP SYN packet received and no session matched—When a non-TCP SYN packet is received and there is no session matched, the packet will be considered as *no match* and will be passed to other maps. If there is no match after all maps have been processed, the packet will be forwarded to a shared collector, if one is configured.
- Out of session—When a TCP SYN packet is received and no free session is available, the packet will be considered as *no match*. Other packets belonging to this session will also be considered as *no match*, as for the special case described above.
- Out of packet storage buffer—When the buffer is full for the first packet of a session, the session will not be created and the packet will be considered as *no match*. When the buffer is full for an existing session, and the APF match has not yet occurred, a flag will be set and the current packet and all buffered packets will be considered as *no match*. Subsequent packets will also be considered as *no match*.
- Exceeded configured buffering limit—When there is no APF match after the configured number of packets have been buffered, all buffered packets and all subsequent packets belonging to this session will be considered as *no match*.

Map Statistics

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

ASF and Buffer ASF Examples

Refer to the following ASF examples (non-buffered):

- *Example 1: ASF, Forward TCP Traffic* on page 1060
- *Example 2: ASF, Forward VNC Traffic* on page 1062
- *Example 3: ASF, Forward Traffic Matching a Pattern* on page 1063
- *Example 4: ASF, Forward GTP Traffic* on page 1066

Refer to the following buffer ASF examples:

- *Example 1: Buffer ASF, Drop Netflix Traffic* on page 1067
- *Example 2: Buffer ASF, Drop YouTube Traffic* on page 1070
- *Example 3: Buffer ASF, Drop Windows Update Traffic* on page 1072
- *Example 4: Buffer ASF, Forward VNC Traffic* on page 1073
- *Example 5: Buffer ASF, Forward HTTPS Traffic on Non-Standard Port* on page 1075

Example 1: ASF, Forward TCP Traffic

In Example 1, ASF is used with GigaSMART Load Balancing and Adaptive Packet Filtering to load balance TCP traffic among multiple tool ports. TCP SYN indicates the start of a connection. Once the TCP SYN packet is detected, subsequent packets belonging to the same TCP connection will be forwarded to a configured tool port. Packets belonging to the same connection will be sent to the same tool port, regardless of the number of connections.

NOTE: This example uses APF to filter TCP packets to capture the SYN packet. Alternatively, use buffer ASF to capture a whole session by buffering packets.

Task	Description	UI Steps
1	Create a flow session.	1. From the device view, select GigaSMART > ASF. 2. Click New. 3. Type asf4 in the Alias field. 4. Select ipv4-tuple from the Session field list. 5. Select outer. 6. Click Save.
2	Create a port group and specify the tool ports for load balancing.	1. Select Ports > Port Groups > All Port Groups. 2. Click New. 3. Type portgrp1 in the Alias field. 4. Select Tool. 5. Select SMART Load Balancing. 6. Click in the Ports field and select the tool ports. For example, 1/1/x6,1/1/x7,1/2/x3, and 1/2/x4.
3	Configure a GigaSMART group and associate it with GigaSMART engine ports.	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type gsgrp1 in the Alias field. 4. Select two engine ports from the Port List field. For example, 1/3/e1 and 1/3/e2 5. Click Save.
4	Configure the combined GigaSMART operation.	1. From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation and create two GigaSMART Operations. 2. Click New. 3. Type gsop1 in the Alias field. 4. Select gsgrp1 from the GigaSMART Groups list. 5. Select the operations. • APF • ASF with asf4 for the ASF profile • Load Balancing with Stateful Type ASF, and Round Robin 6. Click Save.

Task	Description	UI Steps
5	Create a virtual port and associate it with the GigaSMART group.	1. From the device view, select GigaSMART > Virtual ports. 2. Click New. 3. Enter vp1 in the Alias field. 4. Select gsg1 from the GigaSMART Groups list. 5. Click Save.
6	Create a first level map.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Type map11 in the Alias field. • Select First Level for Type. • Select By Rule for Subtype. • Select the network port 1/1/x1 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 4 5. Click Save.
7	Create a second level map. The gsrule captures the first packet of a session.	1. Click New. 2. Configure the map. • Type map22 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 group portgrp1 for the Destination. • Select gsop1 form the GSOP list. 3. Add a rule. a. Click Add a Rule. b. Select Pass. c. Select TCP Control. d. Enter 2 for Value. e. Enter 0 for Mask. f. Set Position to 1. 4. Click Save.

Example 2: ASF, Forward VNC Traffic

In Example 2, traffic from a Virtual Network Computing (VNC) application is forwarded from network port 1/1/x1 to tool port 1/1/x6. Packets will be matched with a VNC signature. Once a packet is matched, subsequent packets with the same IPv4 5tuple will be forwarded to the same destination as the matching packet. By default, both the forward and the reverse traffic of the same session will be captured and forwarded.

Step	Description	Command
1	Create a flow session.	1. From the device view, select GigaSMART > ASF. 2. Click New. 3. Type asf1 in the Alias field. 4. Select ipv4-tuple from the Session field list. 5. Select outer. 6. Click Save.
2	Configure a GigaSMART group and associate it with GigaSMART engine ports.	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type gsgrp1 in the Alias field. 4. Select two engine ports from the Port List field. For example, 1/3/e1 and 1/3/e2 5. Click Save.
3	Configure the combined GigaSMART operation.	1. From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations and create two GigaSMART Operations. 2. Click New. 3. Type gsop1 in the Alias field. 4. Select gsgrp1 from the GigaSMART Groups list. 5. Select the operations. • APF • ASF with asf1 for the ASF profile 6. Click Save.
4	Create a virtual port and associate it with the GigaSMART group.	1. From the device view, select GigaSMART > Virtual ports. 2. Click New. 3. Enter vp1 in the Alias field. 4. Select gsgrp1 from the GigaSMART Groups list. 5. Click Save.

Step	Description	Command
5	Create a first level map.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Type map11 in the Alias field. • Select First Level for Type. • Select By Rule for Subtype. • Select the network port 1/1/x1 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 4 5. Click Save.
6	Create a second level egress map. The gsrule contains the VNC signature.	1. Click New. 2. Configure the map. • Type map22 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 group portgrp1 for the Destination. • Select gsop1 form the GSOP list. 3. Add a rule. a. Click Add a Rule. b. Select Pass. c. Select Pattern Match. d. Select regex for Type and enter ^rfb 00[1-9]\.00[0-9]\x0a$ e. Set Offset from 16 to 1000 4. Click Save.

Example 3: ASF, Forward Traffic Matching a Pattern

In Example 3, the traffic that matches a particular pattern (ymsglypnslyhoo) is forwarded from network port 1/1/x1 to tool port 1/1/x6 after adding a VLAN tag. Packets will be matched with the special signature. Once a packet is matched, subsequent packets with the same source IP, source port, and VLAN ID will be forwarded to the same destination as the matching packet (after the VLAN header is inserted). By

default, both the forward and the reverse traffic of the same session will be captured and forwarded.

Task	Description	UI Steps
1	Create a flow session and other parameters.	1. From the device view, select GigaSMART > ASF. 2. Click New. 3. Type asf2 in the Alias field. 4. Enable Packet Count. 5. Set Number of packets to 50. 6. Set the session field. • Select ipv4-src outer • Select vlan-id position 1 7. Select outer. 8. Click Save.
2	Configure a GigaSMART group and associate it with GigaSMART engine ports.	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type gsgrp1 in the Alias field. 4. Select two engine ports from the Port List field. For example, 1/3/e1 and 1/3/e2 5. Click Save.
3	Configure the GigaSMART operation.	1. From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation and create two GigaSMART Operations. 2. Click New. 3. Type gsop1 in the Alias field. 4. Select gsgrp1 from the GS Groups list. 5. Select the operations. • Adaptive Packet Filtering • Add Header and set VLAN to 1000 • ASF with asf2 for the ASF profile 6. Click Save.
4	Create a virtual port and associate it with the GigaSMART group.	1. From the device view, select GigaSMART > Virtual ports. 2. Click New. 3. Enter vp1 in the Alias field. 4. Select gsgrp1 from the GigaSMART Groups list. 5. Click Save.

Task	Description	UI Steps
5	Create a first level map.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Type map11 in the Alias field. • Select First Level for Type. • Select By Rule for Subtype. • Select the network port 1/1/x1 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 4 5. Click Save.
6	Create a second level map. The gsrule contains the special signature.	1. Click New. 2. Configure the map. • Type map22 in the Alias field. • Select Second Level for Type. • Select By Rule for Subtype. • Select the virtual port vp1 for the Source. • Select the too port 1/1/x6 for the Destination. • Select gsop1 form the GSOP list. 3. Add a rule. a. Click Add a Rule. b. Select Pass. c. Select Pattern Match. d. Select regex for Type and enter (ymsglypnslyhoo) e. Set Offsett from 16 to 1000 4. Click Save.

Example 4: ASF, Forward GTP Traffic

In Example 4, GTP traffic from network port 1/1/x1 is load balanced based on inner IP and tunnel ID to four tool ports: 1/1/x6, 1/1/x7, 1/2/x3, and 1/2/x4. APF filters GTP-u packets. Once a packet is matched, subsequent packets in the same direction with the same gtpu-teid and inner IP will be forwarded to the same destination as the matching packet. In Example 4, both the outer and inner IP are IPv4.

Task	Description	UI Step
1	Create a flow session and other parameters.	1. From the device view, select GigaSMART > ASF. 2. Click New. 3. Type asf3 in the Alias field. 4. Set timeout to 90. 5. Set the session field. • Select gtpu-teid • Select Ipv4 inner 6. Select outer. 7. Click Save.
2	Create a port group and specify the tool ports for load balancing.	1. Select Ports > Port Groups > All Port Groups. 2. Click New. 3. Type portgrp1 in the Alias field. 4. Select Tool. 5. Select SMART Load Balancing. 6. Click in the Ports field and select the tool ports. For example, 1/1/x6,1/1/x7,1/2/x3, and 1/2/x4.
3	Configure a GigaSMART group and associate it with GigaSMART engine ports.	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type gsgrp1 in the Alias field. 4. Select two engine ports from the Port List field. For example, 1/3/e1 and 1/3/e2 5. Click Save.
4	Configure the combined GigaSMART operation.	1. From the device view, select GigaSMART > GigaSMART Operations (GSOPS) > GigaSMART Operation and create two GigaSMART Operations. 2. Click New. 3. Type gsop1 in the Alias field. 4. Select gsgrp1 from the GigaSMART Groups list. 5. Select the operations. • Adaptive Packet Filtering • ASF with asf3 for the ASF profile • Load Balancing with Stateful, Type ASF, and Least Conn 6. Click Save.

Task	Description	UI Step
5	Create a virtual port and associate it with the GigaSMART group.	1. From the device view, select GigaSMART > Virtual ports. 2. Click New. 3. Enter vp1 in the Alias field. 4. Select gsgrp1 from the GigaSMART Groups list. 5. Click Save.
6	Create a first level map.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Type map11 in the Alias field. • Select First Level for Type. • Select By Rule for Subtype. • Select the network port 1/1/x1 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 Protocol and set Value to UDP. d. Select Port Destination and set the port value to 2152 5. Click Save.
7	Create a second level map.	1. Click New. 2. Configure the map. • Type map22 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 group **portgrp1** for the Destination. • Select gsop1 form the GSOP list. 3. Add a rule. a. Click Add a Rule. b. Select Pass. c. Select Pattern Match. d. Select IPv4 Protocol and enter the IPv4 address. Set Position to 1. e. Select Ipv4 Destination and set the port value to 2152. Set Position to 1. 4. Click Save.

Example 1: Buffer ASF, Drop Netflix Traffic

In Example 1, the goal is to drop all Netflix traffic. The flow session is defined by the 5tuple field and the first occurrence of VLAN ID. The Netflix traffic is expected to be

identified in the first 6 packets of a session. (Configure the maximum number of packets buffered before the match to 5.) A maximum of 3 million sessions is specified.

Task	Description	UI Steps
1	Configure a GigaSMART group and associate it with GigaSMART engine ports and Define the maximum number of sessions, in millions.	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type gsgrp1 in the Alias field. 4. Select two engine ports from the Port List field. For example, 1/3/e1 and 1/3/e2 5. Under Params Resource Buffer select ASF and set the Buffer size to 3. 6. Click Save.
2	If needed, reload the GigaSMART line card or module to allocate the resources for buffer ASF.	If you reset the buffer size of an ASF profile in Task 1, go to the Chassis page and select Table View. Under Cards, select the card to reload. From the Actions menu, select Shut Down and then Start Up. You can also issue the following CLI commands to reboot the card (the card is in slot 3 in this example): **(config) # card slot 3 down** **(config) # no card slot 3 down**
3	Create a flow session, specify the buffer count before the match, and enable buffering. **NOTE:** The default protocol is TCP, so it does not need to be specified.	1. From the device view, select GigaSMART > ASF. 2. Click New or select an existing ASF profile then click Edit. 3. Type asf2 in the Alias field if this is a new ASF profile. 4. Enable Buffer. 5. Set Buffer Count before Match to 5. 6. Set the session field. • Select ipv4-5tuple outer • Select vlan-id position 1 7. Click Save.
4	Configure the combined GigaSMART operation.	1. From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation and create two GigaSMART Operations. 2. Click New. 3. Select gsgrp1 from the GigaSMART Groups list. 4. Type gsop1 in the Alias field. 5. Select the operations. • APF • ASF with asf2 for the ASF profile
5	Create a virtual port and associate it with the GigaSMART group.	1. From the device view, select GigaSMART > Virtual Ports. 2. Click New. 3. Enter vp1 in the Alias field. 4. Select gsgrp1 from the GigaSMART Groups list. 5. Click Save.

Task	Description	UI Steps
6	Create a first level map.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Type map11 in the Alias field. • Select First Level for Type. • Select By Rule for Subtype. • Select the network port 1/1/x1 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 Protocol and set Value to UDP. d. Select Port Destination and set the port value to 2152 5. Click Save.
7	Create a second level map. The gsrule specifies the traffic to drop, using keywords. Buffered packets and all subsequent packets will be dropped.	1. Click New. 2. Configure the map. • Type map22 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 form the GSOP list. 3. Add a rule. a. Click Add a Rule. b. Select Drop. c. Select Pattern Match. d. Select regex and enter netflix|nflxvideo|nflximg|Netflix|nflxext. e. Set the offset from 0 to 1460 f. Set Protocol to tcp and set Position to 1. 4. Click Save.

Example 2: Buffer ASF, Drop YouTube Traffic

In Example 2, the goal is to drop all YouTube traffic. The YouTube traffic is expected to be identified in the first 7 packets of a session. (Configure the maximum number of packets buffered before the match to 6.) A maximum of 4 million sessions is specified.

Step	Description	Command
1	Configure a GigaSMART group and associate it with GigaSMART engine ports and define the maximum number of sessions, in millions	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type gsgrp1 in the Alias field. 4. Select two engine ports from the Port List field. For example, 1/3/e1 and 1/3/e2 5. Under Params Resource Buffer, select ASF and set the Buffer Size to 4. 6. Click Save.
2	If needed, reload the GigaSMART line card or module to allocate the resources for buffer ASF.	If you reset the buffer size of an ASF profile in Task 1, go to the Chassis page and select Table View. Under Cards, select the card to reload. From the Actions menu select Shut Down and then Start Up. You can also issue the following CLI commands to reboot the card (the card is in slot 3 in this example): **(config) # card slot 3 down** **(config) # no card slot 3 down**
3	Create a flow session, specify the buffer count before the match, and enable buffering. **NOTE:** The default protocol is TCP, so it does not need to be specified.	1. From the device view, select GigaSMART > ASF. 2. Click New or select an existing ASF profile then click Edit. 3. Type asf2 in the Alias field if this is a new ASF profile. 4. Enable Buffer. 5. Set Buffer Count before Match to 6. 6. Set the session field to ipv4-5tuple outer 7. Click Save.
4	Configure the combined GigaSMART operation.	1. From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations and create two GigaSMART Operations. 2. Click New. 3. Type gsop1 in the Alias field. 4. Select gsgrp1 from the GigaSMART Groups list. 5. Select the operations. • Adaptive Packet Filtering • ASF with asf2 for the ASF profile 6. Click Save.
5	Create a virtual port and associate it with the GigaSMART group.	1. From the device view, select GigaSMART > Virtual Ports. 2. Click New. 3. Enter vp1 in the Alias field. 4. Select gsgrp1 from the GigaSMART Groups list. 5. Click Save.

Step	Description	Command
6	Create a first level map.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Type map11 in the Alias field. • Select First Level for Type. • Select By Rule for Subtype. • Select the network port 1/1/x1 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.
7	Create a second level map. The gsrule specifies the traffic to drop, using keywords. Buffered packets and all subsequent packets will be dropped.	1. Click New. 2. Configure the map. • Type map22 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 form the GSOP list. 3. Add a rule. a. Click Add a Rule. b. Select Drop. c. Select Pattern Match. d. Select regex and enter youtubelytimglyt3.ggphtltubeMogulltmogul. e. Set the offset from 0 to 1460 f. Set Protocol to tcp and set Position to 1. 4. Click Save.

Example 3: Buffer ASF, Drop Windows Update Traffic

In Example 3, the goal is to drop all Windows update traffic. The Windows update traffic is expected to be identified on the HTTP request packet of a session. A maximum of 2 million sessions is specified.

Task	Description	UI Steps
1	Configure a GigaSMART group and associate it with GigaSMART engine ports and define the maximum number of sessions, in millions.	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type gsgrp1 in the Alias field. 4. Select two engine ports from the Port List field. For example, 1/3/e1 and 1/3/e2 5. Under Params Resource Buffer, select ASF and set the Buffer Size to 2. 6. Click Save.
2	If needed, reload the GigaSMART line card or module to allocate the resources for buffer ASF.	If you reset the buffer size of an ASF profile in Task 1, go to the Chassis page and select Table View. Under Cards, select the card to reload. From the Actions menu select Shut Down and then Start Up. You can also issue the following CLI commands to reboot the card (the card is in slot 3 in this example): **(config) # card slot 3 down** **(config) # no card slot 3 down**
3	Create a flow session, specify the buffer count before the match, and enable buffering. **NOTE:** The default protocol is TCP, so it does not need to be specified.	1. From the device view, select GigaSMART > ASF. 2. Click New or select an existing ASF profile then click Edit. 3. Type asf2 in the Alias field if this is a new ASF profile. 4. Enable Buffer. 5. Set Buffer Count before Match to 3. 6. Set the session field to ipv4-5tuple outer 7. Click Save.
4	Configure the combined GigaSMART operation.	1. From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations and create two GigaSMART Operations. 2. Click New. 3. Type gsop1 in the Alias field. 4. Select gsgrp1 from the GigaSMART Groups list. 5. Select the operations. • Adaptive Packet Filtering • ASF with asf2 for the ASF profile 6. Click Save.
5	Create a virtual port and associate it with the GigaSMART group.	1. From the device view, select GigaSMART > Virtual Ports. 2. Click New. 3. Enter vp1 in the Alias field. 4. Select gsgrp1 from the GigaSMART Groups list. 5. Click Save.

Task	Description	UI Steps
6	Create a first level map.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Type map11 in the Alias field. • Select First Level for Type. • Select By Rule for Subtype. • Select the network port 1/1/x1 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.
7	Create a second level map. The gsrule specifies the traffic to drop. Buffered packets and all subsequent packets will be dropped.	1. Click New. 2. Configure the map. • Type map22 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 form the GSOP list. 3. Add a rule. a. Click Add a Rule. b. Select Drop. c. Select Pattern Match. d. Select regex and enter msdownload/update/software. e. Set the offset from 0 to 1460 f. Set Protocol to tcp and set Position to 1. 4. Click Save.

Example 4: Buffer ASF, Forward VNC Traffic

In Example 4, the goal is to forward VNC traffic from network port 1/1/x1 to tool port 1/1/x6. All packets belonging to the TCP connection need to be sent to the tool port.

The first data packet after the TCP handshake is expected to contain the VNC pattern match. A maximum of 2 million sessions is specified.

Task	Description	UI Steps
1	Configure a GigaSMART group and associate it with GigaSMART engine ports and define the maximum number of sessions, in millions.	1. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type gsgrp1 in the Alias field. 4. Select two engine ports from the Port List field. For example, 1/3/e1 and 1/3/e2 5. Under Params Resource Buffer, select ASF and set the Buffer Size to 2. 6. Click Save.
2	If needed, reload the GigaSMART line card or module to allocate the resources for buffer ASF.	If you reset the buffer size of an ASF profile in Task 1, go to the Chassis page and select Table View. Under Cards, select the card to reload. From the Actions menu select Shut Down and then Start Up. You can also issue the following CLI commands to reboot the card (the card is in slot 3 in this example): **(config) # card slot 3 down** **(config) # no card slot 3 down**
3	Create a flow session, specify the buffer count before the match, and enable buffering. **NOTE:** The default protocol is TCP, so it does not need to be specified.	1. From the device view, select GigaSMART > Application Session Filtering. 2. Click New or select an existing ASF profile then click Edit. 3. Type asf1 in the Alias field if this is a new ASF profile. 4. Enable Buffer. 5. Set Buffer Count before Match to 3. 6. Set the session field to ipv4-5tuple outer 7. Click Save.
4	Configure the combined GigaSMART operation.	1. From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation and create two GigaSMART Operations. 2. Click New. 3. Type gsop1 in the Alias field. 4. Select gsgrp1 from the GigaSMART Groups list. 5. Select the operations. • Adaptive Packet Filtering • ASF with asf1 for the ASF profile 6. Click Save.
5	Create a virtual port and associate it with the GigaSMART group.	1. From the device view, select GigaSMART > Virtual Ports. 2. Click New. 3. Enter vp1 in the Alias field. 4. Select gsgrp1 from the GigaSMART Groups list. 5. Click Save.

Task	Description	UI Steps
6	Create a first level map.	1. Select Maps > Maps > Maps. 2. Click New. 3. Configure the map. • Type map11 in the Alias field. • Select First Level for Type. • Select By Rule for Subtype. • Select the network port 1/1/x1 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.
7	Create a second level map. The gsrule specifies the traffic to pass. Buffered packets and all subsequent packets will be passed.	1. Click New. 2. Configure the map. • Type map22 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 form the GigaSMART Operations (GSOP) list. 3. Add a rule. a. Click Add a Rule. b. Select Pass. c. Select Pattern Match. d. Select regex and enter ^rfb 00[1-9]\.00[0-9]\x0a$. e. Set Protocol to tcp and set Position to 1. 4. Click Save.

Example 5: Buffer ASF, Forward HTTPS Traffic on Non-Standard Port

In Example 5, the goal is to forward HTTPS traffic that uses a non-standard Layer 4 port. All packets belonging to the TCP connection need to be sent to the tool port. A maximum of 5 million sessions is specified.

Task	Description	UI Steps
1	Configure a GigaSMART group and associate it with GigaSMART engine ports and define the maximum number of sessions, in millions.	1. GigaSMART > GigaSMART Groups > GigaSMART Groups. 2. Click New. 3. Type gsgrp1 in the Alias field. 4. Select two engine ports from the Port List field. For example, 1/3/e1 and 1/3/e2 5. Under Params Resource Buffer, select ASF and set the Buffer Size to 2. 6. Click Save.

Task	Description	UI Steps
2	If needed, reload the GigaSMART line card or module to allocate the resources for buffer ASF.	If you reset the buffer size of an ASF profile in Task 1, go to the Chassis page and select Table View. Select the card in the table. From the Actions menu select Shut Down and then Start Up. You can also issue the following commands to reboot the card (the card is in slot 3 in this example): **(config) # card slot 3 down** **(config) # no card slot 3 down**
3	Create a flow session, specify the buffer count before the match, and enable buffering. **NOTE:** The default protocol is TCP, so it does not need to be specified.	**1.** From the device view, select GigaSMART > ASF. **2.** Click New or select an existing ASF profile then click Edit. **3.** Type asf2 in the Alias field if this is a new ASF profile. **4.** Enable Buffer. **5.** Set Buffer Count before Match to 3. **6.** Set the session field to ipv4-5tuple outer **7.** Click Save.
4	Configure the combined GigaSMART operation.	**1.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations and create two GigaSMART Operations. **2.** Click New. **3.** Type gsop1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Select the operations. • APF • ASF with asf2 for the ASF profile **6.** Click Save.
5	Create a virtual port and associate it with the GigaSMART group.	**1.** From the device view, select GigaSMART > Virtual Ports. **2.** Click New. **3.** Enter vp1 in the Alias field. **4.** Select gsgrp1 from the GigaSMART Groups list. **5.** Click Save.
6	Create a first level map.	**1.** Select Maps > Maps > Maps. **2.** Click New. **3.** Configure the map. • Type map11 in the Alias field. • Select First Level for Type. • Select By Rule for Subtype. • Select the network port 1/1/x1 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.

Task	Description	UI Steps
7	Create a second level map. The gsrule specifies the traffic to pass. The RegEx expression identifies the traffic as SSL. Buffered packets and all subsequent packets will be passed.	**1.** Click New. **2.** Configure the map. • Type map22 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 form the GSOP list. **3.** Add a rule. **a.** Click Add a Rule. **b.** Select Pass. **c.** Select Pattern Match. **d.** Select regex and enter x16\x03.{3}\x01. **e.** Set Protocol to tcp and set Position to 1. **4.** Click Save.

Display ASF Statistics

To display ASF statistics on the GigaSMART operation, select **GigaSMART > Statistics**.

Refer to *ASF Statistics Definitions* on page 800 for descriptions of these statistics as well as to *GigaSMART Operations Statistics Definitions* on page 794.

GigaSMART NetFlow Generation

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

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 Visibility Platform completely eliminates the risk of using core production network resources in generating this data. Combined with the flexibility offered by Gigamon’s 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.

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 Visibility Platform 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 Visibility Platform 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 30-93.

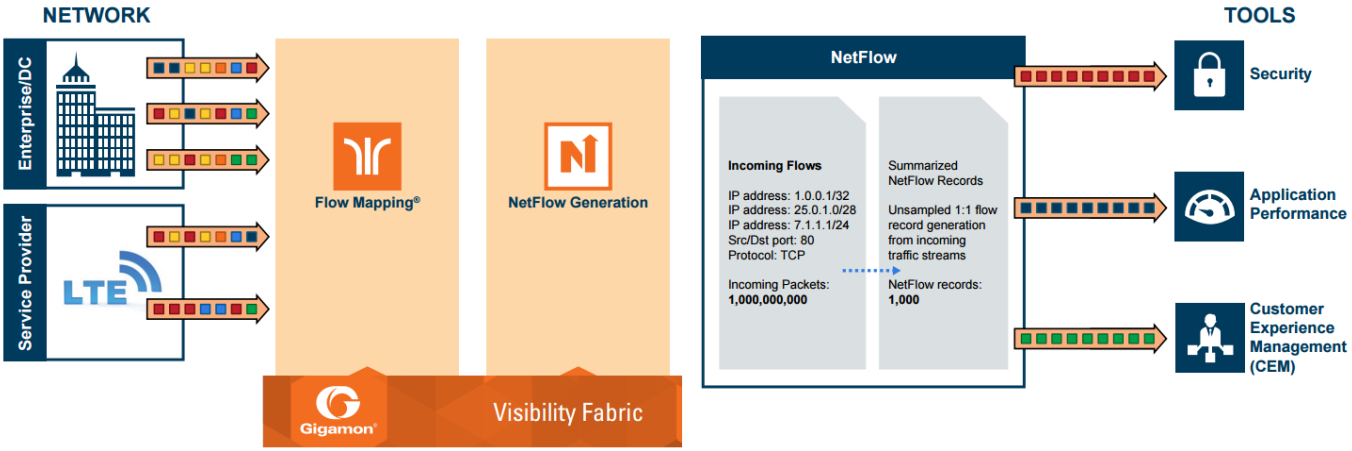

Figure 30-93: NetFlow Generation Gigamon Solution

In Figure 30-93, incoming packets from network(s) enter the Gigamon Visibility Platform 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.

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 30-94 shows the NetFlow Generation components.

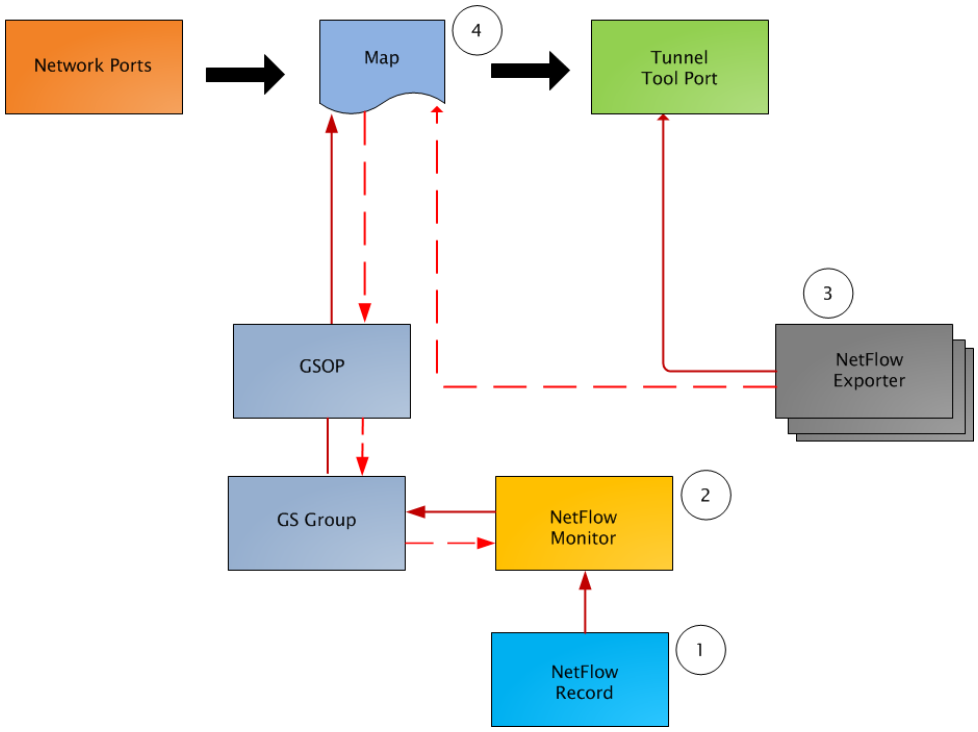

Figure 30-94: NetFlow Generation Components

Figure 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 *Example 1: NetFlow Generation Configuration* on page 1110 and *Example 2: NetFlow Generation Configuration* on page 1116.

For examples of first and second level maps, refer to *Example 3: NetFlow Generation Configuration* on page 1123 and *Example 4: NetFlow Generation Configuration* on page 1128.

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.

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. 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 tIP 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 to 255.255.255.255.
- The maximum number of exporters supported by the IP interface is one.

Enhancements to NetFlow

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

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 *Example 2: NetFlow Generation Configuration* on page 1116.

Remote Interface IDs

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

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

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

Note the following:

- You cannot modify discovery once the map is defined.

- Local port IDs are unique across a cluster. Remote IDs might not be unique. With port discovery enabled, there is a possibility of port ID collisions.

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

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

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

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

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

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Records**.
2. Click **New** to create a new Record or select the recored and click Edit to change an existing record.
3. Click in the field **Non-Key Fields (Collect)** and select **Interface** from the list.
4. Select **Input** and then **Input Name**. Specify an input **Width** as shown in Figure 30-95.
5. Click **Save**.

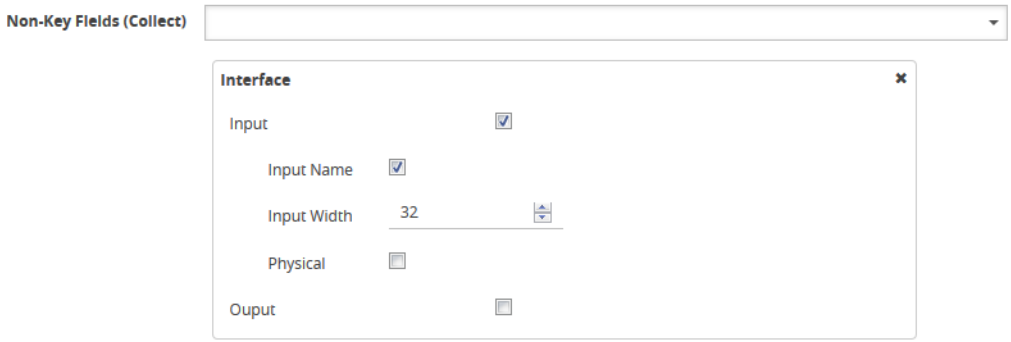

Figure 30-95: Collect Field Interface input Name

NetFlow Option Templates

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

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

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

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

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

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

IPFIX Extension: HTTP Response Code

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

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

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

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

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

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

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

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

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Records**.
2. Click **New**. Refer to Figure 30-96.

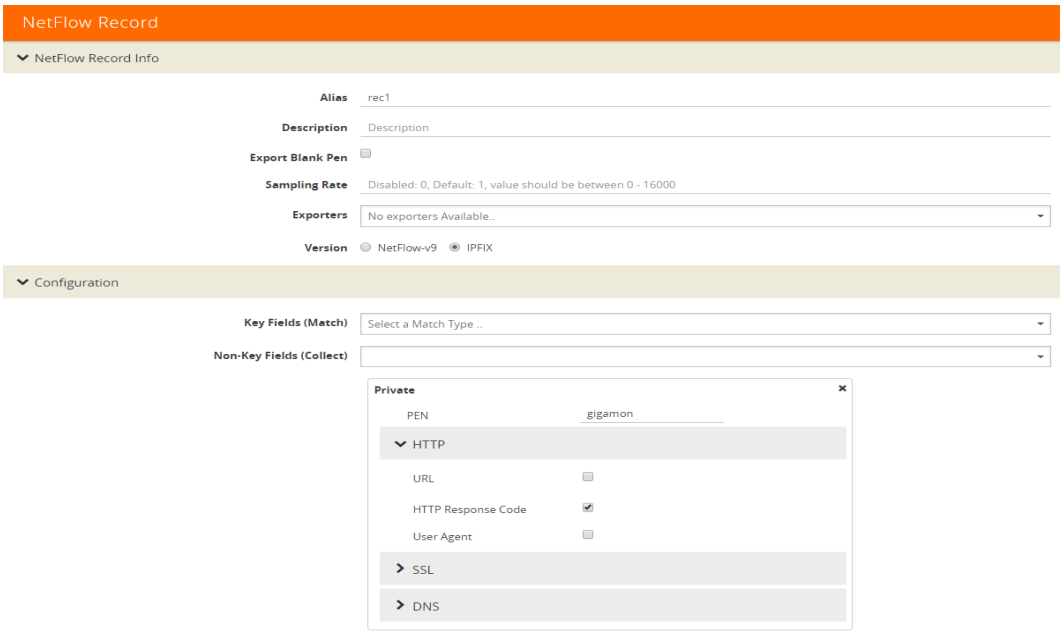

Figure 30-96: NetFlow Record Configuration for HTTP Response Code

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 30-96 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 30-96 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* on page 1134 and *NetFlow Exporter Statistics Definitions* on page 784.

Blank Records for IPFIX

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

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

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

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

Configure DNS Record for IPFIX

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

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

1. Go to **GigaSMART > NetFlow / IPFIX Generation > Records**.

2. Click **New**. Refer to Figure 30-97.

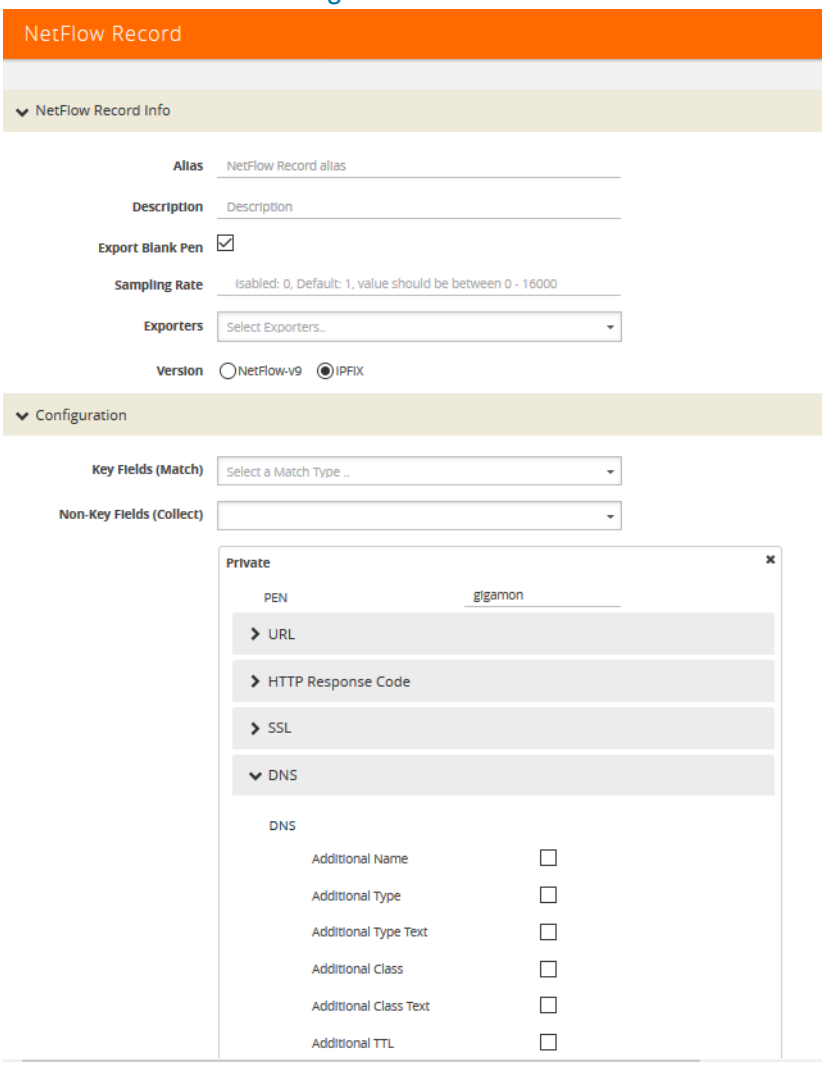

Figure 30-97: NetFlow Record Configuration for DNS

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 30-2.
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 30-2: 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.
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 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.

Method	For more information:
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 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 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 Out-of-Band SSL Decryption* on page 1169. The data is routed to the SSL application first and then to NetFlow. If de-duplication is also enabled, the data is routed from de-duplication to SSL, and then to NetFlow. The SSL application does not decrypt the data.

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

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

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

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

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

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

- Under SSL, select the SSL Certificate and SSL Server parameters and specify a width value in bytes. For example, **Certificate Issuer Common Name** and **Certificate Subject Common Name** as shown in Figure 30-98.

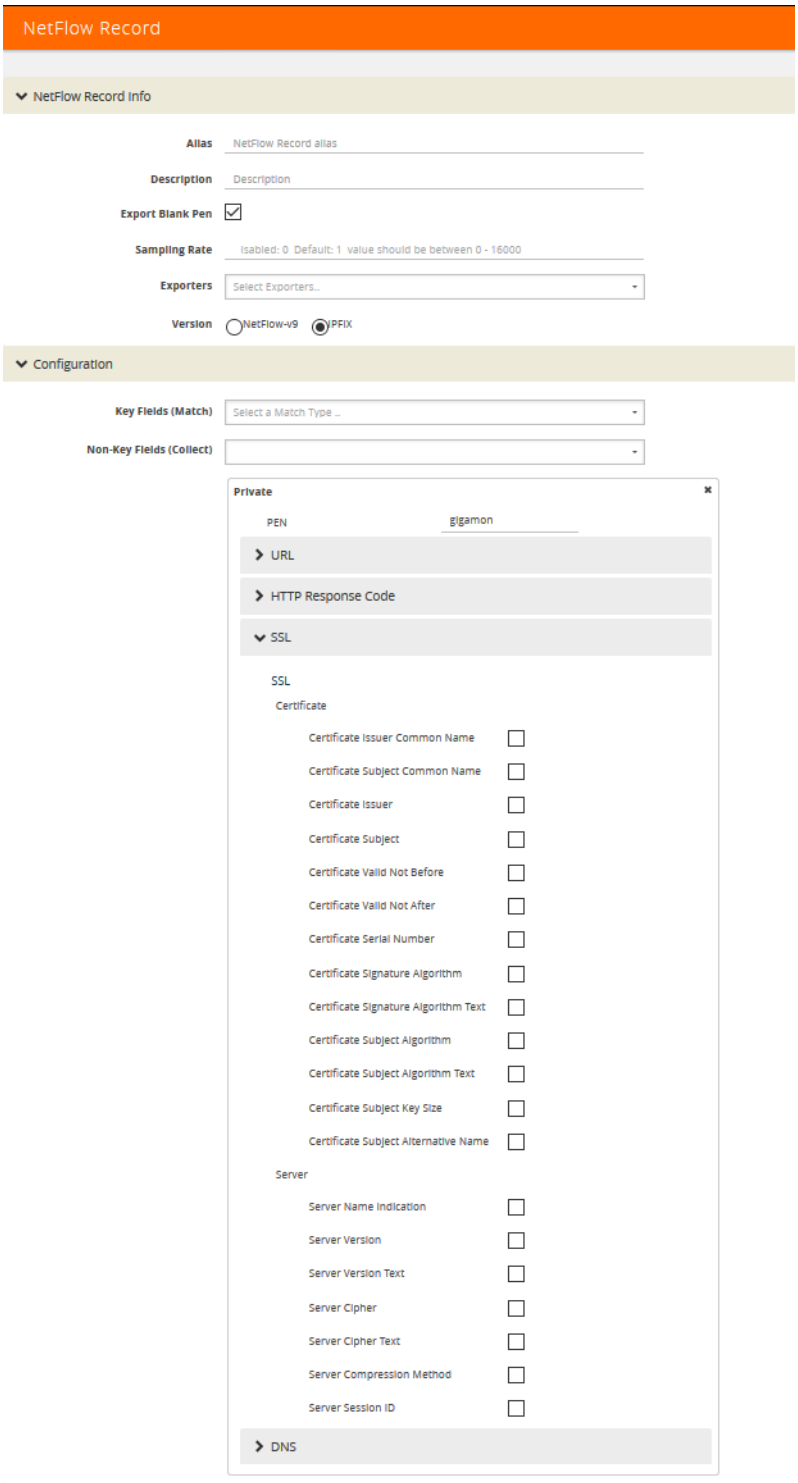

Figure 30-98: NetFlow Record Configuration for SSL

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 **Certificate Issuer**.
- **Certificate Subject Common Name**—the certificate subject common name, which is a subset of **Subject**.
- **Certificate Subject Alternative Name**—the subject alternative name, which allows identities to be bound to the subject of the certificate. This parameter is useful to detect if the certificate claims to sign for anything else and to detect anomalies such as certificates that claim to sign for a wildcard (*). The first subject alternative name present in the certificate is collected.
- **Certificate Valid Not Before** and **Certificate Valid Not After**—the date on which the certificate validity period begins and ends. The certificate validity period is the time interval during which the CA warrants that it will maintain information about the status of the certificate. It is expressed in universal time. The format is YYMMDDHHMMSSZ, where Z is Zulu time (GMT).
- **Certificate Valid Not Before Text** and **Certificate Valid Not After Text**—the text string that corresponds to the hexadecimal value of the date on which the certificate validity period begins and ends. The format is MMM DD HH:MM:SS YYYY GMT.
- **Certificate Serial Number**—the unique number for each certificate issued by a given CA. The issuer name and serial number identify a unique certificate. This parameter is useful to detect any certificate changes or substitutions.
- **Certificate Serial Number Text**—the text string that corresponds to the hexadecimal value of the unique number for each certificate issued by a given CA.

- **Certificate Signature Algorithm**—the identifier for the cryptographic algorithm used by the CA to sign the certificate, defined in ASN.1 format. This parameter is useful to detect servers that are not compliant with an organization's cryptographic standards.
- **Certificate Signature Algorithm Text**—the text string that corresponds to the hexadecimal value of the identifier for the cryptographic algorithm used by the CA to sign the certificate, defined in ASN.1 format. This parameter is useful to detect servers that are not compliant with an organization's cryptographic standards.
- **Certificate Subject Algorithm**—the subject public key algorithm used, defined in ASN.1 format, such as RSA or DSA.
- **Certificate Subject Algorithm Text**—The text string that corresponds to the hexadecimal value of the subject public key algorithm used, defined in ASN.1 format, such as RSA or DSA.
- **Certificate Subject Key Size**—the subject public key size.

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

SSL Server Parameters

The parameters supported for the SSL server are as follows:

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

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

Restrict Ports for NetFlow SSL Sessions

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

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

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

The following are the well-known ports:

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

Configure SSL Certificate and Server Parameters

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

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

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Records**.
2. Click **New**.
3. Enter an alias in the **Alias** field to identify this record. For example, ipfixrec.
4. Select **IPFIX**.
5. From the **Non-Key Field (Collect)** list, select **Private**.
6. Set **PEN** to gigamon. (This is the default.)
7. Under SSL, select any of the following:
    - Certificate issuer Common Name Width 30
    - Certificate Subject Common Name Width 40
    - Certificate issuer Width 150
    - Certificate Subject Width 120
    - Certificate Valid Not Before
    - Certificate Valid Not After
    - Certificate Serial Number
    - Certificate Signature Algorithm
    - Certificate Signature Algorithm Text
    - Certificate Subject Algorithm

- Certificate Subject Algorithm Text
- Certificate Subject Key Size
- Certificate Subject Alternative Name
- Server Name Indication Width 64
- Server Version
- Server Version Text
- Server Cipher
- Server Cipher Text
- Server Compression Method
- Server Session ID

8. Click **Save**.

Best Practices for Collecting SSL Metadata

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

1. From the device view, select **GigaSMART >** NetFlow Record / IPFIX Generation.
2. Select the record ipfixrec and click Edit. (This is the record configured in the previous section *Configure SSL Certificate and Server Parameters* on page 1096.)
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* on page 1096.)
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* on page 1096.)
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* on page 803 for descriptions of these statistics as well as to *GigaSMART Operations Statistics Definitions* on page 794.

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 **Sever Name Indication**, all packets up to the server hello packet must be received.
    - For the **Server Name 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 GigaSMART NetFlow Generation.

External collectors need to recognize GigaSMART NetFlow 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 GigaSMART NetFlow Generation
- with automatic validation, the recipient collector initiates an SNMP query requesting relevant information and GigaSMART NetFlow 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, GigaSMART NetFlow 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 anextension 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 metadat 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=2f433d55532f4f3d446967694365727420496e632f4f553d
7777772e6469676963
6572742e636f6d2f434e3d4469676943657274205348413220486967682041737375
616e636
52053657276 6572204341
GigamonMdataSslCertSubCommonName=*.stickyadstv.com
GigamonMdataSslSub=2f433d55532f53543d4e657720596f726b2f4c3d4e657720596
f726b2f4f3d4672656

5776865656c204d6564696120496e632f4f553d4672656576865656c2f434e3d2a2e7
37469636b796164737 4762e636f6d dst=10.50.22.59 src=38.106.34.118
spt=443
```

Configure NetFlow Generation

The following are the step for setting up a typical NetFlow Generation configuration with H-VUE:

- *Step 1: Configure GigaSMART Group*

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

Step 1: Configure GigaSMART Group

Configure a GigaSMART Group using the following steps. you will use this GigaSMART Group in *Step 6: Add NetFlow Monitor to GigaSMART Group* on page 1107, 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 or select an existing GigaSMART Group and click **Edit**.
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.

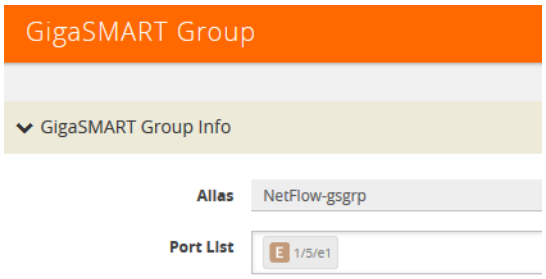

5. Click **Save**.

Notes:

- To use NetFlow, the GigaSMART Group can only contain one GigaSMART engine port.
- Only one NetFlow Generation Monitor can be configured per GigaSMART Group.

Step 2: Configure NetFlow Exporter

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

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

To configure the NetFlow Exporter, do the following:

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Exporters**.
2. Click **New**. The NetFlow Exporters page shown in Figure 30-99 displays.

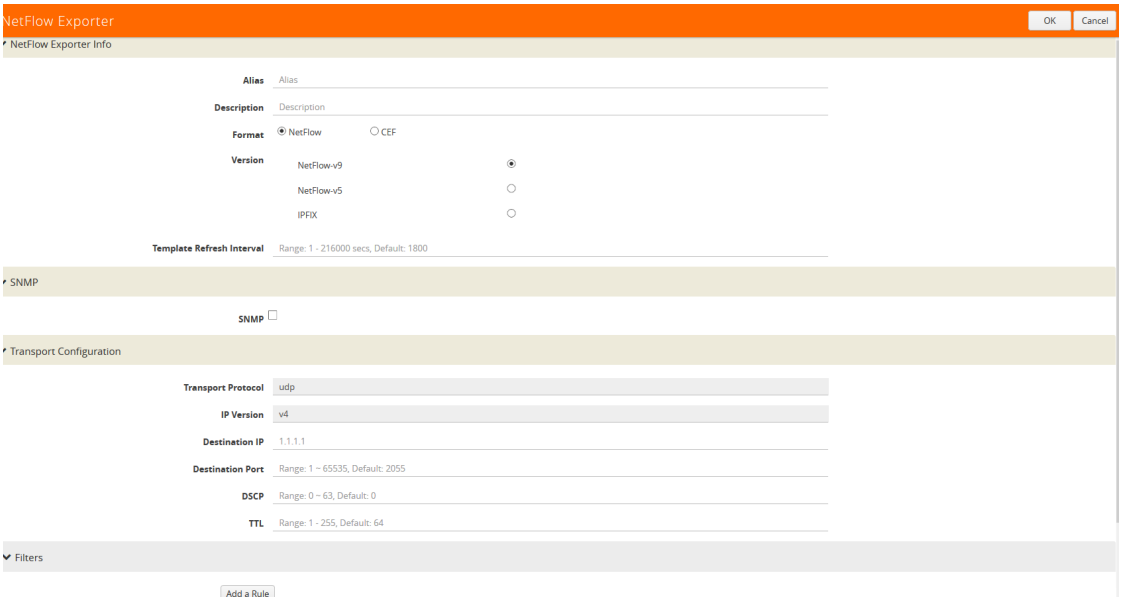

Figure 30-99: NetFlow Exporter Page

3. On the NetFlow Exporter page, enter the information for the exporter. Table 30-3 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 30-3: NetFlow Exporter Configuration 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. Default is set as v4. It cannot be changed.

Table 30-3: NetFlow Exporter Configuration Fields

Field	Description
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 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 **Comment** 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 GigaSMART Group* on page 1102.
   f. From the Exporters drop-down list, select the NetFlow exporter you created in *Step 2: Configure NetFlow Exporter* on page 1102.

Step 4: Configure NetFlow Record

Configure a NetFlow Generation Record that has the following:

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

To access GigaSMART within GigaVUE-FM, access a device that has been added to GigaVUE-FM from the GigaVUE-FM interface. **GigaSMART** appears in the navigation

pane of the device view on supported devices. Refer to the *Access GigaSMART from GigaVUE-FM* for details.

To configure the NetFlow Record, do the following:

1. From the device view, select **GigaSMART > NetFlow / IPFIX Generation > Records**.
2. Click **New**. The NetFlow record page shown in Figure 30-100 displays.

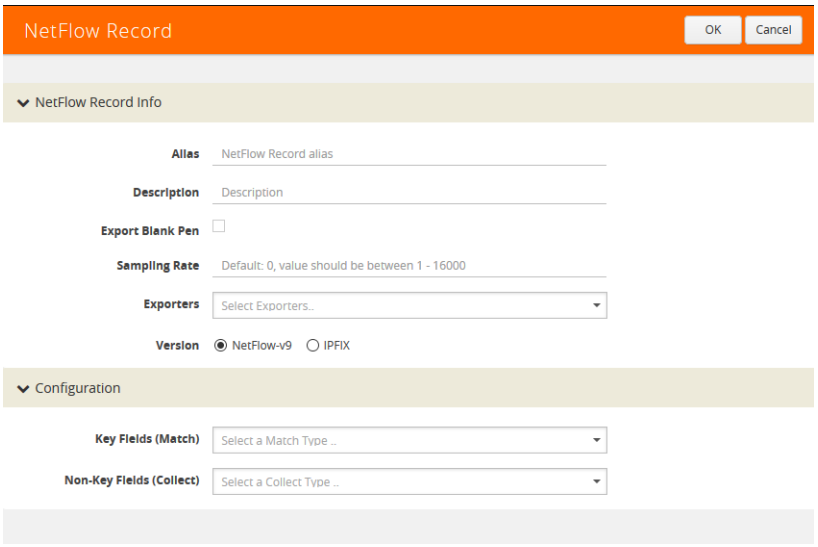

Figure 30-100: NetFlow Record Page

3. On the NetFlow Record page, do the following:
   a. Specify the NetFlow Record information:
      - Enter an **alias** to help identify the record
      - Enter a **Description** (optional)
      - Enter the **Sampling Rate** that you want
      - Select the Exporter that you want from the **Exporters** menu
      - Select the **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.

      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 *Sampled NetFlow Data* on page 1081. The **Sampling Rate** is disabled by 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.

   b. Specify the Configuration:

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.

For details about the match and collect parameters, refer to *NetFlow Generation Match/Key and Collect/Non-Key Elements* on page 1138

Step 5: Configure 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 shown in Figure 30-101 displays.

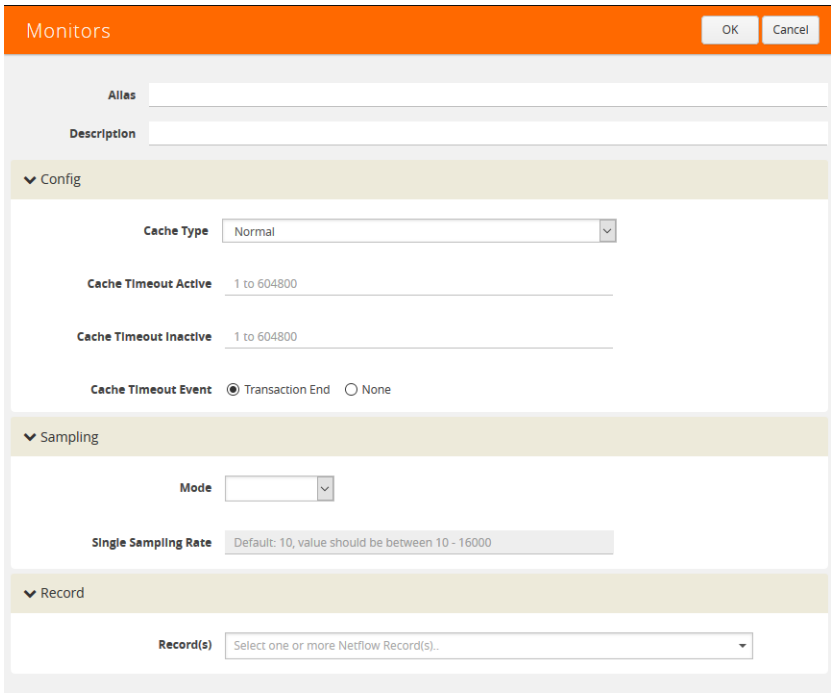

Figure 30-101: NetFlow Monitors Screen

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 30-4.
   d. Configure the **Sampling** parameters. Refer to Table 30-4.
   e. Select the **Record** that you configured in *Step 4: Configure NetFlow Record* on page 1104.

4. Click **Save**.

Table 30-4: 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 *Sampled NetFlow Data* on page 1081.

Step 6: Add NetFlow Monitor to GigaSMART Group

Return to the GigaSMART Group configured in *Step 1: Configure GigaSMART Group* on page 1102and set the NetFlow Monitor to the monitor created in *Step 2: Configure NetFlow Exporter* on page 1102.

1. From the device view, select **GigaSMART > GigaSMART Groups > GigaSMART Groups**.
2. Select the GigaSMART Group configured in *Step 1: Configure GigaSMART Group*, and then click **Edit** a shown in the following figure.

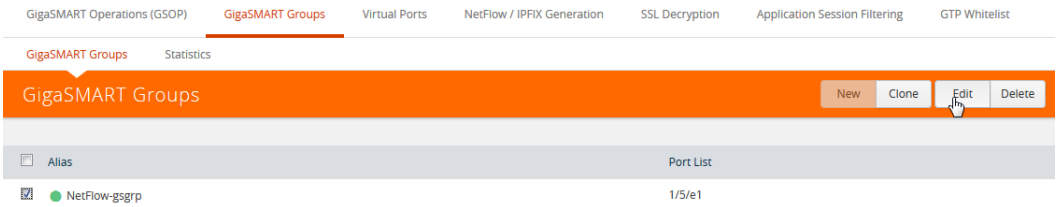

3. Under GigaSMART Parameters, go to NetFlow. Click in the **Monitor** field and select the NetFlow monitor configured in *Step 5: Configure NetFlow Monitor* as shown in the following figure.

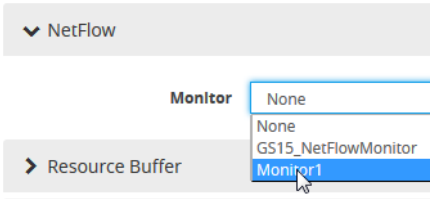

4. Click **Save**.

Step 7: Configure GigaSMART Operation

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

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

To configure the GigaSMART Operation, do the following:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP)**.
2. Click **New**. The GigaSMART Operations (GSOP) page displays. (Refer to Figure 30-102.)
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 GigaSMART Group*.
   c. In the **GigaSMART Operations (GSOP)** field, select **NetFlow**. The NetFlow GigaSMART Operation is enabled by default as shown in Figure 30-102.

Figure 30-102: 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* on page 485 and *Manage Maps* on page 518.

Notes:

- For a single NetFlow GigaSMART Operation, make sure that you create a Regular By Rule map. When combining with APF or Deduplication, 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 Deduplication, 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 IP Interface* on page 1104

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 shown in Figure 30-103 displays.

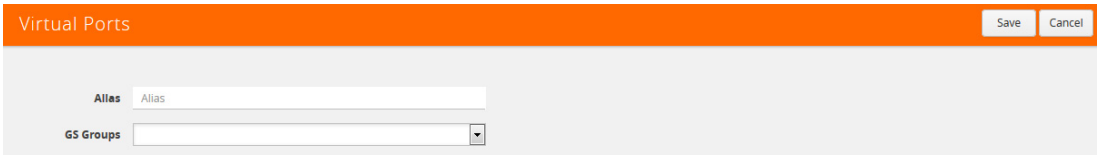

Figure 30-103: Virtual Ports Page

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 GigaSMART Group* on page 1102.
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 GigaSMART Operation* on page 1108.
   d. Click **Add a Rule** to add the rules needed for the map.
4. Click **Save**.

Configure NetFlow Generation Examples

The following sections provide examples of NetFlow Generation. Refer to the following:

- *Example 1: NetFlow Generation Configuration* on page 1110
- *Example 2: NetFlow Generation Configuration* on page 1116

Example 1: NetFlow Generation Configuration

In Example 1, the steps set up a typical NetFlow Generation configuration.

Ex 1, Step 1: Configure GigaSMART Group

Configure a GigaSMART group and associate it with a GigaSMART engine port.

NOTE: To use NetFlow, the gsgroup can only contain one GigaSMART engine port.

Step	Description	UI Steps
1.	Configure the GigaSMART Group	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type gsgp2 in the Alias field to identify this GigaSMART Group. **d.** Click in the Port List field and select an engine port. For example, 1/5/e1. **e.** Click Save.
2.	Display the information about the GigaSMART Group.	• On the GigaSMART Group page, click on the alias gsgp2 to display the GigaSMART Group Quick View.

Ex 1, Step 2: Configure NetFlow Generation Exporter

Configure one or more NetFlow Generation Exporters. There can be up to 6 NetFlow Generation Exporters for each NetFlow Generation Monitor.

The following parameters apply to this example:

- **Transport Protocol** — the UDP port of the collector.
- **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. If no version is specified, version 9 is the default.
- **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 for the NetFlow exporter.
- **Destination IP** — the IP address of the NetFlow/IPFIX collector.
- **DSCP** — the DSCP priority of the packet.
- **TTL** — the Time to Live of the packet.

Step	Description	UI Steps
1.	Configure the exporter. The exporter (exp4) will be used in *Ex 1, Step 8: Configure Mapping Rules to Filter Packets* on page 1115.	**a.** From the device view, select GigaSMART > NetFlow / IPFIX Generation > Exporters. **b.** Click New. **c.** Set NetFlow Exporter Info as follows: Alias: exp4. Version: IPFIX Template Refresh Interval: 60 **d.** Enable SNMP. **e.** Set Transport Configuration. Transport Protocol: UDP IP Version: v4 Destination IP: 20.20.20.20 Destination Port: 2055 DSCP: 10 TTL: 64 **f.** Click Save.
2.	Display the exporter configuration.	**a.** From the device view, select GigaSMART > NetFlow / IPFIX Generation > Exporters. **b.** Click on the row with the alias exp4 to display the details in the Exporter Quick View.

Ex 1, Step 3: Configure IP Interface with Tool Port and Associate with the Exporter

Create an IP interface with tool port. You must associate this IP interface with the NetFlow Exporter you configured in *Ex 1, Step 2: Configure NetFlow Generation Exporter* on page 1110.

Step	Description	UI Steps
1.	Identify the collector port and configure it as a tool port, where the NetFlow collector will be connected.	**a.** Identify the port to use as an IP interface with tool port. For example, 1/1/g3. **b.** Select Ports > All Ports. **c.** Use the Quick Port Editor to configure port 1/1/g3 as a tool port.
2.	Configure the IP interface. The IP address is for the NetFlow interface.	**a.** Select Ports > IP Interfaces. **b.** Click New. **c.** In the Alias and Comment fields, enter the name and description of the IP interface. **d.** From the Port drop-down list, select port 1/1/g3. **e.** Enter 10.10.10.10 in the IP Address field. **f.** Enter 255.255.255.255 in the Mask field **g.** Enter 10.10.10.1 in the Gateway field. **h.** Enter 1500 in the MTU field. **i.** Select the GigaSMART Group created in *Ex 1, Step 1: Configure GigaSMART Group*. **j.** Click Save.

Step	Description	UI Steps
3.	Display the IP interface configuration.	**a.** Select Ports > IP Interfaces. **b.** Click on the row for port 1/1/g3 to display the IP interface details in a Quick View.

Ex 1, Step 4: Configure Record

Configure one or more NetFlow Generation Records, which have the following:

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

NOTE: NetFlow v9 and IPFIX let you configure Match/Key and Collect/Non-Key elements. For details refer to *NetFlow Generation Match/Key and Collect/Non-Key Elements* on page 1138.

The following NetFlow Record parameters apply to this example:

- **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. If no version is specified, version 9 is the default.
- **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.

In this example, the IP source and destination address on the incoming traffic is used to identify network traffic between the unique pair of source and destination addresses. Once unique flows are identified, the following sample parameters are collected and exported for each flow:

- IP source and destination address
- Total number of packets received that match the unique flows
- IPv4 protocol
- Transport source and destination ports
- Input and output interface, plus interface name
- Packet URL
- DNS query name
- Timestamp for the beginning and end of flow

NOTE: Configure the NetFlow version prior to configuring the match and collect parameters because the subsequent parameters depend on the NetFlow version configured. If no version is specified, the version 9 is the default (NetFlow-v9).

Task	Description	UI Steps
1.	Configure the record. The NetFlow version must be the same as the NetFlow version specified in *Ex 1, Step 2: Configure NetFlow Generation Exporter* on page 1110. The record (rec2) will be used in *Ex 1, Step 5: Configure Monitor* on page 1114.	**a.** From the device view, select GigaSMART > NetFlow / IPFIX Generation > Records. **b.** Click New. **c.** Type rec2 in the Alias field. **d.** Select IPFIX for Version. **e.** Click in the Key Fields (Match) field and select IPv4. • Select Source, and then select Address • Select Destination, and then select Address **f.** Click in Non-Key Fields (Collect) and Select IPv4. • Select Source, and then select Address • Select Destination, and then select Address **g.** Click in Non-Key Fields (Collect) and select Counter. • Select Bytes • Select Packets **h.** Click in Non-Key Fields (Collect) and select Transport. • Select Source Port **i.** Click in Non-Key Fields (Collect) and select Interface. • Select Output **j.** Click in Non-Key Fields (Collect) and select Private. • PEN is gigamon • Select URL **k.** Click in Non-Key Fields (Collect) and select Timestamp. • Select Sys-Uptime First • Select Sys-Uptime Last **l.** Click Save.
2.	Configure a second record. The NetFlow version must be the same as the NetFlow version specified in *Ex 1, Step 2: Configure NetFlow Generation Exporter* on page 1110. The match fields must be the same as in Step 1. Each record must have the same match fields but differing collect fields. The record (rec3) will be used in *Ex 1, Step 5: Configure Monitor* on page 1114.	**a.** From the device view, select GigaSMART > NetFlow / IPFIX Generation > Records. **b.** Click New. **c.** Type rec3 in the Alias field. **d.** Select IPFIX for Version. **e.** Click in the Key Fields (Match) field and select IPv4. • Select Source, and then select Address • Select Destination, and then select Address **f.** Click in Non-Key Fields (Collect) and Select IPv4. • Select Source, and then select Address • Select Destination, and then select Address **g.** Click in Non-Key Fields (Collect) and select Private. • PEN is gigamon • Select URL • Select Query Name under DNS **h.** Click in Non-Key Fields (Collect) and select Timestamp. • Select Sys-Uptime First • Select Sys-Uptime Last **i.** Click Save.

Task	Description	UI Steps
3.	Display the record configuration.	• Click on the row with the alias rec2 to display the Record Quick View.

Ex 1, Step 5: Configure Monitor

Configure a NetFlow Generation Monitor and associate the NetFlow Generation Record to the specified NetFlow Generation Monitor.

The following parameters show the binding of the records. The parameters also define the cache (holding statistics for unique flows).

- **Cache Timeout Event** — Set to **Transaction End**. This applies to the TCP flow. The flow is “flushed out” to the Exporter after detecting a FIN or RST.
- **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.
- **Sampling** — Enables sampled NetFlow and defines the sampling rate by specifying a number for 1 in N, where N is the packet count from 10 to 16000.
- **Records** — Records generated for the flow are defined in the record and are stored in the internal cache.

Step	Description	UI Steps
1.	Configure the monitor. The monitor (mon2) will be used in *Ex 1, Step 8: Configure Mapping Rules to Filter Packets* on page 1115. The records (rec2 and rec3) were created in *Ex 1, Step 4: Configure Record* on page 1112. In this example, NetFlow sampling is enabled. The sampling rate is 1 in 1024.	**a.** From the device view, select GigaSMART > NetFlow > Monitors. **b.** Click New. **c.** Type mon2 in the Alias field. **d.** Define Config Cash Type: Normal Cash Timeout Event: Transaction End Cash Timeout Active: 60 Cash Timeout Inactive: 15 Sampling: 1024 **e.** Select Record(s) Select rec2 Select rec3 **f.** Click Save.
2.	Display the monitor configuration.	• Click on the row with the alias mon2 to display the Monitor Quick View.

Ex 1, Step 6: Add Monitor to GigaSMART Group

Add the monitor created in *Ex 1, Step 5: Configure Monitor* on page 1114 to the GigaSMART Group created in step *Ex 1, Step 1: Configure GigaSMART Group* on page 1110.

NOTE: Only one NetFlow Generation Monitor can be configured per gsgroup.

Step	Description	UI Steps
1.	Select the GigaSMART Group configured in *Ex 1, Step 1: Configure GigaSMART Group*.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** On the GigaSMART Groups page, select gsgp2. **c.** Click Edit. **d.** Click in the Monitor field and select the NetFlow Monitor created in *Ex 1, Step 5: Configure Monitor* **e.** Click Save.
2.	Display the GigaSMART Group information.	**a.** From the device view, select GigaSMART Groups > GigaSMART Groups. **b.** On the GigaSMART Groups page, click on the alias to display the information in a Quick View.

Ex 1, Step 7: Configure GigaSMART Operation

Define a GigaSMART Operation to enable NetFlow Generation as follows:

Step	Description	UI Steps
1.	Configure the GigaSMART Operation and associate it with the GigaSMART Group created in *Ex 1, Step 1: Configure GigaSMART Group*	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Type an alias in the Alias field to identify this operation. For example, gsop2. **d.** For GigaSMART Groups, select gsgp2. **e.** For GigaSMART Operations (GSOP), select NetFlow. **f.** Click Save.
2.	Display the configuration GigaSMART Operation.	• On the GigaSMART Operations page, click on the alias gsop2 to open the GigaSMART Operation Quick View.

Ex 1, 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, as follows:

Step	Description	UI Steps
1.	Configure the map. (This is a first level map.)	a. Select Maps > Maps > Maps. b. Click New. c. Type map3 in the Alias field. d. Select Regular for Type and By Rule for Subtype. e. Select the network port for Source. For example 1/1/x11 f. Select 1/1/g3 for Destination. This is the IP interface configured in *Ex 1, Step 3: Configure IP Interface with Tool Port and Associate with the Exporter* on page 1111. g. Click Add a Rule. Select IP Version. Set Version to 4. h. Click Save.
2.	Display the map configuration.	• Click on the row for map3 to display the Quick View for the map.

Example 2: NetFlow Generation Configuration

Starting in software version 4.2, NetFlow exporters can filter NetFlow records. The filtered NetFlow records are sent to the collectors.

In Example 2, there are three exporters, with filtering configured on two of them. Because the second exporter does not have any filtering configured, all the records are sent to the collector. In this example, there are also two tunnels and two maps. Both maps are first level maps.

Ex 2, Step 1: Configure GigaSMART Group

Configure a GigaSMART group and associate it with a GigaSMART engine port, as follows:

Step	Description	UI Steps
1.	Configure the GigaSMART Group	a. From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. b. Click New. c. Type grp in the Alias field to identify this GigaSMART Group. d. Click in the Port List field and select an engine port. For example, 1/5/e1. e. Click Save.
2.	Display the information about the GigaSMART Group.	• On the GigaSMART Group page, click on the alias gsgp2 to display the GigaSMART Group Quick View.

Ex 2, Step 2: Configure Exporters

Configure one or more NetFlow Generation Exporter(s), as follows:

Task	Description	UI Steps
1.	Configure the first exporter.	**a.** From the device view, select GigaSMART > NetFlow / IPFIX Generation > Exporters. **b.** Click New. **c.** Set NetFlow Exporter Info • Alias: exp1. • Version: IPFIX • Template Refresh Interval: Use the default **d.** Select SNMP **e.** Set Transport Configuration. • Transport Protocol: UDP • IP Version: v4 • Destination IP: 1.1.1.1 • Destination Port: 2055 • DSCP: 0 • TTL: 64 **f.** Add Filter Rule 1 • Click Add a Rule • Select IPv4 Destination • Enter 1.1.1.1 for the IPv4 Address • Enter 255.255.255.248 for the Net Mask • Select 0. **g.** Add Filter Rule 2 • Click Add a Rule. • Select VLAN. • Enter 1 for min. • Set Position to 0. **h.** Add Filter Rule 3 • Click Add a Rule. • Select Port Destination. • Enter 1 for min. • Set Position to 0. **i.** Click Save.

Task	Description	UI Steps
2.	Configure the second exporter.	**a.** From the device view, select GigaSMART > NetFlow > Exporters. **b.** Click New. **c.** Set NetFlow Exporter Info • Alias: exp2. • Version: IPFIX • Template Refresh Interval: Use the default **d.** Select SNMP **e.** Set Transport Configuration. • Transport Protocol: UDP • IP Version: v4 • Destination IP: 2.2.2.2 • Destination Port: 2055 • DSCP: 0 • TTL: 64 **f.** Click Save.

Task	Description	UI Steps
3.	Configure the third exporter.	**a.** From the device view, select GigaSMART > NetFlow > Exporters. **b.** Click New. **c.** Set NetFlow Exporter Info • Alias: exp1. • Version: IPFIX • Template Refresh Interval: Use the default **d.** Select SNMP **e.** Set Transport Configuration. • Transport Protocol: UDP • IP Version: v4 • Destination IP: 3.3.3.3 • Destination Port: 2055 • DSCP: 0 • TTL: 64 **f.** Add Filter Rule 1 • Click Add a Rule • Select IPv4 Destination • Enter 3.3.3.3 for the IPv4 Address • Enter 255.255.255.248 for the Net Mask • Select 0. **g.** Add Filter Rule 2 • Click Add a Rule. • Select VLAN. • Enter 3 for min. • Set Position to 0. **h.** Add Filter Rule 3 • Click Add a Rule. • Select Port Destination. • Enter 3 for min. • Set Position to 0. **i.** Click Save.

Ex 2, Step 3: Configure IP Interfaces with Tool Ports and Associate with the Exporters

In this example, create two IP interfaces with tool ports. You must associate these IP interfaces with the NetFlow Exporters in *Ex 2, Step 2: Configure Exporters* on page 1117.

Task	Description	UI Steps
1.	Identify the collector ports and configure them as a tool ports, where the NetFlow collector will be connected	**a.** Identify the ports to use as an IP interface with tool port. For example, 1/1/g1 and 1/1/g2. **b.** Select Ports > All Ports. **c.** Use the Quick Port Editor to configure port 1/1/g1 and 1/1/g2 as tool ports. **d.** Click OK.

Task	Description	UI Steps
2.	Configure the first IP interface.	**a.** Select Ports > IP Interfaces. **b.** Click New. **c.** In the Alias and Comment fields, enter the name and description for the IP interface. **d.** Select the tool port 1/1/g1 from the Port list. **e.** Set the IP address to 1.1.1.1. **f.** Set the IP Mask to 255.255.0.0. **g.** Set the Gateway address to 3.3.3.3. **h.** Set MTU to 2000. **i.** Select the GigaSMART Group grp from the GigaSMART Group list. **j.** From the Exporters drop-down list, select the first Exporter that you configured. **k.** Click OK.
3.	Configure the second IP interface.	**a.** Select Ports > IP Interfaces. **b.** Click New. **c.** In the Alias and Comment fields, enter the name and description of the IP interface. **d.** Select the tool port 1/1/g2 from the Port list. **e.** Set the IP address to 2.2.2.2. **f.** Set the IP Mask to 255.255.255.248. **g.** Set the Gateway address to 4.4.4.4. **h.** Set MTU to 2000. **i.** Select the GigaSMART Group grp from the GigaSMART Group list. **j.** From the Exporters drop-down list, select the third Exporter that you configured. **k.** Click Save.
4.	Display the IP interfaces configurations.	• Select Ports > IP Interfaces. On the IP Interfaces page, click on the IP interfaces to display the details.

Ex 2, Step 4: Configure Record

Configure a NetFlow Generation Record, as follows:

Task	Description	UI Steps
1.	Configure the record.	**a.** From the device view, select GigaSMART > NetFlow / IPFIX Generation > Records. **b.** Click New. **c.** Type rec1 in the Alias field. **d.** Select IPFIX. **e.** Configure Key Fields (Match) • IPv4 TLL • IPv6 Traffic Class **f.** Configure Non-Key Fields (Collect) • IPv6 Transport UDP Source Port • IPv6 Transport • Transport TCP Source Port **g.** Click Save.
2.	Display the record configuration.	• Click on the row with the Alias rec1 to display the Record Quick View.

Ex 2, Step 5: Configure Monitor

Configure a NetFlow Generation Monitor and associate the NetFlow Generation Record to the specified NetFlow Generation Monitor, as follows:

Task	Description	UI Step
1.	Configure the monitor. **NOTE:** In this example, NetFlow sampling is not enabled.	**a.** From the device view, select GigaSMART > NetFlow / IPFIX Generation > Records. **b.** Click New. **c.** Type mon1 in the Alias field. **d.** Select rec1 from the Record list. **e.** Click Save.
2.	Display the monitor configuration.	• Click on the row with the Alias mon1 to display the Monitor Quick View.

Ex 2, Step 6: Add Monitor to GigaSMART Group

Add the monitor created in *Ex 1, Step 5: Configure Monitor* to the GigaSMART Group created in step *Ex 1, Step 1: Configure GigaSMART Group*.

NOTE: Only one NetFlow Generation Monitor can be configured per gsgroup.

Step	Description	UI Steps
1.	Select the GigaSMART Group configured in *Ex 2, Step 1: Configure GigaSMART Group*.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** On the GigaSMART Groups page, select gsgp2. **c.** Click Edit. **d.** Click in the Monitor field and select the NetFlow Monitor created in *Ex 2, Step 5: Configure Monitor* **e.** Click Save.
2.	Display the GigaSMART Group information.	**a.** From the device view, select GigaSMART Groups > GigaSMART Groups. **b.** On the GigaSMART Groups page, click on the alias to display the information in a Quick View.

Ex 2, Step 7: Configure GigaSMART Operation

Define a GigaSMART Operation to enable NetFlow Generation, as follows:

Step	Description	UI Steps
1.	Configure the GigaSMART Operation and associate it with the GigaSMART Group created in *Ex 2, Step 1: Configure GigaSMART Group*	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Type an alias in the Alias field to identify this operation. For example, gsop2. **d.** For GigaSMART Groups, select grp. **e.** For GigaSMART Operations (GSOP), select NetFlow. **f.** Click Save.
2.	Display the configuration GigaSMART Operation.	• On the GigaSMART Operations page, click on the alias gsop2 to open the GigaSMART Operation Quick View.

Ex 2, Step 8: Configure Mapping Rules to Filter Packets

To add flow mapping rules to filter packets that are needed to run NetFlow, configure maps and associate the maps to the IP interfaces with tool ports, as follows:

Task	Description	UI Steps
1.	Configure the first map. (This is a first level 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 Subtype. **e.** Set network ports 1/1/x1 and 1/1/x2 as the Source. **f.** Set tool port 1/1/g1 as the Destination. **g.** From the device view, select GigaSMART Group for this map from the GSOP list. For example, gsop1. **h.** Click Add a Rule. **i.** Select IP Version. **j.** Set Version to v4. **k.** Click Save.
2.	Configure the second map. (This is also a first level map.)	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map2 in the Alias field. **d.** Select Regular for Type and By Rule for Subtype. **e.** Set network ports 1/1/x3 and 1/1/x4 as the Source. **f.** Set tool port 1/1/g1 as the Destination. **g.** From the device view, select GigaSMART Group for this map from the GSOP list. For example, gsop1. **h.** Click Add a Rule. **i.** Select Pass **j.** Select IP Version. **k.** Set Version to v4. **l.** Click Save.
3.	Display the map configuration.	**a.** Select Table View. **b.** Select alias of the port to see the Quick View for the port.

Example 3: NetFlow Generation Configuration

Starting in software version 4.3.01, NetFlow supports both first level and second level maps. In Example 3, there are two maps. However, unlike Example 2, which has two first level maps, in this example, one map is a first level map and the other is a second level map. A virtual port is configured that directs traffic to the second level map.

The configuration of the GigaSMART operation in Example 3 differs from Example 1 and Example 2. The GigaSMART Operation sends traffic to APF first, and then to NetFlow.

In the first level map, the traffic matching the rule is sent to the virtual port. The same traffic is also sent to two tool ports (2/1/g2 and 2/1/g3).

In the second level map, the traffic from the virtual port matching the gsrule is sent to NetFlow and then to the IP interface with tool port, 2/1/g7.

Ex 3, Step 1: Configure GigaSMART Group

Configure a GigaSMART group and associate it with a GigaSMART engine port, as follows:

Step	Description	UI Steps
1.	Configure the GigaSMART Group	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type grp in the Alias field to identify this GigaSMART Group. **d.** Click in the Port List field and select an engine port. For example, 1/5/e1. **e.** Click Save.
2.	Display the information about the GigaSMART Group.	• On the GigaSMART Group page, click on the alias gsgp2 to display the GigaSMART Group Quick View.

Ex 3, Step 2: Configure IP Interface and Associate with the Exporter

Configure the IP interface and associate with the NetFlow Generation Exporter, as follows:

Task	Description	UI Steps
1.	Configure the IP interface.	**a.** Select Ports > IP Interfaces **b.** Click New. **c.** In the Alias and Comment fields, enter the name and description of the IP interface. **d.** Click in the Port field and select the tool port 1/1/g7. **e.** Set the IP address to 10.115.9.5. **f.** Set the IP Mask to 255.255.248.0. **g.** Set the Gateway address to 10.115.8.1. **h.** Set MTU to 1500. **i.** Select the GigaSMART Group created in *Ex 3, Step 1: Configure GigaSMART Group* on page 1124. **j.** From the Exporters drop-down list, select exp1. **k.** Click OK.

Ex 3, Step 3: Configure IP Interface with Tool Port

Create an IP interface with tool port.

Step	Description	UI Steps
1.	Identify the collector port and configure it as a tool port, where the NetFlow collector will be connected.	**a.** Identify the port to use as an IP interface with tool port. For example, 1/1/g7. **b.** Select Ports > All Ports. **c.** Use the Quick Port Editor to configure port 1/1/g3 as tool port.
2.	Configure the IP interface. The IP address is for the NetFlow interface.	**a.** Select Ports > IP Interfaces. **b.** Click New. **c.** In the Alias and Comment fields, enter the name and description of the IP interface. **d.** Click in the Port field and select port 1/1/g3. **e.** Enter 10.115.9.5 in the IP Address field. **f.** Enter 255.255.255.255 in the Mask field **g.** Enter 10.115.9.1 in the Gateway field. **h.** Enter 1500 in the MTU field. **i.** Select the GigaSMART Group created in *Ex 1, Step 1: Configure GigaSMART Group*. **j.** Click Save.
3.	Display the IP interface configuration.	**a.** Select Ports > IP Interfaces. **b.** Click on the row for port 1/1/g7 to display the IP Interfaces details in a Quick View.

Ex 3, Step 4: Configure Record

Configure a NetFlow Generation Record, as follows:

Task	Description	UI Steps
1.	Configure the record.	**a.** From the device view, select GigaSMART > NetFlow > Records. **b.** Click New. **c.** Type rec1 in the Alias field. **d.** Select IPFIX. **e.** Configure Key Fields (Match) Select TOS. Select IPv4 and enable Source. Select Address. **f.** Configure Non-Key Fields (Collect) Select IPv4 and enable Source. Select Interface. **g.** Click Save.

Ex 3, Step 5: Configure Monitor

Configure a NetFlow Generation Monitor and associate the NetFlow Generation Record to the specified NetFlow Generation Monitor, as follows:

Task	Description	UI Steps
1.	Configure the monitor.	**a.** From the device view, select GigaSMART > NetFlow > Records. **b.** Click New. **c.** Type mon1 in the Alias field. **d.** Select rec1 from the Record list. **e.** Click Save.

Ex 3, Step 6: Add Monitor to GigaSMART Group

Add the monitor created in *Ex 3, Step 5: Configure Monitor* to the GigaSMART Group created in step *Ex 3, Step 1: Configure GigaSMART Group*.

NOTE: Only one NetFlow Generation Monitor can be configured per gsgroup.

Step	Description	UI Steps
1.	Select the GigaSMART Group configured in *Ex 3, Step 1: Configure GigaSMART Group*.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** On the GigaSMART Groups page, select grp. **c.** Click Edit. **d.** Click in the Monitor field and select the NetFlow Monitor created in *Ex 3, Step 5: Configure Monitor* **e.** Click Save.
2.	Display the GigaSMART Group information.	**a.** From the device view, select GigaSMART Groups > GigaSMART Groups. **b.** On the GigaSMART Groups page, click on the alias to display the information in a Quick View.

Ex 3, Step 7: Configure GigaSMART Operation

Define a GigaSMART Operation to enable NetFlow Generation, as follows:

Step	Description	UI Steps
1.	Configure the GigaSMART Operation and associate it with the GigaSMART Group created in *Ex 3, Step 1: Configure GigaSMART Group* on page 1124	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Type an alias in the Alias field to identify this operation. For example, gsop2. **d.** For GigaSMART Groups, select grp. **e.** For GigaSMART Operations (GSOP), select NetFlow. **f.** Click Save.

Step	Description	UI Steps
2.	Display the configuration GigaSMART Operation.	• On the GigaSMART Operations page, click on the alias gsop2 to open the GigaSMART Operation Quick View.

Ex 3, Step 8: Configure Virtual Port

For the second level map that you will create in *Ex 3, Step 9: Configure Mapping Rules to Filter Packets* on page 1127, you will need to create a virtual port.

Configure a virtual port and associate it with the GigaSMART group, as follows:

Step	Description	UI Steps
1.	Configure the virtual port for the second level map.	**a.** From the device view, select GigaSMART > Virtual Ports. **b.** On the GigaSMART Groups page, select grp. **c.** Click New. **d.** Type an alias for the virtual port in the Alias field. For example, vp1. **e.** Click in the GigaSMART Group filed and select the GigaSMART Group created in *Ex 3, Step 1: Configure GigaSMART Group*. **f.** Click Save.
2.	Display the virtual port information.	**a.** From the device view, select GigaSMART Groups > Virtual Ports. **b.** On the Virtual Ports page, click on the virtual port alias to display the information in a Quick View.

Ex 3, Step 9: Configure Mapping Rules to Filter Packets

To add flow mapping rules to filter packets that are needed to run NetFlow, configure maps and associate the maps to the IP interface with tool port, as follows:

Task	Description	UI Steps
1.	Configure the first map. (This is a first level map.)	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map1 in the Alias field. **d.** Select First Level for Type and By Rule for Subtype. **e.** Set network port 1/1/g1 as the Source. **f.** Set virtual port v1 and tool ports 21/1/g2 and 2/1/g3 as the Destination. **g.** Click Add a Rule. **h.** Select Pass **i.** Select MAC Destination. **j.** Set the MAC address and mask to 00:00:00:00:00:00 **k.** Click Save.

Task	Description	UI Steps
2.	Configure the second map. (This is a second level map.)	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map2 in the Alias field. **d.** Select Second Level for Type and By Rule for Subtype. **e.** Set virtual port v1 as the Source. **f.** Set tool port v1 2/1/g7 as the Destination. **g.** Select gsop_apf_netflow from the GSOP list. **h.** Click Add a Rule. **i.** Select Pass **j.** Select MAC Destination. **k.** Set the MAC address and mask to 00:00:00:00:00:00 **l.** Click Save.

Example 4: NetFlow Generation Configuration

Starting in software version 4.3.01, NetFlow supports both first level and second level maps. In Example 4, there are three maps. One map is a first level map and the other two are second level maps. Two virtual ports are configured that direct traffic to the second level maps.

Two GigaSMART operations are configured. One gsop sends traffic to masking. The other gsop sends traffic to APF and then to NetFlow.

In the first level map, the traffic matching the rule is sent to two virtual ports. The same traffic is also sent to a tool port (11/1/g3).

In the first second level map, the traffic from the first virtual port, vp1, that matches the gsrule, is sent to masking and then to the tool port 11/1/g2.

In the next second level map, the traffic from the second virtual port, vp2, that matches the gsrule, is sent to NetFlow and then to the IP interface with tool port, 11/1/g4.

Ex 4, Step 1: Configure GigaSMART Group

Configure two GigaSMART groups and associate them with a GigaSMART engine port, as follows:

Step	Description	UI Steps
1.	Configure the first GigaSMART Group.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type an grp1 in the Alias field to identify this GigaSMART Group. **d.** Click in the Port List field and select an engine port. For example, 1/5/e1. **e.** Click Save.

Step	Description	UI Steps
2.	Configure the second GigaSMART Group	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** Click New. **c.** Type an grp2 in the Alias field to identify this GigaSMART Group. **d.** Click in the Port List field and select an engine port. For example, 1/5/e1. **e.** Click Save.
3.	Display the information about the GigaSMART Group.	• On the GigaSMART Group page, click on the alias gsgp2 to display the GigaSMART Group Quick View.

Ex 4, Step 2: Configure IP Interface with Tool Port

Create an IP interface with tool port. You must associated this IP interface with the NetFlow Exporter in *Ex 4, Step 3: Configure IP Interface and Associate with Exporter* on page 1130.

Step	Description	UI Steps
1.	Identify the collector port and configure it as a tool port, where the NetFlow collector will be connected.	**a.** Identify the port to use as an IP interface with tool port. For example, 1/1/g3. **b.** Select Ports > All Ports. **c.** Use the Quick Port Editor to configure port 1/1/g3 as tool port.
2.	Configure the IP interface. The IP address is for the NetFlow interface.	**a.** Select Ports > IP Interfaces. **b.** Click New. **c.** In the Alias and Comment fields, enter the name and description of the IP interface. **d.** Click in the Port field and select port 1/1/g3. **e.** Enter 10.10.10.10 in the IP Address field. **f.** Enter 255.255.255.255 in the Mask field **g.** Enter 10.10.10.1 in the Gateway field. **h.** Enter 1500 in the MTU field. **i.** Select the GigaSMART Group grp2 created in *Ex 4, Step 1: Configure GigaSMART Group*. **j.** Click Save.
3.	Display the IP interfaces configuration.	**a.** Select Ports > IP Interfaces. **b.** Click on the row for port 1/1/g3 to display the details in a Quick View.

Ex 4, Step 3: Configure IP Interface and Associate with Exporter

Configure the IP interface and associate the NetFlow Generation Exporter, as follows:

Step	Description	Command
1.	Configure the IP interface.	**a.** Select Ports > IP Interfaces. **b.** Click New. **c.** In the Alias and Comment fields, enter the name and description of the IP interface. **d.** Click in the Port field and select the IP interface configured in *Ex 4, Step 2: Configure IP Interface with Tool Port* on page 1129. **e.** Set the IP address to 10.115.9.6. **f.** Set the IP Mask to 255.255.248.0. **g.** Set the Gateway address to 10.115.8.1. **h.** Set MTU to 1500. **i.** Select the GigaSMART Group grp2 from the GigaSMART Group list. **j.** From the Exporter drop-down list, select ex21. **k.** Click Save.

Ex 4, Step 4: Configure Record

Configure a NetFlow Generation Record, as follows:

Task	Description	UI Steps
1.	Configure the record.	**a.** From the device view, select GigaSMART > NetFlow / IPFIX Generation > Records. **b.** Click New. **c.** Type rec1 in the Alias field. **d.** Select IPFIX. **e.** Configure Key Fields (Match) Select IPV4 Select TOS and then select Source. **f.** Configure Non-Key Fields (Collect) Select IPv4 and then select Protocol and Source Select Interface **g.** Click Save.

Ex 4, Step 5: Configure Monitor

Configure a NetFlow Generation Monitor and associate the NetFlow Generation Record to the specified NetFlow Generation Monitor, as follows:

Step	Description	Command
1.	Configure the monitor.	**a.** From the device view, select GigaSMART > NetFlow / IPFIX Generation > Records. **b.** Click New. **c.** Type mon1 in the Alias field. **d.** Select Normal for Cache Type. **e.** Set Cache Timeout Active to 2. **f.** Set Cache Timeout Inactive to 2 **g.** Select rec1 from the Record list. **h.** Click Save.

Ex 4, Step 6: Add Monitor to GigaSMART Group

Add the monitor created in *Ex 4, Step 5: Configure Monitor* to the GigaSMART Group created in step *Ex 4, Step 1: Configure GigaSMART Group*.

NOTE: Only one NetFlow Generation Monitor can be configured per gsgroup.

Step	Description	UI Steps
1.	Select the GigaSMART Group configured in *Ex 4, Step 1: Configure GigaSMART Group*.	**a.** From the device view, select GigaSMART > GigaSMART Groups > GigaSMART Groups. **b.** On the GigaSMART Groups page, select grp2. **c.** Click Edit. **d.** Click in the Monitor field and select the NetFlow Monitor created in *Ex 4, Step 5: Configure Monitor* **e.** Click Save.
2.	Display the GigaSMART Group information.	**a.** From the device view, select GigaSMART Groups > GigaSMART Groups. **b.** On the GigaSMART Groups page, click on the alias grp2 to display the information in a Quick View.

Ex 4, Step 7: Configure GigaSMART Operation

Define the two GigaSMART operation to enable masking and NetFlow, as follows:

Step	Description	UI Steps
1.	Configure the first GigaSMART Operation for APF and masking	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Type gsop_mask_aa in the Alias field **d.** For the GigaSMART Group, select grp. **e.** Click in th GigaSMART Operations (GSOP) field and select Adaptive Packet Filtering. **f.** Click in th GigaSMART Operations (GSOP) field and select Masking. Set Masking as follows: Protocol: None Offset: 50 Pattern: aa. Length: 100 **g.** Click Save.
2.	Configure the second GigaSMART Operation for NetFlow	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation. **b.** Click New. **c.** Type an alias in the Alias field to identify this operation. For example, gsop2. **d.** For GigaSMART Group, select grp **e.** For GigaSMART Operations (GSOP), select NetFlow. **f.** Click Save.
3.	Display the configuration GigaSMART Operation.	• On the GigaSMART Operations page, click on the alias gsop2 to open the GigaSMART Operation Quick View.

Ex 4, Step 8: Configure Virtual Port

For the second level map that you will create in *Ex 4, Step 9: Configure Mapping Rules to Filter Packets*, you will need to create virtual ports.

Configure a virtual port and associate it with the GigaSMART group, as follows:

Step	Description	UI Steps
1.	Configure the first virtual port for the second level map.	**a.** From the device view, select GigaSMART > Virtual Ports. **b.** On the GigaSMART Groups page, select grp. **c.** Click New. **d.** Type vp1 n the Alias field. **e.** Click in the GigaSMART Group field and select grp1. **f.** Click Save.

Step	Description	UI Steps
1.	Configure the second virtual port for the second level map.	**a.** From the device view, select GigaSMART > Virtual Ports. **b.** On the GigaSMART Groups page, select grp. **c.** Click New. **d.** Type vp2 in the Alias field. **e.** Click in the GigaSMART Group field and select the grp2. **f.** Click Save.
2.	Display the virtual port information.	**a.** From the device view, select GigaSMART Groups > Virtual Ports. **b.** On the Virtual Ports page, click on the virtual port alias to display the information in a Quick View.

Ex 4, Step 9: Configure Mapping Rules to Filter Packets

To add flow mapping rules to filter packets that are needed to run NetFlow, configure maps and associate the maps to the IP interfaces with tool ports, as follows:

Task	Description	UI Steps
1.	Configure the first map. (This is a first level map.)	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map1 in the Alias field. **d.** Select First Level for Type and By Rule for Subtype. **e.** Set network port 11/1/g1 as the Source. **f.** Set virtual ports v1 and v2 plus tool port 1/1/g3 as the Destination. **g.** Click Add a Rule. **h.** Select Pass **i.** Select MAC Destination. **j.** Set the MAC address and mask to 00:00:00:00:00:00 **k.** Click Save.
2.	Configure the second map. (This is a second level map.)	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map2 in the Alias field. **d.** Select Second Level for Type and By Rule for Subtype. **e.** Set virtual port v1 as the Source. **f.** Set tool port v1 11/1/g2 as the Destination. **g.** Select gsop_mask_aa from the GSOP list. **h.** Click Add a Rule. **i.** Select Pass **j.** Select MAC Destination. **k.** Set the MAC address and mask to 00:00:00:00:00:00 **l.** Click Save.

Task	Description	UI Steps
3.	Configure the third map. (This is also a second level map.)	**a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map2 in the Alias field. **d.** Select Second Level for Type and By Rule for Subtype. **e.** Set virtual port v2 as the Source. **f.** Set tool port v1 11/1/g4 as the Destination. **g.** Select gsop_apf_netflow from the GSOP list. **h.** Click Add a Rule. **i.** Select Pass **j.** Select MAC Destination. **k.** Set the MAC address and mask to 00:00:00:00:00:00 **l.** Click Save.

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* on page 784 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* on page 783 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* on page 785 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 30-104.

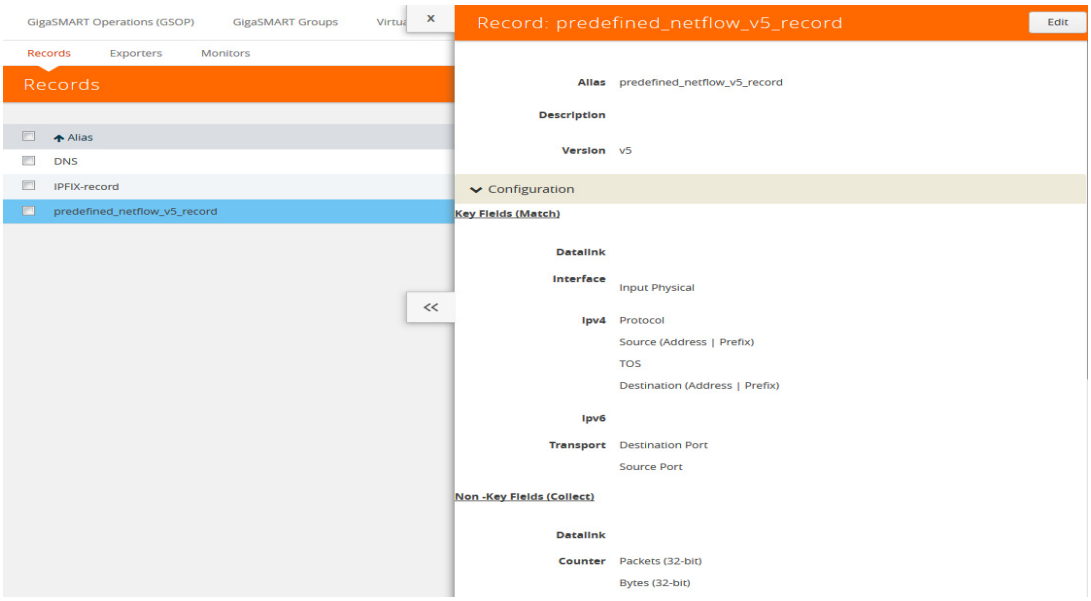

Figure 30-104: NetFlow Record predefined_netflow_v5_record

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.
	Option Map			Supported only for IPFIX.
	Precedence			Supported only for IPFIX.
	Protocol			Supported for v9 and IPFIX.

Match Type	Parameters			Description
	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.
	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.

Match Type	Parameters			Description
		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 | disable>	[ACK] | [CWR] | [ECE] | [FIN] | [PSH] | [RST] | [SYN] | [URG]	Supported only for v9 and IPFIX.
		Header Length		Supported only for IPFIX.
		Sequence Number		Supported only for IPFIX.
		Source Port		Supported only for IPFIX.
		Urgent Pointer		Supported only for IPFIX.
		window-size		Supported only for IPFIX.
	UDP	Destination Port		Supported only for IPFIX.
		Message Length		Supported only for IPFIX.
		Source Port		Supported only for IPFIX.

Collect/Non-Key Syntax

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

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

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

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

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

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

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

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

Collect Type	Parameters	Size		Description
		Prefix	**<netmask | mask_length>**	Configures a prefix for the IPv4 destination address as a non-key field. Supported only for IPFIX.
	Traffic Class			Supported for v9 and IPFIX.
	Version			Supported for v9 and IPFIX.

Collect Type	Parameters	Size		Description
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> | 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>] |**	Supported only for IPFIX.

Collect Type	Parameters	Size		Description
Private (continued)	**PEN <pen name>**	**DNS**	**query-type-text [number-of-collects <1-10>] | response-class [number-of-collects <1-10>] | response-class-text [number-of-collects <1-10>] | response-code | response-ipv4-addr [number-of-collects <1-10>] | response-ipv4-addr-text [number-of-collects <1-10>] | response-ipv6-addr [number-of-collects <1-10>] | response-ipv6-addr-text [number-of-collects <1-10>] | response-name [number-of-collects <1-10>] | 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>]>**	Supported only for IPFIX.
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.

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

Collect Type	Parameters	Size	Description
	UDP	**Destination Port**	Supported only for IPFIX.
		Message Length	Supported only for IPFIX.
		Source Port	Supported only for IPFIX.

GigaSMART Load Balancing

GigaSMART Load Balancing does not 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-HC2, and GigaVUE-HC3), and Tunneling license (GigaVUE-HC1 and GigaVUE-HB1)

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* on page 757 for details.

The following sections describe the available load balancing schemes:

- *Stateful Load Balancing* on page 1147
- *Stateless Load Balancing* on page 1155
- *Enhanced Load Balancing* on page 1163

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.

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.

GTP, Application Session Filtering (ASF), and tunnel are the currently supported stateful applications.

For information on GTP, refer to *GigaSMART GTP Correlation* on page 885.

For information on ASF, refer to *GigaSMART Application Session Filtering (ASF) and Buffer ASF* on page 1054.

For information on Layer 2 GRE tunnel encapsulation, refer to *GigaSMART Layer 2 GRE Tunnel Encapsulation/Decapsulation* on page 828.

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

- *Stateful Load Balancing Metrics* on page 1149
- *Hashing Key Support* on page 1150
- *Configure Stateful Load Balancing* on page 1150

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**, and **Tunnel** are the currently supported stateful applications.
7. Specify a load balancing metric. For example, **Hashing** as shown in Figure 30-105.

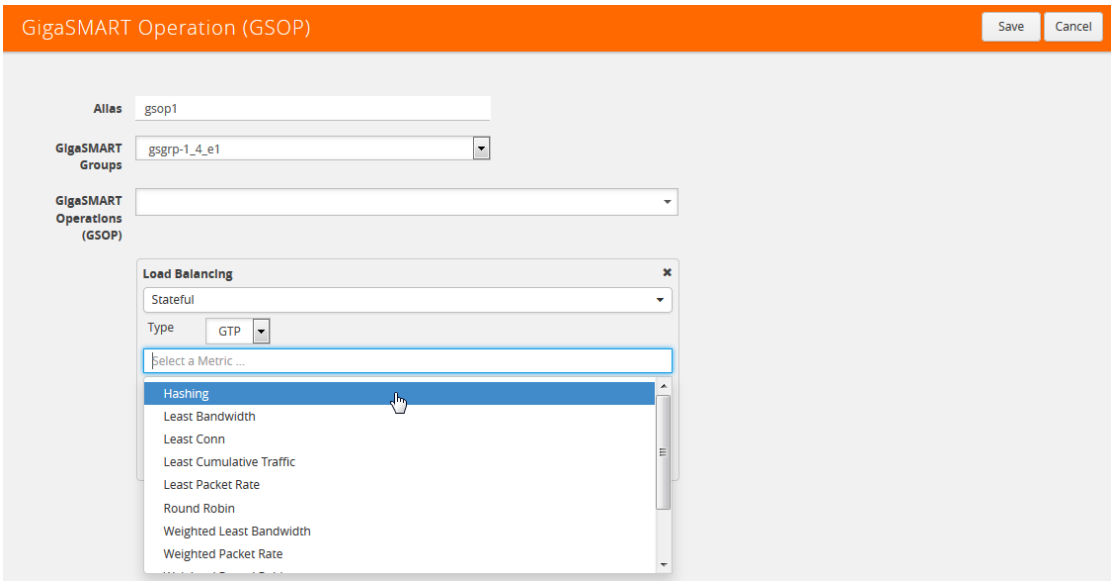

Figure 30-105: Configuring Stateful Load Balancing

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. If the following command is enabled, link speed is considered in addition to the bandwidth of the port. If the following command is disabled, 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 the following command is enabled, link speed is considered with packet rate. If the following command is disabled, 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 the following command is enabled, link speed is considered with Round Robin. If the following command is disabled, 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.

Metric	Description
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 the following command is enabled, link speed is considered with Least Connection. If the following command is disabled, 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 the following command is enabled, link speed is considered with least cumulative traffic. If the following command is disabled, the weight configured for each port in the port group is considered with Least Cumulative Traffic.
Hashing	A tool port is selected from a port list based on hashing of data provided by the GSOP application, which is normally extracted from the packet. The values for hashing key are: IMSI GTP key (imsi), IMEI GTP key (imei), and MSISDN GTP key (msisdn). The hashing key only applies to the GTP stateful application. Refer to *Hashing Key Support* on page 672 for details.

Hashing Key Support

The following table describes the support for GTP hashing key.

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

Configure Stateful Load Balancing

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

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

1. Select **Ports > Port Groups > All Port Groups**.
2. Click **New**.
3. In the Alias filed, enter an alias. For example, load-balgrp.
4. Select **SMART Load Balancing.**

5. Use the Ports field to select the ports for this port group as shown in Figure 30-106. Click Save when you are done.

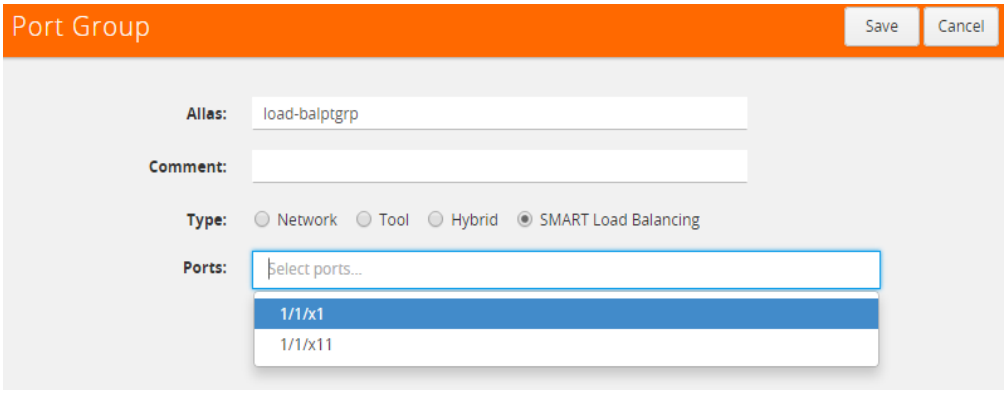

Figure 30-106: Port Group Page to Set Ports for Load Balancing

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.

Refer to the following examples:

- *Example 1: GigaSMART Stateful Load Balancing* on page 1151
- *Example 2: GigaSMART Stateful Load Balancing* on page 1153

For an example of load balancing on L2GRE encapsulation tunnel, refer to *Example 2 – GigaSMART L2GRE Tunnel Encap Stateful LB* on page 833.

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.** Select SMART 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: • Select Stateful. • Select GTP for Type. • Select Weighted Least Bandwidth for the metric.
4.	Create a virtual port and associate it with the GigaSMART group.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > Virtual Ports. **b.** Click New. **c.** Type vp1 in the Alias field. **d.** Select gsgrp1 from the GigaSMART Group list. **e.** Click Save.

Task	Description	UI Steps
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. • Click Add a Rule • Select Pass • Select Port Destination for the condition. • Set the port value to 2123. **h.** Add Rule 2. • Click Add a Rule • Select Pass • Select Port Destination for the condition. • 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 • Select Pass • Select GTP IMSI for the condition. • Enter 234567* for IMSI. • 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.

Step	Description	Command
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.** Select SMART 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 Save.
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 Save.
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: • Select Stateful. • Select GTP for Type. • Select Hashing for the metric. • Select IMEI **h.** Click Save.
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 Save.

Step	Description	Command
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. • Click Add a Rule • Select Pass • Select Port Destination for the condition. • Set the port value to 2123. **h.** Add Rule 2. • Click Add a Rule • Select Pass • Select Port Destination for the condition. • 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 • Select Pass • Select GTP IMSI for the condition. • Enter 234567* for IMSI. • 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.

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

- *Stateless Load Balancing Metrics* on page 1156
- *Configure Stateless Load Balancing* on page 1157

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 Figure 30-105.
6. Click **Save**.

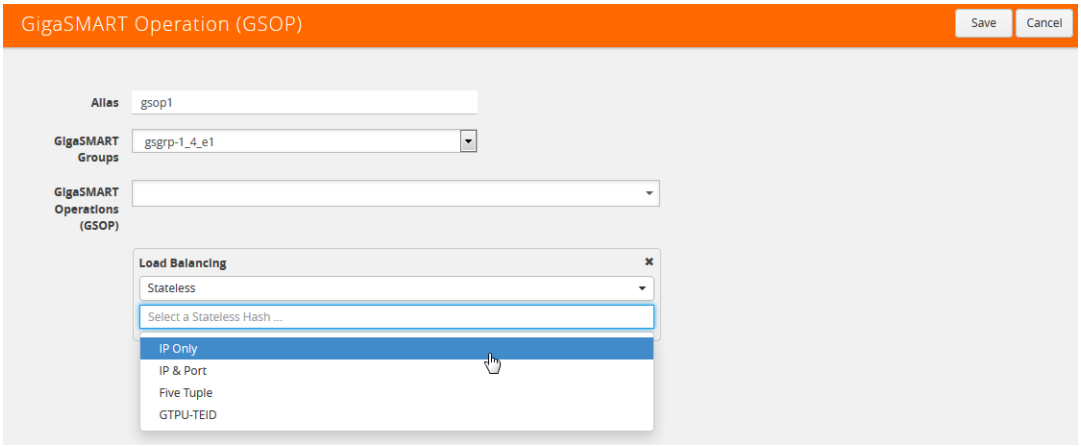

Figure 30-107: Selecting Stateless Load Balancing

For more information, refer to the following sections:

- *Stateless Load Balancing Metrics* on page 1156
- *Configure Stateless Load Balancing* on page 1157

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 30-5 describes each of the metrics.
4. Click **Save**.

Table 30-5: Stateless Load Balancing Metrics

Metric	Description
IP Only	The source IP and destination IP addresses.
IP & Port	The source IP and destination IP addresses, and Layer 4 source port and destination port numbers.
Five Tuple	The source IP and destination IP addresses, source port and destination port numbers, and protocol field in the IP header.
GTPU-TEID	The GTP-u tunnel identifier (ID). **NOTE:** There is no inner or outer field location for **GTPU-TEID**.
outer	The first occurrence of header or field. For example, **IP Only outer** is the first IP header in the packet, which could be IPv4 or IPv6.
inner	The second occurrence of header or field. For example, **ip-only inner** is the second IP header in the packet. The first IP header could be IPv4 or IPv6, as follows: • IPv4-IPv4—**IP Only inner** is the IP addresses in the second IPv4 header • IPv6-IPv4—**IP Only inner** is the IP addresses in the IPv4 header • IPv4-IPv6—**IP-Only inner** is the IP addresses in the IPv6 header The supported IP encapsulation types are: IP-in-IP, VXLAN, GTP, GRE, and ERSPAN.

Configure Stateless Load Balancing

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

1. Select **Ports > Port Groups > All Port Groups**.
2. Click **New**.
3. Type an alias in the **Alias** field. For example, load-balgrp.
4. Select **SMART Load Balancing.**
5. Use the Ports field to select the ports for this port group as shown in Figure 30-106. Click **Save** when you are done.

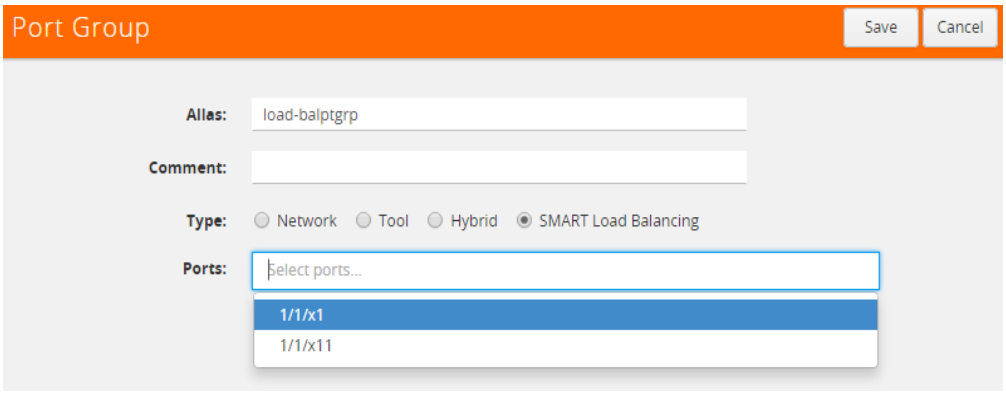

Figure 30-108: Port Group Page to Set Ports for Load Balancing

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.

Refer to the following examples:

- *Example 1: GigaSMART Stateless Load Balancing* on page 1158
- *Example 2: GigaSMART Stateless Load Balancing* on page 1159
- *Example 3: GigaSMART Stateless Load Balancing* on page 1161

For an example of load balancing on L2GRE encapsulation tunnel, refer to *Example 3 – GigaSMART L2GRE Tunnel Encap Stateless LB* on page 836.

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.

Step	Description	Command
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.** Select SMART 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 Save.
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 Save.

Step	Description	Command
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 Slicing. **f.** Configure Slicing as follows: • Select Stateful. • Select None for protocol • Set Offset to 70. **g.** From the device view, select GigaSMART Operation Load Balancing. **h.** Configure Load Balancing as follows: • Select Stateless • Select IP Only Outer for the hash metric **i.** Click Save.
4.	Create a first level map. Note the following: • You can specify only one port group as part of the map tool port in the **to** statement. • You can define the same load balancing port group in multiple maps, however, the load balancing metrics defined in the GSOPs on those maps have to be the same. • You cannot use a shared collector map for load balancing.	**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 • Select Pass • Select IP Version for the condition. • 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.

Step	Description	Command
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.** Select SMART 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 Save.

Step	Description	Command
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 Save.
3.	Create a GSOP, including load balancing metric.	**a.** From the device view, select GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operations. **b.** Click New. **c.** Type gsop1 in the Alias field **d.** Select gsgrp1 from the GigaSMART Groups list. **e.** Select the GigaSMART Operation Load Balancing. **f.** Configure Load Balancing as follows: • Select Stateless • Select GTPU-TEID for the hash metric **g.** Click Save.
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 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 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 Save.

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.

Step	Description	Command
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.** Select SMART 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 Save.
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 Save.
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. • Select Stateless • Select IP Only Inner for the hash metric **k.** Click Save.

Step	Description	Command
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 Regular for Type and By Rule for Subtype. **e.** Select port 1/1/x1 for the Source. **f.** Select port group portgrp1 for the Destination. **g.** Click Add a Rule, and then select Pass **h.** Select IPv4 Protocol for the first condition. **i.** Select UDP for Value **j.** Select Port Destination for the second condition. **k.** Enter 2123 for the port value. **l.** Click Save. Create the second map: **a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map1 in the Alias field. **d.** Select First Level for Regular and By Rule for Subtype. **e.** Select port 1/1/x1 for the Source. **f.** Select port vp1 for the Destination. **g.** Select gsop1 from the GSOP list. **h.** Click Add a Rule, and then select Pass **i.** Select IPv4 Protocol for the first condition. **j.** Select UDP for Value **k.** Select Port Destination for the second condition. **l.** Enter 2152 for the port value. **m.** Click Save.
5.	Create the second level map.	Create the third map: **a.** Select Maps > Maps > Maps. **b.** Click New. **c.** Type map22 in the Alias field. **d.** Select Second Level for Type and By Rule for Subtype. **e.** Select virtual vp1 for the Source. **f.** Select port group portgrp1 for the Destination. **g.** Select gsop1 from the GSOP list. **h.** Click Add a Rule, and then select Pass **i.** Select IPv4 Destination for the first condition. **j.** Enter 80 for the destination value. **k.** Select 2 for position. **l.** Click Save.

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
  - No tool port within the port group receives more than 5% of average traffic based on the enhanced load balancing metric defined

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.

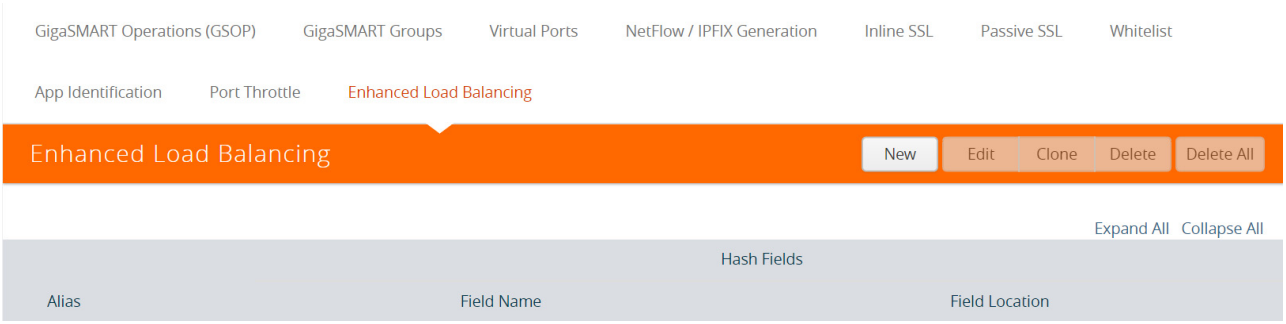

4. Click **New**. The enhanced load balancing screen appears.

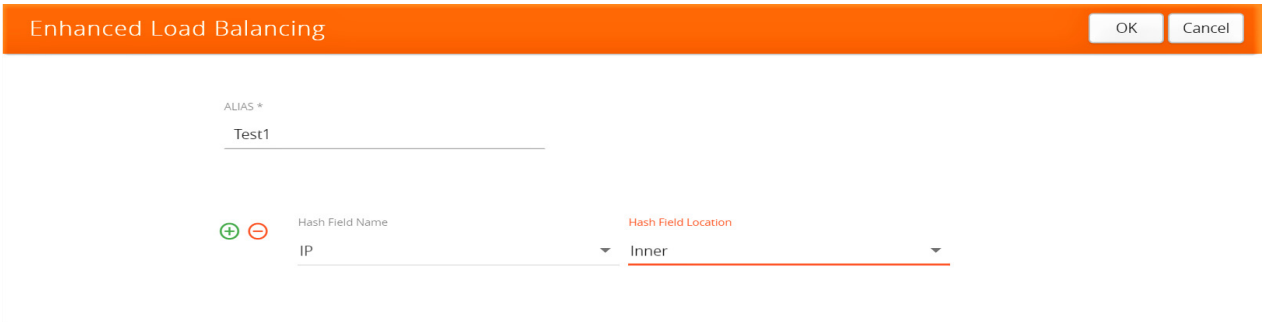

5. Enter an Alias.
6. Select the Hash Field Name. The following options are available:
   - IP
   - L4 Port
   - GTP-U TEID
7. Select Hash Field Location. Options:
   - inner
   - outer

   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.

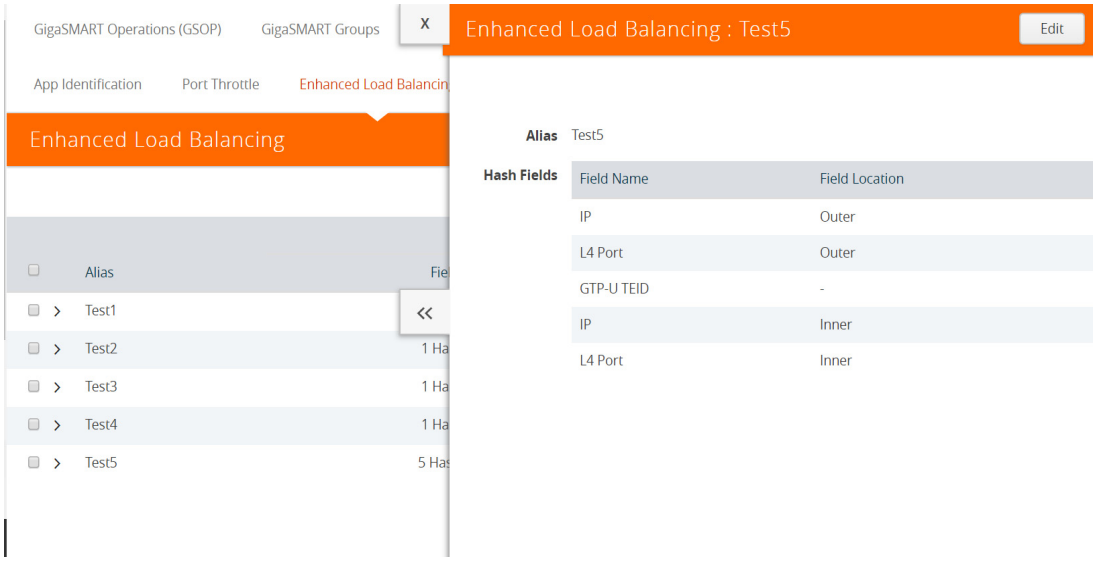

8. Click **OK**.

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

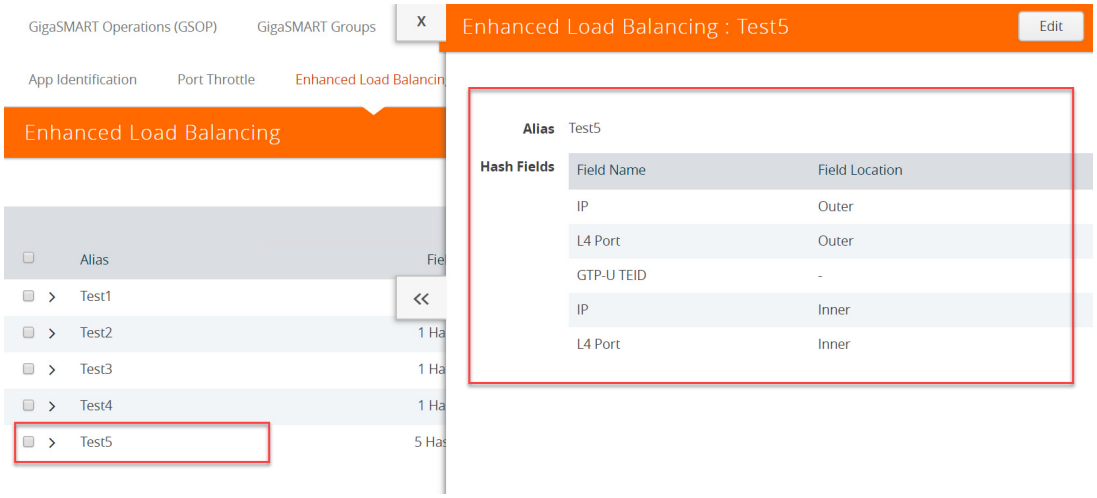

Configure GigaSMART Operation (GSOP)

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.

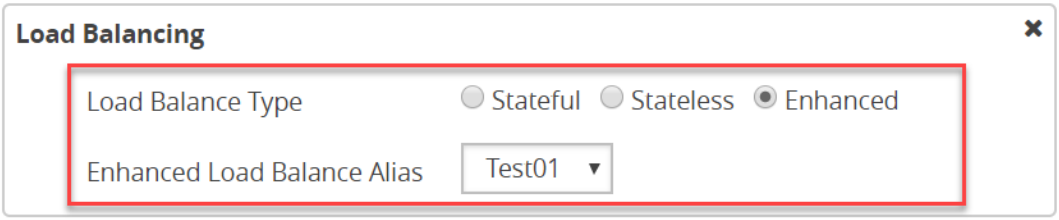

10. Click **OK**. The enhanced load balance GSOP is now available on the GSOP page.

GigaSMART MPLS Traffic Performance Enhancement

GigaSMART MPLS Traffic Performance Enhancement does not require a separate license.

The GigaSMART MPLS 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 30-109 shows the MPLS example.

Figure 30-109: 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.

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.

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.

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 Group**s.
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* on page 1167
- *Flow Masking Example 2* on page 1168

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 MAC address + 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 MAC address + 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 Out-of-Band SSL Decryption

Required License: SSL Decryption for Out-of-Band

GigaVUE H Series nodes support Secure Sockets Layer (SSL) decryption. SSL is a cryptographic protocol that adds security to 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. Out-of-band SSL decryption delivers decrypted traffic to out-of-band tools that can then detect threats entering the network.

Out-of-band SSL decryption is a pillar of the GigaSECURE Security Delivery Platform. For an overview of GigaSECURE, refer to *GigaSECURE Security Delivery Platform* on page 556.

On GigaVUE H Series nodes, GigaSMART line cards or modules perform the decryption of 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 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 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.

Out-of-band SSL decryption is supported on the following GigaVUE H Series products with GigaSMART line cards or modules installed:

- GigaVUE-HC3
- GigaVUE-HC2
- GigaVUE-HC1
- GigaVUE-HB1

Use out-of-band SSL decryption on the GigaSMART line card or module with passive or offline traffic. Tap the traffic to and from a server and pass it to the GigaVUE H Series node with the GigaSMART line card or module.

Out-of-band SSL decryption operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to *Groups of GigaSMART Engine Ports* on page 757 for details.

For secure storage of private keys, Thales Hardware Security Module (HSM) is integrated with out-of-band SSL decryption. Refer to *Thales HSM for SSL Decryption for Out-of-Band Tools* on page 1181 for details.

Gigamon also offers inline SSL decryption, which inspects SSL encrypted traffic inline. Refer to the *Inline SSL Decryption Guide* for details.

Overview of Out-of-Band SSL Decryption

SSL encryption secures traffic between a client and a server, such as a Web server. 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 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 30-110 shows the basic SSL handshake between a client and server to establish a session. The messages are unencrypted up to step 6 in Figure 30-110. The messages are encrypted after step 6, including the step 9 Finished message.

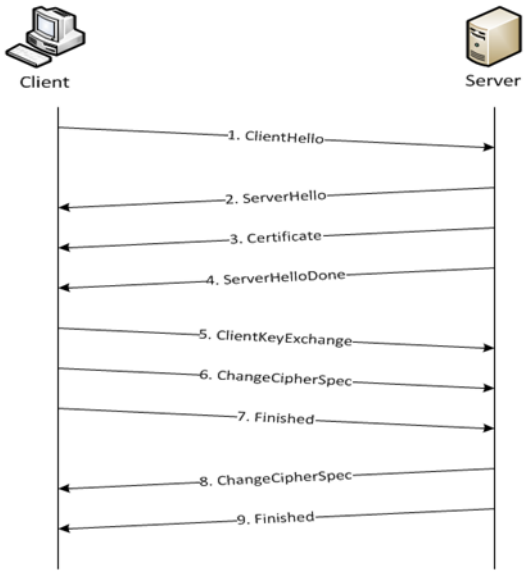

Figure 30-110: Basic SSL Handshake

Once a session has been established, the keys are saved so a session can be resumed efficiently later. Figure 30-111 shows the resumed SSL handshake, with fewer steps.

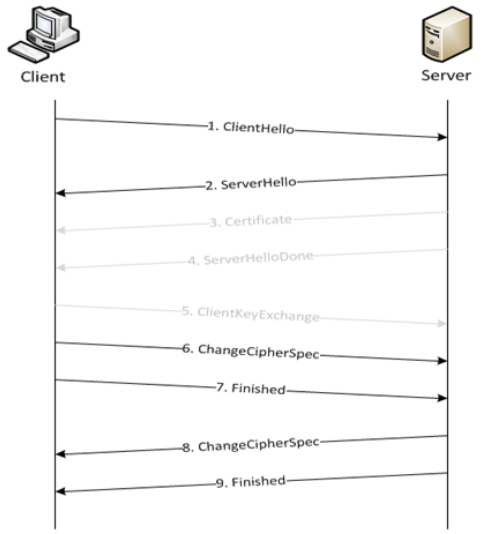

Figure 30-111: Resumed SSL Handshake

Out-of-band SSL decryption can be deployed close to the server, as shown in Figure 30-112.

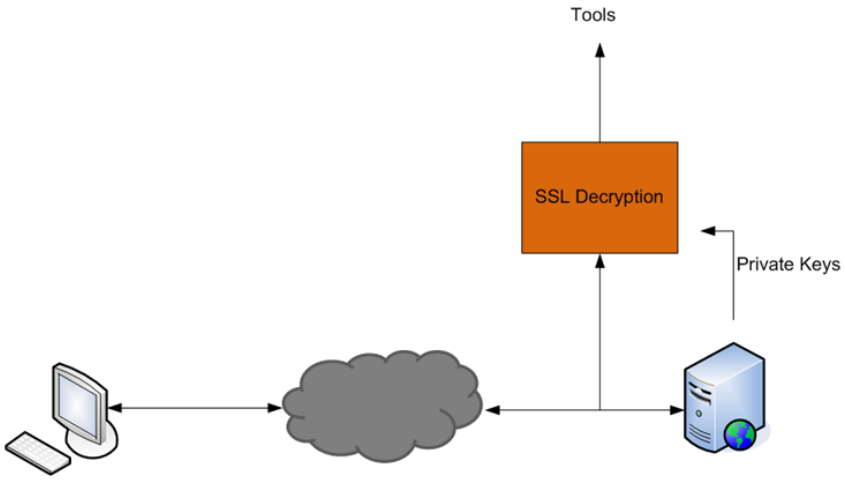

Figure 30-112: Inbound (Server Side)

Out-of-band SSL decryption can also be deployed close to an SSL proxy, with the server in the Enterprise domain as shown in Figure 30-113.

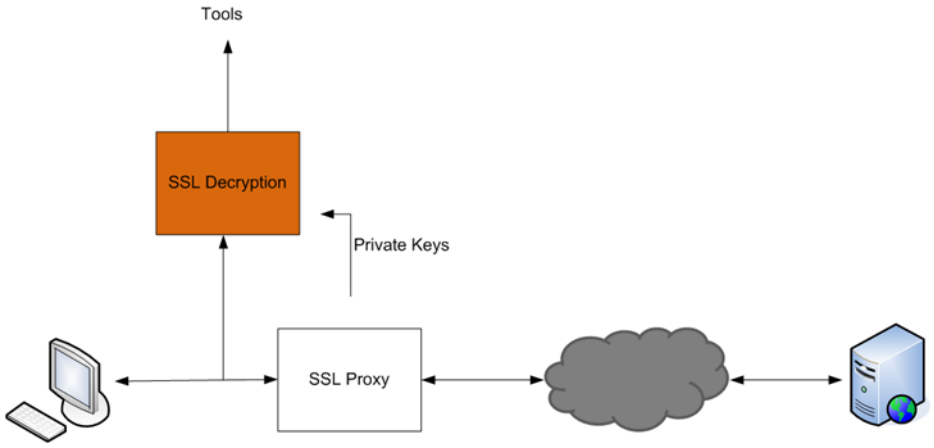

Figure 30-113: Outbound (Client/Enterprise Side)

In Figure 30-112, tap traffic to the server and then send it for decryption. In Figure 30-113, 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 out-of-band SSL decryption on GigaSMART:

- *Supported Protocols, Algorithms, and Ciphers* on page 1172
- *Limitations* on page 1174
- *Create and Reset Passwords* on page 1175
- *Work with Keys and Services* on page 1176

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 30-6.

Table 30-6: Supported Ciphers for Out-Of-Band 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

Table 30-6: Supported Ciphers for Out-Of-Band SSL decryption

Cipher Name	Kx	Au	Enc	Bits	Mac
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 out-of-band SSL decryption are FIPS 140-2 compliant.

All key URLs must point to an RSA private key stored in the PEM or PKCS12 format, as follows:

- http://keyserver.domain.com/path/keyfile.pem
- https://keyserver.domain.com/path/keyfile.pem
- ftp://keyserver.domain.com/path/keyfile.pem
- tftp://keyserver.domain.com/path/keyfile.pem
- scp://username[:password]@keyserver.domain.com/path/keyfile.pem

The supported applications are as follows:

- HTTPS
- FTPS
- SMTP, IMAP, and POP3 with StartTLS

Limitations

The following are limitations of out-of-band SSL decryption:

- 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* on page 1172. If an SSL session cannot be decrypted due to having a non-supported protocol or cipher, the packets will be forwarded to the tool without decryption if the GS Parameter **SSL Decryption** has **Decrypt Fail Action** is set to **Pass to Tool Port**. Non-supported ciphers and protocols include SSL 2.0, TLS 1.3, Diffie-Hellman (DHE keys), Ephemeral keys, Elliptic Curves Extension, compression, and 8K key size.

Licensing

The GigaSMART license for out-of-band SSL decryption is installed as any other license.

There are no limits to the number of out-of-band SSL decryption sessions or the number of users.

Create and Reset Passwords

To perform the configuration in the following section, you must have an admin level access role.

Before uploading keys or configuring 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* on page 1178.

NOTE: When uploading SSL keys, make sure that you are not creating a duplicate key. Adding a duplicate key can cause errors.

To create an SSL keychain password, use the following steps:

1. From the device view, select **GigaSMART > Passive SSL > Key Store**.
2. Click **Keychain Password**.

   The set-up page shown in figure Figure 30-114 displays.

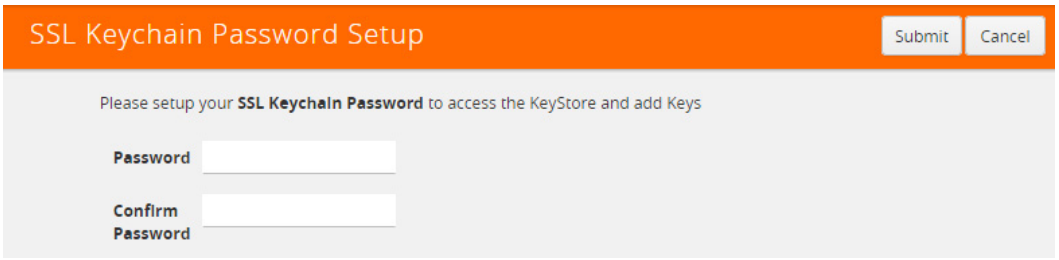

Figure 30-114: SSL Keychain Password Setup Page

3. Enter a password in the **Password** and **Confirm Password** fields.

   You can only configure a strong password. A strong password has at least ten (10) characters and at least three (3) of the following:

   - uppercase letters
   - lowercase letters
   - numbers
   - special characters
4. 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 SSL keychain password is lost, it can be reset, but all existing private keys will be revoked. When there are keys installed on the node, a warning is displayed before you are prompted for the new password.

Once you have a new password, you will have to upload the keys again.

Work with Keys and Services

This section describes working with private keys as well as services. Keys must be uploaded to the GigaVUE H Series node using a unique alias. Services must be defined for each server destination that needs decryption.

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* on page 1178.

NOTE: When uploading SSL keys, make sure that you are not creating a duplicate key. Adding a duplicate key can cause errors.

Upload SSL Private Keys

To upload an SSL private key, do the following:

1. From the device view, select **GigaSMART >Passive SSL > Key Store** to open the Key Store page shown in Figure 30-115.

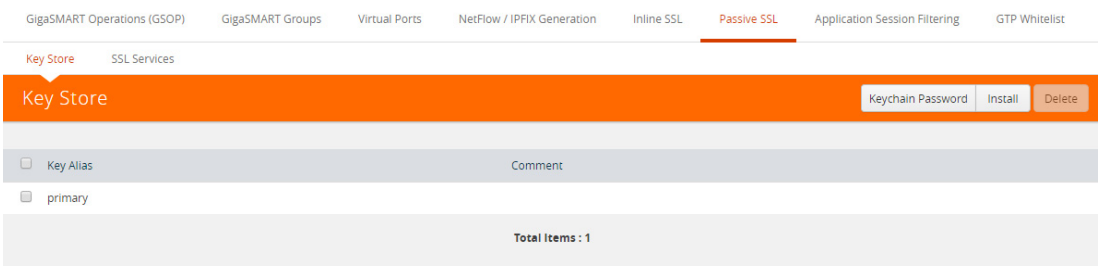

Figure 30-115: SSL Key Store Page

2. Click **Install**.

   The SSL Key page shown in Figure 30-116 displays.

Figure 30-116: SSL Key Page

3. Enter an alias for the SSL key in the Alias field.

4. Select the upload type.
5. Choose the file.

   The URL can be downloaded using HTTP, HTTPS, FTP, TFTP, SCP, and SFTP. Using a secure protocol, such as HTTPS is recommended. Only PEM encoded keys are supported.
6. Click **Save**.

Delete SSL Key

To delete a particular SSL private key, select the key on the SSL Keys page, and then select Delete. To delete all SSL private keys, select multiple keys.

Create 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 SSL Private Keys* on page 1187
- Create GigaSMART Group with 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 displays as shown in Figure 30-117.

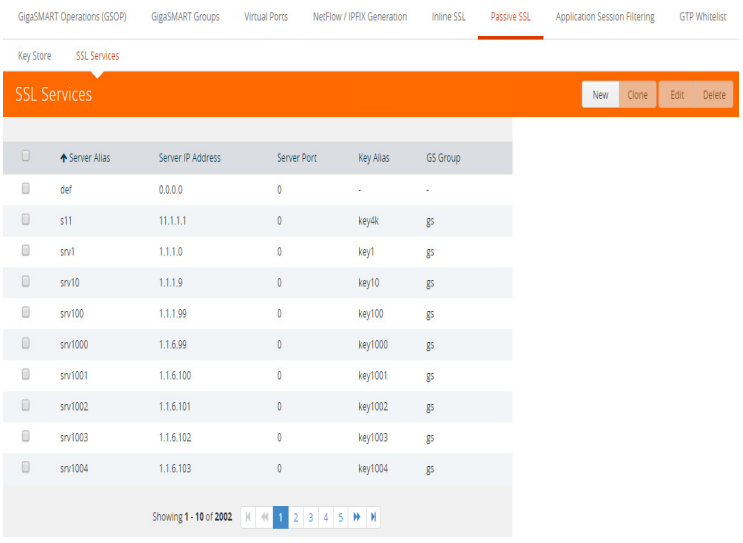

Figure 30-117: SSL Services Page

2. Click **New**.

3. On the SSL Service configuration page, do the following:
   - Enter an alias.
   - Enter the information for the service: IP Address, Server Port.
   - Select the alias of SSL Key previously uploaded. For the steps, refer to *Upload SSL Private Keys* on page 1187.
   - Select the GigaSMART Group with SSL decryption enabled to associate with this SSL service.

   Figure 30-118 shows an example of an SSL Service.
4. Click **OK**.

Figure 30-118: SSL Service Settings

Delete Service

To delete a particular SSL service select the service on the SSL Services page, and then select Delete. To delete all 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, out-of-band SSL decryption is not working.
- Key content cannot be displayed.
- Keys that are synchronized across a cluster are encrypted.

ECODES for Troubleshooting Out-of-Band SSL Decryption

Use the following table of ECODE messages to assist with troubleshooting out-of-band SSL decryption:

ECODE	Description
81	TCP flow errors detected. Make sure you see the complete TCP flow. Use the de-duplication GigaSMART operation with out-of-band 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 SSL version. An SSL handshake processing error occurred. Use the de-duplication GigaSMART operation with out-of-band SSL decryption.
221	Unknown SSL version. An unsupported 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.

Display Out-of-Band SSL Decryption Flow Ops Report

To display the Flow Ops report for out-of-band SSL decryption:

1. From the device view, select **GigaSMART > GigaSMART Groups > Report**. The Report page displays as shown in Figure 30-119.

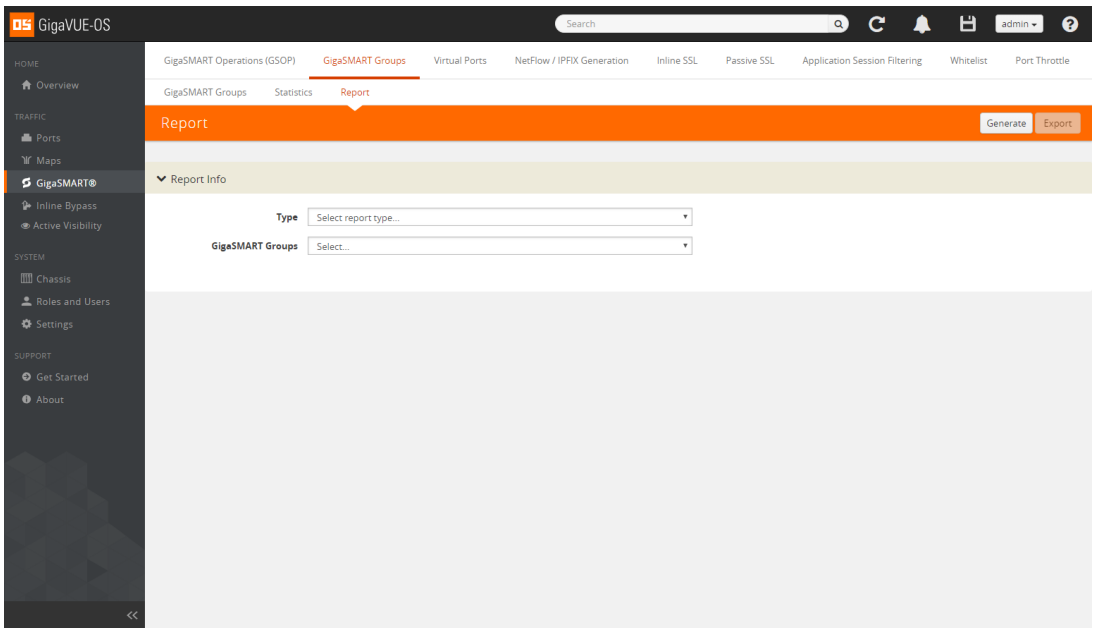

Figure 30-119: Report Page

2. Select report type **SSL Decryption** and select a GigaSMART group from the drop-down menu.
3. Click **Generate**. The SSL Decryption Report Summary displays as shown in Figure 30-120.

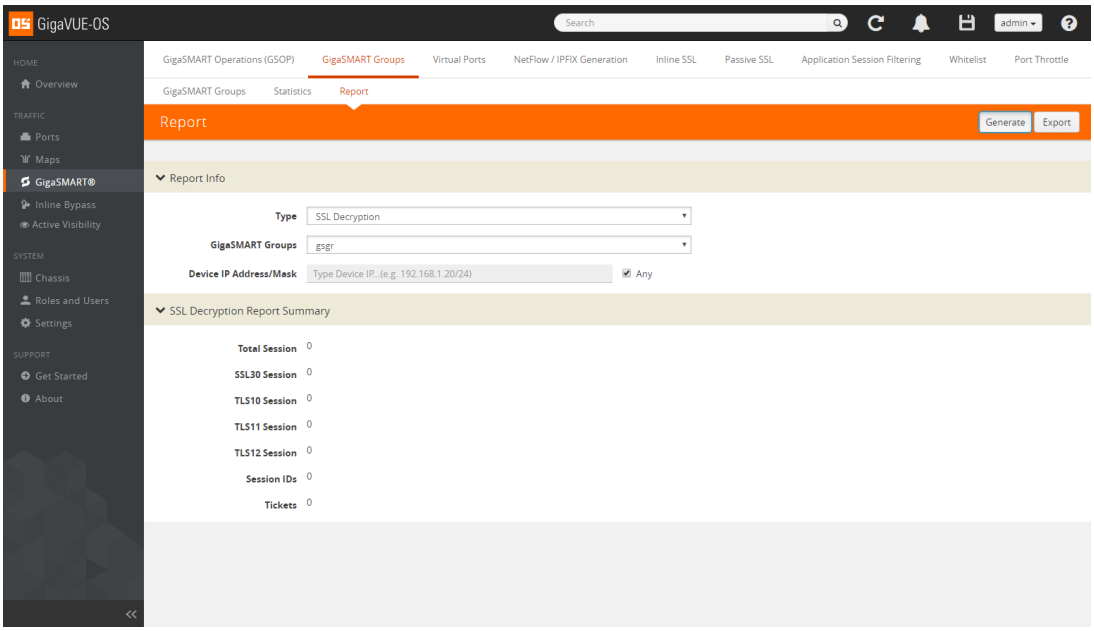

Figure 30-120: SSL Decryption Report Summary

The fields in the SSL Decryption Report Summary are as follows:

- Total Session—The total number of out-of-band SSL decryption sessions.
- SSL30 Session—The cumulative total number of SSL 3.0 sessions.
- TLS10 Session—The cumulative total number of TLS 1.0 sessions.
- TLS11 Session—The cumulative total number of TLS 1.1 sessions.
- TLS12 Session—The cumulative total number of TLS 1.2 sessions.
- Session IDs—The number of concurrent session IDs.
- Tickets—The number of current TLS tickets.

Thales HSM for SSL Decryption for Out-of-Band Tools

Required License: Included with SSL Decryption for Out-of-Band Tools

Starting in software version 5.3, Thales Hardware Security Module (HSM) is integrated with out-of-band SSL decryption. Hardware Security Modules offer secure storage, management, and operation of cryptographic material, such as private keys and passphrases.The HSM stores and manages the keys in a safe and secure environment. Since the keys reside on the HSM in the network, they are offloaded from an application on a network device.

The application could be a web server or a database server, but, in the case of 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 the HSM and the application.

Keys are added to the HSM by an administrator. When an application's key is on the 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 the HSM to use the key. In this way, the use of keys by the application is secure because only key tokens are exchanged.

Thales HSM is supported on GigaVUE-HC1, GigaVUE-HC2, and GigaVUE-HC3.

Refer to the following limitations:

- 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.
- Keys are module-protected. With module-protection, the application is a registered client that does not need to log in to the HSM.
- 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.
- Only IPv4 addresses are supported.
- Each GigaSMART card that interfaces with the Thales HSM will use one Thales license.

- Clustering is not supported.

Refer to the following sections for details:

- *Configure HSM* on page 1182

Configure HSM

This section provides information about the steps to configure HSM. The following topics are covered:

Create HSM Appliance

Configure at least one HSM 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 configure an HSM appliance, do the following:

1. **From the device view, select GigaSMART> Passive SSL > HSM.**
2. Click **Add**.

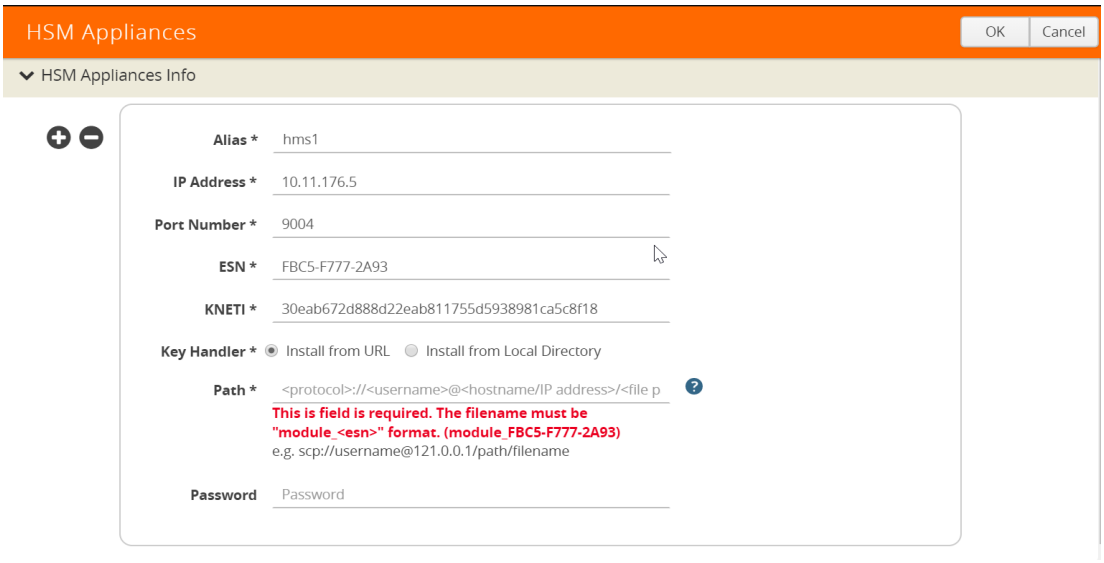

Figure 30-121: Creating New HSM

3. Click **Add**.
4. Type **hsm1** in the Alias field.
5. Enter valid **IP address**.
6. Enter Port number.
7. Type **ESN**
8. Type **KNET**
9. Select **Key Handler >Install from URL.**
10. Type the Path of the Key handler file.
11. Click **OK**.

Configure Set Key Handler

1. From the device view, select GigaSMART**> Passive SSL > HSM**.
2. Select the **HSM appliance** you just created.
3. Click **Configure**.

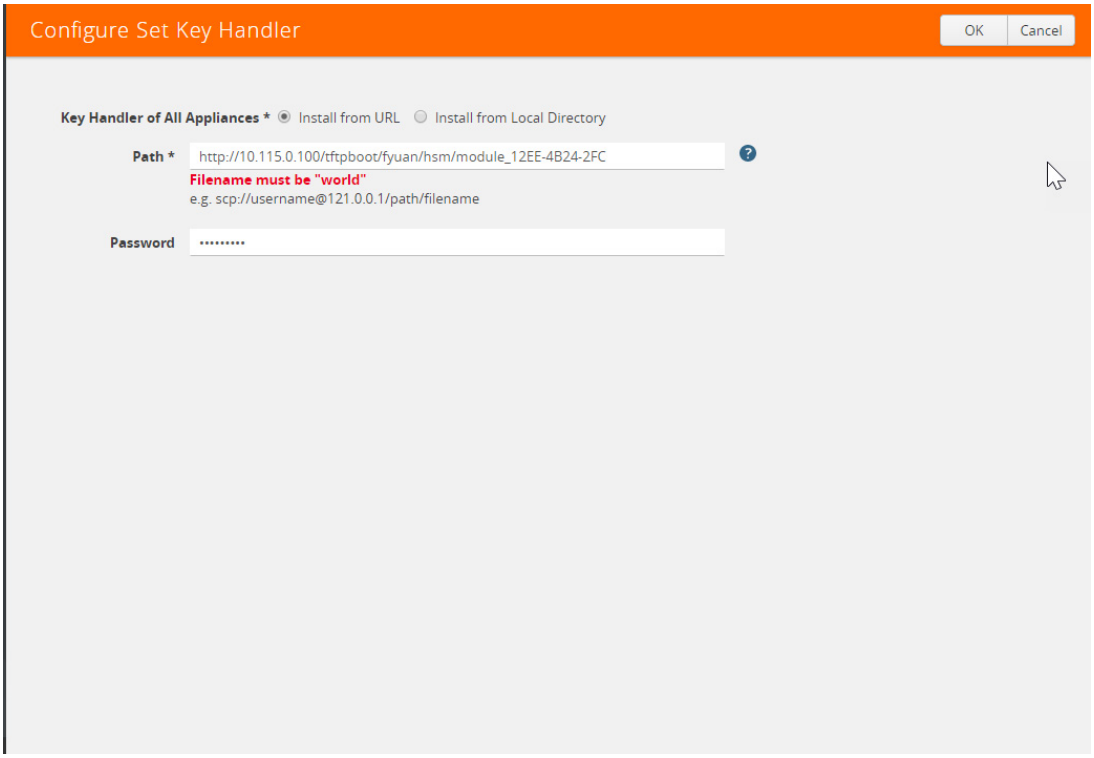

Figure 30-122: HSM-Configure Key Handler

4. Type the Path. Path filename must include world.
5. Click **OK**.

Configure Passive SSL Network Access

In this step you need to configure Passive SSL Network Access along with an valid IP address for GigaSMART.

1. From the device view, select GigaSMART**> Passive SSL > Network Access.**
2. Select the **GigaSMART appliance.**
3. Click **Edit**.

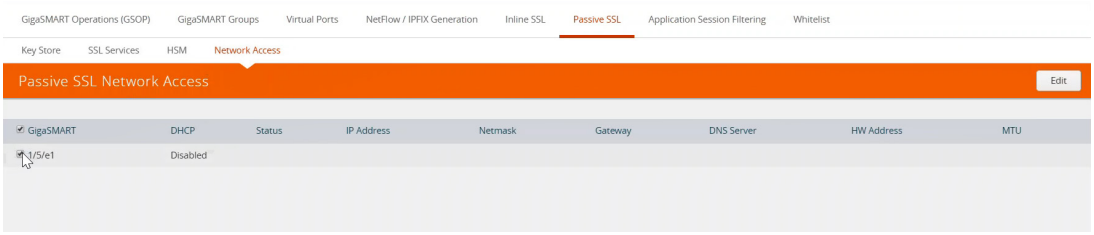

Figure 30-123: Passive SSL Network - Edit

4. Select **IP Address**

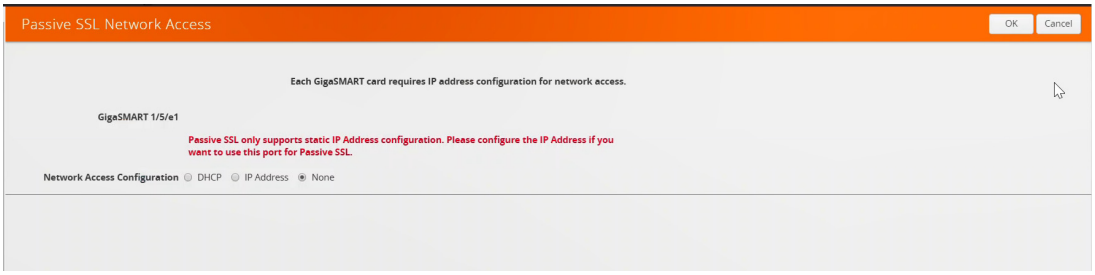

Figure 30-124: Passive SSL Network Access

5. Click **OK**.

Upload SSL Private Keys

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.

Each GigaSMART card requires IP address configuration for network access. To configure IP address details, do the following:

1. Enter **IP Address, Network, Gateway, DNS, MTU and VLAN** parameters.

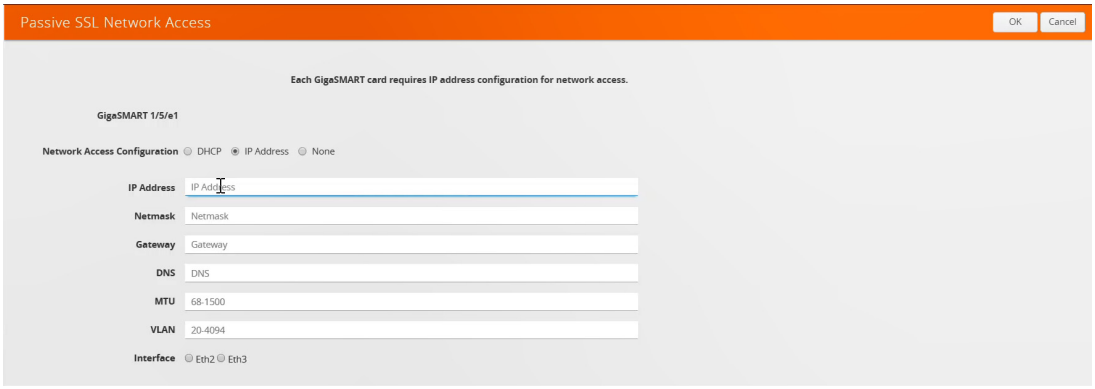

Figure 30-125: Passive SSL Network Access - IP Configuration

2. Click **OK**.

Configure GigaSMART Group

1. From the device view, select GigaSMART**> Passive SSL > HSM.**
2. Click **New**.
3. Type **GS** in Alias field.

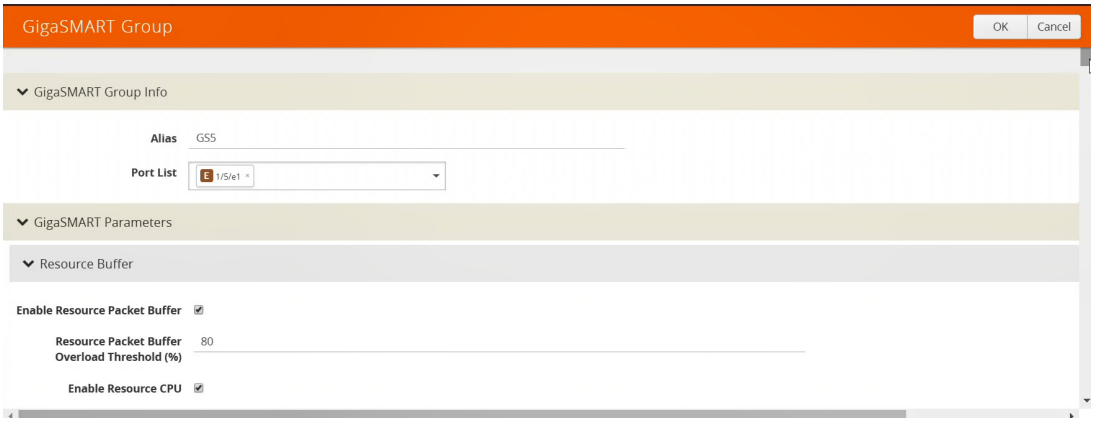

Figure 30-126: GigaSMART Group Setup Page

4. Select the **Port** you want to associate with this group.
5. Scroll down the page and click **Enable HSM**.
6. Click **OK**.

Create GigaSMART Operations (GSOP)

To create a GigaSMART operation with an SSL Decryption component, do the following:

1. From the device view, select **GigaSMART > GigaSMART Operations (GSOP) > GigaSMART Operation**.
2. Click **New**.
3. In the Alias field, enter **hsm** for the GigaSMART Operation.
4. From the GigaSMART Groups list, select a **GigaSMART group**.
5. From the GigaSMART Operations (GSOP) list, **select SSL Decryption**.

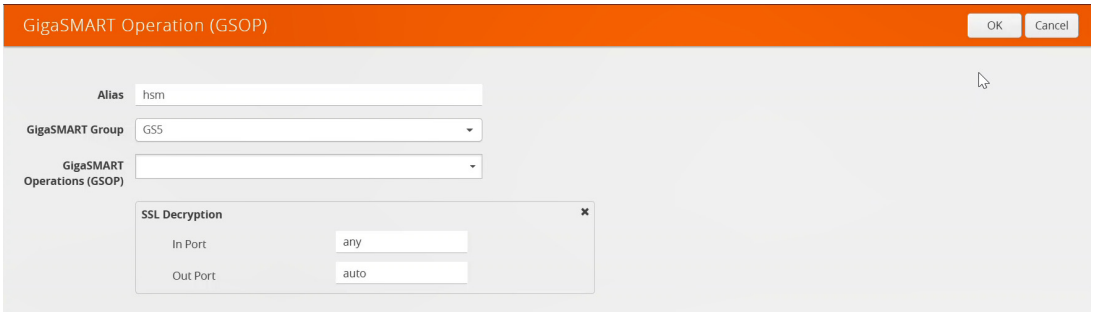

Figure 30-127: GigaSMART Operations - Setup Page

6. Click **OK**.

Configure Keys Residing on HSM

Before uploading keys or configuring SSL, you must create an SSL keychain password. The password is used to encrypt the private keys that you upload to the node.

NOTE: When uploading SSL keys, make sure that you are not creating a duplicate key. Adding a duplicate key can cause errors.

To create an SSL keychain password, use the following steps:

1. From the device view, select **GigaSMART > Passive SSL > Key Store.**
2. Click **Keychain** Password.

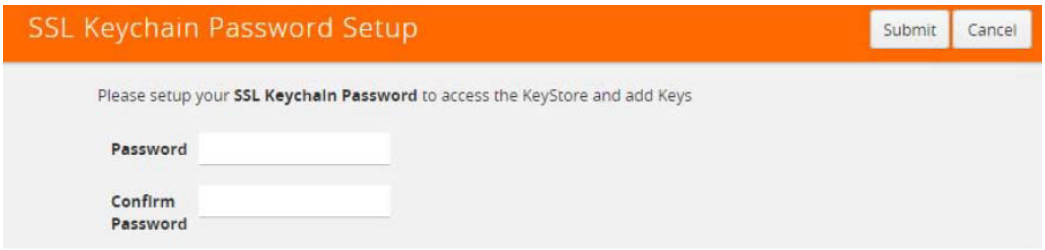

Figure 30-128: SSL Keychain Password Setup Page

3. Enter a **password** in the Password and Confirm Password fields. You can only configure a strong password. A strong password has at least ten (10) characters and at least three (3) of the following:
   - uppercase letters
   - lowercase letters
   - numbers
   - special characters
4. Click **Submit**.

Upload SSL Private Keys

To upload an SSL private key, do the following:

1. From the device view, select **GigaSMART >**Passive SSL > Key Store to open the Key Store page.

Figure 30-129: SSL Key Store Page

2. Click **Install**.
3. The SSL Key page displays.

Figure 30-130: SSL Key Page

4. Enter an **Alias** for the SSL key in the Alias field.
5. Select the **upload type.** Keys residing on the HSM, the type selected must be **PKCS #11**.
6. Choose the file. The URL can be downloaded using **HTTP, HTTPS, FTP, TFTP, SCP, and SFTP**. Using a secure protocol, such as **HTTPS** is recommended.
7. Click **Save**.

Configure GigaSMART Operation for Out-of-Band SSL Decryption

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 a service, do the following:

1. From the device view, select **GigaSMART > Passive SSL > SSL Services.**
2. Click **New**.

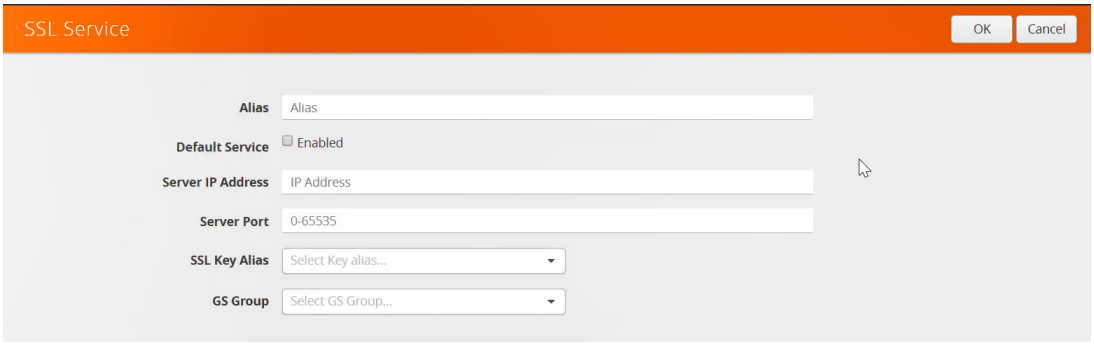

Figure 30-131: SSL Service

3. On the SSL Service configuration page, do the following:

- Enter an **alias**.
- Enter the information for the service: **IP Address, Server Port.**
- Select the **alias** of SSL Key previously uploaded.
- Select the **GigaSMART Group** with SSL decryption enabled to associate with this SSL service.

4. Click **OK**.

Configure Maps

1. Select **Maps > Maps.**
2. Click **New**.
3. Configure the map.

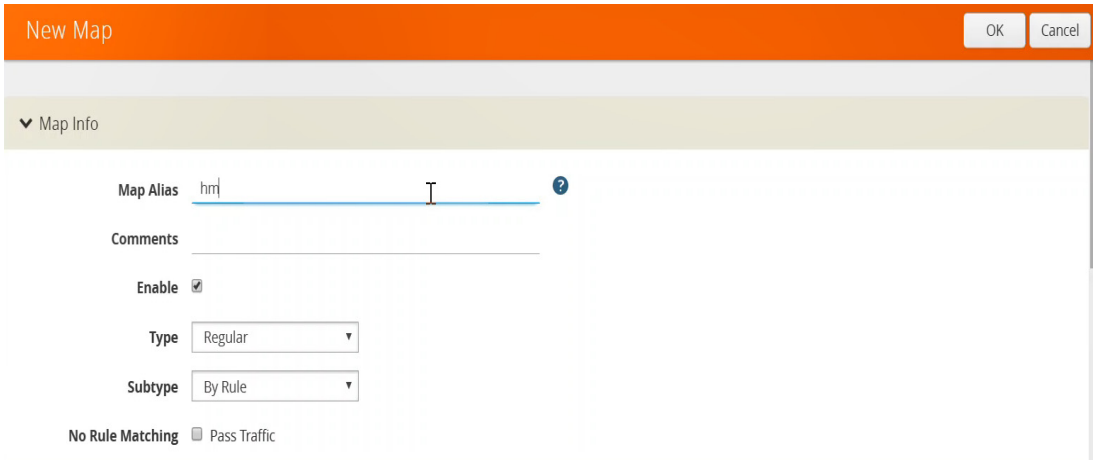

Figure 30-132: Create New Map

- Type **map11** inthe Alias field.
- Select Regular for **Type**.
- Select **By Rule** for **Subtype**.
- Select the network port for the Source.
- Select **Tool port/Hybrid port** for Destination.

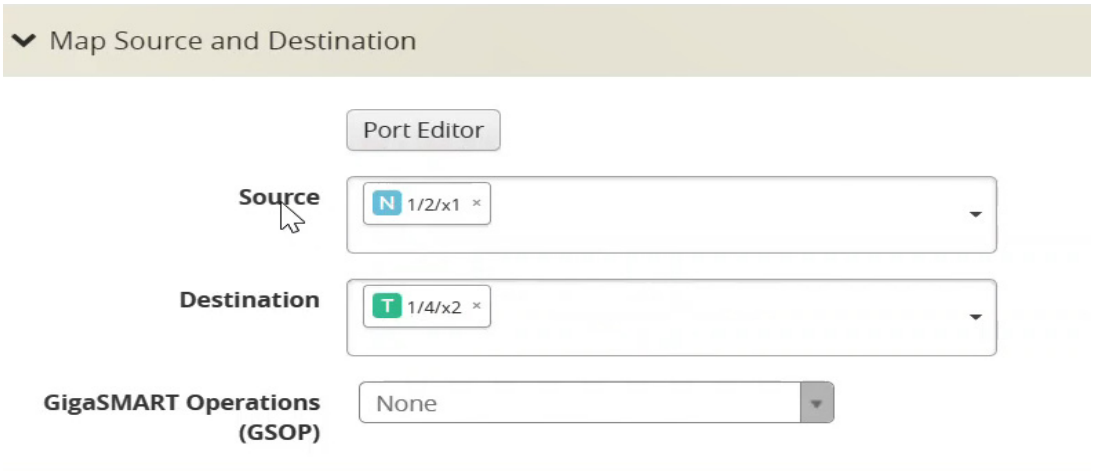

Figure 30-133: Configure Map Details

4. Add a **Rule**.

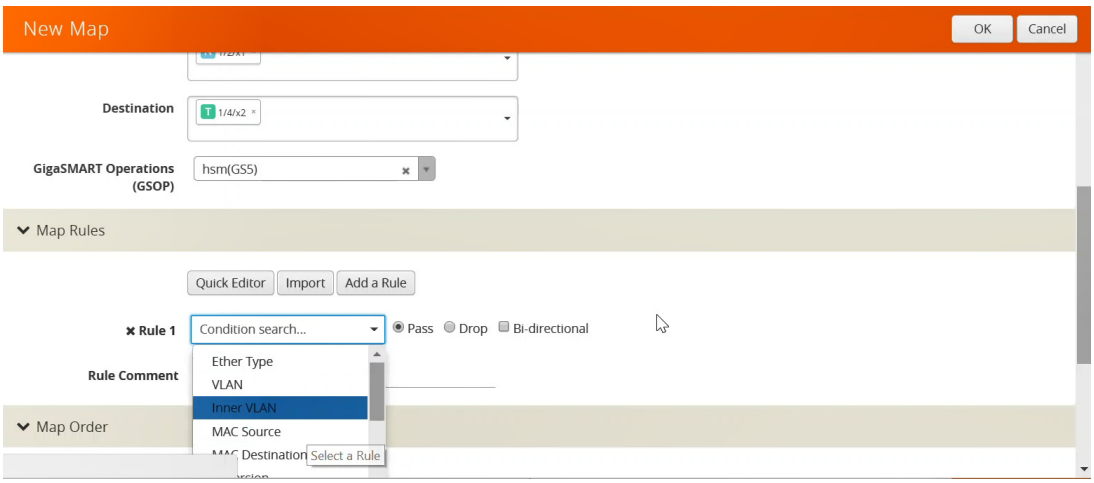

Figure 30-134: Figure 20-123: Map Details - Create Rule

a. Click **Add a Rule**.

b. Select **Pass**.

c. Select **IPv4 Version** and set **Version to v4.**

4. Click **Save**.

GigaSMART SSL Decryption for Inline and Out-of-Band Tools

Required License: SSL Decryption for Inline and Out-of-Band Tools

SSL decryption for inline and out-of-band tools is described in the following document: *Inline SSL Decryption Guide*. It is only supported on GigaVUE-HC2 and GigaVUE-HC3 in this software version.

GigaSMART Trailers

Required License: Base

GigaSMART operations can add the GigaSMART Trailer to packets, providing metadata on the packet and how it was processed.

GigaSMART Trailers are optional for some GigaSMART operations. For example, trailers can be included with Masking but not with Slicing. Refer to *How to Combine GigaSMART Operations* on page 778 for the valid combinations.

If a trailer is included, it can optionally include the original packet's CRC as one field and a Source ID as another. The Source ID indicates where the packet entered the GigaVUE H Series node and how it was processed. Refer to *About Source ID Field* on page 1191 for information included in the Source ID field.

Trailer operations can be assigned to GigaSMART groups consisting of multiple engine ports. Refer to *Groups of GigaSMART Engine Ports* on page 757 for details.

About Source ID Field

When you enable the **Source ID** field for the GigaSMART Trailer, the trailer includes an additional field that identifies the platform type, box ID, slot number, and port number of the network port where the packet entered the GigaVUE H Series node. Refer to *Format of GIGAMON_SRCID TLV* on page 1199 for the exact details.

Keep in mind the following when configuring GigaSMART operations with a **Source ID** argument:

Feature	Description
Source ID Field in GigaSMART Trailer	The **Source ID** field in the Gigamon Trailer includes the following values: • **Platform Type –** The type of GigaVUE node where the packet was first seen. • **Group ID** – The cluster ID/group ID configured for the node on which the packet was received. The GigaVUE H Series uses cluster IDs; the GigaVUE G Series uses group IDs. • **Box ID** – The box ID configured for the GigaVUE node on which the packet was received. Box IDs are used for unique identification of nodes in a cluster. **NOTE:** The box ID field in the Gigamon Trailer supports box ID values from 0-63, inclusive. • **Slot ID** – The slot ID for the port on which the packet was received. • **Port ID** – The physical port number on which the packet was received.

Example – GigaSMART Source Labeling with GigaSMART Trailer

This example creates a GigaSMART operation named **src_headermask** with **Masking** and **Trailer** components. This operation will mask packets using a static masking offset of 148 bytes that continues for the next 81 bytes, writing over the existing data with an FF pattern. Then it attaches a GigaSMART Trailer indicating the original size of the packet before masking, the original packet's CRC, and the box ID, slot ID, and port ID of the physical input port on the GigaVUE H Series node. Figure 30-135shows the GigaSMART operation with masking an trailer components.

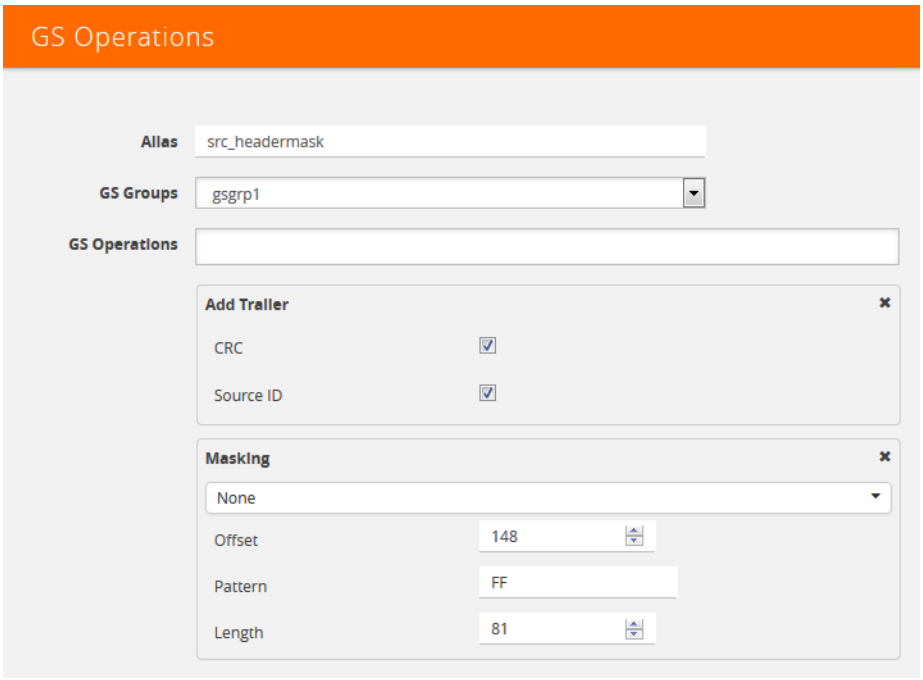

Figure 30-135: GigaSMART Operation with Masking and Trailer Components

Remove GigaSMART Trailers

You can also construct GigaSMART operations that remove the GigaSMART Trailer from packets. These operations are useful in cases where you have cascade connections – a tool port receiving packets with a GigaSMART trailer is physically cabled to a GigaVUE H Series network port, sending the packets received on the tool port back into a GigaVUE H Series node. You may want to remove the GigaSMART trailer before the packets are forwarded to other tools – that is when the special **Remove Trailer** argument comes in handy. Figure 30-136 shows Remove Trailer enabled for a GigaSMART operation.

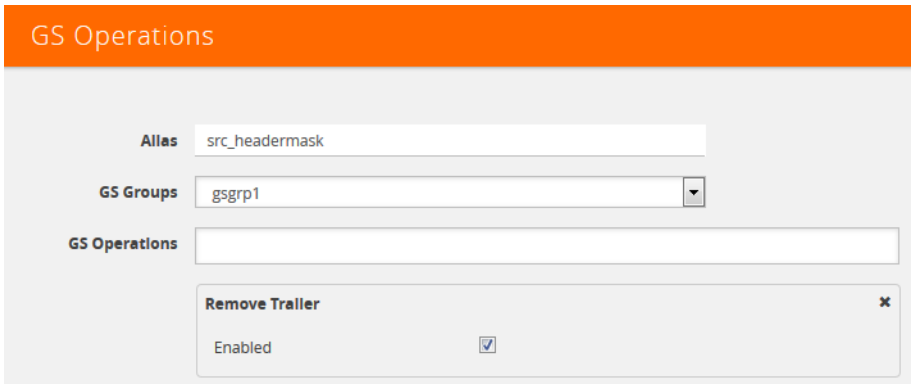

Figure 30-136: Remove Trailer Enabled

Example: Remove GigaSMART Trailers in a Cascade

Figure 30-137 and Figure 30-138 illustrate a situation where a GigaSMART operation that removes the GigaSMART trailer would be useful. Consider the physical deployment shown in Figure 30-137:

- The green illustrates a one-way cascade between a tool (output) port (port 1/1/x12) and a network (input) port (port 1/1/x9).
- If traffic arriving on port 1/1/x12 includes a GigaSMART Trailer, you may want to use a **Trailer Remove** GigaSMART operation to remove it before logically forwarding traffic from port 1/1/x9 to another port on the GigaVUE.

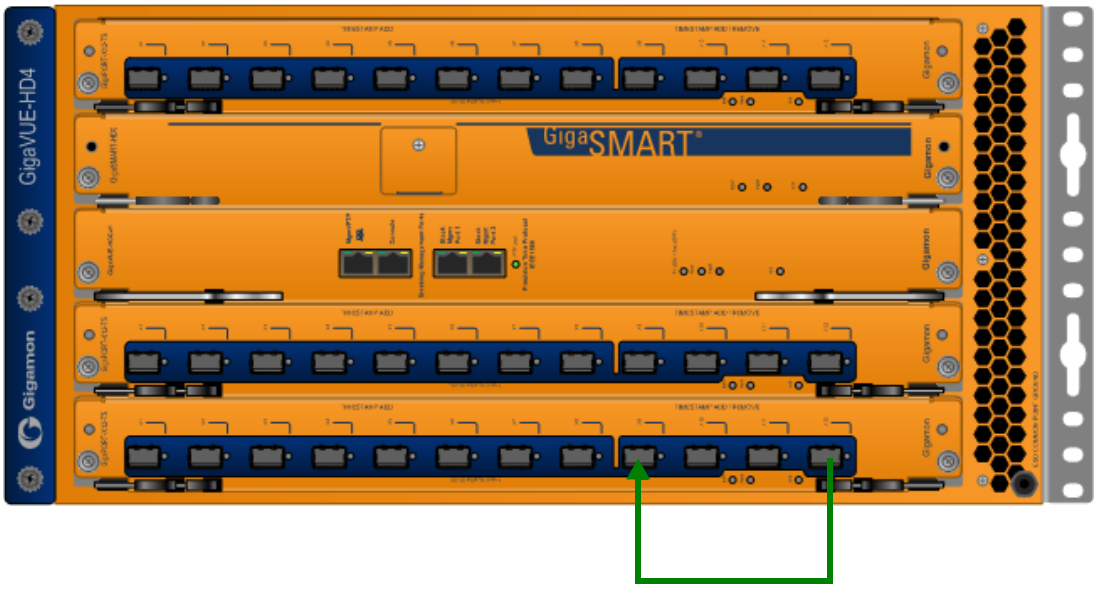

Cascade physically cabled from tool port 1/1/x12 to network port 1/1/x9 feeds traffic back into the GigaVUE H Series node for logical distribution.

Figure 30-137: Cascade Physically Cabled from Tool to Network Port

For example, consider the packet distribution shown in Figure 30-138:

- The map named **add_trailer** is bound to network port 1/1/x1..x2. It adds the GigaSMART trailer and sends it to tool ports 1/1/x5..x8.
- Tool ports 1/1/x5..x7 are all connected to tools that expect the extra data in the GigaSMART Trailer.
- Tool port 1/1/x8 is physically cabled to network port 1/1/x4 in a cascade. To remove the GigaSMART Trailer from packets arriving on this port before they are forwarded to tool port 1/1/x9, we have bound a map named **no_trailer** to network port 1/1/x4 that is configured to remove the GigaSMART Trailers from all arriving packets.
- Tool port 1/1/x9 receives packets without the GigaSMART Trailer attached.

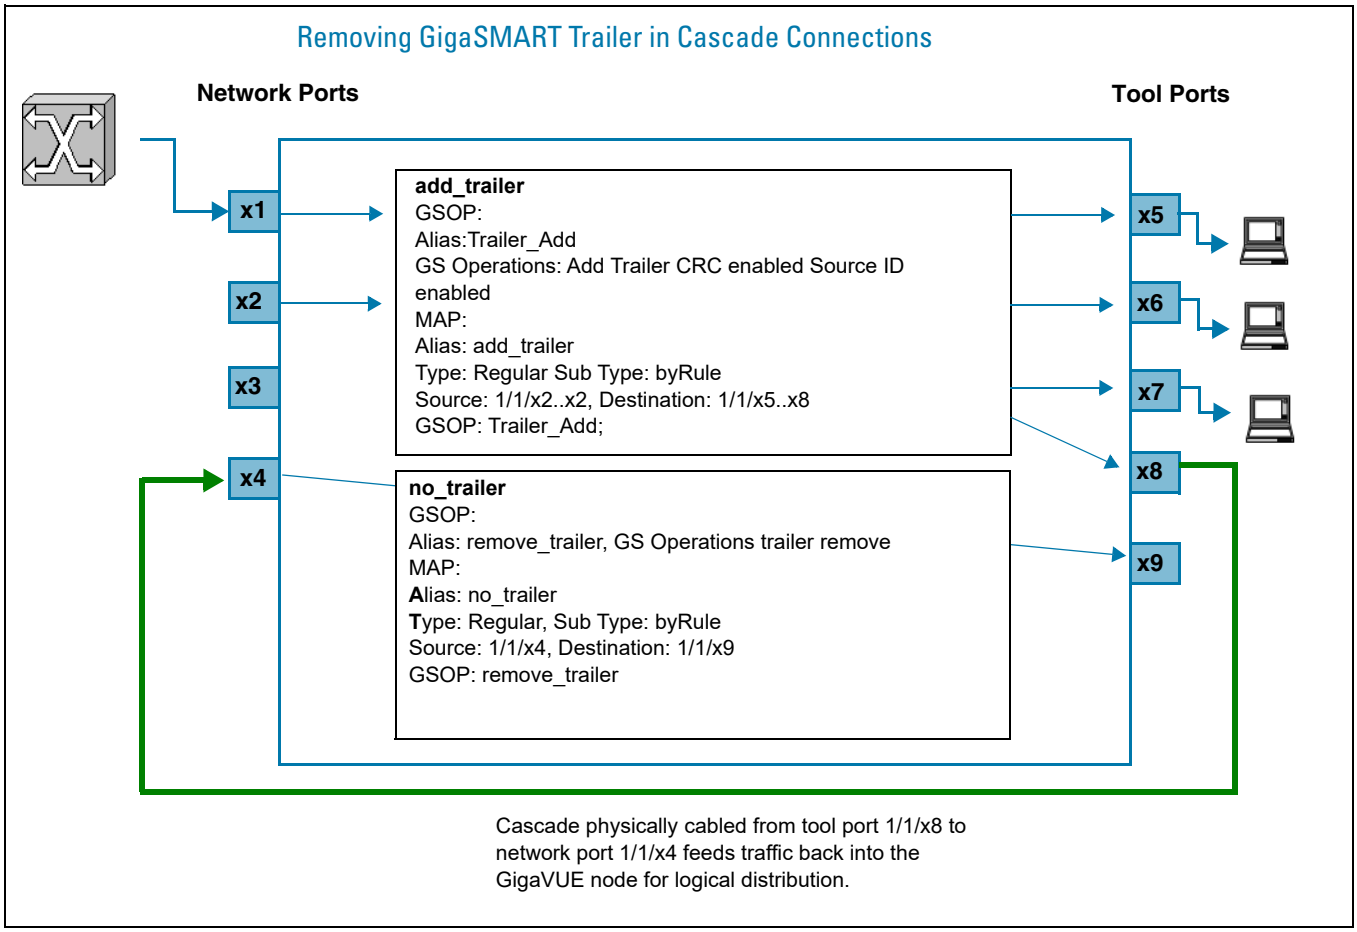

Figure 30-138: Removing the GigaSMART Trailer from a Cascade Connection

Multiple GigaSMART Trailers and Cascade Connections

Cascade connections also make it possible for multiple GigaSMART trailers to be attached to the same packet. For example, consider the cascade shown in Figure 30-138 and suppose that instead of the **no_trailer** map removing the GigaSMART Trailer on packets arriving over the cascade physically cabled from tool port 1/1/x8 that there is a second GigaSMART operation adding another trailer. In cases like this, the GigaSMART adds the most recent trailer at the end of the packet.

The same principle works for the **Remove Trailer** operations:

- The most recent trailer is removed from the end of the packet. Any other trailers are left intact by a single Remove operation.

How to Interpret GigaSMART Trailer

The trailer inserted by the GigaSMART line card can be interpreted using a recent version of the Wireshark® Protocol Analyzer. Refer to *GigaSMART Trailer Reference* on page 1196 for details on the GigaSMART Trailer and its TLVs.

GigaSMART Trailer Reference

This section provides reference information on the format, position, and contents of the Gigamon Ethertype and GigaSMART Trailer fields in a packet processed by the GigaSMART-HD0 line card.

Refer to *GigaSMART Trailers* on page 1191 for details on how the GigaSMART Trailer is used in packets.

GigaSMART Trailer Format

This section describes the format of the GigaSMART Trailer. Figure 30-139 summarizes the position and contents of the GigaSMART Trailer and the Gigamon Ethertype field (0x22E5).

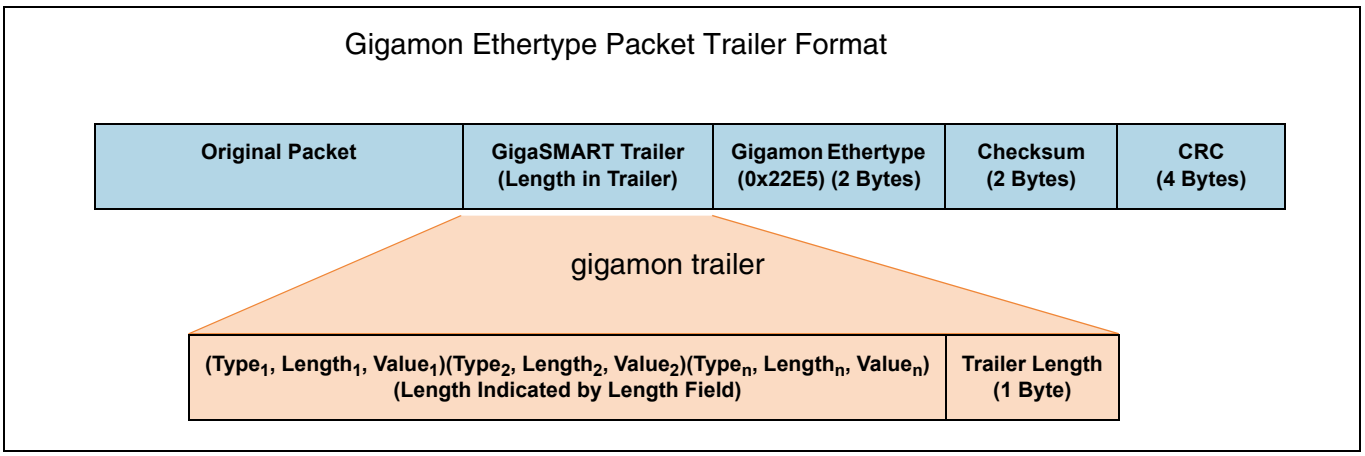

Figure 30-139: GigaSMART Trailer Format

GigaSMART Trailer Format

GigaSMART operations can insert metadata as a trailer at the end of the packet. As shown in the following figure, the GigaSMART Trailer consists of one or more TLVs, followed by the Trailer Length Field. The GigaSMART Trailer is followed by the Gigamon Ethertype, a Checksum, and a recalculated CRC for the packet.

$(Type_1, Length_1, Value_1)(Type_2, Length_2, Value_2)(Type_n, Length_n, Value_n)$ (Length Indicated by Length Field)	Trailer Length (1 Byte)	Gigamon Ethertype (0x22E5) (2 Bytes)	Checksum (2 Bytes)

Table 30-7 lists and describes each of the fields in the GigaSMART Trailer:

Table 30-7: Gigamon Trailer Format

Field	Description
TLVs	TLVs are used to carry the metadata inserted by GigaSMART operations. Within a TLV, the Type and Length fields are each one byte long. The size of the Value field is indicated by the Length field. Refer to *GigaSMART Trailer TLVs* on page 1197 for a summary of the available TLVs and their formats.
Length	Specifies the size of the GigaSMART trailer. The Trailer Length includes the Length byte itself.
Gigamon Ethertype	Two-byte field identifying packet modified by GigaSMART line card. Value is 0x22E5. Protocol analysis equipment can use the Gigamon Ethertype in the trailer to find the start of the trailer: • Move one byte to left to find Trailer Length Field. • Read the number of bytes for the Trailer Length and move the number of bytes specified to the left to find the start of the Trailer.
Checksum	Two-byte field used to validate that extra data is in fact a trailer and not random data with the Gigamon Ethertype.

GigaSMART Trailer TLVs

This section lists and describes the format of the Gigamon TLVs used in this release. TLVs are used to carry the metadata inserted by GigaSMART operations. Within a TLV, the Type and Length fields are each one byte long. The size of the Value field is indicated by the Length field.

Table 30-8: Gigamon TLVs

Tag Type	TLV ID	Value Field Length (Bytes)	Description
GIGAMON_PKT_LEN	1	2	**Original Packet Length** This TLV is included in any packet with a GigaSMART Trailer – adding the Trailer changes the original packet length.
GIGAMON_SRCID_G	2	3	**Original Packet Source Identifier (GigaVUE G Series)** This TLV is included in any packet processed by a GigaSMART operation configured to include the Source ID field as part of its trailer. Refer to *Format of GIGAMON_SRCID TLV* on page 1199 for a description of how the physical input source is encoded.

Table 30-8: Gigamon TLVs

Tag Type	TLV ID	Value Field Length (Bytes)	Description
GIGAMON_TIMESTAMP	3	8	**Timestamp – System Clock** **Timestamp** A Timestamp TLV is included in any packet processed by a timestamp GigaSMART operation. The Timestamp always indicates the UTC time at which the last byte of the packet was seen by the GigaSMART line card or module. In addition, the exact type of Timestamp TLV indicates the time source used (NTP, GPS, and so on). Note that timestamp GigaSMART operations are not yet supported in this release of the GigaVUE H Series node.
GIGAMON_TIMESTAMP_NTP	4	8	**Timestamp – NTP Clock Source**
GIGAMON_TIMESTAMP_GPS	5	8	**Timestamp – GPS Clock Source**
GIGAMON_TIMESTAMP_1588	6	8	**Timestamp Based on PTP (IEEE 1588)**
GIGAMON_CRC	7	4	**Original Packet CRC** This TLV is included in any packet processed by a GSOP configured to include the original packet's CRC as part of its trailer (the **CRC** option is selected in the **Trailer Add** argument).
GIGAMON_SRCID	8	4	**Original Packet Source Identifier (GigaVUE H Series)** This TLV is included in any packet processed by a GigaSMART line card on a GigaVUE H Series node configured to include the Source ID field as part of its trailer. Refer to *Format of GIGAMON_SRCID TLV* on page 1199 for a description of how the physical input source is encoded.
GIGAMON_TIMESTAMP_1588	6	8	**Timestamp – PTP (IEE 1588) Clock Source**

Format of GIGAMON_SRCID TLV

The GIGAMON_SRCID TLV is included in any packet processed by a **Source ID** GigaSMART operation. The GIGAMON_SRCID TLV consists of 3 bytes indicating the platform type, group ID, box ID, and port ID for the physical port where the packet entered the GigaVUE H Series node:

GIGAMON_SRCID TLV (32 Bits/4 Bytes)				
Platform (6 Bits)	**Group ID (4 Bits)**	**Box ID (6 Bits)**	**Slot ID (6 Bits)**	**Port ID (10 Bits)**

Name	Description	Bits
Platform	The type of GigaVUE node where the packet was first seen. Can be one of the following: • 0 – Unknown • 1 – GigaVUE-2404 • 2 – GigaVUE-420 • 3 – GigaVUE-MP • 4 – GigaVUE-212 • 5 – GigaVUE-HB1 • 6 – GigaVUE-HC2 • 7 – GigaVUE-TA1 • 8 – GigaVUE-TA10 • 9 – GigaVUE-TA40 • 10 – GS_CHASSIS_TYPE_LY2 (white box) • 11 – GigaVUE-TA100 • 12 – GigaVUE-TA100 CXP • 13 – Reserved for internal use • 14 – GigaVUE-TA200 • 15 – Reserved for internal use • 16 – GigaVUE-HC1 • 17 – GigaVUE-HC3	6
Group ID	Group IDs are used to identify a particular cluster. If a packet enters a cluster with a group ID configured, it is reflected here.	4
Box ID	Box IDs are used to uniquely identify nodes in a cluster. Standalone systems typically have the default box ID of **1** here. **NOTE:** The box ID field in the Gigamon Trailer supports box ID values from 0-63, inclusive.	6
Slot ID	The number of the slot including the physical port for the packet.	6
Port ID	The physical input port number for the packet. Refer to *Port ID Values by Line Card Type* on page 1200 for a summary of the values used by GigaVUE H Series line card.	10

Port ID Values by Line Card Type

The following table summarizes the values inserted in the GIGAMON_SRCID TLV for port ID by line card.

Line Card	Ports	Port IDs Inserted in Gigamon Trailer
PRT-H00-X12G04	x1..x12	1..12
	g1..g4	21..24
PRT-H00-X04G44	g1..g44	1..44
	x1..x4	45..48
PRT-H00-X12TS	x1..x12	1..12
PRT-H00-Q02X32 – 2q Mode	x5..x28	5..28
	q1..q2	41..42
PRT-H00-Q02X32 – 32x Mode	x1..x32	1..32
PRT-H00-Q08	q1..q8	1..8
GigaVUE-HB1 Node	g1..16	1..16
	x1..x4	17..20
PRT-HC0-X24	x1..x24	1..24
PRT-HC0-Q06	q1..q6	1..6

GigaSMART Trailer Example

The following figures show a sample GigaSMART trailer with the packet length and source ID TLV included. The total length of the trailer in this example is 10 bytes, plus another 5 bytes for the Trailer Length, Ethertype, and Checksum fields.

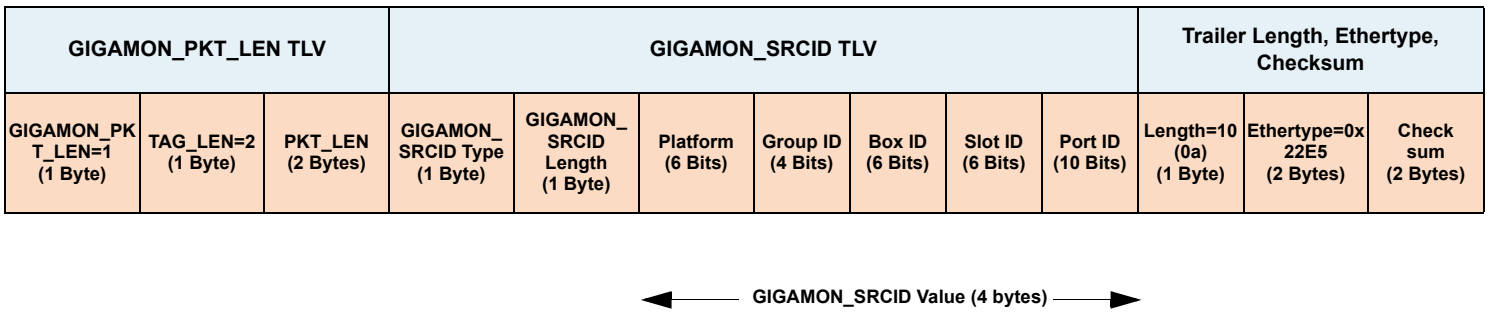

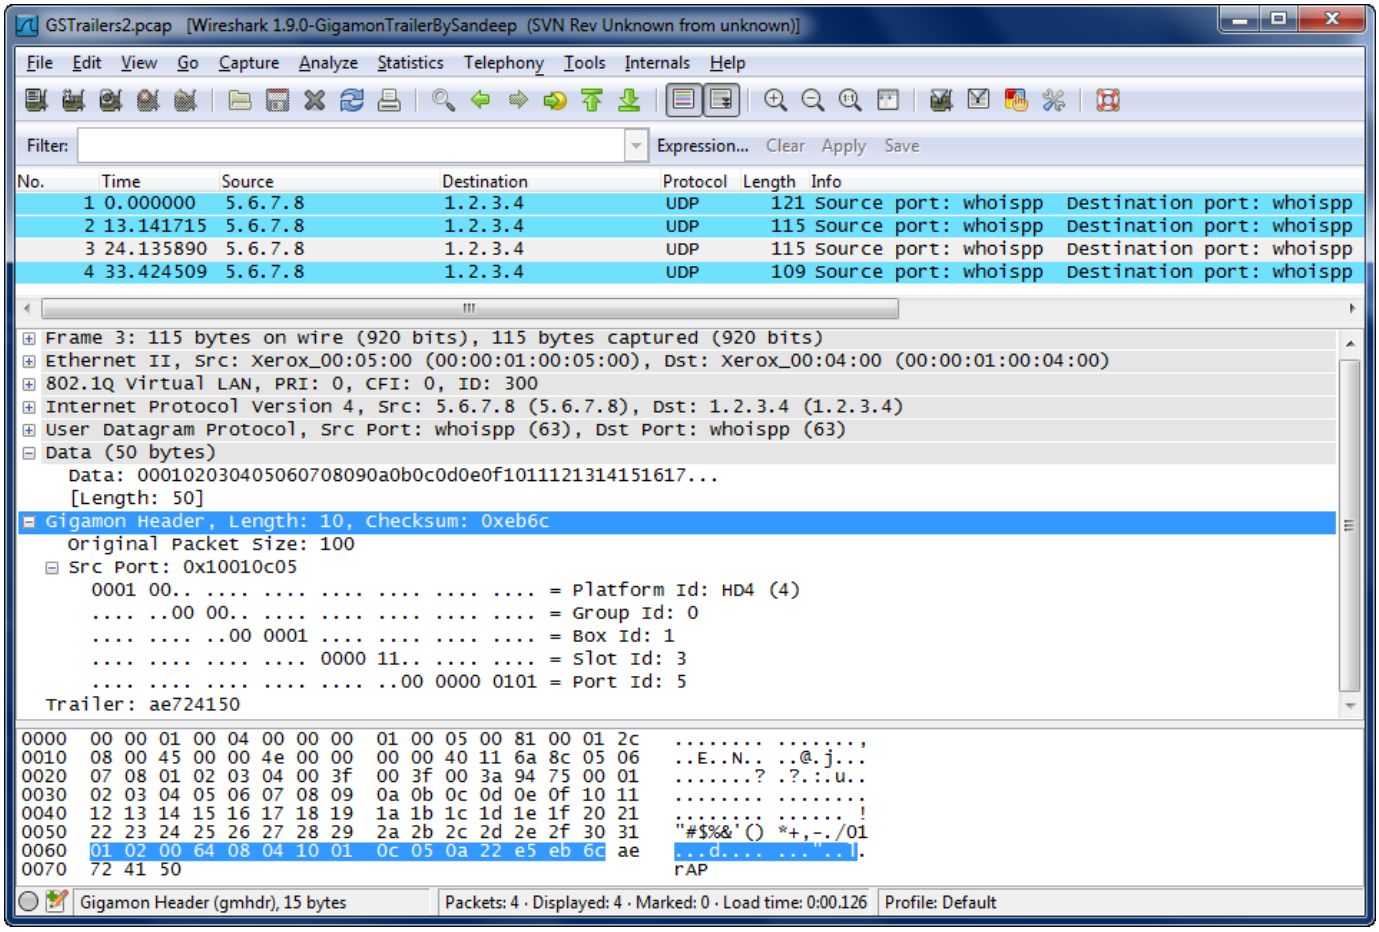

Figure 30-140: Wireshark Decode of GigaSMART Trailer

31 GigaSMART Logs

As of 5.4, GS Log files enable you to generate and download application-specific logging information to use for troubleshooting problematic applications. Gigamon Support can use these files for root cause analysis.

GS Log levels are applied at the process level by specifying *info* or *debug* when configuring the logging level for the application's processes. Logging levels modified at the application level will take priority over the system level setting.

To access Logs, click **Administration** on the top navigation link. On the left navigation pane, select **System > Logs**.

Create GS Log file

To create a log file that Gigamon can use for analysis, do the following:

1. Access a device: Physical > Physical Node > click the Cluster ID for a node.
2. Under Support in the navigation pane, select **Debug > GS Logs**.
3. Click **Generate**.
4. Select the Host Name and Ports.

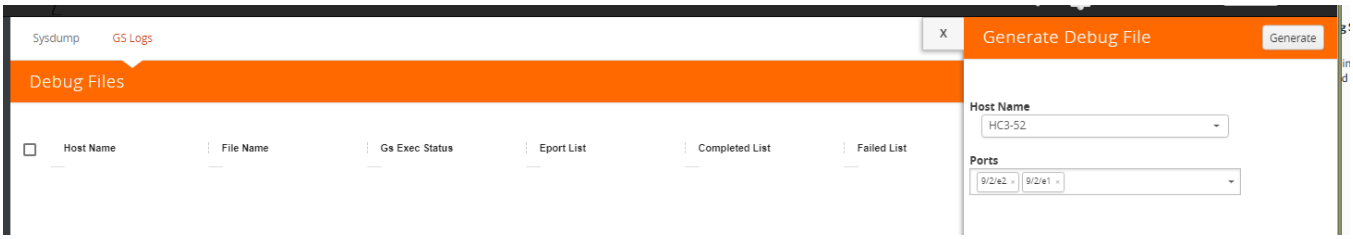

The system generates a new GS Log file.

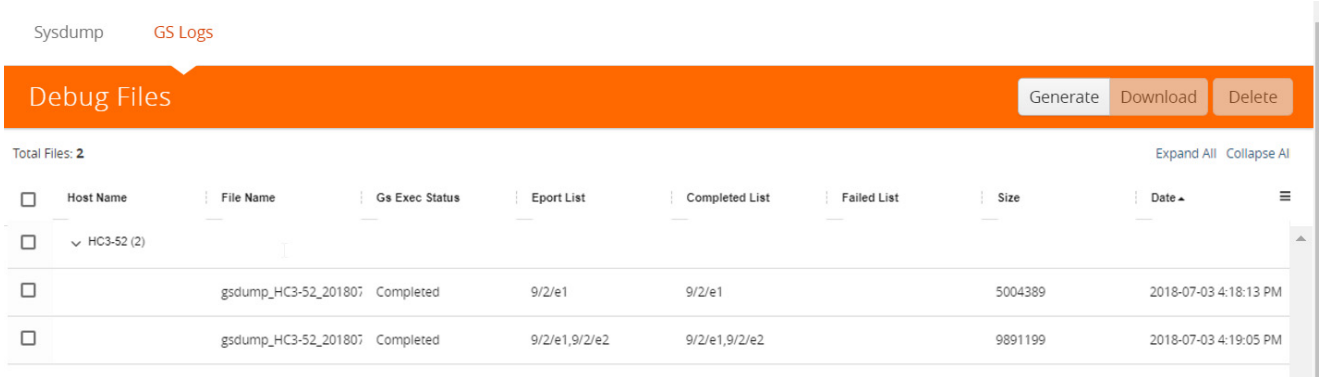

5. Select the GS Log file to download, and then click **Download**.

   You can only download one file at a time.

   The system downloads the file to your local environment with a name like, gsdump_<hostname>_<date>_<time>.tgz.gpg. The file is in a compressed and encrypted format that you can provide to Gigamon.

Delete Log File

To delete the GS Log files for clearing up the disk space:

1. Select **Debug** > **GS Logs**.
2. Select the GS Logs that you want to delete and click **Delete**.

Part 7: Fabric Maps

This section provides information about how to use GigaVUE-FM to configure fabric maps for the visibility nodes.

The following topics are covered:

32 About Fabric Maps

Fabric maps simplify how you create flow mapping across multiple clusters. Using GigaVUE-FM, you can create and manage fabric maps automatically. You can provide cross-cluster flow mapping parameters by creating a fabric map. GigaVUE-FM then allocates the required circuit ID resources, along with generating and deploying cluster specific maps and fabric paths (refer *About Circuit-ID Tunnels* on page 473). Fabric paths can be defined as the link between the devices through which the traffic flows. With a successfully deployed fabric map, you can save time and effort because you no longer need to replicate rules at each hop on the network to manage cross-cluster traffic.

Supported Topologies

Starting with software release 5.5, fabric map support is available for the following topologies.

- *Multiple Access or Aggregation Clusters* on page 1208
- *Multi-Cluster Mesh* on page 1209

Multiple Access or Aggregation Clusters

In this topology, a set of access switches is connected to network ports. This is the access cluster. The access cluster is then connected to a tool-hosting cluster using Ethernet interfaces. The tool-hosting cluster is connected to the tool ports or to GigaSMART operations (GSOP).

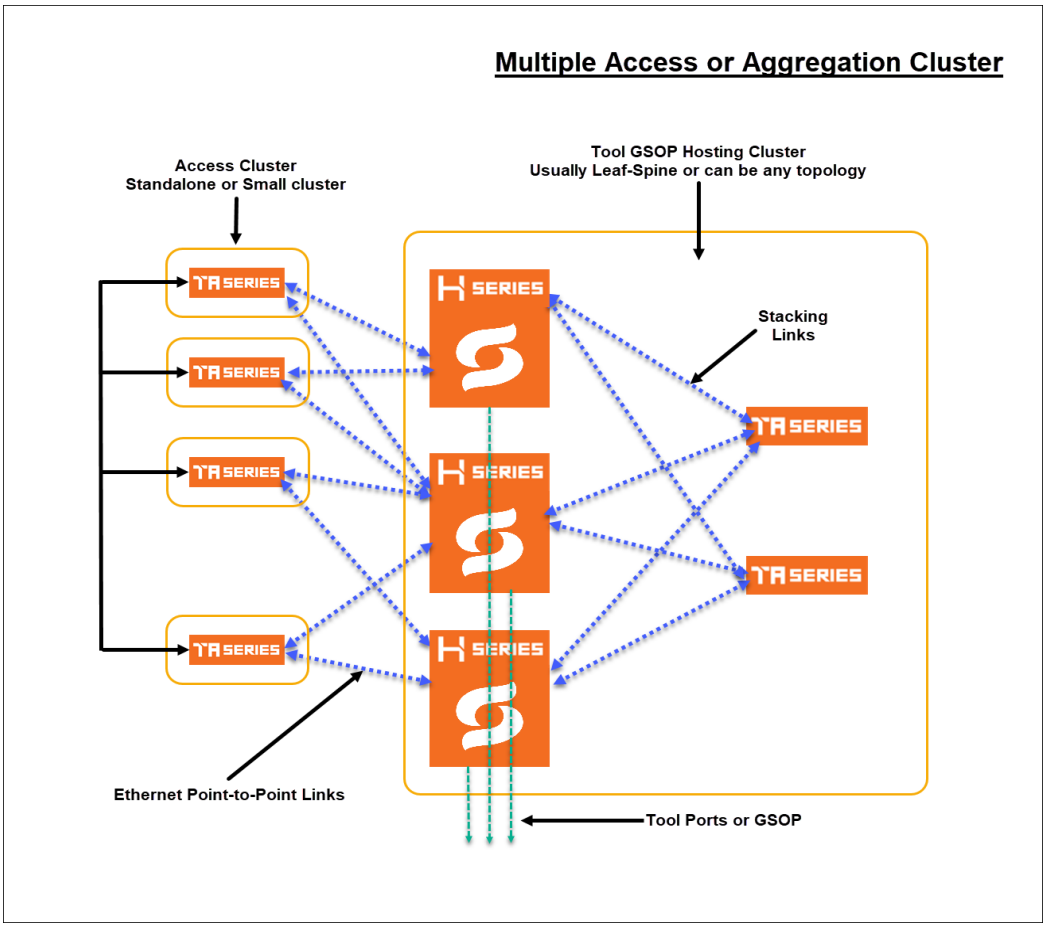

Figure 32-1: Multiple Access Clusters

The access nodes are connected to the tools cluster using point-to-point Ethernet links, which offer redundant connections to multiple nodes within the tool cluster.

This topology illustrates how ports can be created from any network port to any tool port as a single fabric-wide map.

Multi-Cluster Mesh

The following is an illustration of multi-cluster mesh topology. In this topology there are three clusters, (A, B and C) with each cluster containing network ports, tool ports and applications (GSOP).

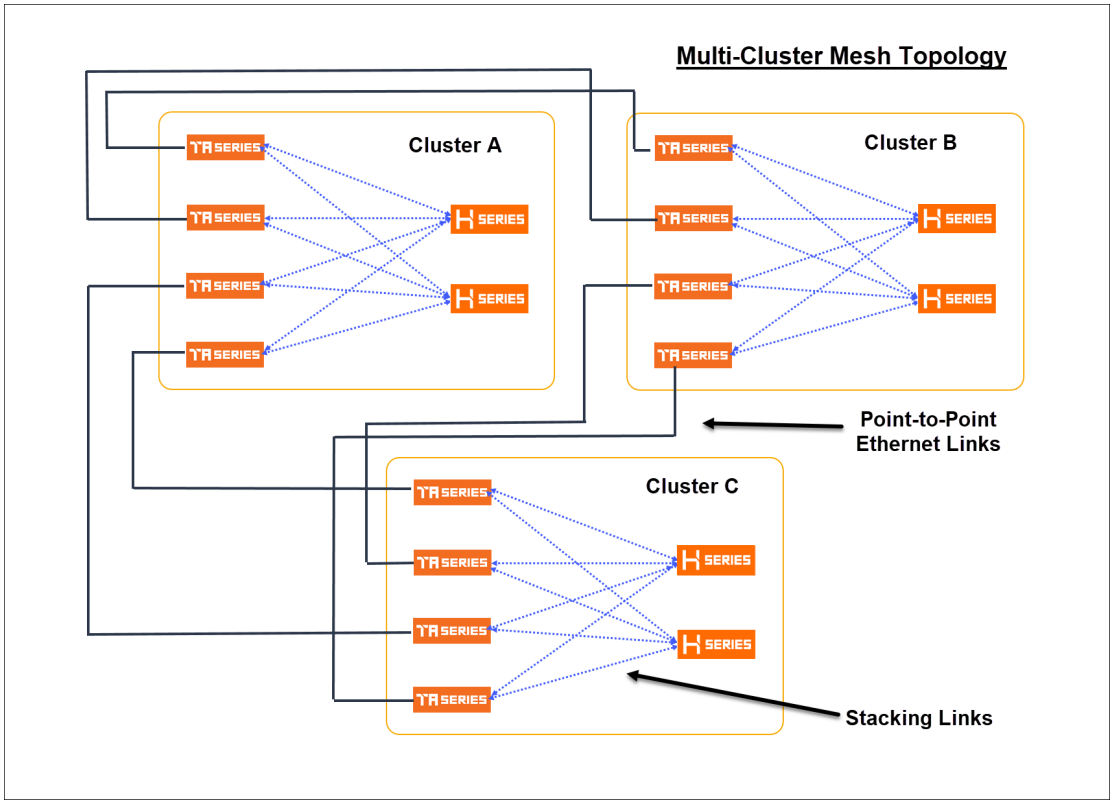

Figure 32-2: Multi-Cluster Mesh

Fabric Maps Prerequisites

To deploy fabric maps, you must complete the following prerequisites:

- Designate inter-cluster connection ports as circuit ports. Refer to *Configure Ports* on page 408.
- Configure a regular circuit GigaStream with the required circuit ports. Refer to *Configure Regular GigaStream* on page 434.
- Create links between clusters. Refer to *Create Links between Clusters* on page 1210.

Using GigaVUE-FM, you can create connections between devices using manual topology links, or you can have the devices discovered through the Gigamon Discovery (GDP). After you create physical links between devices, the device configuration needs to be added into GigaVUE-FM, so that fabric maps can use it. Once the circuit and GigaStream ports are created and connected to the devices, you can start creating fabric maps. See *Enable Gigamon Discovery on Chassis* on page 300 for more information on GDP.

Create Links between Clusters

You can either connect the devices manually or you can have them automatically discovered by Gigamon Discovery Protocol (GDP). If a user's ports are already physically connected, then the link between those ports will be displayed if GDP is enabled at both ends and the physical links are up. GDP must be enabled in both the source and destination ports to have the devices displayed in the topology map.

- *Create Manual Links* on page 1210
- *Create Gigamon Discovery Protocol (GDP) Based Links* on page 1211

Create Manual Links

To create links manually, do the following.

1. Go to **Physical >Topology**.
2. Select **Add Link(s)** option from the **Add** tab drop-down menu on the Topology menu bar.
3. Select the **Source** device and the circuit port which you already created.
4. Select the **Destination** device and the circuit port which you already created.
5. Click **Submit** to connect the devices.

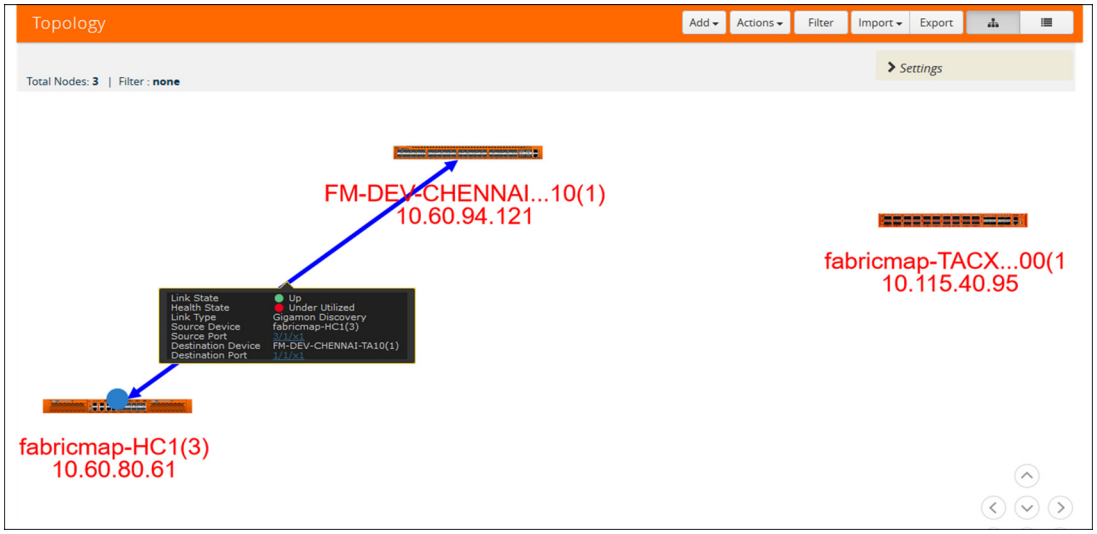

Figure 32-3: Connecting Devices

6. Click the Topology link from left-navigation menu to view the connection.

After you have successfully connected your devices, the topology map is displayed with the connections. If you click on the connected devices link, you can view the connection details of the ports you created.

Create Gigamon Discovery Protocol (GDP) Based Links

To create Gigamon Discovery Protocol (GDP) based links:

1. Go to **Physical > Physical Nodes**.
2. Select the node (Source or Destination).
3. Select > **Ports**. On the Ports screen, place a check mark next to the circuit port.
4. Click **Edit** from the top menu bar. The Port screen is displayed for the device you selected.

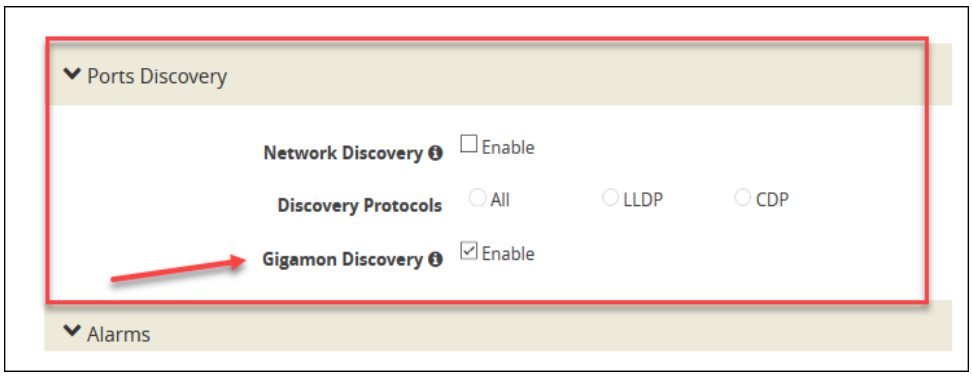

Figure 32-4: Gigamon Discovery Protocol

5. Click **OK** to submit configuration. You must then enable GDP on the device chassis as follows:

   a. Select > **Chassis**. The chassis screen is displayed showing the Gigamon hardware.

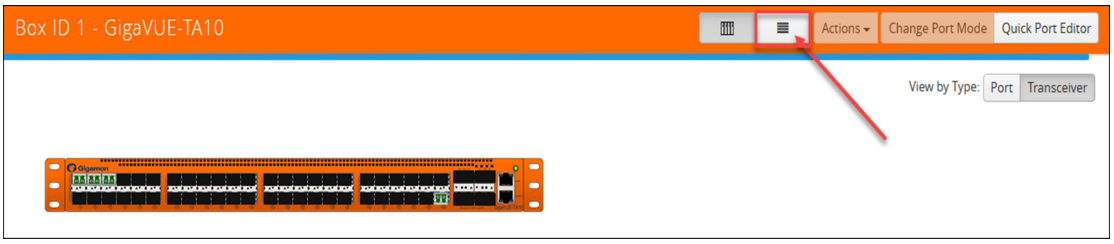

Figure 32-5: Chassis Screen

   b. Click the **List View** icon located at the top of the page. Details about the chassis is displayed in a list view format.

   c. Select **Enable Gigamon Discovery** from the **Actions** menu at the top of the screen.

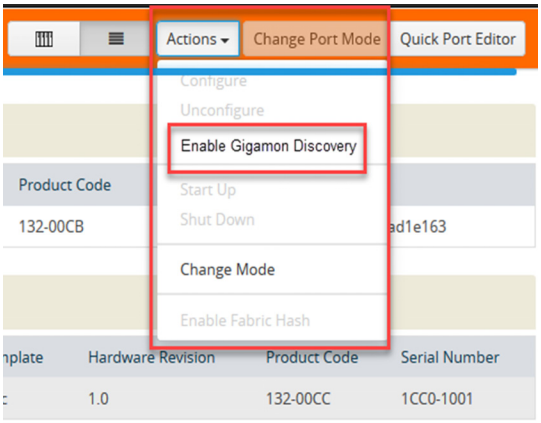

Figure 32-6: Enabling Gigamon Discovery

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

This completes the prerequisites steps for creation of fabric maps. Refer to section *Create Fabric Maps* on page 1212 for details on creating fabric maps.

Create Fabric Maps

To create Fabric Maps:

1. Click **Physical > Fabric Maps.**

The following table describes the parameters displayed in the fabric map list view.

Counter	Description
Alias	Alias name of the fabric map
Source	Number of source ports associated with fabric map

Counter	Description
Destination	Number of destination ports associated with fabric map
Rules	Rules for fabric map
GigaSMART Alias	Alias name of the GigaSMART operation
Type	Port type
Enabled	Indicates if fabric map is enabled – Yes/No
Status	Health status of the fabric map

2. Click **Create**.The create map screen is displayed.
3. Enter map parameters.
   a. Enter **Alias** name and **Description**.
   b. Select map **Type:**
   - Regular
   - First Level
   - Second Level
   c. Select subtype**:**
   - By Rule
   - Pass All
   - Collector
4. Select a GigaSMART OPERATION node and operation.
5. Enable Non-Matching Rules, if map type is regular.
6. Specify the source port. You can either:

- Enter the port alias/port id, or
- Specify the source node and source ports. Use the +/- icon to add new nodes or remove source ports.

Use the toggle bar to toggle between these two options.

7. Specify the destination port. You can either:

- Enter the port alias/port id, or
- Specify the destination node and destination ports. Use the +/- icon to add new nodes or remove destination ports.

  Use the toggle bar to toggle between these two options.

You can click the **Tool Finder** option to search for the available destination ports.

8. Click **ADD A RULE** to add a description and select one or more conditions.

NOTE: Using this option, you can only add rules one at a time. Click the **+/-** icon to add or delete rules and conditions.

9. Click the **RULES EDITOR** to add a set of rules.

Rule Type options:

Rule Type	Description
Batch Rules	Use this option to enter multiple fabric map rules at the same time.
Copying from Map Template	Copy fabric map rules from and existing fabric map template.

10. After you enter all the fabric map parameters, click **Save**. The new fabric map is added to the list view.

Edit and Delete Fabric Maps

You can edit and delete fabric maps that have been created. Refer to the following section for details:

- *Edit Fabric Maps* on page 1214
- *Delete Fabric Maps* on page 1215

Edit Fabric Maps

To edit fabric Maps:

1. Select **Physical** > **Fabric Maps**.
2. Select a fabric map. Each fabric map has a summary page. The top half of the summary page shows the fabric map component interactions and traffic flow.

The bottom half of the screen shows a list of the rules associated with the fabric map and a link to the section where you can access fabric maps statistics.

Clicking on any component in the fabric map displays more details about the component.

Edit Fabric Map Component

1. Click the **Fabric Map component** icon. The fabric map component summary screen appears.
2. Click the option menu and select **edit**.
3. Edit fabric map parameters.
4. Click **Save** to update the fabric map components.

Delete Fabric Maps

1. Select the fabric map you want to delete from the fabric maps summary page.
2. Click the **Trash Can** icon. Confirm delete notification and click **OK**.

After the fabric map is deleted you will see an acknowledgment on the screen. Refresh the fabric map summary page to confirm the fabric map is deleted.

Prioritize Fabric Maps

You can prioritize the traffic flowing through the fabric maps. The priority is handled at the cluster level.

To prioritize the traffic flow through fabric maps:

1. Select **Physical** > **Fabric Maps**.
2. Select a fabric map and click **View Details** either from the **Task** drop-down menu or from the fabric map selected.
3. Scroll down to view the current priority setting of the fabric maps.
4. Click **Rearrange Order**.
5. Use the drag and drop icon to change the order of priority.

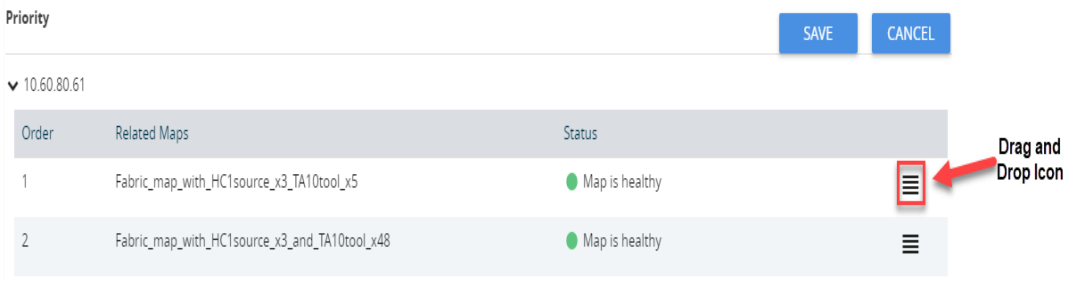

6. Click **Save** to save the changed priority.

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

Fabric Maps Statistics

GigaVUE-FM provides the ability to view detailed information about fabric maps configured between the devices including packets received and transmitted by the following ports:

- Network ports
- Tool ports
- Hybrid ports
- Circuit ports
- Inline-network ports
- Inline-tool ports
- GigaSMART engine ports

Display Fabric Map Statistics

Using the fabric map statistics page, you can view the statistics associated with the port types. If there are packet drops in the ports, you can use the statistics page to investigate the cause of the packet drops. Fabric map statistical data is generated from the rules you specify when you create a fabric map.

Display Fabric Map Details

To display the fabric map details:

1. Select **Physical> Fabric Maps.**
2. Click on a **Fabric Map** from the main page. The fabric map list view is displayed.
3. Click the **View Statistics** from the options menu. The statistics screen is displayed. Using this screen, you can view fabric map traffic data rates intervals.

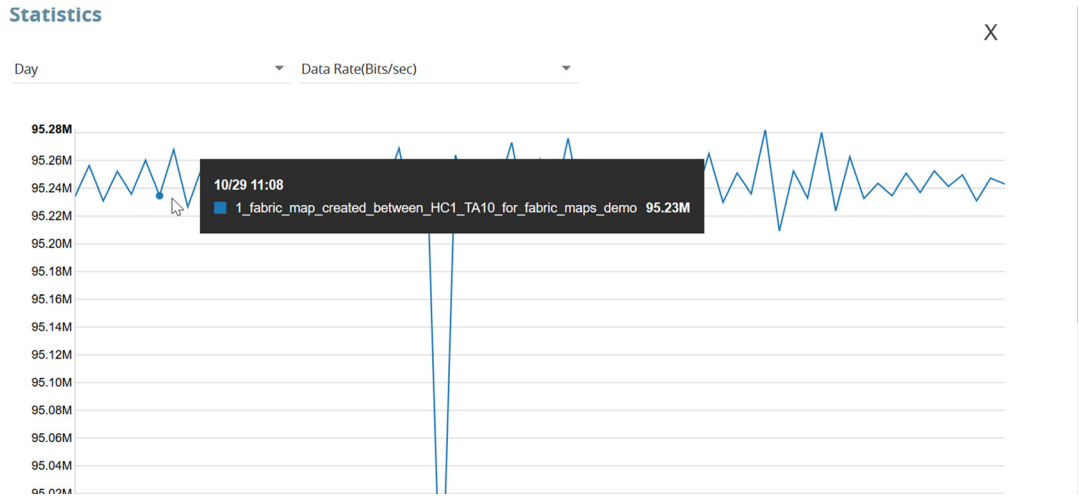

Figure 32-7: Fabric Maps Statistics

NOTE: Statistics and data rates can be displayed per hour, day, week, or month.

4. Click the current statistics label, (Hour, Day, Week or Month), to display the options menu where you can select a different statistics view.

Filter Fabric Maps List View

Using the filter function, you can search and narrow down the fabric maps you wanted to be displayed on the fabric map list view page.

1. To use the filtering functionality, click the **Filter** icon.

The Filter quick view dialog is displayed.

2. Enter the parameters to specify the data you want to display:

Fabric map filter parameter options:

Criteria	Description
Alias	Alias name of the Fabric Map
Source	Source port(s) of the Fabric Map
Destination	Destination port(s) of the Fabric Map
Status	Displays health of the fabric map. Options are: Healthy, Unhealthy, and Warning

Troubleshooting

There may be situations when fabric fails because of a configuration error or you need to investigate an unhealthy GigaStream or port associated with fabric map. Fabric map maintains the configuration status and health state.

The configuration statuses are SUCCESS, FAILED, and PARTIAL_SUCCESS. You can investigate issues related to fabric maps and once issues are corrected, you can run resubmit to push changes. Health state is consolidated from generated components healthStates.

Fabric Maps health states consist of the cluster level maps, (ports), and circuit tunnels (ports and links). After configuration and health state issues are corrected you can sync the Fabric configStatus and healthStates.

To use the troubleshooting feature:

1. Select **Physical > Fabric Maps**.
2. Select a **Fabric Map** to troubleshoot.
3. Select **View Details** from the **Task** drop-down menu.

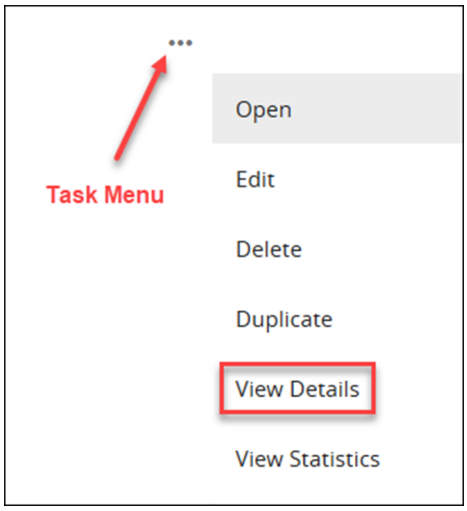

Figure 32-8: View Details.

The Details screen displays fabric map component details.

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

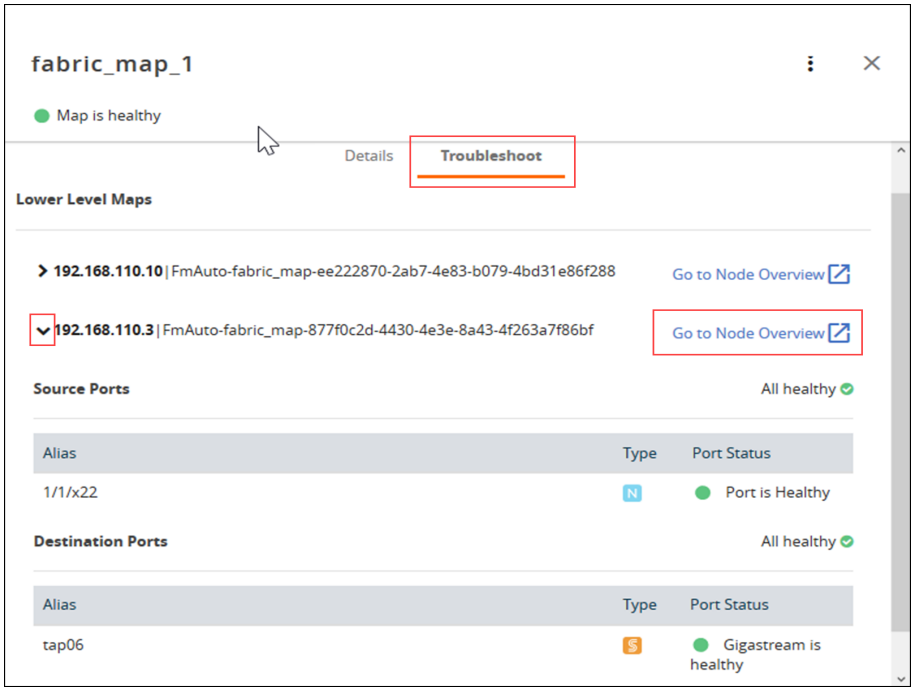

Figure 32-9: Troubleshooting

6. Click the **> arrow** next to the fabric map to display the map details. Use this screen to investigate configuration issues related to your fabric map.

Fabric Map Configuration Status

The following table helps to troubleshoot a failed or partially successful fabric map configuration.

Configuration Status	Definition	Suggested Action
SUCCESS	Fabric map was accepted at GigaVUE-FM level and traffic flow deployed correctly.	• No action required
FAILED	Fabric map was not configured correctly and not accepted or deployed at GigaVUE-FM level.	• Go to the topology page and ensure there is a link between the cluster. • Make sure the circuit GigaStreams are successfully created. • If the clusters are no longer connected, connect the clusters and reapply the configuration. • If the circuit GigaStreams are configured correctly, make the necessary changes and reapply the configuration.
PENDING	• This means no fabric path exists between 2 nodes. **or** • There is a failed operation on the device and further deployment of cluster level maps are not allowed in that device.	• Create a fabric path and reapply fabric map.
PARTIAL_SUCCESS	This means a portion of the fabric map was accepted at GigaVUE-FM level, but devices cannot provision the generated components.	• See *How to Troubleshoot Partial Success Errors* on page 1222 for more details.

Fabric Map Health State

The following table helps to troubleshoot the health state of your fabric map.

Health State	Definition	Suggested Action
Healthy	All fabric map components (nodes, ports, GigaStreams) are all healthy.	• No action required.
Unhealthy	An unhealthy state means there could be a port that is down or packets dropped.	• Make sure no ports are down and no packets dropped. • Correct any configuration connection issues, and re-sync the fabric map.

How to Troubleshoot Partial Success Errors

If your fabric map displays a red status light on the list view page, this means a portion of the fabric map was accepted at GigaVUE-FM level, but devices cannot provision the generated components.

The error message associated with fabric map appears when you hover over the fabric map status column. You can use the error message to help you troubleshoot and identify the components that are in conflict or mis-configured.

Example 1:

In the following example, the status error message indicates that the ports are mis-configured. You can use the following workflow as a best practice for investigating and correcting map configuration problems.

1. Hover over the fabric map status column displaying a red status light to view the error message. The error message provides information on the possible issues present with fabric map components.

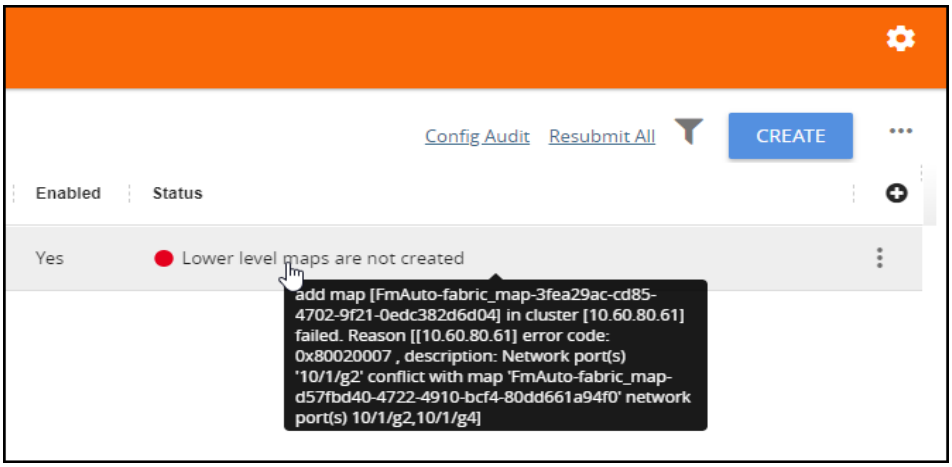

Figure 32-10: Troubleshooting

This error indicates that the input network ports overlapped with the existing map, and therefore it is required to correct this issue and remove the overlapping ports. Use this screen to investigate the exact ports causing the problem.

2. Click **View Details** from the options menu. The fabric map detail screen is displayed.

Figure 32-11: Fabric Maps Details

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

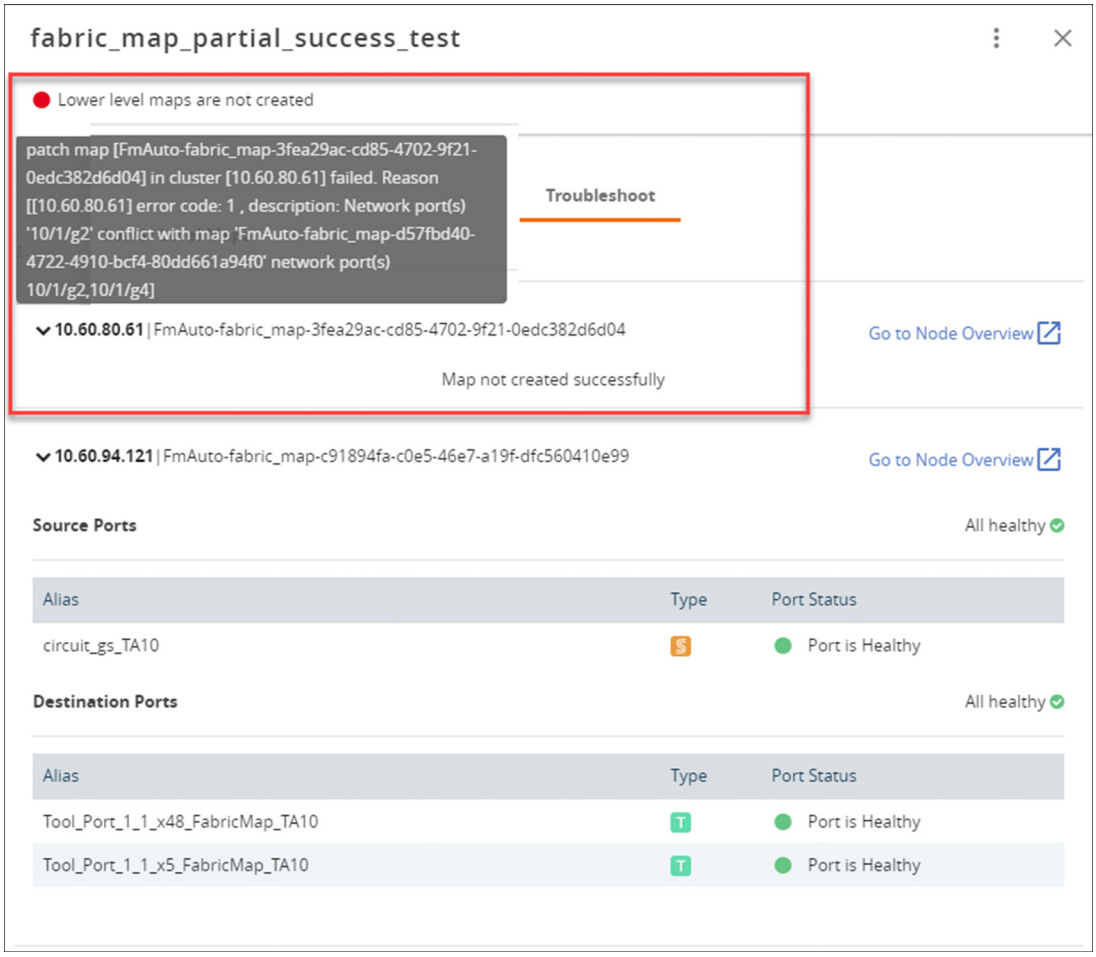

Figure 32-12:

The network ports that you tried to add to your fabric map were not available and contributed to the error in the fabric map configuration. Next, you need to edit the network ports.

4. Click **Edit** from the options menu to edit the fabric map. The Edit Fabric Map screen is displayed.

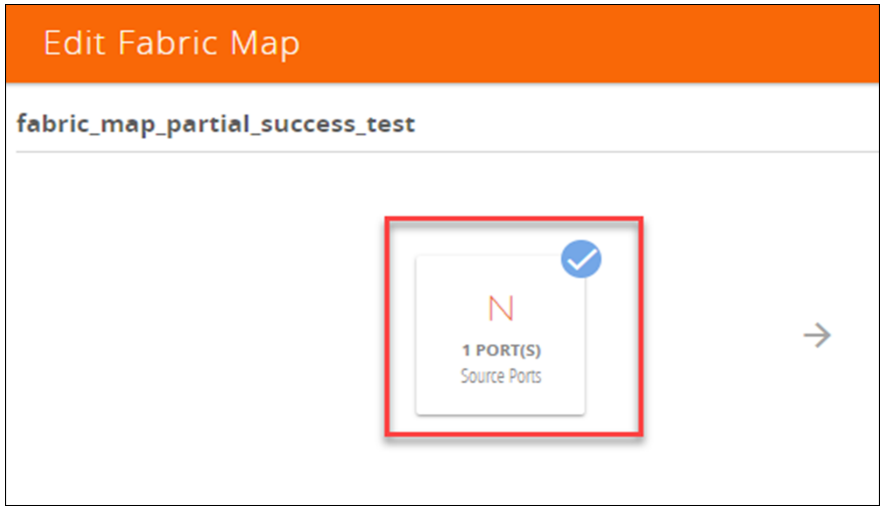

Figure 32-13: Editing Fabric Maps

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

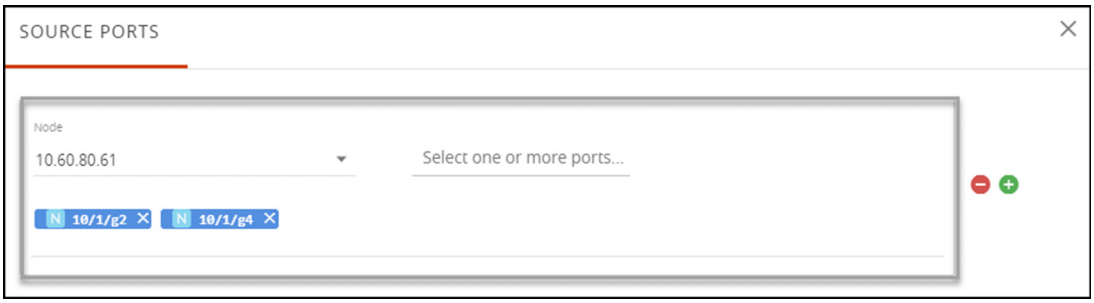

Figure 32-14: Source Port Component

6. The Source Ports screen appear.
7. Edit the source ports with correct information and save your configuration.
8. Click **Submit** on the **Edit Fabric Map** page.

After you have successfully updated the fabric map component causing the problem, the fabric map list view displays a healthy status symbol.

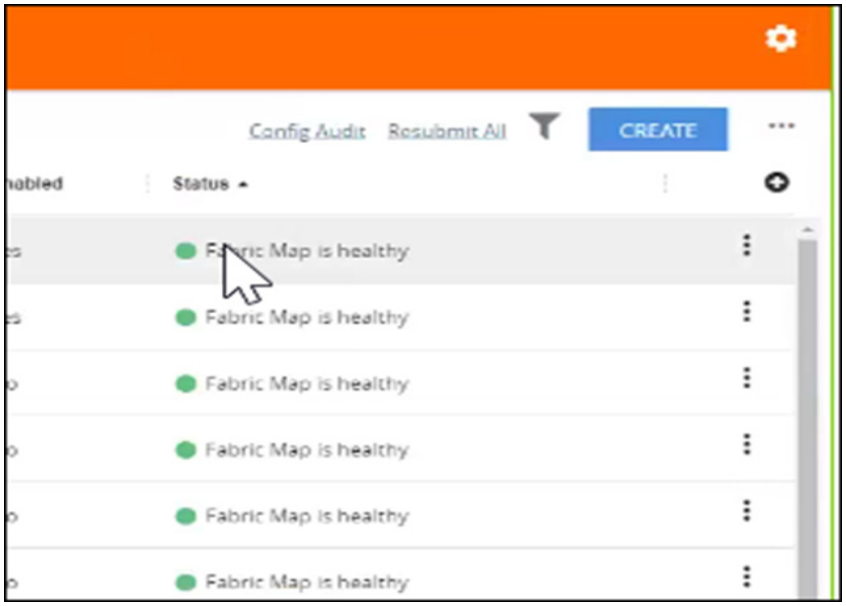

Figure 32-15:

If a problem with the fabric map still exists, repeat the steps above to identify and track down the problem.

Config Audit

Use the Config Audit option to conduct an audit of your fabric map configuration. This process audits the configuration including the components of your fabric map. Config audit verifies that the fabric map configuration is active, the devices are connected, and traffic is flowing through the fabric map with no issues.

To start the configuration audit process:

1. Select **Physical > Fabric Maps.**
2. Click the **Config Audit** link at the top of the screen.

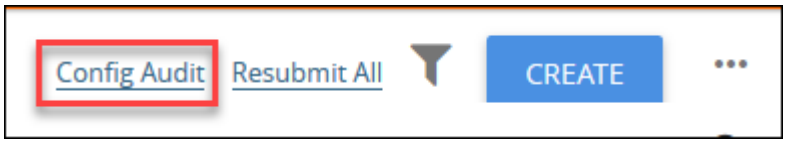

After the config audit is complete, a notification appears with the audit results.

If the config audit is not successful, check the fabric map configuration error message(s) for help in resolving issues.

Limitations of Fabric Maps

The following are the limitations of fabric maps:

- Circuit ports, circuit GigaStreams and manual topology links (or GDP must be enabled on circuit ports to discover the connected clusters) must be configured before creating fabric map.
- Fabric map configurations are saved in GigaVUE-FM and not in the individual devices. Therefore, if you change the GigaVUE-FM instance that manages the devices, then fabric map configurations cannot be rediscovered from the devices. In this scenario, the GigaVUE-FM configuration that has been saved and backed-up from the original GigaVUE-FM instance must be restored to the new GigaVUE-FM instance to continue.
- To avoid FAILED or PENDING state, or no traffic stats found during deployment of fabric maps, you must ensure that the associated devices are connected to GigaVUE-FM and also interconnected for traffic paths.

Starting in software version 5.5, fabric maps support directly connected clusters. Refer *Supported Topologies* on page 1207.

Part 8: Administration

This chapter describes the Administration activities that can be performed in GigaVUE-FM. The administration activities include connecting to your servers, creating sites and tags, notifying users of alarms and events, providing information about the nodes you want to monitor, and adding users to have access to GigaVUE-FM.

The following topics are covered in this section:

33 Authentication

This chapter describes how to configure authentication and authorization settings for GigaVUE-FM.

This section covers of the following main topics:

Overview of Authentication

Authentication pages are used to configure authentication and authorization settings for GigaVUE-FM. In general, configuring authentication consists of specifying the login methods accepted, the order in which they are tried, the local user account to map to external logins, whether to accept roles specified by the AAA server, and the configuration of the external authentication server itself.

Click **Administration** on the top navigation link. On the left navigation pane, click **Authentication**.

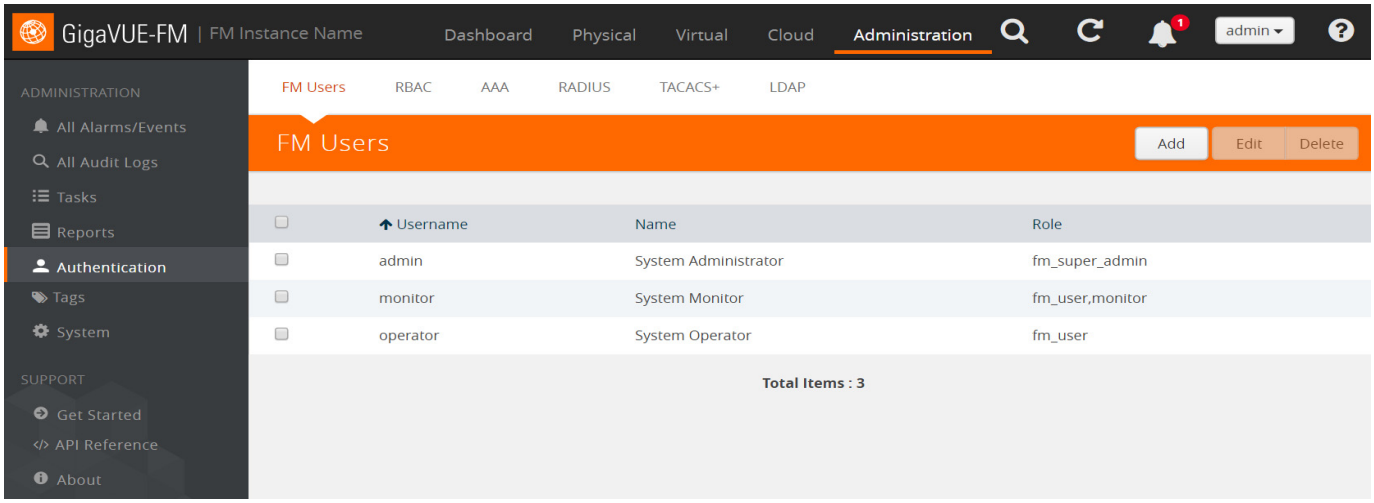

Figure 33-1: Authentication Pages

The following table describes pages available when **Authentication** is selected from the main navigation pane.

Page	Description
FM Users	Manage local user accounts. From here, you can add new accounts, edit existing accounts, or create new ones. Refer to *FM Users* on page 1231 for details.
RBAC	Controls the admin level privileges versus the read-only option. Refer to *RBAC* on page 1234 for details.
Authentication **(AAA, RADIUS, TACACS+, LDAP)**	GigaVUE-FM can authenticate users against a local user database or against the database stored on an external authentication server (RADIUS, TACACS+, or LDAP). Users with the **fm_super_admin** role assigned can specify the authentication methods used for H Series logins using the options in the *AAA (Authentication, Authorization and Accounting)*, *RADIUS*, *TACACS+*, and *LDAP*. The *AAA (Authentication, Authorization and Accounting)* specifies *which* authentication methods to use to configure the methods themselves in their individual tabs (RADIUS, TACACS+, and LDAP).

See the following sections for details on the tabs in the **Authentication** page:

- *FM Users* on page 1231

FM Users

FM Users page lets you manage the local database of GigaVUE FM users. You can configure which local user accounts have access to Fabric Manager and the level of privileges they have (their Role). Through this table you can add a user, make changes to a user, and delete users. You can also view all users that have been created.

Accounts and credentials configured in the **FM Users** page are saved in a local file on the machine where Fabric Manager is installed.

GigaVUE-FM is preconfigured with one user with the **fm_super_admin** role assigned (user name - **admin**, password - **admin123A!**). You should change this password before you start using Fabric Manager.

Notes:

1. To fully take advantage of GigaVUE-FM, Gigamon highly recommends that you have the same user name and password (with roles) registered with the physical node(s). In doing so, GigaVUE-FM provides the ability to manage and monitor physical devices with all of its features.
2. If a user has full access (super admin or admin) on GigaVUE-FM but limited access on the node, they will be able to view the traffic and all the ports from the Dashboard page, Audit logs and Reports but will not be able to configure the node itself.
3. If the user with the same name is created on GigaVUE-FM and the node but the passwords are different, the user will be able to view all the ports on the node from GigaVUE-FM but will not be able to configure the node from GigaVUE-FM. In order to have full access, it is required that both the username and passwords be identical on the node as well as GigaVUE-FM. To avoid such situations it is recommended to use centralized authorization servers such as LDAP, RADIUS or TACACS+.

FM Users table has the following buttons that allow you to manage the users and the information displayed in the table: **Add**, **Edit**, and **Delete**.

Controls	Description
Add	Click to add users to Fabric Manager. Opens a dialog allowing you to specify the name, user name, password, and a role for the user. **Role**: Determines the level of user privileges. They are as follows: • **fm_super_admin** - Allows a user to do everything in Fabric Manager, including adding or modifying users and configuring all AAA settings in the RADIUS, TACACS+, and LDAP tabs. This user can change password for all users. • **fm_admin** - Allows a user to do everything in Fabric Manager except add or modify users and change AAA settings. This user can only change own password. • **fm_user** - Allows a user to view everything in Fabric Manager, including AAA settings, but cannot make any changes.
Edit	Click to change the selected user in the list. Must be a user with the **fm_super_admin** role assigned. Note that the **Change Password** link at the upper right of the GigaVUE-FM interface allows the currently logged in user to change their own password at any time – see *Changing Your Password* on page 1233 for details.
Delete	Click to delete a user. Must be a user with the **fm_super_admin** role assigned.

On the FM Users page, click **Add**.

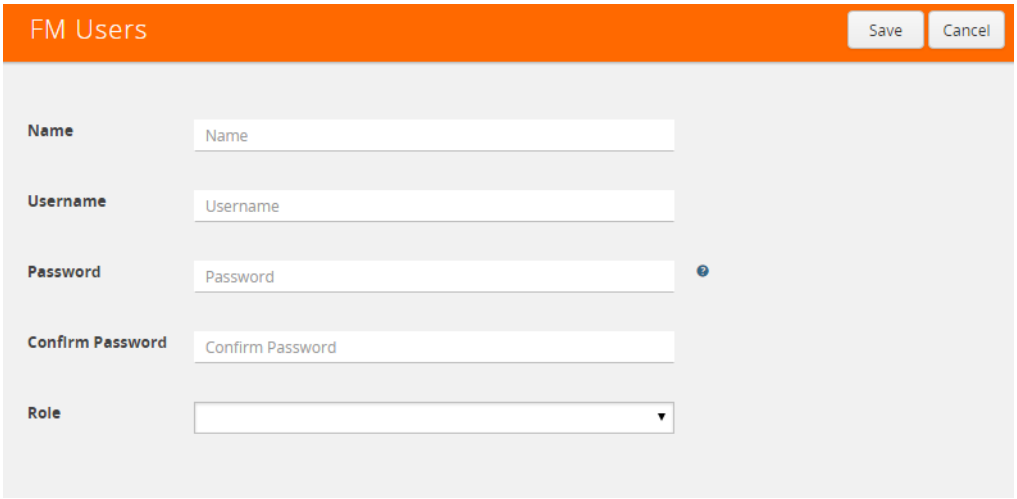

Figure 33-2: FM Local Users Page

Changing Your Password

Users authenticated against GigaVUE-FM's local user database can always change their own passwords. GigaVUE-FM passwords must conform to the following minimum standards:

- One numerical character
- One uppercase character
- One lowercase character
- One special character (**!**, **@**, **#**, and so on)

NOTE: Users authenticated against an external AAA server cannot change their password using this link. External passwords must be changed in the external authentication server itself.

The following are the steps for changing your password:

1. Click on the button in the upper right-hand corner of GigaVUE-FM, where your user name is displayed, and select **Change Password**.

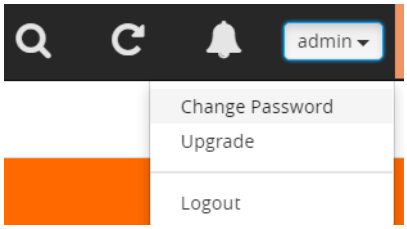

   The Change Password page displays.

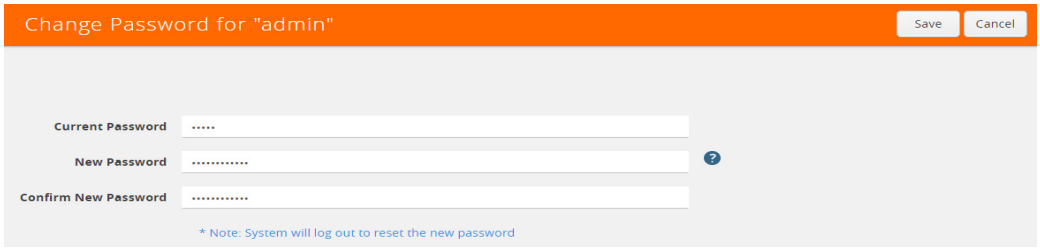

2. On the Change Password page, do the following:
   - Enter your current password in the **Current Password** field.
   - Enter the new password in the **New Password** and **Confirm Password** fields.
3. Click **Save**.

   GigaVUE-FM logs out to reset the password. Enter your new password to log in again.

RBAC

RBAC controls the admin level privileges versus the read-only option. To disable RBAC mode in GigaVUE-FM, check the box as shown in Figure 33-3, and then click **Save**.

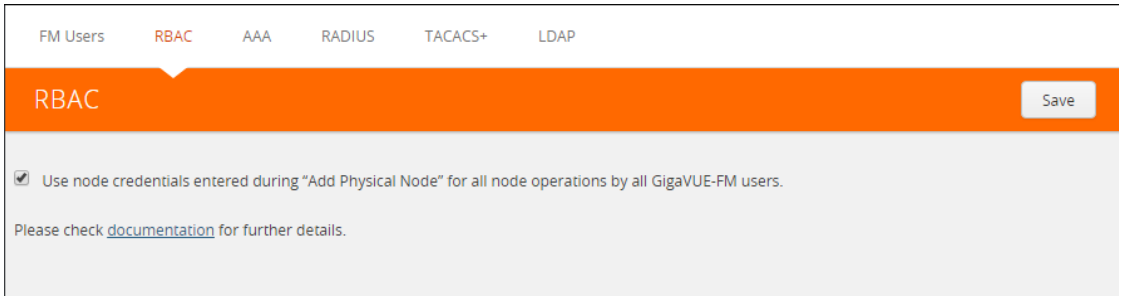

Figure 33-3: RBAC Mode Disabled on GigaVUE-FM

The checkbox on this page does the following:

- If the checkbox is unchecked (which is the default), for a GigaVUE-FM user to be able to manage a node, the user should have the same credentials (username and password) in both GigaVUE-FM and the node. If the number of nodes and/or devices is large, it is recommended that LDAP or similar mechanism be used to ease user management.
- If the checkbox is checked, for a GigaVUE-FM user to be able to manage a node, the user does NOT have to have the exact same credentials in GigaVUE-FM and the node. GigaVUE-FM will use the credentials entered when the node was added to GigaVUE-FM for all operations performed by all users.

In both the cases, GigaVUE-FM Role Based Access Control will be enforced. For example, a GigaVUE-FM user with fm_user role will not be able to modify anything on the node.

For more information about RBAC, refer to *Roles and Users in GigaVUE-FM* on page 1341.

AAA (Authentication, Authorization and Accounting)

You use the **Authentication > AAA** to configure how user logins are authenticated. The appliance can authenticate users against the local user database configured in the *FM Users* or against the database stored on an external authentication server (LDAP, RADIUS, or TACACS+). You can specify a hierarchy of login options, allowing failover in case a given method fails. As shown in the figure above, the AAA tab also shows you how the currently logged-in user was authenticated – **Local**, **RADIUS**, **TACACS**+, or **LDAP**.

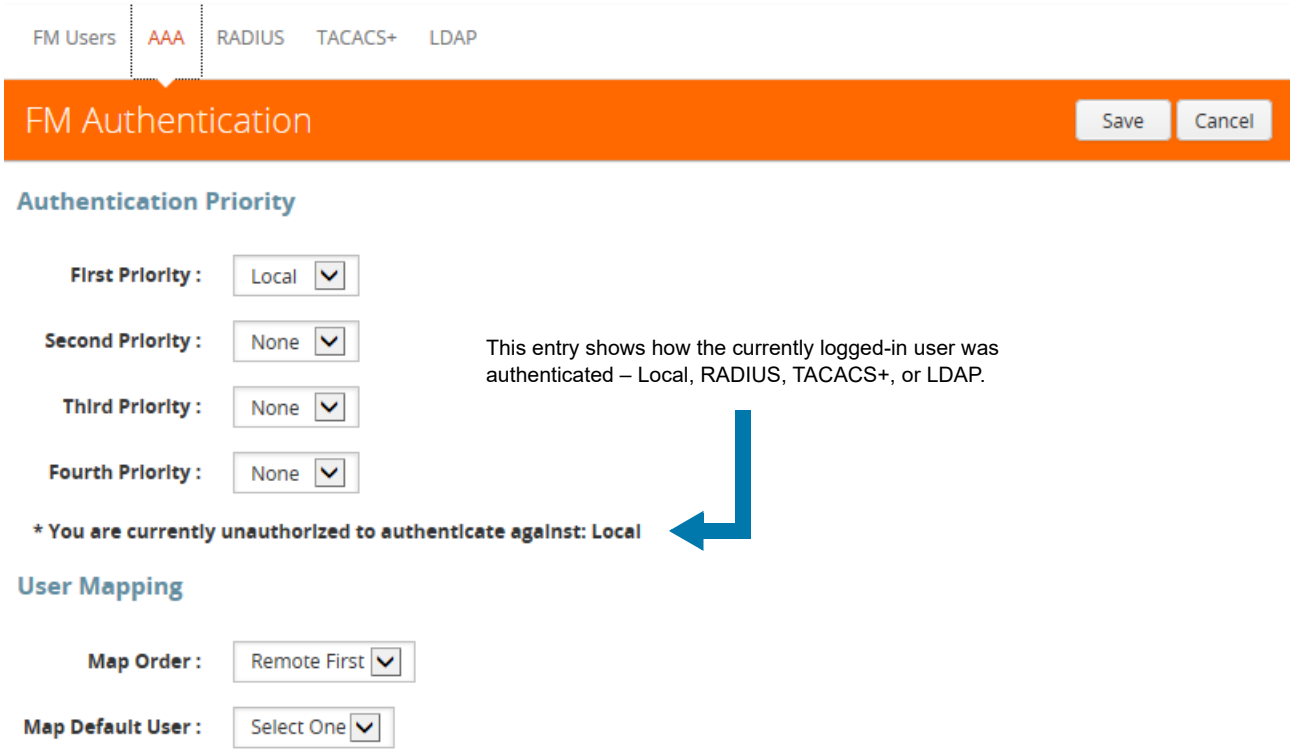

Figure 33-4: AAA Page

Configuring Authentication Priority

GigaVUE-FM supports each of the following authentication methods:

- Local database
- External authentication server
  - TACACS+
  - RADIUS
  - LDAP

You can enable all of these authentication methods at the same time using the Authentication Priority List. If you enable more than one method, GigaVUE-FM uses the methods in the same order in which they are specified (**First Priority**, **Second Priority**, **Third Priority**, **Fourth Priority**), falling back as necessary. If all servers using the first method are unreachable, GigaVUE-FM falls back to the secondary method, and so on. One of the authentication methods specified must always be Local to prevent lockouts.

For example, you could use an external authentication server as your primary authentication method with local authentication as a fallback. The fallback is used when an authentication server is unreachable.

If a server responds to a login attempt with an authentication reject, no further servers using that method are tried. Instead, the next method is tried until either the user's login is granted or all specified methods are exhausted.

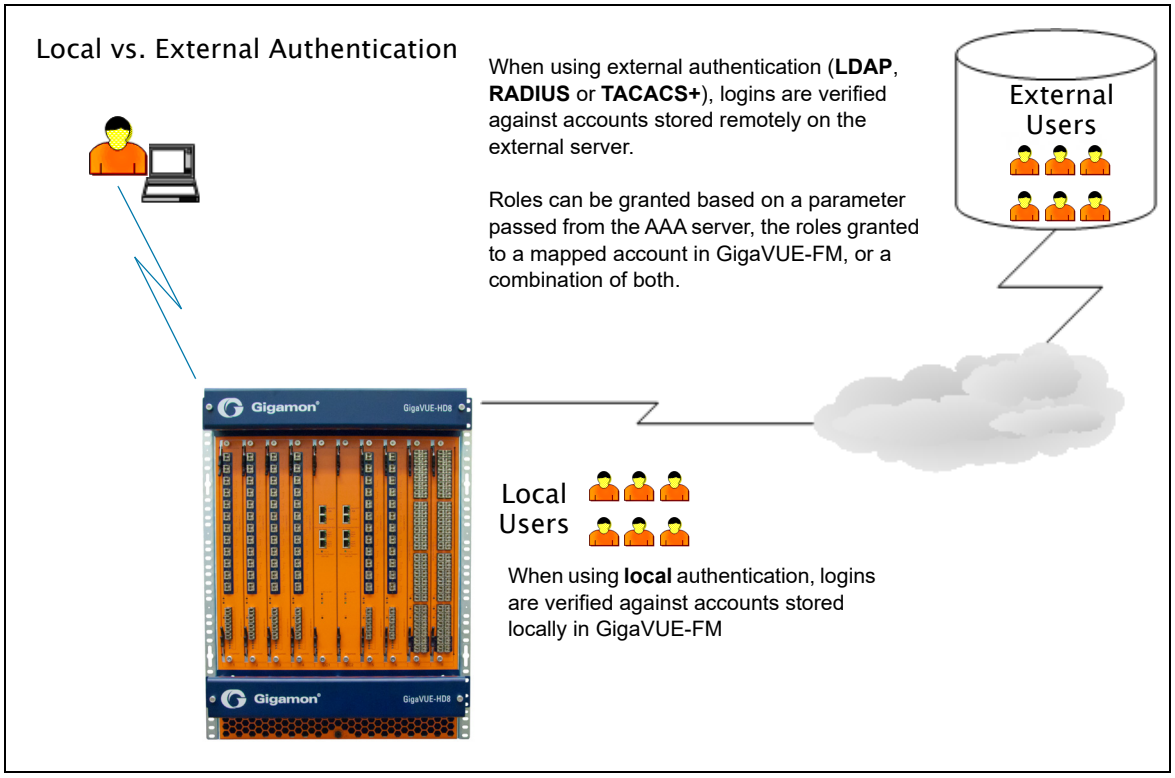

Figure 33-5: Local vs. External Authentication

Configuring User Mapping

Configure the **User Mapping** options to specify how externally authenticated users should be granted privileges in GigaVUE-FM. You can map all external logins to a specific local account, use matching accounts in the local database, or reject external logins unless they have a matching account in the local database.

Setting	Description
Map Order	• Specifies how externally authenticated logins (RADIUS, TACACS+, or LDAP) are mapped to local accounts • **Remote First** – Externally authenticated logins are mapped to the matching local account, if present. If there is no matching local account, the account name specified by the **Map Default User** option, below, is used. This is the default behavior. • **Remote Only** – Externally authenticated logins are only accepted if there is a matching account name in the local database. If there is no matching local account, the externally authenticated login is denied. • **Local Only** – All externally authenticated logins are mapped to the user specified by the **Map Default User** option.
Map Default User	Specifies the account to which externally authenticated logins are mapped when Map Order is set to **Remote First** (if there is no matching local account) or **Local Only**. The drop-down lists all users in the local database except for the default **admin** user provided with all GigaVUE-FM installations.

Next Steps

If you enable RADIUS, TACACS+, or LDAP, you must also do the following:

- Add the RADIUS, TACACS+, or LDAP server to the appliance's list using the corresponding page in **Authentication** (*RADIUS*, *TACACS+*, and *LDAP*).
- Set up GigaVUE-FM installations and users within the external authentication server itself. Depending on your authorization model, you can grant privileges to externally authenticated users based on the roles assigned to a corresponding account on the local node, the roles passed from the AAA server, or a combination of both. Refer to *Grant Roles with External Authentication Servers* on page 1248 for details.

NOTE: You must also have defined a valid DNS server for AAA authentication to work successfully. DNS servers are usually specified during the initial configuration of the node with the Jump Start script.

RADIUS

Users with the **fm_super_admin** role assigned can use the **Authentication > RADIUS** page to add entries to GigaVUE-FM's list of available RADIUS authentication servers.

You can add multiple RADIUS servers. Servers are used as fallbacks in the same order they are specified – if the first server is unreachable, the second is tried, and so on, until all named servers have been used. If a server is reachable and authentication fails, the authentication process terminates.

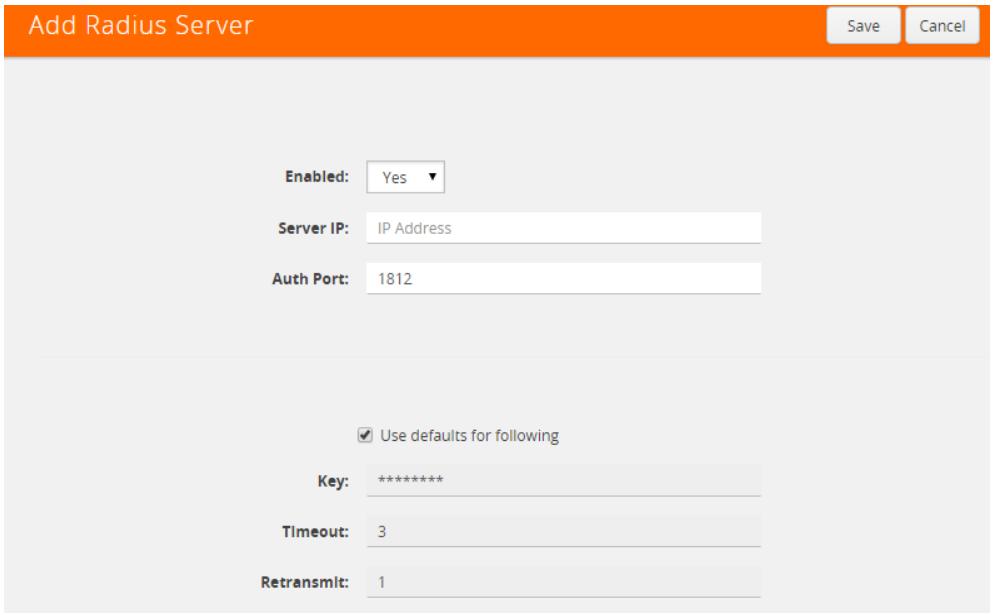

Figure 33-6: Adding Radius Server

NOTE: If you are deploying GigaVUE-FM inside AWS, make sure to provide the private IP address of GigaVUE-FM to the Radius server for authentication and not its public IP address. For more information about AWS, refer to the *Gigamon Visibility Platform for AWS Getting Started Guide*.

Supported RADIUS Servers

GigaVUE-FM has been tested with the RADIUS implementation provided by Cisco Secure ACS v5.4.0.46.0. Although other versions and implementations may operate acceptably, they have not been tested.

RADIUS Page Controls and Fields

RADIUS page has five buttons that allow you to manage the information that appears in the table. **Add**, **Edit**, **Delete**, **and Default**.

Controls	Description
Add	Allows you to add a new RADIUS Server to the list. See *Adding a New RADIUS Server* on page 1238 for details.
Edit	Allows you to change the settings for an existing RADIUS Server entry. Select a server's entry and click **Edit** to open a dialog where you make the changes.
Delete	Allows you to delete a RADIUS Server entry.
Edit Default	Allows you to set default **Key**, **Timeout**, and **Retransmit** options for RADIUS Servers. When you add a new RADIUS Server to the list, you have the option of accepting these default settings or providing custom values. See *Setting Default Key, Timeout, and Retransmit Options for RADIUS Servers* on page 1239 for details.

Adding a New RADIUS Server

Add a new RADIUS Server to GigaVUE-FM's list by clicking the **Add** button and setting the options shown in Figure 33-7 on page 1238.

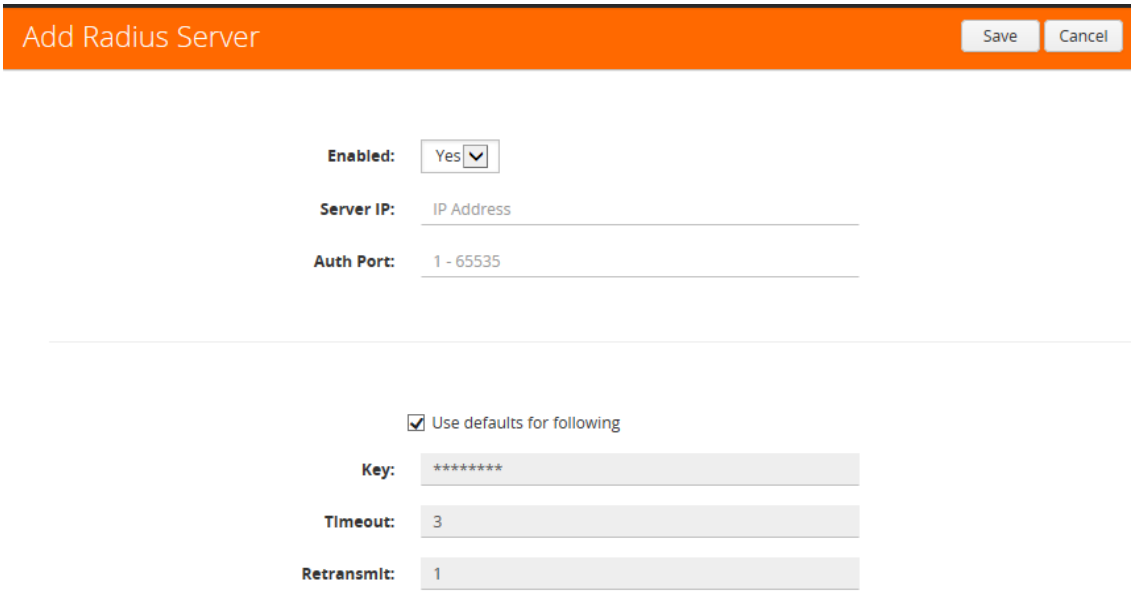

Figure 33-7: Adding Radius Server

The following table describes the settings on the Add Radius Server page.

Setting	Description
Enabled	Specifies whether this server is currently enabled for use with authentication requests
Server IP	Specifies the IP address of the RADIUS server. The same IP address can be used for more than one RADIUS server so long as they use different **Auth Port** values.
Auth Port	Specify the UDP port number on which the RADIUS server is running. If not specified, the port is set to the default RADIUS port number of 1812.
Use defaults for following	Leave this box checked to accept the default values for the **Key**, **Timeout**, and **Retransmit** options configured by clicking the **Edit Default** button at the top of the **RADIUS** page. Alternatively, you can leave this box unchecked and set custom values for the **Key**, **Timeout**, and **Retransmit** options using the respective fields.
Key	Specifies a shared secret string to be used for encryption of authentication packets sent between GigaVUE-FM and this RADIUS server.
Timeout	Specifies how long GigaVUE-FM will wait for a response from this RADIUS server to an authentication request before declaring a timeout failure. The valid range is 0-60 seconds; default value is five seconds.
Retransmit	Specifies the number of times GigaVUE-FM will attempt to authenticate with this RADIUS server before moving on to the next authentication server or method. The valid range is 0-5; default is two. Set to 0 to disable retransmissions.

Setting Default Key, Timeout, and Retransmit Options for RADIUS Servers

Click **Edit Default** to open the Edit Radius Default Settings page shown in the following figure. Use this page to set default **Key**, **Timeout**, and **Retransmit** options available for use with all new RADIUS server entries.

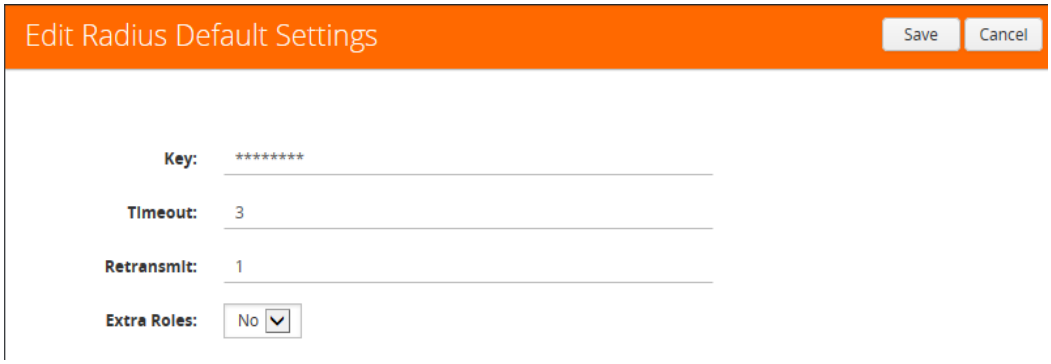

The following table describes the settings.

Setting	Description
Key	Specifies a default shared secret string to be used for encryption of authentication packets sent between GigaVUE-FM and all RADIUS servers. Can be overridden with the key specified for a specific RADIUS Server when the server is added.
Timeout	Specifies a default value for how long GigaVUE-FM should wait for a response from a RADIUS server to an authentication request before declaring a timeout failure. This can be overridden with the timeout value specified for a specific RADIUS Server when the server is added. The valid range is 0-60 seconds. The default value is five seconds.
Retransmit	Specifies a default value for the number of times GigaVUE-FM will attempt to authenticate with a RADIUS server. Can be overridden with the retransmit value specified for a specific RADIUS Server when the server is added. The valid range is 0-5; default is two. Set to 0 to disable retransmissions.
Extra Roles	Specifies whether GigaVUE-FM accepts user roles assigned in the RADIUS server. Refer to *Grant Roles with External Authentication Servers* on page 1248 and *Configure Cisco ACS: RADIUS Authentication* on page 1250 for details.

TACACS+

Users with the **fm_super_admin** role assigned can use the **Authentication > TACACS+** page to add entries to GigaVUE-FM's list of available TACACS+ authentication servers.

You can add multiple TACACS+ servers. Servers are used as fallbacks in the same order they are specified – if the first server is unreachable, the second is tried, and so on, until all named servers have been used. If a server is reachable and authentication fails, the authentication process terminates.

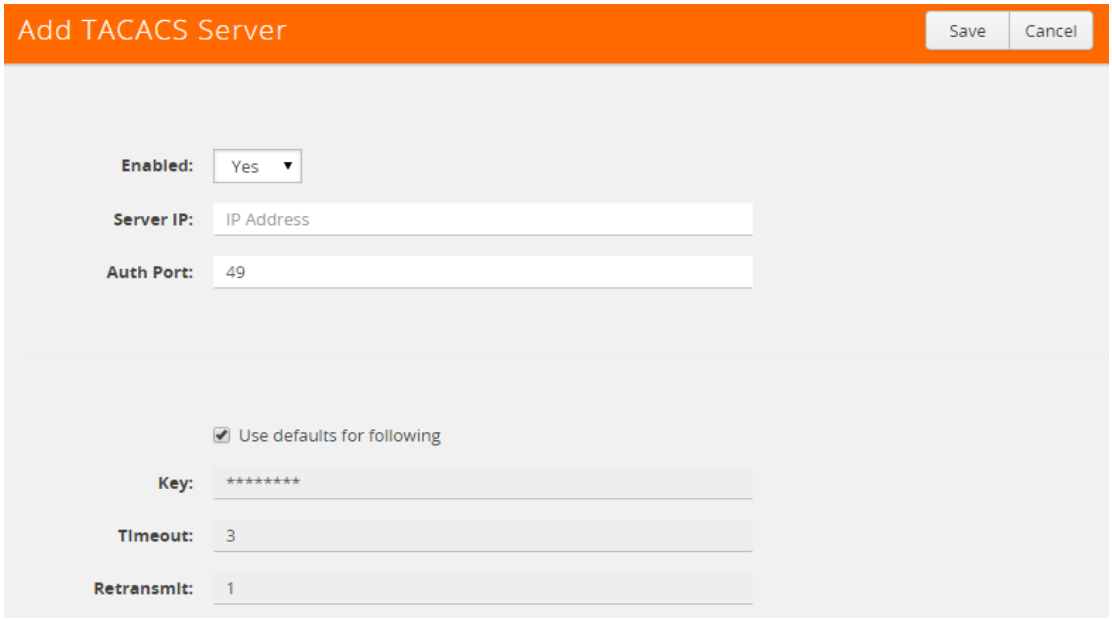

Figure 33-8: TACACS+ Page

NOTE: If you are deploying GigaVUE-FM inside AWS, make sure to provide the private IP address of GigaVUE-FM to the TACACS+ server for authentication and not its public IP address. For more information about AWS, refer to the *Gigamon Visibility Platform for AWS Getting Started Guide*.

Supported TACACS+ Servers

GigaVUE-FM has been tested with the TACACS+ implementation provided by Cisco Secure ACS v5.4.0.46.0. Although other versions and implementations may operate acceptably, they have not been tested.

TACACS+ Page Controls and Fields

TACACS+ tab has five buttons that allow you to manage the information that appears in the table. **Add**, **Edit**, **Delete**, **Edit Default,** and **Refresh**.

Controls	Description
Add	Allows you to add a new TACACS+ Server to the list. See *Adding a New TACACS+ Server* on page 1242 for details.
Edit	Allows you to change the settings for an existing TACACS+ Server entry. Select a server's entry and click **Edit** to open a dialog where you make the changes.
Delete	Allows you to delete a TACACS+ Server entry.
Edit Default	Allows you to set default **Key**, **Timeout**, and **Retransmit** options for TACACS+ Servers. When you add a new TACACS+ Server to the list, you have the option of accepting these default settings or providing custom values.
Refresh	Refreshes the list with information from the GigaVUE-FM database.
Server IP	The IP address configured for this TACACS+ Server entry.
Auth Port	The UDP port number configured for this TACACS+ server entry. The default TACACS+ port number is 49.
Timeout	Indicates how long GigaVUE-FM will wait for a response from the TACACS+ server to an authentication request before declaring a timeout failure.
Retransmit	Indicates the number of times GigaVUE-FM will attempt to authenticate with this TACACS+ server before moving on to the next authentication server or method.
Enabled	Indicates whether this server is currently enabled for use with authentication requests.

Adding a New TACACS+ Server

Add a new TACACS+ Server to GigaVUE-FM's list by clicking **Add** and setting the options on the Add TACACS Server page shown in Figure 33-9 on page 1242.

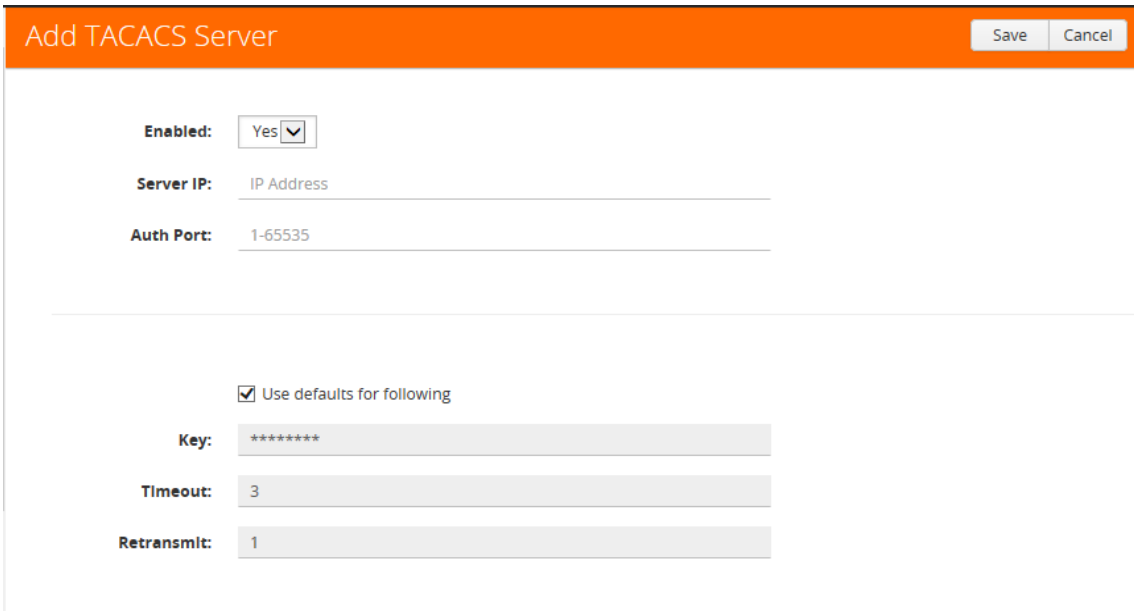

Figure 33-9: Adding TACACS+ Server Settings

The following table describes the settings.

Setting	Description
Enabled	Specifies whether this server is currently enabled for use with authentication requests
Server IP	Specifies the IP address of the TACACS+ server. The same IP address can be used for more than one TACACS+ server so long as they use different **Auth Port** values.
Auth Port	Specify the UDP port number on which the TACACS+ server is running. If not specified, the port is set to the default TACACS+ port number of 49.
Use defaults for following	Leave this box checked to accept the default values for the **Key**, **Timeout**, and **Retransmit** options configured by clicking the **Edit Default** button at the top of the **TACACS+**. Alternatively, you can leave this box unchecked and set custom values for the **Key**, **Timeout**, and **Retransmit** options with the respective fields.
Key	Specifies a shared secret string to be used for encryption of authentication packets sent between GigaVUE-FM and this TACACS+ server.
Timeout	Specifies how long GigaVUE-FM will wait for a response from this TACACS+ server to an authentication request before declaring a timeout failure. The valid range is 0-60 seconds; default value is five seconds.

Retransmit	Specifies the number of times GigaVUE-FM will attempt to authenticate with this TACACS+ server before moving on to the next authentication server or method. The valid range is 0-5; default is two. Set to 0 to disable retransmissions.

LDAP

Users with the **fm_super_admin** role assigned can use the **Authentication > LDAP** page to add entries to GigaVUE-FM's list of available LDAP authentication servers.

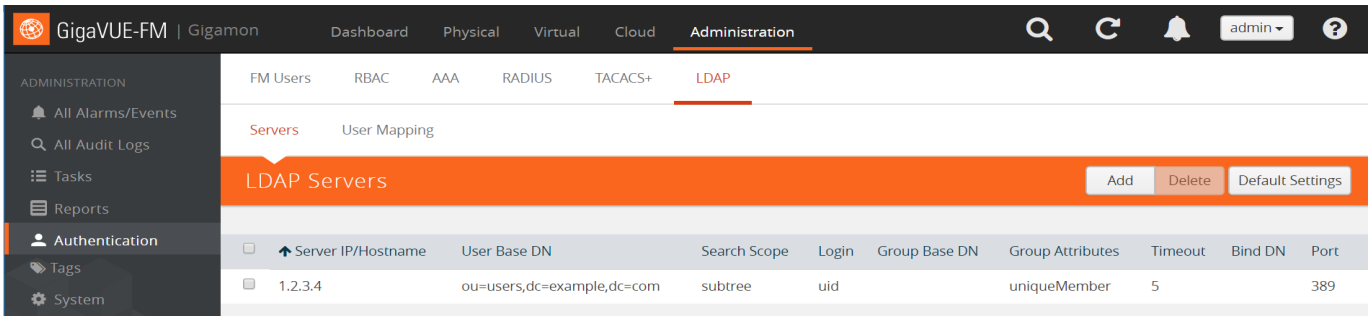

Figure 33-10: LDAP Page

You can add multiple LDAP servers. Servers are used as fallbacks in the same order they are specified – if the first server is unreachable, the second is tried, and so on, until all named servers have been used. If a server is reachable and authentication fails, the authentication process terminates.

Supported LDAP Servers

GigaVUE-FM has been tested with the LDAP implementation provided by Apache Directory Studio v2.0.0.v20130308. Although other implementations may operate acceptably, they have not been tested. GigaVUE-FM does not support the LDAP implementation provided by Active Directory with SSL in this release.

LDAP Page Controls and Fields

LDAP has the following buttons that allow you to manage the information**.**

Controls	Description
Add	Add a new LDAP Server to the list. See *Adding a New LDAP Server* on page 1244 for details.
Delete	Delete a LDAP Server entry.
Default Settings	Set default options for LDAP Servers. When you add a new LDAP Server to the list, you have the option of accepting these default settings or providing custom values. See *Setting Default Options for LDAP Servers* on page 1244 for details.

Adding a New LDAP Server

Select **Authentication > LDAP** and click **Add**. The Add LDAP Server page is displayed. Refer to Figure 33-11. A new LDAP Server is added to the GigaVUE-FM's list.

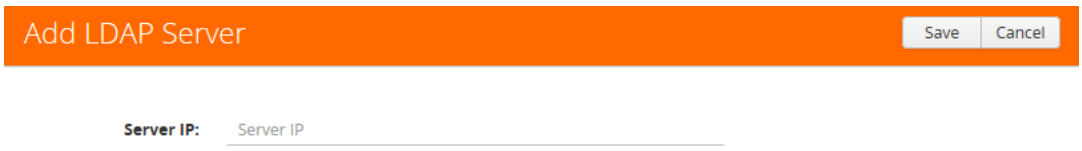

Figure 33-11: Adding LDAP Server

All other settings for LDAP servers are inherited from the defaults configured by clicking the **Edit Default** button at the top of the **LDAP** page. Refer to *Setting Default Options for LDAP Servers* on page 1244 for details.

Setting Default Options for LDAP Servers

Click **Edit Default** to set configuration options for use with all new LDAP server entries, and then set the following options for LDAP servers. Note that these options are all global options and cannot be configured on a per-host basis.

Setting	Description
User Base DN	Identifies the base distinguished name (location) of the user information in the LDAP server's schema. Provide the value as a string with no spaces.
User Search Scope	Specifies the search scope for the user under the base distinguished name (DN): **Subtree** (default) – Searches the base DN and all of its children. **One-Level** – Searches only the immediate children of the base DN.
Login UID	Specify the name of the LDAP attribute containing the login name. The default is **sAMAccountName**. You can also specify a custom string or uid (for User ID).
Bind Password	Provides the credentials to be used for binding with the LDAP server. If Bind DN is left undefined for anonymous login (the default), Bind Password should be left undefined, too.
Group Base DN	Set this option to require membership in a specific Group Base DN for successful login to the appliance. By default, the **Group Base DN** is left empty – group membership is not required for login to the system. If you do specify a **Group Base DN**, the attribute specified by the **Group Login Attribut**e option must contain the user's distinguished name as one of the values in the LDAP server or the user will not be logged in.
Bind DN	Specifies the distinguished name (DN) on the LDAP server with which to bind. By default, this is left empty for anonymous login.
Group Login Attribute	Use this argument to specify the name of the attribute to check for group membership. If you specify a value for **Group Base DN**, the attribute you name here will be checked to see whether it contains the user's distinguished name as one of the values in the LDAP server.

LDAP Version	Specify which version of LDAP to use. The default of Version 3 is the current standard; some older servers still use Version 2.
Port	Specify the port number on which the LDAP server is running. If you do not specify a port, the default LDAP authentication port number of 389 is used.
Timeout	Specifies how long the appliance should wait for a response from the LDAP server to an authentication request before declaring a timeout failure. The valid range is 0-60 seconds; default value is five seconds.
Extra Roles	Specifies whether GigaVUE-FM accepts user roles assigned in the LDAP server. Refer to *Grant Roles with External Authentication Servers* on page 1248 and *Configure LDAP Authentication* on page 1253 for details.
SSL Mode	Enables SSL or TLS to secure communications with LDAP servers as follows: • None—Does not use SSL or TLS to secure LDAP • SSL—Secures LDAP using SSL over the SSL port. • TLS—Secures LDAP using TLS over the default server port.
SSL Port	Specifies the LDAP SSL port number.
Referrals	Specifies the type of user information search in the LDAP servers. • Yes—Searches the user information in all the LDAP servers. • No—Searches the user information in the selected LDAP server.
SSL Certificate Check	Enables LDAP SSL/TLS certificate verification. Use Off to disable.
SSL CA List	Configures LDAP to use a supplemental CA list. • Default CA List—Configures CA list with the Secure Cryptography. • None—Does not use a supplemental list.
Search Timeout	Specifies how long the appliance should wait for a response from the LDAP server over SSL/TLS port before declaring a timeout failure. The valid range is 0-60 seconds; default value is five seconds.

Group Based Role Assignment

GigaVUE-FM provides the ability to assign roles to the members of a group based on their existing directory server group membership.

User Mapping enables you to assign a role of a particular user to the members of a specific group. Mapping a remote user group to a local user account provides a granular way the roles are assigned to a group when they log in to GigaVUE-FM. Moreover, this eliminates the need to create specific roles on the remote server, since a remote user group can be mapped to a local user account.

NOTE: Only users with **fm_super_admin** role assigned can enable or disable User Mapping.

Refer to the following steps to enable User Mapping:

1. Under **LDAP > User Mapping,** click on **New.**

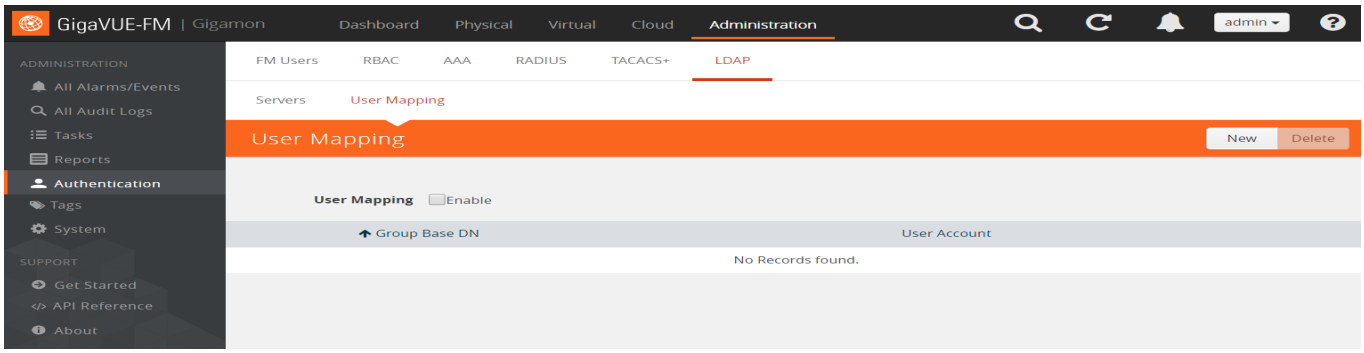

Figure 33-12: User Mapping

2. Enter the **Remote Group Base DN** and select a local user account you want the remote user group to map to.

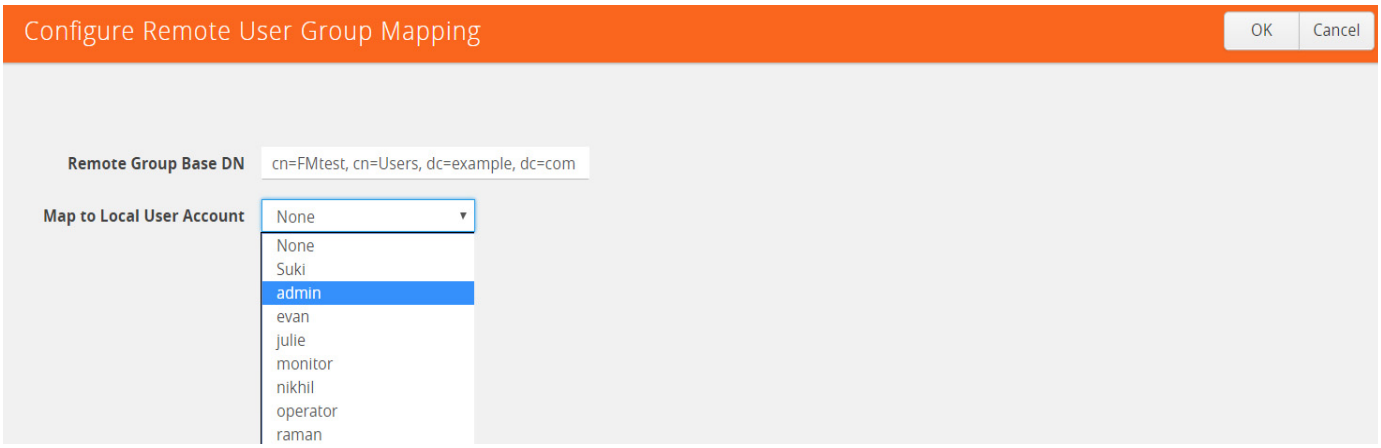

Figure 33-13: Selecting Local User Account

The following table describes the settings.

Setting	Description
Remote Group Base DN	Specifies the user mapping for a specific Remote Group Base.
Map to Local User Account	Specifies local user accounts that a remote group can be mapped to.

NOTE: Group Base DN is case-sensitive. **CN=FMtest** is different from **cn=FMtest**.

3. Click **OK** to configure remote user group mapping.
4. Check **User Mapping** to enable it.

Figure 33-14: User Mapping Enabled

Now when a remote user logs in, they would be given the role of user admin.

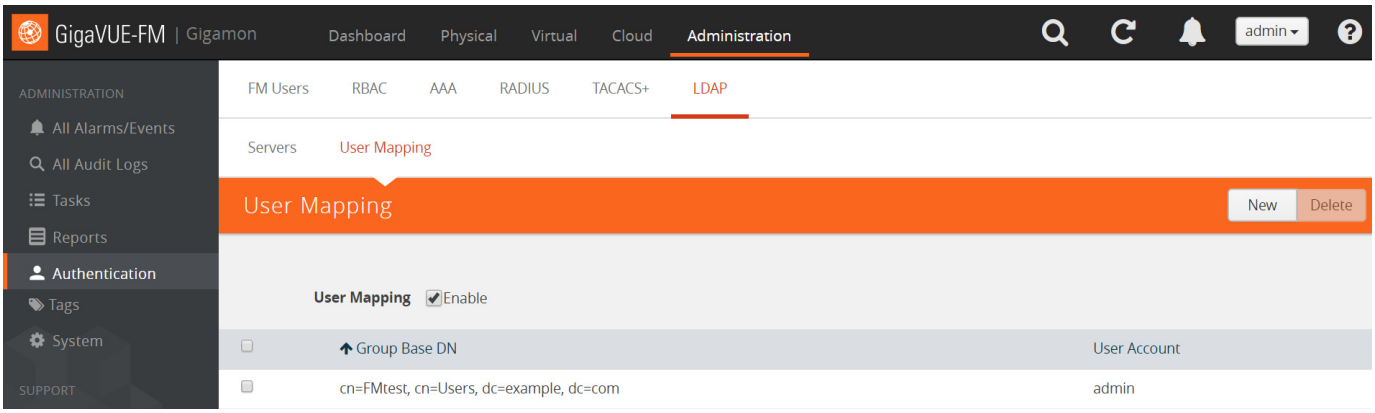
GigaVUE-FM | Gigamon
Dashboard
Physical
Virtual
Cloud
Administration
admin
ADMINISTRATION
All Alarms/Events
All Audit Logs
Tasks
Reports
Authentication
Tags
System
SUPPORT
FM Users
RBAC
AAA
RADIUS
TACACS+
LDAP
Servers
User Mapping
User Mapping
New
Delete
User Mapping
Enable
Group Base DN
User Account
cn=FMtest, cn=Users, dc=example, dc=com
admin

Grant Roles with External Authentication Servers

GigaVUE-FM provides the following user roles:

- **fm_super_admin** — Allows a user to do everything in GigaVUE-FM, including adding or modifying users and configuring all AAA settings in the RADIUS, TACACS+, and LDAP pages. Can change password for all users.
- **fm_admin** — Allows a user to do everything in GigaVUE-FM except add or modify users and change AAA settings. Can only change their own password.
- **fm_user** — Allows a user to view everything in GigaVUE-FM, including AAA settings, but not make any changes. Can only change their own password.

These roles can be assigned using any of the following techniques:

Role Assignment Technique	See this section:
Roles assigned in local GigaVUE-FM database.	*How to Use Local Role Assignments* on page 1248
Roles assigned in external authentication server.	*How to Use AAA Server Role Assignments* on page 1248
Roles assigned using a combination of local roles and the external authentication server.	*How to Use Combination of Local and AAA Role Assignments* on page 1249

How to Use Local Role Assignments

In this model, an externally authenticated user is granted the roles assigned to the account in GigaVUE-FM itself. This can take place either by a matching account name (the same account name is specified both in the AAA server and GigaVUE-FM), or by using the **Local Only** option in the **User Mapping** section of the **AAA** page to map all externally authenticated users to a specific account in GigaVUE-FM (refer to *AAA (Authentication, Authorization and Accounting)* on page 1234 for information on the **Local Only** option).

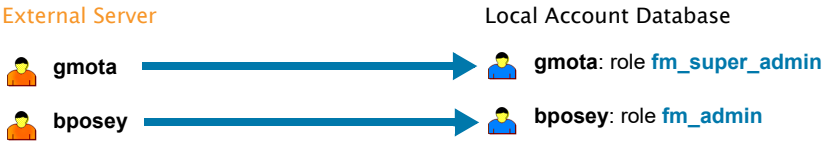

In this model, matching accounts are configured both in the external server and the local account database. The AAA account automatically receives all roles assigned to the matching account in GigaVUE-FM.

How to Use AAA Server Role Assignments

In this model, you configure GigaVUE-FM to accept roles passed from the AAA server by setting the **Extra Roles** option in the Default Settings dialog box for your AAA server to **Yes**. Then, you set up a **local-user-name** attribute for the account in the AAA server to pass a reserved account name (**operator**) and one or more roles to GigaVUE-FM. In

this case, the roles are fully assigned in the AAA server and there are no matching accounts in GigaVUE-FM.

In this model, there are no matching accounts configured in GigaVUE-FM.
The local-user-name attribute configured in the AAA server specifies a special reserved **operator** account to be used in GigaVUE-FM with the roles assigned.

How to Use Combination of Local and AAA Role Assignments

In this model, you configure GigaVUE-FM to accept roles passed from the AAA server. Then, you set up a **local-user-name** attribute for the account in the AAA server that maps it to an existing local user account in GigaVUE-FM. The **local-user-name** attribute can optionally include additional roles to be assigned to the user in addition to those already assigned to the targeted local user account.

For example, in the figure below, the **gmota** account does not exist in GigaVUE-FM's local database. It has a **local-user-name** attribute that specifies the account should be mapped to the local user account **mcain**. The **admin** role is already locally assigned to **mcain**; the **fm_super_admin** role comes from the AAA server with the **role-fm_super_admin** argument.

In this model, the roles assigned are a combination of those from the AAA server and those from the local account database:

- **gmota** is mapped to local user **mcain**. He receives both the role configured in the AAA server (**fm_super_admin**) and the role locally assigned to **mcain** (**fm_admin**).

- **bposey** is mapped to local user **psandoval** with no additional roles specified. He receives only the role locally assigned to the **psandoval** account (**fm_admin**).

Assign Roles in AAA Servers

Refer to *Configure Roles in External Authentication Servers* on page 1250 for instructions on how to set up users with local-user-name attributes in RADIUS, TACACS+, and LDAP AAA servers.

Configure Roles in External Authentication Servers

This section describes how to set up RADIUS, TACACS+, and LDAP servers to perform authentication for GigaVUE-FM, including how to include a local user mapping attribute that GigaVUE-FM can use to assign roles to an externally-authenticated user. See the following sections for details:

- *Grant Roles with External Authentication Servers* on page 1248
- *Configure Cisco ACS: RADIUS Authentication*
- *Configure Cisco ACS: TACACS+ Authentication*
- *Configure LDAP Authentication*

Configure Cisco ACS: RADIUS Authentication

Use the following steps to configure Cisco ACS 5.x (RADIUS) to grant extra roles to externally authenticated users in GigaVUE-FM.

Enable Extra Roles for RADIUS in GigaVUE-FM

1. Configure GigaVUE-FM to accept extra roles in the response from the AAA server:
   a. Click **Edit Default** in the **Authentication > RADIUS**
   b. Set the **Extra Roles** option to **Yes**.
   c. Click **Update**.

Assign the Class Attribute in RADIUS Authorization Profile (ACS 5.x)

2. Navigate to **Policy Elements > Authorization and Permissions > Network Access > Authorization Profiles** and click **Create** to add a new authorization profile.
3. Give the profile a name and description and click on the **RADIUS Attributes**.
4. Leave **Dictionary Type** set to **RADIUS-IETF** and click the **Select** button adjacent to the **RADIUS Attribute** field.
5. Select the **Class** attribute from the dialog that appears and click **OK**.
6. Leave the **Attribute Type** and **Attribute Value** fields at their default value (**String** and **Static**, respectively).

7. Supply the local user mapping and optional roles using the following syntax, as shown in the figure below:

   <mapping_local_user>[:role-<mapping_local_role_1> [role-<mapping_local_role_2>[...]]]

   NOTE: The extra role specified in the authentication server must match a role already available in GigaVUE-FM – **fm_super_admin**, **fm_admin**, or **fm_user**.

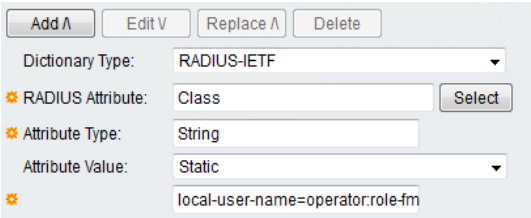

8. Click the **Add** button  to add this attribute to the authorization profile.
9. Assign this authorization profile to a group and populate it with GigaVUE-FM users.

Figure 33-15 shows these settings in a CiscoSecure ACS 5.x authorization profile.

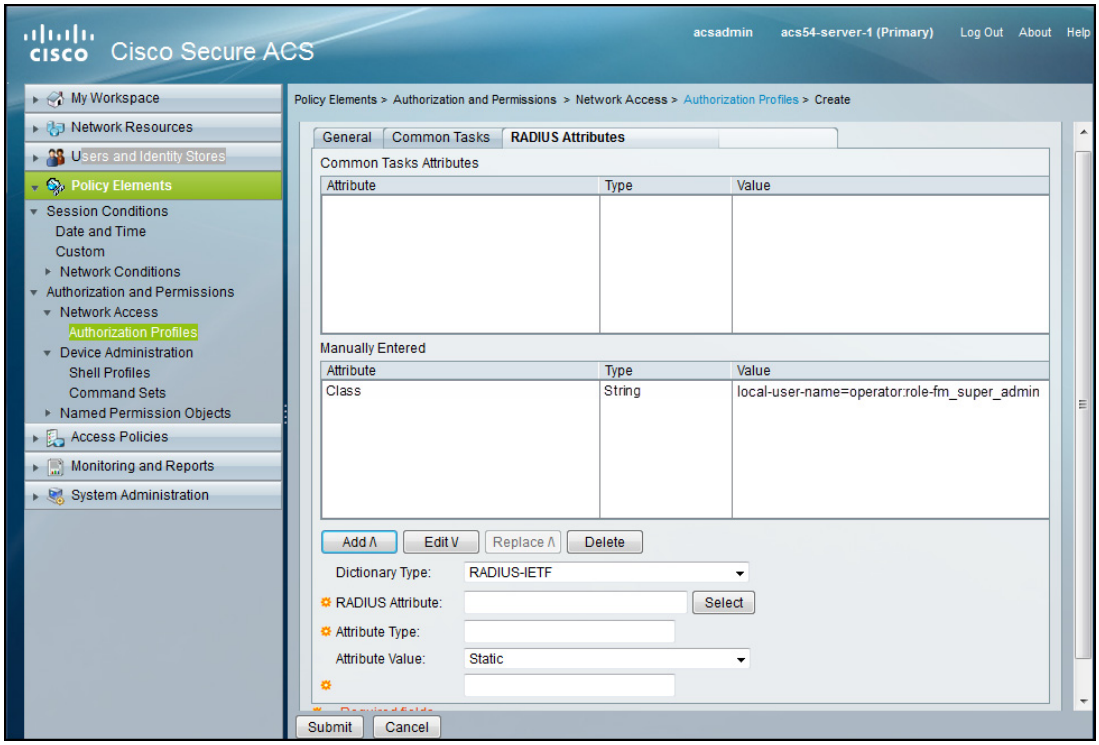

Figure 33-15: Supplying the Class Field for RADIUS (ACS 5.x)

Configure Cisco ACS: TACACS+ Authentication

Use the following steps to configure Cisco ACS 5.x (TACACS+) to grant extra roles to externally authenticated users in GigaVUE-FM.

Enable Extra Roles for TACACS+ in GigaVUE-FM

1. Configure GigaVUE-FM to accept extra roles in the response from the AAA server:
   a. Click **Edit Default** in the **Authentication > TACACS+**.
   b. Set the **Extra Roles** option to **Yes**.
   c. Click **Update**.

Assign local-user-name to Shell Profile (ACS 5.x)

2. Navigate to **Policy Elements > Authorization and Permissions > Device Administration > Shell Profiles** and click **Create** to add a new shell profile.
3. Give the profile a name and description in the **General** page.
4. Click the **Custom Attributes** page.
5. Set the **Attribute** field to **local-user-name**.
6. Leave the **Requirement** and **Attribute Value** fields at their default value (**Mandatory** and **Static**, respectively).
7. Supply the local user mapping and optional roles using the following syntax, as shown in the figure below:

   <mapping_local_user>[:role-<mapping_local_role_1> [role-<mapping_local_role_2>[...]]]

   NOTE: The extra role specified in the authentication server must match a role already available in GigaVUE-FM – **fm_super_admin**, **fm_admin**, or **fm_user**.

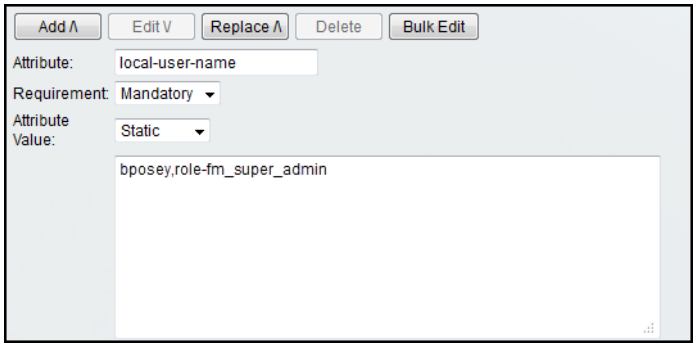

8. Click the **Add** button to add this attribute to the shell profile.
9. Click **Submit** to finalize this shell profile.
10. Create Service Selection Rules that will assign this shell profile to desired GigaVUE users.

Figure 33-16 shows the an example of a shell profile for TACACS+ in ACS 5.x with the local-user-name attribute supplied.

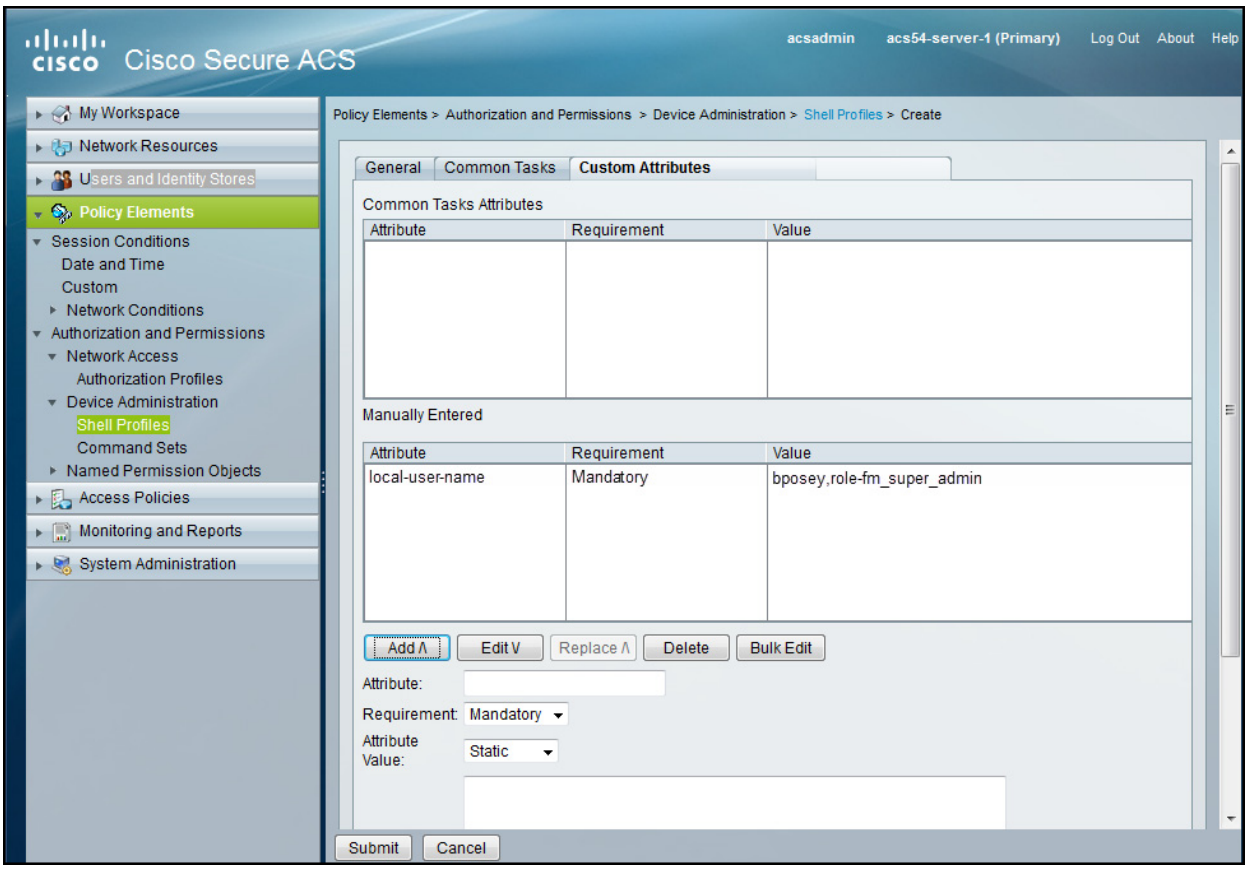

Figure 33-16: Supplying local-user-name and Roles in ACS 5.x for TACACS+

Configure LDAP Authentication

Use the following steps to configure an LDAP server (for example, Apache Directory Server) to grant extra roles to externally authenticated users in GigaVUE-FM.

Enable Extra Roles for LDAP in GigaVUE-FM

1. Configure GigaVUE-FM to accept extra roles in the response from the AAA server:
   a. Click **Edit Default** in the **Authentication > LDAP**
   b. Set the **Extra Roles** option to **Yes**.
   c. Click **Update**.

Assign local-user-name to Shell Profile (ACS 5.x)

2. Add an **employeeType** attribute to the **InetOrgPerson** user object.

   The attribute format is as follows:

   <mapping_local_user>[:role-<mapping_local_role_1> [role-<mapping_local_role_2>[...]]]

   NOTE: The extra role specified in the authentication server must match a role already available in GigaVUE-FM – **fm_super_admin**, **fm_admin**, or **fm_user**.

   Figure 33-17 on page 1254 shows an example of this.

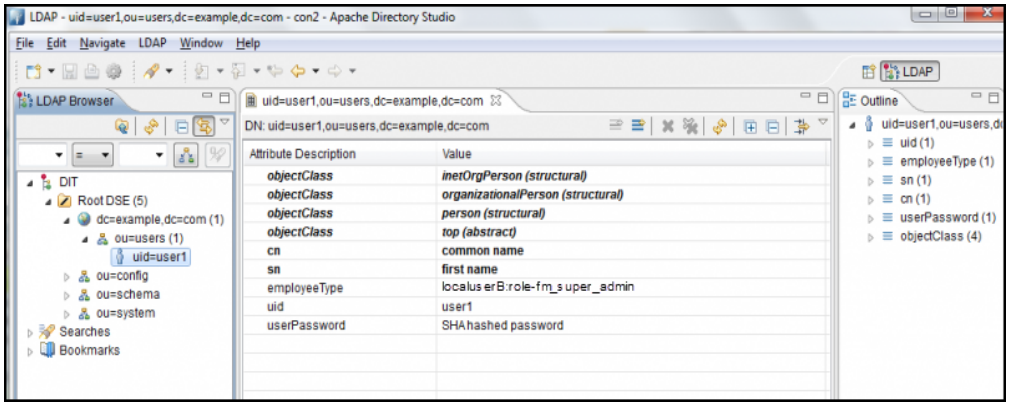

Figure 33-17: Adding the employeeType Attribute

34 Sites and Tags

This chapter describes how to use sites and tags to group clusters, ports, port groups, GigaSMART groups, GigaStreams, and port pairs.

This section covers the following main topics:

Introduce Sites and Tags

Managing hundreds of clusters and thousands of nodes in a cluster can be a daunting challenge. Using sites and tags, GigaVUE-FM lets you group similar types of clusters and objects such as ports, port groups, GigaSMART groups, GigaStreams, and port pairs. By associating clusters to a site, you can quickly view the related topology, alarms and events, and audit logs of the clusters associated to the site. In the physical dashboard, you can select the site and view the highest and lowest traffic, most and least utilized traffic, node and unhealthy map status summary, and inventory of the clusters associated to the site. By associating ports, port groups, GigaSMART groups, GigaStream, and port pairs to a tag, you can compare the aggregated traffic flowing through a list of ports using Traffic Comparison by Tags widget.

Sites can be associated to only clusters whereas user-defined tags can be associated to clusters as well as other objects.

A cluster can have only one site associated to it. However, multiple tags can be associated to a cluster, port, port group, GigaSMART group, GigaStream, and port pairs.

The Sites and Tags feature is supported only on GigaVUE H series. You must have admin privileges to create sites and tags.

Sites

A site is a predefined tag name. You can associate multiple values to the same site. When a site is associated to a cluster, all the standalone nodes within the cluster automatically inherit the same site and value.

In this example, the GigaVUE-FM instance is managing multiple clusters. You could categorize these clusters based on the geographic regions—Austin, Santa Clara, Seattle, and so on as shown in Figure 34-1 on page 1256.

After categorizing the clusters, you can select the Santa Clara site from the Physical Nodes page and view only those clusters and the standalone nodes associated to Santa Clara. Refer to Figure 34-1 on page 1256. You can perform actions such as image upgrade, backup, restore, and reboot on the clusters and standalone nodes associated to Santa Clara.

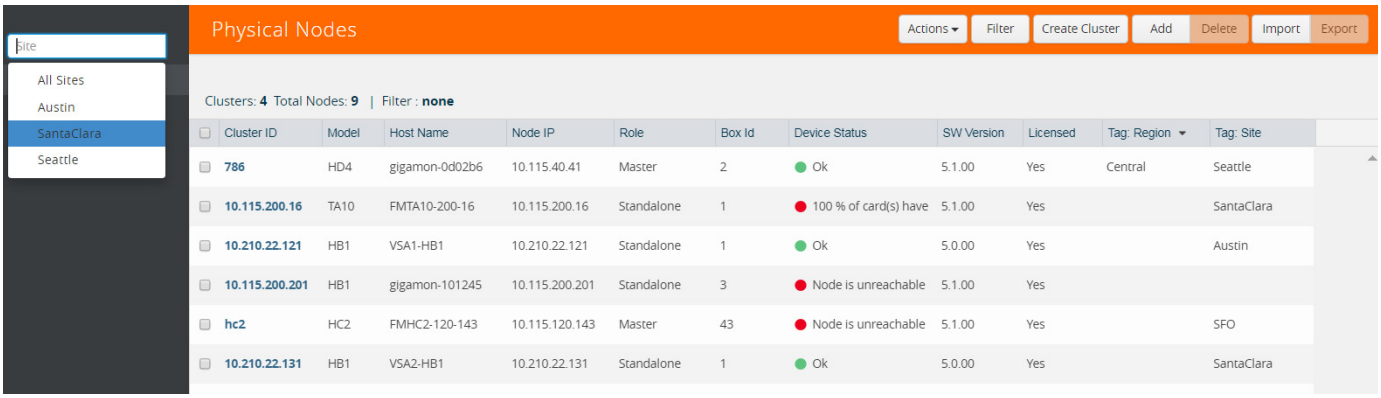

Figure 34-1: Site Selection in Physical Nodes Page

On the left navigation pane under Physical tab, you can select Santa Clara as the site and then select Topology, Alarms/Events, or Audit Logs to view the respective information for the clusters and standalone nodes associated to Santa Clara. Refer to Figure 34-2 on page 1256.

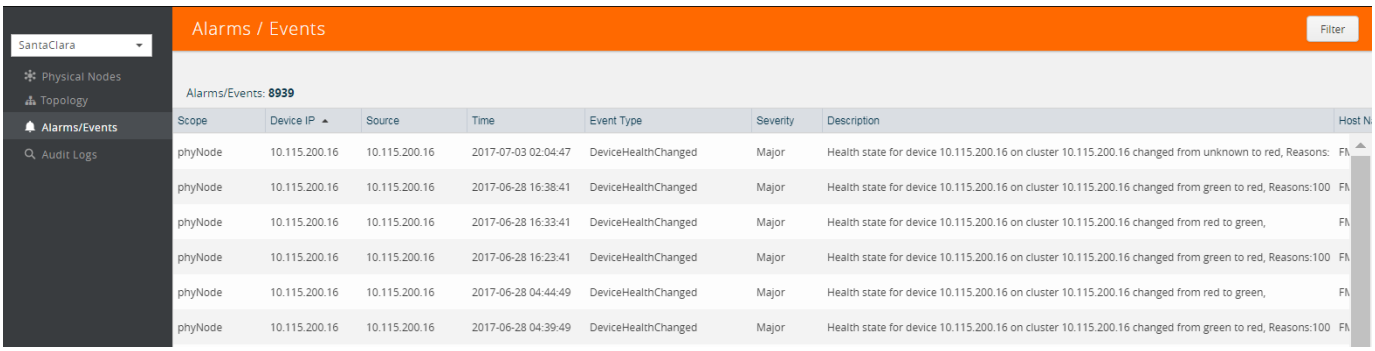

Figure 34-2: Alarms/Events For the Selected Site

In the Physical Dashboard, you can select Santa Clara and create widgets to view the highest traffic, lowest traffic, most utilized and least utilized traffic, inventory, and status summary. Refer to Figure 34-3 on page 1257.

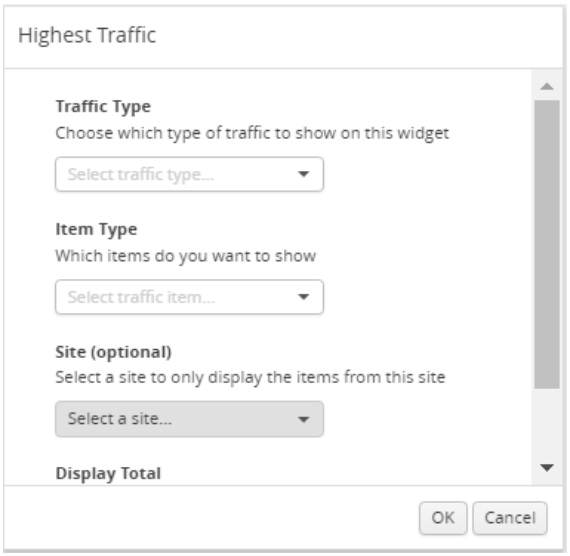

Figure 34-3: Dashboard Widget Creation Selecting a Site

User-defined Tags

A user-defined tag can be associated to clusters, ports, port groups, GigaSMART groups, GigaStream, and Port Pair. To create your own (user-defined) tag, you can define a tag key and a value name. You can create a single tag key and associate multiple values to the same key.

In this example, the traffic is forwarded from GigaVUE-TA10 to GigaVUE-HC2. The traffic from GigaVUE-HC2 is then forwarded to FireEye and WireShark. Refer to Figure 34-4 on page 1257.

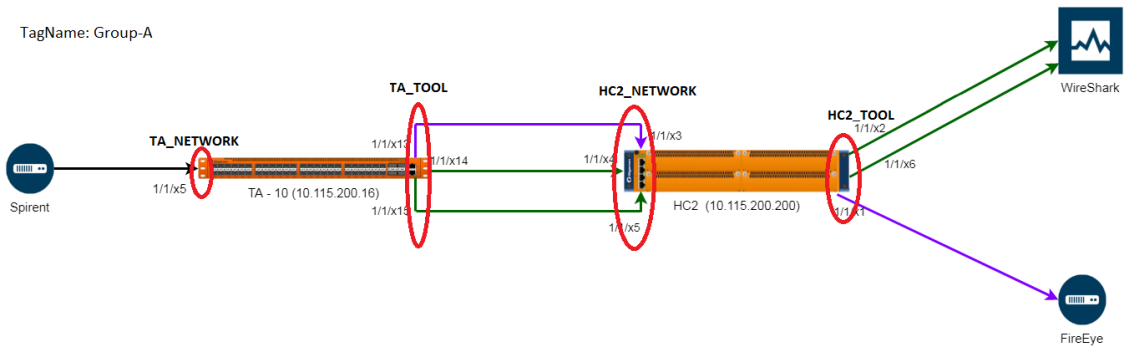

Figure 34-4: Example For Tag Creation

Using tags, you can easily view and compare the aggregated traffic flowing through a list of ports. To analyze the aggregated traffic flowing through the ports highlighted in red in Figure 34-4 on page 1257, you can create a tag with the name Group-A and assign the following tag values:

Ports	Tag Value
1/1/x5	TA_NETWORK
1/1/x13, 1/1/x14, 1/1/x15	TA_TOOL
1/1/x3, 1/1/x4, 1/1/x5	HC2_NETWORK
1/1/X2, 1/1/x6, 1/1/x1	HC2_TOOL

In the physical dashboard, you can create Traffic Comparison By Tags widget to quickly compare the aggregated traffic flowing through the ports associated with TA_NETWORK with the traffic flowing through the ports associated with HC2_NETWORK. Refer to Figure 34-5 on page 1258.

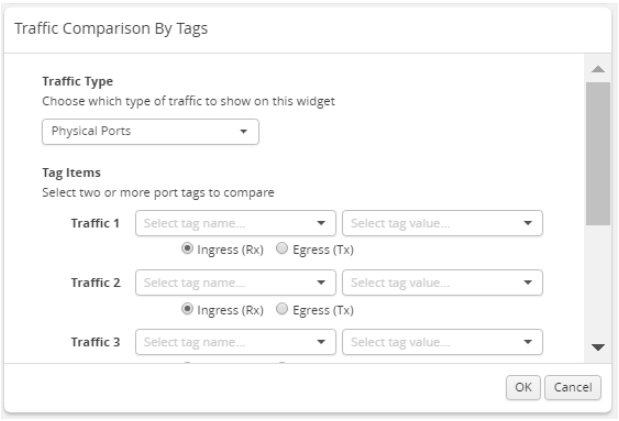

Figure 34-5: Example For Port Traffic Comparison By Tags widget

It is convenient to devise a set of tag keys and values that meet your needs for the type of clusters and objects you are managing. Using a consistent set of tag keys and values makes it easier for you to manage your resources efficiently.

Work with Sites and Tags

To view the tags, click **Administration** on the top navigation link. On the left navigation pane, select **Tags**. The existing tags are displayed in the Tags page. Refer to Figure 34-6 on page 1259.

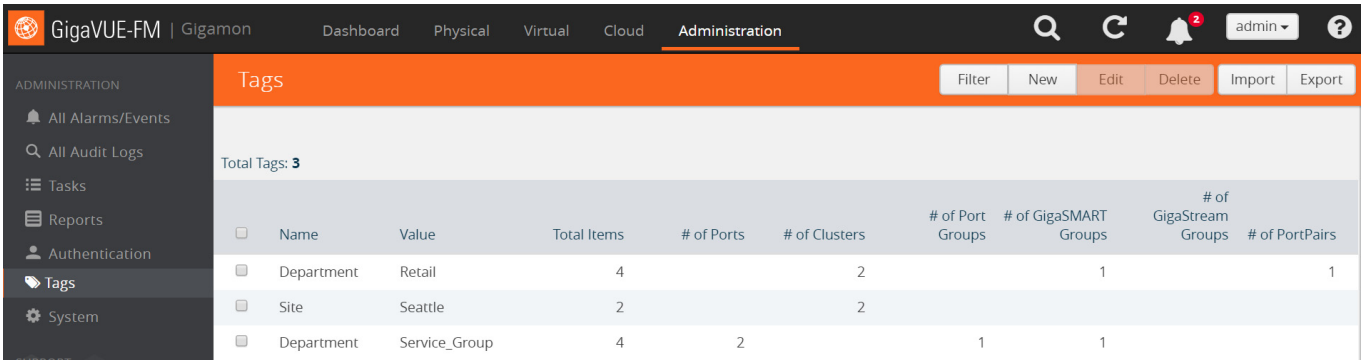

Figure 34-6: Tags Home Page

The following options are displayed in the Tags page.

Field	Description
Filter	Filters the sites and tags available in the Tags page. For more information, refer to *Filter Sites and Tags* on page 1266.
New	Creates a new site or a tag. For more information, refer to *Create Site* on page 1261 and *Create User-defined Tag* on page 1262.
Edit	Edits an existing site or a tag. For more information, refer to *Edit a Site or a Tag* on page 1265.
Delete	Deletes an existing site or a tag.
Import	Imports an existing file with sites and tags information. For more information, refer to *Import Sites and Tags* on page 1267.
Export	Exports the existing sites and tags. For more information, refer to *Export Sites and Tags* on page 1268.

The following columns are displayed in the Tags page.

Field	Description
Name	The name of the site or tag.
Value	The value of the site or tag.
Total Items	The total number of clusters, ports, port groups, GigaSMART groups, GigaStream, and port pairs associated to the site or tag.
# of Ports	The number of ports associated to the tag.
# of Clusters	The number of clusters associated to the site or tag.
# of Port Groups	The number of port groups associated to the tag.
# of GS Groups	The number of GigaSMART groups associated to the tag.
# of GigaStream Groups	The number of GigaStream groups associated to the tag.

Field	Description
# of Port Pairs	The number of port pairs associated to the tag.

In the Tags page, click on a tag or a site to view the quick view.

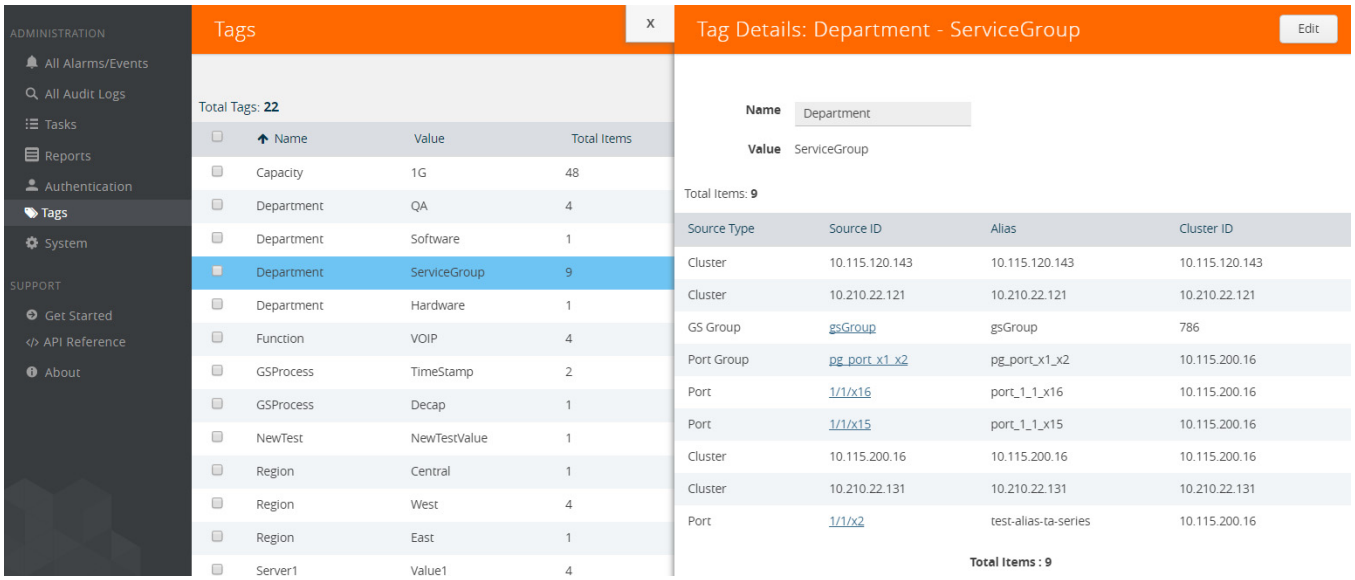

Figure 34-7: Tags Quick View

The Tag Details quick view provides information about the source types associated to the tag. They could be clusters, ports, port groups, GigaSMART groups, GigaStream, and port pairs. You can click the numbers in the columns to open the quick view and view the respective details.

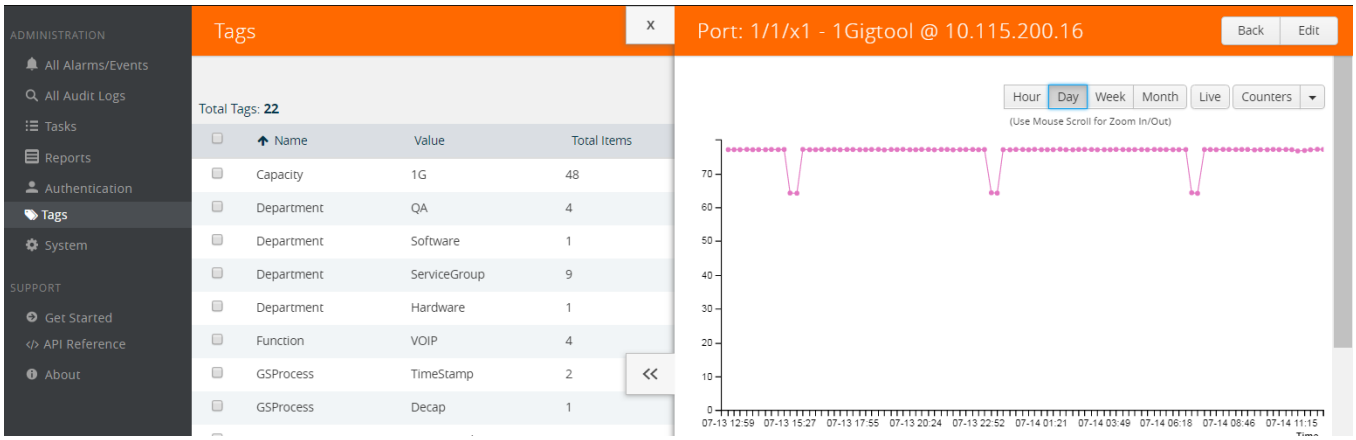

Figure 34-8: Tags - Port Details Quick View

After viewing the information, click **Back** to return to the Tag Details quick view.

Create Site

To create a site:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **Tags**.
3. In the New Tag page, select **Site**. Refer to Figure 34-9 on page 1261.

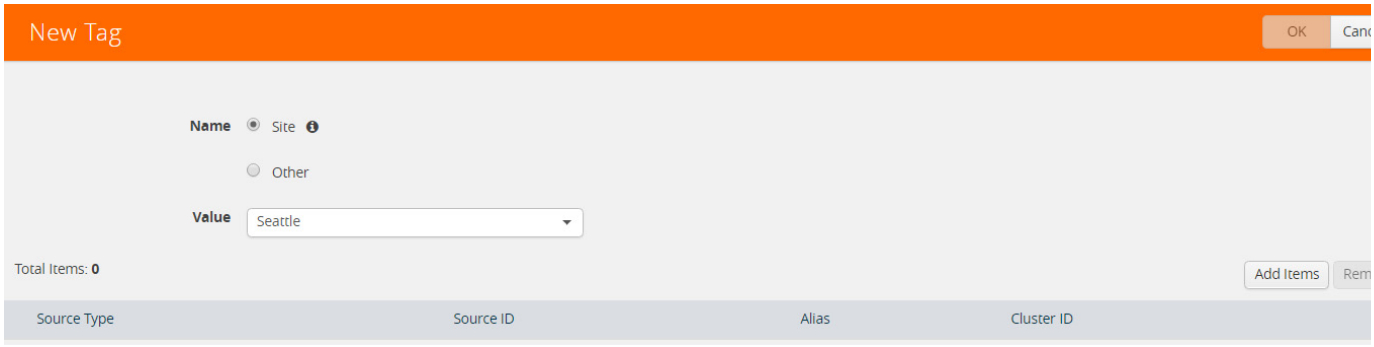

Figure 34-9: New Site Creation

4. In the **Value** text box, type the tag value and press **Enter**.
5. Click **Add Item**. The Add Items page is displayed.

   The item type and Cluster ID is grayed out. Cluster is selected as the item type by default.
6. In the Search text box, enter the cluster ID to search for the clusters or select the clusters to be associated to the site from the table.

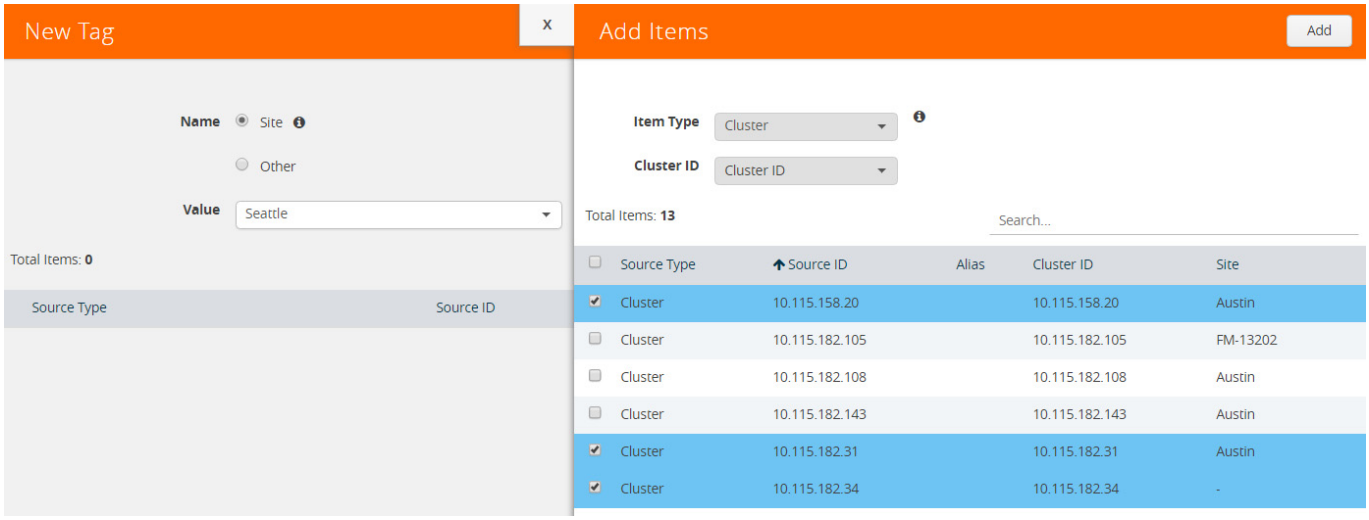

Figure 34-10: Clusters selected for Site Creation

7. Click **Add**. The clusters associated to the site are displayed in the New Tag page. Refer to Figure 34-11 on page 1262.

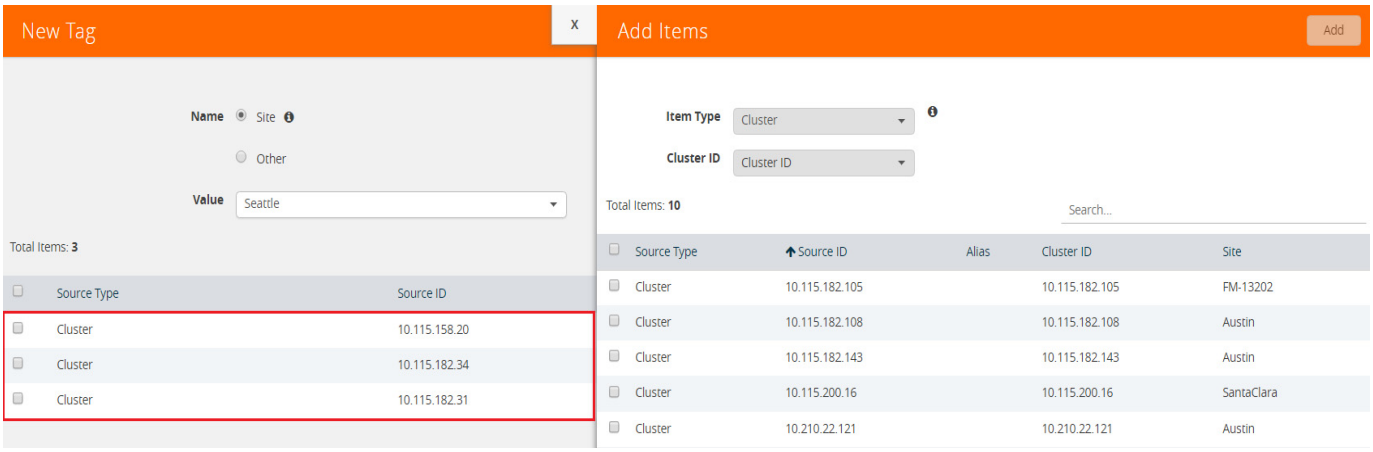

Figure 34-11: Sites Added to the New Tag Page

8. Click **OK**. Refer to Figure 34-12 on page 1262.

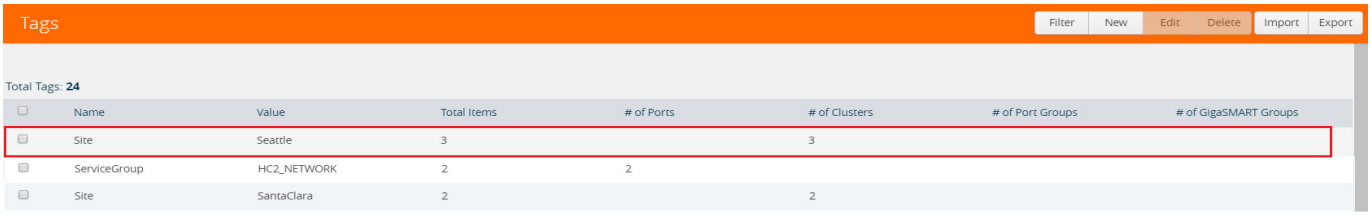

Figure 34-12: Sites Added to the Tags Page

Create User-defined Tag

To create a user-defined tag:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **Tags**.

3. In the New Tag page, select **Other**. Refer to Figure 34-13 on page 1263.

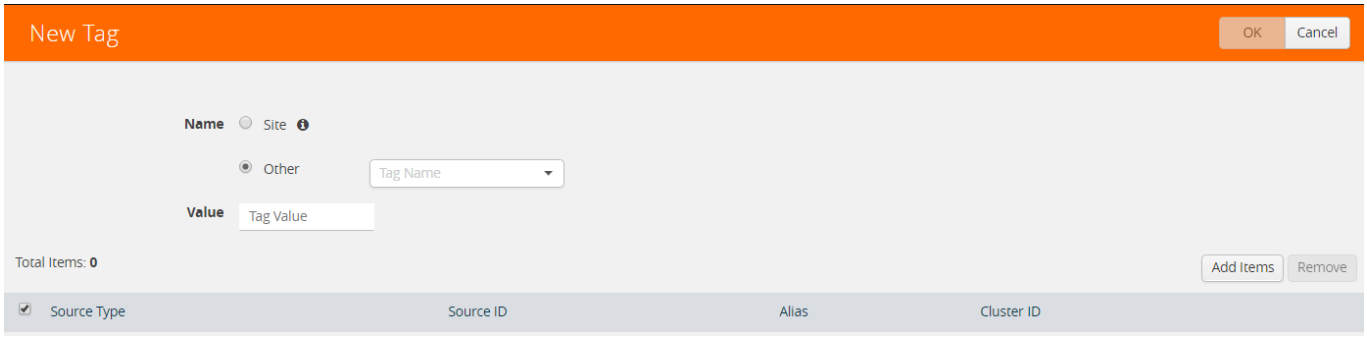

Figure 34-13: New Tag Creation

4. In the **Tag Name** box, enter the name of the tag or select an existing tag.
5. In the **Value** box, type the tag value and press **Enter**. Refer to Figure 34-14 on page 1263.

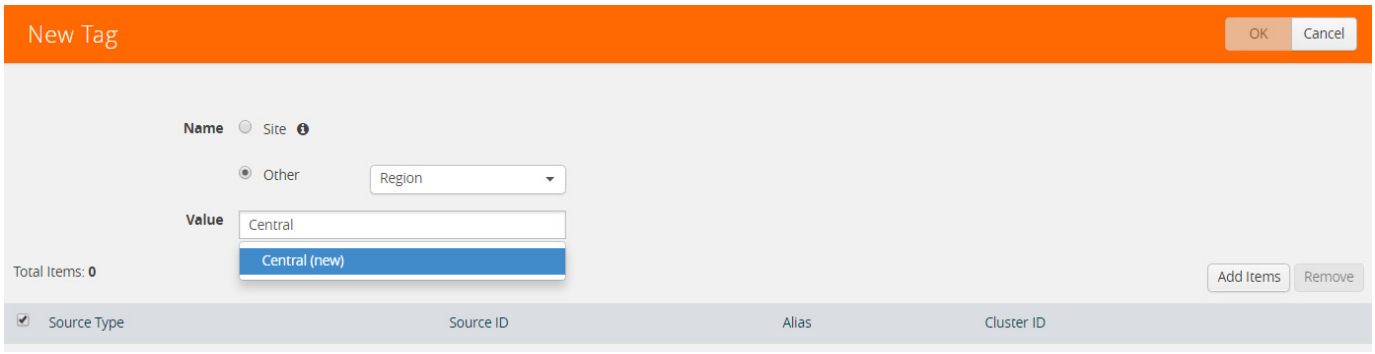

Figure 34-14: Enter the New Tag Value

6. Click **Add Items**. The Add Items page is displayed.
7. To add clusters, ports, port groups, GigaSMART groups, GigaStream, or port pairs or all of them to the new tag:

   To associate clusters to the new tag:

   a. From the **Item Type** drop-down list, select **Cluster**. All the clusters managed by GigaVUE-FM is displayed.

b. Select the clusters from the table. Refer to Figure 34-15 on page 1264.

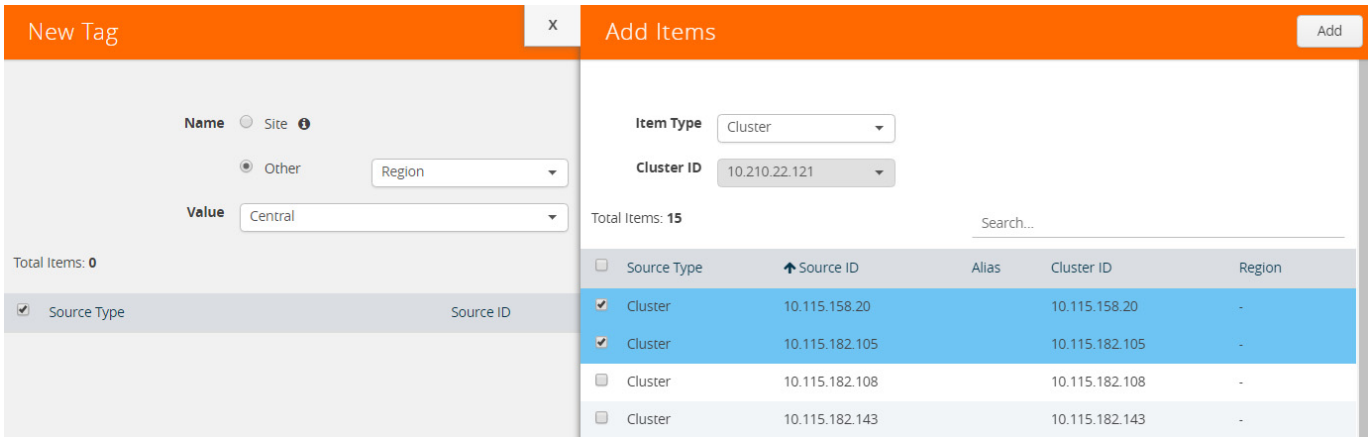

Figure 34-15: Select the Clusters

c. Click **Add**. The clusters are added to the New Tag page. Refer to Figure 34-16 on page 1264.

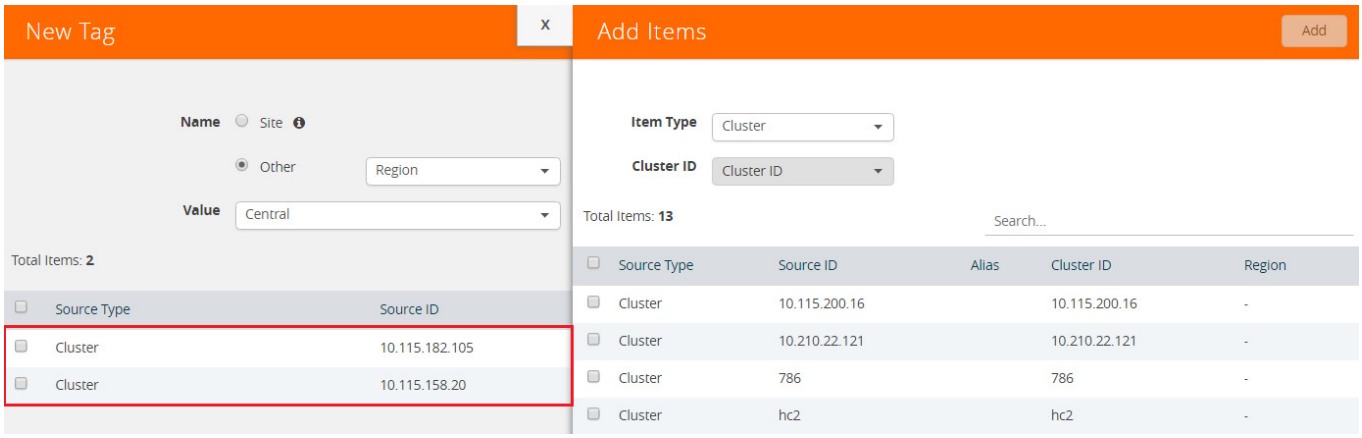

Figure 34-16: Add the Clusters to the New Tag Page

To add ports, port groups, GigaSMART groups, GigaStream, and port pairs to the new tag:

a. From the **Item Type** drop-down list, select the type.

b. From the **Cluster ID** drop-down list, select a cluster ID. The source type associated to the selected cluster is displayed.

c. Select the source types from the table. Refer to Figure 34-17 on page 1265.

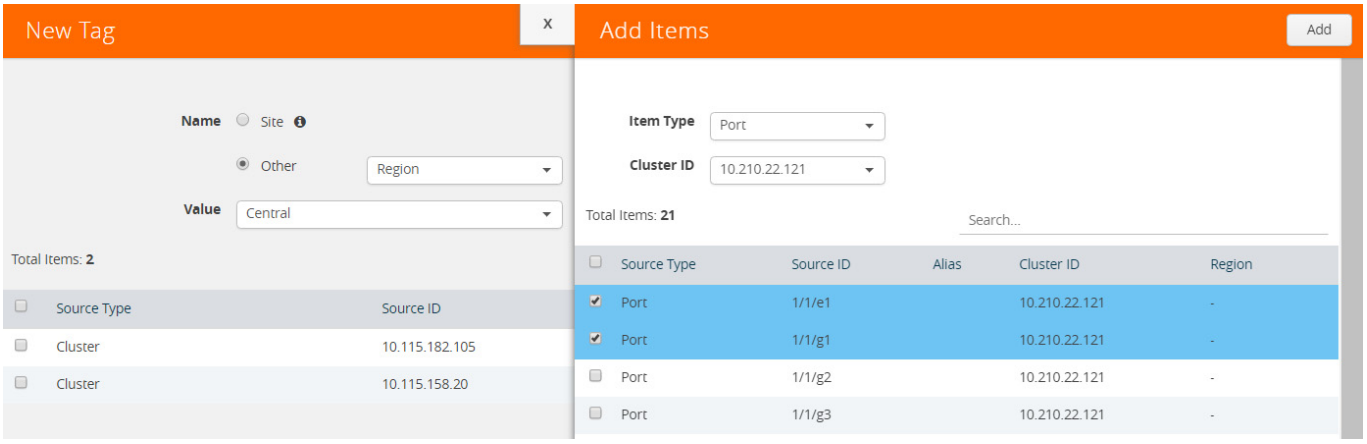

Figure 34-17: Select the Source Types

d. Click **Add**. The selected source types are added to the New Tag page. Refer to Figure 34-18 on page 1265.

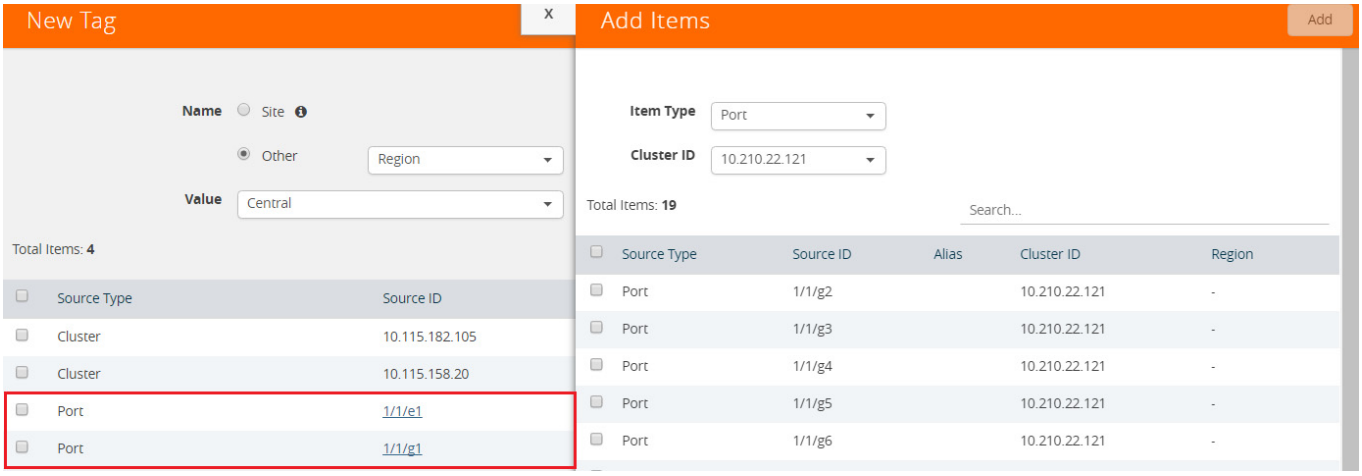

Figure 34-18: Add the Source Types to the New Tag Page

8. Click **OK**. The new tag is added to the Tags page.

Edit a Site or a Tag

To change the items associated to an existing site or tag:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **Tags**.

3. In the Tags page, select a site or a tag you want to edit and click **Edit**. Refer to Figure 34-19 on page 1266.

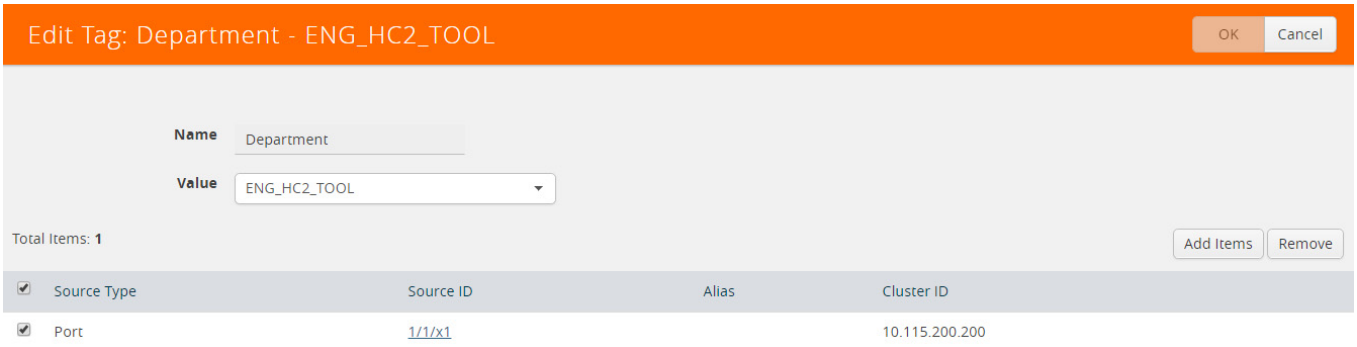

Figure 34-19: Edit Tag Page

4. In the Edit Tag page, click **Add Items** to add more items to the list or select the existing items and click **Remove** to remove the existing items from the table.
5. Click **OK**.

Filter Sites and Tags

To filter the sites or tags:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **Tags**.
3. In the Tags page, click **Filter** to filter the sites and tags. The Filter quick view is displayed. Refer to Figure 34-20 on page 1266.

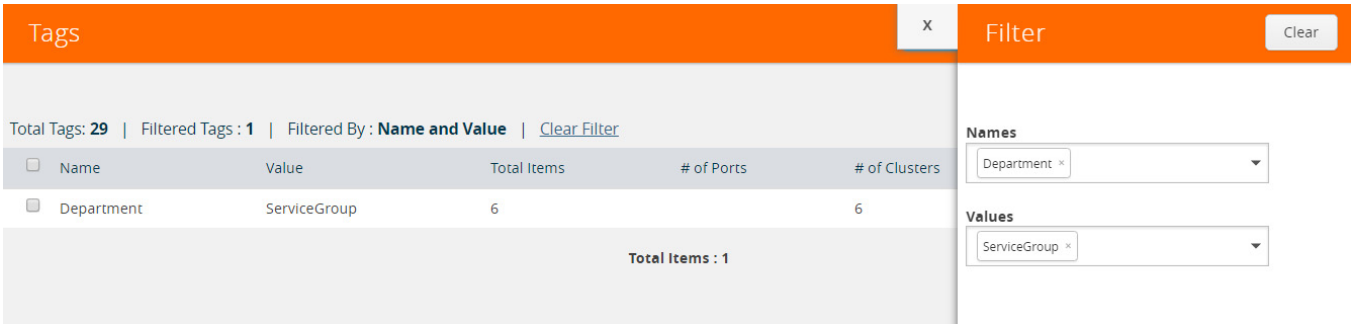

Figure 34-20: Tag Filters

4. From the **Names** drop-down list, select the name of the site or tag that you want to search. You can select multiple tags. The respective site name and tag values are displayed in the Values drop-down list.

NOTE: Either the tag name or the tag value can be selected to filter the available sites and tags.

5. From the **Values** drop-down list, select the name of the site or tag value. The results matching the filter criteria is displayed in the Tags page.

Import Sites and Tags

The sites and tags can be added to the Tags page by importing an Excel spreadsheet that contains the custom sites or tags. To view a sample Excel containing custom sites and tags, refer to Figure 34-21 on page 1267. The Excel spreadsheet must include the following columns:

- Resource type—The type can be cluster, port, port group, GS group, GigaStream, or port pairs.
- Resource ID—The resource ID depends on the type of the resource. If the resource type is Cluster, the resource ID is the cluster ID. If the resource type is Port, the resource ID is the port ID.
- Cluster ID—The IP address of the cluster or the standalone node.
- Tag Key—The site or the tag name.
- Tag value—The site value or the tag value.

	A	B	C	D	E
1	resourceType	resourceId	clusterId	tagKey	tagValue
2	Port	2/3/g28	786	Capacity	1G
3	Port	2/3/g29	786	Capacity	1G
4	Port	2/3/g30	786	Capacity	1G
5	Port	2/3/g31	786	Capacity	1G
6	Port	1/1/x5	10.115.200.16	ServiceGroup	MSO_TA_NETWORK
7	Cluster	786	786	Department	Hardware
8	Cluster	10.115.120.143	10.115.120.143	Server1	Value1
9	Cluster	10.210.22.121	10.210.22.121	Server1	Value1
10	Cluster	10.115.200.16	10.115.200.16	Server1	Value1
11	Port	1/1/x4	10.115.200.200	ServiceGroup	MSO_HC2_NETWORK
12	Port	1/1/x5	10.115.200.200	ServiceGroup	MSO_HC2_NETWORK
13	Port	1/1/x13	10.115.200.16	ServiceGroup	MSO_TA_TOOL
14	Port	1/1/x14	10.115.200.16	ServiceGroup	MSO_TA_TOOL
15	Port	1/1/x15	10.115.200.16	ServiceGroup	MSO_TA_TOOL
16	Cluster	10.210.22.121	10.210.22.121	Site	Austin
17	Port	1/1/x3	10.115.200.200	ServiceGroup	MSO_HC2_TOOL

Figure 34-21: Import Sites and Tags Excel Format

To import an Excel spreadsheet with sites and tags information:

1. Click **Import**. The Import Tags page is displayed. Refer to Figure 34-22 on page 1267.

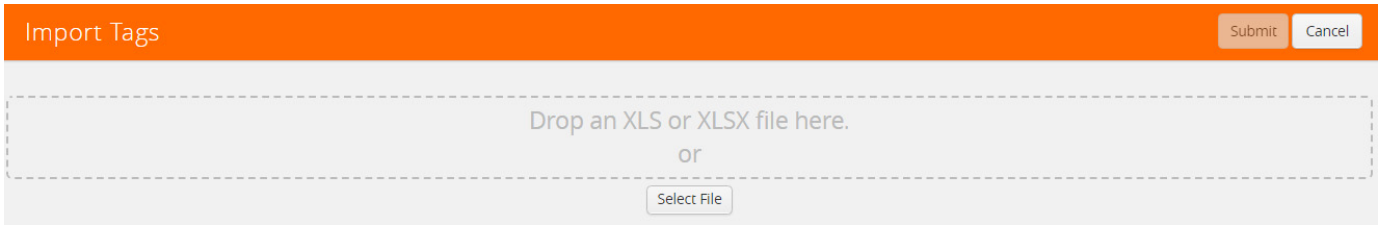

Figure 34-22: Import Tags

2. On the Import Tags page, do either of the following:
   - Click **Select File** and navigate to the file you want to import.

- Drag and drop the file onto the page.

The sites and tags added from the imported Excel spreadsheet are displayed in the Import Tags page. Only the newly added sites and tags are selected to highlight that they are newly added.

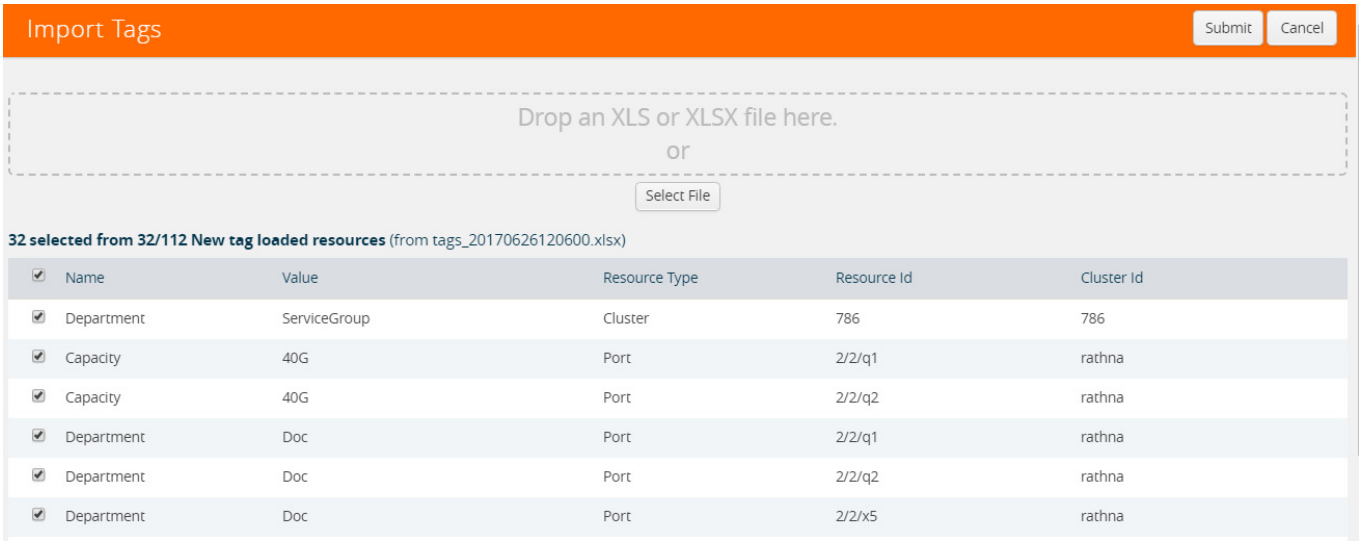

3. Click **Submit** to apply the spreadsheet information to the Tags page.

Export Sites and Tags

To export the sites and tags, click **Export**, which downloads an Excel spreadsheet. To identify the file, the filename ends in a timestamp string that is in the format tags_yyymmddhhmmss.

35 All Alarms/Events

GigaVUE-FM keeps track of all alarms and events that occur in the system. The alarms and events lists all notifiable events that have occurred in the physical, virtual, and cloud. A variety of filters are also available to filter what alarms and events are displayed on the page.

This chapter covers the following topics:

Overview of All Alarms/Events

The All Alarms/Events display the alarms and events generated from GigaVUE nodes or clusters, GigaVUE-VM virtual traffic visibility nodes, and cloud such as AWS that are stored in the GigaVUE-FM database. Refer to Figure 35-1 on page 1270.

You can also manage the records by archiving them or purging them on a regular basis. Refer to *Archive or Purge Alarm/Event Records* on page 1273.

Click **Administration** on the top navigation link. On the left navigation pane, select **All Alarms/Events**.

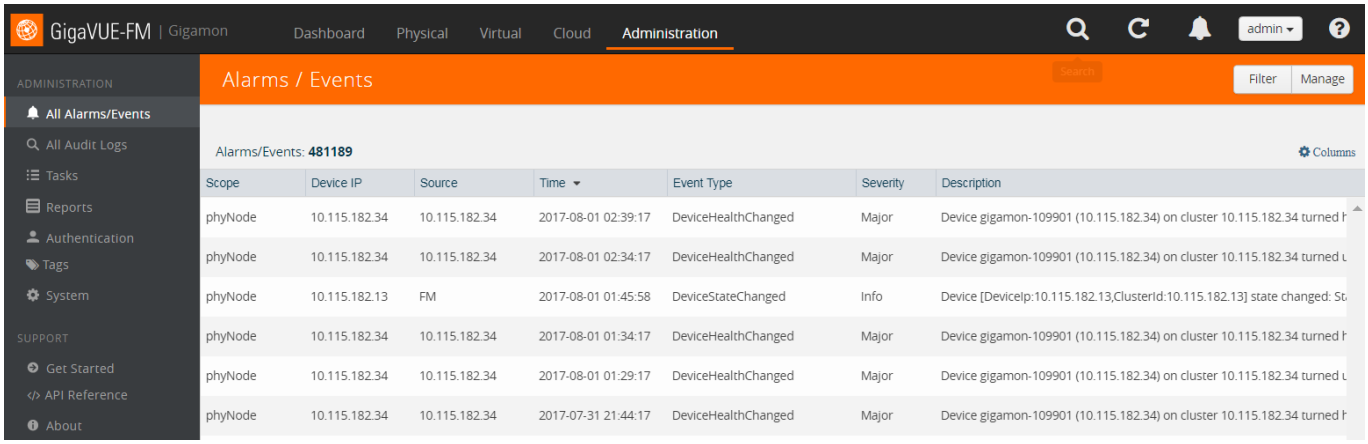

Figure 35-1: Alarms / Events Page

Table 35-1 describes the parameters recording for each alarm or event. You can also use filters to narrow down the results. Refer to *Filter Alarms/Events* on page 1272.

Table 35-1: All Alarm/Event Parameters

Controls/ Parameters	Description
Filter	Opens the Filter quick view for narrowing down the alarms and events to view the desired results.
Manage	Opens the Manage Alarms/Event page for exporting and selecting records for archiving or purging. For more information, refer to *Filter Alarms/Events* on page 1272. **NOTE:** This option is not available in the Physical and Virtual Alarms/Events page.
Settings	Opens the Health Monitor Threshold Settings page for setting the Health Monitor thresholds. For more information about the Health Monitor and setting the thresholds, refer to *How to Set Health Monitor Alarm Thresholds and Notifications* on page 167. **NOTE:** This option is not available in the Physical and Virtual Alarms/Events page.

Table 35-1: All Alarm/Event Parameters

Controls/ Parameters	Description
Source	The source from where the alarms and events are generated. The criteria can be as follows: • FM - indicates the event was flagged by the Fabric Manager. • IP address - is the address of the GigaVUE H Series or GigaVUE G Series node that detected the event. For a node to be able to send notifications to the Fabric Manager, the SNMP_TRAP must be configured with the Fabric Manager's IP address specified as a host. Refer to the GigaVUE-OS User's Guides for instructions on adding a destination for SNMP traps. • VMM - indicates the event was flagged by the Virtual Machine Manager. • FM Health - indicates the event was flagged due to the health status change of GigaVUE-FM. • FTA- indicates the event was flagged by the FabricVUE Traffic Analyzer. For example, if there is a change in the netflow processing statistics, then the source is displayed as FTA.
Time	The timestamp when the event occurred. **IMPORTANT:** Timestamps are shown in the time zone of the client browser's computer and not the timezone of the node reporting the event. The timestamp is based on the correctly configured clock on the GigaVUE-FM server and converted from UTC to the client computer's configured timezone.
Scope	The category to which the alarms or events belong. Alarms and events can belong to the following category: Domain, Node, Card, Port, Stack, Cluster, Chassis, GigaVUE-FM, GigaVUE-VM, and so on. For example, if there is a notification generated for port utilization low threshold, the scope is displayed as Physical Node.
Event Type	The type of event that generated the alarms and events. The type of events can be CPU utilization high, cluster updated, device discovery failed, fan tray changed, netflow statistics, and so on.
Severity	The severity is one of Critical, Major, Minor, or Info. Info is informational messages. For example, when power status change notification is displayed, then the message is displayed as Info.
Affected Entity Type	The resource type associated with the alarm or event. The resource type is displayed only for ports, cards, fans, and boxes. For example, when low disk space notification is generated, Box is displayed as the affected entity type.
Affected Entity	The resource ID of the affected entity type. For example, when low disk space notification is generated, the IP address of the node with the low disk space is displayed as the affected entity.
Description	The description of the event, which includes any of the possible notifications with additional identifying information where appropriate (such as reporting nodes IP address, user name, and so on).
Device IP	The IP address of the device.
Host Name	The host name of the device.

NOTE: The columns in the Alarms / Events page can be customized based on the type of content you want to view in the table. For customizing the columns, refer to *Table View Customization* on page 45.

Filter Alarms/Events

The alarms and events can be filtered based on the following criteria:

Controls/ Parameters	Description
Source	Displays the alarms and events generated by a specific source. **NOTE:** This option is not available in the Virtual Alarms/Events page.
Start Date **End Date**	Displays the alarms and events occurred within a specific time range.
Scope	Displays the alarms and events associated with the selected category. For example, physical node, physical port, appliance server, and so on.
Event Type	Displays the events associated with the selected event type.
Severity	Displays the alarms and events that match the selected severity level.
Affected Entity Type	Displays the alarms and events associated with the affected entity type. The affected entity type can be ports, cards, fans, or boxes.
Affected Entity	Displays the alarms and events associated with the affected entity. The affected entity can be port ID, slot label, fan name, and so on.
Device IP	Displays the alarms and events associated with the IP address of the device. Partial IP addresses may be entered to display the results containing the specified octets. For example, if the last 2 octets of the IP address entered is 46.100, the IP addresses listed will include all those that end with 46.100.
Host Name	Displays the alarms and events associated with the host name of the device. Partial host name may be entered to filter the events. For example, if the first portion of the host name entered is GIMO, the host names listed will include all those that contain GIMO.

To filter the alarms and event:

1. Click **Filter**.

   The Filter quick view appears.

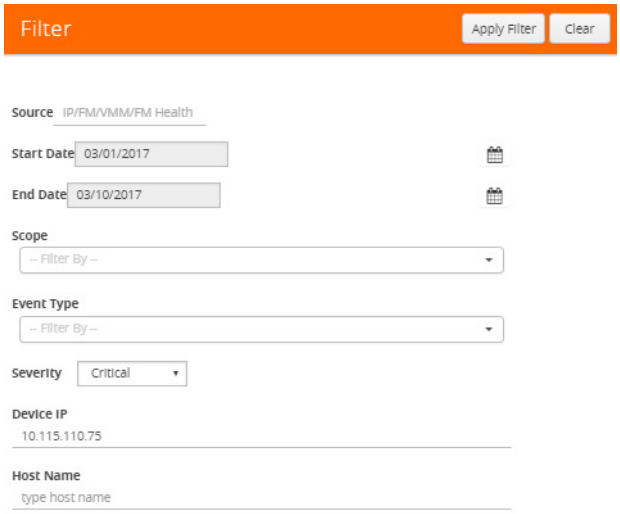

2. Specify the filter criteria, then click **Apply Filter**.

Archive or Purge Alarm/Event Records

Alarm/Events are saved in the FM database. Alarms/Events records continues to grow over time. GigaVUE-FM allows you to archive and purge these records based on a specific date. Records older than that date will be exported to an SFTP server.

When archiving, the records are archived as a CSV file with a timestamp appended. For example, audit_20151005105607.csv. The file is compressed to a zip file before exporting to the server.

The archive and purge option for alarms/events records is only available to super_admin users. The audit and purge action for alarms/events is also recorded to the audit log.

Archive Alarm/Event Records

To archive the alarm/events records, do the following:

1. Select **Alarms/Events** in the navigation pane.
2. Click **Manage**.
3. Click the Calendar icon and select a date. Records older than this date will be exported.
4. Select **Export Records** and specify the following:
   - The address of the SFTP Server to which the logs will be exported.
   - The user name and password for the SFTP server.
   - The file path on the server where the files will be archived. The file should be in the .zip format.
5. Click **OK** to export to the records to the SFTP server.

Purge Alarms/Events Records

The alarms/events data continues to grow over time. You can purge the records, by doing the following:

1. Select **Alarms/Events** in the navigation pane.
2. Click **Manage**.
3. Click the Calendar icon and select a date. Records older than this date will be purged.
4. Select **Purge Selected Records**.
5. Click **OK** to purge the records.

Archive and Purge Alarms/Events Records

Audit log records can be exported and purged at the same time by doing the following:

1. Select **Alarms/Events** in the navigation pane.
2. Click **Manage**.
3. Click the **Calendar** icon and select a date. Records older than this date will be exported.
4. Select **Export Records** and specify:
   - The address of the SFTP Server to which the records will be exported.
   - The user name and password for the SFTP server.
   - The file path on the server where the files will be archived. The file should be in the .zip format.
5. Select **Purge Selected Records**.
6. Click **OK** to export the records to the SFTP server, and then purge the records.

36 All Audit Logs

This section describes the Audit Logs page and provides information about filtering and managing the logs. The topics covered are:

Overview of Audit Logs

The Audit Logs page captures audit logs for all users connected to the given GigaVUE-FM. There are 10 results shown by default on every page. The logs can also be further filtered to view specific information. Unlike the zipped logs under **Admin**> **System > Logs,** the audit logs can be seen by users. For more information about filtering, refer to *Filtering Audit Logs* on page 1276.

The Audit Logs have the following parameters:

Parameters	Description
Time	Provides the timestamp on the log entries.
User	Provides the logged user information.
Operation Type	Provides specific entries that are logged by the system such as: • Log in and Log out based on users. • Create/Delete/Edit tasks, GS operations, maps, virtual ports, and so on.
Source	Provides details on whether the user was in FM or on the node when the event occurred.
Status	Success or Failure of the event.
Description	In the case of a failure, provides a brief update on the reason for the failure.

NOTE: Ensure that the GigaVUE-FM time is set correctly to ensure accuracy of the trending data that is captured.

Filtering Audit Logs

Filtering the audit logs allows you to display only those items of interest. You can filter based on any of the following:

- When—display logs that occurred within a specified time range.
- Who—display logs related a specific user or users.
- What—display logs for one or more operations, such as Create, Read, Update, and so on.
- Where—display logs for GigaVUE-FM or devices.
- Result—display logs for success or failure.

To filter the audit logs, do the following:

1. Click **Filter**.

   The quick view for Audit Log Filters displays.

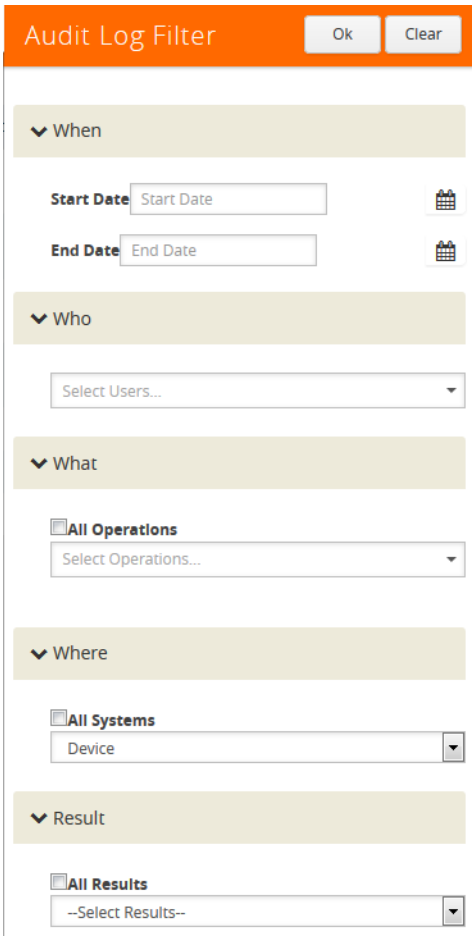

2. Specify any or all of the following:
   - **Start Date** and **End Date** to display logs within a specific time range.
   - **Who** limits the scope of what displays on the Audit Logs page to a specific user or users.
   - **What** narrows the logs to the types of operation that the log is related to. You can select multiple operations. Select **All Operations** to apply all operation types as part of the filter criteria.
   - **Where** narrows the logs to particular of system that the log is related to, either FM or device. Select **All Systems** apply both FM and device to the filter criteria.
   - **Result** narrows the logs related to failures or successes. Select All Results to apply both success and failure to the filter criteria.
3. Click **OK** to apply the selected filters to the Audit Logs page.

Archive or Purge Audit Log Records

Audit logs are save to the FM database. Audit log records continues to grow over time. GigaVUE-FM allows you to archive these records based on a specific date. Records older than that date will be exported to an SFTP server.

The records are output are archived as a CSV file with a timestamp appended. For example, audit_20151005105607.csv. The file is compressed to a zip file before exporting to the server.

The archive and purge option for audit log records is only available to super_admin users. The audit and purge action for audit logs is also recorded to the audit log. The Purge action for the audit log never purges the purge entry.

Archive Audit Logs

To archive the audit log records, do the following:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **All Audit Logs** in the navigation pane.
3. Click **Manage**.
4. Click the **Calendar** icon and select a date. Records older than this date will be exported.
5. Select **Export Records** and specify:
   - The address of the SFTP server to which the records will be exported.
   - The user name and password for the SFTP server.
   - The file path on the server where the files will be archived.
6. Click **OK** to exported to the records to the SFTP server.

Purge Audit Log Records

The audit log data continues to grow over time. You can purge the audit log records, by doing the following:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **All Audit Logs** in the navigation pane.
3. Click **Manage**.
4. Click the **Calendar** icon and select a date. Records older than this date will be purged.
5. Select **Purge Selected Records**.
6. Click **OK** to purge the records.

Archive and Purge Audit Log Records

Audit log records can be exported and purged at the same time by doing the following:

1. Select **Audit Logs** in the navigation pane.
2. Click **Manage**.
3. Click the Calendar icon and select a date. Records older than this date will be exported.
4. Select **Export Records** and specify:
   - The address of the SFTP Server to which the records will be exported.
   - The user name and password for the SFTP server.
   - The file path on the server where the files will be archived.
5. Select **Purge Selected Records**.
6. Click **OK** to exported the records to the SFTP server, and then purge the records.

37 Tasks

The Tasks page provides access to the Admin Tasks and Scheduled Tasks pages. The Admin Task page displays any administrative tasks waiting to occur on the nodes managed by GigaVUE-FM. The Scheduled Tasks page displays the scheduled reoccurring task on the nodes

This section covers the following topics:

Admin Tasks

Currently, the only tasks that can be scheduled are node image installs, node upgrades, and node reboots. Once a task listed in this table executes, it also appears in the Events list. To view the Admin Tasks, click **Administration** on the top navigation link. On the left navigation pane, click Tasks.

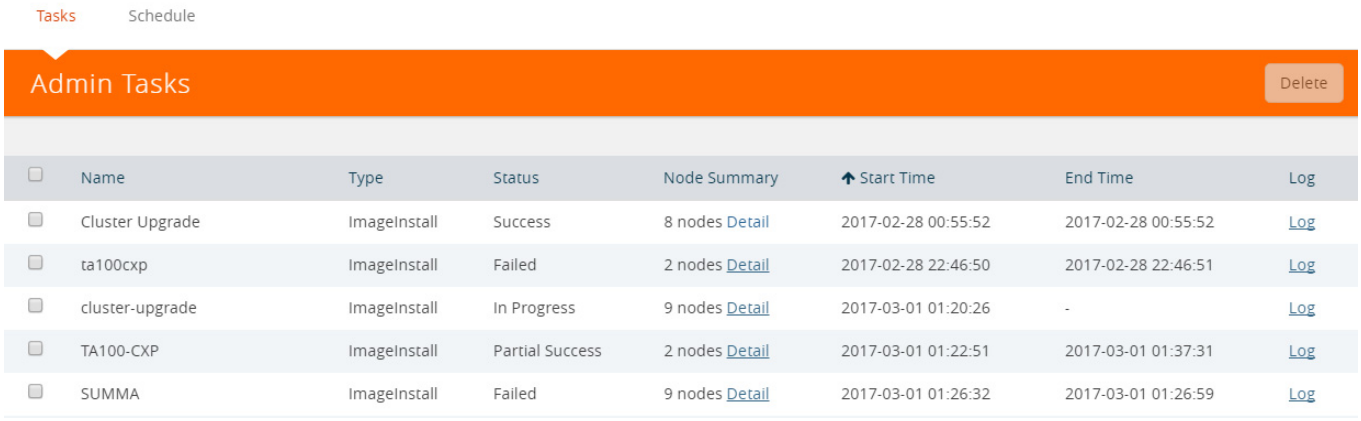

Figure 37-1: Admin Tasks Page

The Admin Tasks page displays the following information:

Parameters	Description
Name	The name of the task.
Type	The type of task that is scheduled, for example, **Node Reboot**.
Status	The status of the task. They are **In Progress**, **Success**, or **Failure**.
Node Summary	The total number of nodes available in the clusters. Click Details to view the post upgrade sanity check of all the available configuration objects in the cluster. For more information, refer to *View Upgrade Sanity Check* on page 1283.
Start Time	The start time of the scheduled task.
End Time	The end time of the scheduled task. **IMPORTANT:** Timestamps are shown in the time zone of the client browser's computer and not the timezone of the node where the task is scheduled. The timestamp is based on the correctly configured clock on the GigaVUE-FM server and converted from UTC to the client computer's configured timezone.
Log	The log records every step performed with the timestamps. Click **Log** for a detailed view of every step that occurred during the upgrade process. Refer to Figure 37-2 on page 1283.

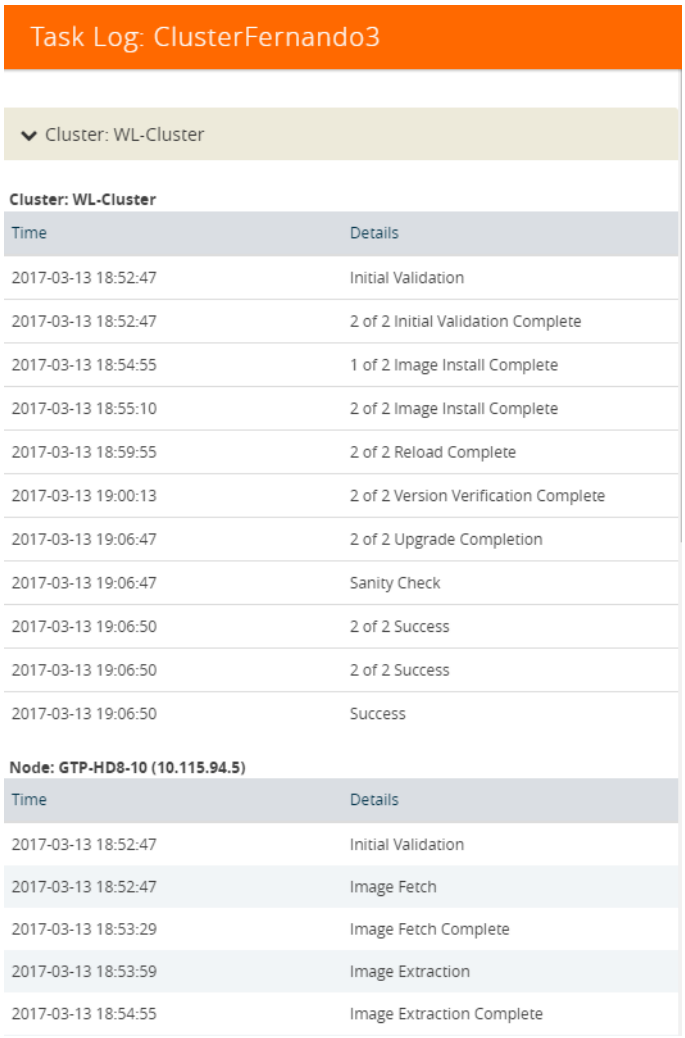

Figure 37-2: Admin Tasks Log Quick View

View Upgrade Sanity Check

The information in the Node Summary Details page is grouped based on clusters. Each cluster displays the configuration objects and their state before and after the upgrade. For example, if cards are down after the upgrade, the number of cards that are down are displayed in the Result column. Click the number to view more details about the cards that are down.

To view the upgrade sanity check:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **Tasks**.

3. In the Admin Tasks page, click the Detail link in the Node Summary column. Refer to Figure 37-3 on page 1284.

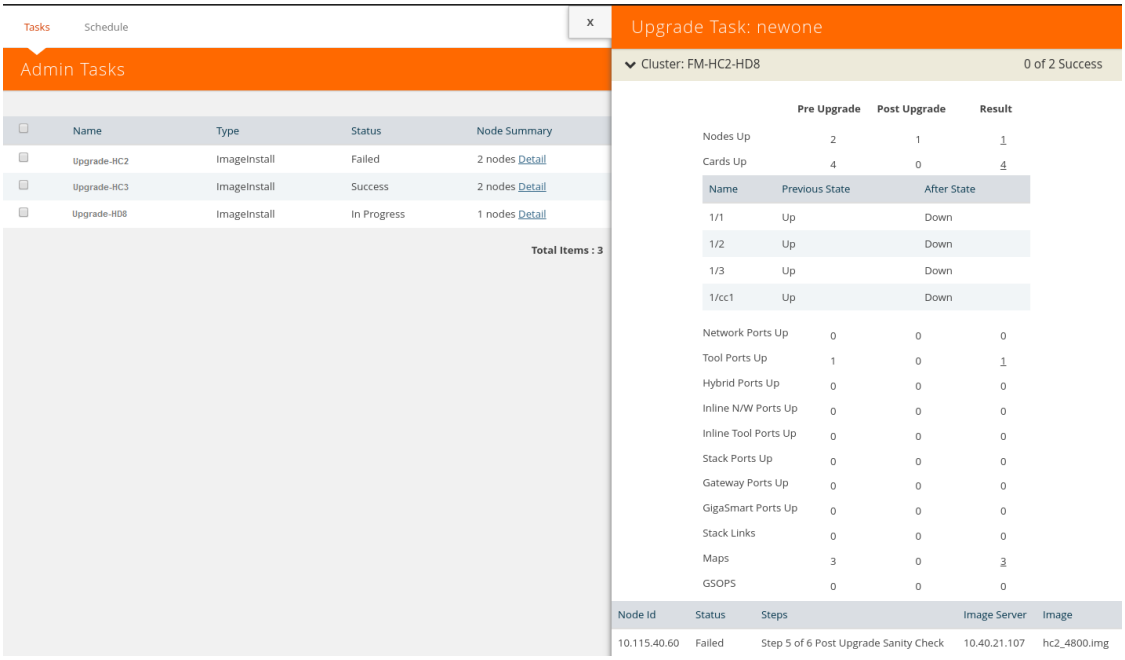

Figure 37-3: Admin Tasks - Detail Page

4. In the Result column, click the number and view the detailed information about the configuration objects.

Delete Admin Task

To delete a scheduled task from the Admin Tasks, do the following:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **Tasks**.
3. From the Admin Tasks page, select a task from the list.
4. Click **Delete**.

   The task is unscheduled and stopped the task from happening.

Scheduled Tasks

The Scheduled Tasks page displays tasks that have been set to reoccur at scheduled times. Currently, the only tasks that can be scheduled are device backups, GigaVUE nodes upgrade, and GigaVUE-FM configuration data.

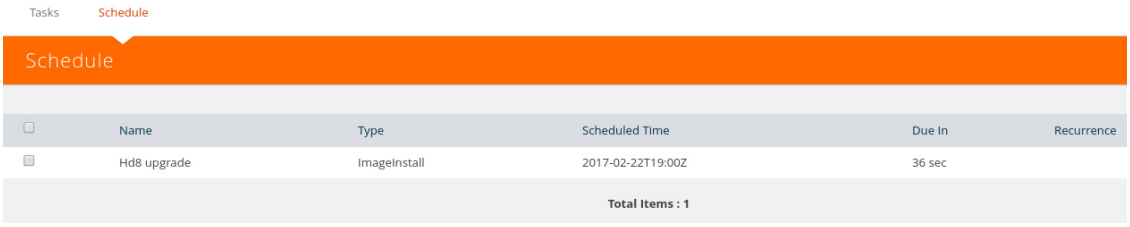

Figure 37-4: Scheduled Tasks Page

The Scheduled Task page displays the following information.

Parameters	Description
Name	The name of the scheduled task.
Type	The type of the scheduled task.
Scheduled Time	The timestamp when the task is scheduled to begin.
Due In	The time left for the scheduled task to begin.
Recurrence	The frequency of the scheduled task. For example, daily at 4 hours 35 minutes.

Delete Scheduled Task

To reschedule a task, do the following:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **Tasks > Scheduled Task**.
3. Select a task from the list.

   The task is either a configuration data backup of GigaVUE-FM (FMServerConfigBackUp) or a backup of a device (configBackup).
4. Click **Delete**.

   Clicking Delete unschedules and stops the task from happening.

38 Reports

The Reports page lists different templates that you can use to generate reports. This section covers the following topics:

Overview of Reports

The reports can be downloaded in PDF or HTML format to your local drive.

- Only one report can be selected for each generate and download option.
- The report layout and format is not customizable.
- Reports page is available based on the GigaVUE-FM and GigaVUE-VM licenses installed on the system. See the *Licensing* section for more details.
- Reports can be stored or deleted on the GigaVUE-FM.
- Reports are polled live and therefore can change each time they are generated.
- Each report appears with the timestamp on when the report was generated.

To view the reports, click **Administration** on the top navigation link. On the left navigation pane, select **Reports**. Click **Generate New**.

Figure 38-1: Reports Page View

NOTE: To view the reports directly from the GigaVUE-FM settings, ensure that the pop-up blocker settings on your browser are disabled. This will allow you to view the reports without downloading. The reports will be available on a separate page.

After the report is generated, if you wish to view it, the browser will try to open a new window. However, if you have a pop-up blocker enabled, you will need to disable the pop-up blocker to view the pages.

Report Templates

This section describes the report templates available for generating reports:

- *Template 1: Visibility Fabric Performance Report* on page 1289 provides traffic analysis information.
- *Template 2: Visibility Fabric Node Details Report* on page 1290 provides specific details relating to the physical nodes (includes H Series, G Series and TA Series).
- *Template 3: GigaVUE-VM Report* on page 1291 provides for GigaVUE-VM traffic analysis information.
- *Template 4: Visibility Fabric Inventory Report* on page 1292 provides a summary of all physical inventory (includes H Series, G Series and TA Series) that is visible on GigaVUE-FM.
- *Template 5: GigaSMART Performance Report* on page 1293 summary on GigaSMART performance for all H Series nodes with GigaSMART functionality.

Template 1: Visibility Fabric Performance Report

This multi-page template provides you with printable format for Traffic analysis including:

- Top N / Bottom N Ports
- Top N / Bottom N Traffic Maps
- Overlay Traffic Maps / Ports
- Top N / Bottom N VM rule stats
- Top N/ Bottom N Logical Network stats

Figure 38-2 shows an example of a report for Visibility Fabric Performance.

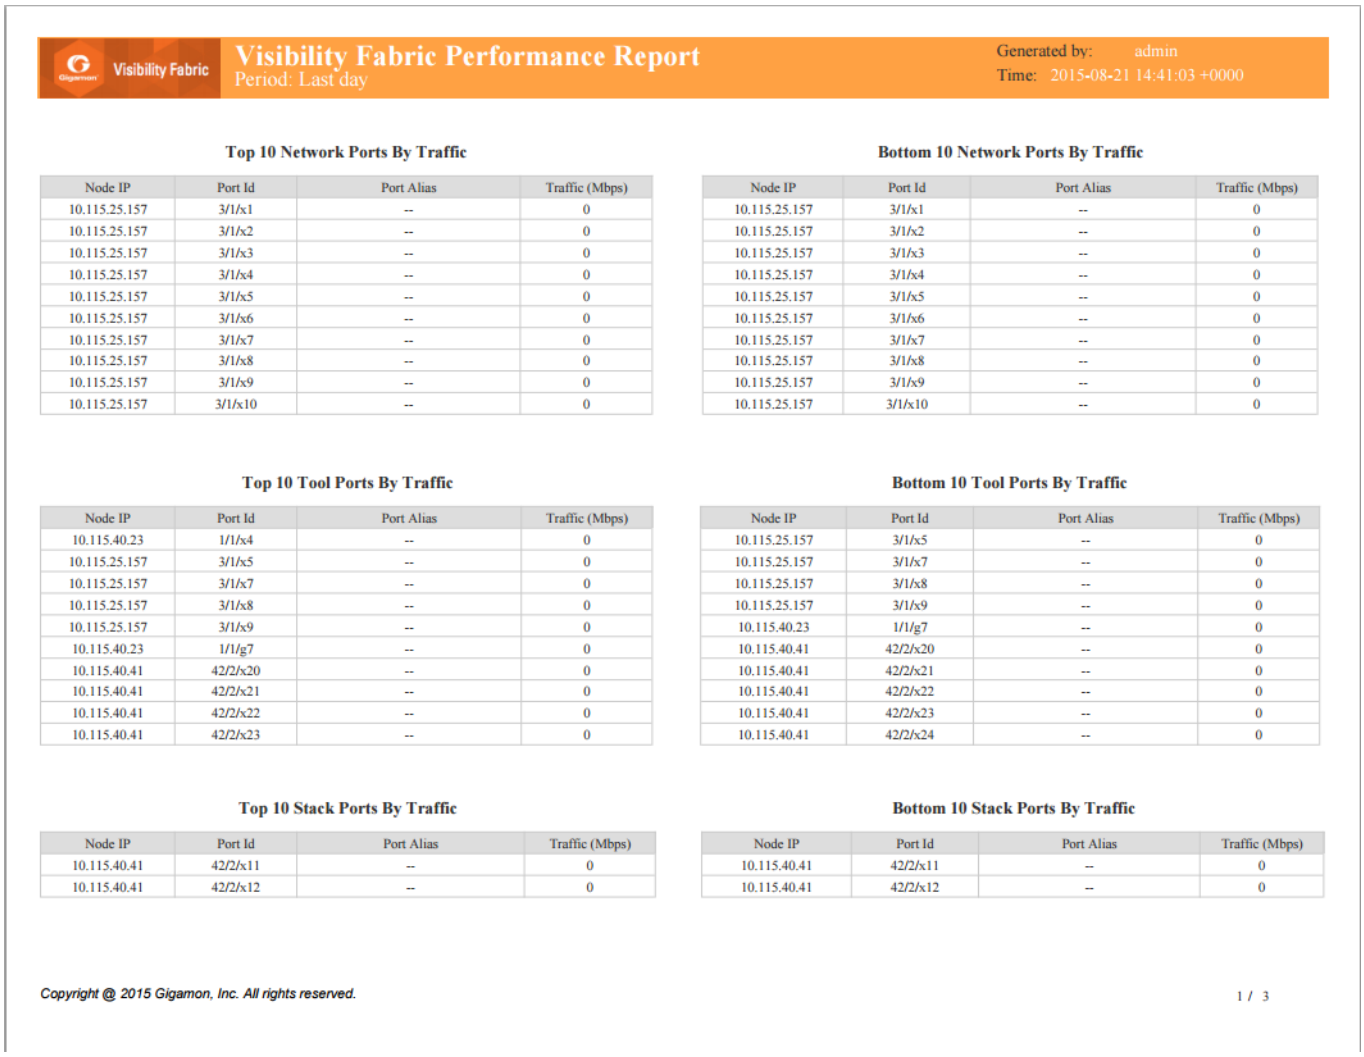

Visibility Fabric

Visibility Fabric Performance Report
Period: Last day

Generated by: admin
Time: 2015-08-21 14:41:03 +0000

Top 10 Network Ports By Traffic

Node IP	Port Id	Port Alias	Traffic (Mbps)
10.115.25.157	3/1/x1	--	0
10.115.25.157	3/1/x2	--	0
10.115.25.157	3/1/x3	--	0
10.115.25.157	3/1/x4	--	0
10.115.25.157	3/1/x5	--	0
10.115.25.157	3/1/x6	--	0
10.115.25.157	3/1/x7	--	0
10.115.25.157	3/1/x8	--	0
10.115.25.157	3/1/x9	--	0
10.115.25.157	3/1/x10	--	0

Bottom 10 Network Ports By Traffic

Node IP	Port Id	Port Alias	Traffic (Mbps)
10.115.25.157	3/1/x1	--	0
10.115.25.157	3/1/x2	--	0
10.115.25.157	3/1/x3	--	0
10.115.25.157	3/1/x4	--	0
10.115.25.157	3/1/x5	--	0
10.115.25.157	3/1/x6	--	0
10.115.25.157	3/1/x7	--	0
10.115.25.157	3/1/x8	--	0
10.115.25.157	3/1/x9	--	0
10.115.25.157	3/1/x10	--	0

Top 10 Tool Ports By Traffic

Node IP	Port Id	Port Alias	Traffic (Mbps)
10.115.40.23	1/1/x4	--	0
10.115.25.157	3/1/x5	--	0
10.115.25.157	3/1/x7	--	0
10.115.25.157	3/1/x8	--	0
10.115.25.157	3/1/x9	--	0
10.115.40.23	1/1/g7	--	0
10.115.40.41	42/2/x20	--	0
10.115.40.41	42/2/x21	--	0
10.115.40.41	42/2/x22	--	0
10.115.40.41	42/2/x23	--	0

Bottom 10 Tool Ports By Traffic

Node IP	Port Id	Port Alias	Traffic (Mbps)
10.115.25.157	3/1/x5	--	0
10.115.25.157	3/1/x7	--	0
10.115.25.157	3/1/x8	--	0
10.115.25.157	3/1/x9	--	0
10.115.40.23	1/1/g7	--	0
10.115.40.41	42/2/x20	--	0
10.115.40.41	42/2/x21	--	0
10.115.40.41	42/2/x22	--	0
10.115.40.41	42/2/x23	--	0
10.115.40.41	42/2/x24	--	0

Top 10 Stack Ports By Traffic

Node IP	Port Id	Port Alias	Traffic (Mbps)
10.115.40.41	42/2/x11	--	0
10.115.40.41	42/2/x12	--	0

Bottom 10 Stack Ports By Traffic

Node IP	Port Id	Port Alias	Traffic (Mbps)
10.115.40.41	42/2/x11	--	0
10.115.40.41	42/2/x12	--	0

1 / 3

Figure 38-2: Visibility Fabric Performance Report

Template 2: Visibility Fabric Node Details Report

This multi-page template provides you with printable format for specific details relating to the Physical Nodes (includes H Series, G Series and TA Series) similar to what you would see under Chassis/Device pages. The report includes:

- Pie Chart of total Nodes, total (collective) ports and card types (collective)
- Table format showing each Node associated with this instance of FM
- Detailed report similar to as shown on Chassis Page including clustered nodes

Figure 38-3 shows an example of a report for Visibility Fabric Node Details.

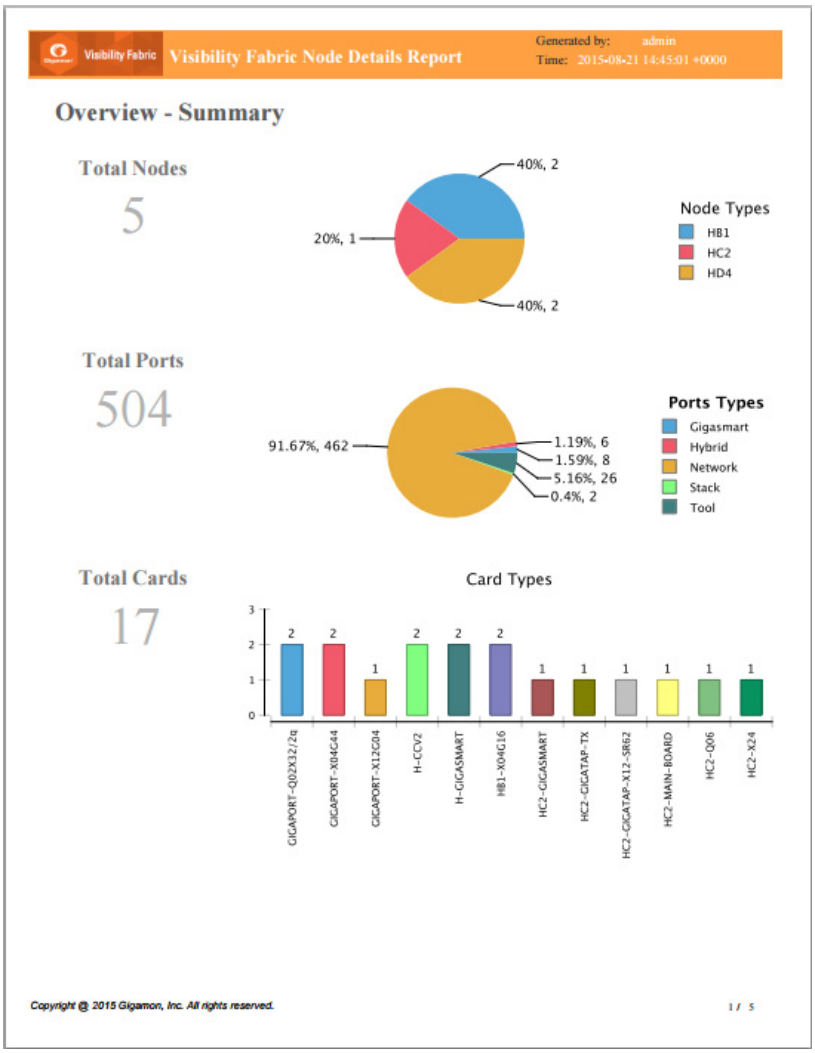

Figure 38-3: Inventory Details Report

Template 3: GigaVUE-VM Report

This multi-page template provides you with printable format for GigaVUE-VM traffic analysis including:

- Summary of GigaVUE-VM virtual centers
- Details on the virtual centers
- Top N / Bottom N Ports
- Top N / Bottom N Traffic Maps

Figure 38-4 show an example of a report for GigaVUE-VM.

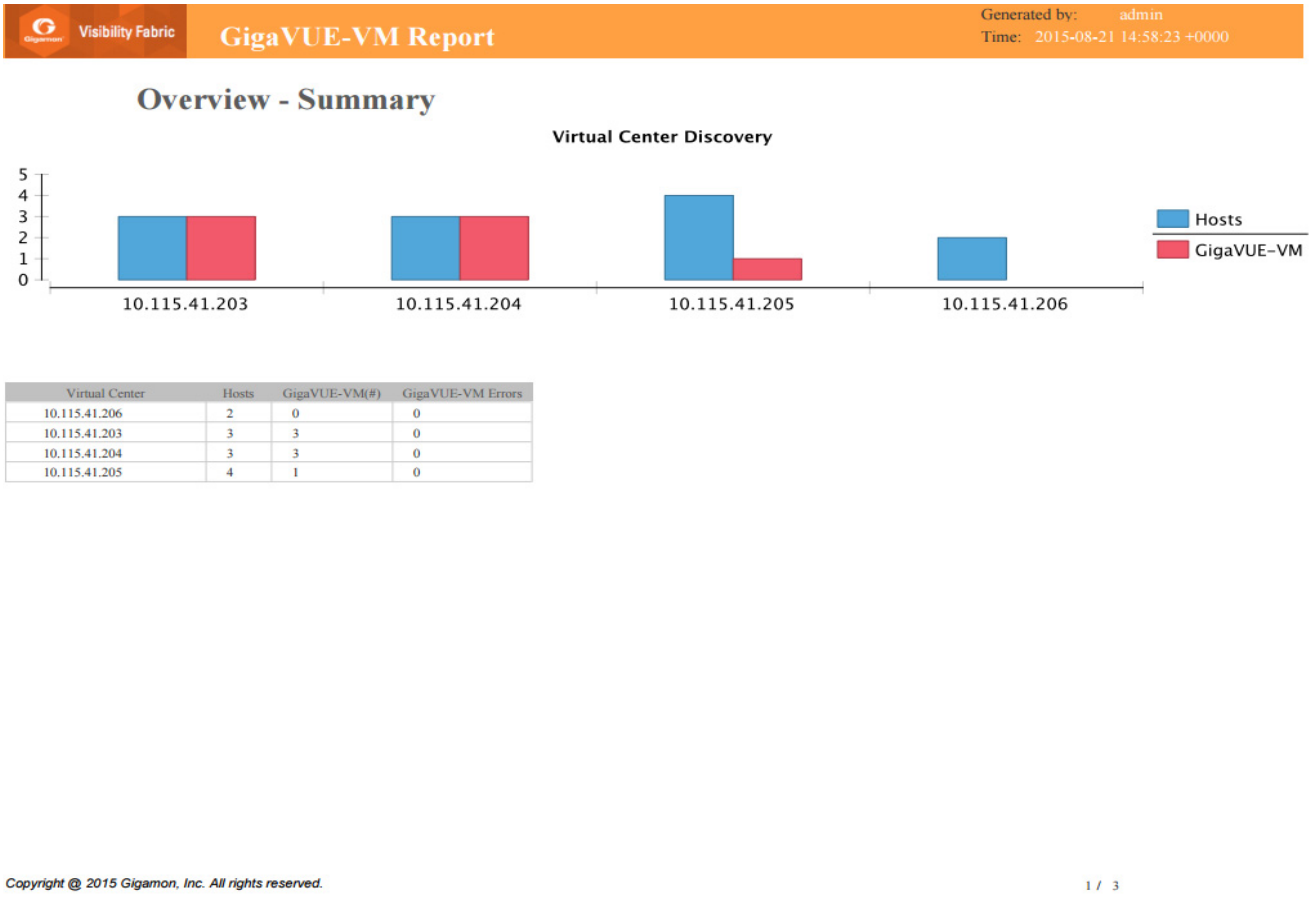

Figure 38-4: Report Pages Available for GigaVUE-VM

Template 4: Visibility Fabric Inventory Report

This multi-page template provides you with printable format of summary on all your Physical inventory (includes H Series, G Series and TA Series) that is visible on that GigaVUE-FM.

- Pie chart format for Status, Cluster Status, Node Types, Network and Tool Port Status and SW Versions.
- Table Format with all the Device IP with associated parameters such as Model, Status, Box ID, SW Version, Serial #, and so on.

Figure 38-5 show an example of a report for Visibility Fabric Inventory.

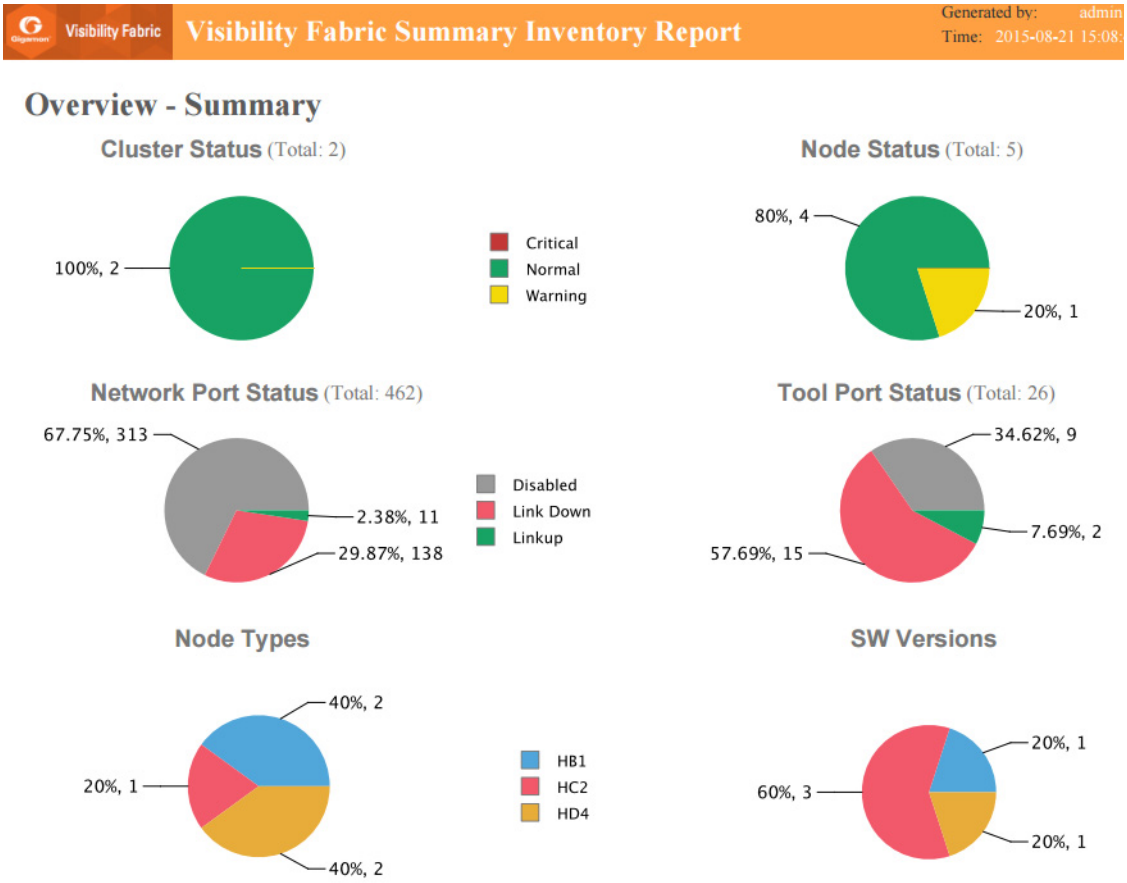

Figure 38-5: Inventory Summary Report

Template 5: GigaSMART Performance Report

This template provides you with printable format of summary on GigaSMART performance for all H Series nodes with GigaSMART functionality.

The report includes GigaSMART statistics for the following:

- Top/Bottom 10 GigaSMART (GS) Groups by Traffic: This information will indicate which GS groups are heavily utilized. To ensure to capture all the relevant information it is good to have the GS groups be description in the GS Groups alias names.
- Top/Bottom 10 GigaSMART (GS) Operations by Traffic: This information will indicate which GS operations are heavily utilized. To ensure to capture all the relevant information it is good to have the GS Operations be description in the GSOP alias names.
- Top/Bottom 10 GigaSMART (GS) Virtual Ports by Traffic: This information will indicate which virtual ports might be over-utilized and which are under-utilized.

Figure 38-6 show an example of a report for GigaSMART Performance.

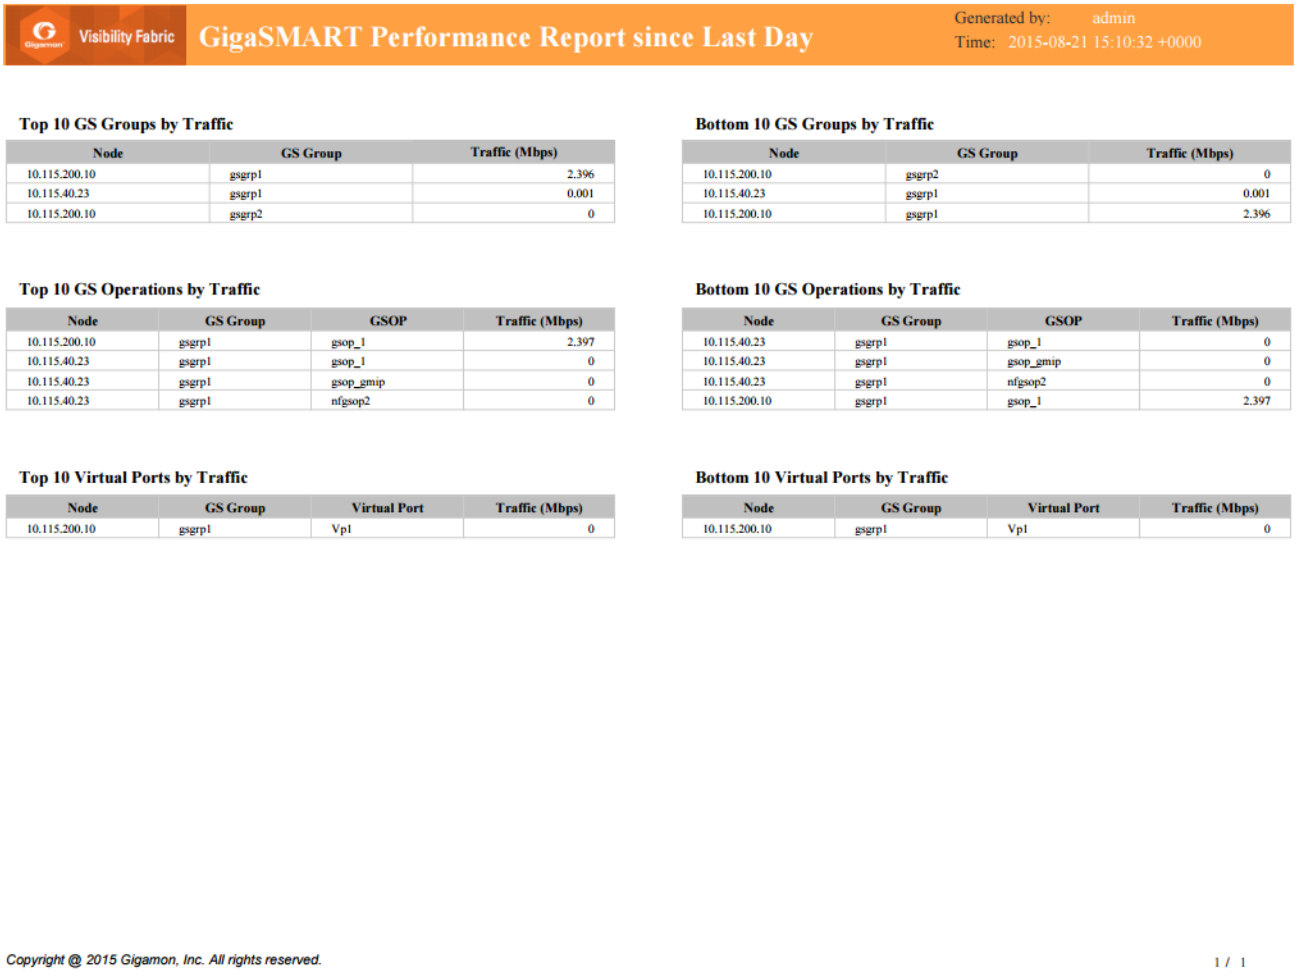

Visibility Fabric

GigaSMART Performance Report since Last Day

Generated by: admin
Time: 2015-08-21 15:10:32 +0000

Top 10 GS Groups by Traffic

Node	GS Group	Traffic (Mbps)
10.115.200.10	gsgrp1	2.396
10.115.40.23	gsgrp1	0.001
10.115.200.10	gsgrp2	0

Bottom 10 GS Groups by Traffic

Node	GS Group	Traffic (Mbps)
10.115.200.10	gsgrp2	0
10.115.40.23	gsgrp1	0.001
10.115.200.10	gsgrp1	2.396

Top 10 GS Operations by Traffic

Node	GS Group	GSOP	Traffic (Mbps)
10.115.200.10	gsgrp1	gsop_1	2.397
10.115.40.23	gsgrp1	gsop_1	0
10.115.40.23	gsgrp1	gsop_gmip	0
10.115.40.23	gsgrp1	nfgsop2	0

Bottom 10 GS Operations by Traffic

Node	GS Group	GSOP	Traffic (Mbps)
10.115.40.23	gsgrp1	gsop_1	0
10.115.40.23	gsgrp1	gsop_gmip	0
10.115.40.23	gsgrp1	nfgsop2	0
10.115.200.10	gsgrp1	gsop_1	2.397

Top 10 Virtual Ports by Traffic

Node	GS Group	Virtual Port	Traffic (Mbps)
10.115.200.10	gsgrp1	Vp1	0

Bottom 10 Virtual Ports by Traffic

Node	GS Group	Virtual Port	Traffic (Mbps)
10.115.200.10	gsgrp1	Vp1	0

1 / 1

Figure 38-6: Report for GigaSMART Performance Indicators

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 anextension 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=2f433d55532f4f3d446967694365727420496e632f4f553d
7777772e6469676963

6572742e636f6d2f434e3d446967694365727420534841322048696768204173737572
616e636

52053657276 6572204341
GigamonMdataSslCertSubCommonName=*.stickyadstv.com

GigamonMdataSslSub=2f433d55532f53543d4e657720596f726b2f4c3d4e657720596
f726b2f4f3d4672656

5776865656c204d6564696120496e632f4f553d46726565776865656c2f434e3d2a2e7
37469636b796164737 4762e636f6d dst=10.50.22.59 src=38.106.34.118
```

39 System

The **System** pages provides a variety of options allowing you to set up key features of GigaVUE-FM. These pages allow you to configure licenses for GigaVUE-FM and GigaVUE-VM activation, set up notifications for events and their email recipients, and view event logs.

To access the system pages, click **Administration** on the top navigation link. On the left navigation pane, click **System**.

System provides access to the following pages:

Preferences

The **Preferences** page displays the user profile and general settings for the current instance of GigaVUE-FM. Users with fm_admin and fm_super_admin role can only edit the Preferences.

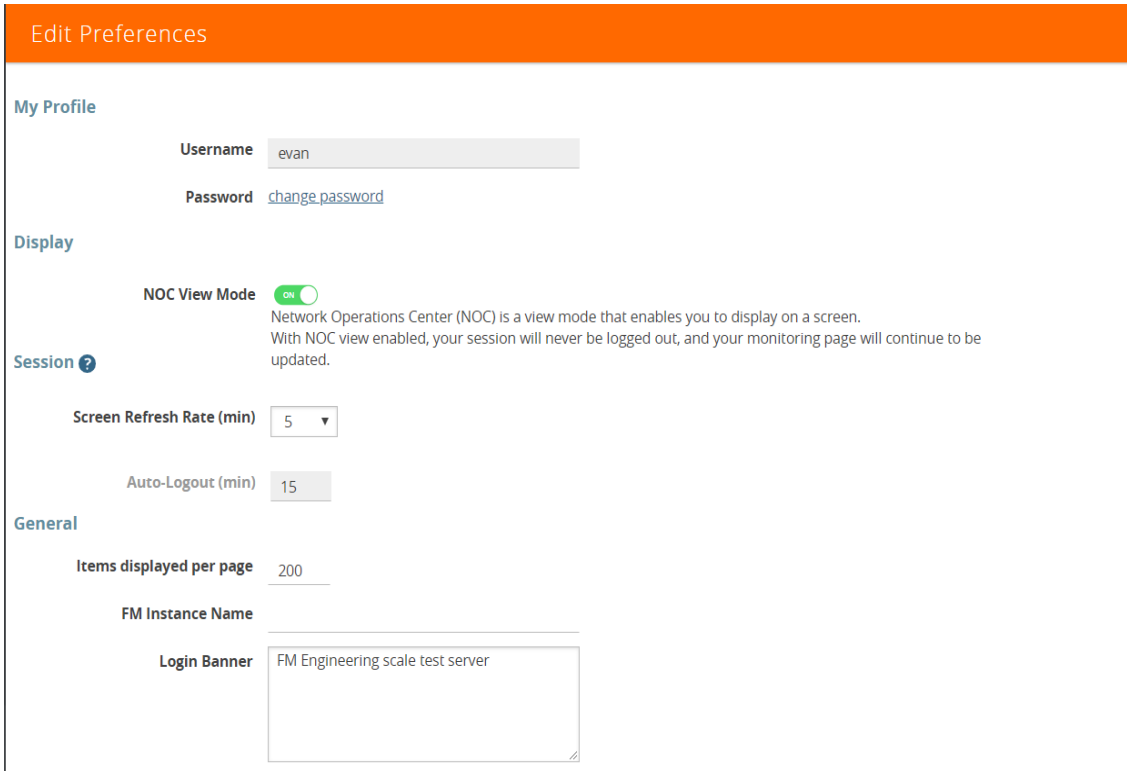

Figure 39-1: Preferences for GigaVUE-FM

To change the GigaVUE-FM preferences:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, go to **System** > **Preferences**.
3. Click **Edit**.
4. On the Edit Preferences page, perform any of the following:
   - Change the user name in the **Username** field.
   - Click the **change password** link to change the password. (For more information changing the password, refer to *Changing Your Password* on page 1233.)
   - Set the frequency of screen refresh from the Screen Refresh Rate (min) drop-down list. You can select from 0.5 to 5 minutes.
   - From the Auto-Logout (min) drop-down list, select the maximum duration GigaVUE-FM can be inactive before it is logged out automatically. By default, the auto-logout time is set to 30 minutes. You can select 15, 45, or 60 minutes.

   NOTE: If the **NOC View Mode** is set to on, then you cannot set the auto-logout time. Therefore, the session will never be logged-out and the screen gets refreshed continuously.

- Select the number of items to be displayed on a page by entering a value in the **Items displayed per page** field.
- Enter a name for the GigaVUE-FM instance in the **FM Instance Name** box. The GigaVUE-FM instance name is displayed in the browser tab as well as beside the GigaVUE-FM logo. Refer to *How to Add the GigaVUE-FM Instance Name* on page 49.
- Configure a pre-login banner which states the security policy of your company or organization in the Login Banner box. For more information about configuring a custom banner, refer to *Configure a Custom Banner* on page 44.

Thresholds

You can perform the following configurations from the Thresholds page:

- *Traffic Health Thresholds* on page 1299
- *SNMP Throttling* on page 1301

Traffic Health Thresholds

Using Traffic Health Thresholds, you can configure the packet error and packet drop threshold values for computing the health status of port types such as hybrid, network, stack, tool, inline tool, inline network, and gateway port. You can also configure the threshold values for GigaSMART engine port packet correlation, packet drops, and packet errors for computing the health status of GigaSMART engine port.

GigaVUE-FM checks the health status of the ports every 5 min. If the port packet correlation, errors, or drops increment every 5 min and exceeds the configured threshold for a specified time interval, then the port becomes unhealthy. When the port becomes unhealthy, the related maps also become unhealthy.

NOTE: When a new node or a cluster is added, GigaVUE-FM does not compute the traffic health immediately after the first configuration synchronizing cycle is completed. It takes the next synchronizing cycle to compute the traffic health based on the traffic health thresholds described in this section. While the traffic health is still being computed, the health status of a map is shown as gray (unknown state) in the **Physical Dashboards** > **Unhealthy Maps** widgets.

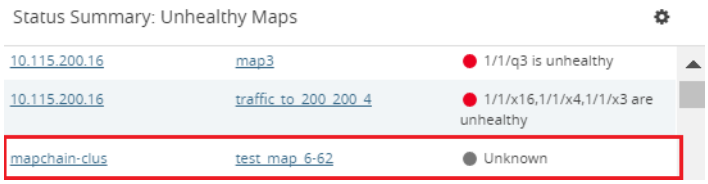

Figure 39-2: Unknown State of a Map

For more information about port health status, refer to *Port Health Status* on page 1362.

You must have fm_super_admin role to configure the traffic health threshold.

To set the traffic health threshold:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, click **System** > **Traffic Health Thresholds**. Refer to Figure 39-3 on page 1300.

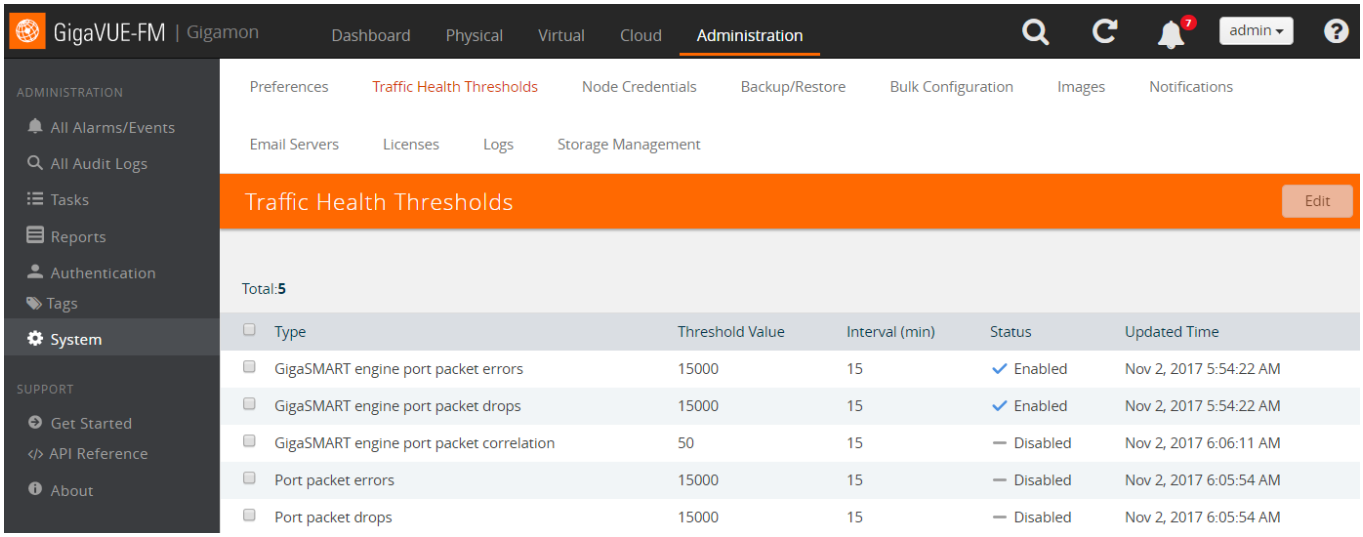

Figure 39-3: Traffic Health Thresholds

3. In the Traffic Health Thresholds page, select any of the following check boxes:

- **GigaSMART engine port packet errors:** The cumulative number of packet errors coming into a GigaSMART engine port.
- **GigaSMART engine port packet drops:** The cumulative number of packets dropped due to over subscription of a GigaSMART engine port.
- **GigaSMART engine port packet correlation:** The percentage (%) of packet correlation seen in a GigaSMART engine port. The GigaSMART engine port packet correlation is calculated based on the following factors:
  - the cumulative number of packets coming into a GigaSMART group
  - the cumulative number of packets going out of a GigaSMART interface
  - the cumulative number of packets dropped at a GigaSMART operation for a map
- **Port Packet Drop**: The cumulative number of packets dropped due to over subscription of a port.
- **Port Packet Error**: The cumulative number of packet errors coming into a port.

4. Click **Edit**. By default, all threshold types are enabled.

   To disable the thresholds, clear the **Enabled** check box.

5. In the **Threshold Value** box, enter the threshold value. The minimum and the maximum threshold value that can be entered are displayed in the threshold value box. Refer to Figure 39-4 on page 1301.

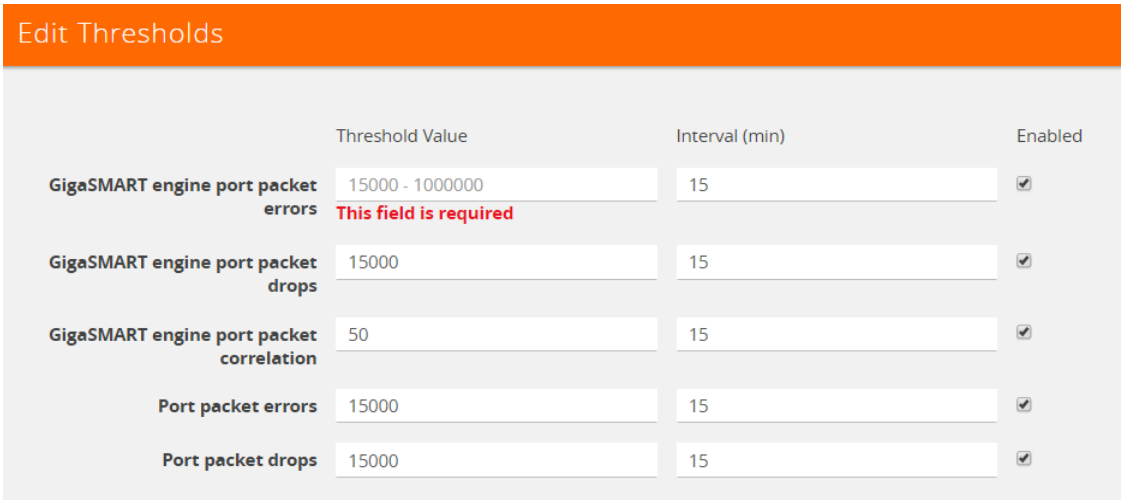

Figure 39-4: Edit Traffic Health Thresholds

6. In the Interval box, specify the time interval that the threshold value must exceed for the port to be considered unhealthy. By default, the time interval is set to 15 min.
7. Click **Save**. The threshold value and time interval are displayed as shown in Figure 39-3 on page 1300.

SNMP Throttling

Using SNMP Throttling, you can reduce the flooding of SNMP traps. You can manage the flooding by configuring the nodes with appropriate parameters for the trap events.

To configure SNMP Throttling:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, go to **System**, click **Thresholds** > **SNMP Throttling**.
3. The **SNMP Throttling** page is displayed as shown in the following figure:

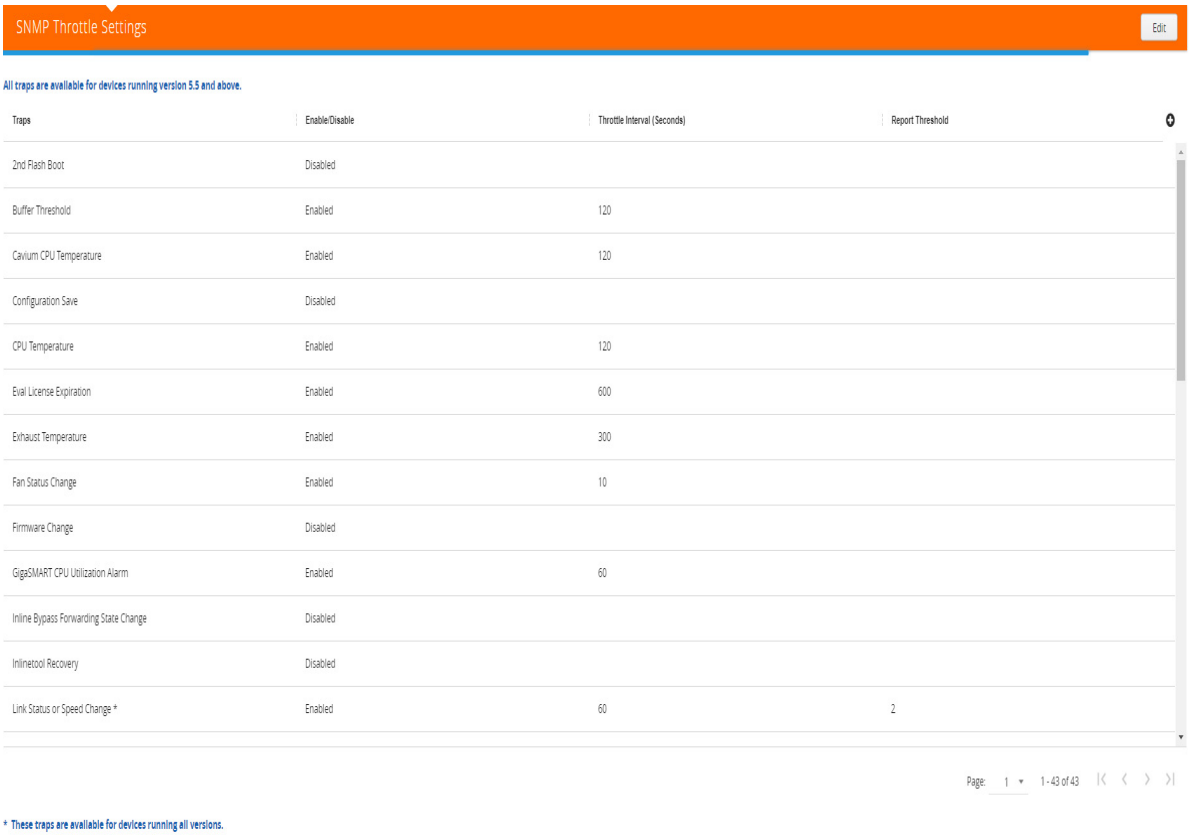

Figure 39-5: SNMP Throttle Settings Page

4. Click **Edit** to configure the following throttling settings for the traps:

- **Disable Throttle**: Allows you to disable the throttle for the required traps. If you select the **Disable Throttle** checkbox in the header, then throttling is disabled for all the traps.
- **Interval**: Allows you to configure the throttling interval. The throttling interval is configured by default for some of the traps (which is displayed in the page).
- **Report Threshold**: Allows you to configure the threshold limit for each of the traps based on which a throttle report trap is sent at the end of the interval. You can view the report in the Alarms and Events page.

5. Click **Save**.

NOTE: SNMP throttling is available for all traps for all devices running version 5.5 and above. For devices running earlier versions, SNMP throttling is available only for the following three traps:

- Link Status or Speed Change
- Packet Drop
- Packet Rx/Tx Error

Node Credentials

The **Node Credentials** page includes the credentials for every node managed by GigaVUE-FM. Figure 39-6 shows an example. For each node, you need to provide a user name and password that allow administrator privileges on the node.

To access **Node Credentials**, click **Administration** on the top navigation link. On the left navigation pane, go to **System** > **Node Credentials**.

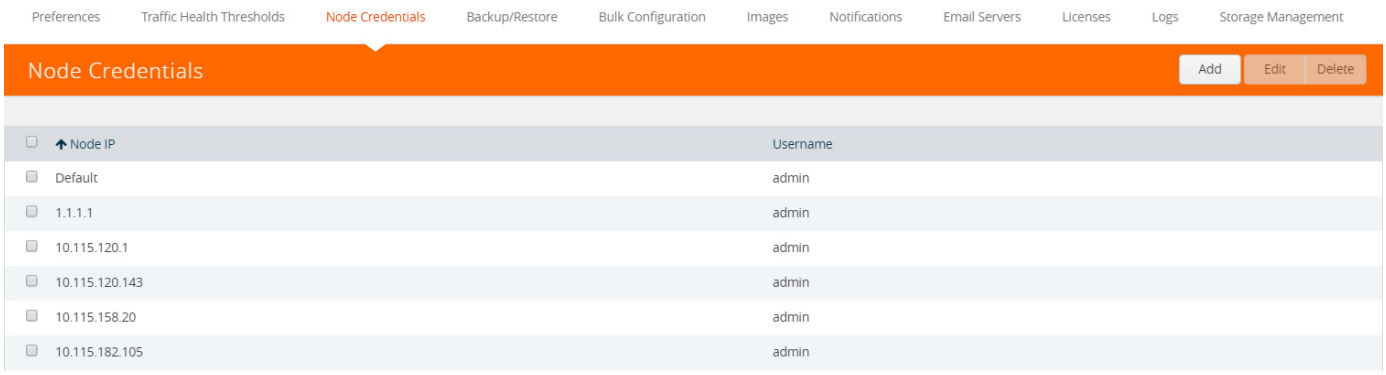

Figure 39-6: Add and Manage Node Credentials for Physical Nodes

NOTE: To ensure that users have the same privileges whether using GigaVUE-FM or H-VUE, it is advised that you use a centralized authentication server such as LDAP, instead of storing the username and password locally.

The list of node credentials is maintained in a local database and is accessed whenever GigaVUE-FM needs to connect to a node. Also, GigaVUE-FM will use the credentials in this page to log into any node added with the **Add** button in the **Physical Node** page.

Using the "Default" Credentials Effectively

The **Node Credentials** page includes both a **Default** entry as well as entries for specific IP addresses. The **Default** credentials make it easier to add multiple GigaVUE nodes that use the same username/password quickly. Instead of adding node-specific credentials for each system, you can just set the **Default** credentials to match the username/password in use on multiple nodes, and then add all the IP addresses that use those credentials in the same **Add Node(s)** dialog box.

Node Credentials Page Controls and Fields

Node Credentials table has following buttons that allow you to manage the information that appears in the table, **Add**, **Edit**, **and Delete**. To Edit or Delete a Node, click on the check box to the left of the IP address that needs to be modified.

Controls	Description
Add	Allows you to add a node and its login credentials. • Clicking **Add** opens a dialog where you specify the node IP address, a User name, and a Password. • Only one node can be added each time. **NOTE:** The user name and password you provide must have administrator privileges on the node.
Edit	Allows you to change the credentials for a node. • Select a node and click Edit to open a dialog where you make the changes. • Multiple IP addresses cannot be selected for editing.
Delete	Allows you to delete a node and its credentials. • The Delete Option will have a validation option to select as a pop-up prior to deleting a node. • Multiple IP addresses can be selected for deletion.

Backup/Restore

The Backup/Restore page allows you to backup and restore the configuration data for GigaVUE-FM, Physical Nodes, and add Archive Servers used for back up.

GigaVUE-FM Appliance

GigaVUE-FM includes a backup-and-restore feature for saving configuration data. You can use the saved data to restore an instance of GigaVUE-FM or provide a copy of the configuration data and have it available for a new instance of GigaVUE-FM. This is useful restoring the configuration on an appliance or when migrating to a GigaVUE-FM hardware appliance.

You can schedule GigaVUE-FM for an immediate backup or schedule a backup to occur once at a specified time or on a reoccurring basis. For example, you can schedule a backup for a particular day, week, month, or date at regular intervals.

Notes:

- After restore, you will need to reconfigure the RADIUS and TACACS+ passwords and regenerate the licenses.
- Backup and restore of GigaVUE-FM is only supported for users with the role fm_super_admin.

Data Saved When Backing Up GigaVUE-FM

When you back up GigaVUE-FM, the following information is saved:

- List of standalone nodes and clusters of directly under the management of the Fabric Manager.
- User credential needed a access the nodes
- Node level user account and RBAC configurations
- vMaps
- GigaVUE-FM credentials and preferences
- Other information, such as node level Radius, TACACS, SSH servers and SNMP or email notification configurations

The backup does not include the following data:

- GigaVUE-FM appliance host/IP configuration
- DHCP, NTP, and DNS configurations

These are configured through the jump-start configuration when configuration a new GigaVUE-FM.

Back Up Immediately

To do an immediate back up of a GigaVUE-FM, do the following:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **System > Backup/Restore > GigaVUE FM Appliance.**
3. Click **Backup**.
4. Select **Immediate**.
5. Select the archive server for the backup file. Refer to Figure 39-7.

   To add an archive server, refer to *Add Archive Server* on page 1310.
6. Click **OK**.

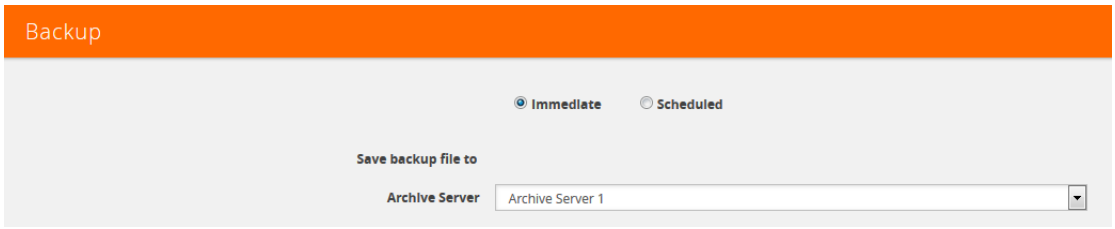

Figure 39-7: Immediate Backup to an Archive Server

Schedule Backups

To create a schedule for backing up GigaVUE-FM, do the following:

1. Click **Administration** on the top navigation link.

2. On the left navigation pane, select **System > Backup/Restore > GigaVUE FM Appliance.**
3. Click **Backup**.
4. Select **Scheduled**.

   The GigaVUE-FM time clock is based on ESX host time. Make sure that you have synchronized clock before any scheduling operation.
5. To repeat backups, use **Recurrences**, **Start Date**, and **Start Time** to set how often the backup occurs, at what time, and when the backup schedule will end.

   If you want to schedule a single backup for a specific data and time, select Once Only for **Recurrence**.

   Figure 39-8 shows an example scheduled backup. In this example, the weekly backup to archive server Archive Server 1 starts on March 17 and occurs every Saturday at 9:00 pm until March 31.

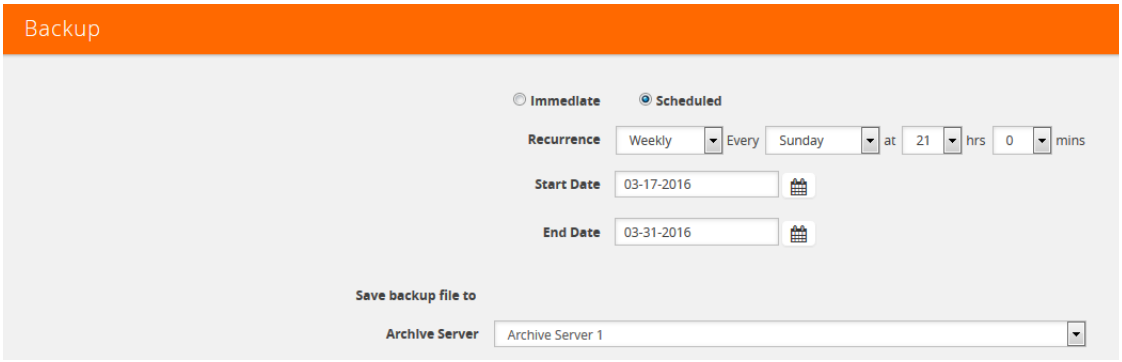

Figure 39-8: Scheduled Backup for GigaVUE-FM

6. Click **OK**. To monitor the progress of the event, select All Alarms/Events in the left navigation pane.

Once you have scheduled a recurring backup, the scheduled backup appears as a scheduled task on the Scheduled Tasks page. To view tasks, select **Tasks > Scheduled Tasks**. Figure 39-9 shows an example where a single backup task is scheduled. After the backup has completed the outcome of the task is displayed on the Alarm/Events page.

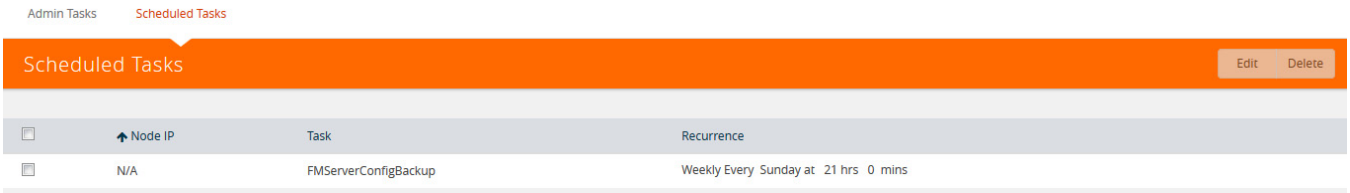

Figure 39-9: Scheduled Backup Task

Restore GigaVUE-FM Configuration Files

To restore a GigaVUE-FM configuration from a backup file, do the following:

1. Click **Administration** on the top navigation link.

2. On the left navigation pane, select **System > Backup/Restore > GigaVUE-FM Appliance.**
3. Click **Restore**.

   The Restore page displays, showing the file names from which to restore.
4. Select the Archive Server from which to retrieve the backup file.
5. Select the configuration to restore by clicking the check box next to the file name. Only one configuration can be selected with an restore action.
6. Click **OK**.

Physical Nodes

The **Physical Nodes** page lists the backup files currently saved in local storage on the machine where GigaVUE-FM is installed. You can also change the Do not Purge setting for the file and down load the files.

NOTE: You can backup multiple configuration files. The default is 10 per cluster. This file will be kept during automatic purge.

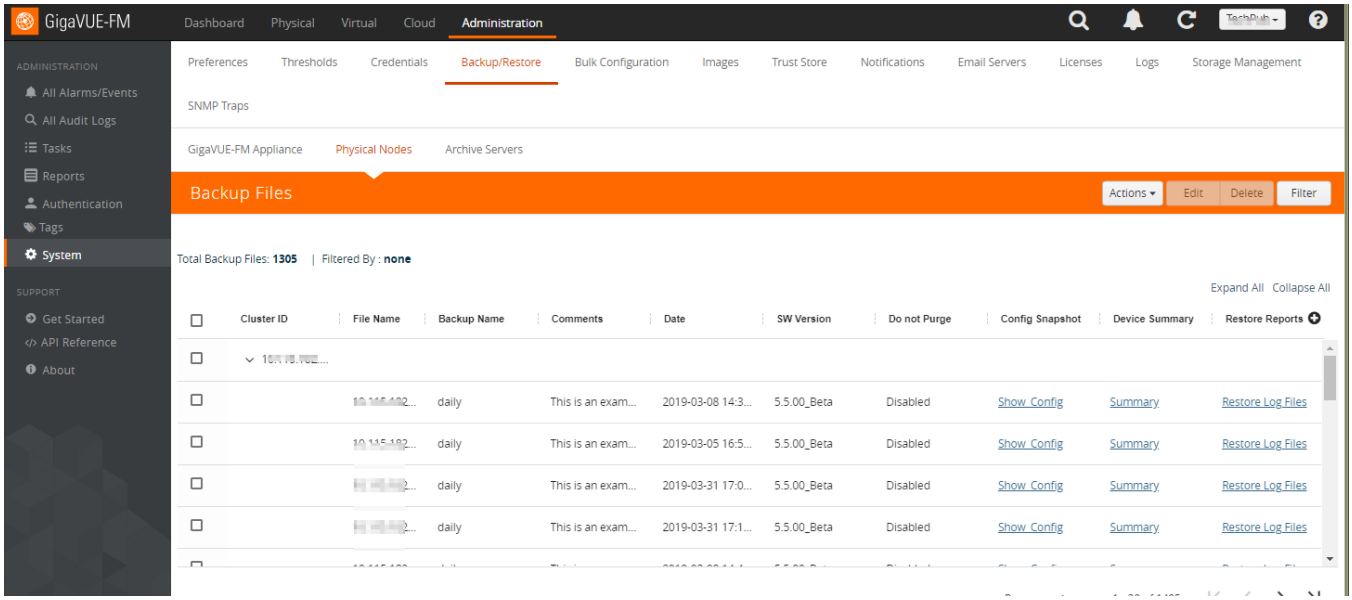

Figure 39-10: Backup Files Page

Enable Do Not Purge

To set Do Not Purge for a backup file, do the following:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **System > Backup/Restore > Physical Nodes**.
3. On the **Backup Files** page, select the backup file or files.
4. Click **Actions**, and then select **Enable Do Not Purge**.

   The Do Not Purge column will display a check mark for each backup file that has Do Not Purge enabled.

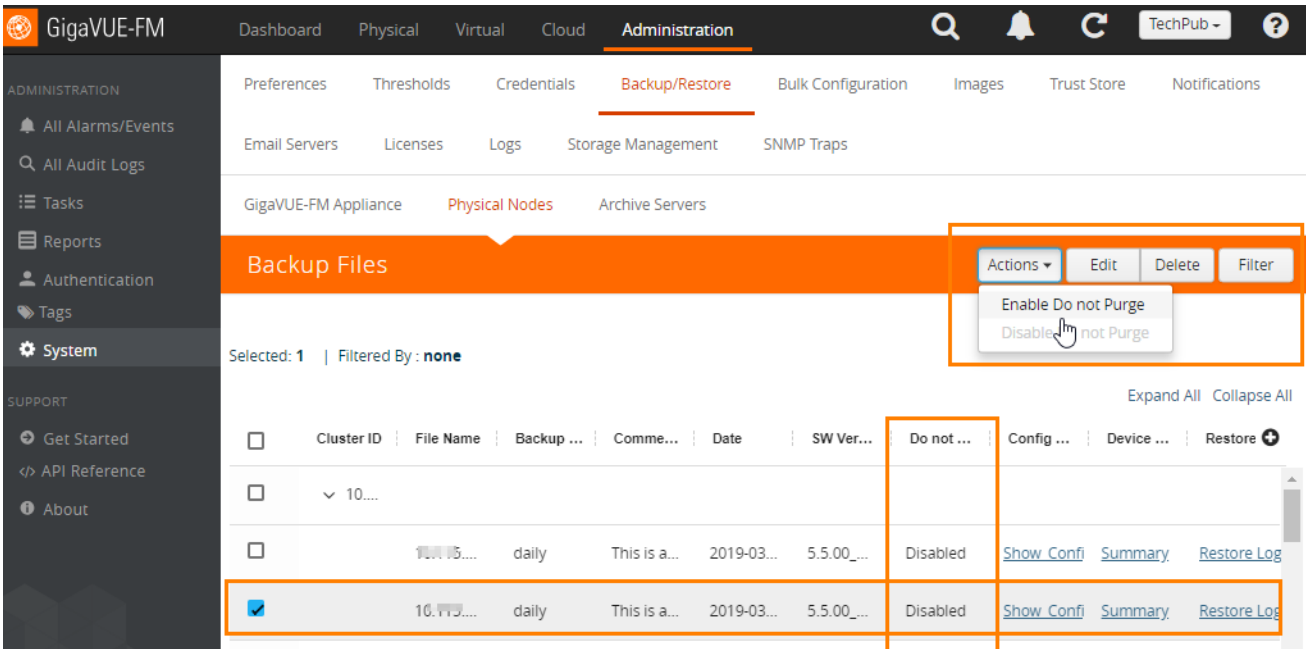

Figure 39-11: Do Not Purge Enabled

Disable Do Not Purge

To disable Do Not Purge for a backup file, do the following:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **System > Backup/Restore > Physical Nodes**.
3. On the **Backup Files** page, select the backup file or files.
4. Click **Actions**, and then select **Disable Do Not Purge**.

Download Backup Files.

You can also download the backup files by doing the following:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **System > Backup/Restore > Physical Nodes**.
3. On the **Backup Files** page, click **Show_Config** for the file to download the backup.
4. Click **Download**.

   GigaVUE-FM downloads the file. The filename includes the node's IP address and a timestamp.

Archive Servers

The Archive Servers page displays the archive servers currently available for backing up GigaVUE-FM. The page displays the following information:

- The alias to help identify the server
- The IP address of the server

- The type of server, either SCP or SFTP
- The username for logging in to the server
- The path on the server to the backup files

Add Archive Server

The Backup/Restore feature of GigaVUE-FM requires an archive server for saving and restoring the configuration files. To add an archiver server to GigaVUE-FM, do the following:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **System > Backup/Restore > Archive Servers**.
3. Click **Add**.
4. Enter the following information about the server:
   - Alias—An name to help identify the archive server.
   - Server Address—The IP address of the server.
   - Type—The type of archive server. The only type available is SCP.
   - File Path—The path to the backup files on the server
   - Username—The login user name for the server.
   - Password—The login password for the server.
5. Click **Save**.

Edit Archive Server

To make changes to an archive server, do the following:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **System > Backup/Restore > Archive Servers**.
3. Click **Edit**.
4. On the Edit Archive Servers page, make changes to the server information.
5. Click **Save**.

Delete Archive Server

To delete an archive serve, do the following:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **System > Backup/Restore > Archive Servers**.
3. On the Archive Servers page, select the server to delete.
4. Click **Delete**.

Device Configuration Backup

In GigaVUE-FM v5.500, GigaVUE-FM retrieves and stores the device configuration in binary and text formats. GigaVUE-FM restores the device configurations in binary format and allows you to view the configurations in text format. Previously the configuration backup was text based, which was intended to backup only traffic related configs. The advantage of the binary format is that it backs up all state parameters, including system parameters.

Version Compatibility:

- For GigaVUE devices 5.1 version onwards, GigaVUE-FM will take a binary backup of those devices. For devices prior to version 5.1, GigaVUE-FM will continue to take text-based backups.
- For backups taken in text format in GigaVUE-FM v5.4.01 or above, GigaVUE-FM will allow those configurations to be restored.

Points to Remember

- **Describe your backup**: Provide a meaningful name and comments while taking the backup, to help track the configuration while restoring.
- **View your backup**: You can view and download the text format of binary contents for readability. Some details in the binary content might not appear in text format.
- **View the restore report**: After performing a restore operation, a restore report displays the results of the restore operation such as the success or failure as well as all the logs from the master device.

Restore Configuration (RMA'ed device)

The following are the steps to be performed for Standalone and Cluster Nodes:

Standalone Devices (RMA):

1. Make a backup on Device A, which will contain Device A details (Serial Number, chassis ID, GUID, etc.)
2. Device A has now failed.
3. Order and get a new device, Device B (identical hardware inventory, except all the hardware serial numbers are different).
4. Power up Device B and assign it a hostname and IP address.
5. The hostname should be the same as the previous device.
6. The IP address should be the same. (Different IP addresses are supported, but not recommended.)
7. On GigaVUE-FM, restore the backup of Device A onto this device:
    a. If the IP address is the same, then it will get discovered as Device A.
    b. If the IP address is not the same, restore the Device A data onto the Device B (IP address, etc.)
    c. Now the backup taken on device A is pushed to device B.
    d. When the backup is complete, GigaVUE-FM invokes a new “Migration” API from the node.
    e. When the process is complete, Device B is restored to Device A’s configuration.

A node of the cluster (RMA):

1. The binary backup of the cluster is available in GigaVUE-FM.
2. Node A fails.
3. Replace with node B.

4. Configure the node B with the IP address, hostname, cluster ID, cluster VIP etc.
5. Node B joins into the cluster.
6. Cluster master will push the config to the new node, which will not apply to its hardware since its serial number does not match.
7. GigaVUE-FM will now discover node B back in the inventory.
8. Instruct GigaVUE-FM to migrate the configuration of node B, from the old serial number to the new one.
9. This will be sent to the Master of the cluster (and new API that will be provided same as 7 above).
10. Cluster master will do the migration and push the configuration to the new node

NOTE:

- GigaVUE-FM handles the UUID stored in GigaVUE-FM.
- GigaVUE-FM has a dependency on the device API to migrate configuration of RMA box to a new serial number.

Bulk Configuration

The Bulk Configuration page allows you upload and send a configuration file to one or more G Series nodes or clusters at the same time, replicating the configuration on each node or cluster. Bulk Configuration is not supported on H Series nodes.

The configuration file is a text-based file. This means that you can create a custom configuration file and upload it, or you can make a backup of a node and then edit the backup file to create a new configuration.

Bulk configuration is only supported on G Series models GV2404 and GV420. GV212 and GV216 are not supported. If unsupported device models are in a G Series stack, the entire stack is disregarded for configuration.

Important: GigaVUE-FM does not validate the configuration file before pushing it to the specified node or nodes during bulk configuration. If any errors occur, they are logged in the configuration log files.

Replicate Configuration Files

Use the following steps to replicate a configuration across nodes and clusters. If you are creating a new configuration file for bulk configuration, go directly to step 3.

1. Create a backup file.
   a. Click Physical on the top navigation link.
   b. On the Physical Nodes page, select a node.
   c. Go to **Actions > Backup**.
   d. On the Backup page, select **Immediate**.
   e. Click **OK**.
2. Download the backup file created in Step 1.
   a. Click **Administration** on the top navigation link.
   a. Select **Backup/Restore > Physical Nodes**.
   b. Select the backup file of the node you want to replicate.
   c. Select **Actions > Download**.
   d. Select **Immediate**.
   e. Click **OK** and then save the file.

      GigaVUE-FM downloads the configuration as a text file.
3. Open the configuration file in a text editor edit to make any needed changes to the configuration.

   The configuration file is expected to have header information that is based on the device type. If you are creating a configuration file from scratch, you need to provide the correct header. Table 39-1 on page 1315 provides the headers for each device type that is supported. In the header, version is the software version and file is the filename of the device image.

4. Upload the configuration file:
   a. Select **System > Bulk Configuration**.

      The Bulk Configuration Files page displays. An example is shown in the following figure.

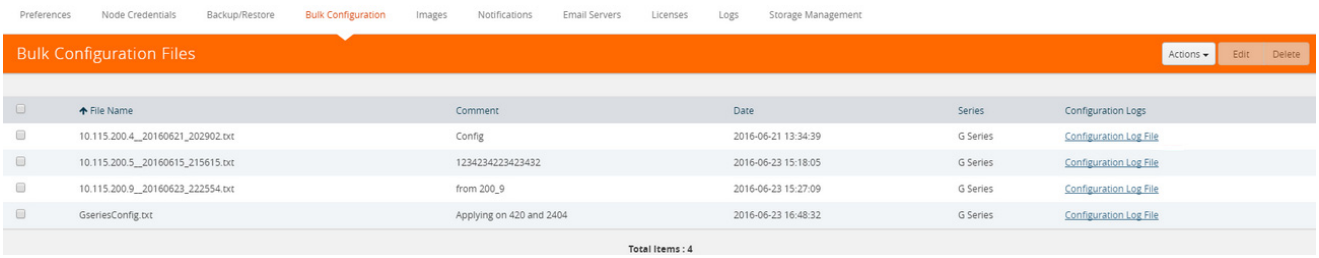

   b. Select **Actions > Upload**.

      The Upload Configuration File page displays. The page is shown in the following figure.

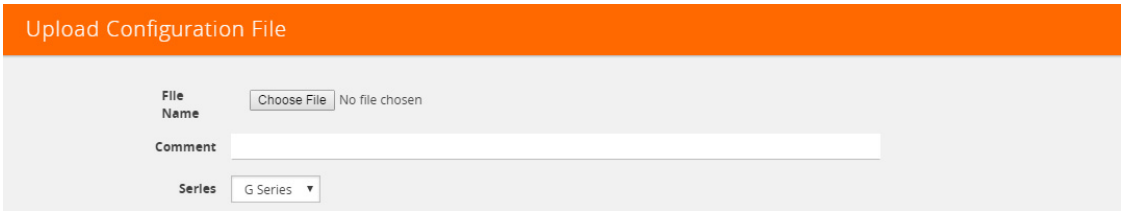

   c. Click **Choose File** to upload the file downloaded and edited in Step 2.
   d. (Optional) Enter a comment about the file in the **Comment** field.
   e. For **Series**, select G Series. (Only G Series nodes are supported in the current release.)
   f. Click **OK**.

      The uploaded file appears on the Configuration File page.

5. Replicate the file on the node or cluster.
   a. On the Bulk Configuration page, select the file uploaded in Step 4.
   b. Select **Actions > Replicate**.

      The Replicate Configuration File page displays. The page shows the selected configuration file, comment entered on the Upload Configuration File page, and the list of nodes that you can select for replication. Figure 39-12 shows an example.

      When **Autosave Backup Configuration** is selected, GigaVUE-FM takes a backup prior to applying the configuration changes. Configuration changes are not be applied if backup fails.

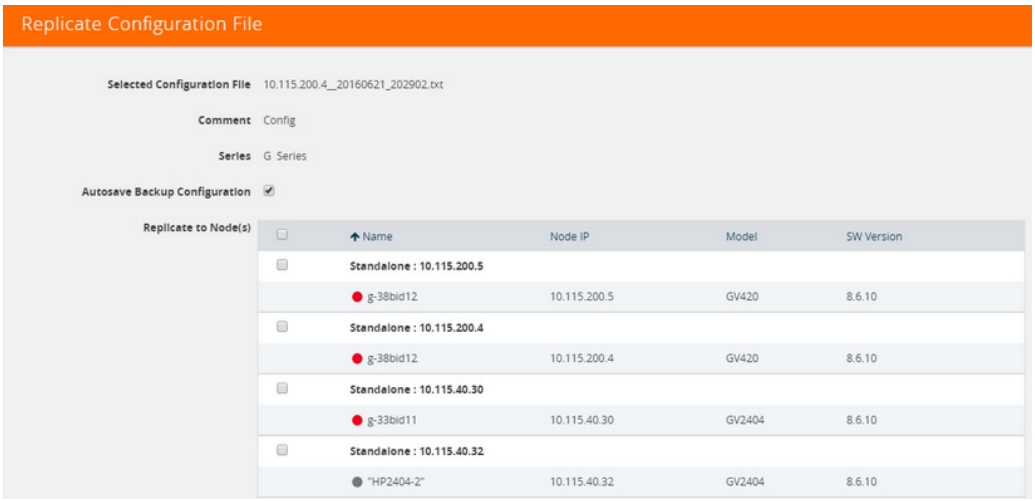

Figure 39-12: Replicate Configuration File Page

c. Select the nodes to which you want to replicate the configuration file

d. Click **OK**.

6. To view the progress of the configuration, select **All Alarms/Event** in the left navigation pane.

Table 39-1: Headers for G Series Configuration Files

Device	Header
GigaVUE-420	#========== #Platform: GigaVUE-420 #Software version/file: 8.6.10/gvb86.01_07 #==========
GigaVUE-2404	#========== #Platform: GigaVUE-2404 #Software version/file: 8.6.10/gvc86.11_04 #==========

View Configuration Log Files

When a configuration file is applied to a physical node, the node returns response messages that are recorded in a log file. These log files are useful for identifying any errors if the configuration fails. The log file is a text field that contains the list of CLI commands applied in during the configuration and the results.

The log file for a configuration file applied to a node has the following format:

```
<config-filename>_<device-ip>_<date>_<time>.txt
```

For example, if the configuration log file is named GseriesConfig__10.10.10.10_20160621_203610.txt, the filename is interpreted as follows:

- Configuration filename: GseriesConfig.txt
- Applied to device IP: 10.10.10.10
- Date Applied: 20160521 (May 21, 2016)
- Time applied (FM server time): 203610

To view a log for a configuration file, do the following:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **System > Bulk Configuration**.

   The Bulk Configuration Files page displays. An example is shown in the following figure.

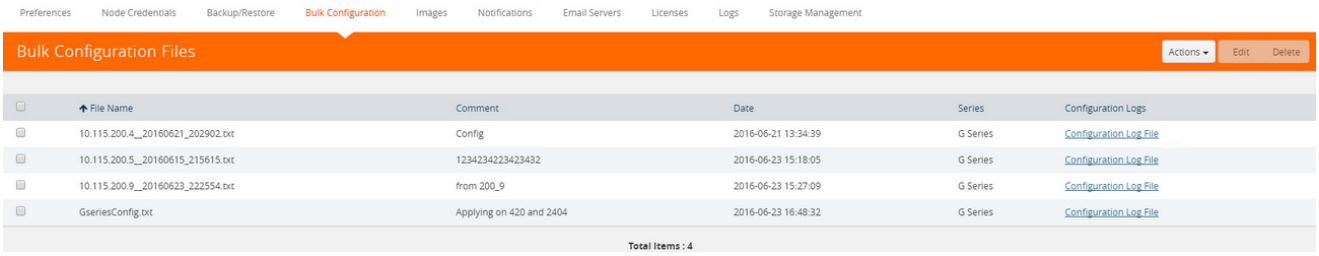

3. Under **Configuration Logs**, click the Configuration Log File link for the configuration file log you want to view.

   The Configuration Logs page displays.
4. Select the configuration file, and then click **Download**. In the following figure, the file selected for download is 10.115.200.4__2010621_202902_10.115.200.4_20160521_203610.txt.

5. Open the downloaded configuration file in a text editor to review the contents.

Images

The Images page is used to specify the servers where you will store image files for upgrading your nodes. You obtain images for your nodes by contacting Technical Support. Once you have the images, you can use an external server or use GigaVUE-FM as the image server.

To access Images, click **Administration** on the top navigation link. On the left navigation pane, select **System > Images**.

Internal Image Files

If you use GigaVUE-FM for the image files, the files used to upgrade the physical nodes to the latest software version are stored on your local system and uploaded to GigaVUE-FM from the Upload Internal Image Files page. To access this page, go to **System > Images > Internal Image Files**.

After obtaining the image files, copy them to your local system. Use the **Browse** button to upload the files. Figure 39-13 shows an image file for a Gigamon-HC2 node selected for uploading. To upload the file, click **OK**.

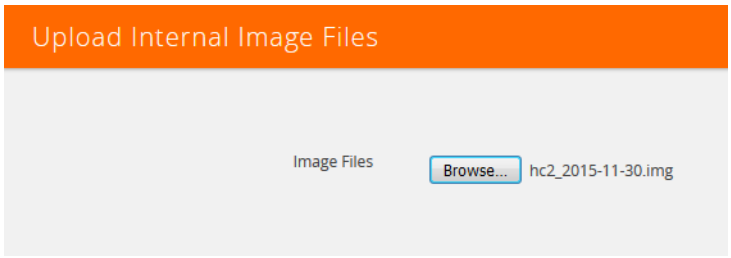

Figure 39-13: Image File Uploaded

After the uploading has completed, the image file is shown on the Internal Image Files page as shown in Figure 39-14. Use the **Download** button to download images stored on GigaVUE-FM to your local system. Use the **Delete** button to remove image files.

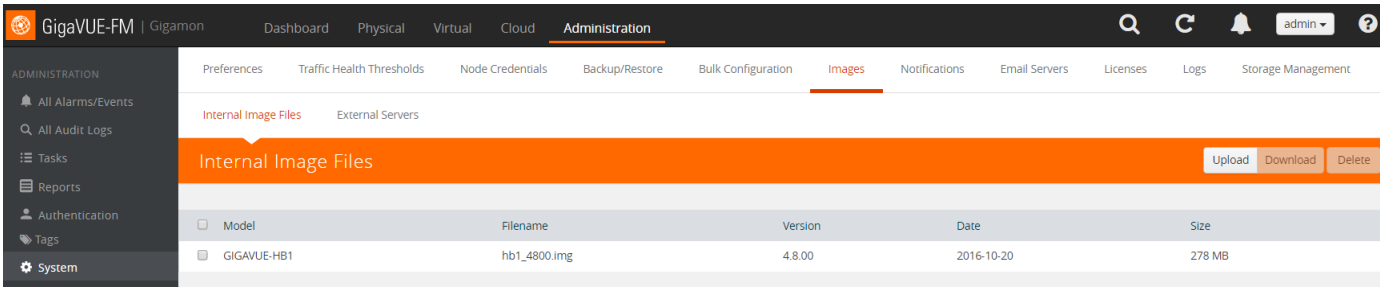

Figure 39-14: Internal Image Files Page

External Servers

If you use an external server for the image files, the files used to upgrade the physical nodes or GigaVUE-FM to the latest software version are stored on an external Image servers. To access the External Servers page, go to **System > Images > External Servers**. The External Servers page has buttons used to set up and manage external image servers. These buttons are described in Table 39-2. For information on how to upgrade from an external server, refer to *Upgrade from an External Image Server* on page 230.

Table 39-2: Controls on External Servers Page

Controls	Description
Add	Allows you to specify where server images will be stored. The page is shown in Figure 39-15. Clicking **Add** opens the Image Server Details dialog, where you specify: • **Alias** — A name to identify the server. • **Server Address** — The IP address of the server. • **Base Image Directory** — The base path where image files are stored. Images can be placed in subdirectories of this base directory. **NOTE:** Images can be updated using SCP, FTP or TFTP. • **Username and Password** —The user name and password that will be used to log into the server to store the image file. **NOTE:** A username and password are not required if using TFTP or SCP.
Edit	Select a server and click **Edit** to open the Image Server Details dialog, where you can modify the values specified for the server. • Same options are to be filled as noted for **Add**.
Delete	Select a server and click **Delete** to delete the server specified. • The Delete Option will have a validation option to select as a pop-up prior to deleting a node. • Multiple IP addresses can be selected for deletion.

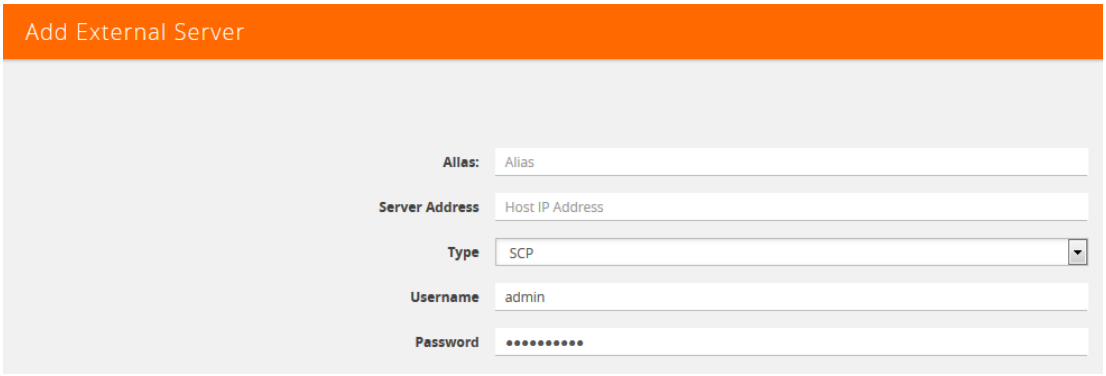

Figure 39-15: Add External Server Page

Trust Store

The SSL Certificate Enhancement feature in GigaVUE-FM ensures secure communication between GigaVUE-FM and the devices added to GigaVUE-FM. The Trust Store page in GigaVUE-FM enables security by maintaining a list of certificates provided by the devices. To add new devices to GigaVUE-FM and to manage the existing devices, you must add the root CA certificate of the respective devices to the Trust Store.

The Trust Store page lets you toggle between enabling and disabling security:

- If you enable security, GigaVUE-FM performs the following:
  - Verifies if the root CA certificate of the device is available in GigaVUE-FM.
  - Adds the device only if the certificate is signed by an authorized CA.
  - Verifies the chain of custom certificates, as required.
- If you disable security, GigaVUE-FM adds the devices without any validation.

IMPORTANT RECOMMENDATION: Prior to adding the certificates to the Trust Store, you must ensure to do the following:

- Login to the devices and add the private key and certificate of the devices through CLI/Console into each of the devices.
- Login to GigaVUE-FM and add the private key and certificate of GigaVUE-FM through CLI/Console (into GigaVUE-FM).

Use the `crypto` CLI command for adding the keys and certificates. Refer to the GigaVUE-OS CLI User's Guide for detailed information.

To access the Trust Store Page, click **Administration** on the top navigation link. On the left navigation pane, select **System** > **Trust Store**.

To add a certificate to GigaVUE-FM:

1. Click **Add** on the Trust Store page. The Add Certificate page appears.
2. Enter an **Alias** for the certificate.
3. Click **Choose File** to upload the certificate.
4. Click **OK**.

The certificate is added to the list view.

Notifications

GigaVUE-FM provides powerful email notification capabilities, automatically sending emails to specified addresses when any of a wide variety of events take place on the node. Gigamon strongly recommends that you configure this feature so you have immediate visibility of events affecting node health.

To configure automatic email notification, you will need to configure the email notification settings, the events about which to be notified, and the recipient or recipients for the notifications.

To access the Notifications page, click **Administration** on the top navigation link. On the left navigation pane, select **System > Notifications**.

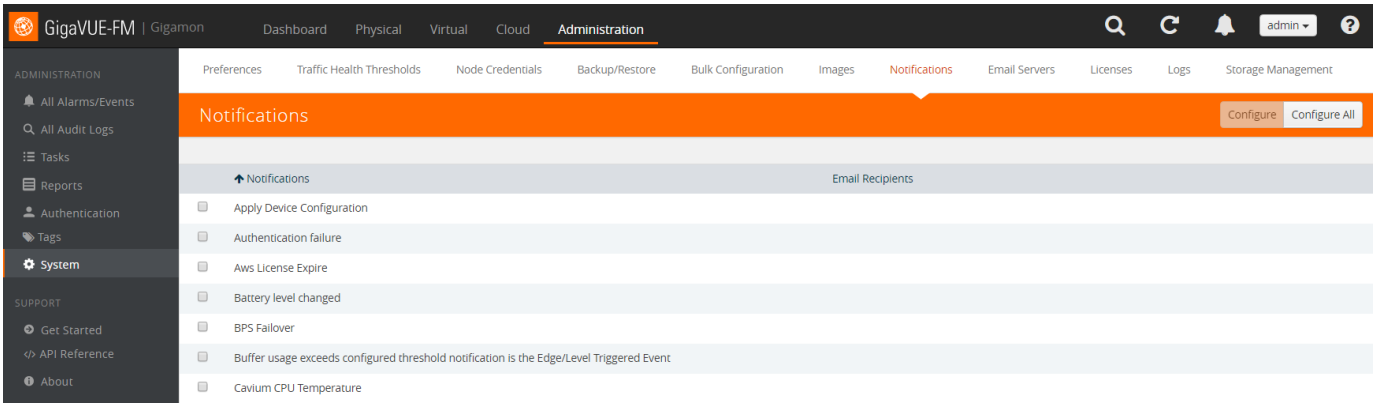

Figure 39-16: Notifications Page to Set Email Recipients for Alarms/Events

Some of these events are detected by GigaVUE H Series, TA Series, and G Series nodes, and the notifications are forwarded to the Fabric Manager. For a node to be able to send notifications to the Fabric Manager, the node's SNMP notifications must be configured with the Fabric Manager's IP address. For information about adding a destination for SNMP notifications, refer to Configuring SNMP Notifications in GigaVUE-OS CLI User's Guide.

Configure Email Notifications

To configure the automatic email notifications:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **System > Notifications**.
3. Do one of the following:

- To set the automatic email notification to all the events, click **Configure All**. The Configure all notifications page is displayed. Enter the email addresses of the recipient or recipients in the Recipient(s) field. Click **Save**.

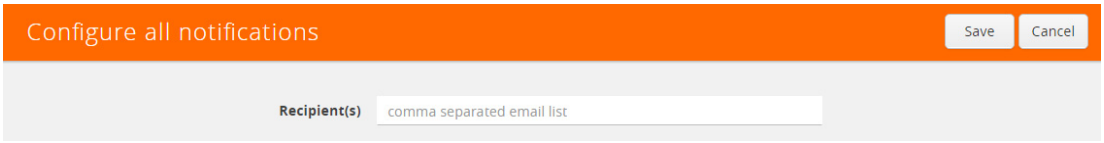

Figure 39-17: Configure all notifications

- To set the automatic email notification to a single event, select the Enter the email addresses of the recipient or recipients in the Recipient(s) field.

The Notifications page has two buttons that allow you to modify your recipients lists that appears in the table: **Configure** and **Configure All**.

Controls	Description
Configure	Allows you to add recipients to a single notification. Select a notification from the Description list and then click **Configure** to open the Configure Notification page, where you specify the recipient or recipients in the **Recipient(s)** field for that notification. After specifying the recipient or recipients you can send a test email by clicking **Send Test Email**.
Configure All	Allows you to add recipients to all notifications. Click **Configure All** to open the Configure all notifications page, where you enter the email addresses for the recipient or recipients in the **Recipient(s)** field that you want added to all notifications. After specifying the recipient or recipients you can send a test email by clicking **Send Test Email**.
Notifications	The list of notifications that can be emailed to recipients as alerts. Table 39-3 provides more information about these notifications. You can toggle notifications to display in descending or ascending order by clicking the label.
Email Recipients	Email addresses of the people to be notified. Each email address should be separated by a comma. You can toggle the order of the email recipients to display in descending or ascending order by clicking the label.

The following table describes all GigaVUE-FM notifications.

Table 39-3: GigaVUE-FM Notifications

Notifications	Description	Node Series
Apply Device Configuration	The configuration was applied to a managed node.	H Series
Authentication failure	A user login attempt failed on the indicated GigaVUE G Series node.	G Series
AWS License Expire	The AWS license is close to expiry.	
Battery level changed	The battery charge in a G-TAP A Series tap changed. Traps are generated at 25% increments as the available battery charge falls - 75%, 50%, and 25%. Traps are also generated when the available battery charge falls to 15% and the system closes the tap relays, falling back to passive mode.	A Series

Table 39-3: GigaVUE-FM Notifications

Notifications	Description	Node Series
BPS Failover	An inline bypass failover has occurred.	H Series
Buffer usage exceeds configured threshold notification is the Edge/Level Triggered Event	Buffer threshold can be set by going to **Node IP > Ports > Ports > Configure**. When using the drop down on Configure, you can set the buffer threshold for each port.	H Series
GigaSMART CPU Temperature	GigaSMART CPU temperature is above an acceptable level	H_Series
Cluster added	A cluster is added to the GigaVUE-FM instance.	H Series TA Series
Cluster creation failed	The cluster failed to form successfully.	H Series TA Series
Cluster creation started	The cluster creation has started.	H Series TA Series
Cluster image install finished	Installation of a new image on the indicated cluster of nodes began.	H Series TA Series
Cluster image install started	Installation of a new image on the indicated cluster of nodes began.	H Series TA Series
Cluster reboot finished	A reboot of the specified cluster completed.	H Series TA Series
Cluster reboot started	A reboot of the specified cluster was initiated.	H Series TA Series
Cluster Removed	The indicated cluster of nodes was removed from GigaVUE-FM.	H Series TA Series
Cluster Updated	The indicated cluster of nodes was updated in GigaVUE-FM.	H Series TA Series
Configuration saved	The configuration of a node was saved to local storage (for example, by using the **write memory** command).	Any
CPU Temperature	The CPU temperature is above the threshold limit.	
CPU utilization is high	CPU utilization on the indicated node exceeded a hard-coded threshold.	Any
Device config backup	GigaVUE-FM performed a configuration backup for the indicated node(s).	GigaVUE-FM
Device config deleted	GigaVUE-FM deleted a backed-up configuration file for the indicated node(s).	GigaVUE-FM
Device config restore	GigaVUE-FM restored a configuration file to the indicated node(s).	GigaVUE-FM
Device Health changed	The health status of the device is changed based on the health status of ports, cards, fan tray, power module, memory utilization, and CPU utilization.	H Series TA Series
Device image install failed	Installation of a new image on the indicated node failed.	Any
Device image install finished	Installation of a new image on the indicated node completed at the indicated time.	Any

Table 39-3: GigaVUE-FM Notifications

Notifications	Description	Node Series
Device image install started	Installation of a new image on the specified node began at the indicated time.	Any
Device Operational mode changed	The cluster or standalone node is in Safe or Limited mode.	H Series
Device reboot finished	Reboot of the specified node began at the indicated time.	Any
Device reboot started	Reboot of the specified node finished at the indicated time.	Any
Disk space low	The available disk space on the indicated node fell below a hard-coded threshold.	Any
Evaluation License Expire	The evaluation license for GigaVUE-FM has expired.	GigaVUE-FM
Exhaust Temperature	Exhaust temperature is above the acceptable level.	H Series
Fabric Node Down	The GigaVUE V Series node is down	
Fabric Node Reboot Failed	The GigaVUE V Series node has failed to reboot.	
Fabric Node Rebooted	The GigaVUE V Series node is rebooted.	
Fabric Node Replacement Launch Failed	The GigaVUE V Series node upgrade failed to launch the new version.	
Fabric Node Replacement Launched	The new version of GigaVUE V Series node is launched and the old version is removed.	
Fabric Node Restart Failed	The GigaVUE V Series node failed to restart.	
Fabric Node Restarted	The GigaVUE V Series node restarted.	
Fabric Node Unreachable	The GigaVUE V Series node is unreachable.	
Fabric Node Up	The GigaVUE V Series node is up.	
Fan tray changed	The Fan Tray in GigaVUE H Series node was removed and reinserted.	H Series
Firmware changed	The system booted and detected that its firmware has been updated from the previous boot.	G Series
FM image install Finished	Installation of the image file for GigaVUE-FM completed installation	GigaVUE-FM
FM Image Upgrade Completed	GigaVUE-FM was successfully upgraded from the installed image.	GigaVUE-FM
FM Image Upgrade Failed	Upgrade of GigaVUE-FM from the installed image failed.	GigaVUE-FM
FM Image Upgrade Started	An upgrade of GigaVUE-FM has started.	GigaVUE-FM
FM Server config backup	Backup of GigaVUE-FM has completed.	GigaVUE-FM
Gigamon Discovery	An email notification is sent to all configured destinations each time a new Gigamon discovery neighbor is discovered or Gigamon discovery information for an existing neighbor is changed or expired.	
GigaSMART Application Core Crash	A GigaSMART application core crash occurs due to a back trace trigger or a soft reset being initiated.	H_Series
GigaSMART CPU Utilization	GigaSMART CPU utilization is above the rising threshold.	H_Series
GigaSMART Packet Drop	A packet drop was detected by GigaSMART.	H_Series

Table 39-3: GigaVUE-FM Notifications

Notifications	Description	Node Series
GigaVUE-VM came online	The indicated GigaVUE-VM node came online and was detected by GigaVUE-FM.	GigaVUE-VM
GigaVUE-VM for Datacenter install completed	A bulk deploy of GigaVUE-VM nodes from GigaVUE-FM completed.	GigaVUE-VM
GigaVUE-VM for Datacenter install interrupted	A bulk deploy of GigaVUE-VM nodes was interrupted before the installation completed.	GigaVUE-VM
GigaVUE-VM for Datacenter install started	A bulk deploy of GigaVUE-VM nodes from GigaVUE-FM began.	GigaVUE-VM
GigaVUE-VM install completed	Installation of the indicated GigaVUE-VM node completed.	GigaVUE-VM
GigaVUE-VM install started	Installation of the indicated GigaVUE-VM node began.	GigaVUE-VM
GigaVUE-VM pinned to host	GigaVUE-VM was pinned to the host.	GigaVUE-VM
GigaVUE-VM unpinned from host	GigaVUE-VM is no longer pinned to the host.	GigaVUE-VM
Inline Bypass State Change	Forwarding state changed on an inline network.	H Series
Inline Tool Recovery	An inline tool recovered from failover.	H Series
Link state changed	An GigaVUE H Series node detected that either: • A port's link status has changed from up to down or vice-versa. • A port's speed has changed. **NOTE:** This trap is **not** sent when the Management port's link status changes. **NOTE:** The link state polling interval is 1 second. If a link state change is detected during the poll, a trap is generated.	H Series
Manual link added	A link was added to the Topology manually.	GigaVUE-FM
Manual link removed	A manually added link was removed from the Topology.	GigaVUE-FM
Manual link updated	A manually added link in the Topology was changed.	GigaVUE-FM
Manual node added	A node was added to the Topology manually.	GigaVUE-FM
Manual node removed	A manually added node was removed from the Topology.	GigaVUE-FM
Manual node updated	A manually added node in the Topology was changed.	GigaVUE-FM
Module changed	A GigaVUE node has detected a change in line card/module type from the last polling interval. This typically happens when a line card/module is pulled from a slot or inserted in an empty slot.	G Series H Series
Node added	A new physical node was added to GigaVUE-FM.	GigaVUE-FM
Node cold start	A GigaVUE node restarted with a possible configuration change.	Any
Node failed to join cluster	The nodes failed to join the cluster.	
Node failed to remove from cluster	The nodes failed to leave the cluster.	
Node joined to cluster	The nodes joined the cluster.	
Node Link down	The link status on the indicated port changed from up to down.	Any

Table 39-3: GigaVUE-FM Notifications

Notifications	Description	Node Series
Node Link up	The link status on the indicated port changed from up to down.	Any
Node removed	A physical node was removed from GigaVUE-FM.	GigaVUE-FM
Node removed from cluster	The nodes are removed from the cluster.	
Node state changed	The status of a physical node in GigaVUE-FM changed from up to down (or vice-versa).	Any
Node warm start	A node restarted without changing its configuration.	Any
Packets dropped	A node detected dropped packets on the indicated port.	Any
Port link changed	A G Series node detected that a port's link status changed from up to down or vice-versa. **NOTE:** The **portlinkchange** trap is **not** sent when the Management port's link status changes. **NOTE:** The link state polling interval is 1 second. If a link state change is detected during the poll, a trap is generated.	G Series
Port Optics Temperature	Port optics temperature is above the acceptable level.	H Series
Port utilization below threshold Change	The utilization of the port is below the threshold limit.	
Power module changed	A G Series node detected either: • One of the two power supplies was powered on or off. • Power was lost or restored to one of the two power supplies.	G Series
Power source changed	The power source used by an A Series tap changed (for example, from Primary AC to Battery).	A Series
Process CPU utilization is high	An email notification is sent to all configured destinations each time the control card CPU utilization exceeds the pre-configured process threshold values.	H Series TA Series
Process Memory utilization is high	An email notification is sent to all configured destinations each time the control card memory utilization exceeds the pre-configured process threshold values.	H Series TA Series
System CPU utilization is high	An email notification is sent to all configured destinations each time the control card CPU utilization exceeds the pre-configured system threshold values.	H Series TA Series
System Memory utilization is high	An email notification is sent to all configured destinations each time the control card memory utilization exceeds the pre-configured system threshold values.	H Series TA Series
Resource Extreme High Usage Problem Cleared	An email notification is sent to all configured destinations each time the disk usage percentage falls below the extreme high usage threshold.	GigaVUE-FM
Resource Extreme High Usage Problem Detected	An email notification is sent to all configured destinations each time the disk usage percentage exceeds the extreme high usage threshold.	GigaVUE-FM
Resource High Usage Problem Cleared	An email notification is sent to all configured destinations each time the CPU and memory utilization falls below the high usage threshold.	GigaVUE-FM
Resource High Usage Problem Detected	An email notification is sent to all configured destinations each time the CPU and memory utilization exceeds the high usage threshold over the period of 2 minutes.	GigaVUE-FM

Table 39-3: GigaVUE-FM Notifications

Notifications	Description	Node Series
Scheduled task [%s] created	Any new task that is created under **All Nodes > Inventory > More Actions**, will trigger a notification.	H Series
Secondary Flash Boot	A boot partition next operation has occurred.	GigaVUE-FM
Service Status Down Detected	A service monitored by GigaVUE-FM is down.	GigaVUE-FM
Service Status Up Detected	A service monitored by GigaVUE-FM is up.	GigaVUE-FM
Stack image install finished	Install complete of a new image on the indicated stack of G Series nodes.	G Series
Stack image install started	Install start of a new image on the indicated stack of G Series nodes.	G Series
Stack reboot finished	A reboot of the specified stack of G Series nodes completed.	G Series
Stack reboot started	A reboot of the specified stack of G Series nodes was initiated.	G Series
Switch CPU Temperature	The switch CPU temperature is over the threshold limit.	
System Reset	A node has started, either as a result of cycling the power or a soft reset initiated by the **reload** command (H Series) or the **reset system** command (G Series).	Any
System Reset by Watchdog timer	A node detected that the watchdog monitor had to reset a failed process on the system.	H Series
TAP Relay changed	A G Series node detected a GigaTAP-Tx module's relays switched from active to passive or passive to active, as a result of the **config port-params taptx** command.	G Series
Thresholds exceeded	The utilization threshold configured for a port was exceeded: • **GigaVUE G Series** – The threshold configured with **config port-alarm** was exceeded for five consecutive seconds. • **GigaVUE H Series** – The threshold configured with **port <Port list> alarm utilization-threshold <percentage>** was exceeded for six consecutive seconds.	G Series H Series
Transmit / Receive error	G Series node received one of the following physical errors on a data port: • Undersize error • Fragment • Jabber • CRC or Alignment errors • Unknown errors	G Series
Unexpected shutdown	A H Series node shut down unexpectedly (for example, because power was lost and subsequently restored).	H Series
User authentication failed	A user login has failed on a H Series node.	H Series
VM Instance Running	A VM instance is running	
VM Instance Stopped	A VM instance is stopped.	
VM Instance Terminated	A VM instance is terminated	
Vmm Error	An error related to VMware ESXi inventory occurred.	GigaVUE-VM

Email Servers

The Email Servers page displays email hosts currently configured used for to send notifications and the email address for the From filed in email notifications.

To access the Email Servers page, click **Administration** on the top navigation link. On the left navigation pane, select **System > Email Servers**.

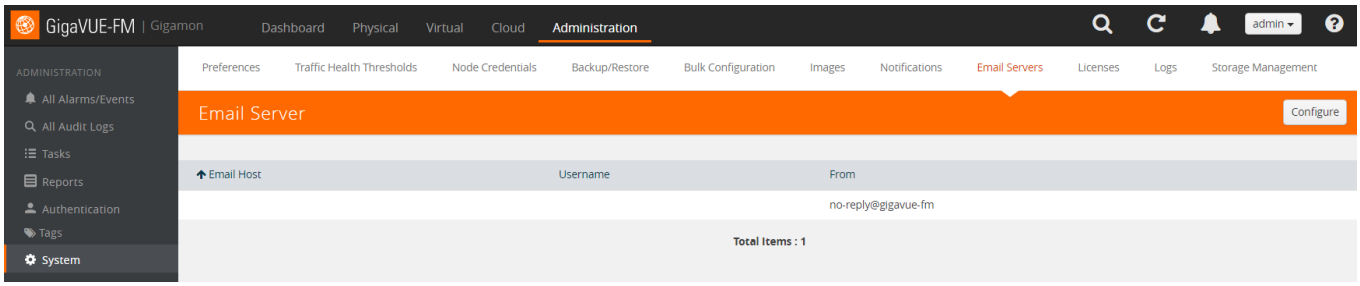

Figure 39-18: Setting Email Addresses for Alarm/Event Notifications

Click the Configure button on the Email Server page to open the configuration page shown in Figure 39-19 on page 1327. Table 39-4 on page 1327 describes the field on the Configure Email Server page.

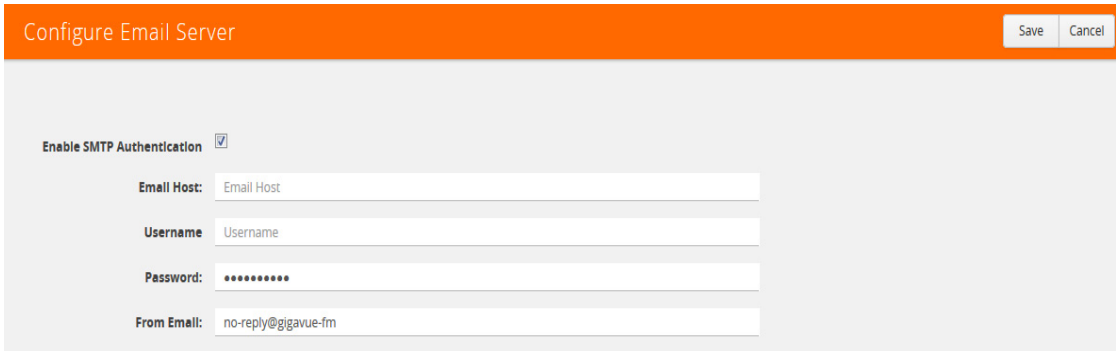

Figure 39-19: Configuring Email Servers for Notifications

Table 39-4: Configure Email Server Fields

Field	Description
Enable SMTP Authentication	The user's credentials are used for SMTP authentication when this option is select. When the option is not selected, SMTP authentication is disabled.
Email Host	The email server to be used for sending notification emails.
Username	The user name to login to the email server.
Password	The password for the user name.
From Email	The address you want to have show up in the From field of the notification emails.

Licenses

The **Licenses** page lets you review and apply licenses for the following components:

- GigaVUE-FM and GigaVUE-VM nodes using the **FM** tab
- Devices managed by GigaVUE-FM using the **Node** tab

GigaVUE-FM License

To access the GigaVUE-FM license(s) page, select **System > Licenses > FM**. You can generate, add and delete GigaVUE-FM licenses from this page. Figure 39-20 shows an example with four licenses currently installed.

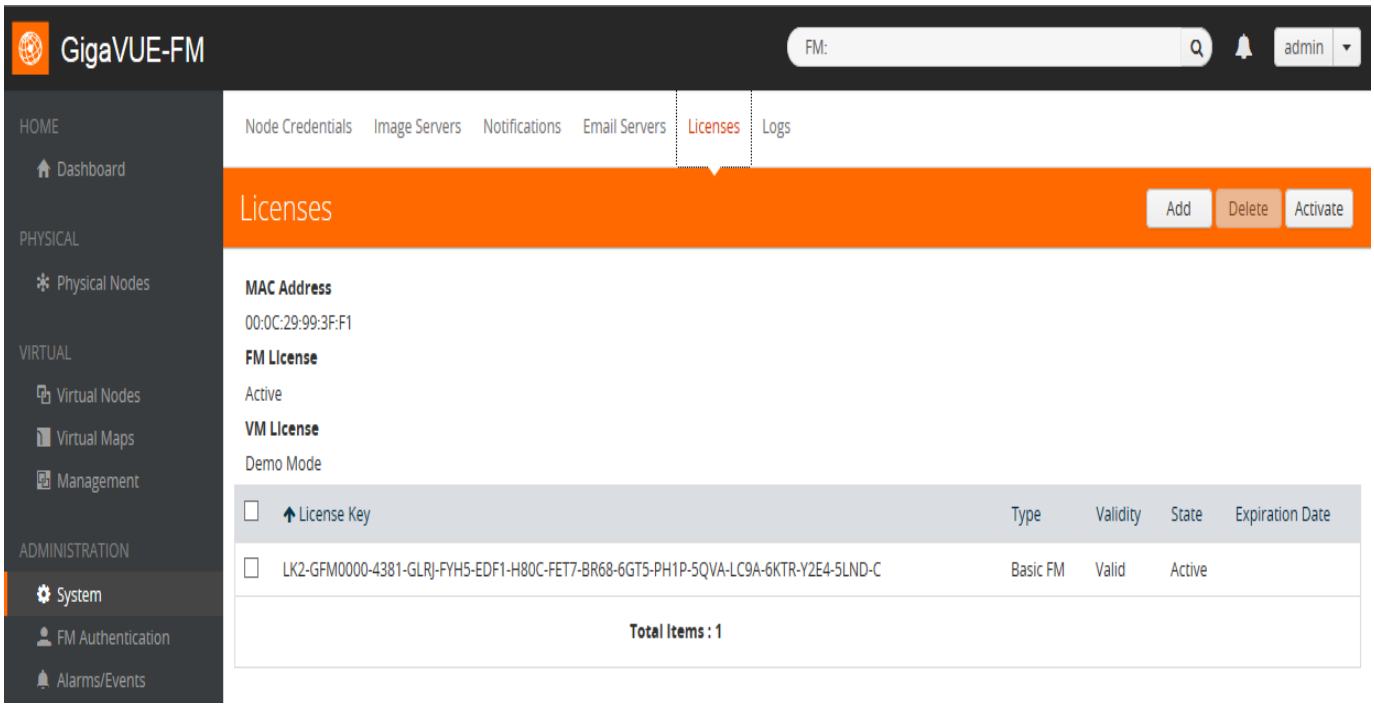

Figure 39-20: GigaVUE-FM Licenses Page

NOTE: If you use the GigaVUE-FM CLI command **show license**, the command may show an active Prime license as unrecognized.

Generate License

To generate or find the license or licenses, do the following:

1. Get the MAC address for your instance of GigaVUE-FM. To get the MAC address, click **Administration** on the top navigation link. On the left navigation pane, click **About**. The address is in the **MAC Address** field. Note the address for the next steps.
2. Go back to **System > Licenses** and click the **Generate** button.

   This will route you to the Gigamon Licensing portal or you can go to *Licensing Portal*.

3. To generate the license key or keys, navigate to the option for generating the GigaVUE-FM/VM licenses and then use the GIK provided in the email from Gigamon and the MAC address from Step 1. Refer to Figure 39-21 on page 1329.

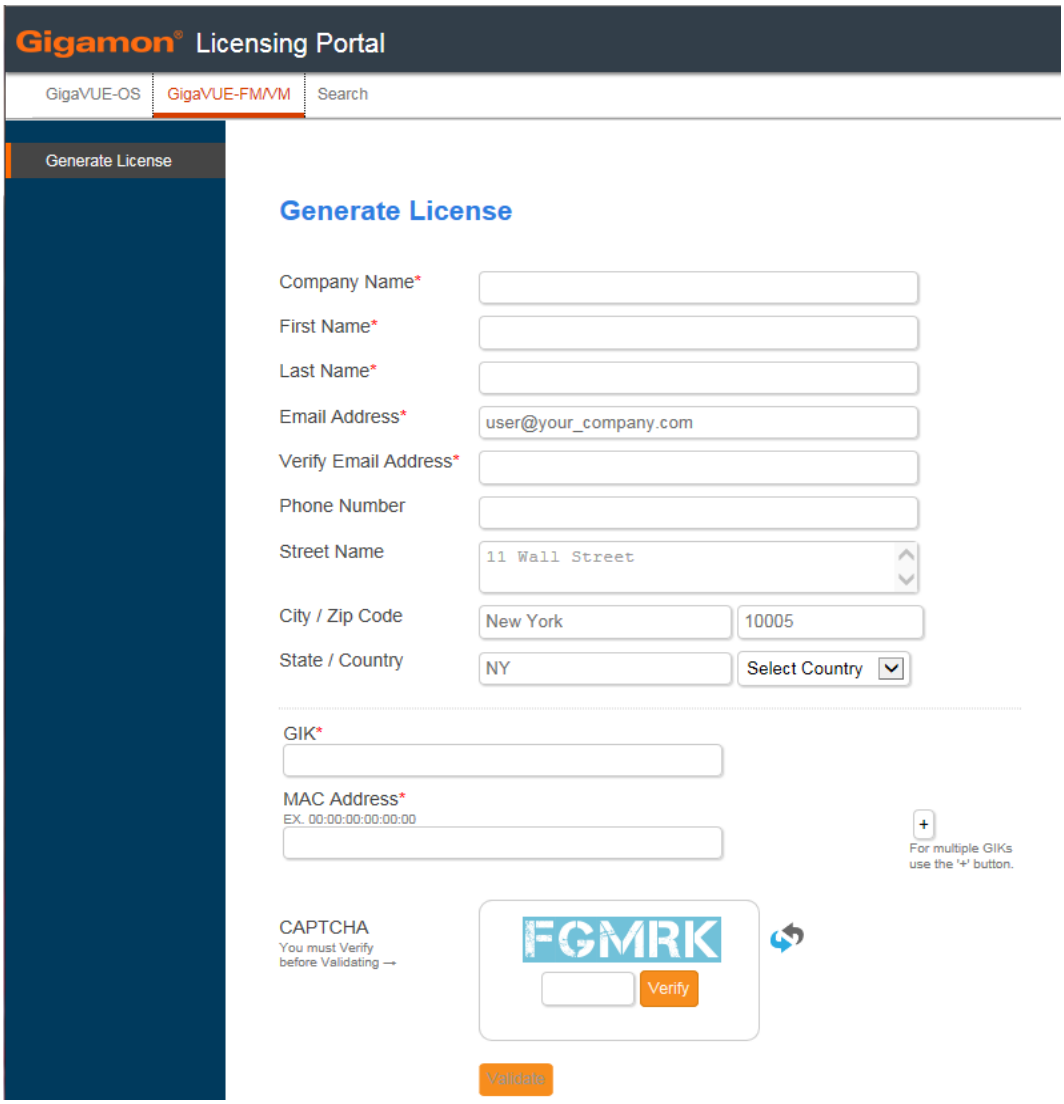

Figure 39-21: Gigamon Licensing Portal

To find the license key or keys of an already generated key, use the GIK in the Search tab. Refer to Figure 39-22.

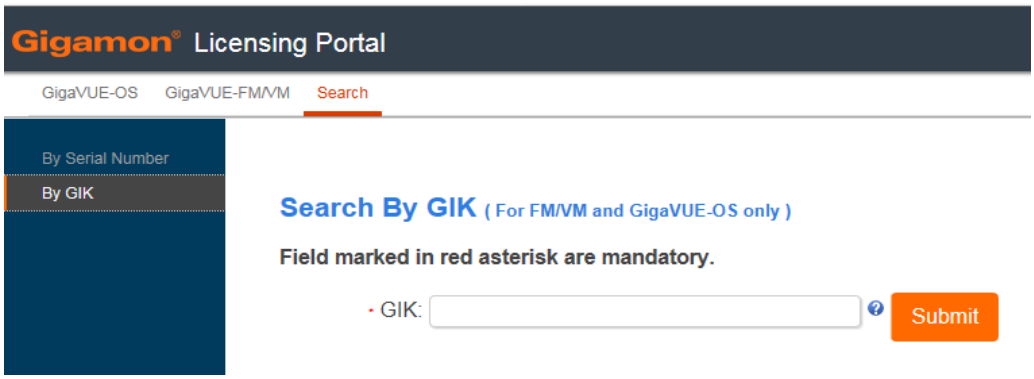

Figure 39-22: Search Option on Gigamon Licensing Portal

NOTE: To enable and view licenses for GigaVUE TA Series port enablement or clustering, or for GigaSMART licenses for GigaVUE H Series nodes, log in to the H-VUE or CLI for that node to apply the licenses. Currently GigaVUE-FM does not allow for this option.

4. After you have obtained the license key, follow the steps described in *Add License* on page 1330.

Add License

To add a license to GigaVUE-FM, do the following:

1. Generate a license key as described in *Generate License* on page 1328.
2. Select **System > Licenses**.
3. Click **Add**. The Add License page is displayed.

4. Enter the license key in the License Key field.

   To add more than one license, click the **+** button to add additional License Key field.
5. Click **Save**.

   The license and its description is added to the Licenses page.

Delete License

To delete a license, do the following:

1. Select **System > Licenses**.
2. On the Licenses page, select the license key for the license you want to delete.
3. Click **Delete** to remove the license.

Node License

To access the node license(s) page, select **System > Licenses > Node**.

NOTE: When you login to GigaVUE-FM, a notification is displayed on the top of the page with license expiry details for the nodes which have validity less than 30 days. Click the **Go To Licenses** option to go to the **Licenses** page.

This page displays the node and cluster details along with the license expiry information. The **Expires** option indicates the status of the license:

- **Never**: License is a life time license and will never expire
- **Expired:** License has already expired
- **Exact date:** Date when the license will expire (for example Oct 12, 2019)

You can generate and add licenses for the devices from this page.

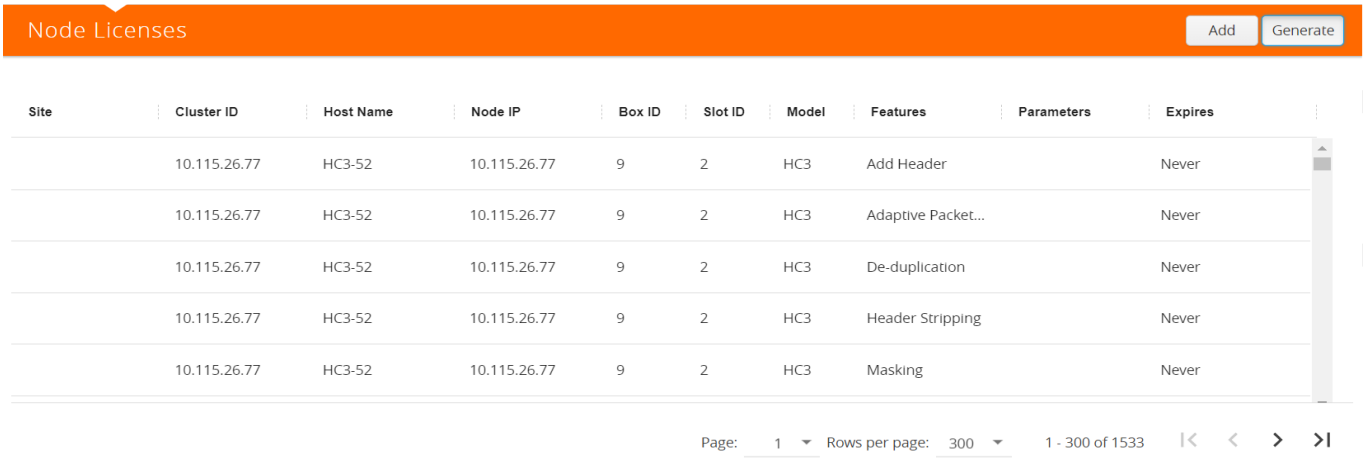

Figure 39-23: Node Licenses Page

Generate License

To generate or find the license or licenses, do the following:

1. Click **Generate.**
2. This will route you to the **Gigamon Licensing portal** (https://licensing.gigamon.com). Click the GigaSMART tab.
3. Enter the Serial Number of the device and the GIK (received through email). Select the end-user license agreement.
4. Click **Submit** to generate.

NOTE: Please contact technical support team for assistance.

Add License

To add a license to the device:

1. Generate the license as described in section *Generate License* on page 1331.
2. Select the required Cluster ID and Box ID.
3. Enter the License Key.
4. Click **Save.**

The license is added to the Licenses page.

System Logs

You can generate log files that contain information about the system. Gigamon support can use these files for root cause analysis. Click the **Download** button to download the compressed files.

To access Logs, click **Administration** on the top navigation link. On the left navigation pane, select **System > Logs**.

Create Log file

To create a log file that Gigamon can use for analysis, do the following:

1. Select **System > Logs**.

   The Logs page displays, which shows a list of log files.

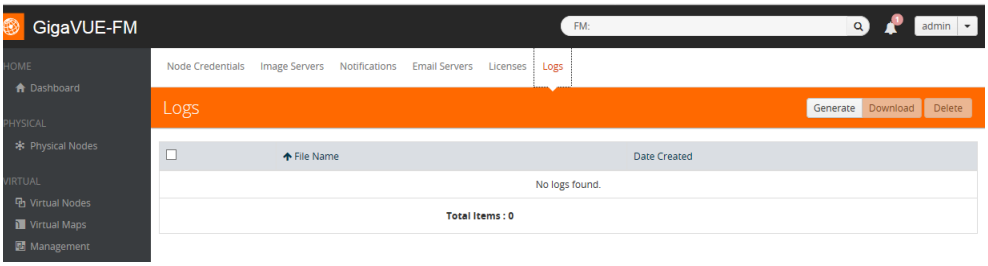

Figure 39-24: Logs Page

2. Click **Generate**.

   The system generates a new log file and displays an event message. Figure 39-25 shows an example of a successfully generated log file.

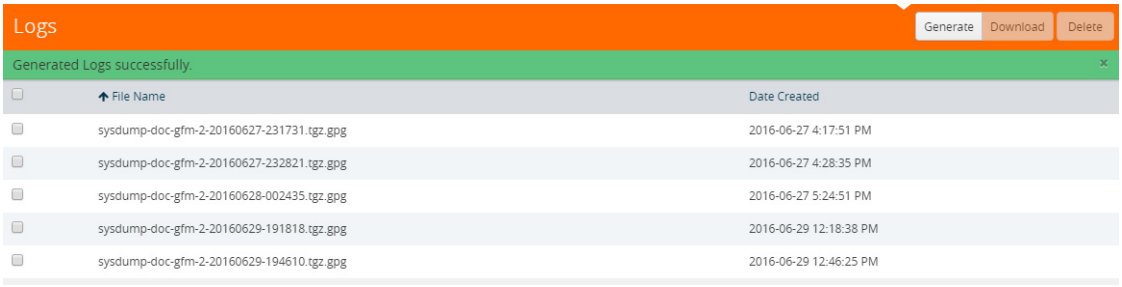

Figure 39-25: Successfully Generated Log File

3. Select the log file to download, and then click **Download**.

   The system downloads the file to your local environment. The file is in a compressed and encrypted format that you can provide to Gigamon.

Delete Log File

To delete the log files for clearing up the disk space:

1. Select **System** > **Logs**.

   The Logs page displays a list of log files. Refer to Figure 39-25.

2. Select the Logs that you want to delete and click **Delete**.

Storage Management

The Storage Management page shows the available and used storage space shown under /var file system of the GigaVUE-FM appliance. Table 39-26 shows that 2588MB of storage is used and 25750MB is available only 9 percent full. This information is pulled from the same file system irrespective of the virtual environment where the appliance is installed.

This information is useful when collecting NetFlow records for FabricVUE Traffic Analyzer, reports, and audit logs because the appliance may run out of storage and there might be a degradation in performance. Generally, if everything is functioning well, the NetFlow records would be transferred to /config file system and this issue may never arise.

GigaVUE-FM Storage Management allows you to define how the stored logs are managed. You can specify a schedule for purging old device logs. You can also specify an SFTP server to export the log records prior to purging. Storage Management is used for all storage settings, including device logs, alarm/event notifications, and statistics. Refer to *All Alarms/Events* on page 1269.

To access Storage Management, click **Administration** on the top navigation link. On the left navigation pane, select **System > Storage Management**.

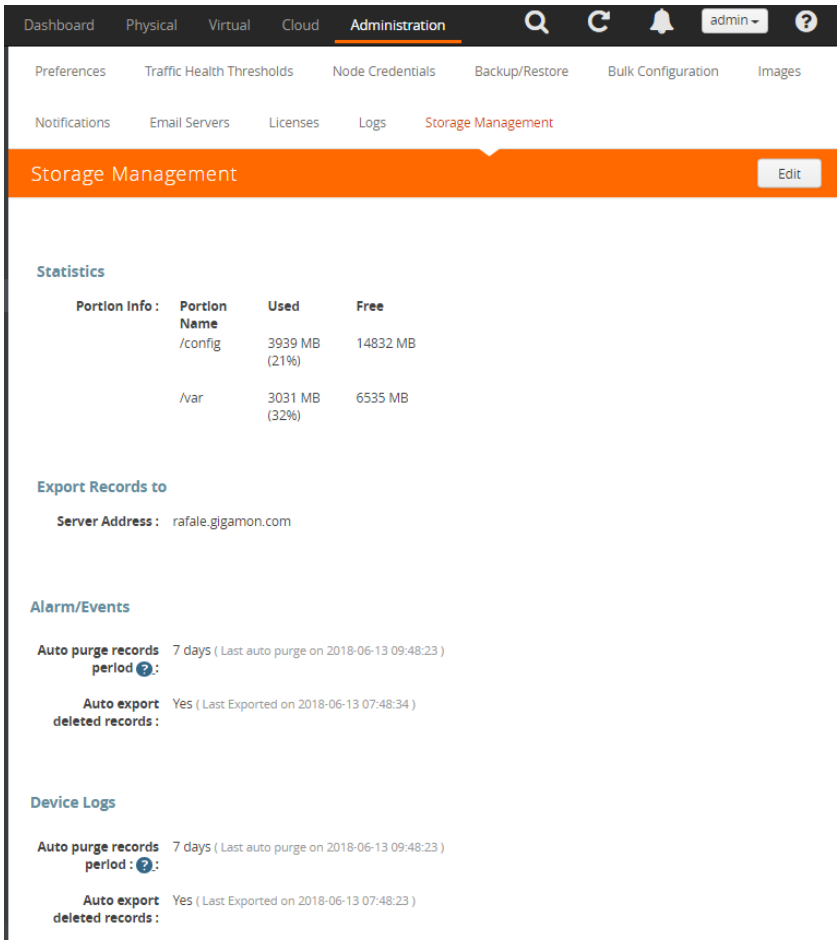

Figure 39-26: Storage Management Page

If needed, you can free the used storage older than a specified date by doing the following:

1. Select **Administration** from the top menu.
2. Navigate to **System > Storage Management**. The Storage Management page is displayed.
3. Click **Edit** to edit the settings.

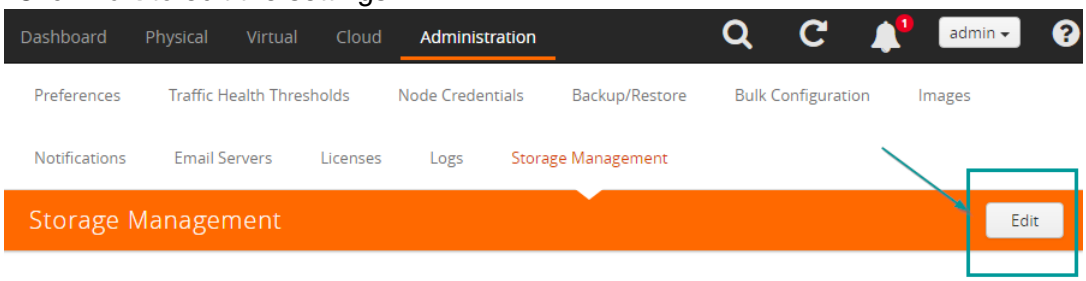

The Edit Storage Management page appears.

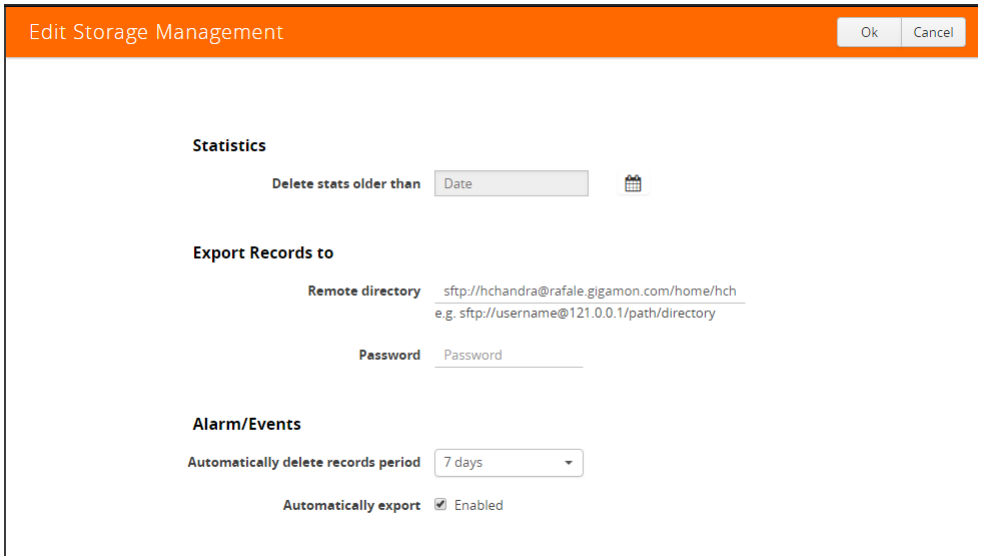

4. Specify the Storage Management settings for each type of record:

Setting	Description
Statistics	
Delete stats older than	Click the Calendar icon to select a date. This will specify a cut off date for deleting statistics. Statistics prior to the specified date will be deleted, immediately, when you click **OK**. **Important**: If a date is specified, the records will be immediately and permanently deleted from the database when you click **OK**.
Export Records To	
Remote directory	If the “automatically export” check box is selected under Alarm/Events or Device Logs, the records will be exported to a CSV file at the specified interval. Use this field to specify the ftp/sftp location to send the export records. For example: `sftp://username@121.0.0.1/path/directory`
Password	Specify a password for accessing the remote server.
Alarm/Events	
Automatically delete records period	Specify how often to delete Alarm/Event records. Options are 7 days, 30 days, 60 days or custom (in days). (7 days is the default). When you click **OK**, the records older than the specified duration will get deleted immediately. Records will be purged regularly while maintaining records for the specified duration. This means that, at any given time, records of the specified duration will be available to view from the GigaVUE-FM Alarms/Events page. Click the **Automatically export** check box to enable exporting the records to the specified location on a periodic basis. (It is enabled by default.) When selected, records will be exported once before the immediate purge, and then again in the number of days specified in the records period. Records will continue to be exported once every set period of days.

Setting	Description
Device Logs	
Automatically delete records period	Specify how often to delete Device Log records. Options are 7 days, 30 days, 60 days or custom (in days). (7 days is the default). When you click **OK**, the records older than the specified duration will get deleted immediately. Records will be purged regularly while maintaining records for the specified duration. This means that, at any given time, records of the specified duration will be available to view from the GigaVUE-FM Logs page for the node. Click the **Automatically export** check box to enable exporting the records to the specified location on a periodic basis. (It is enabled by default.) When selected, records will be exported once before the immediate purge, and then again in the number of days specified in the records period. Records will continue to be exported once every set period of days.

e. To permanently remove the records from the database based on the specified settings, click **OK**.

Caution: There is no undo. Statistics records prior to the date specified will be immediately and permanently deleted from the database when you click **OK**. Alarm/Event records will be permanently deleted from the database at the specified scheduled interval.

SNMP Traps

The SNMP Traps page shows the configuration settings applied to the devices managed by GigaVUE-FM. This page also allows you to configure the settings that need to be applied to all the devices managed by that GigaVUE-FM instance.

When the GigaVUE-FM instance starts, the following SNMP traps are enabled by default (for all the devices):

- Link Status or Speed Change
- Port Link Change
- Module Change
- Fan Status Change
- Power Supply Status Change
- Inline Bypass Forwarding State Change

To access SNMP Traps, click **Administration** on the top navigation link. On the left navigation pane, select **System > SNMP Traps**.

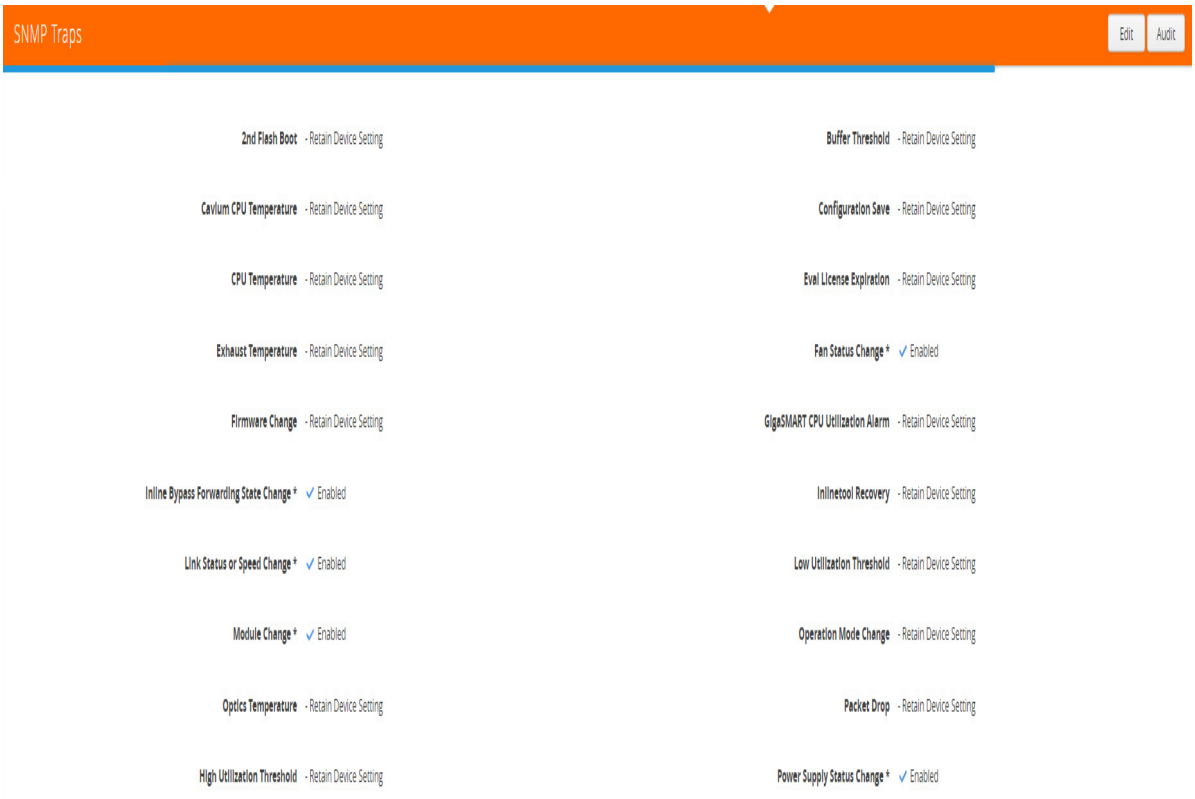

Figure 39-27: SNMP Traps Page

The SNMP Traps page allows you to configure the following:

- Enable/disable all the traps for all the devices using the **Enable All** and **Disable All** options
- Enable/disable specific traps for all the devices using the **Enable** and **Disable** options for each of the traps
- Retain device settings for all the traps using the **Retain Device Settings for All Traps** option

To configure the SNMP traps:

1. Click the **Edit** button on the top right corner.
2. Configure the required setting. For example, to retain the individual device level setting for all traps, select the **Retain Device Setting for All Traps** checkbox.

NOTE: You can also retain the device settings for specific traps by selecting the **Retain Device Setting** checkbox against the required traps.

3. Click **Save**

With this functionality, the following configuration settings are applied to all the devices:

- Specific configuration type changes
- Audit configuration changes

NOTE: If a new device is added to GigaVUE-FM, then the global configuration setting is applied to the new device. If for some reason, the configuration setting is not applied to a device, then an event is raised with the appropriate details in the Alarms/Events page.

If a trap has been forcefully enabled/disabled on a device because of the global configuration setting, then an event is raised with the appropriate details in the Alarms/Events page.

40 Roles and Users in GigaVUE-FM

This chapter provides basic information about role-based access and the procedures to manage roles and users in GigaVUE-FM along with assigning access permissions. The following topics are covered:

About Role-Based Access

GigaVUE H Series and TA Series nodes use role-based access to manage access to the Gigamon Visibility Fabric. Creating roles and assigning users to those roles from H-VUE or CLI allows for the partition of separate sets of tool ports for different groups of users while different sets of network ports are shared. This makes it possible to provide different groups of users with different analysis needs to have full access to the packets they need for their tools.

GigaVUE-FM provides a mode that controls whether GigaVUE-FM manages devices under a single admin-level device account provided during the Add Physical Node process. The devices managed by GigaVUE-FM should be able to validate user credentials with the credentials provided to the device during the login from GigaVUE-FM. If the user account does not exist on the device or the passwords do no match GigaVUE-FM cannot log in to the device. To minimize this type of problem, it is recommended that both GigaVUE-FM and the managed device validate user credentials against a common authentication service (such as LDAP, RADIUS, or TACACS+).

GigaVUE-FM has an RBAC mode is that is set by default. When in RBAC mode, GigaVUE-FM ensure that the users are added to the local server or the central server (LDAP, RADIUS, or TACACS+) on the GigaVUE-FM with the same node credentials as the device. To set the mode, do the following:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **Authentication > RBAC**.

   The RBAC page shown in Figure 40-1 displays.

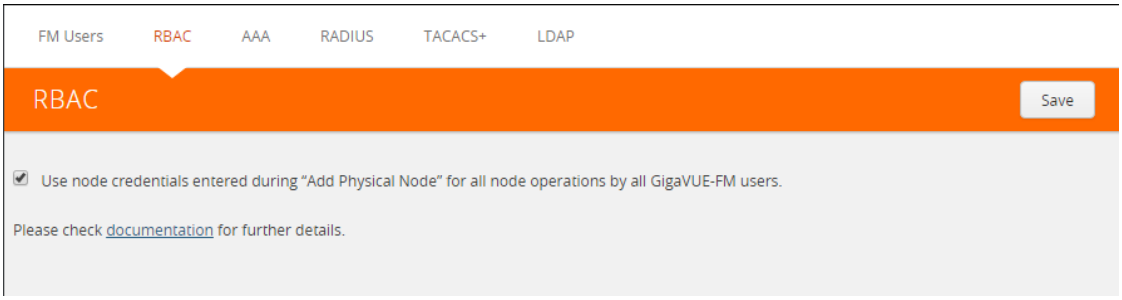

Figure 40-1: Enabling or Disabling RBAC Mode on GigaVUE-FM

3. Select or clear the checkbox to set the mode as follows:
   - If the checkbox is cleared (which is the default), for a GigaVUE-FM user to be able to manage a node, the user should have the same credentials (username and password) in both GigaVUE-FM and the node. If the number of nodes and/or devices is large, it is recommended that LDAP or similar mechanism be used to ease user management.
   - If the checkbox is selected, for a GigaVUE-FM user to be able to manage a node, the user does NOT have to have the exact same credentials in GigaVUE-FM and the node. GigaVUE-FM will use the credentials entered when the node was added to GigaVUE-FM for all operations performed by all users.

In both the cases, GigaVUE-FM RBAC is enforced. For example, a GigaVUE-FM user with the role fm_user will not be able to modify anything on the node.

NOTE: Selecting or clearing the checkbox has no impact on the following operations performed by GigaVUE-FM:

- Rediscovery
- Configuration sync
- Statistics collection

Irrespective of whether the checkbox is selected/cleared, GigaVUE-FM uses the credentials stored in the GigaVUE-FM database to managing the devices.

4. Click **Save** to set the mode.

For more detailed information related to role-based access, refer to the following sections:

- *Access Levels on GigaVUE-FM* on page 1343
- *Role-Based Access and Flow Mapping* on page 1345

Access Levels on GigaVUE-FM

When in RBAC mode in GigaVUE-FM, any admin level user can add a node to the GigaVUE-FM. However, when adding the node credentials, if the node credentials do not match the privileges on the node, the admin user from GigaVUE-FM will not be able to manage the node.

To manage the node from GigaVUE-FM, the admin username and password has to match exactly the username and password on the node with the same privileges.

To understand more on the level of privileges and how to manage them on GigaVUE-FM versus on the node, refer to the Figure 40-2, Table 40-1, and Table 40-2

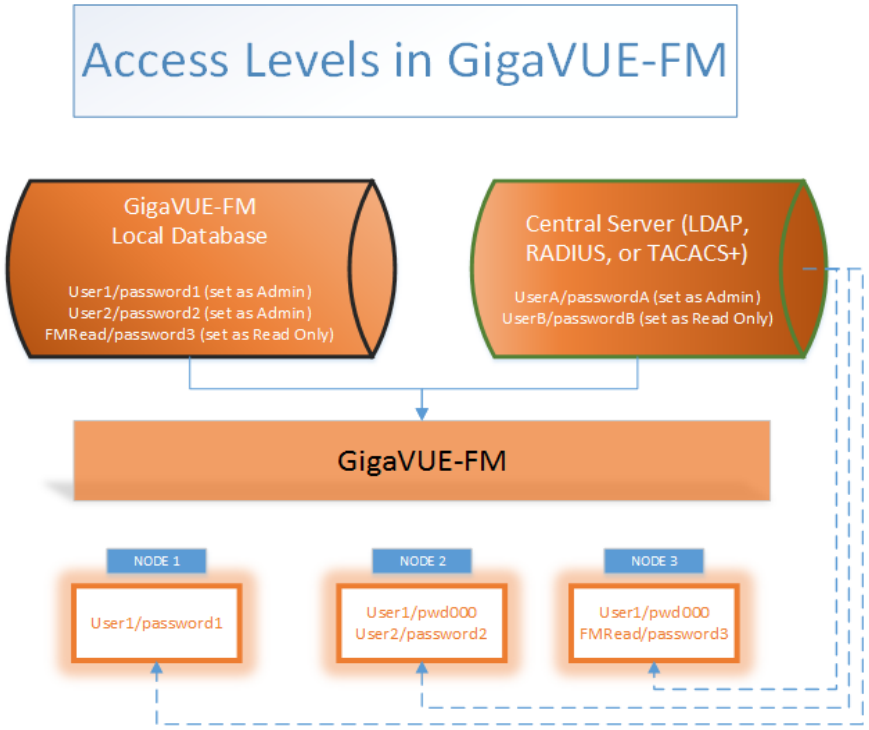

Figure 40-2: Graphical Representation of User Databases and Management

Table 40-1: Access Levels from GigaVUE-FM for Scenario in Figure 40-2 when "Use Node Credentials Entered During Add Physical Node" Button is NOT Selected

User Privileges from GigaVUE-FM	GigaVUE-FM	Node1	Node2	Node3
GigaVUE-FM Local User Database				
User1/password1 (set as Admin on GigaVUE-FM)	Admin	Admin	Read Only	No Access
User2/password2 (set as Admin on GigaVUE-FM)	Admin	No Access	Admin	No Access
FMRead/password3 (set as Read-Only on GigaVUE-FM)	Read Only	No Access	No Access	Read Only
Users on the Central Server (LDAP, RADIUS, TACACS+)				
UserA/passwordA (set as Admin on GigaVUE-FM)	Admin	Admin	Admin	Admin
UserB/passwordB (set as Read Only on GigaVUE-FM)	Read Only	Read Only	Read Only	No Access
Local Node Database				
User1/pwd000 (set as Admin on local Node but not set up on GigaVUE-FM)	No Access	No Access	No Access	No Access

Table 40-2: Access Levels from GigaVUE-FM for Scenario in Figure 40-2 when "Use Node Credentials Entered During Add Physical Node" Button is Selected

User Privileges from GigaVUE-FM	GigaVUE-FM	Node1	Node2	Node3
GigaVUE-FM Local User Database				
User1/password1 (set as Admin on GigaVUE-FM)	Admin	Admin	Admin	Admin
User2/password2 (set as Admin on GigaVUE-FM)	Admin	Admin	Admin	Admin
FMRead/password3 (set as Read-Only on GigaVUE-FM)	Read Only	Read Only	Read Only	Read Only
Users on the Central Server (LDAP, RADIUS, TACACS+)				
UserA/passwordA (set as Admin on GigaVUE-FM)	Admin	Admin	Admin	Admin
UserB/passwordB (set as Read Only on GigaVUE-FM)	Read Only	Read Only	Read Only	Read Only
Local Node Database				
User1/pwd000 (set as Admin on local Node but not set up on GigaVUE-FM)	No Access	No Access	No Access	No Access

Role-Based Access and Flow Mapping

Flow Mapping allows different users to share network ports. Because Flow Mapping sends a packet matching multiple maps to the destination specified by the map with the highest priority, you must exercise caution when adjusting maps on shared network ports. Administrators can change the priority of maps to ensure that packets are sent to the desired destination.

Permission can also be associated with maps based on roles. For more information about map permissions, refer to *Set Map-Sharing Permission Levels* on page 1347

Configure Role-Based Access and Setting Permissions

Configuring RBAC in GigaVUE-FM consists of the following tasks:

- *Add Users* on page 1346
- *Create Roles* on page 1347
- *Set Map-Sharing Permission Levels* on page 1347

Add Users

This section describes provides the steps for adding users. Users are also assigned to roles that set there access permissions. For the step to create roles, refer to *Create Roles* on page 1347.

To add users, do the following:

1. Click **Administration** on the top navigation link.
2. On the left navigation pane, select **Authentication** > **FM Users**. The **FM Users** page displays.

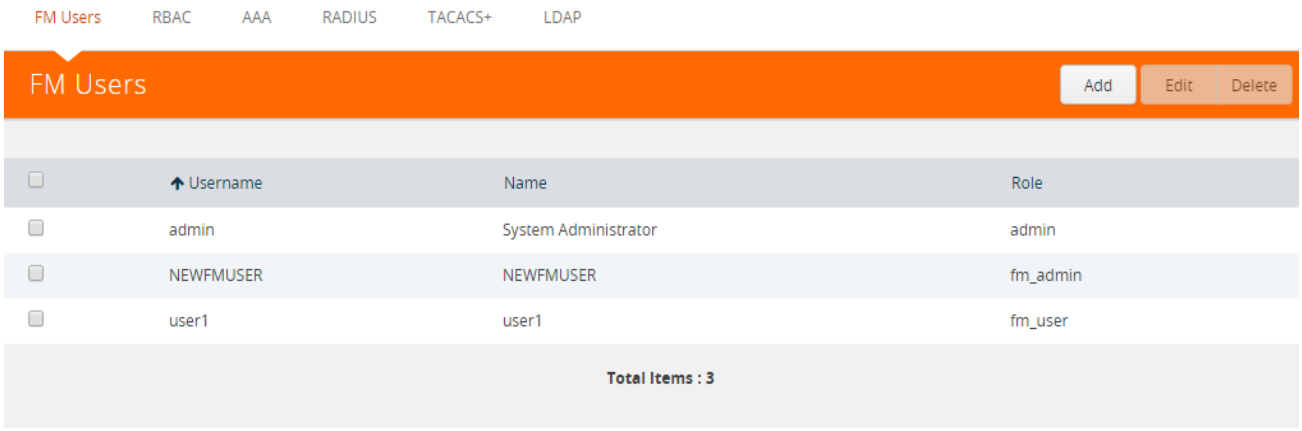

Figure 40-3: FM Users Page

3. Click **Add**. The **Add New User** page displays.

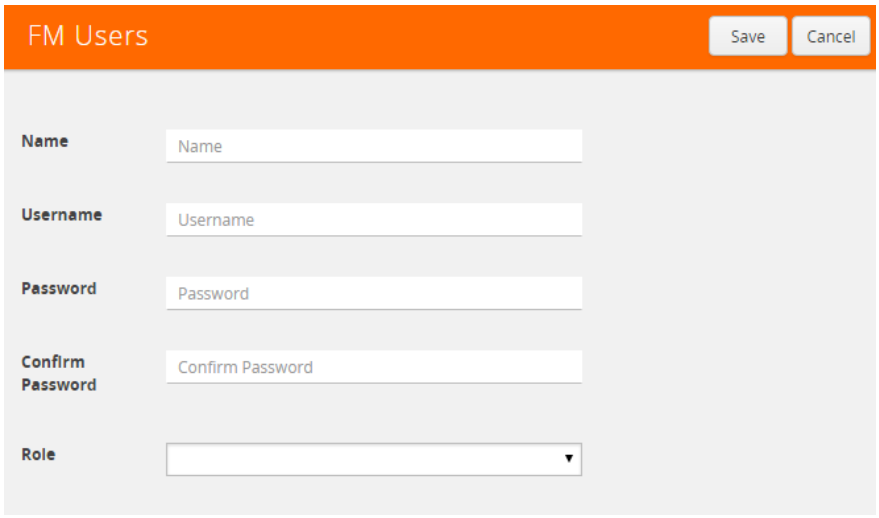

Figure 40-4: Add User Page

4. On the Add New User page, do the following:
   - Enter the user's actual name in the **Name** field.
   - Enter a user name in **User Name** field.

- Enter a password for the user in the **Password** field and in the **Confirm Password** field.
- Assign a role to the user by clicking in **Capability** field and selecting a role from the drop-down list. For the steps to create a role, refer to *Create Roles* on page 1347.

5. Click **Save**.

Create Roles

This section describes the steps for creating roles and assigning user to those roles using the GigaVUE-FM UI. Note that the users created through the CLI have additional roles available including the FM roles.

NOTE: Before creating roles, refer to *About Role-Based Access* on page 1342.

GigaVUE-FM UI has three built-in roles for specifying which users have access to a given port. These roles are:

- **fm_super_admin** — Allows a user to do everything in Fabric Manager, including adding or modifying users and configuring all AAA settings in the RADIUS, TACACS+, and LDAP tabs. Can change password for all users.
- **fm_admin** — Allows a user to do everything in Fabric Manager except add or modify users and change AAA settings. Can only change own password.
- **fm_user** — Allows a user to view everything in Fabric Manager, including AAA settings, but not make any changes.

GigaVUE-FM CLI has three built-in roles in addition to the roles available from the UI. These roles are:

- **Admin**
  - This role is essential for GigaVUE-FM upgrades. All GigaVUE-FM software come with the default admin role. This role is also available on the UI.
  - Take care not to delete this role.
- **Operator**
- **Monitor**

Set Map-Sharing Permission Levels

Map-Sharing Permissions are only available to node level users. FM users cannot be added for this level of permissions. When creating maps on the node from GigaVUE-FM, admin users can add map permissions pertaining to the node level users. There are four map-sharing permission levels:

Permission Level	Description
View	Role can view the map but cannot make any changes.
Listen	Role can add or remove tool ports they own[1]. This is equivalent to *subscribing* to a map.

Permission Level	Description
Edit	Role can delete and edit the map, can remove any network ports, can add network ports they own[1], and can add or remove tool ports they own[1].
Owner	Role can perform all the Read/Write functions and assign map sharing permission levels.

1. Requires Level 2 or Level 3 access, based on the user's role membership.

To set permissions for a map, do the following:

1. Click **Physical** on the top navigation link.
2. On the Physical Nodes page, click on a physical node.
3. On the left navigation pane of the GigaVUE node, click **Maps**.
4. Select a map.

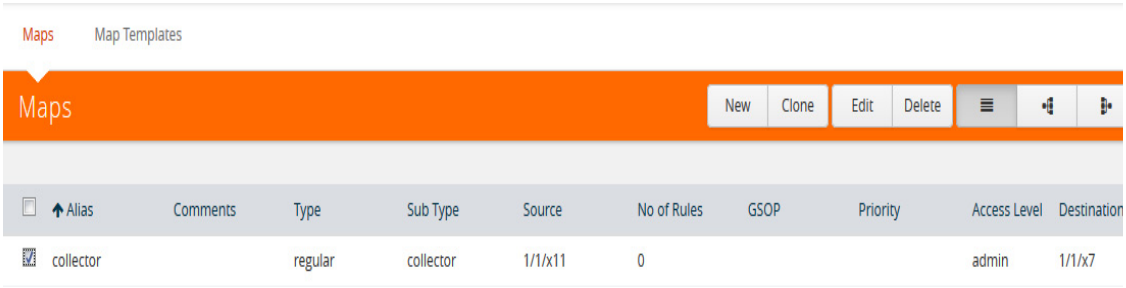

Figure 40-5: Maps Page

5. Click **Edit**.
6. Go to the **Map Permissions** section of the **Edit Map** page.
7. Click in the **Owner**, **Edit**, **Listen**, or **View** field and select roles from the drop-down list.

Part 9: Appendixes

This section provides additional information useful for GigaVUE-FM. The following topics are covered:

A Disk Size on GigaVUE-FM

This section describes how to increase the storage space available on GigaVUE-FM in VMware ESXi, Microsoft Hyper-V. and KVM environments. It also explains how to clear the space in the /var directory.

After installing GigaVUE-FM, the size of /config can be increased by adding disk and then rebooting GigaVUE-FM. *Increase Disk Size on a New or Existing GigaVUE-FM Installation* on page 1352 describes increasing the disk size for a new or existing GigaVUE-FM installation as well as for an upgrade from previous releases. Ensure that you apply this procedure only to GigaVUE-FM version 3.1 or higher.

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

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

1. To verify the disk size through the GigaVUE-FM CLI, use the following command:

```
config(#) show files system
Statistics for /config file system
  Space Total              19777 MB
  Space Used               1275 MB
  Space Free               18501 MB
  Space Available          17497MB
  Space Percent Free       93%
   Inodes Percent Free     99%

Statistics for /var filesystem
  Space Total              10079 MB
  Space Used               359 MB
  Space Free               9721 MB
  Space Available          9209 MB
  Space Percent Free       96%
  Inodes Percent Free      99%
```

2. From the vSphere Client, select **Edit Settings > Hardware > Add**.
3. Select Hard Disk as shown in Figure 0-1.

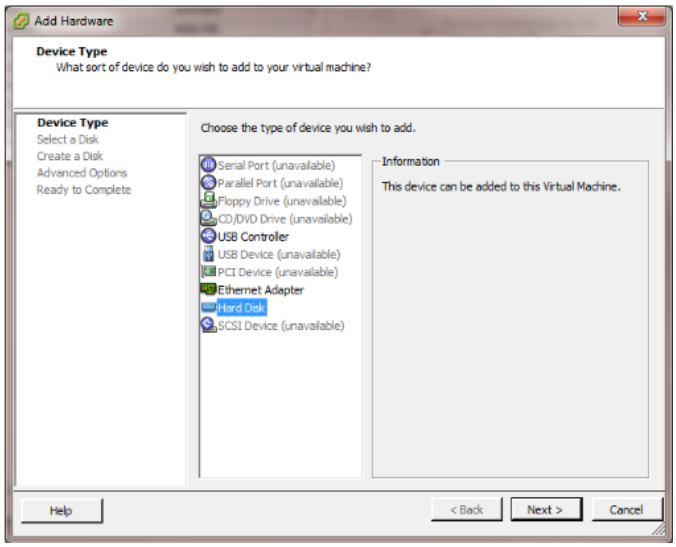

Figure 0-1: Select Hard Disk from vSphere Client

4. Click **Next**.
5. Select the size of the disk as desired. In Figure 0-2, the size is set to 16GB. Ensure only Thick Provisioning is selected.

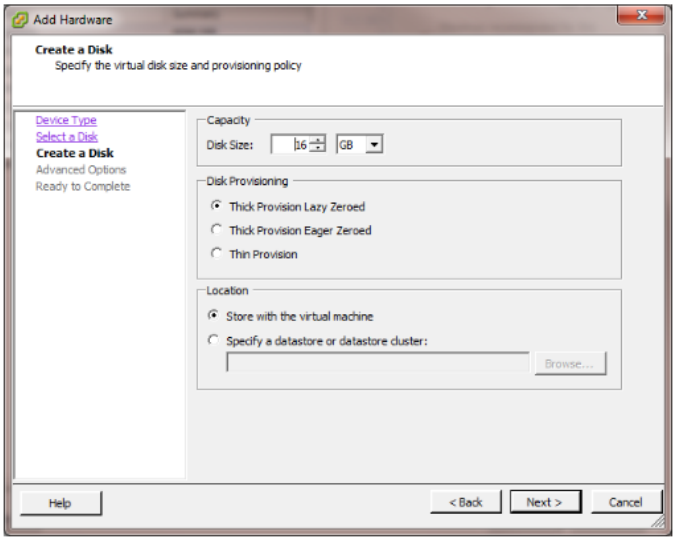

Figure 0-2: Example Shows the Disk Size Set as 16GB

6. Click **Next.**
7. The default for the Virtual Device Node is displayed in the dialog as shown in Figure 0-3. It is best to keep the default.

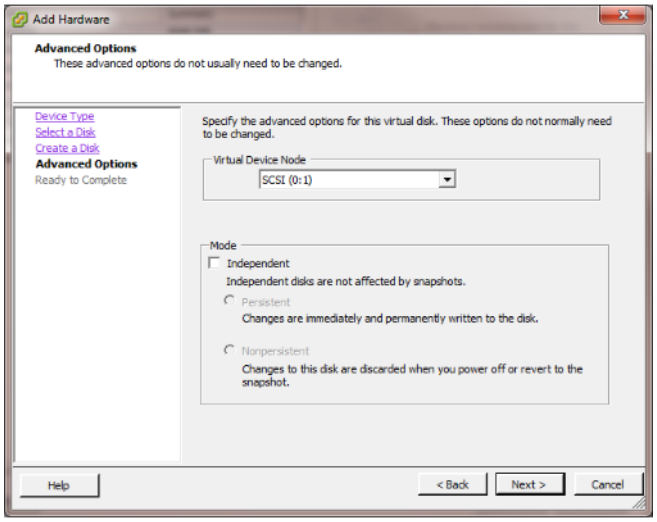

Figure 0-3: Default Selection for Virtual Device Node

8. Click **Next** for summary and complete by clicking **Finish**. Figure 0-4 shows an example of the Summary.

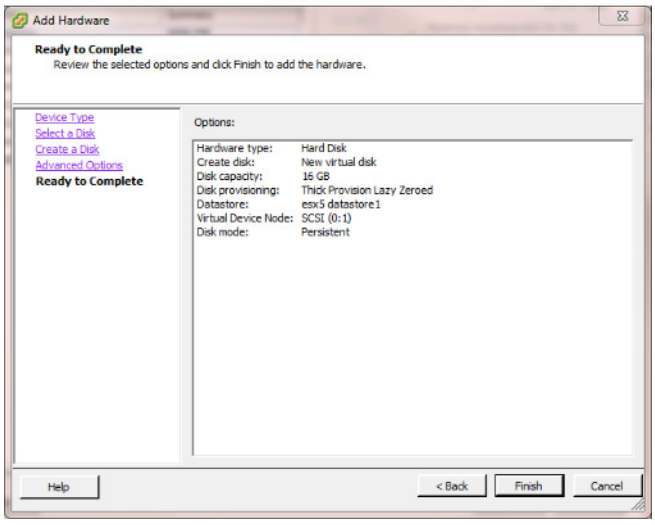

Figure 0-4: Summary Screen for Disk Capacity

After clicking **Finish**, the Edit Settings dialog reopens where you will see the new hard disk information in the list.

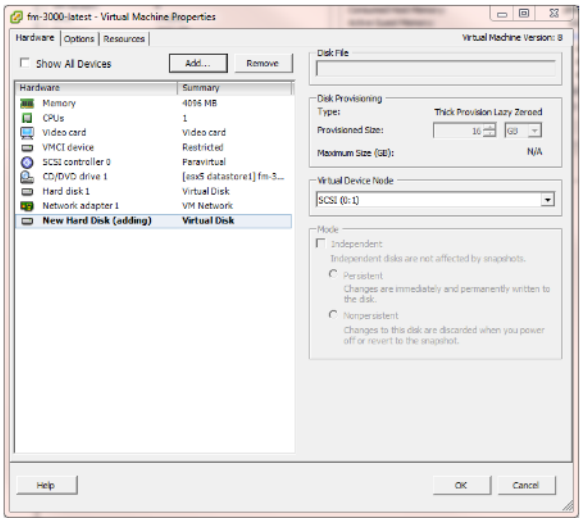

Figure 0-5: Edit Settings Screen with the New Hard Disk Displayed

9. After verifying the information in the Edit Settings dialog, click **OK** and reboot the GigaVUE-FM using **Restart Guest OS** in vSphere.

   As GigaVUE-FM reboots, messages will appear on the screen showing the change in the disk size. Figure 0-6 shows an example.

   NOTE: For /var, the size is fixed at 10GB. When the size is increased by adding a new disk, only the size of /config is increased.

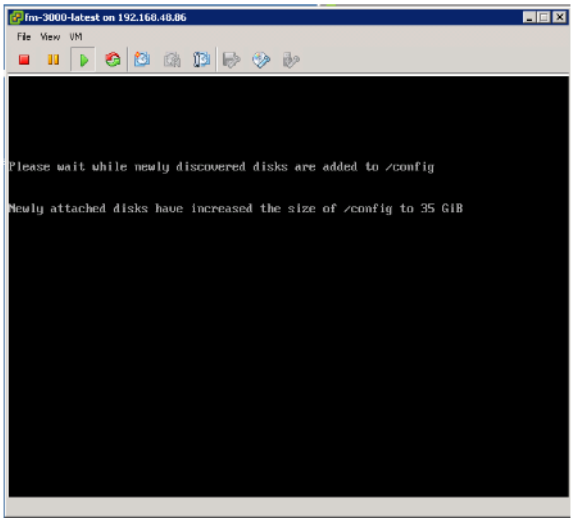

Figure 0-6: GigaVUE-FM Rebooting Shows the Increased Disk Size

10. To verify the disk size after a disk size increase through the GigaVUE-FM CLI, use the show files system command.

    The following is an example after an upgrade.

```
config(#) show files system
Statistics for /var filesystem:
  Space Total          10079 MB
  Space Used           358 MB
  Space Free           9721 MB
  Space Available      9209 MB
  Space Percent Free   96%
  Inodes Percent Free  99%

 Statistics for /config filesystem:
   Space Total          35900 MB
   Space Used           1279 MB
   Space Free           34621 MB
   Space Available      32797 MB
   Space Percent Free   96%
   Inodes Percent Free  99%
```

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

The /var directory can sometimes run out of storage space and result in performance degradation of GigaVUE-FM. The best practice is to periodically check the storage space (refer to *Storage Management* on page 1334) and clear up the disk space to avoid system misbehavior.

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

- Purging the statistics older than a certain date. For more information, refer to *Storage Management* on page 1334.
- Removing the unused internal images. For more information, refer to *Internal Image Files* on page 1317.
- Deleting the logs. For more information, refer to *Delete Log File* on page 1333.

B Data Transfer Rate Units

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

Table b-1: Decimal Multiple of Data Transfer Rate in Bits

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

NOTE: 8 kilobits = 1 kilobyte

Table b-2: Decimal Multiple of Data Transfer Rate in Bytes

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

C Open Ports in the Firewall

This appendix provides information about the default ports open in the firewall for GigaVUE-FM. The *Open Default Ports* table is sorted by Protocol and then Port Number.

Table C-1: Open Default Ports

Port Number	Protocol	Service	Traffic Direction	Description
80	HTTP	GigaVUE-FM GUI	Bidirectional traffic between Web Browser and GigaVUE-FM	Used for redirecting the traffic internally to port 443.
443	HTTPS	GigaVUE-FM GUI	Bidirectional traffic between Web Browser and GigaVUE-FM Bidirectional traffic between GigaVUE-FM and GigaVUE-VM	Used for normal GigaVUE-FM user interaction.
389	LDAP	AAA	Bidirectional traffic between LDAP server and GigaVUE-FM	Used for accessing and maintaining distributed directory information services over the Internet Protocol (IP) network.
636	LDAP	AAA	Bidirectional traffic between LDAP server and GigaVUE-FM	Used for secure LDAP protocol over SSL for accessing and maintaining distributed directory information services over the Internet Protocol (IP) network.
1812/ 1813 1645/ 1646	Radius	AAA	Bidirectional traffic between Radius server and GigaVUE-FM	Used for running the client/server protocol in the application layer. They can use either TCP or UDP as the transport protocol.
49	TACACS	AAA	Bidirectional traffic between TACACS server and GigaVUE-FM	Used for communicating with the authentication server in order to determine if you have access to the network.
22	TCP	SSH	Bidirectional traffic between Putty and GigaVUE-FM	Used for GigaVUE-FM admin user login. Also, used for initial GigaVUE-FM IP configuration. Used by the web browser to communicate with GigaVUE-VM for troubleshooting purposes.
514	TCP	Logstash	Unidirectional traffic from physical nodes to GigaVUE-FM	Used for sending device log messages via SSL from devices to GigaVUE-FM.
5672	TCP	RabbitMq	Unidirectional traffic from physical nodes to GigaVUE-FM	Used for sending event notifications via SSL from devices to GigaVUE-FM.
5671	TCP/SSL	RabbitMq	Unidirectional traffic from physical nodes to GigaVUE-FM	Used for sending event notifications via SSL from devices to GigaVUE-FM.

Table C-1: Open Default Ports

Port Number	Protocol	Service	Traffic Direction	Description
53	UDP	DNS	Bidirectional traffic between a DNS server and GigaVUE-FM	Used to resolve Fully Qualified Domain Names (FQDNs).
68	UDP	DHCP	Bidirectional traffic between GigaVUE-FM and DHCP server	Used only if DHCP is enabled on the GigaVUE-FM appliance.
123	UDP	NTP	Bidirectional traffic between a Network Time Protocol (NTP) server and GigaVUE-FM	Used only if GigaVUE-FM is configured to use NTP.
162	UDP	SNMP	Unidirectional traffic from managed appliances to GigaVUE-FM	Used to process incoming traps sent from managed appliances to GigaVUE-FM.
2055	UDP	NetFlow	Unidirectional traffic from managed Appliances to GigaVUE-FM	Used for receiving NetFlow traffic.

NOTE: For migration and provisioning purposes, Port 902 must be open between the VMware vCenter server and the VMware ESXi hosts. Else, the FM bulk deploy fails.

Port Number	Protocol	Service	Traffic Direction	Description
902	TCP/UDP	ESXi Host	Bidirectional traffic between VMware vCenter and ESXi hosts	Used for sending data from vCenter Server to the ESXi host. The ESXi host uses this port to send regular heartbeat to the vCenter Server system.

D Health Status

This appendix provides the health status information of the following:

Node Health Status

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

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

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

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

A node is considered unhealthy if:

- at least 50% of the cards are down
- at least 50% of the ports in a card are down
- at least 1 power module is down
- the average memory usage over the past one hour is more than 70%

- the CPU load per core is more than 50% overloaded

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

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

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

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

Port Health Status

The health of a port is determined by multiple factors:

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

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

In the **Administration > System > Traffic Health Thresholds**, the packet drops and errors are enabled by default for computing the health status of a port. For information about setting the traffic health thresholds, refer to *Traffic Health Thresholds* on page 1299.

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

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

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

Inline Networks

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

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

Map Health Status

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

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

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

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

Color	Health Status
Gray	Unknown

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

GigaSMART Map Health Status

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

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

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

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

GigaSMART Group Health Status

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

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

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

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

For information about setting traffic health thresholds, refer to *Traffic Health Thresholds* on page 1299. All threshold types are enabled by default for computing the health status of the GigaSMART engine ports.

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

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

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

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

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

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

vPort Health Status

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

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

GigaSMART Operations Health Status

GigaSMART Operation consists of one or more advanced processing applications.

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

IP Interfaces Health Status

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

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

For information on how to calculate the map health, refer to *Map Health Status* on page 1363.

Flow Health Status

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

- Health of the priority maps in the flow
- Health of the maps that constitute the priority map set

- Gigamon Discovery or Manual links

Priority Map Set Health

A priority map set consists of more than one map configured with the same source ports in priority order. The health of such map is determined by the aggregated health of the constituted maps. For information about how map health is computed, refer to *Map Health Status* on page 1363.

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

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

Flow Health Computation

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

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

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

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

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

E Additional Sources of Information

This appendix provides additional sources of information. Refer to the following sections for details:

Documentation

Gigamon provides additional documentation for GigaVUE-FM on the Gigamon Customer Portal.

Document	Summary
GigaVUE-FM Release Notes	Summarizes new features and known issues in this release for the respective product.
GigaVUE-VM Release Notes	
GigaSECURE Cloud Release Notes	
GigaVUE-VM Configuration Guide	
GigaSECURE Cloud for AWS Configuration Guide	Provides instructions on configuring the GigaSECURE Cloud components and setting up traffic monitoring sessions for the respective Cloud platform.
GigaSECURE Cloud for Azure Configuration Guide	
GigaSECURE Cloud for OpenStack Configuration Guide	
REST API Getting Started Guide	Introduction to the Application Program Interfaces (APIs) for GigaVUE-FM and provides an overview of these REST APIs, basic work flows, and use cases. The APIs are implemented with the Representational State Transfer (REST) architecture. (Deprecation announcement: This has not been updated since 5.4. The content will be merged into the GigaVUE-FM User's Guide in a subsequent release.)

Documentation Feedback

To send feedback and report issues in our documentation, complete the short survey at the following link:

https://www.surveymonkey.com/r/gigamondocumentationfeedback

Contacting Technical Support

See https://www.gigamon.com/support-and-services/contact-support for Technical Support hours and contact information. You can also email Technical Support at support@gigamon.com.

Contacting Sales

Use the following information to contact sales:

Telephone: +1.408.831.4025

Sales: inside.sales@gigamon.com

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 Gigamon Community

The Gigamon 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 Gigamon 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.
- Submit and vote on feature enhancements and share product feedback. (Customers only)
- Open support tickets (Customers only)

- Download the latest product updates and documentation (Customers only)

The Gigamon 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