



V SERIES

GigaSECURE[®] Cloud for OpenStack Configuration Guide

Version 5.6

Document Version: 2.0 (*Change Notes*)

COPYRIGHT

Copyright © 2019 Gigamon Inc. All Rights Reserved. No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language in any form or by any means without Gigamon's written permission.

TRADEMARK ATTRIBUTIONS

Copyright © 2019 Gigamon Inc. All rights reserved. Gigamon and the Gigamon logo are trademarks of Gigamon in the United States and/or other countries. Gigamon trademarks can be found at www.gigamon.com/legal-trademarks. All other trademarks are the trademarks of their respective owners.

DOCUMENT REVISION – 4/12/19

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 | 04/12/2019 | Removed the following section: <ul style="list-style-type: none">Shared Controller and GigaVUE V Series Node Configuration |

Contents

| | |
|--|-----------|
| About this Guide | 5 |
| Audience | 5 |
| Licensing Information | 5 |
| Obtaining a New License | 6 |
| Applying Licensing | 6 |
| Chapter 1 Overview | 7 |
| Introduction to the GigaSECURE® Cloud for OpenStack Cloud | 7 |
| Traffic Capturing Mechanism | 8 |
| G-vTAP Agent | 8 |
| Gigamon GigaSECURE® Cloud Components | 9 |
| Chapter 2 Configuring the Components in OpenStack | 11 |
| Before You Begin | 11 |
| Supported Hypervisor | 11 |
| Minimum Compute Requirements | 12 |
| Network Requirements | 12 |
| Virtual Network Interface Cards (vNICs) | 13 |
| Security Group | 13 |
| Creating a Security Group | 15 |
| Key Pairs | 17 |
| Uploading the Images | 18 |
| Launching the GigaVUE-FM Instance | 19 |
| Initial GigaVUE-FM Configuration | 25 |
| G-vTAP Agents | 27 |
| Single vNIC Configuration | 27 |
| Multiple vNICs Configuration | 27 |
| Installing the G-vTAP Agents | 28 |
| Installing from an Ubuntu/Debian Package | 28 |
| Installing from an RPM package | 29 |
| Installing IPsec on G-vTAP Agent | 29 |
| Installing from an Ubuntu/Debian Package | 30 |
| Installing from Red Hat Enterprise Linux and CentOS | 30 |
| Installing from Red Hat Enterprise Linux and CentOS with Selinux Enabled | 31 |
| Configuring the GigaSECURE® Cloud in OpenStack | 33 |
| Pre-Configuration Checklist | 33 |
| Logging in to GigaVUE-FM | 34 |

| | |
|---|------------|
| Connecting to OpenStack | 36 |
| Configuring the G-vTAP Controllers | 40 |
| Configuring the GigaVUE V Series Controllers | 44 |
| Configuring the GigaVUE V Series Node | 46 |
| Chapter 3 Configuring Monitoring Sessions | 51 |
| Overview of Visibility Components | 51 |
| Creating Tunnel Endpoints | 55 |
| Creating a Monitoring Session | 57 |
| Creating a New Session | 57 |
| Cloning a Monitoring Session | 59 |
| Splitting a Monitoring Session | 60 |
| Creating a Map | 61 |
| Agent Pre-filtering | 66 |
| Adding Applications to the Monitoring Session | 69 |
| Sampling | 69 |
| Slicing | 70 |
| Masking | 72 |
| NetFlow | 73 |
| Deploying the Monitoring Session | 94 |
| Adding Header Transformations | 98 |
| Viewing the Statistics | 101 |
| Viewing the Topology | 104 |
| Configuring the OpenStack Settings | 108 |
| Appendix 4 Compatibility Matrix | 111 |
| GigaVUE-FM Version Compatibility | 111 |
| Supported Features in GigaVUE V Series Nodes | 111 |
| Supported Features in G-vTAP Agents | 112 |
| Appendix B Troubleshooting | 113 |
| OpenStack Connection Failed | 113 |
| Handshake Alert: unrecognized_name | 113 |
| GigaVUE V Series Node or G-vTAP Controller is Unreachable | 114 |
| Appendix C Additional Sources of Information | 115 |
| Documentation | 115 |
| Documentation Feedback | 115 |
| Contacting Technical Support | 115 |
| Contacting Sales | 116 |
| The Gigamon Community | 116 |

About this Guide

This guide describes how to install, configure, and deploy the GigaSECURE[®] Cloud for OpenStack Cloud. Use this document for instructions on configuring the GigaSECURE[®] Cloud components and setting up the traffic monitoring sessions for OpenStack.

Audience

This guide is intended for users who have basic understanding of OpenStack and OpenStack terminologies. This document expects the users to be familiar with the following Openstack terminologies that are used in this guide:

- Cloud
- Flavor
- Floating IP
- Multi-project (Multi-tenant)
- Project (Tenant)

For a detailed list of OpenStack terms and definitions, refer to the OpenStack Glossary: <https://docs.openstack.org/contributor-guide/common/glossary.html>

Licensing Information

You can purchase a license that is based on the number of TAP points. All GigaVUE-FM are available with the base option of 1 free G-vTAP. No licenses are required to activate this option.

Additional TAP points are available for purchase. [Table 0-1](#) summarizes the available packages and support features with each package.

Table 0-1: G-vTAP License Packages

| Features | FM-Base (Free-of-Charge) | 100-Pack | 250-Pack | 1000-Pack |
|--------------------|-----------------------------|-----------|-----------|------------|
| Audit, Events Logs | Yes | Yes | Yes | Yes |
| Virtual Tap Points | 10 | Up to 100 | Up to 250 | Up to 1000 |
| Trending Data | 1 Day | 1 Month | 1 Month | 1 Month |

To purchase new licenses, refer to [Obtaining a New License on page 6](#).

Obtaining a New License

Contact your Sales representative to obtain a new license for the Gigamon GigaSECURE® Cloud for OpenStack. Refer to [Contacting Sales on page 116](#).

Applying Licensing

To apply the purchased licenses, refer to the “Applying Licenses” section in the *GigaVUE-FM User’s Guide*.

Overview

This chapter introduces the Gigamon GigaSECURE® Cloud for OpenStack Cloud and the supported architecture. Refer to the following sections for details:

- [Introduction to the GigaSECURE® Cloud for OpenStack Cloud on page 7](#)
- [Gigamon GigaSECURE® Cloud Components on page 9](#)
- [Traffic Capturing Mechanism on page 8](#)

Introduction to the GigaSECURE® Cloud for OpenStack Cloud

The OpenStack software is designed for multi-tenancy (multiple projects), where a common set of physical compute and network resources are used to create project domains that provide isolation and security. Characteristics of a typical OpenStack deployment include the following:

- Projects are unaware of the physical hosts on which their instances are running.
- A project can have several virtual networks and may span across multiple hosts.

In a multi-project OpenStack cloud, where project isolation is critical, the Gigamon solution extends visibility for the project's workloads without impacting others by doing the following:

- Support project-wide monitoring domains—a project may monitor any of its instances.
- Honor project isolation boundaries—no traffic leakage from one project to any other project during monitoring.
- Monitor traffic without needing cloud administration privileges. There is no requirement to create port mirror sessions and so on.
- Monitor traffic activity of one project without adversely affecting other projects.

Traffic Capturing Mechanism

GigaSECURE® Cloud for OpenStack captures traffic in OpenStack cloud using G-vTAP agents, as described in this section.

G-vTAP Agent

A G-vTAP agent is a tiny footprint user-space agent (G-vTAP) that is deployed in a project instance. This agent mirrors the traffic from a source interface to a destination mirror interface. The mirrored traffic is then sent to the GigaVUE® V Series node. [Figure 1-1 on page 8](#) shows a high level architecture of Gigamon GigaSECURE® Cloud for OpenStack using G-vTAP agents as the source for acquiring the traffic.

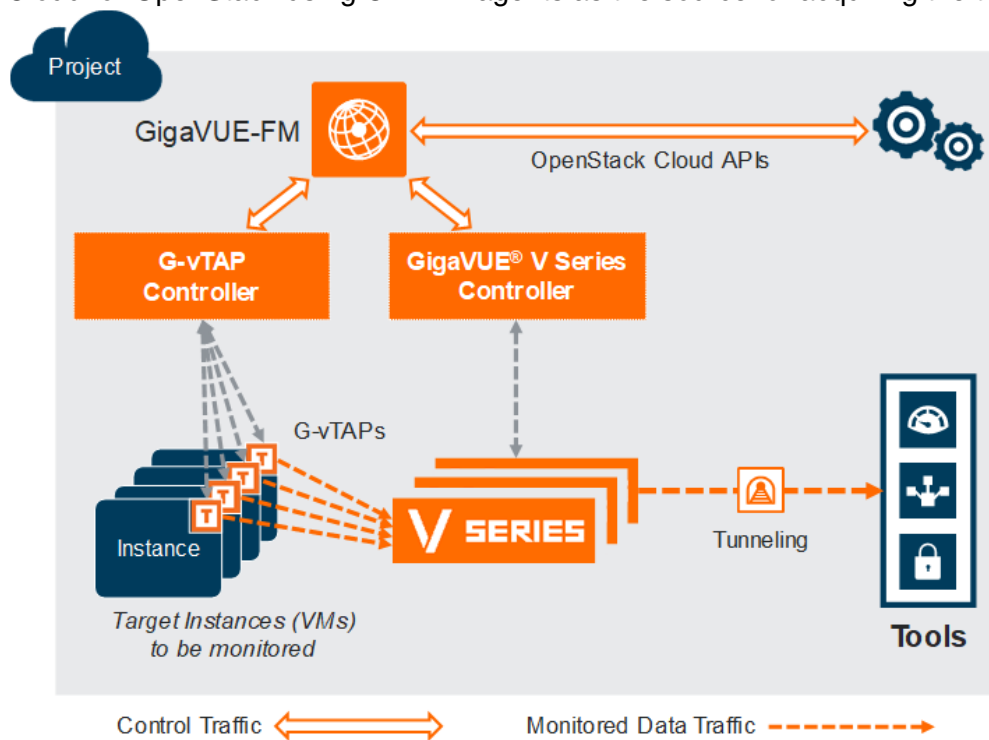


Figure 1-1: GigaSECURE® Cloud Components for OpenStack using G-vTAP

A G-vTAP agent is deployed by installing the agent in the virtual instances. When a G-vTAP agent is installed, a G-vTAP Controller must be configured in your environment. A G-vTAP Controller orchestrates the flow of mirrored traffic from G-vTAP agents to the GigaVUE V Series nodes. A single G-vTAP Controller can manage up to 100 G-vTAP agents deployed in the cloud.

By using G-vTAP agents for mirroring traffic, the monitoring infrastructure is fully contained within the virtual machine being monitored. This agent is agnostic of the underlying virtual switch. Also, the cost of monitoring a virtual machine is borne by the same virtual machine.

Gigamon GigaSECURE® Cloud Components

The GigaSECURE® Cloud for OpenStack includes the following components:

- **GigaVUE® Fabric Manager (GigaVUE-FM)** is a web-based fabric management and orchestration interface that provides a single pane of glass visibility, management, and orchestration of both the physical and virtual traffic that form the GigaSECURE® Cloud.

GigaVUE-FM can be installed on-premise or launched from an OpenStack image. GigaVUE-FM manages the configuration of the following visibility components in your OpenStack project:

- GigaVUE® V Series nodes
- GigaVUE® V Series Controllers
- G-vTAP Controllers (only if you are using G-vTAP agent as the traffic acquisition method)
- **G-vTAP Controller** manages multiple G-vTAP agents and orchestrates the flow of mirrored traffic to GigaVUE V Series nodes. GigaVUE-FM uses one or more G-vTAP Controllers to communicate with the G-vTAP agents.
- **GigaVUE® V Series Controller** manages multiple V Series nodes and orchestrates the flow of traffic from GigaVUE V Series nodes to the monitoring tools. GigaVUE-FM uses one or more GigaVUE V Series Controllers to communicate with the GigaVUE V Series nodes.
- **GigaVUE® V Series Node** is a visibility node that aggregates mirrored traffic from multiple G-vTAP agents. It applies filters, manipulates the packets using GigaSMART applications, and distributes the optimized traffic to cloud-based tools or backhaul to GigaSECURE® Cloud using GRE or VxLAN tunnels.

You can choose one of the following two options for configuring the components described above:

Table 1-1: Configuration options for Controllers and Nodes

| | |
|--|--|
| Option 1: Standard Configuration | GigaVUE V Series nodes, GigaVUE V Series controllers and G-vTAP controllers are launched in all the projects |
| Option 2: Shared Controller Configuration | <ul style="list-style-type: none">• GigaVUE V Series nodes are launched in all the projects• GigaVUE V Series controllers and G-vTAP controllers are launched in a shared project |

Configuring the Components in OpenStack

This chapter describes how to configure GigaVUE® Fabric Manager (GigaVUE-FM), G-vTAP Controllers, GigaVUE V Series Controllers, and GigaVUE V Series nodes in your OpenStack Cloud (Project). Refer to the following sections for details:

- [Before You Begin](#) on page 11
- [Uploading the Images](#) on page 18
- [Launching the GigaVUE-FM Instance](#) on page 19
- [Installing the G-vTAP Agents](#) on page 28
- [Configuring the GigaSECURE® Cloud in OpenStack](#) on page 33

Before You Begin

This section describes the requirements and prerequisites for configuring the GigaSECURE® Cloud for OpenStack. Refer to the following section for details.

- [Supported Hypervisor](#) on page 11
- [Network Requirements](#) on page 12
- [Virtual Network Interface Cards \(vNICs\)](#) on page 13
- [Security Group](#) on page 13
- [Key Pairs](#) on page 17

Supported Hypervisor

Table 2-1 lists the hypervisor with the supported versions for G-vTAP.

Table 2-1: Hypervisor for OpenStack

| Hypervisor | Version |
|------------|---|
| KVM | G-vTAP—Pike, Queens, Ocata, Newton, Mitaka, and Liberty |

Minimum Compute Requirements

In OpenStack, flavors set the vCPU, memory, and storage requirements for an image. Gigamon recommends that you create a flavor that matches or exceeds the minimum recommended requirements listed in [Table 2-2](#).

Table 2-2: Minimum Compute Requirement

| Compute Instances | vCPU | Memory | Disk Space | Description |
|---------------------|--------|--------|------------|--|
| G-vTAP Agent | 2 vCPU | 4GB | N/A | Available as rpm or debian package. Instances can have a single vNIC or dual vNICs configured for monitoring the traffic. |
| G-vTAP Controller | 1 vCPU | 4GB | 8GB | Based on the number of agents being monitored, multiple controllers will be required to scale out horizontally. |
| V Series Node | 2 vCPU | 3.75GB | 20GB | NIC 1: Monitored Network IP; Can be used as Tunnel IP NIC 2: Tunnel IP (optional) NIC 3: Management IP |
| V Series Controller | 1 vCPU | 4GB | 8GB | Based on the number of GigaVUE V Series nodes being monitored, multiple controllers will be required to scale out horizontally |
| GigaVUE-FM | 4 vCPU | 16GB | 41GB | GigaVUE-FM must be able to access the controller instance for relaying the commands. |

Network Requirements

[Table 2-3](#) lists the recommended requirements to setup the network topology.

Table 2-3: Types of Networks

| Network | Purpose |
|------------|--|
| Management | Identify the Network Interface Card (NIC) that GigaVUE-FM uses to communicate with the GigaVUE V Series nodes and controllers. |
| Data | Identify the Network Interface Card (NIC) that receives the mirrored GRE tunnel traffic from the monitored instances. This is applicable only for G-vTAP agents. |

Virtual Network Interface Cards (vNICs)

OpenStack Cloud Instances can be configured with one or more vNICs.

- **Single vNIC**—If there is only one interface configured on the instance with the G-vTAP agent, the G-vTAP agent sends the mirrored traffic out using the same interface.
- **Multiple vNICs**—If there are two or more interfaces configured on the instance with the G-vTAP agent, the G-vTAP agent monitors any number of interfaces. It provides an option to send the mirrored traffic out using any one of the interfaces or using a separate, non-monitored interface.

Security Group

A security group defines the virtual firewall rules for your instance to control inbound and outbound traffic. When you launch GigaVUE-FM, GigaVUE V Series Controllers, GigaVUE V Series nodes, and G-vTAP Controllers in your project, you add rules that control the inbound traffic to instances, and a separate set of rules that control the outbound traffic.

It is recommended to create a separate security group for each component using the rules and port numbers listed in [Table 2-4 on page 14](#).

Table 2-4 on page 14 lists the rules and port numbers for each component.

Table 2-4: Security Group Rules

| Direction | Ether Type | Protocol | Port | CIDR | Purpose |
|------------------------------------|----------------------|----------|------|--|--|
| GigaVUE-FM | | | | | |
| Inbound | HTTPS | TCP | 443 | Any IP address | Allows G-vTAP Controllers, GigaVUE V Series Controllers, and GigaVUE-FM administrators to communicate with GigaVUE-FM |
| Inbound | IPv4 | UDP | 68 | Any IP address | Allows GigaVUE-FM to communicate with DHCP server for assigning IP addresses and other related configuration information such as the subnet mask and default gateway |
| Inbound | IPv4 | UDP | 53 | Any IP address | Allows GigaVUE-FM to communicate with DNS server for resolving the host name of the cloud controller for OpenStack |
| G-vTAP Controller | | | | | |
| Inbound | IPv4 | TCP | 9900 | GigaVUE-FM IP address | Allows GigaVUE-FM to communicate with G-vTAP Controllers |
| G-vTAP Agent | | | | | |
| Inbound | IPv4 | TCP | 9901 | G-vTAP Controller IP address | Allows G-vTAP Controllers to communicate with G-vTAP agents |
| GigaVUE V Series Controller | | | | | |
| Inbound | IPv4 | TCP | 9902 | GigaVUE-FM IP address | Allows GigaVUE-FM to communicate with GigaVUE V Series Controllers |
| GigaVUE V Series node | | | | | |
| Inbound | Custom TCP Rule | TCP(6) | 9903 | GigaVUE V Series Controller IP address | Allows GigaVUE V Series Controllers to communicate with GigaVUE V Series nodes |
| GRE Traffic | | | | | |
| Inbound | Custom Protocol Rule | GRE (47) | 47 | Any IP address | Allows mirrored traffic from G-vTAP agents to be sent to GigaVUE V Series nodes using the L2 GRE or VXLAN tunnel Allows monitored traffic from GigaVUE V Series nodes to be sent to the monitoring tools using the L2 GRE or VXLAN tunnel |

NOTE:

- [Table 2-4 on page 14](#) lists only the ingress rules. Make sure the egress ports are open for communication.

- Along with the ports listed in [Table 2-4 on page 14](#), make sure the suitable ports required to communicate with Service Endpoints such as Identity, Compute, and Cloud Metadata are also open.

Creating a Security Group

To create an inbound security group for a component:

1. In OpenStack, click **Access & Security**.
2. Click the **Security Groups** tab. Refer to [Figure 2-1 on page 15](#).

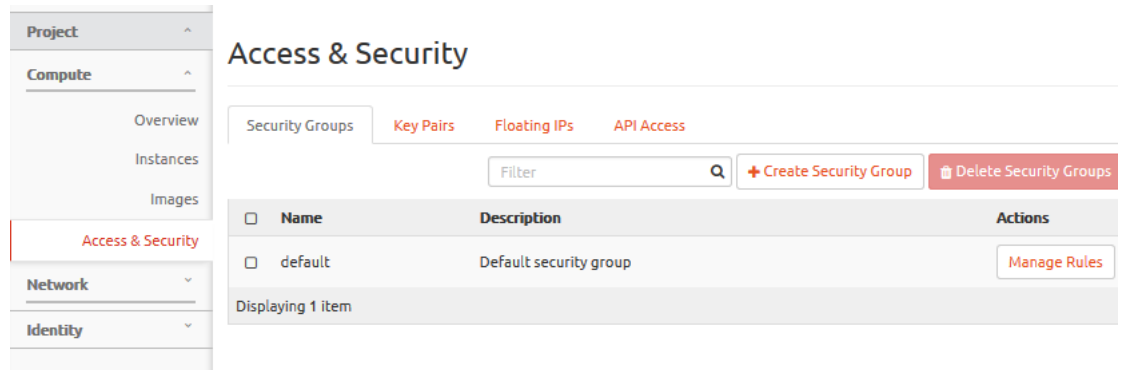


Figure 2-1: Access & Security page > Security Group Tab

3. Click **Create Security Group**.
4. Enter a name and description in the respective fields and click **Create Security Group**. Refer to [Figure 2-2 on page 15](#).

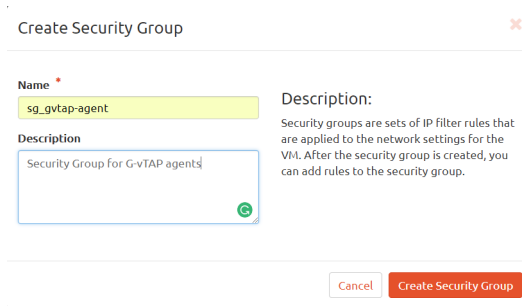


Figure 2-2: Access & Security page > Create Security Group

The security group is created and added to the Access & Security page. Refer to [Figure 2-3 on page 16](#).

Access & Security

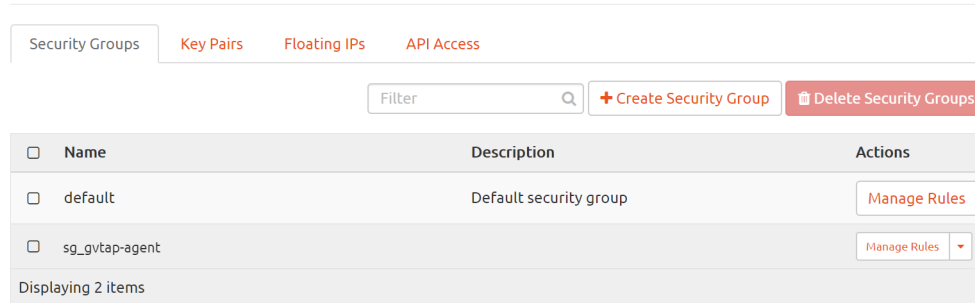


Figure 2-3: Access & Security page > Security Group Created

- For the new security group added, click **Manage Rules**. The Manage Security Group Rules page is displayed. Refer to [Figure 2-4 on page 16](#).

Manage Security Group Rules: security-group-gigamon (a0dfcc02-ed64-4614-9840-c0a08eb74724)

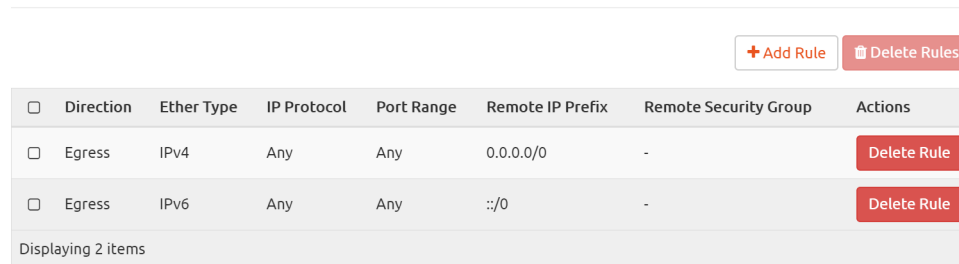


Figure 2-4: Access & Security page > Manage Rules

- Click **Add Rule**. The Add Rule page is displayed. Refer to [Figure 2-5 on page 16](#).

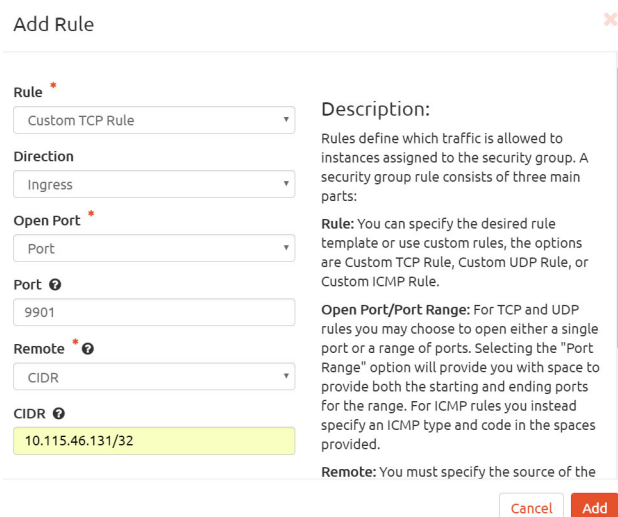


Figure 2-5: Access & Security - Add Rule

7. Enter the appropriate values in the respective fields as shown in [Figure 2-5 on page 16](#).
8. Click **Add**. The Manage Rules page is displayed with the newly added rule. Refer to [Figure 2-6 on page 17](#).

Project / Compute / Access & Security / Manage Security Group Rules

Manage Security Group Rules: sg_gvtap-agent (5e2c05fb-2cd3-42f5-9333-18f9e8beeb7e4)

[+ Add Rule](#) [Delete Rules](#)

| <input type="checkbox"/> | Direction | Ether Type | IP Protocol | Port Range | Remote IP Prefix | Remote Security Group | Actions |
|--------------------------|-----------|------------|-------------|------------|------------------|-----------------------|-----------------------------|
| <input type="checkbox"/> | Egress | IPv6 | Any | Any | ::/0 | - | Delete Rule |
| <input type="checkbox"/> | Egress | IPv4 | Any | Any | 0.0.0.0/0 | - | Delete Rule |
| <input type="checkbox"/> | Ingress | IPv4 | TCP | 9901 | 10.115.46.131/32 | - | Delete Rule |

Displaying 3 items

Figure 2-6: Access & Security - Manage Rules

9. Repeat steps 2 to 8 to create security groups for all the components. Refer to [Figure 2-7 on page 17](#).

Filter [+ Create Security Group](#) [Delete Security Groups](#)

| <input type="checkbox"/> | Name | Description | Actions |
|--------------------------|-------------------------------|--|--------------------------------|
| <input type="checkbox"/> | default | Default security group | Manage Rules |
| <input type="checkbox"/> | sg_gigavue-fm | Security Group for GigaVUE-FM | Manage Rules ▼ |
| <input type="checkbox"/> | sg_gigavue-vseries-controller | Security Group for V Series Controller | Manage Rules ▼ |
| <input type="checkbox"/> | sg_gigavue-vseries-node | Security Group for GigaVUE V Series Node | Manage Rules ▼ |
| <input type="checkbox"/> | sg_gre-traffic | Security Group for GRE Traffic | Manage Rules ▼ |
| <input type="checkbox"/> | sg_gvtap-agent | Security Group for G-vTAP Agent | Manage Rules ▼ |
| <input type="checkbox"/> | sg_gvtap-controller | Security Group for G-vTAP Controller | Manage Rules ▼ |

Displaying 7 items

Figure 2-7: Security Group Tab

Key Pairs

A key pair consists of a public key and a private key. You must create a key pair and specify the name of this key pair when you launch the G-vTAP Controllers, GigaVUE V Series nodes, and GigaVUE V Series Controllers in your instance. Then, you must provide the private key to connect to these instances. For information about creating a key pair, refer to OpenStack documentation.

Uploading the Images

First, you must fetch the images from [Gigamon Customer Portal](#) using FTP, TFTP, SCP, or other desired method and copy it to your cloud controller. After fetching the images, you must source the credentials file and then upload the qcow2 images to Glance.

For example, you can source the credentials file with admin credentials using the following command:

```
$ source admin_openrc.sh
```

To upload the qcow2 images to Glance, use one of the following commands:

```
glance image-create --disk-format qcow2 --visibility public
--container-format bare --progress -name
gigamon-gigavue-vseries-cntlr-1.x-x -file
gigamon-gigavue-vseries-cntlr-1.x-x.qcow2
```

OR

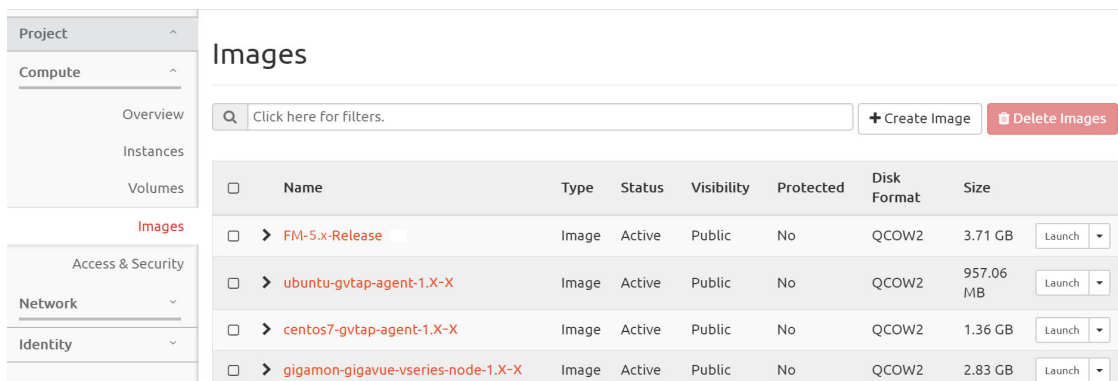
```
openstack image create --disk-format qcow2 -public --container-format
bare --file gigamon-gigavue-vseries-cntlr-1.x-x
gigamon-gigavue-vseries-cntlr-1.x-x.qcow2
```

NOTE: The 1.x-x represents the version number of the image. Enter an appropriate version in the above commands.

While uploading images to OpenStack, the names of the image files should be of the following format:

- gigamon-gigavue-vseries-node-1.x-x
- gigamon-gigavue-vseries-cntlr-1.x-x
- gigamon-gigavue-gvtap-cntlr-1.x-x

Once the images are uploaded, they are displayed under **Compute > Images**. Refer to [Figure 2-8 on page 18](#).



The screenshot shows the OpenStack Images page. On the left is a navigation sidebar with 'Project' selected, and 'Compute' expanded to show 'Overview', 'Instances', and 'Volumes'. The 'Images' link is highlighted in red. Below the sidebar are 'Access & Security', 'Network', and 'Identity' sections. The main content area is titled 'Images' and contains a search bar with the text 'Click here for filters.', a '+ Create Image' button, and a 'Delete Images' button. Below this is a table with columns: Name, Type, Status, Visibility, Protected, Disk Format, and Size. The table lists four images: 'FM-5.x-Release' (3.71 GB), 'ubuntu-gvtap-agent-1.X-X' (957.06 MB), 'centos7-gvtap-agent-1.X-X' (1.36 GB), and 'gigamon-gigavue-vseries-node-1.X-X' (2.83 GB). Each row has a 'Launch' button with a dropdown arrow.

| <input type="checkbox"/> | Name | Type | Status | Visibility | Protected | Disk Format | Size | |
|--------------------------|--------------------------------------|-------|--------|------------|-----------|-------------|-----------|----------|
| <input type="checkbox"/> | > FM-5.x-Release | Image | Active | Public | No | QCOW2 | 3.71 GB | Launch ▾ |
| <input type="checkbox"/> | > ubuntu-gvtap-agent-1.X-X | Image | Active | Public | No | QCOW2 | 957.06 MB | Launch ▾ |
| <input type="checkbox"/> | > centos7-gvtap-agent-1.X-X | Image | Active | Public | No | QCOW2 | 1.36 GB | Launch ▾ |
| <input type="checkbox"/> | > gigamon-gigavue-vseries-node-1.X-X | Image | Active | Public | No | QCOW2 | 2.83 GB | Launch ▾ |

Figure 2-8: Instances Uploaded in Images Page

Launching the GigaVUE-FM Instance

To launch the GigaVUE-FM instance inside the cloud:

1. Login to Horizon.
2. From the Horizon GUI, select the appropriate project, and select **Compute > Images**. The list of existing images is displayed, as shown in [Figure 2-9 on page 19](#).

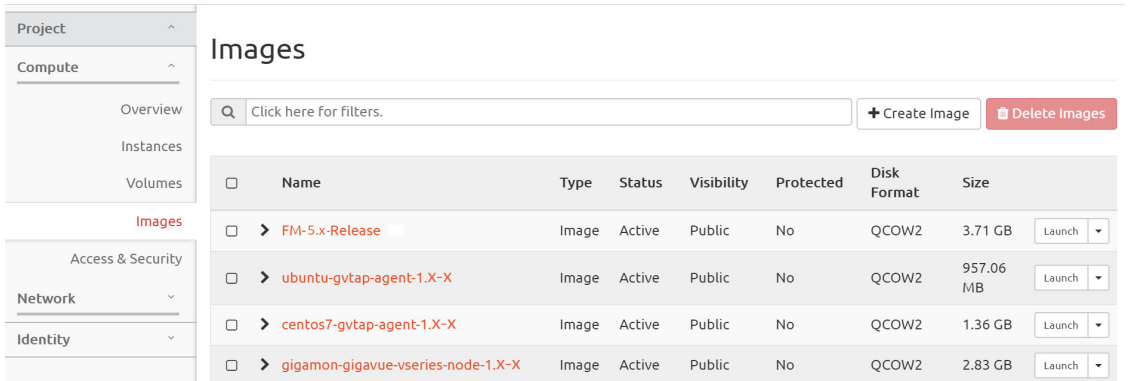


Figure 2-9: List of Available Images

3. Select the GigaVUE-FM image and click **Launch**. The Launch Instance dialog box is displayed. Refer to [Figure 2-10 on page 19](#).

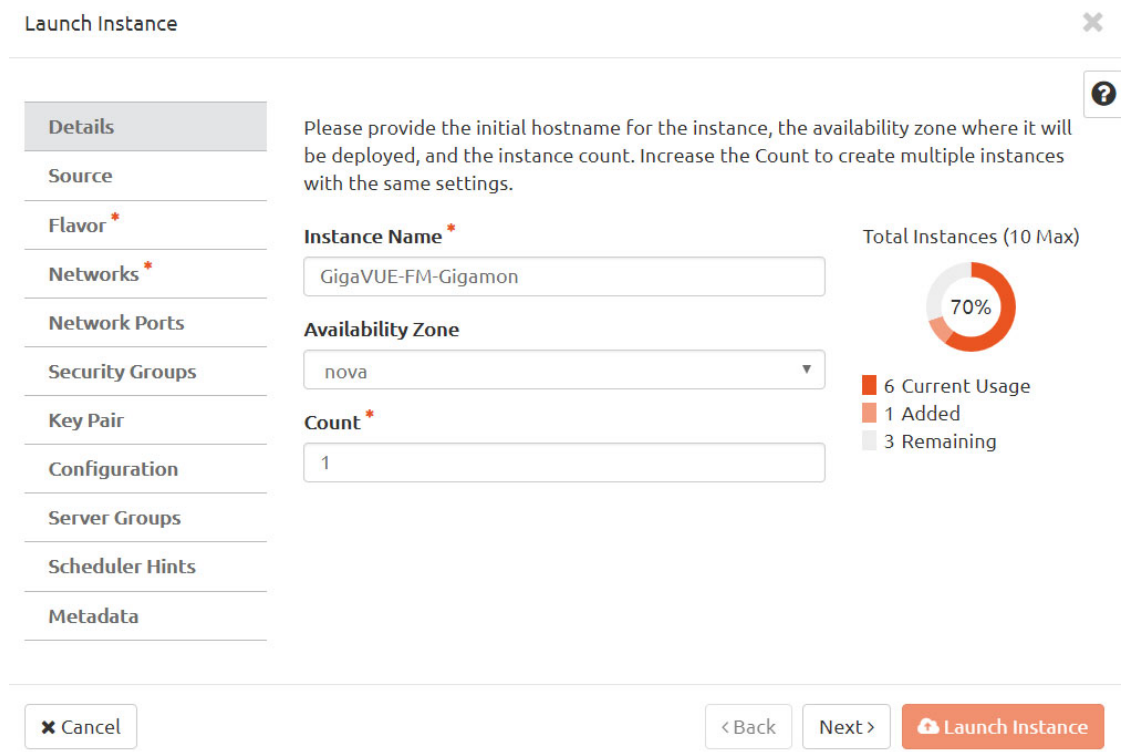


Figure 2-10: Launch Instance Page

4. In the **Details** tab, enter the information as shown in [Figure 2-10 on page 19](#).
 - **Instance Name**—Enter the initial hostname for the instance.
 - **Availability Zone**—Select the availability zone where the image will be deployed.
 - **Count**—Enter the number of instances to be launched.
5. Click **Next**.
6. In the **Source** tab, select the GigaVUE-FM image source file from the Available list and then click +(Plus). The selected GigaVUE-FM image is displayed under Allocated. Click **Next**. Refer to [Figure 2-11 on page 20](#).

Launch Instance ✕

- Details
- Source
- Flavor *
- Networks *
- Network Ports
- Security Groups
- Key Pair
- Configuration
- Server Groups
- Scheduler Hints
- Metadata

Instance source is the template used to create an instance. You can use a snapshot of an existing instance, an image, or a volume (if enabled). You can also choose to use persistent storage by creating a new volume.

Select Boot Source

Image
▼

Create New Volume

Yes

No

Allocated

| Name | Updated | Size | Type | Visibility | |
|-------------|-----------------|---------|-------|------------|---|
| > FM-5.x-GA | 8/17/17 4:02 PM | 3.63 GB | qcow2 | Public | - |

▼ Available 18 Select one

| Name | Updated | Size | Type | Visibility | |
|-----------------|----------------|-----------|-------|------------|---|
| > WireShark-RDP | 9/5/17 1:39 PM | 1.77 GB | qcow2 | Public | + |
| > spirent-v4.64 | 9/5/17 1:06 PM | 564.25 MB | qcow2 | Public | + |

✕ Cancel

< Back

Next >

Launch Instance

Figure 2-11: Source tab

7. In the **Flavor** tab, select a flavor with a specific compute, memory, and storage capacity from the Available list and then click +(plus sign). The selected

GigaVUE-FM flavor is displayed under Allocated. Click **Next**. Refer to [Figure 2-12 on page 21](#).

Launch Instance ✕

- Details
- Source
- Flavor
- Networks
- Network Ports
- Security Groups
- Key Pair
- Configuration
- Server Groups
- Scheduler Hints
- Metadata

Flavors manage the sizing for the compute, memory and storage capacity of the instance.

Allocated

| Name | VCPUS | RAM | Total Disk | Root Disk | Ephemeral Disk | Public | |
|-------------|-------|------|------------|-----------|----------------|--------|---|
| > FM_Flavor | 2 | 4 GB | 41 GB | 41 GB | 0 GB | Yes | - |

▼ Available 4 Select one

| Name | VCPUS | RAM | Total Disk | Root Disk | Ephemeral Disk | Public | |
|-------------|-------|--------|------------|-----------|----------------|--------|---|
| > m1.nano | 1 | 64 MB | 1 GB | 1 GB | 0 GB | Yes | + |
| > m1.small | 1 | 512 MB | 20 GB | 20 GB | 0 GB | Yes | + |
| > m1.medium | 2 | 4 GB | 40 GB | 40 GB | 0 GB | Yes | + |
| > m1.large | 4 | 8 GB | 80 GB | 80 GB | 0 GB | Yes | + |

✕ Cancel
< Back
Next >
Launch Instance

Figure 2-12: Flavor tab

8. In the **Networks** tab, select the specific network for the GigaVUE-FM instance from the Available list and then click +(plus sign). For information about the

requirements, refer to [Network Requirements on page 12](#). The selected network is displayed under Allocated. Click **Next**. Refer to [Figure 2-13 on page 22](#).

Launch Instance

Details

Source

Flavor

Networks

Network Ports

Security Groups

Key Pair

Configuration

Server Groups

Scheduler Hints

Metadata

Networks provide the communication channels for instances in the cloud.

Select networks from those listed below.

▼ Allocated 1

| Network | Subnets Associated | Shared | Admin State | Status |
|----------|--------------------|--------|-------------|----------|
| 1 > mgmt | mgmt-subnet | No | Up | Active - |

Select at least one network

▼ Available 2

Click here for filters.

| Network | Subnets Associated | Shared | Admin State | Status |
|----------|--------------------|--------|-------------|----------|
| > tunnel | tunnel-subnet | No | Up | Active + |
| > data | data-subnet | No | Up | Active + |

✕ Cancel

< Back

Next >

Launch Instance

Figure 2-13: Networks tab

9. In the **Network Ports** tab, click **Next** again.
10. In the **Security Groups** tab, select the appropriate security group for the GigaVUE-FM instance from the Available list and then click +(plus sign). For information about the security groups, refer to [Security Group on page 13](#). The

selected security group is displayed under Allocated. Click **Next**. Refer to [Figure 2-14 on page 23](#).

Launch Instance

Source

Flavor

Networks

Network Ports

Security Groups

Key Pair

Configuration

Server Groups

Scheduler Hints

Metadata

▼ Allocated ¹

| Name | Description | |
|-----------|------------------------|---|
| > default | Default security group | - |

▼ Available ⁶ Select one or more

Q Click here for filters.

| Name | Description | |
|---------------------------------|--|----------|
| > sg_gigavue-vseries-node | Security Group for GigaVUE V Series Node | + |
| > sg_gvtap-controller | Security Grou for G-vTAP Controller | + |
| > sg_gre-traffic | Security Group for GRE Traffic | + |
| > sg_gvtap-agent | Security Group for G-vTAP Agent | + |
| > sg_gigavue-fm | Security Group for GigaVUE-FM | + |
| > sg_gigavue-vseries-controller | Security Group for V Series Controller | + |

✕ Cancel < Back Next > Launch Instance

Figure 2-14: Security Groups tab

11. In the **Key Pair** tab, select the existing key pair from the Available list and then click + (plus sign) or create a new key pair. For information about the key pairs, refer to

[Key Pairs on page 17](#). The selected key pair is displayed under Allocated. Click **Next**. Refer to [Figure 2-15 on page 24](#).

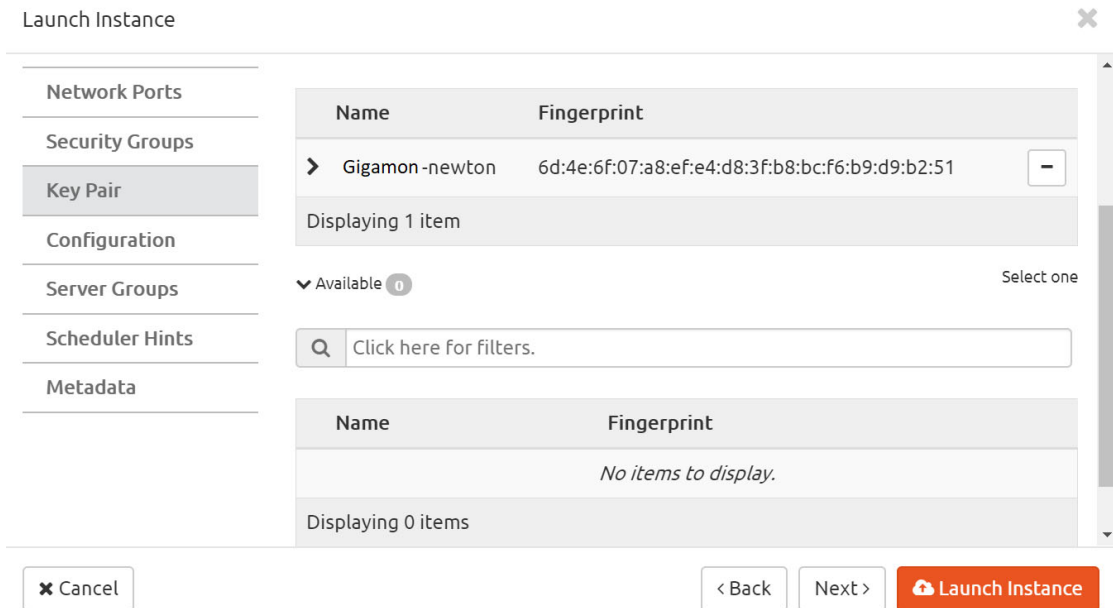


Figure 2-15: Key Pair tab

- In the Configuration, Server Groups, Scheduler Hints, and Metadata tabs, click **Next**.
- Click **Launch Instance**. The GigaVUE-FM instance takes a few minutes to fully initialize. Refer to [Figure 2-16 on page 24](#).

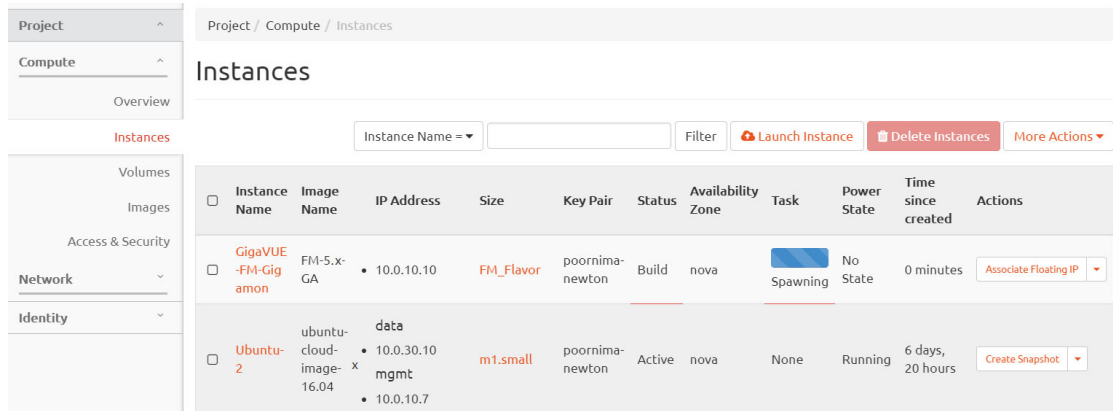


Figure 2-16: Instance Launched

- During the initial boot-up sequence, click **Allocate Floating IP**. The Manage Floating IP Associations dialog box is displayed. Refer to [Figure 2-17 on page 25](#).

Figure 2-17: Assigning Floating Address

- In the Manage Floating IP Associations dialog box, enter the following information as shown in [Figure 2-17 on page 25](#):
 - From the IP Address drop-down list, select an IP address.
 - From the Port to be associated drop-down list, select an appropriate port for the GigaVUE-FM instance.
- Click **Associate**. The Floating IP is displayed in the IP Address column. Refer to [Figure 2-18 on page 25](#).

| | | | | | | | | | | | | | | | | | | | |
|--------------------------|-----------|-----------|------------------|-----------|-----------|--------|------|------|---------|--|--|--|--|--|--|--|--|--|--|
| <input type="checkbox"/> | GigaVUE-F | | • 10.0.10.10 | | | | | | | | | | | | | | | | |
| | M-Gigamo | FM-5.x-GA | Floating IPs: | FM_Flavor | poornima- | Active | nova | None | Running | | | | | | | | | | |
| | n | | • 10.210.219.144 | | | | | | | | | | | | | | | | |

Figure 2-18: Floating IP of GigaVUE-FM Instance

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 is a one time activity that must be performed for each GigaVUE-FM instance deployed.

- In the Instances page, click the GigaVUE-FM instance name. The GigaVUE-FM instance Overview tab is displayed by default.

2. Click the **Console** tab. Refer to [Figure 2-19 on page 26](#).

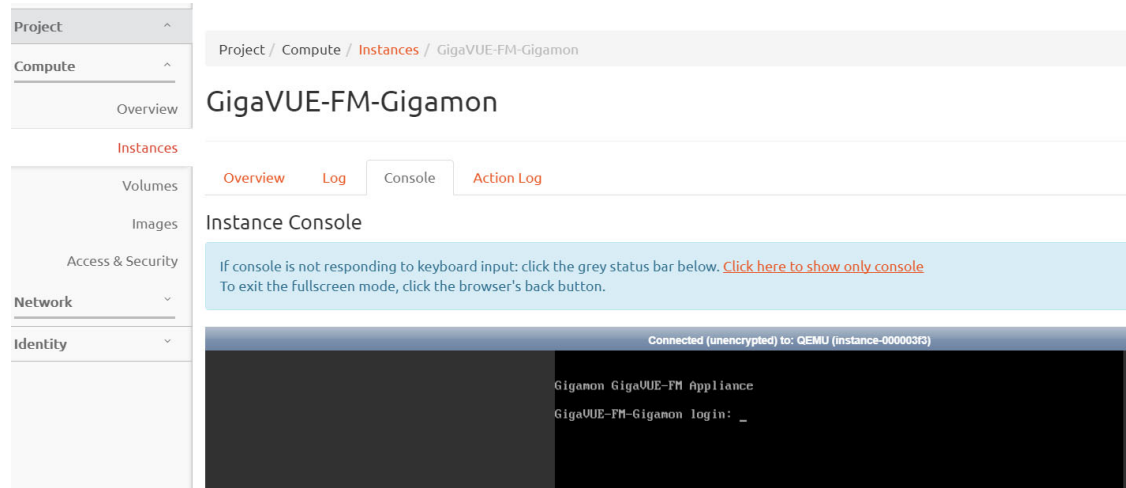


Figure 2-19: GigaVUE-FM Instance Console

3. Log in as admin with password as admin123A! The jump start configuration for GigaVUE-FM starts automatically.
4. For the hostname, enter 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. Refer to [Figure 2-20 on page 26](#).




Figure 2-20: GigaVUE-FM Jump Start Configuration

5. To enable DHCP on eth0 interface, type yes and press enter.

NOTE: DHCP must be enabled for GigaVUE-FM.

6. For admin password, enter an appropriate password for your environment or just press Enter to leave the password unchanged.

NOTE: For admin user, GigaVUE-FM requires a password.

7. For additional domain name server IP addresses, enter the address of any additional name servers required. The names must be provided as a set of IP addresses with spaces.
8. For additional DNS domains, enter the DNS domain name.
9. To enable NTP, type yes.
10. Enter the NTP Server IP address and the NTP server version.
11. Click **Enter** to save the configuration and exit the console.

G-vTAP Agents

A G-vTAP agent is a tiny footprint user-space agent (G-vTAP) that is deployed on each instance that you want to monitor. This agent mirrors the selected traffic from a source interface to a destination mirror interface. The mirrored traffic is encapsulated using GRE or VXLAN tunneling and then sent to the GigaVUE® V Series node.

A source interface can be configured with one or more ENIs. While configuring a source interface, you can specify the direction of the traffic to be monitored in the instance. The direction of the traffic can be egress or ingress or both.

Single vNIC Configuration

A single vNIC acts both as the source and the destination interface. A G-vTAP agent with a single vNIC configuration lets you monitor the ingress or egress traffic from the vNIC. The monitored traffic is sent out using the same vNIC.

For example, assume that there is only one interface eth0 in the monitoring instance. In the G-vTAP configuration, you can configure eth0 as the source and the destination interface, and specify both egress and ingress traffic to be selected for monitoring purpose. The egress and ingress traffic from eth0 is mirrored and sent out using the same interface.

Using a single vNIC as the source and the destination interface can sometimes cause increased latency in sending the traffic out from the instance.

Multiple vNICs Configuration

A G-vTAP agent lets you configure multiple vNICs. One or many vNICs can be configured as the source interface. The monitored traffic can be sent out using any one of the vNICs or using a separate, non-monitored vNIC.

For example, assume that there is eth0 and eth1 in the monitoring instance. In the G-vTAP agent configuration, eth0 can be configured as the source interface and egress traffic can be selected for monitoring purpose. The eth1 interface can be configured as the destination interface. So, the mirrored traffic from eth0 is sent to eth1. From eth1, the traffic is sent to the GigaVUE V Series node.

Installing the G-vTAP Agents

This is applicable only if you are using G-vTAP agent as the source of acquiring traffic. You must have sudo/root access to edit the G-vTAP agent configuration file. Before installing the G-vTAP agents, you must have launched the GigaVUE-FM instance.

You can install the G-vTAP agents either from Debian or RPM packages as follows:

- [Installing from an Ubuntu/Debian Package](#)
- [Installing from an RPM package](#)

Installing from an Ubuntu/Debian Package

To install from a Debian package:

1. Download the latest version of G-vTAP Agent Debian (.deb) package from the [Gigamon Customer Portal](#).
2. Copy this package to your instance. Install the package with root privileges, for example:

```
ubuntu@ip-10-0-0-246:~$ ls gvtap-agent_1.x-x_amd64.deb
ubuntu@ip-10-0-0-246:~$ sudo dpkg -i
gvtap-agent_1.x-x_amd64.deb
```

NOTE: The 1.x-x represents the version number of the G-vTAP agent. Enter the appropriate version in the configuration file.

3. Once the G-vTAP package is installed, modify the file `/etc/gvtap-agent/gvtap-agent.conf` to configure and register the source and destination interfaces.

The file contains an example, which you can use by uncommenting the last two lines. The following example registers `eth0` as the mirror source for both ingress and egress traffic and `eth1` as the destination for this traffic:

[Example 1—Configuration example to monitor ingress and egress traffic at interface eth0 and use the same interface to send out the mirrored packets](#)

```
# eth0 mirror-src-ingress mirror-src-egress mirror-dst
```

[Example 2—Configuration example to monitor ingress and egress traffic at interface eth0 and use eth1 to send out the mirrored packets](#)

```
# eth0 mirror-src-ingress mirror-src-egress
# eth1 mirror-dst
```

[Example 3—Configuration example to monitor ingress and egress traffic at interface eth0 and eth 1; use eth1 to send out the mirrored packets](#)

```
# eth0 mirror-src-ingress mirror-src-egress
# eth1 mirror-src-ingress mirror-src-egress mirror-dst
```

4. Save the file.
5. Reboot the instance.

The instance should have two interfaces. The G-vTAP agent status will be displayed as running. Check the status using the following command:

```
ubuntu@ip-10-0-0-246:~$ sudo service gvtap-agent status
G-vTAP Agent is running
```

Installing from an RPM package

To install from an RPM (.rpm) package on a Redhat, Centos, or other RPM-based system:

1. Download the G-vTAP Agent RPM (.rpm) package from the [Gigamon Customer Portal](#).
2. Copy this package to your instance. Install the package with root privileges, for example:

```
[user@ip-10-0-0-214 ~]$ ls
gvtap-agent_1.x-x_x86_64.rpm
[user@ip-10-0-0-214 ~]$ sudo rpm -i
gvtap-agent_1.x-x_x86_64.rpm
```

NOTE: The 1.x-x represents the version number of the G-vTAP agent. Enter the appropriate version in the configuration file.

3. Modify the file /etc/gvtap-agent/gvtap-agent.conf to configure and register the source and destination interfaces.

The file contains an example, which you can use by uncommenting the last two lines. The following example registers the eth0 as the mirror source for both ingress and egress traffic and registers eth1 as the destination for this traffic as follows:

[Example 1—Configuration example to monitor ingress and egress traffic at interface eth0 and use the same interface to send out the mirrored packets](#)

```
# eth0 mirror-src-ingress mirror-src-egress mirror-dst
```

[Example 2—Configuration example to monitor ingress and egress traffic at interface eth0 and use eth1 to send out the mirrored packets](#)

```
# eth0 mirror-src-ingress mirror-src-egress
# eth1 mirror-dst
```

[Example 3—Configuration example to monitor ingress and egress traffic at interface eth0 and eth 1; use eth1 to send out the mirrored packets](#)

4. Save the file.
5. Reboot the instance.

Check the status with the following command:

```
[user@ip-10-0-0-214 ~]$ sudo service gvtap-agent status
G-vTAP Agent is running
```

6. Save the G-vTAP agent running on an instance as an image. Install more number of G-vTAP agents on the deployed instances as needed.

Installing IPsec on G-vTAP Agent

IPsec can be used to establish a secure connection between G-vTAP agents and GigaVUE V series nodes. If IPsec is used to establish a secure connection, then you must install IPsec on G-vTAP agent instances.

NOTE: Secure Tunnel configuration is supported only on the following operating systems:

- CentOS
- Red Hat Linux
- Ubuntu

To install IPsec on G-vTAP agent you need the following files:

- **StrongSwan binary installer TAR file:** The TAR file contains strongSwan binary installer for different platforms. Each platform has its own TAR file. Refer to <https://www.strongswan.org/> for more details.
- **IPsec package file:** The package file includes the following:
 - CA Certificate
 - Private Key and Certificate for G-vTAP Agent
 - IPsec configurations

Refer to the following sections for installing IPsec on G-vTAP Agent:

- [Installing from an Ubuntu/Debian Package on page 30](#)
- [Installing from Red Hat Enterprise Linux and CentOS on page 30](#)
- [Installing from Red Hat Enterprise Linux and CentOS with Selinux Enabled on page 31](#)

Installing from an Ubuntu/Debian Package

1. Launch the G-vTAP agent AMI.
2. Copy the G-vTAP package files and strongSwan TAR file to the G-vTAP agent:

- [strongswan5.3.5-1ubuntu3.8_amd64-deb.tar.gz](#)
- [gvtap-agent_1.6-1_amd64.deb](#)
- [gvtap-ipsec_1.6-1_amd64.deb](#)

3. Install the G-vTAP agent package file:

```
sudo dpkg -i gvtap-agent_1.6-1_amd64.deb
```

4. Edit gvtap-agent.conf file to configure the required interface as source/destination for mirror:

```
eth0# mirror-src-ingress mirror-src-egress mirror-dst
sudo /etc/init.d/gvtap-agent restart
```

5. Install strongSwan:

```
tar -xvf strongswan5.3.5-1ubuntu3.8_amd64-deb.tar.gz
cd strongswan-5.3.5-1ubuntu3.8_amd64/
sudo sh ./swan-install.sh
```

6. Install IPsec package:

```
sudo dpkg -i gvtap-ipsec_1.6-1_amd64.deb
```

Installing from Red Hat Enterprise Linux and CentOS

1. Launch RHEL/CentOS agent AMI image.

2. Copy the following package files and strongSwan TAR files to the G-vTAP agent:

- [strongswan-5.7.1-1.el7.x86_64.tar.gz](#) for rhel7/centOS7
- [gvtap-agent_1.6-1_x86_64.rpm](#)
- [gvtap-ipsec_1.6-1_x86_64.rpm](#)

3. Install G-vTAP agent package:

```
sudo rpm -ivh gvtap-agent_1.6-1_x86_64.rpm
```

4. Edit gvtap-agent.conf file to configure the required interface as source/destination for mirror:

```
# eth0 mirror-src-ingress mirror-src-egress mirror-dst
# sudo /etc/init.d/gvtap-agent restart
```

5. Install strongSwan:

```
tar -xvf strongswan-5.7.1-1.el7.x86_64.tar.gz
cd strongswan-5.7.1-1.el7.x86_64
sudo sh ./swan-install.sh
```

6. Install IPsec package:

```
sudo rpm -i gvtap-ipsec_1.6-1_x86_64.rpm
```

NOTE: You must install IPsec package after installing StrongSwan.

Installing from Red Hat Enterprise Linux and CentOS with Selinux Enabled

1. Launch the RHEL/CentOS agent AMI image.

2. Copy package files and strongSwan TAR file to G-vTAP agent.

- [strongswan-5.7.1-1.el7.x86_64.tar.gz](#) for rhel7/centOS7
- [gvtap-agent_1.6-1_x86_64.rpm](#)
- [gvtap-ipsec_1.6-1_x86_64.rpm](#)
- gvtap.te and gvtap_ipsec.te files (type enforcement files)

3. checkmodule -M -m -o gvtap.mod gvtap.te

```
semodule_package -o gvtap.pp -m gvtap.mod
```

```
sudo semodule -i gvtap.pp
```

4. checkmodule -M -m -o gvtap_ipsec.mod gvtap_ipsec.te

```
semodule_package -o gvtap_ipsec.pp -m gvtap_ipsec.mod
```

```
sudo semodule -i gvtap_ipsec.pp
```

5. Install G-vTAP agent package:

```
sudo rpm -ivh gvtap-agent_1.6-1_x86_64.rpm
```

6. Edit gvtap-agent.conf file to configure the required interface as source/destination for mirror:

```
# eth0 mirror-src-ingress mirror-src-egress mirror-dst
# sudo /etc/init.d/gvtap-agent restart
```

7. Install strongSwan:

```
tar -xvf strongswan-5.7.1-1.el7.x86_64.tar.gz
cd strongswan-5.7.1-1.el7.x86_64
sudo sh ./swan-install.sh
```

8. Install IPsec package:

```
sudo rpm -i gvtap-ipsec_1.6-1_x86_64.rpm
```


Configuring the GigaSECURE[®] Cloud in OpenStack

First, you must establish a connection between GigaVUE-FM and your OpenStack environment. Then, GigaVUE-FM lets you launch the G-vTAP Controllers or V Series Controllers and V Series nodes in the specified project.

Pre-Configuration Checklist

Table 2-5 on page 33 provides information that you would need while launching the visibility components using GigaVUE-FM. Obtaining this information will ensure a successful and efficient deployment of the GigaSECURE[®] Cloud for OpenStack:

Table 2-5: Pre-configuration Checklist

| Required Information | |
|--------------------------|--|
| <input type="checkbox"/> | Authentication URL |
| <input type="checkbox"/> | Project Name |
| <input type="checkbox"/> | Peering NOTE: Peering must be active between the projects within the same monitoring domain. This is required only when shared controller option is chosen for configuring the components. |
| <input type="checkbox"/> | Floating IP |
| <input type="checkbox"/> | Region name for the Project |
| <input type="checkbox"/> | Domain |
| <input type="checkbox"/> | SSH Key Pair |
| <input type="checkbox"/> | Networks |
| <input type="checkbox"/> | Security groups |

Logging in to GigaVUE-FM

To login to GigaVUE-FM, do the following:

1. Enter the Floating IP address of GigaVUE-FM into a browser. The GigaVUE-FM login page is displayed. Refer to [GigaVUE-FM Login Page on page 34](#).



Figure 2-21: GigaVUE-FM Login Page

NOTE: GigaVUE-FM must be able to resolve the hostname of the cloud controller for OpenStack, either through DNS or by manually adding it through the GigaVUE-FM CLI, using the `ip host <hostname> < ip address>` command.

2. Enter `admin` as the user name and `admin123A!` as the password. If the password is changed during the jump-start configuration as described in [Initial GigaVUE-FM Configuration on page 25](#), enter the changed password.
3. Click **Log In**. The GigaVUE-FM Dashboard page is displayed. Refer to [GigaVUE-FM on page 35](#).

The screenshot shows the GigaVUE-FM dashboard interface. The top navigation bar includes 'Dashboard', 'Physical', 'Virtual', 'Cloud', and 'Administration'. The left sidebar contains 'OVERVIEW' and 'PROFILES' sections. The main content area is titled 'Profile: Default' and features two status summary widgets.

Status Summary: Unhealthy Maps

| Cluster | Map Alias | Status |
|------------------------|------------------------------|-------------------------------|
| fm-upg | inline_map_1 | ● Alias(s) in_1 are unhealthy |

Status Summary: Port Links

| Port Type | Total | ● Up | ● Down | ● Disabled |
|--------------------------|-------|------|--------|------------|
| G Gateway | - | - | - | - |
| H Hybrid | 2 | 2 | - | - |
| IN Inline Network | 11 | 3 | 8 | - |
| IT Inline Tool | 3 | 3 | - | - |
| N Network | 176 | 25 | 9 | 142 |
| S Stack | 11 | 6 | 5 | - |
| T Tool | 3 | 3 | - | - |

Figure 2-22: GigaVUE-FM

Connecting to OpenStack

NOTE: In order for GigaVUE-FM to make a connection to an OpenStack tenant, GigaVUE-FM **must** be able to resolve the hostname of the OpenStack controller, even if using an IP address in the Identity URL. For example, if GigaVUE-FM is configured to use DNS, and that controller hostname is in the DNS, this will work and no further configuration will be needed. If not, then you must add a host entry to GigaVUE-FM.

You can login to GigaVUE-FM and use the CLI command: `ip host <controller-hostname> <ip-address of the controller>`. (For example: `ip host os-controller1 192.168.2.3`.) Then, add the connection to the OpenStack tenant.

To create a new connection:

1. Click **Cloud** in the top navigation.

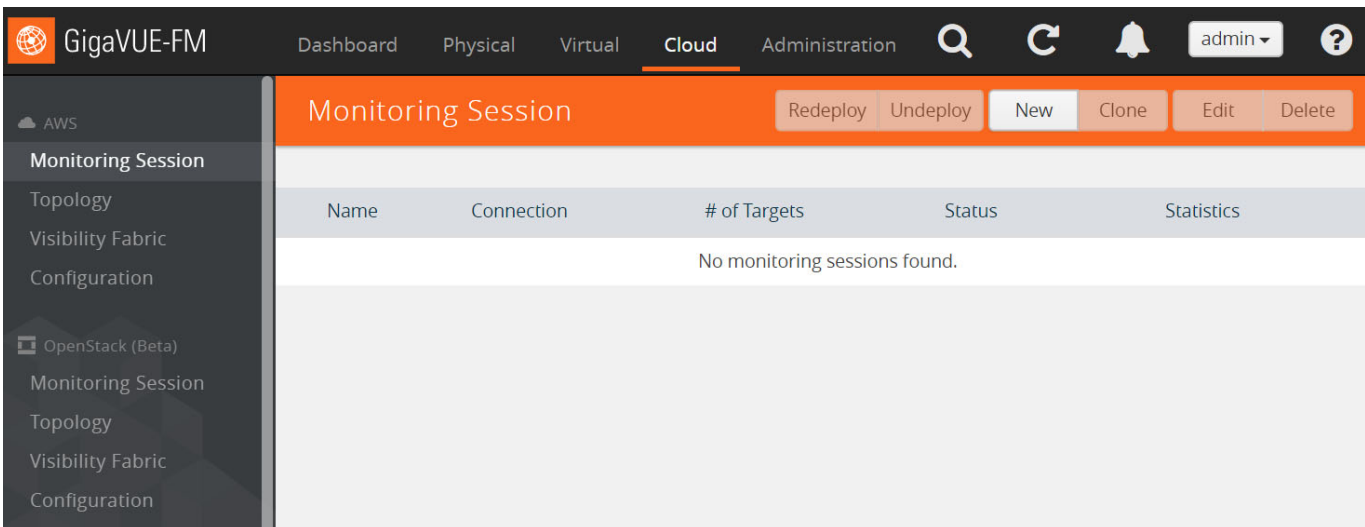


Figure 2-23: Cloud Page

2. Under OpenStack on the left navigation pane, select **Configuration** and then click the **New** drop-down menu. You can either create a new monitoring domain or a new connection.
 - If you select **Monitoring Domain**, then the **Create Monitoring Domain** dialog box is displayed. Enter the alias that is used to identify the monitoring domain.
 - If you select **Connection**, then the Openstack Connection page is displayed as shown in [Figure on page 37](#).

Connection
Save Cancel

Alias

URL

Project Name

Domain Name

Region

Username

Password

Tap Method

3. Enter or select the appropriate information as shown in [Table 2-6 on page 37](#).

Table 2-6: OpenStack Connection

| Field | Description |
|--------------------------|---|
| Alias | An alias used to identify the connection to OpenStack. |
| Monitoring Domain | An alias used to identify the monitoring domain. You can either create a new monitoring domain or select an existing monitoring domain that is already created. NOTE: Monitoring domain consists of set of connections. |

URL The authentication URL is the Keystone URL of the OpenStack cloud. This IP address must be DNS resolvable.

To get the authentication URL from the OpenStack dashboard:

- Login to Horizon.
- Go to **Compute > Access & Security**.
- Click the **API Access** tab and copy the Identity URL. Refer to [Figure 2-24 on page 37](#).

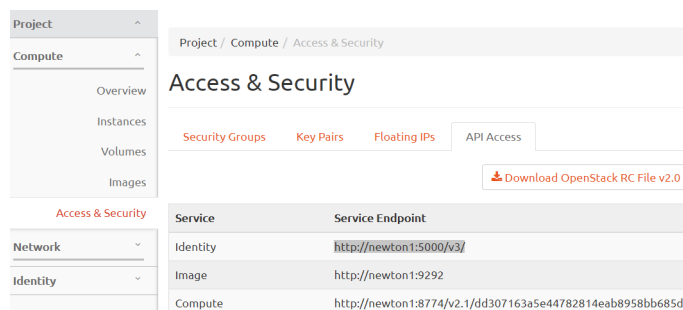


Figure 2-24: Copying the Identity URL

Paste the Identity URL into the URL field.

| | |
|---------------------|-------------------------------------|
| Project Name | The name of the Project. |
| Domain Name | The DNS domain name of the project. |

Table 2-6: OpenStack Connection

| Field | Description |
|-----------------|--|
| Region | The region where the Project resides. You can find your region by running one of these commands, depending on your OpenStack version. keystone endpoint-list or openstack endpoint list |
| Username | The user name used to connect to the OpenStack cloud. |
| Password | The password for the OpenStack cloud. |

4. Click **Save**.

If GigaVUE-FM connects to OpenStack successfully, the status is displayed as Connected in the Status column as shown in [Figure 2-25 on page 38](#). GigaVUE-FM discovers the inventory of the cloud in the background.




Figure 2-25: OpenStack Connection in GigaVUE-FM

If GigaVUE-FM fails to connect to OpenStack, an error message is displayed specifying the cause of failure.

The connection status is also displayed in **Cloud > Audit Logs**. Refer to [Figure 2-26 on page 39](#).

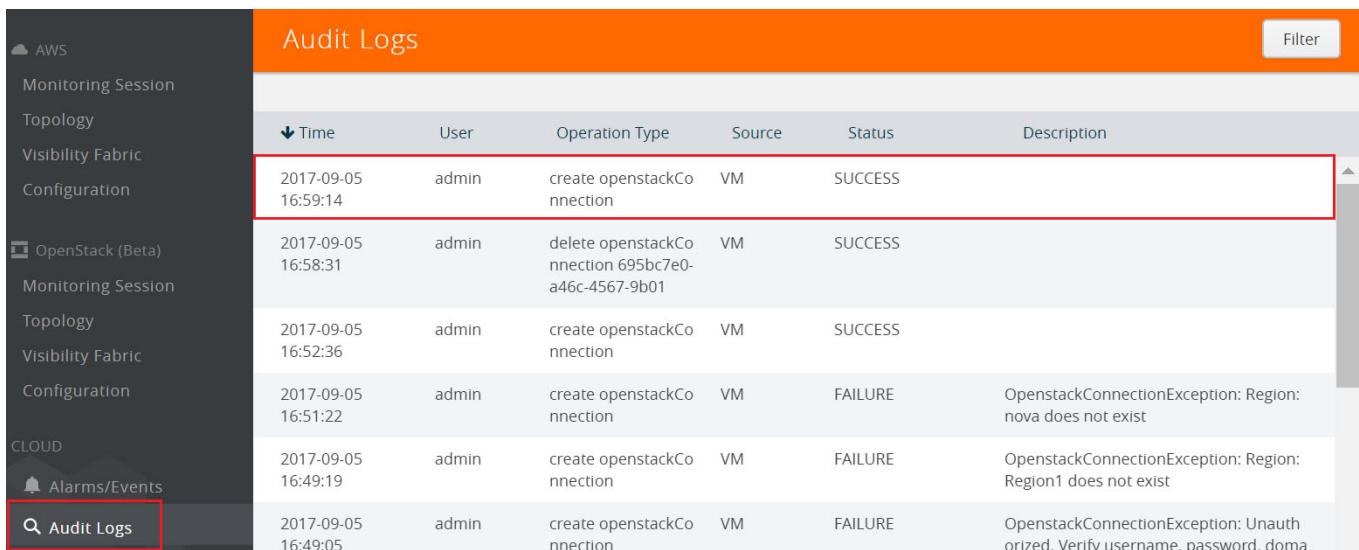


Figure 2-26: Audit Logs

The Connections page has the following controls:

| Control | Description |
|---------|---|
| Action | Allows to connect, disconnect, or rediscover connections. |
| New | Opens the page for specifying the connection details for a new connection. |
| Edit | Allows to make changes to a connection. |
| Delete | Deletes the connection. NOTE: Deleting a connection destroys all GigaVUE V Series Nodes, G-vTAP Controllers, and the virtual maps on the project. |

Configuring the G-vTAP Controllers

Only if G-vTAP agents are used for capturing traffic, then the G-vTAP Controllers must be configured in the OpenStack cloud. If TaaS is used for capturing the traffic, then skip to [Configuring the GigaVUE V Series Controllers on page 44](#).

A G-vTAP Controller manages multiple G-vTAP agents and orchestrates the flow of mirrored traffic to GigaVUE V Series nodes.

A G-vTAP Controller can only manage G-vTAP agents that has the same version. For example, the G-vTAP Controller v1.3 can only manage G-vTAP agents v1.3. So, if you have G-vTAP agents v1.2 still deployed in the instances, you must configure both G-vTAP Controller v1.2 and v1.3.

While configuring the G-vTAP Controllers, you can also specify the tunnel type to be used for carrying the mirrored traffic from the G-vTAP agents to the GigaVUE V Series nodes. The tunnel type can be L2GRE or VXLAN.

To configure the G-vTAP Controllers:

1. Click **Cloud** on the top navigation pane.
2. On the left navigation pane, select **OpenStack > Configuration**.
3. Select the **G-vTAP Controllers** tab. The **G-vTAP Controllers** page is displayed. Refer to [Figure 2-27 on page 40](#).

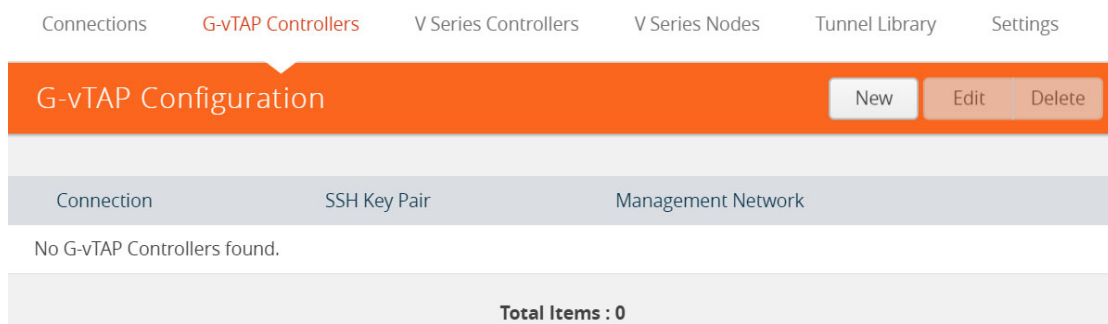


Figure 2-27: G-vTAP Controllers Page

- Click **New**. The OpenStack G-vTAP Controller Configuration profile page is displayed. [Figure 2-28](#) shows an example of a G-vTAP Controller Configuration profile.

Figure 2-28: OpenStack G-vTAP Controller Configuration profile page

- Configure the profile for the G-vTAP Controller by clicking in each field and selecting the configuration item from the drop-down list.

Table 2-7: Fields for G-vTAP Configuration

| Fields | Description |
|---------------------------|---|
| Connection | The name of the connection. NOTE: For shared controller configuration, you must select the required connection for configuring the G-vTAP Controller. Peering must be active in the selected connection to allow the rest of the connections containing the V-series nodes to be monitored. |
| SSH KeyPair | The SSH key pair for the G-vTAP Controller. For more information about SSH key pair, refer to Key Pairs on page 17 . |
| Security Groups | The security group created for the G-vTAP Controller. For example, sg_gvtap-controller. For more information, refer to Security Group on page 13 . |
| Management Network | The management network that GigaVUE-FM uses to communicate with G-vTAP Controllers, GigaVUE V Series Controllers, and GigaVUE V Series Nodes. |

Table 2-7: Fields for G-vTAP Configuration

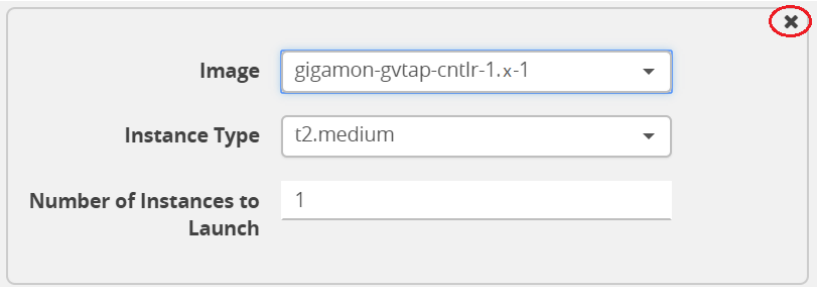
| Fields | Description |
|--|---|
| IP Address Type | <p>The type of IP address GigaVUE-FM needs to communicate with G-vTAP controllers:</p> <ul style="list-style-type: none"> • Private—A private IP can be used when GigaVUE-FM, the G-vTAP Controller, or the GigaVUE V Series Controller reside inside the same project. • Floating—A floating IP is needed only if GigaVUE-FM is not in the same project in the cloud or is outside the cloud. GigaVUE-FM needs a floating IP to communicate with the controllers from an external network. Make sure that this floating IP will not be used by other instances in the cloud. <p>NOTE: If GigaVUE-FM resides inside the same project, no floating IPs are necessary for the controllers.</p> |
| Controller Version(s) | <p>The G-vTAP Controller version.</p> <p>The G-vTAP Controller version you configure must always have the same version number as the G-vTAP agents deployed in the instances. This is because the G-vTAP Controller v1.2-1 can only manage G-vTAP agents v1.2-1. Similarly, the G-vTAP Controller v1.3-1 can only manage G-vTAP agents v1.3-1.</p> <p>If there are multiple versions of G-vTAP agents deployed in the instances, then you must configure multiple versions of G-vTAP Controllers that matches the version numbers of the G-vTAP agents.</p> <p>NOTE: If there is a version mismatch between the G-vTAP controllers and G-vTAP agents, GigaVUE-FM cannot detect the agents in the instances.</p> <p>To add multiple versions of G-vTAP Controllers:</p> <ol style="list-style-type: none"> Under Controller Versions, click Add. From the Image drop-down list, select a G-vTAP Controller image that matches with the version number of G-vTAP agents installed in the instances. From the Flavor down-down list, select a flavor for the G-vTAP Controller. In Number of Instances to Launch, specify the number of G-vTAP Controllers to launch. The minimum number you can specify is 1. |
| Controller Version(s) (continued) | <p>An older version of G-vTAP Controller can be deleted once all the G-vTAP agents are upgraded to the latest version.</p> <p>To delete a specific version of G-vTAP Controller, click x (delete) next to its G-vTAP Controller image.</p>  <p>The screenshot shows a configuration panel for a G-vTAP Controller. It includes three fields: 'Image' with a dropdown menu showing 'gigamon-gvtap-cntlr-1.x-1', 'Instance Type' with a dropdown menu showing 't2.medium', and 'Number of Instances to Launch' with a text input field containing '1'. A red 'x' icon in a circle is located in the top right corner of the configuration panel, indicating a delete action.</p> |

Figure 2-29: Delete a G-vTAP Controller Version

Once you delete a G-vTAP Controller image from the G-vTAP Configuration page, all the G-vTAP Controller instances of that version are deleted.

Table 2-7: Fields for G-vTAP Configuration

| Fields | Description |
|-----------------------------|---|
| Additional Subnet(s) | <p>(Optional) If there are G-vTAP agents on subnets that are not IP routable from the management subnet, additional subnets must be specified so that the G-vTAP Controller can communicate with all the G-vTAP agents.</p> <p>Click Add to specify additional subnets, if needed. Also, make sure that you specify a list of security groups for each additional subnet.</p> |
| Tag(s) | <p>(Optional) The key name and value that helps to identify the G-vTAP Controller instances in your environment. For example, you might have G-vTAP Controllers deployed in many regions. To distinguish these G-vTAP Controllers based on the regions, you can provide a name that is easy to identify such as us-west-2-gvtap-controllers. To add a tag:</p> <ol style="list-style-type: none"> Click Add. In the Key field, enter the key. For example, enter Name. In the Value field, enter the key value. For example, us-west-2-gvtap-controllers. <p>When the G-vTAP Controllers are launched in the VPC, they appear as shown in Figure 2-30 on page 43:</p> |




Figure 2-30: G-vTAP Controllers with Custom Tag Name

| | |
|---|--|
| Agent Tunnel Type | The type of tunnel used for sending the traffic from G-vTAP agents to GigaVUE V Series nodes. The options are GRE or VXLAN tunnels. |
| G-vTAP Agent MTU (Maximum Transmission Unit) | <p>The Maximum Transmission Unit (MTU) is the maximum size of each packet that the tunnel endpoint can carry from the G-vTAP agent to the GigaVUE V Series node.</p> <p>For GRE, the default value is 1450.</p> <p>For VXLAN, the default value is 1400. However, the G-vTAP agent tunnel MTU should be 50 bytes less than the agent's destination interface MTU size.</p> |

6. Verify that the G-vTAP Controller instance is created and running:
 - a. Log in to Horizon and select **Project > Instances** to verify the launch of the G-vTAP Controller. Refer to the example in [Figure 2-31 on page 44](#).

| Instance Name | Image Name | IP Address | Size | Key Pair | Status | Availability Zone | Task | Power State | Time since created | Actions |
|--------------------|-------------------|------------|-----------|----------------|--------|-------------------|------|-------------|--------------------|-----------------|
| G-vTapController-1 | gvtap-cntlr_1.x-1 | | m1.medium | mitaka_dan_key | Active | nova | None | Running | 0 minutes | Create Snapshot |
| FM 56203 | FM 5.x | | FM_Flavor | mitaka_dan_key | Active | nova | None | Running | 1 week, 5 days | Create Snapshot |

Figure 2-31: G-vTAP Controller Instance in OpenStack Horizon

- b. In GigaVUE-FM, select **OpenStack > Visibility Fabric > G-vTAP Controllers** to verify the launch of the G-vTAP Controller.

The G-vTAP Controller is displayed with the status as OK as shown in [Figure 2-32 on page 44](#).

| G-vTAP Controller Name | Management IP | Version | Status |
|------------------------|---------------|-------------|--------|
| G-vTapController-1 | | cntlr_1.x-1 | OK |

Total Items : 1

* Note: If configured G-vTap Controller Instances do not show on this page, please check Alarms/Events page for more details.

Figure 2-32: Cloud > OpenStack > Visibility Fabric > G-vTAP Controllers

The launch of the G-vTAP Controller is also displayed in **Cloud > Audit Logs**.

Configuring the GigaVUE V Series Controllers

The GigaVUE V Series Controller Configuration page defines the parameters for a GigaVUE V Series Controller. Creating a GigaVUE V Series Controller profile automatically launches the controllers.

To configure a GigaVUE V Series Controller:

1. Click **Cloud** in the top navigation pane.
2. On the left navigation pane, select **OpenStack > Configuration**.

3. Select the **V Series Controllers** tab. The V Series Controller Configuration page is displayed.

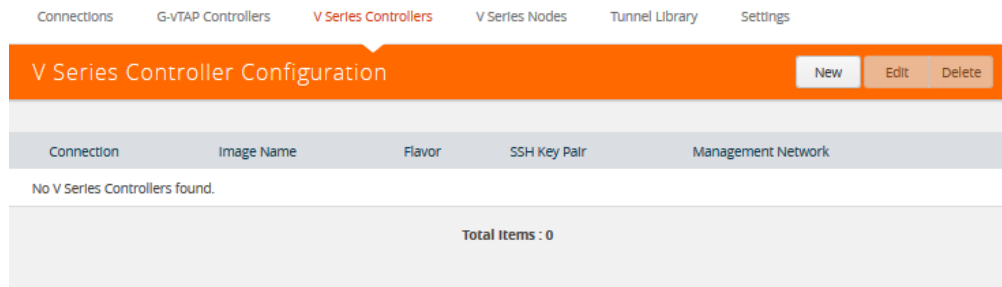


Figure 2-33: GigaVUE V Series Controller Configuration

4. Click **New**. The V Series Controller Configuration Profile page is displayed. [Figure 2-34](#) shows an example of a V Series Controller profile.

NOTE: For shared controller configuration, you must select the required connection for configuring the V Series Controller. Peering must be active in the selected connection to allow the rest of the connections to be monitored.

Figure 2-34: V Series Controller Configuration Profile

5. Configure the profile for the V Series Controller by clicking in each field and selecting the configuration item from the drop-down list.
For a description of the fields, refer to step 5 in [Configuring the G-vTAP Controllers on page 40](#).
6. Verify that the V Series Controller instance is created and running:
 - a. Log in to Horizon and select **Project > Instances** to verify the launch of the GigaVUE V Series Controller. Refer to the example in [Figure 2-35 on page 46](#).

The screenshot shows the OpenStack Horizon 'Instances' page. On the left is a navigation sidebar with categories: Project, Compute, Overview, Instances (highlighted), Images, Access & Security, Network, and Identity. The main content area has a search bar for 'Instance Name', a 'Filter' button, and action buttons for 'Launch Instance', 'Delete Instances', and 'More Actions'. Below is a table of instances:

| Instance Name | Image Name | IP Address | Size | Key Pair | Status | Availability Zone | Task | Power State | Time since created | Actions |
|---------------------|-------------------------------------|------------|-----------|----------------|--------|-------------------|------|-------------|--------------------|-----------------|
| VSeriesController-1 | gigamon-gigavue-vseries-cntlr-1.x-1 | | m1.medium | mitaka_dan_key | Active | nova | None | Running | 1 minute | Create Snapshot |
| G-vTapController-1 | gvtap-cntlr_1.x-1 | | m1.medium | mitaka_dan_key | Active | nova | None | Running | 19 minutes | Create Snapshot |

Figure 2-35: V Series Controller Instance in OpenStack Horizon

- b. In GigaVUE-FM, select **OpenStack > Visibility Fabric > V Series Controllers** to verify the launch of the V Series Controller.

The V Series Controller is displayed with the status as OK as shown in [Figure 2-32 on page 44](#).

The screenshot shows the 'V Series Controllers' page in GigaVUE-FM. It has a navigation bar with tabs for 'G-VTAP Controllers', 'V Series Controllers' (selected), 'V Series Nodes', and 'Tunnel Endpoints'. Below the tabs is a table with columns: 'V Series Controller Name', 'Management IP', 'Version', and 'Status'. The table contains one entry:

| V Series Controller Name | Management IP | Version | Status |
|--------------------------|---------------|---------|--------|
| VSeriesController-1 | | 1.x-1 | OK |

Below the table, it says 'Total Items : 1' and includes a note: '* Note: If configured V Series Controller instances do not show on this page, please check Alarms/Events page for more details.'

Figure 2-36: OpenStack > Visibility Fabric > V Series Controllers

Configuring the GigaVUE V Series Node

GigaVUE® V Series node is a visibility node that aggregates mirrored traffic from multiple G-vTAP agents. It applies filters, manipulates the packets using GigaSMART applications, and distributes the optimized traffic to cloud-based tools or backhaul to GigaSECURE® Cloud using the standard IP GRE tunnels.

GigaVUE V Series nodes can be successfully launched only after GigaVUE V Series Controller is fully initialized and the status is displayed as OK.

The V Series Node Configuration page defines the parameters for a GigaVUE V Series node. Creating a GigaVUE V Series node profile automatically launches the V Series node.

To configure a GigaVUE V Series node profile:

1. Click **Cloud** in the top navigation pane.
2. On the left navigation pane, select **OpenStack > Configuration**.

3. Select V Series Nodes.

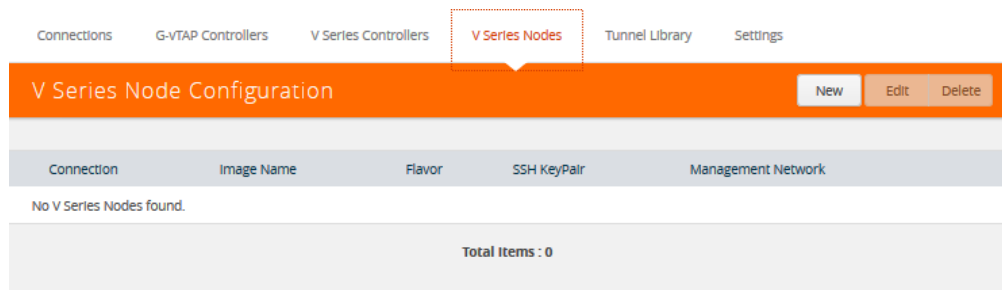


Figure 2-37: GigaVUE V Series Node Configuration

4. Click New.

The V Series Node Configuration Profile page is displayed. Figure 2-38 shows an example of a V Series Node profile.

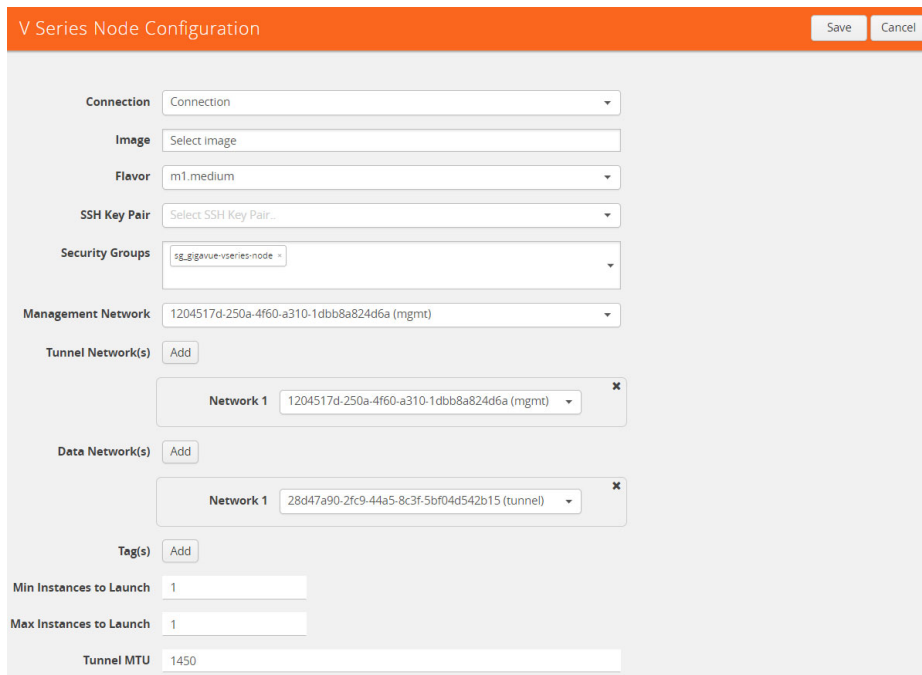


Figure 2-38: V Series Node Configuration Profile page

5. Configure the profile for the V Series Node by clicking in each field and selecting the configuration item from the drop-down list.

For a description of the fields, refer to step 5 in [Configuring the G-vTAP Controllers on page 40](#). For additional options in step 5, refer to Table 2-8 on page 48.

Table 2-8: Fields for GigaVUE V Series Node Launch Configuration

| Parameter | Description |
|--|---|
| Tunnel Subnet(s) | The network that the GigaVUE V Series node uses to communicate with the monitoring tools or GigaVUE H Series node. The tunnel network can be same as the management network. |
| Data Subnet(s) | The network that receives the mirrored GRE tunnel traffic from the G-vTAP agents or TaaS. |
| Min Instances to Launch | The minimum number of GigaVUE V Series nodes to be launched in OpenStack. The minimum number can be 0. Note: The GigaVUE-FM will delete the nodes if they are idle for over 15 minutes. |
| Max Instances to Launch | The maximum number of GigaVUE V Series nodes that can be launched in OpenStack. |
| MTU (Maximum Transmission Unit) | The Maximum Transmission Unit (MTU) is applied on the outgoing tunnel endpoints of the GigaVUE V Series node when a monitoring session is deployed. The default value is 1450. |

6. Verify that the V Series Node instance is created and running:
 - a. Log in to Horizon and select **Project > Instances** to verify the launch of the GigaVUE V Series node. Refer to the example in [Figure 2-39 on page 48](#).

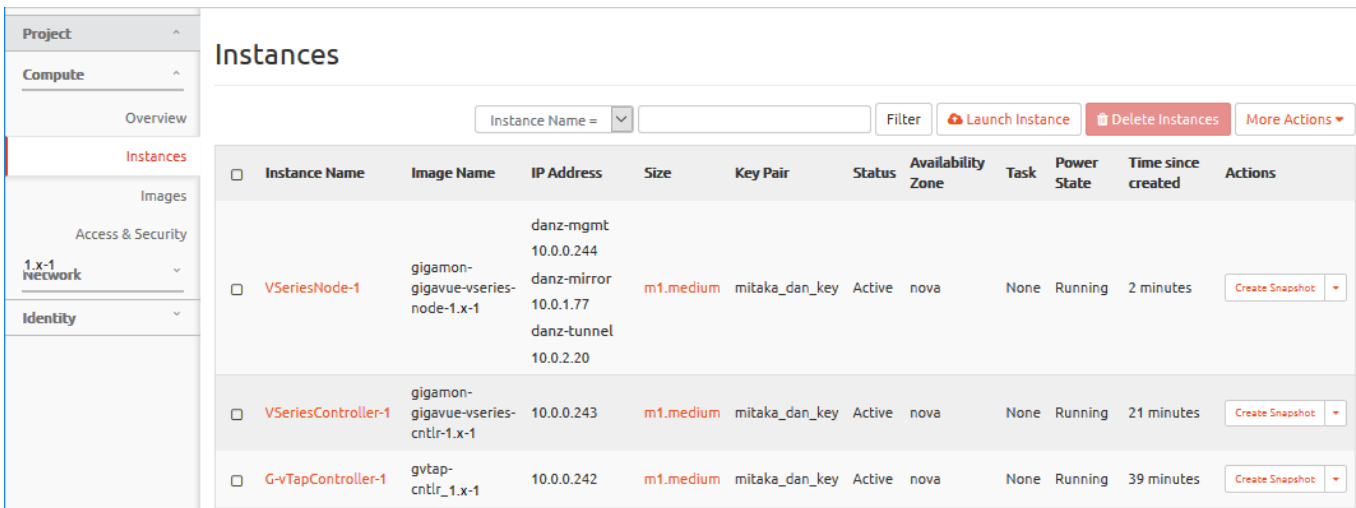


Figure 2-39: V Series Node Instance in OpenStack Horizon

- b. In GigaVUE-FM, select **OpenStack > Visibility Fabric > V Series Node** to verify the launch of the GigaVUE V Series node.

The V Series Node is listed with a Status of OK as shown in the example in [Figure 2-40 on page 49](#).

G-vTAP Controllers V Series Controllers **V Series Nodes** Tunnel Endpoints

V Series Nodes

| V Series Node Name | Management IP | Version | Status |
|---|---------------|---------|--------|
| Connection1 | | | |
| VSeriesNode-1 | | 1.1-1 | ● OK |
| Total Items : 1 | | | |
| * Note: If configured V Series Node instances do not show on this page, please check Alarms/Events page for more details. | | | |

Figure 2-40: Cloud > Visibility Fabric > V Series Nodes

Configuring Monitoring Sessions

This chapter describes how to setup tunnel endpoints in a monitoring session to receive and send traffic to the GigaVUE V Series node. It also describes how to filter, manipulate, and send the traffic from the V Series node to monitoring tools or to a GigaVUE H Series node.

Refer to the following sections for details:

- [Overview of Visibility Components on page 51](#)
- [Creating Tunnel Endpoints on page 55](#)
- [Creating a Monitoring Session on page 57](#)
- [Configuring the OpenStack Settings on page 108](#)

Overview of Visibility Components

The GigaVUE V Series node aggregates the traffic from multiple G-vTAP agents or TaaS and filters them using maps. It applies intelligence and optimization to the aggregated traffic using GigaSMART applications such as Flow Mapping™, sampling, slicing, and masking, and distributes them to the tunnel endpoints.

[Table 3-1 on page 52](#) lists the components of the monitoring session:

Table 3-1: Components of Traffic Visibility Sessions

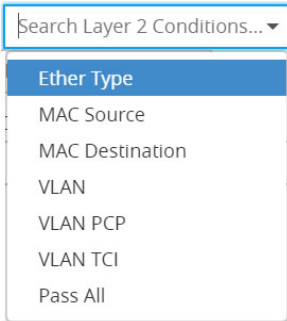
| Parameter | Description |
|-----------------|---|
| Map | A map (M) is used to filter the traffic flowing through the V Series node. It is a collection of one or more rules (R). The traffic passing through a map can match one or more rules defined in the map. |
| Rule | <p>A rule (R) contains specific filtering criteria that the packets must match.</p> <p>The filtering criteria lets you determine the target instances and the (egress or ingress) direction of tapping the network traffic.</p> <p>The rules must contain the appropriate Layer 2 (L2) to Layer 4 (L4) filters defined in them. For example, if you want to filter the traffic for HTTP Port 80, you must select the following criteria:</p> <ul style="list-style-type: none"> • Layer 2—Ethertype IPv4 • Layer 3—Protocol TCP • Layer 4—Port Destination 80 <p>By default, a rule always displays conditions based on the attributes of L2. Refer to Figure 3-1 on page 52.</p>  |
| Priority | <p>A rule is also associated with priority and action set.</p> <p>A priority determines the order in which the rules are executed. The greater the value, the higher the priority.</p> <p>The priority value can range from 0 to 99.</p> |

Figure 3-1: Default Rule Conditions

Table 3-1: Components of Traffic Visibility Sessions

| Parameter | Description |
|-------------------|--|
| Action Set | <p>An Action Set is an exit point in a map that you can drag and create links to the other maps, applications, and monitoring tools. A single map can have multiple action sets. A single action set can have multiple links connecting to maps and applications. You can create an Action Set when you create a rule for a map.</p> <p>In the following example (refer to Figure 3-2), Map 1 has two action sets: Action Set 0 and Action Set 1. The packets that match the rules in Action Set 0 are forwarded to monitoring tools. The packets that match the rules in Action Set 1 are forwarded to Map 2.</p> |

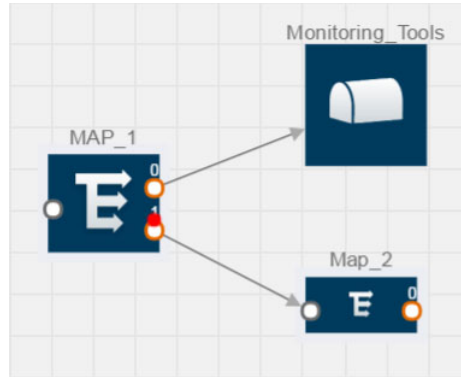


Figure 3-2: Action Set

A single action set can have up to 8 links connecting the same destination point. The same packets from the map are replicated in 8 different links. Refer to [Figure 3-3 on page 53](#).




Figure 3-3: Action Set with Multiple Links

| | |
|----------------------|--|
| Link | <p>A link directs the packets to flow from a map to the destination. The destination could be the other maps, applications, and the monitoring tools. In Figure 3-2 on page 53, the link originating from action set 0 is moving the traffic from MAP_1 to Monitoring_Tools.</p> <p>A link lets you add header transformation to the packets passing through it before they are sent to the destination. This transformation is supported only with GigaVUE V Series node v1.2-1 and above. For more information about Header Transformation, refer to Adding Header Transformations on page 98.</p> |
| Group | A group is a collection of maps that are pre-defined and saved in the map library for reuse. |
| Application | An application performs operations such as sampling, slicing, and masking on the traffic. |
| Inclusion Map | An inclusion map determines the instances to be included for monitoring. This map is used only for target selection. |
| Exclusion Map | An exclusion map determines the instances to be excluded from monitoring. This map is used only for target selection. |

Table 3-1: Components of Traffic Visibility Sessions

| Parameter | Description |
|---|---|
| Target | A target determines the instances that are to be monitored. Targets are determined based on the following formula: $\text{Target} = (\text{Maps} \cap \text{Inclusion map}) - \text{Exclusion map}$ |
| Automatic Target Selection (ATS) | A built-in feature that automatically selects the cloud instances based on the rules defined in the maps, inclusion maps, and exclusion maps in the monitoring session. |
| Tunnel | A tunnel lists the monitoring tools to which the traffic matching the filtered criteria is routed. |

Creating Tunnel Endpoints

Traffic from the V Series node is distributed to tunnel endpoints in a monitoring session. A tunnel endpoint can be created using a standard L2 Generic Routing Encapsulation (GRE) tunnel or a Virtual Extensible LAN (VXLAN) tunnel.

To create a tunnel endpoint:

1. In GigaVUE-FM, on the top navigation pane, select **Cloud**.
2. On the left navigation pane, select **OpenStack > Configuration**.
3. Select the **Tunnel Spec Library** tab. The Tunnel Library page appears. Refer to [Figure 3-4 on page 55](#).

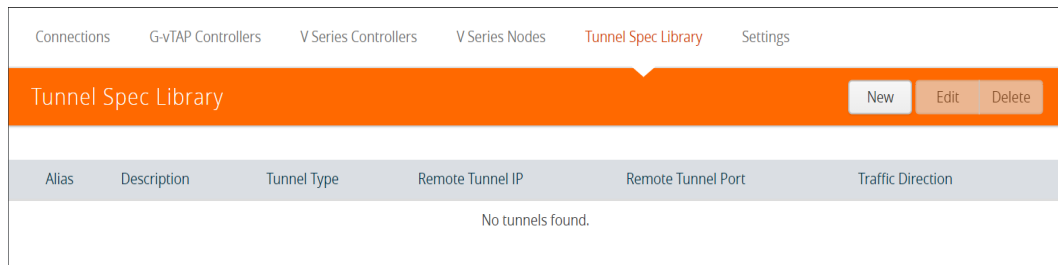


Figure 3-4: Tunnel Library

4. Click **New**. The Edit Tunnel page appears. Refer to [Figure 3-5 on page 55](#).

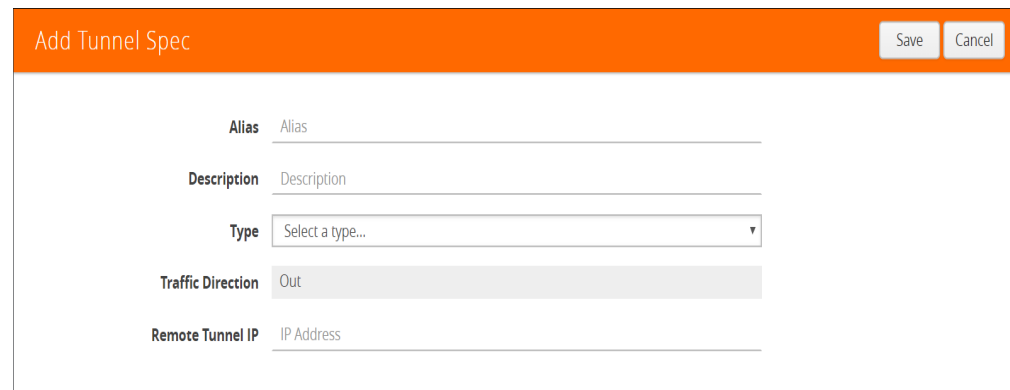


Figure 3-5: Adding a Tunnel Endpoint

5. On the Edit Tunnel page, select or enter the appropriate information in the fields. Refer to [Figure 3-5 on page 55](#) and [Table 3-2 on page 56](#).

Table 3-2: Field Descriptions for Tunnel Endpoint

| Field | Description |
|--------------------------|---|
| Alias | The name of the tunnel endpoint. NOTE: Do not enter spaces in the alias name. |
| Description | The description of the tunnel endpoint. |
| Type | The type of the tunnel. Select L2GRE or VXLAN to create a tunnel. If you choose VXLAN, you must enter the remote IP interface. |
| Traffic Direction | The direction of the traffic flowing through the V Series node. Choose Out for creating a tunnel from the V Series node to the destination endpoint. NOTE: Traffic Direction In is not supported in the current release. |
| Remote Tunnel IP | The IP address of the tunnel destination endpoint. NOTE: You cannot create two tunnels from a V Series node to the same IP address. |

- Click **Save**.
- Select **OpenStack > Visibility Fabric > Tunnel Endpoints** and verify the tunnel endpoint added to GigaVUE-FM. Refer to [Figure 3-6 on page 56](#).

| Tunnel Spec | Connection | Type | Traffic Direction | Remote Tunnel IP | Remote Tunnel Port | Local IP | Fabric Node ID | Fabric Node Name | Fabric Node Tunnel Name | Monitoring Sessions |
|-------------|-------------|------|-------------------|------------------|--------------------|-----------|--------------------------------------|------------------|-------------------------|------------------------|
| Tunnel1 | Connection1 | GRE | out | 54.106.153.87 | 0 | 10.0.2.20 | 22c287ec-cd7a-453f-92e8-7b5aa4ada0a7 | VSeriesNode-1 | tep0 | Connection1_Monitoring |

Total Items : 1

Figure 3-6: OpenStack > Visibility Fabric > Tunnel Endpoints

Creating a Monitoring Session

GigaVUE-FM automatically collects inventory data on all target instances available in your OpenStack environment. You can design your monitoring session to include or exclude the instances that you want to monitor. You can also choose to monitor egress, ingress, or all traffic.

When a new target instance is added to your OpenStack environment, GigaVUE-FM automatically detects and adds the instance into your monitoring session. Similarly, when an instance is removed, it updates the monitoring sessions to show the removed instance.

To design your monitoring session, refer to the following sections:

- [Creating a New Session on page 57](#)
- [Cloning a Monitoring Session on page 59](#)
- [Splitting a Monitoring Session on page 60](#)
- [Creating a Map on page 61](#)
- [Adding Applications to the Monitoring Session on page 69](#)
- [Deploying the Monitoring Session on page 94](#)
- [Viewing the Statistics on page 101](#)
- [Viewing the Topology on page 104](#)

Creating a New Session

You can create multiple monitoring sessions within a single project connection.

To create a new session:

1. In GigaVUE-FM, on the top navigation pane, select **Cloud**.
2. Select **OpenStack > Monitoring Session**. The Monitoring Sessions page appears. Refer to [Figure 3-7 on page 57](#).




Figure 3-7: Monitoring Sessions

3. Click **New**.

The Monitoring Session configuration page appears. Refer to [Figure 3-8](#) on page 58.

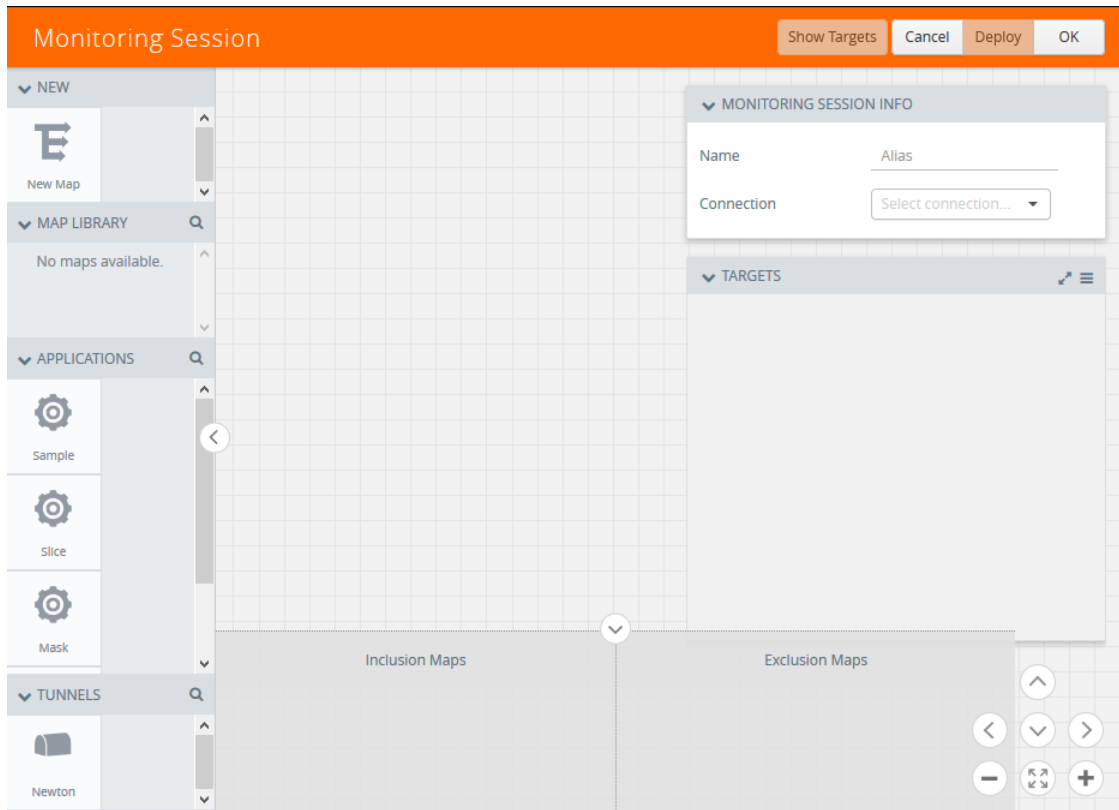


Figure 3-8: Creating Monitoring Session

4. Enter the appropriate information in the **Create a New Monitoring Session Info** dialog box as shown in [Table 3-3](#).

Table 3-3: Fields for Session Info

| Field | Description |
|----------------------------|---|
| Alias | The name of the monitoring session. |
| Monitoring Domain | The name of the monitoring domain. |
| Connection | The OpenStack connection(s) that are to be included as part of the monitoring domain. You can select the required connections that need to be part of the monitoring domain. |
| Agent Pre-filtering | When enabled, traffic is filtered at the G-vTAP agent-level, before mirroring to the V Series Nodes, which reduces the load on the V Series Nodes and the Cloud networks. Refer to Agent Pre-filtering. |

Cloning a Monitoring Session

You can clone an existing monitoring session.

To clone a monitoring session:

1. Select the monitoring session that you need to clone from the **Monitoring Sessions** page.
2. Click **Clone**.
3. Enter the appropriate information in the **Clone Monitoring Session** dialog box as shown in Table 3-3 on page 58.

Table 3-4: Fields for Cloning the Monitoring Session.

| Field | Description |
|-------------------|-------------------------------------|
| Alias | The name of the monitoring session. |
| Monitoring Domain | The name of the monitoring domain. |

4. Click **Create** to create the cloned monitoring session.
5. Once the monitoring session is created, click **Edit** to add the connections to the cloned monitoring session.

Splitting a Monitoring Session

You can split a monitoring session.

To split a monitoring session:

1. Select the monitoring session that you need to split from the **Monitoring Sessions** page.
2. Click **Split**.
3. Enter the appropriate information in the **Split A Monitoring Session** dialog box as shown in Table 3-3 on page 58.

Table 3-5: Fields for Splitting the Monitoring Session.

| Field | Description |
|------------------------------------|---|
| Original Monitoring Session | Alias: The name of the original monitoring session from which a split monitoring session is to be created. Connections: Connections that belong to the original monitoring session. |
| New Monitoring Session | Alias: The name of the new monitoring session that is to be created. Connections: Connections that have been added to the new monitoring session. NOTE: You can use the arrow to move the connections from the original monitoring session to the split the monitoring session and vice-versa. Use the Search filter to search for the required connections. |

4. Click **Split**.

Creating a Map

Each map can have up to 32 rules associated with it. [Table 3-7](#) lists the various rule conditions that you can select for creating a map, inclusion map, and exclusion map.

Table 3-6: Conditions for the Rules

| Conditions | Description |
|---|--|
| L2, L3, and L4 Filters | |
| Ether Type | <p>The packets are filtered based on the selected ethertype. The following conditions are displayed:</p> <ul style="list-style-type: none"> • IPv4 • IPv6 • ARP • RARP • Other <p>L3 Filters</p> <p>If you choose IPv4 or IPv6, the following L3 filter conditions are displayed:</p> <ul style="list-style-type: none"> • Protocol • IP Fragmentation • IP Time to live (TTL) • IP Type of Service (TOS) • IP Explicit Congestion Notification (ECN) • IP Source • IP Destination <p>L4 Filters</p> <p>If you select TCP or UDP protocol, the following L4 filter conditions are displayed:</p> <ul style="list-style-type: none"> • Port Source • Port Destination |
| MAC Source | The egress traffic from the instances or ENIs matching the specified source MAC address is selected. |
| MAC Destination | The ingress traffic from the instances or ENIs matching the specified destination MAC address is selected. |
| VLAN | All the traffic matching the specified IEEE 802.1q Virtual LAN tag is filtered. Specify a number from 0 to 4094. |
| VLAN Priority Code Point (PCP) | All the traffic matching the specified IEEE 802.1q Priority Code Point (PCP) is filtered. Specify a value between 0 to 7. |
| VLAN Tag Control Information (TCI) | All the traffic matching the specified VLAN TCI value is filtered. Specify the exact TCI value. |
| Pass All | All the packets coming from the monitored instances are passed through the filter. When Pass All is selected, the L3 and L4 filters are disabled. |

When you select a condition without source or destination specified, then both egress and ingress traffic is selected for tapping the traffic. For example, if you select Ether

Type as IPv4, TCP as the protocol, and do not specify IPv4 source or destination, then both egress and ingress traffic is selected for monitoring purpose.

When you select a condition with either source or destination specified, it determines the direction based on the selection.

NOTE: You can create Inclusion and Exclusion Maps using all default conditions except Ether Type and Pass All.

To create a new map:

1. Select **OpenStack > Monitoring Session**.
2. Click **New**. The Monitoring Sessions page is displayed.
3. Create a new session. Refer to [Creating a New Session on page 57](#).
4. From **Maps**, drag and drop a new map template to the workspace.
5. Click on the map, then click details. Refer to [Figure 3-9 on page 62](#).




Figure 3-9: Map Details

The map rules quick view is displayed as shown in [Figure 3-10](#).

A screenshot of a web form titled 'Map_1'. At the top right are 'Save' and 'Add to Library' buttons. The form contains several sections: 'Alias' with the value 'Map_1'; 'Comments' with the value 'Comments'; 'Map Rules' with an 'Add a Rule' button; 'Rule 1' configuration including search boxes for Layer 2, 3, and 4 conditions, and 'Priority' and 'ActionSet' both set to 0; and a 'Rule Comment' section with a 'Comment' field. Below the comment field are two rule conditions: 'Ether Type' with a value of 'IPv4' and '0x0800', and 'Protocol' with a value of 'TCP' and '6'. Each condition has a close button (X).

Figure 3-10: Creating a New Map

6. Enter the appropriate information for creating a new map as shown in [Table 3-7](#).

Table 3-7: Fields for Creating a New Map

| Parameter | Description |
|-----------------|---|
| Alias | The name of the new map. NOTE: The name can contain alphanumeric characters with no spaces. |
| Comments | The description of the map. |
| Rule Conditions | The rules for filtering the traffic in the map. |
| Map Rules | To add a map rule: |

- a. Click **Add a Rule**.
- b. Select a condition from the **Search L2 Conditions** drop-down list and specify a value. Based on this selection, the Search L3 Conditions drop-down list is automatically updated. Refer to [Figure 3-11 on page 63](#).




Figure 3-11: L2 Conditions

- c. Select a condition from the **Search L3 Conditions** drop-down list and specify a value. Refer to [Figure 3-12 on page 63](#).

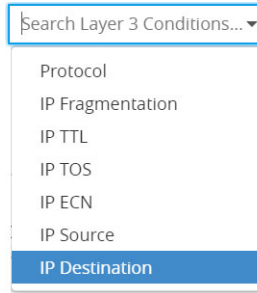


Figure 3-12: L3 Conditions

Table 3-7: Fields for Creating a New Map

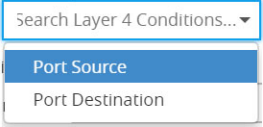
| Parameter | Description |
|-----------|--|
| | <p>d. (Optional) If you have selected TCP or UDP as the protocol in the L3 conditions, then select Port Source or Port Destination from the Search L4 Conditions drop-down list and specify a value. If you have selected conditions other than TCP or UDP, then the Search L4 Conditions drop-down list is disabled. Refer to Figure 3-13 on page 64.</p>  |
| | <p>e. (Optional) In the Priority and Action Set box, assign a priority and action set.</p> <p>f. (Optional) In the Rule Comment box, enter a comment for the rule.</p> <p>NOTE: Repeat steps b through f to add more conditions.</p> <p>NOTE: Repeat steps a through f to add nested rules.</p> |

Figure 3-13: L4 Conditions

NOTE: Do not create duplicate map rules with the same priority.

7. To reuse the map, click **Add to Library**. Save the map using one of the following ways:

- Select an existing group from the **Select Group** list and click **Save**.
- Enter a name for the new group in the **New Group** field and click **Save**.

NOTE: The maps saved in the Map Library can be reused in any monitoring session created in the project.

8. Click **Save**.

To edit or delete a map, click a map and select **Details** to edit the map or **Delete** to delete the map as shown in [Figure 3-14](#).






Figure 3-14: Editing or Deleting a Map

Click the **Show Targets** button to view the monitoring targets highlighted in blue. Refer to [Figure 3-15](#).



Figure 3-15: Viewing the Topology

Click on  to expand the **Targets** dialog box. Click on  to change the view from the Topology view to the Instances view. To view details about an Instance, click the arrow next to the Instance. Refer to [Figure 3-16](#).

The screenshot shows the 'Instances' view with a filter set to 'all'. It lists two instances:



| Instances (Total: 2) | Selected |
|--|-------------------------------------|
|  3c8af658-bf9a-420e-b57b-5f9332c5ac3d (CentOS-7 G-VTAP Agent) | <input checked="" type="checkbox"/> |
|  9678b632-a7e5-4008-b436-426e001e0712 (Ubuntu-G-VTAP Agent) | <input checked="" type="checkbox"/> |
| Tag(s) | |
| vNIC ID: | fa:16:3e:b5:29:02 |
| vNIC IP: | 10.0.0.128 10.210.219.217 |
| Direction: | ["Ingress","Egress"] |

Figure 3-16: Viewing the Instances

Click on the Filter icon to filter Instances based on the Instance Name Prefix, IP address, or MAC address. Refer to [Figure 3-17](#).

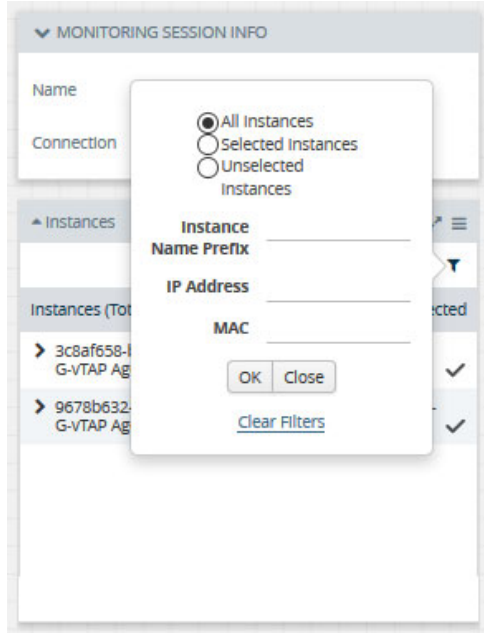


Figure 3-17: Filtering Instances

Agent Pre-filtering

The G-vTAP agent pre-filtering option filters traffic before mirroring it from G-vTAP agent to the V Series Nodes.

Agent pre-filtering is performed directly at the packet capturing point. By filtering at this point, unnecessary traffic is prevented from reaching the fabric nodes that perform filtering and manipulation functions. Preventing this traffic reduces the load on the V Series nodes and the underlying network.

Agent Pre-filtering Guidelines

In cloud environments, there will be limits on how much traffic could be sent out per instance/single or double network interface.

Traffic will be passed if a network packet matches one or more of these rules:

- Only filters from traffic maps will be considered for G-vTAP filters. Inclusion and exclusion maps are purely for ATS (automatic target selection); not for G-vTAP.
- Filters from the first-level maps of the monitoring session will only be used to create G-vTAP maps.
- User-entered L2-L4 filters in the monitoring-session maps must be in the format that V Series Node currently accepts. Non L2-L4 filters are used purely by ATS to select the targets; not for G-vTAP.
- Both egress and ingress maps with filters are supported on G-vTAP.
- Both single and dual network interfaces for G-vTAP agent VMs are supported.

Agent Pre-filtering Capabilities and Benefits

G-vTAP agent pre-filtering has the following capabilities and benefits:

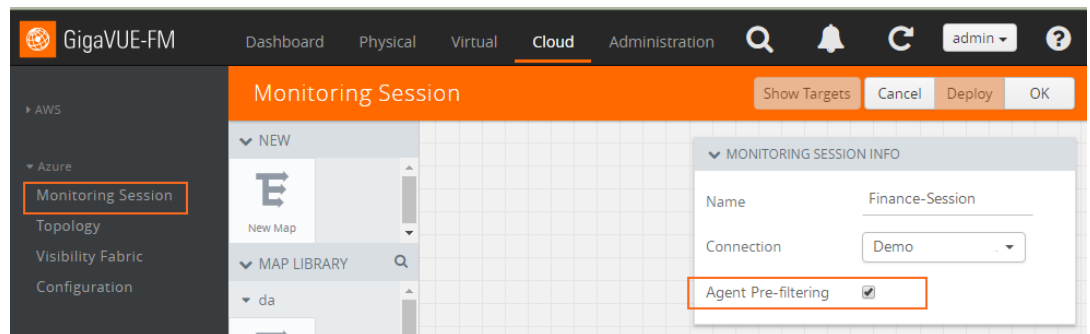
- The agent pre-filtering option can be enabled or disabled at the monitoring-session level and is enabled by default.
- When enabled, traffic is filtered at the G-vTAP agent-level, before mirroring to the V Series Nodes. Consequently, traffic flow to the V Series Nodes is reduced, which reduces the load/cost on the Cloud networks.
- Only rules from first-level maps are pushed to the agents.
- Pass rules are supported 100%.
- Drop rules are supported for only simple cases or single-drop rules with a pass all case.
- Rules that span all monitoring sessions will be merged for an G-vTAP agent, if applicable
- If the max-rule limit of 16 is reached, then all the traffic is passed to the V Series node; no filtering will be performed.

Enable/Disable G-vTAP Agent Pre-filtering

Agent pre-filtering can be enabled or disabled by the user at the monitoring-session level. This ensures that we provide a knob to the user to turn it on or off at the G-vTAP level according to the requirements.

To change the G-vTAP Agent Pre-filtering option setting:

1. **Cloud > OpenStack> Monitoring Session**
2. Open a monitoring session by doing one of the following:
 - a. Click **New** to create a new session.
 - b. Click the check box next to a session and then click **Edit** to edit an existing session.



3. Select or deselect the **Agent Pre-filtering** check box in the MONITORING SESSION INFO box to change the setting. It is enabled by default.
 - a. Deselect the check box to disable it.
 - b. Select the check box to enable it.
4. Click **OK**.

The Monitoring Session view displays the setting in the Agent Pre-filtering column

Monitoring Session

[Deploy](#) [Undeploy](#) [New](#) [Clone](#) [Edit](#)

| <input type="checkbox"/> | Name | Connection | # of Targets | Status | Statistics | Pre-capture Filtering |
|--------------------------|-----------------|------------|--------------|--|----------------------|-----------------------|
| <input type="checkbox"/> | Finance-Session | Demo | 4 | ● Success | View | Yes |
| <input type="checkbox"/> | HR-Session | Demo | 4 | ● Success | View | No |

Adding Applications to the Monitoring Session

Gigamon supports the following GigaSMART applications:

- [Sampling on page 69](#)
- [Slicing on page 70](#)
- [Masking on page 72](#)
- [NetFlow on page 73](#)

You can optionally use these applications to optimize the traffic sent from your instances to the monitoring tools.

Sampling

Sampling lets you sample the packets randomly based on the configured sampling rate and then forwards the sampled packets to the monitoring tools.

To add a sampling application:

1. Drag and drop **Sample** from **APPLICATIONS** to the graphical workspace.



Figure 3-18: Dragging the Sample Application

2. Click **Sample** and select **Details**.

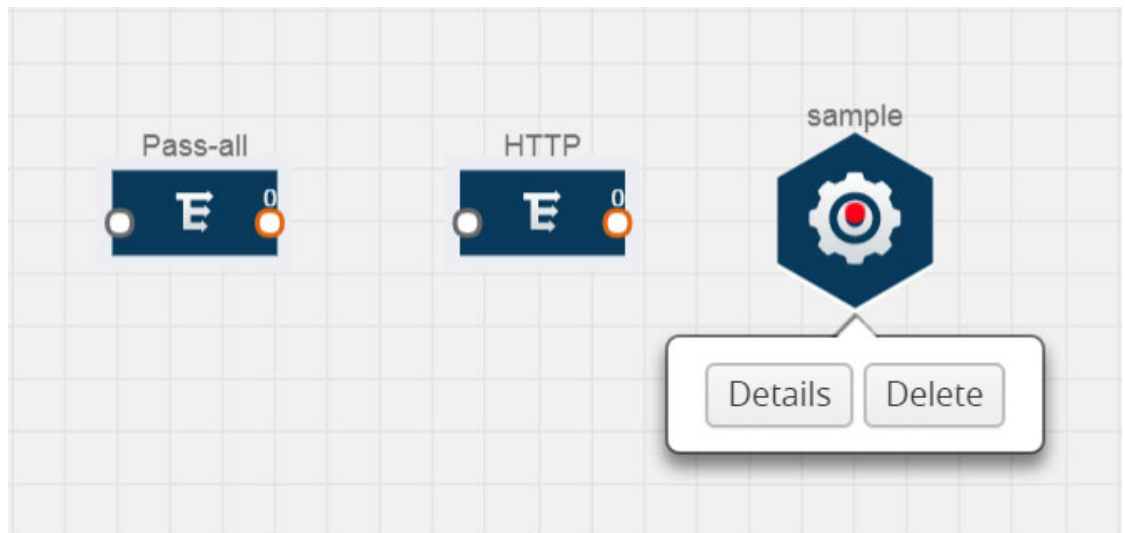


Figure 3-19: Selecting Details

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

The screenshot shows a configuration window titled "Application" with a "Save" button in the top right corner. The window contains the following fields:

- Application:** sample
- Alias:** sample
- State:** On Off
- Sampling Type:** Random Simple (dropdown menu)
- Sampling Rate:** 1:1

On the left side, there is a sidebar with a gear icon and the text "sample".

Figure 3-20: Viewing Sample Application Quick View

4. For State, select the **On** check box to determine that the application is sampling packets randomly. Select the **Off** check box to determine that the application is not currently sampling the packets. The state can be changed at anytime whenever required.
5. From the Sampling Type drop-down list, select the type of sampling:
 - **Random Simple** — The first packet is selected randomly. The subsequent packets are also selected randomly based on the rate specified in the **Sampling Rate** field.
For example, if the first packet selected is 5 and the sampling rate is 1:10, after the 5th packet a random 10 packets are selected for sampling.
 - **Random Systematic** —The first packet is selected randomly. Then, every nth packet is selected, where n is the value specified in the **Sampling Rate** field.
For example, if the first packet selected is 5 and the sampling rate is 1:10, then every 10th packet is selected for sampling: 15, 25, 35, and so on.
6. In the **Sampling Rate** field, enter the ratio of packets to be selected. The default ratio is 1:1.
7. Click **Save**.

Slicing

Packet slicing lets you truncate packets after a specified header and slice length, preserving the portion of the packet required for monitoring purposes.

To add a slicing application:

1. Drag and drop **Slice** from **APPLICATIONS** to the graphical workspace.



Figure 3-21: Dragging the Slice Application

2. Click the Slice application and select **Details**.

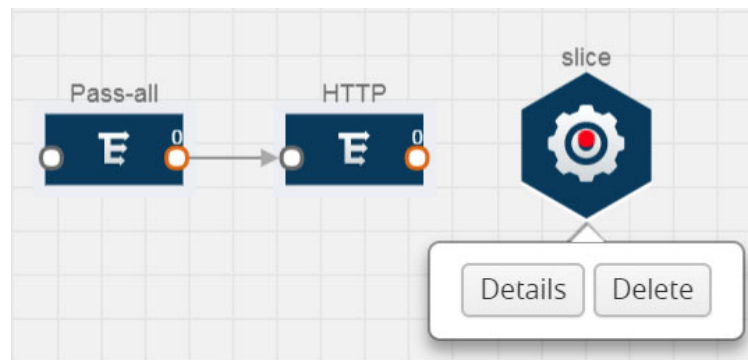


Figure 3-22: Selecting Details

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

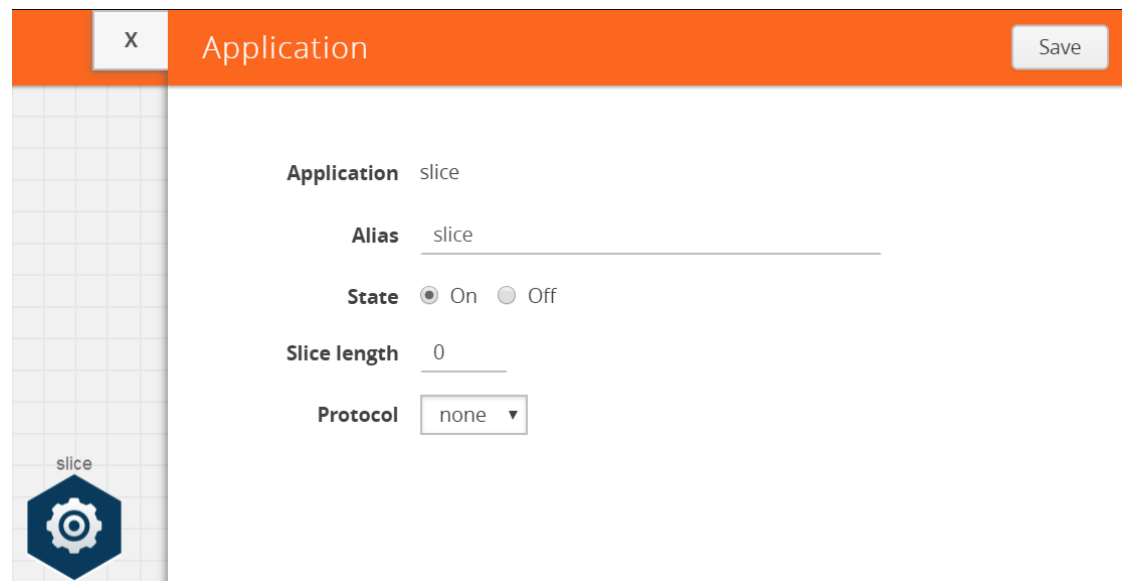


Figure 3-23: Viewing Slice Application Quick View

4. For State, select the **On** check box to determine that the application is slicing packets. Select the **Off** check box to determine that the application is not currently slicing the packets. The state can be changed at a later time whenever required.

5. In the Slice Length field, specify the length of the packet that must be sliced.
6. From the Protocol drop-down list, specify an optional parameter for slicing the specified length of the protocol. The options are as follows:
 - None
 - IPv4
 - IPv6
 - UDP
 - TCP
7. Click **Save**.

Masking

Masking lets you overwrite specific packet fields with a specified pattern so that sensitive information is protected during network analysis.

To add a masking application:

1. Drag and drop **Mask** from **APPLICATIONS** to the graphical workspace.

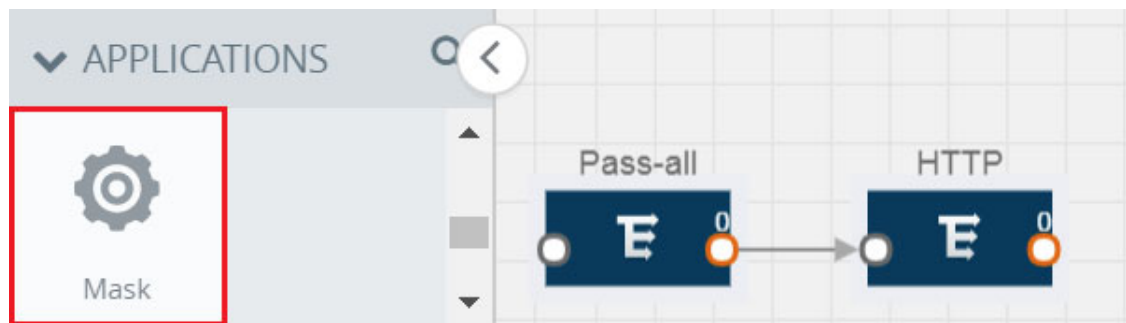


Figure 3-24: Dragging the Mask Application

2. Click the Mask application and select **Details**.

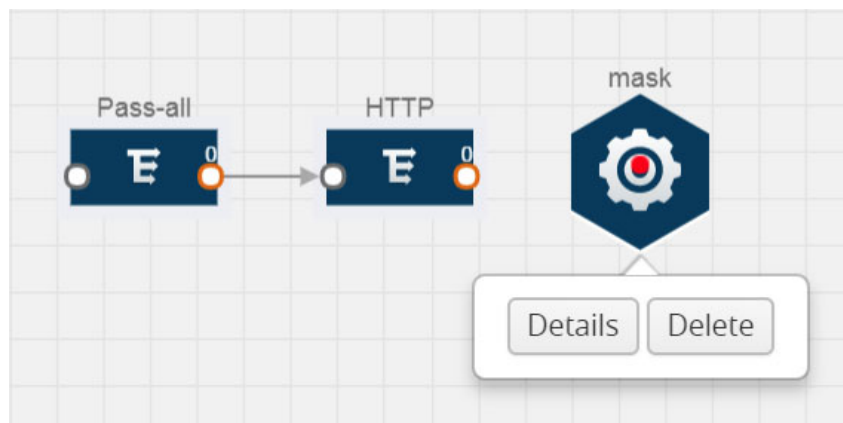


Figure 3-25: Selecting Details

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

The screenshot shows a configuration window titled "Application" with a close button (X) and a "Save" button. The configuration fields are as follows:

- Application:** mask
- Alias:** mask
- State:** On (selected) / Off
- Mask offset:** 0
- Mask length:** 1
- Mask pattern:** 0
- Protocol:** none

On the left side, there is a grid with a gear icon and the label "mask".

Figure 3-26: Viewing Mask Application Quick View

4. For State, select the **On** check box to determine that the application is masking packets. Select the **Off** check box to determine that the application is not currently masking the packets. The state can be changed at anytime whenever required.
5. In the Mask offset field, enter the offset from which the application should start masking data following the pattern specified in the Pattern field.
The value can be specified in terms of either a static offset, that is, from the start of the packet or a relative offset, that is, from a particular protocol layer as specified in the Protocol field.
6. In the Mask length field, enter the length of the packet that must be masked.
7. In the Mask pattern field, enter the pattern for masking the packet. The value of the pattern is from 0 to 255.
8. From the Protocol drop-down list, specifies an optional parameter for masking packets on the data coming from the selected protocol.
9. Click **Save**.

NetFlow

NetFlow collects IP network traffic on all interfaces where NetFlow monitoring is enabled. It gathers information about the traffic flows and exports the NetFlow records, which includes data and templates, to at least one NetFlow collector. The application that serves as a NetFlow collector receives the NetFlow data sent from exporters, processes the information, and provides data visualization and security analytics.

The following are the key benefits of NetFlow application:

- Compresses network information into a single flow record.
- Facilitates up to 99% reduction in data transferred.

- Accelerates the migration of mission-critical workloads.
- Provides summarized information on traffic source and destination, congestion, and class of service.
- Identifies and classifies DDOS attacks, viruses, and worms in real-time.
- Secures network against internal and external threats.
- Identifies top consumers and analyzes their statistics.
- Reduces the cost of security monitoring.
- Analyzes the network flows based on algorithms and behavior rather than signature matching.
- Analyzes east-west traffic between flows within and across VPCs.

The NetFlow application contains key elements that specify what to match in the flow, such as all packets with the same source and destination port, or the packets that come in on a particular interface. For information about Match/Key fields, refer to [Match/Key Fields on page 75](#). A NetFlow record is the output generated by NetFlow. A flow record 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 information about Match/Key fields, refer to [Collect/Non-Key Fields on page 77](#).

[Figure 3-27 on page 74](#) shows an example of a NetFlow application created on a GigaVUE V Series node in the monitoring session.

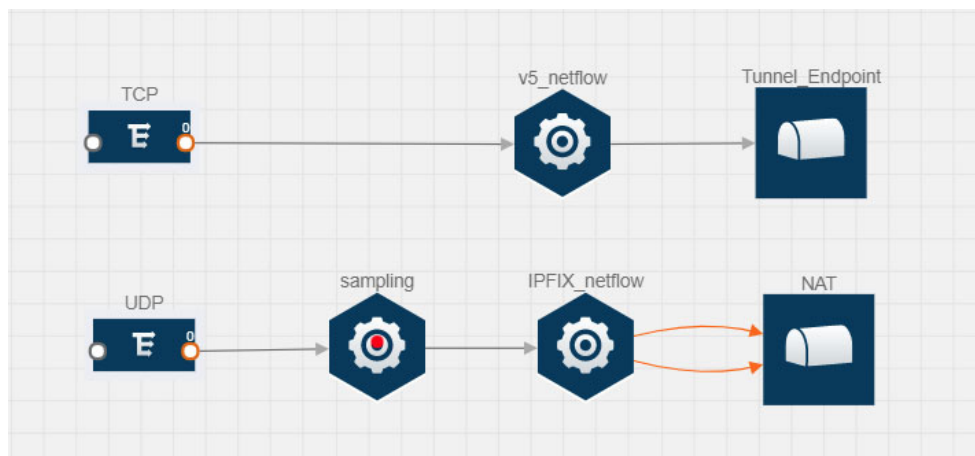


Figure 3-27: NetFlow on GigaVUE V Series Node

The NetFlow record generation is performed on GigaVUE V Series node running the NetFlow application. In [Figure 3-27 on page 74](#), incoming packets from G-vTAP agents are sent to the GigaVUE V Series node. In the GigaVUE V Series node, one map sends the TCP packet to the version 5 NetFlow application. Another map sends the UDP packet to a sampling application. The map rules and applications such as slice, mask, and sample can only be applied prior to sending the data to NetFlow.

A NetFlow application examines the incoming packets and creates a single or multiple flows from the packet attributes. These flows are cached and exported based on the active and inactive cache timeout specified in the Netflow application configuration.

The flow records can be sent to a tunnel or to a NAT device for flow inspection. NAT allows the NetFlow records to be directly transmitted to a collector without a tunnel. For more information about NAT, refer to [Network Address Translation \(NAT\) on page 83](#).

The Netflow application exports the flows using the following export versions:

- version 5—The fields in the NetFlow record are fixed.
- version 9—The fields are configurable, thus a template is created. The template contains information on how the fields are organized and in what order. It is sent to the collector before the flow record, so the collector knows how to decode the flow record. The template is sent periodically based on the configuration.
- IPFIX—The extended version of version 9 supports variable length fields as well as enterprise-defined fields.

Match/Key Fields

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

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

Table 3-8: Match/Key Elements

| Match Type | Description | Supported NetFlow Versions |
|-------------------|---|----------------------------|
| Data Link | | |
| Destination MAC | Configures the destination MAC address as a key field. | v9 and IPFIX |
| Egress Dest MAC | Configures the post Source MAC address as a key field. | IPFIX |
| Ingress Dest MAC | Configures the IEEE 802 destination MAC address as a key field. | IPFIX |
| Source MAC | Configures the IEEE 802 source MAC address as a key field. | v9 and IPFIX |
| IPv4 | | |
| ICMP Type Code | Configures the type and code of the IPv4 ICMP message as a key field. | v9 and IPFIX |
| IPv4 Dest IP | Configures the IPv4 destination address in the IP packet header as a key field. | v9 and IPFIX |
| IPv4 ICMP Code | Configures the code of the IPv4 ICMP message as a key field. | IPFIX |
| IPv4 ICMP Type | Configures the type and code of the IPv4 ICMP message as a key field. | IPFIX |
| IPv4 Options | Configures the IPv4 options in the packets of the current flow as a key field. | IPFIX |
| IPv4 Src IP | Configures the IPv6 source address in the IP packet header as a key field. | v9 and IPFIX |
| IPv4 Total Length | Configures the total length of the IPv4 packet as a key field. | IPFIX |

Table 3-8: Match/Key Elements

| Match Type | Description | Supported NetFlow Versions |
|---------------------|---|----------------------------|
| Network | | |
| IP CoS | Configures the IP Class Of Service (CoS) as a key field. | v9 and IPFIX |
| IP DSCP | Configures the value of a Differentiated Services Code Point (DSCP) encoded in the Differentiated Services field as a key field. | IPFIX |
| IP Header Length | Configures the length of the IP header as a key field. | IPFIX |
| IP Precedence | Configures the value of the IP Precedence as a key field. | IPFIX |
| IP Protocol | Configures the value of the protocol number in the IP packet header as a key field. | v9 and IPFIX |
| IP Total Length | Configures the total length of the IP packet as a key field. | IPFIX |
| IP TTL | For IPv4, configures the value of Time to Live (TTL) as a key field. For IPv6, configures the value of the Hop Limit field as a key field. | IPFIX |
| IP Version | Configures the IP version field in the IP packet header as a key field. | v9 and IPFIX |
| IPv6 | | |
| IPv6 Dest IP | Configures the IPv6 destination address in the IP packet header as a key field. | v9 and IPFIX |
| IPv6 Flow Label | Configures the value of the IPv6 flow label field in the IP packet header as a key field. | v9 and IPFIX |
| IPv6 ICMP Code | Configures the code of the IPv6 ICMP message as a key field. | IPFIX |
| IPv6 ICMP Type | Configures the type of the IPv6 ICMP message as a key field. | IPFIX |
| IPv6 ICMP Type Code | Configures the type and code of the IPv6 ICMP message as a key field. | IPFIX |
| IPv6 Payload Length | Configures the value of the payload length field in the IPv6 header as a key field. | IPFIX |
| IPv6 Src IP | Configures the IPv6 source address in the IP packet header as a key field. | v9 and IPFIX |
| Transport | | |
| L4 Dest Port | Configures the destination port identifier in the transport header as a key field. | v9 and IPFIX |

Table 3-8: Match/Key Elements

| Match Type | Description | Supported NetFlow Versions |
|-------------------|---|----------------------------|
| L4 Src Port | Configures the source port identifier in the transport header as a key field. | v9 and IPFIX |
| TCP AcK Number | Configures the acknowledgment number in the TCP header as a key field. | IPFIX |
| TCP Dest Port | Configures the destination port identifier in the TCP header as a key field. | IPFIX |
| TCP Flags | Configures the TCP control bits observed for the packets of this flow as a key field. | v9 and IPFIX |
| TCP Header Length | Configures the length of the TCP header as a key field. | IPFIX |
| TCP Seq Number | Configures the sequence number in the TCP header as a key field. | IPFIX |
| TCP Src Port | Configures the source port identifier in the TCP header as a key field. | IPFIX |
| TCP Urgent | Configures the urgent pointer in the TCP header as a key field. | IPFIX |
| TCP Window Size | Configures the window field in the TCP header as a key field. | IPFIX |
| UDP Dest Port | Configures the destination port identifier in the UDP header as a key field. | IPFIX |
| UDP Src Port | Configures the source port identifier in the TCP header as a key field. | IPFIX |

Collect/Non-Key Fields

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

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

Table 3-9: Collect/Non-Key Elements

| Match Type | Description | Supported NetFlow Versions |
|------------------|--|----------------------------|
| Counter | | |
| Byte Count | Configures the number of octets since the previous report in incoming packets for the current flow as a non-key field. | v9 and IPFIX |
| Packet Count | Configures the number of incoming packets since the previous report for this flow as a non-key field. | v9 and IPFIX |
| Data Link | | |
| Destination MAC | Configures the destination MAC address as a non-key field. | v9 and IPFIX |

Table 3-9: Collect/Non-Key Elements

| Match Type | Description | Supported NetFlow Versions |
|---------------------|--|----------------------------|
| Egress Des MAC | Configures the post source MAC address as a non-key field. | IPFIX |
| Ingress Des MAC | Configures the IEEE 802 destination MAC address as a non-key field. | IPFIX |
| Source MAC | Configures the IEEE 802 source MAC address as a non-key field. | v9 and IPFIX |
| Timestamp | | |
| Flow End Millisec | Configures the absolute timestamp of the last packet of current flow in milliseconds as a non-key field. | IPFIX |
| Flow End Sec | Configures the flow start SysUp time as a non-key field. | IPFIX |
| Flow End Time | Configures the flow end SysUp time as a non-key field. | v9 and IPFIX |
| Flow Start Millisec | Configures the value of the IP Precedence as a non-key field. | IPFIX |
| Flow Start Sec | Configures the absolute timestamp of the first packet of this flow as a non-key field. | IPFIX |
| Flow Startup Time | Configures the flow start SysUp time as a non-key field. | v9 and IPFIX |
| Flow | | |
| Flow End Reason | Configures the reason for Flow termination as a non-key field. | IPFIX |
| IPv4 | | |
| ICMP Type Code | Configures the type and code of the IPv4 ICMP message as a non-key field. | v9 and IPFIX |
| IPv4 Dest IP | Configures the IPv4 destination address in the IP packet header as a non-key field. | v9 and IPFIX |
| IPv4 ICMP Code | Configures the code of the IPv4 ICMP message as a non-key field. | IPFIX |
| IPv4 ICMP Type | Configures the type of the IPv4 ICMP message as a non-key field. | IPFIX |
| IPv4 Options | Configures the IPv4 options in the packets of the current flow as a non-key field. | IPFIX |
| IPv4 Src IP | Configures the IPv6 source address in the IP packet header as a non-key field. | v9 and IPFIX |
| IPv4 Total Length | Configures the total length of the IPv4 packet as a non-key field. | IPFIX |
| Network | | |

Table 3-9: Collect/Non-Key Elements

| Match Type | Description | Supported NetFlow Versions |
|-------------------|---|----------------------------|
| IP CoS | Configures the IP Class Of Service (CoS) as a key field. | v9 |
| IP Protocol | Configures the value of the protocol number in the IP packet header as a key field. | v9 |
| IP Version | Configures the IP version field in the IP packet header as a key field. | v9 |
| IPv6 | | |
| IPv6 Dest IP | Configures the IPv6 destination address in the IP packet header as a key field. | v9 |
| IPv6 Flow Label | Configures the value of the IPv6 flow label field in the IP packet header as a key field. | v9 |
| IPv6 Src IP | Configures the IPv6 source address in the IP packet header as a key field. | v9 |
| Transport | | |
| L4 Dest Port | Configures the destination port identifier in the transport header as a non-key field. | v9 and IPFIX |
| L4 Src Port | Configures the source port identifier in the transport header as a non-key field. | v9 and IPFIX |
| TCP Ack Number | Configures the acknowledgment number in the TCP header as a non-key field. | IPFIX |
| TCP Dest Port | Configures the destination port identifier in the TCP header as a non-key field. | IPFIX |
| TCP Flags | Configures the TCP control bits observed for the packets of this flow as a non-key field. | v9 and IPFIX |
| TCP Header Length | Configures the length of the TCP header as a non-key field. | IPFIX |
| TCP Seq Number | Configures the sequence number in the TCP header as a non-key field. | IPFIX |
| TCP Src Port | Configures the source port identifier in the TCP header as a non-key field. | IPFIX |
| TCP Urgent | Configures the urgent pointer in the TCP header as a non-key field. | IPFIX |
| TCP Window Size | Configures the window field in the TCP header as a non-key field. | IPFIX |
| UDP Dest Port | Configures the destination port identifier in the UDP header as a non-key field. | IPFIX |
| UDP Src Port | Configures the source port identifier in the UDP header as a non-key field. | IPFIX |

Adding a Version 5 NetFlow Application

To add a version 5 NetFlow application:

1. Drag and drop **NetFlow** from **APPLICATIONS** to the graphical workspace.

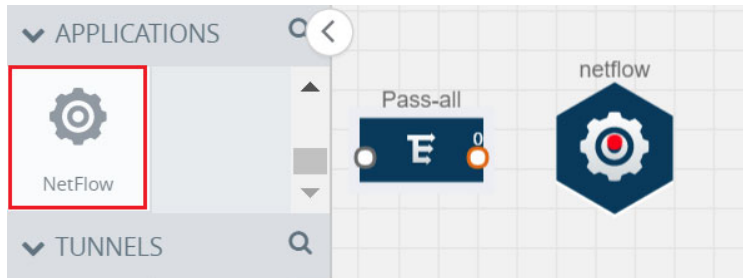


Figure 3-28: Dragging the NetFlow Application

2. Click the NetFlow application and select **Details**. A quick view is displayed for configuring the NetFlow application.

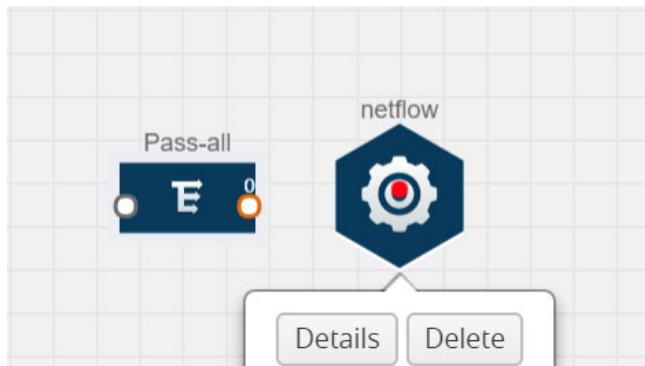


Figure 3-29: Selecting Details

3. In the **Alias** field, enter a name for the v5 NetFlow application.

ApplicationSave

| | |
|------------------------|---|
| Application | netflow |
| Alias | <input type="text" value="Netflow_V5"/> |
| State | <input checked="" type="radio"/> On <input type="radio"/> Off |
| NetFlow version | <input type="text" value="v5"/> |
| Active cache timeout | <input type="text" value="1800"/> |
| Inactive cache timeout | <input type="text" value="15"/> |

Figure 3-30: Viewing v5 NetFlow Application Quick View

4. For State, select the **On** check box to determine that the application is currently running. Select the **Off** check box to determine that the application is currently not running. The state can be changed at anytime whenever required.

5. From the **NetFlow version** drop-down list, select v5.
6. In **Active cache timeout**, enter the number of seconds that an active flow record must remain in the cache before it is exported and removed. The default value is 1800 seconds.
7. In **Inactive cache timeout**, enter the number of seconds an inactive flow record must remain in the cache before it times out. The default value is 15 seconds.
8. Click **Save**.

For some examples demonstrating the NetFlow application configuration in the GigaVUE V Series nodes, refer to [NetFlow Examples on page 86](#).

Adding a Version 9 and IPFIX NetFlow Application

To add a v9 and IPFIX NetFlow application:

1. Drag and drop **NetFlow** from **APPLICATIONS** to the graphical workspace.

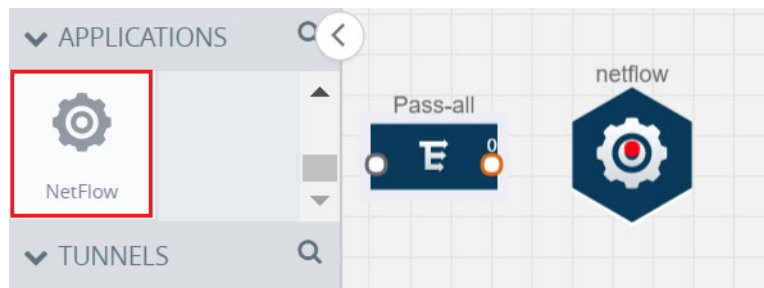


Figure 3-31: Dragging the NetFlow Application

2. Click the NetFlow application and select **Details**. A quick view is displayed for configuring the NetFlow application.



Figure 3-32: Selecting NetFlow Details

3. In the **Alias** field, enter a name for the NetFlow application.

The screenshot shows a configuration interface for a NetFlow application. At the top, there is an orange header with the word "Application" and a "Save" button. Below the header, the configuration is organized into several sections:

- Application:** A text field containing "netflow".
- Alias:** A text field containing "Netflow_IPFIX".
- State:** Two radio buttons, "On" (selected) and "Off".
- NetFlow version:** A dropdown menu currently showing "IPFIX".
- Source Id:** A text field containing "1".
- Match fields:** A dropdown menu showing "L4 Src Port" and "IPv4 Src IP".
- Collect fields:** A dropdown menu showing a list of fields: "Byte Count", "Packet Count", "TCP Flags", "IPv4 Src IP", "Source MAC", "Destination MAC", "IP Version", "Flow Start Sec", "UDP Src Port", "UDP Dest Port", "IP Header Length", "IPv4 Total Length", and "IP Total Length".
- Active cache timeout:** A text field containing "1800".
- Inactive cache timeout:** A text field containing "15".
- Template refresh interval:** A text field containing "1800".

Figure 3-33: Viewing NetFlow Application Quick View

4. For State, select the **On** check box to determine that the application is generating NetFlow records from the packets coming from the G-vTAP agents. Select the **Off** check box to determine that the application is not currently generating NetFlow records. The state can be changed at anytime whenever required.
5. From the **NetFlow version** drop-down list, select the version you want to use to generate the NetFlow records. The default version selected is v5.
6. In the **Source ID** field, enter the observation domain to isolate the traffic. The NetFlow application uses source ID to segregate the records into categories. For example, you can assign source ID 1 for traffic coming over TCP. This results in generating a separate NetFlow record for TCP data. Similarly, you can assign Source ID 2 for traffic coming over UDP. This results in generating a separate NetFlow record for UDP data.
7. From the **Match fields** drop-down list, select the parameters that identify what you want to collect from the incoming packets. The Match fields displayed in the drop-down list are based on the NetFlow version selected in step 5. Refer to [Match/Key Fields on page 75](#).
8. From the **Collect fields** drop-down list, select the parameters that identify what you want to collect from the NetFlow records. The Collect fields displayed in the drop-down list are based on the NetFlow version selected in step 5. Refer to [Collect/Non-Key Fields on page 77](#).
9. In **Active cache timeout**, enter the number of seconds that an active flow record must remain in the cache before it is exported and removed. The default value is 1800 seconds.

10. In **Inactive cache timeout**, enter the number of seconds an inactive flow record must remain in the cache before it times out. The default value is 15 seconds.
11. In **Template refresh interval**, enter the frequency at which the template must be sent to the tool. The default value is 1800 seconds.
12. Click **Save**.

For some examples demonstrating the NetFlow application configuration in the GigaVUE V Series nodes, refer to [NetFlow Examples on page 86](#).

Network Address Translation (NAT)

NAT allows the NetFlow records to be directly transmitted to a collector without a tunnel. It lets you configure the destination IP of one or more collectors and the source IP of the GigaVUE V Series node interface through which the NetFlow records are sent out. The NetFlow records are exported to the collector over UDP protocol with the configurable source IP and destination IP.

NOTE: Only one NAT can be added per monitoring session.

Adding NAT

To add a NAT device:

1. Drag and drop **NAT** to the graphical workspace.

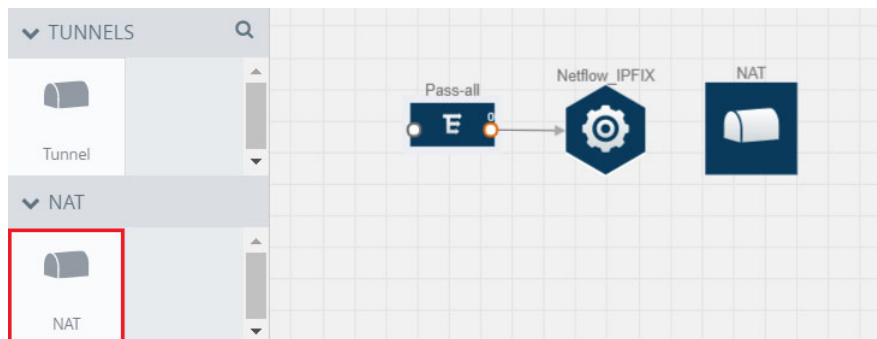


Figure 3-34: Adding NAT

2. Click **NAT** and select **Details**. A quick view is displayed for configuring a NAT device.

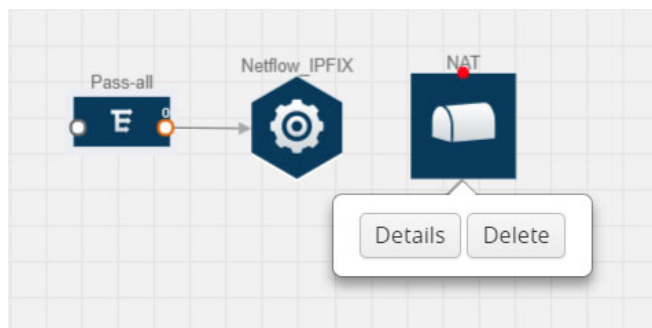


Figure 3-35: Selecting Details

3. In the **Alias** field, enter a name for the NAT device.

Alias: NAT-to-Splunk

Local Subnet: 192.168.50.0/24

Routes:

| | Destination IP | Node Interface Subnet Cidr |
|---|----------------|----------------------------|
| + | | |
| x | 10.0.2.189 | 10.0.1.0/24 |

Figure 3-36: Configuring NAT

4. (Optional) In **Local Subnet**, enter a local subnet IP address that you want to assign to the NetFlow record. By default, GigaVUE V Series node auto generates a default local subnet. The subnet that you enter will override the default subnet.
5. (Optional) In **Routes**, define the routes to send the flow records to NetFlow collectors. Enter the following:
 - a. In **Destination IP**, enter the IP address of the NetFlow collector. For example, if Splunk is the NetFlow collector, enter the IP address of Splunk.
 - b. In **Node Interface Subnet CIDR**, enter the GigaVUE V Series node interface subnet Cidr for routing the NetFlow records out from GigaVUE V Series node.
 - c. Click **+** to add more routes. Repeat steps a and b to enter the destination IP and node interface CIDR.
6. Click **Save**.

Linking a NetFlow Application to NAT

To create a link from a NetFlow application to a NAT device:

1. Drag and drop a link from the NetFlow application to a NAT device. A Link quick view is displayed. It is a header transformation operation that lets you configure the IPv4 destination IP of the NetFlow collector.

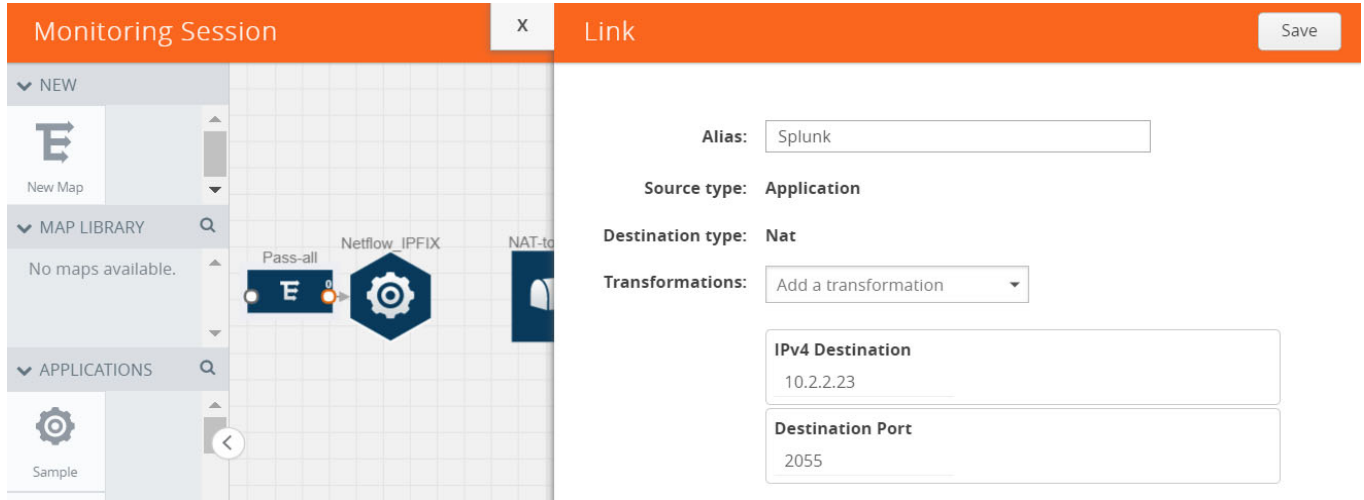


Figure 3-37: Creating a Link from NetFlow to NAT

2. In the **Alias** field, enter a name for the link.
3. From the **Transformations** drop-down list, select any one of the header transformations:
 - IPv4 Destination
 - ToS
 - Destination Port

NOTE: Only the above three header transformations are allowed on the link from the NetFlow application to a NAT device.

4. In **IPv4 Destination**, enter the IP address of the NetFlow collector.
5. (Optional) By default, the Destination Port is 2055. To change the destination port, enter a port number.
6. Click **Save**. The transformed link is displayed in Orange.

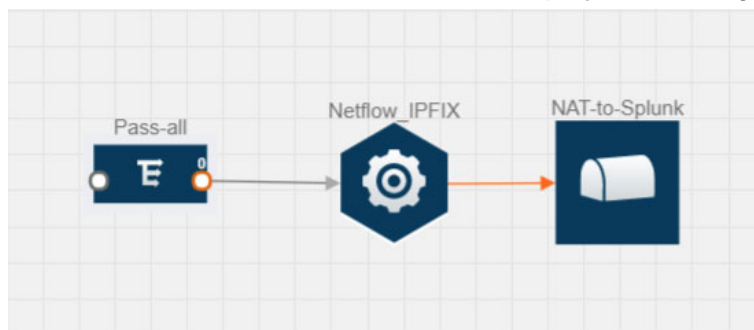


Figure 3-38: Linking NetFlow to NAT

- Repeat steps 7 to 10 to send additional NetFlow records to NAT.

NetFlow Examples

This section provides some examples demonstrating the NetFlow application configuration in the GigaVUE V Series nodes.

- [Example 1 on page 86](#)
- [Example 2 on page 90](#)

Example 1

In this example, a pass all map is created and the entire traffic from a VPC is sent to a tool for full packet inspection. At the same time, a NetFlow application is added to generate flow records for flow inspection.

- Create a monitoring session. For steps, refer to [Creating a Monitoring Session on page 57](#).

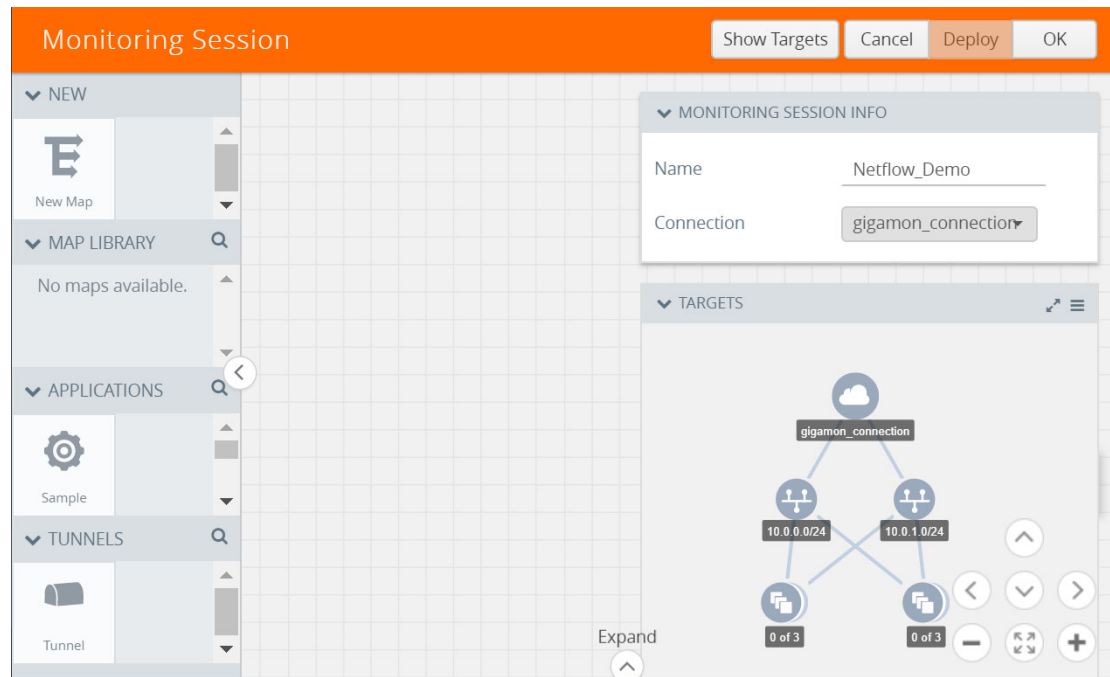


Figure 3-39: Creating a Monitoring Session

- In the monitoring session, create a Pass all map. A pass all map sends all the traffic received from the G-vTAP agents to the tunnel endpoint or NAT. For steps, refer to [Creating a Map on page 61](#).

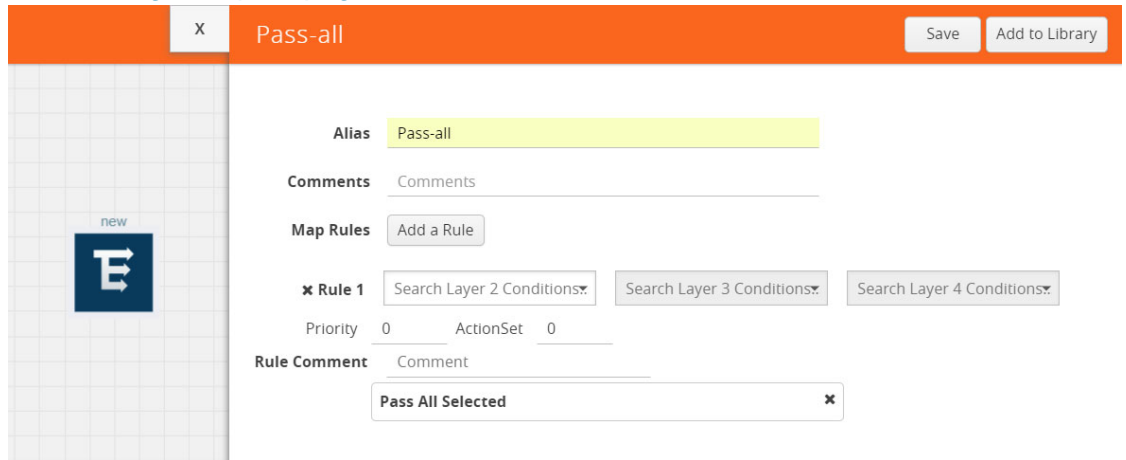


Figure 3-40: Creating a Pass All Map

- Drag and drop a tunnel from **Tunnels**. A tunnel encapsulates the flow records and then sends them to the tools for full packet inspection.



Figure 3-41: Adding a Tunnel

- Create a link from the Pass-all map to the tunnel endpoint. The traffic from the Pass-all map is forwarded to the tunnel endpoint that is connected to a tool.

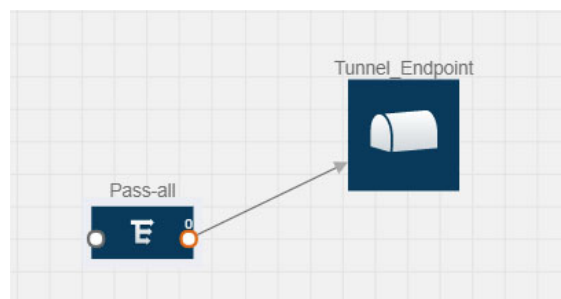


Figure 3-42: Creating a Link from Pass-all Map to Tunnel_Endpoint

5. Drag and drop a v5 NetFlow application.




Figure 3-43: Adding a link from Pass-all Map to Tunnel_Endpoint

6. Click the NetFlow application and select **Details**. The Application quick view is displayed. For steps to configure the v5 NetFlow application, refer to [Adding a Version 5 NetFlow Application on page 80](#).

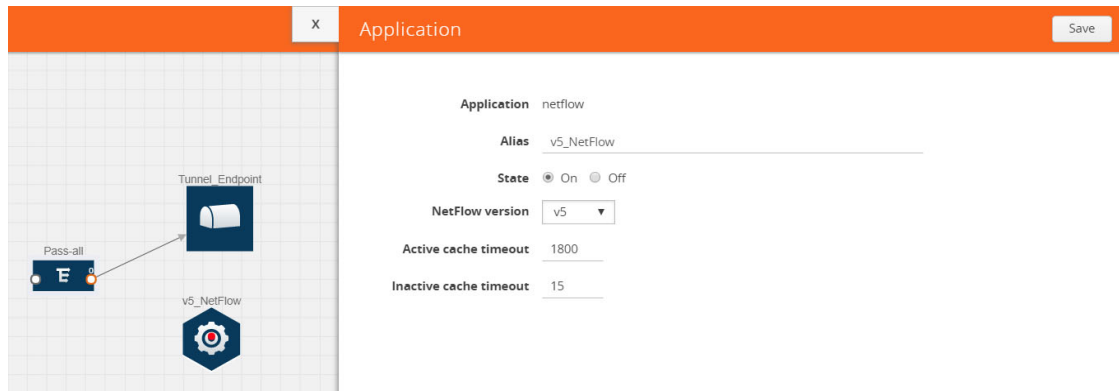


Figure 3-44: Configuring the NetFlow Application

7. Create a link from the Pass all map to the v5 NetFlow application.

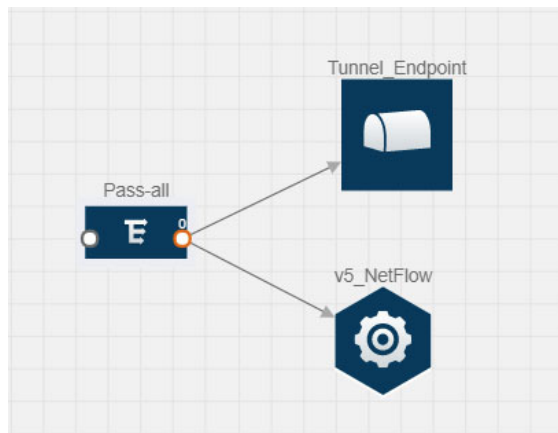


Figure 3-45: Adding a link from Pass-all Map to v5_NetFlow

8. Drag and drop **NAT** to the graphical workspace. A quick view to configure the NAT device is displayed. For steps to configure the NAT device, refer to [Adding NAT on page 83](#).

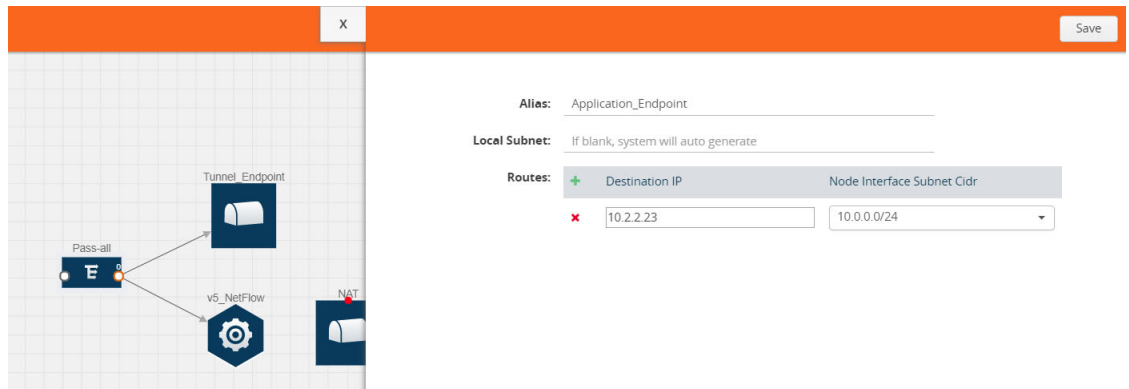


Figure 3-46: Adding a NAT Device

9. Create a link from the v5 NetFlow application to NAT. The link must be configured with the destination IP address of the NetFlow collector and the GigaVUE V Series node interface. For steps to configure the link, refer to [Linking a NetFlow Application to NAT on page 85](#).

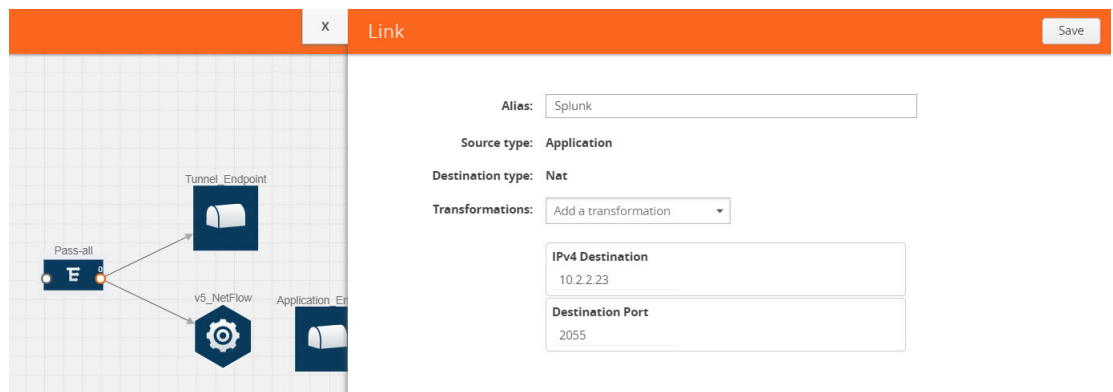


Figure 3-47: Adding a Link from v5 NetFlow Application to NAT

- Click on the link created from the v5 NetFlow application to NAT. The information about the NetFlow collector destination IP and port is displayed.

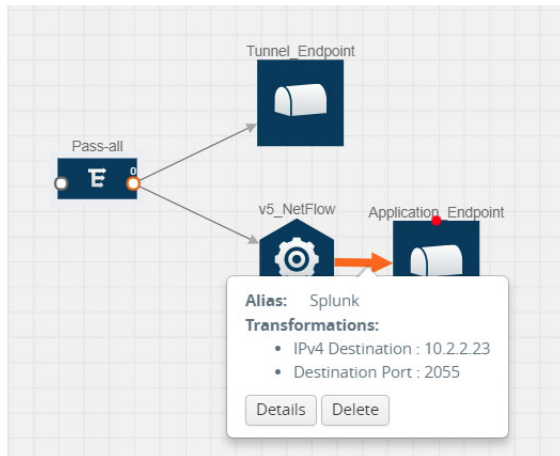


Figure 3-48: Viewing the Transformation Dialog Box

Example 2

In this example, two different versions of NetFlow applications are created. One map is configured to send the TCP packets to the v9 NetFlow application. Another map is configured to send the UDP packets to the IPFIX NetFlow application. The flow records generated from v9 and IPFIX NetFlow application are sent to NAT.

- Create a monitoring session. For steps, refer to [Creating a Monitoring Session on page 57](#).

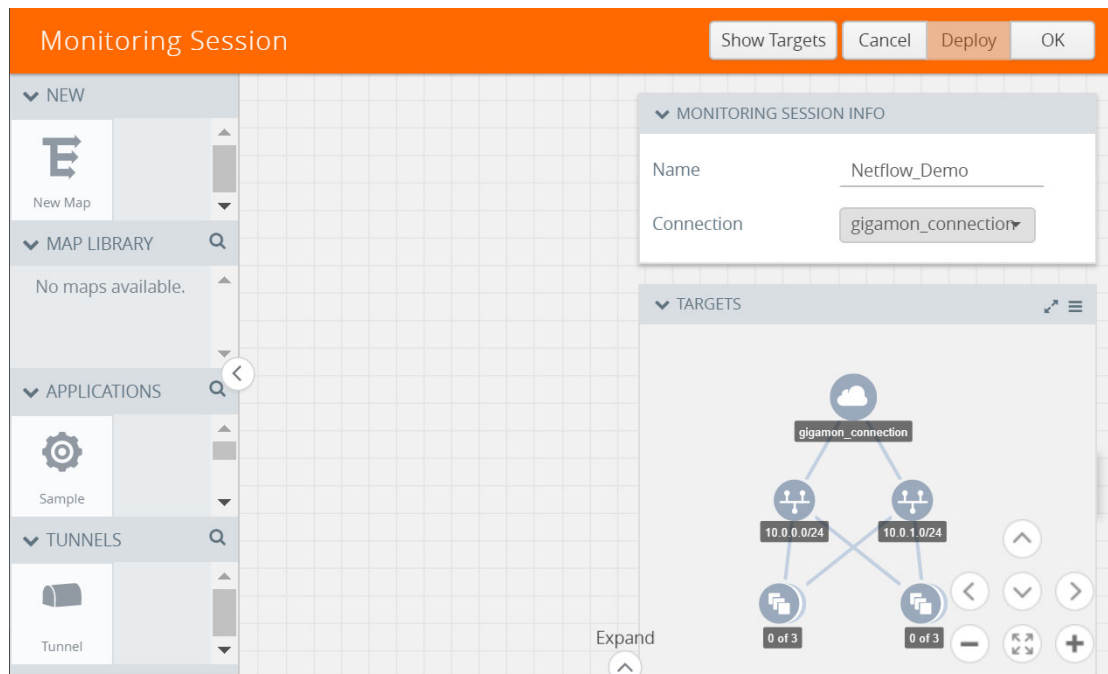


Figure 3-49: Creating a Monitoring Session

2. Create a map rule to filter the TCP packets. For steps on creating a map, refer to [Creating a Map on page 61](#).

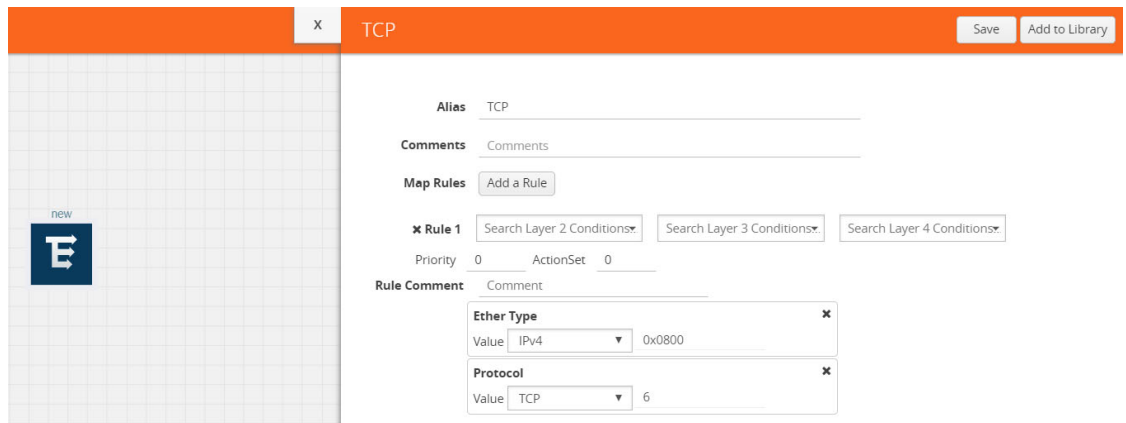


Figure 3-50: Creating a TCP Map

3. Create another map rule to filter the UDP packets.

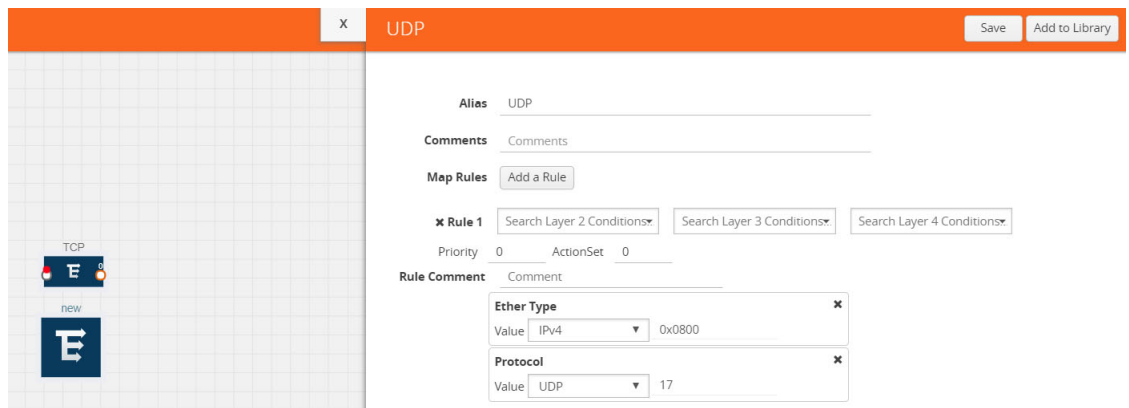


Figure 3-51: Creating a UDP Map

4. Create another map rule to filter the UDP packets.

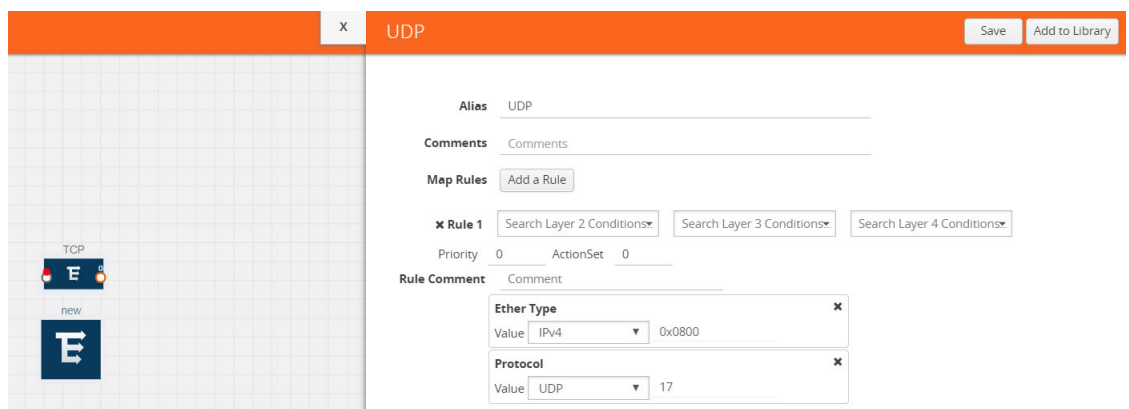


Figure 3-52: Creating a Map to Filter UDP Packets

5. Drag and drop a NetFlow application. Choose v9 as the NetFlow version. Select the match and the collect fields. For steps to configure the v9 NetFlow application, refer to [Adding a Version 9 and IPFIX NetFlow Application on page 81](#).

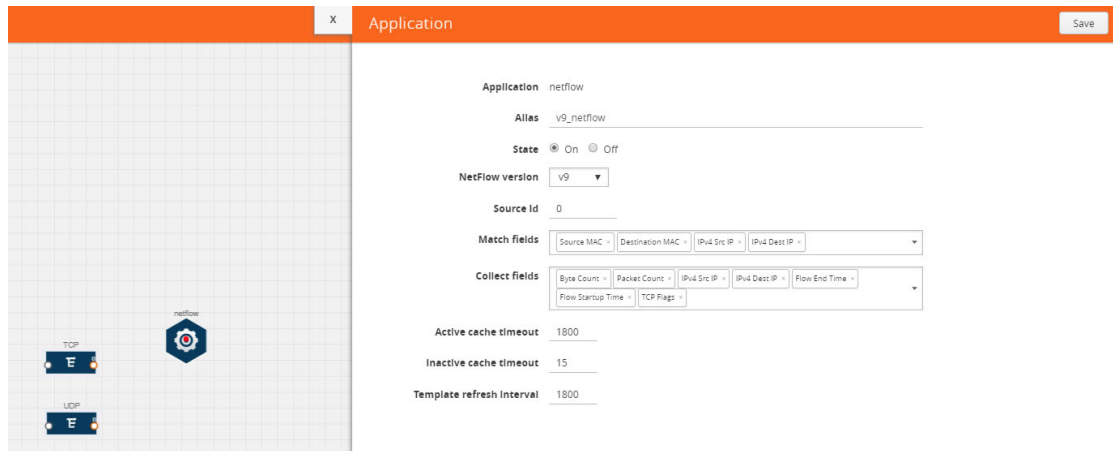


Figure 3-53: Adding a v9 NetFlow Application

6. Create a link from the TCP map to the v9 NetFlow application.




Figure 3-54: Adding a Link from the TCP Map to v9_netflow Application

7. Drag and drop a NetFlow application. Choose IPFIX as the NetFlow version. Select the match and the collect fields. For steps to configure the IPFIX NetFlow application, refer to [Adding a Version 9 and IPFIX NetFlow Application on page 81](#).

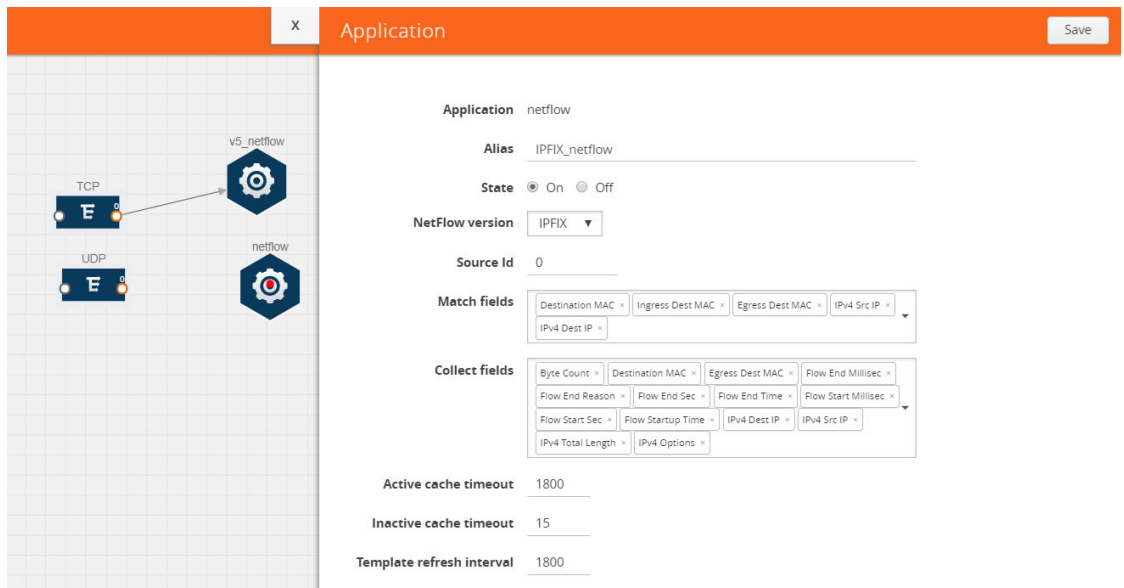


Figure 3-55: Adding a IPFIX NetFlow Application

8. Create a link from the UDP map to the IPFIX NetFlow application.

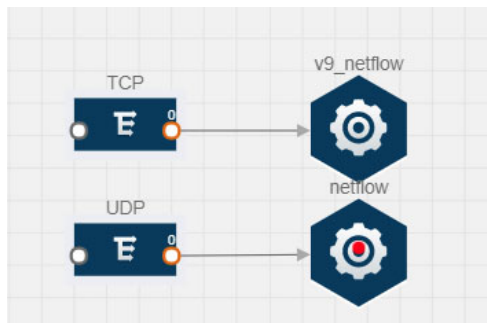


Figure 3-56: Adding a Link from the UDP Map to the IPFIX NetFlow Application

9. Drag and drop a NAT. A quick view to configure the NAT is displayed. For steps to configure a NAT, refer to [Adding NAT on page 83](#).

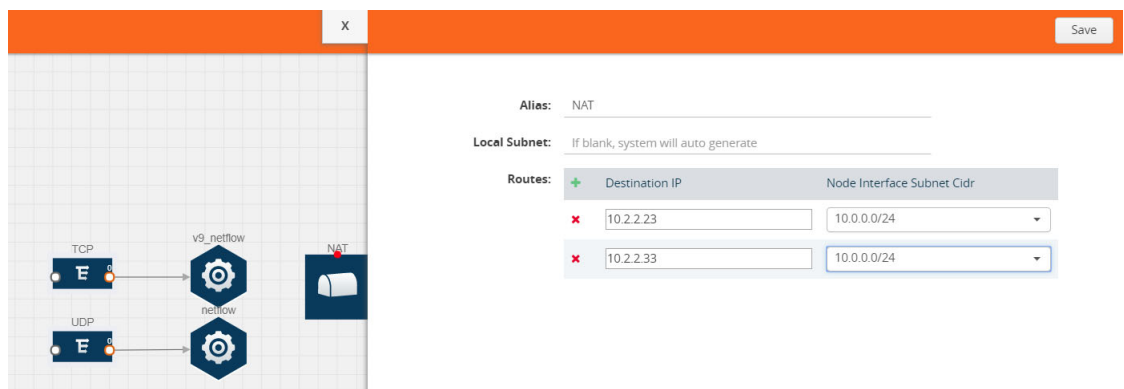


Figure 3-57: Adding a NAT

10. Create a link from the v9 NetFlow application to the NAT.

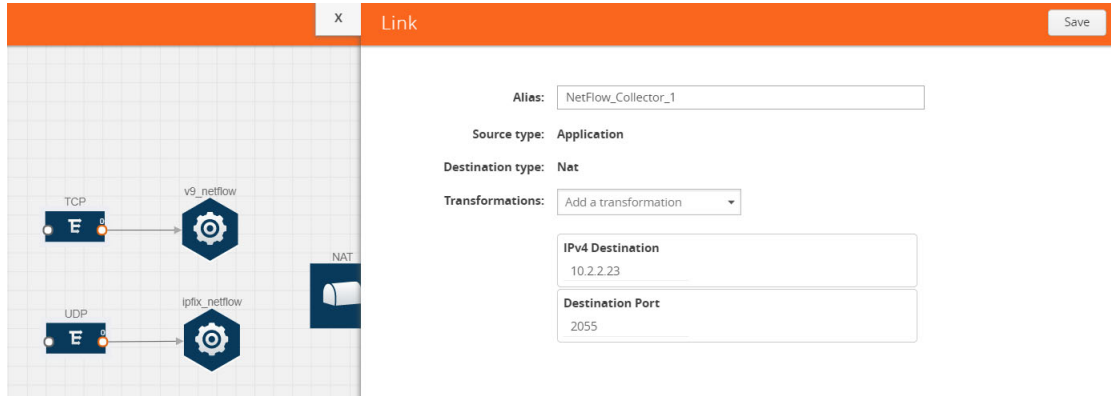


Figure 3-58: Adding a Link from NetFlow Application to NAT

11. Create another link from the IPFIX NetFlow application to the NAT.

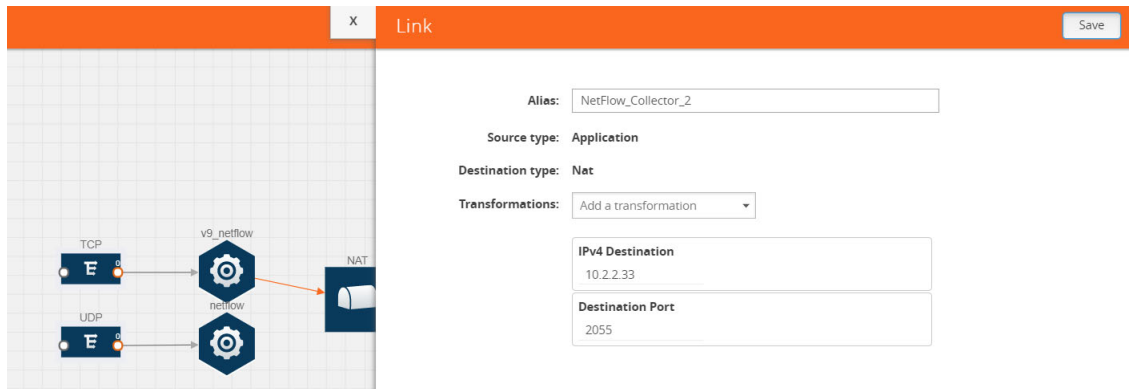


Figure 3-59: Adding a Link from NetFlow Application to NAT

12. Click on the link connecting the NetFlow application to the NAT.

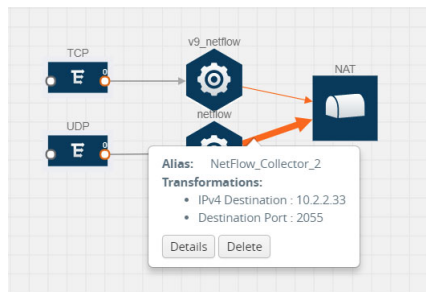


Figure 3-60: Viewing the Header Transformation

Deploying the Monitoring Session

To deploy the monitoring session:

1. Drag and drop one or more maps from the **MAP Library** to the workspace.
2. (Optional) To add Inclusion and Exclusion maps, drag and drop the maps from the Map Library to their respective section at the bottom of the workspace.

3. (Optional) Drag and drop one or more applications from the APPLICATIONS section to the workspace.

NOTE: For information about adding applications to the workspace, refer to [Adding Applications to the Monitoring Session on page 69](#).

4. Drag and drop one or more tunnels from the TUNNELS section.

Figure 3-61 illustrates three maps, one exclusion map, one application, and two tunnel endpoints that have been dragged and dropped to the workspace. The tunnel endpoints are named `Monitoring_Tool_1` and `Monitoring_Session_2`.

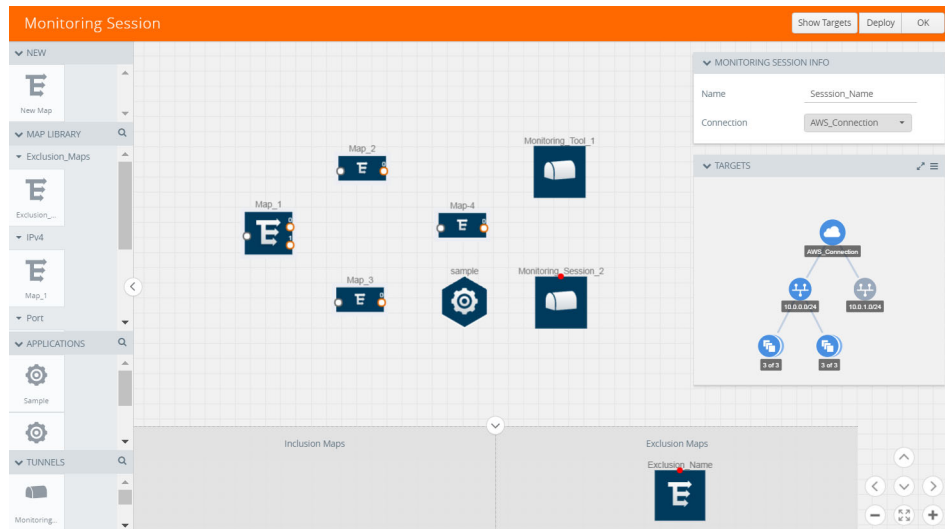


Figure 3-61: Dragging and Dropping the Maps, Applications, and Monitoring Tools

5. Hover your mouse on the map, click the red dot, and drag the arrow over to another map, application, or tunnel. Refer to [Figure 3-62 on page 96](#).

NOTE: You can drag multiple arrows from a single map and connect them to different maps and applications.

6. Hover your mouse on the application, click the red dot, and drag the arrow over to the tunnel endpoints.

In [Figure 3-62](#), the traffic matching the rules in each action set is routed to maps, applications, or monitoring tools.

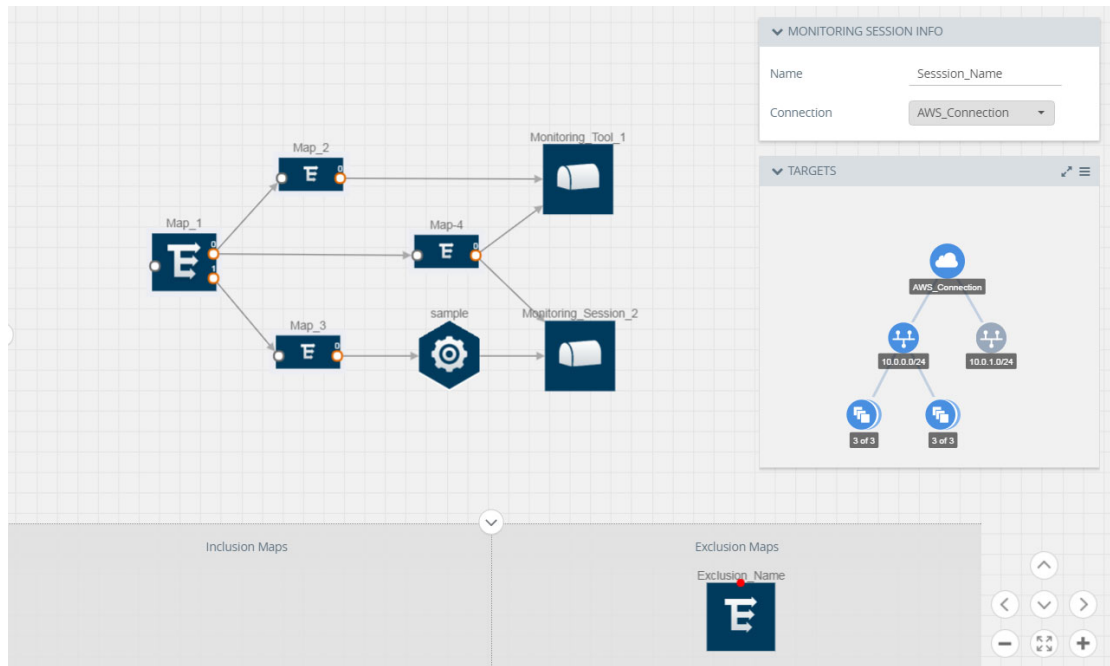


Figure 3-62: Connecting the Maps, Applications, and Monitoring Tools

7. Click **Show Targets** to view details about the subnets and monitoring instances. The instances and the subnets that are being monitored are highlighted in blue.
8. Click **Deploy** to deploy the monitoring session.

The status is displayed as **Success** in the Monitoring Sessions page. The session is successfully deployed on all V Series nodes and G-vTAP agents or TaaS. Click on the status link in the Status column on the Monitoring Session page to view the Monitoring Session Deployment Report.

| Monitoring Sessions | | | | | |
|--------------------------|------------------------|-------------|--------------|--|----------------------|
| <input type="checkbox"/> | Name | Connection | # of Targets | Status | Statistics |
| <input type="checkbox"/> | Connection1_Monitoring | Connection1 | 2 | ● Success | View |

Total Items : 1

When you click on the Status link, the Deployment Report is displayed. Refer to Figure 3-63.

| Name | Connection |
|---|------------|
| <input type="checkbox"/> Example1_Monitor | Example1 |

| | |
|--|---------------------|
| Monitoring Session Alias : | Example1_Monitor |
| Deployment Status : | Success |
| Operation : | deploy |
| Start Time : | 2017-08-08 15:14:58 |
| End Time : | 2017-08-08 15:14:58 |
| General Failure Messages : | |
| NONE | |
| Selected Targets : | 2 |
| Target Deployment Successes : | 2 |
| Target Deployment Failures : | 0 |
| Nic License Failures : | 0 |
| V-Series Node Deployment Successes : | 1 |
| V-Series Node Deployment Failures : | 0 |
| Unselected Targets : | 0 |
| Target Undeployment Successes : | 0 |
| Target Undeployment Failures : | 0 |
| V-Series Node Undeployment Successes : | 0 |
| V-Series Node Undeployment Failures : | 0 |

Figure 3-63: Monitoring Session Deployment Report

If the monitoring session is not deployed properly, then one of the following errors is displayed in the Status column.

- **Partial Success**—The session is not deployed on one or more instances due to G-vTAP or TaaS or V Series node failure.
- **Failure**—The session is not deployed on any of the V Series nodes and G-vTAP agents or TaaS.

If there was an error in deploying, the Monitoring Session Deployment Report will display the information about it.

The Monitoring Session page also has the following buttons:

- **Redeploy**—Redeploys the selected monitoring session.
- **Undeploy**—Undeploys the selected monitoring session.
- **Clone**—Duplicates the selected monitoring session.
- **Edit**—Opens the Edit page for the selected monitoring session.
- **Delete**—Deletes the selected monitoring session.

Adding Header Transformations

Header transformation is performed on a link in a monitoring session. You can select a link and modify the packet header before they are sent to the destination. The header transformation feature is supported only with GigaVUE V Series node version 1.3-1 and above.

Header transformations are used to perform many simple operations on the network packets. The source and destination MAC addresses, port numbers, and IP addresses can be masked to prevent the information from being exposed to the monitoring tools.

The monitoring tools cannot always distinguish the traffic coming from multiple VPCs with the same subnet range. You can add VLAN ID, VLAN priority, and DSCP bits to the header for distinguishing the traffic coming from multiple VPCs with the same subnet range.

In addition to header transformation, GigaVUE V Series node allows you to add multiple links to the same destination. Using multiple links, you can send duplicate packets or various transformed packets to the same destination. For example, you can add different L2GRE or VXLAN tunnel IDs to the packets and send them to different applications within the same tool.

In [Figure 3-64 on page 98](#), the filtered packets from the ICMP map are sent to the same tunnel endpoint in four different links. In each link, you can apply one or more header transformations. A link with the header transformation applied is displayed in orange. When you mouse over the orange link, a detailed information about the alias and the type of transformation is displayed.



Figure 3-64: Action Set with Multiple Links

GigaVUE V Series node supports the following header transformations:

Table 3-10: Header Transformations

| Option | Description |
|-----------------|--|
| MAC Source | Modify the Ethernet source address. |
| MAC Destination | Modify the Ethernet destination address. |

Table 3-10: Header Transformations

| Option | Description |
|------------------|---|
| VLAN Id | Specify the VLAN ID. |
| VLAN PCP | Specify the VLAN priority. |
| Strip VLAN | Strip the VLAN tag. |
| IPv4 Source | Specify the IPv4 source address. |
| IPv4 Destination | Specify the IPv4 destination address. |
| ToS | Specify the DSCP bits in IPv4 traffic class. |
| Source Port | Specify the UDP, TCP, or SCTP source port. |
| Destination Port | Specify the UDP, TCP, or SCTP destination port. |
| Tunnel ID | Specify the tunnel ID. The tunnel ID header transformation can only be applied on the links with the tunnel endpoint destination. Using Tunnel ID header transformation, the filtered packets can be sent to different applications or programs within the same monitoring tool. |

To add a header transformation:

1. On the Monitoring Session, click the link and select **Details**. The Link quick view is displayed.



Figure 3-65: Opening the Link Quick View

2. From the **Transformations** drop-down list, select one or more header transformations.

NOTE: Do not apply VLAN Id and VLAN PCP transformation types with the Strip VLAN ID transformation type on the same link.

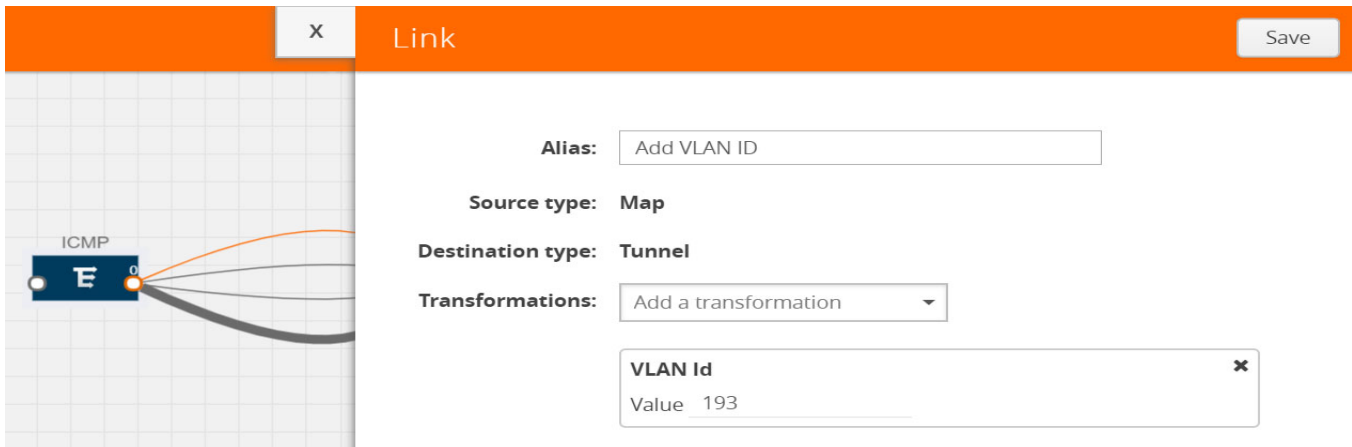


Figure 3-66: Adding Transformation

3. Click **Save**. The selected transformation is applied to the packets passing through the link.
4. Click **Deploy** to deploy the monitoring session.

Viewing the Statistics

The Monitoring Session Statistics page lets you analyze the incoming and outgoing traffic on an hourly, daily, weekly, and monthly basis. The traffic can be viewed based on kilobits/second, megabits/second, or gigabits/second.

On the Monitoring Sessions page, click **View** in the Statistics column to view the Monitoring Session Statistics page.

| Monitoring Sessions | | | | | |
|---|------------------------|-------------|--------------|--|----------------------|
| Redeploy Undeploy New Clone Edit Delete | | | | | |
| <input type="checkbox"/> | Name | Connection | # of Targets | Status | Statistics |
| <input type="checkbox"/> | Connection1_Monitoring | Connection1 | 2 | ● Success | View |

Total Items : 1

The Monitoring Session Statistics page appears where you can analyze incoming and outgoing traffic.


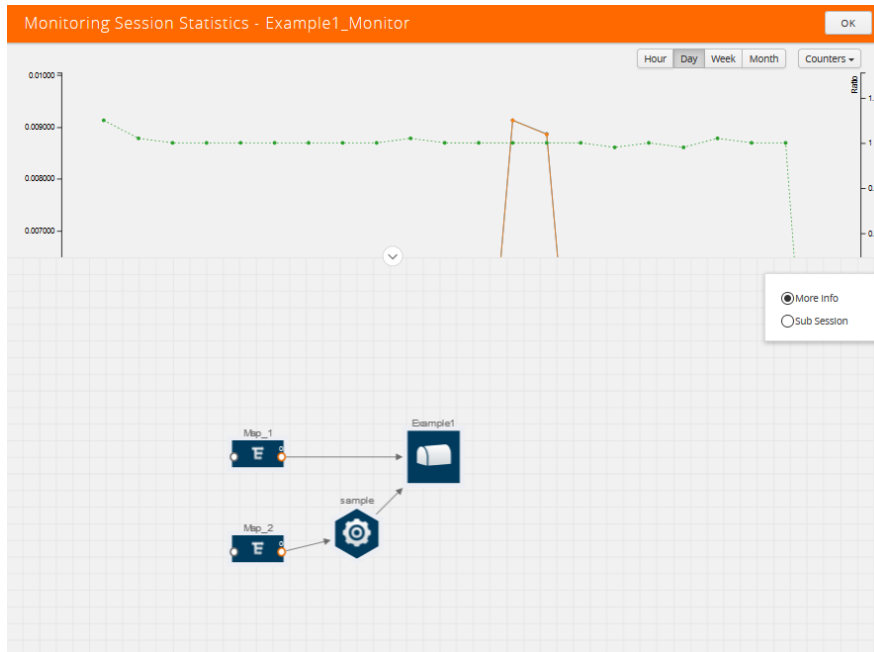


Figure 3-67: Monitoring Session Statistics View

Directly below the graph, you can click on **Incoming Maps**, **Outgoing Maps**, or **Ratio (Out/In)** to view the statistics individually.

At the bottom of the Monitoring Session Statistics page, you can click on **View Monitoring Session Diagram**.

The Monitoring Session Diagram page appears.



On the **Monitoring Session Diagram** page, you can expand any map, application, or tunnel to open a Quick View for that item to see more details about the incoming and outgoing traffic for that item. Refer to [Figure 3-68](#).

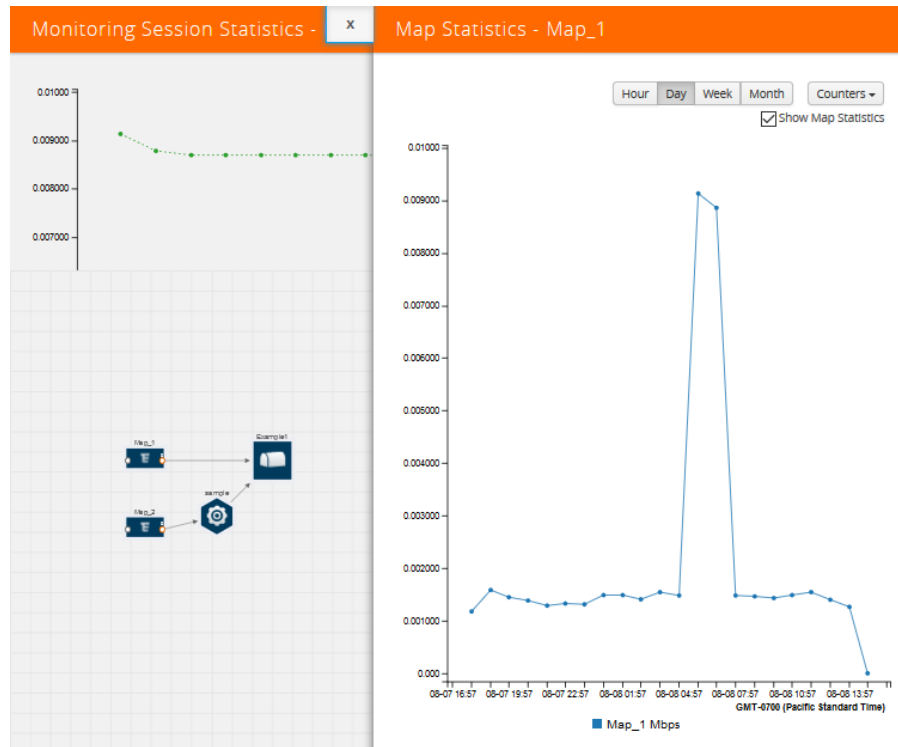


Figure 3-68: Viewing the Map Statistics

[Figure 3-68](#) shows the Map Statistics Quick View with a graph of the traffic for Map_1. You can also scroll down the Map Statistics Quick View to see the Map Rules, Action Sets, and Map Info for this map. You can select Map Rules or Action Sets to view the traffic matching the selected rule on the graph in the Quick View.

Map Rules

| RULE | PRIORITY | ACTION SET | CONDITIONS |
|--|----------|------------|------------------------------|
| <input checked="" type="checkbox"/> Rule 0 | 0 | 0 | etherType 0x0800 IpProto 1 |

Action Sets

Action Set 0

Map Info

Map Alias Map_2

Comment

Viewing the Topology

You can have multiple project connections in GigaVUE-FM. Each project can have multiple monitoring sessions configured within them. You can select the connection and the monitoring session to view the selected subnets and instances in the topology view.

To view the topology diagram:

1. Select **OpenStack > Topology**. The Topology page appears.
2. Select a connection from the **Select connection...** drop-down list. The topology view of the subnets and instances is displayed.




Figure 3-69: Selecting the Connection and Monitoring Session

3. (Optional) Select a monitoring session from the **Select Monitoring Session...** drop-down list. The monitored subnets and instances change to blue.
4. Select one of the following check boxes:
 - **Source**— Displays the topology view of the source target interfaces that are being monitored.
 - **Destination**—Displays the topology view of the destination target interfaces where the traffic is being mirrored.

- **Other**—Displays the topology view of the non-G-vTAP agents such as GigaVUE V Series Controllers, G-vTAP Controllers, monitoring tools, and instances that are being used in the connection.

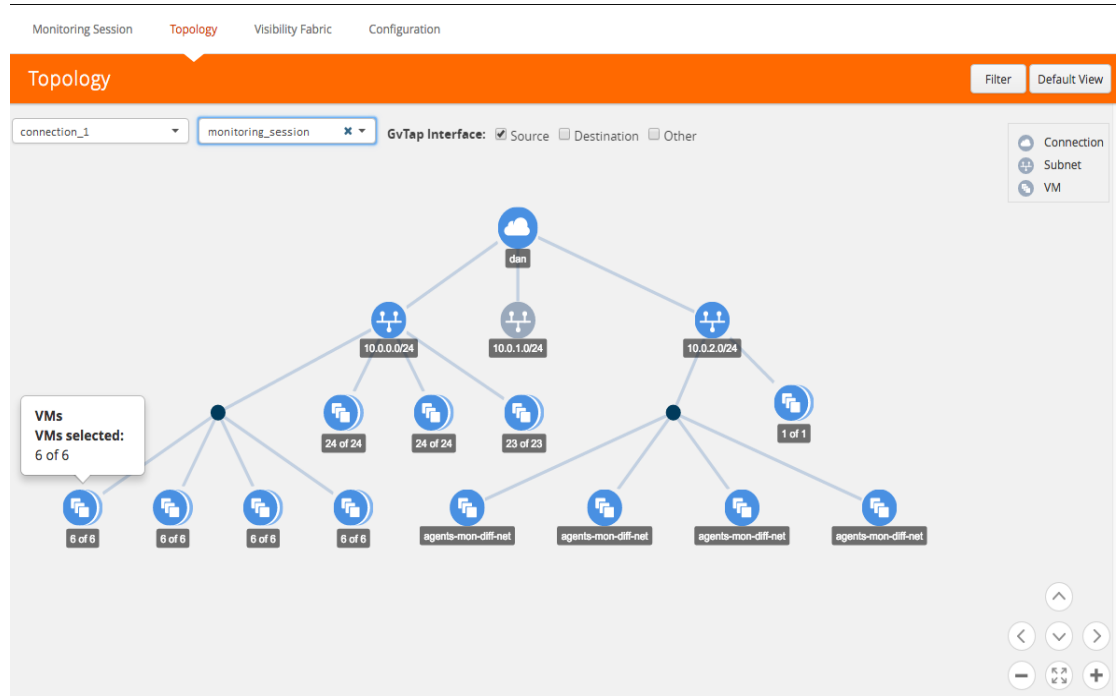


Figure 3-70: Viewing the Topology

5. (Optional) Hover over or click the subnet or VM Group icons to view the subnets or instances present within the group.

In the topology page, you can also do the following:

- Use the **Filter** button to filter the instances based on the VM name, VM IP, Subnet ID, or Subnet IP, and view the topology based on the search results. Refer to [Figure 3-71 on page 106](#).

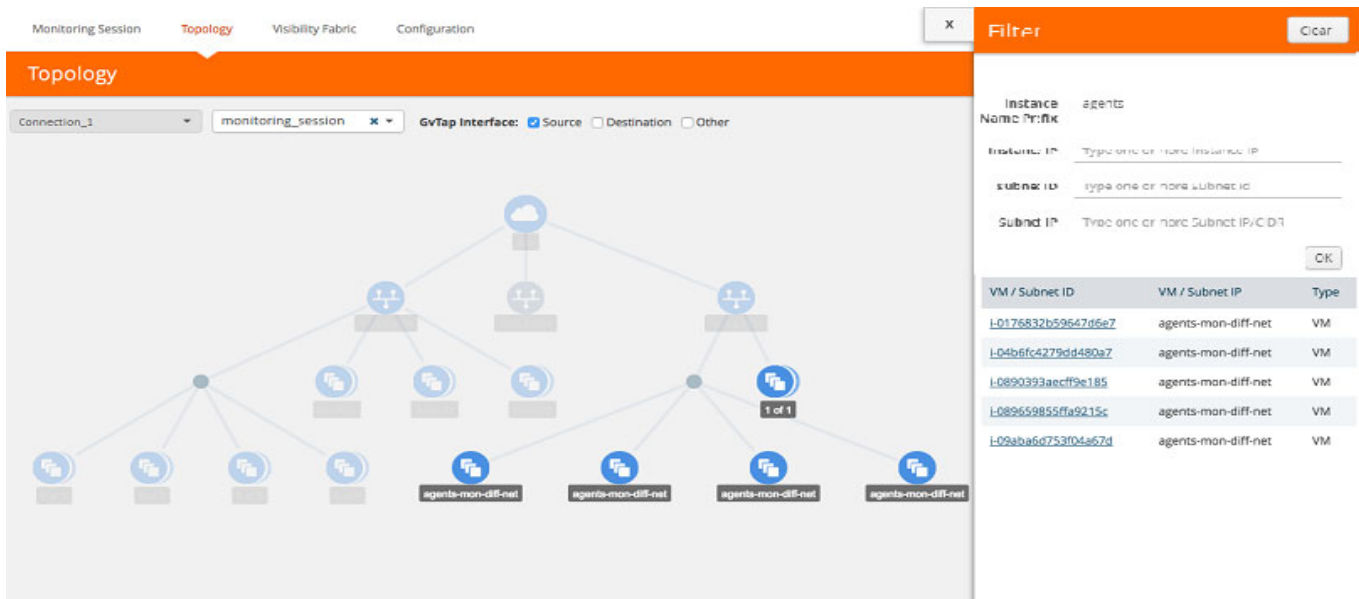


Figure 3-71: Filtering in Topology View

- Use the **Default View** button to view the topology diagram based on the source interfaces of the monitoring instances.
- Use the arrows at the right-bottom corner to move the topology page up, down, left, or right. Click the **Fit-to-Width** icon to fit the topology diagram according to the width of the page.
- Use + or - icons to zoom in and zoom out the topology view.

At the right-bottom corner of the Topology page, there are arrows to move the page up, down, left, or right. There are also plus, minus, and full screen icons to zoom in and zoom out.

On the Topology page, you can also use the **Filter** button to filter instances based on the Instance Name Prefix, Instance IP, Subnet ID, or Subnet IP to view the topology based on the filtered results. Refer to [Figure 3-72 on page 107](#).

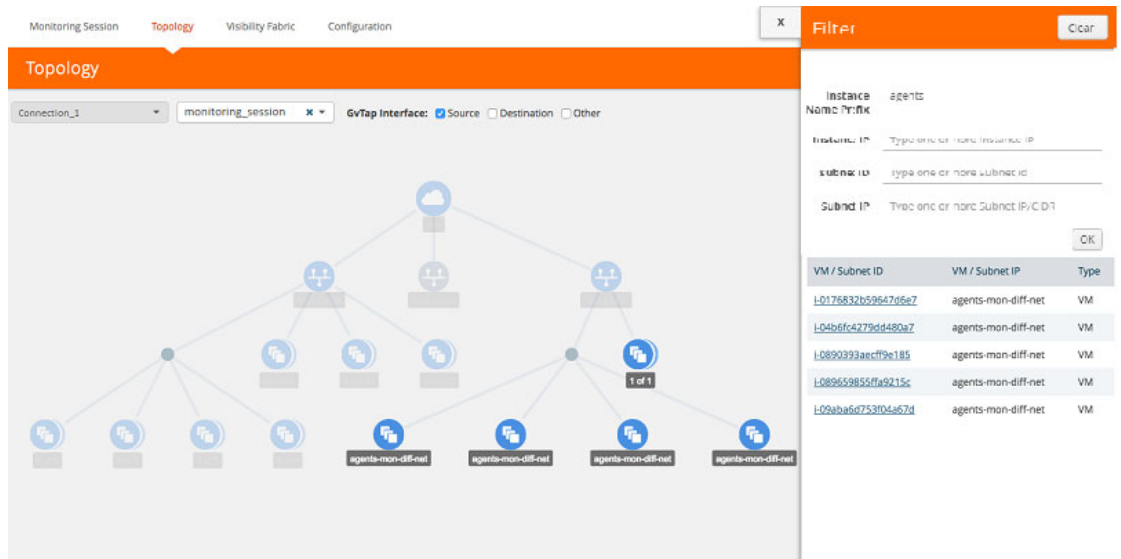



Figure 3-72: Filtering the Topology page

To remove a filter, click the **Clear Filter** button.



Configuring the OpenStack Settings

To configure the OpenStack Settings:

1. In GigaVUE-FM, on the top navigation pane, select **Cloud**.
2. On the left navigation pane, select **OpenStack > Configuration**.
3. Select **Settings** to edit the OpenStack settings. The **Settings** page appears. Refer to [Figure 3-73](#).

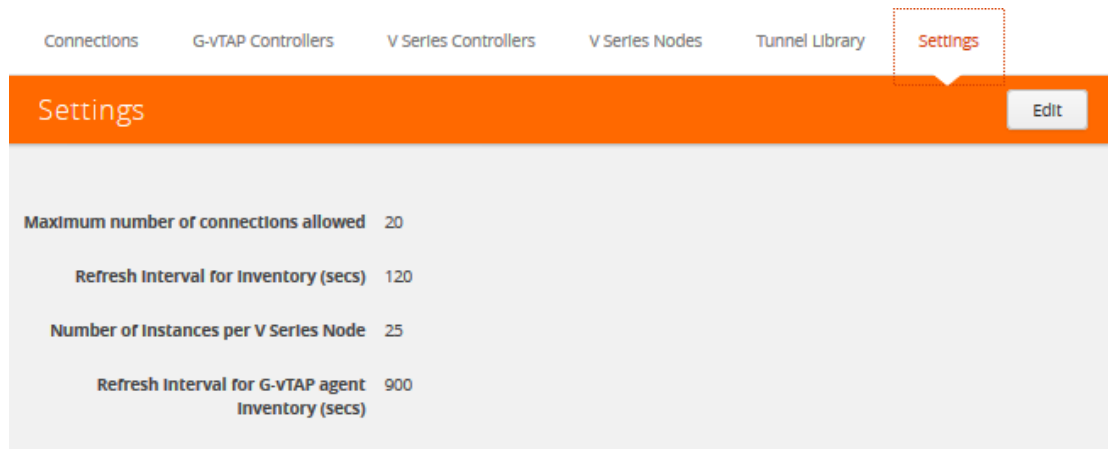


Figure 3-73: Cloud > OpenStack > Connections > Settings

4. Click **Edit** to edit the Settings fields. Refer to [Table 3-11 on page 109](#) for descriptions of the Settings fields:

Table 3-11: OpenStack Settings

| Settings | Description |
|---|--|
| Maximum number of connections allowed | Specifies the maximum number of project connections you can establish in GigaVUE-FM. |
| Refresh interval for instance inventory (secs) | Specifies the frequency for updating the state of cloud instances in OpenStack. |
| Number of instances per V Series Node | Specifies the maximum number of instances that can be assigned to the V Series node. |
| Refresh interval for G-vTAP agent inventory (secs) | Specifies the frequency for discovering the G-vTAP agents available in the project. This is applicable for G-vTAP agents only. |

Compatibility Matrix

This appendix provides information about GigaVUE-FM version compatibility and the features supported in various versions of GigaVUE V Series nodes and G-vTAP agents.

Refer to the following sections for details:

- [GigaVUE-FM Version Compatibility on page 111](#)
- [Supported Features in GigaVUE V Series Nodes on page 111](#)
- [Supported Features in G-vTAP Agents on page 112](#)

GigaVUE-FM Version Compatibility

The following table lists the different versions of GigaSECURE® Cloud solution components available with different versions of GigaVUE-FM.

| GigaVUE-FM | G-vTAP Agent Version | G-vTAP Controller Version | GigaVUE V Series Controller | GigaVUE-V Series Nodes |
|------------|----------------------|---------------------------|-----------------------------|------------------------|
| 5.3.01 | v1.4-1 | v1.4-1 | v1.4-1 | v1.4-1 |
| 5.4.00 | v1.4-1 | v1.4-1 | v1.4-1 | v1.4-1 |

Supported Features in GigaVUE V Series Nodes

The following table lists the features supported in various versions of GigaVUE V Series nodes:

| Features | GigaVUE V Series v1.4-x |
|-----------------------|-------------------------|
| Header Transformation | Yes |
| Multi-link Support | Yes |
| NetFlow Application | Yes |
| NAT Support | Yes |

Supported Features in G-vTAP Agents

The following table lists the features supported in various versions of G-Tap Agents:

| Features | G-vTAP Agent v1.4-x |
|---------------------|---------------------|
| Dual ENI Support | Yes |
| Single ENI Support | Yes |
| VXLAN Support | Yes |
| Agent Pre-filtering | Yes |

Troubleshooting

This section provides the information needed to troubleshoot GigaVUE-FM integration with OpenStack.

OpenStack Connection Failed

The `connFailed` state indicates that the OpenStack connection has failed. Check the following troubleshoot tips to restore the connection:

- Verify if GigaVUE-FM is able to reach the OpenStack cloud controller.
- Check if the OpenStack cloud controller is DNS resolvable from GigaVUE-FM.
- Verify if the region name provided while launching the instance is accurate.
- Ensure that all the security group rules required for communication between GigaVUE-FM and OpenStack cloud controller OR GigaVUE-FM and DNS server are accurately setup.
- Check if the Compute Servers that the nova API returns are reachable from GigaVUE-FM. Refer to [Handshake Alert: unrecognized_name](#) on page 113.

Handshake Alert: unrecognized_name

When setting up the OpenStack connection in GigaVUE-FM, the GigaVUE-FM logs might show a handshake alert: `unrecognized_name` error. This error is related to a Server Name Indication (SNI) error. Starting with Java 7, the JDK does not ignore the unrecognized name warning. To resolve this issue, perform either of the following:

- Fix the configuration on the server where the error is occurring.
- Ignore the warning on the client side (GigaVUE-FM server) by using the Java system property `--Djsse.enableSNIExtension=false` while launching GigaVUE-FM.

Contact support for information on how to use the Java system property. However, this is not recommended for security reasons.

GigaVUE V Series Node or G-vTAP Controller is Unreachable

If GigaVUE V Series node or G-vTAP controller is unreachable, verify the following:

- The correct version of the image is uploaded.
- The network is reachable.

Additional Sources of Information

This appendix provides additional sources of information. Refer to the following sections for details:

- [Documentation](#) on page 115
- [Documentation Feedback](#) on page 115
- [Contacting Technical Support](#) on page 115
- [Contacting Sales](#) on page 116

Documentation

Gigamon provides additional documentation for the GigaVUE H Series on the [Gigamon Customer Portal](#):

| Document | Summary |
|--------------------------------|---|
| GigaVUE-FM User's Guide | Describes how to install, deploy, and operate the GigaVUE® Fabric Manager (GigaVUE-FM) |
| GigaVUE VM User's Guide | Describes how to install, deploy, and operate the GigaVUE® Virtual Manager (GigaVUE-VM) |

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

Refer to <http://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

Table i shows how to reach the Sales Department at Gigamon.

Table i: Sales Contact Information

| | |
|------------------|--|
| Telephone | +1 408.831.4025 |
| Sales | inside.sales@gigamon.com |

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