



**V** **SERIES**

# GigaVUE Cloud Suite for OpenStack Configuration Guide

*Version 5.7*

#### 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](http://www.gigamon.com/legal-trademarks). All other trademarks are the trademarks of their respective owners.

DOCUMENT REVISION – 8/12/19

# CONTENTS

GigaVUE Cloud Suite for OpenStack .....	5
Gigamon GigaVUE Cloud Components .....	6
Traffic Capturing Mechanism .....	8
G-vTAP Agent .....	8
OpenVSwitch (OVS) Mirroring .....	9
OpenVSwitch (OVS) Mirroring + DPDK .....	12
Configuring the Components in OpenStack .....	13
Before You Begin .....	13
Uploading the Images .....	22
Launching the GigaVUE-FM Instance .....	23
G-vTAP Agents .....	33
Configuring the GigaVUE Cloud in OpenStack .....	38
Configuring Monitoring Sessions .....	65
Overview of Visibility Components .....	65
Creating Tunnel Endpoints .....	69
Creating a Monitoring Session .....	71
Configuring the OpenStack Settings .....	134
Compatibility Matrix .....	136
GigaVUE-FM Version Compatibility .....	136
Supported Features in GigaVUE V Series Nodes .....	136
Supported Features in G-vTAP Agents .....	137
Troubleshooting .....	138
OpenStack Connection Failed .....	138
Handshake Alert: unrecognized_name .....	138
GigaVUE V Series Node or G-vTAP Controller is Unreachable .....	139
Additional Sources of Information .....	140
Documentation .....	140
Documentation Feedback .....	140
Contacting Technical Support .....	141
Contacting Sales .....	141
The Gigamon Community .....	141





## GigaVUE Cloud Suite for OpenStack

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.

## Gigamon GigaVUE Cloud Components

The GigaVUE Cloud Suite 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 GigaVUE 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. G-vTAP Controllers
- **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 GigaVUE Cloud using GRE or VxLAN tunnels.

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

*Table 1: Configuration options for Controllers and Nodes*

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

# Traffic Capturing Mechanism

GigaVUE Cloud Suite 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: GigaVUE Cloud Components for OpenStack using G-vTAP on page 8](#) shows a high level architecture of Gigamon GigaVUE Cloud Suite for OpenStack using G-vTAP agents as the source for acquiring the traffic.

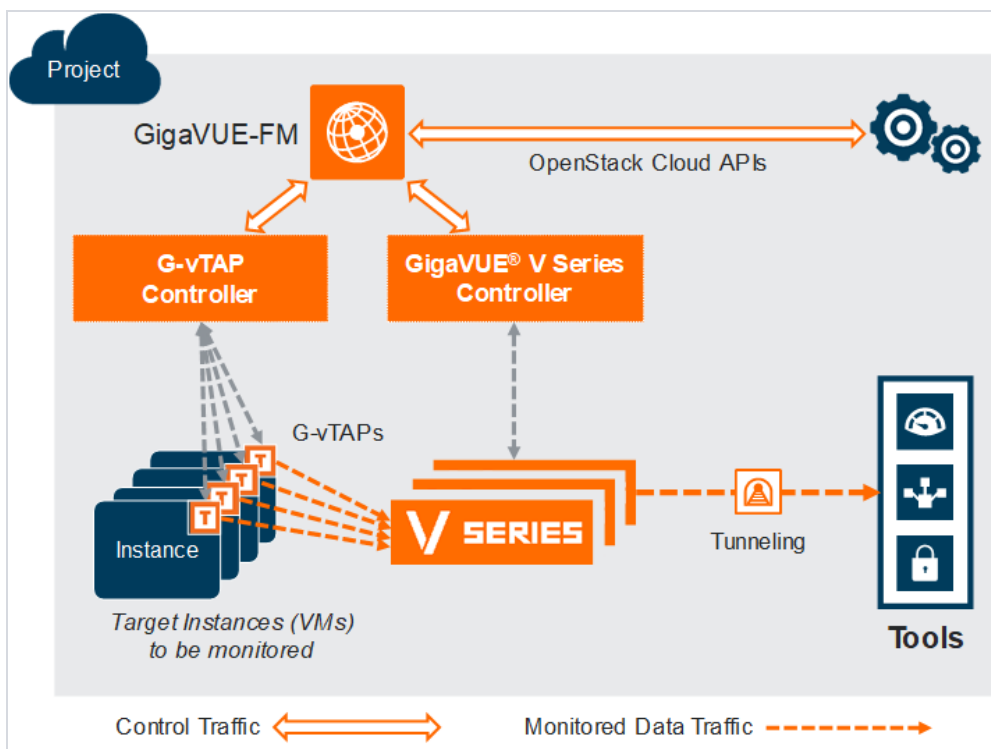


Figure 1: GigaVUE Cloud Components for OpenStack using G-vTAP



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.

## OpenVSwitch (OVS) Mirroring

When deploying OpenVSwitch (OVS) Mirroring, a G-vTAP agent is installed on the hypervisor where the VMs you wish to monitor are located. 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 OVS Mirroring or OVS Mirroring + DPDK, the mirroring infrastructure is fully contained within the hypervisors. This G-vTAP agent must be on OpenVSwitch.



**Note:** OVS Mirroring also supports OpenVSwitch with DPDK as a preview. The configuration steps for OVS Mirroring and OVS Mirroring with DPDK are the same.

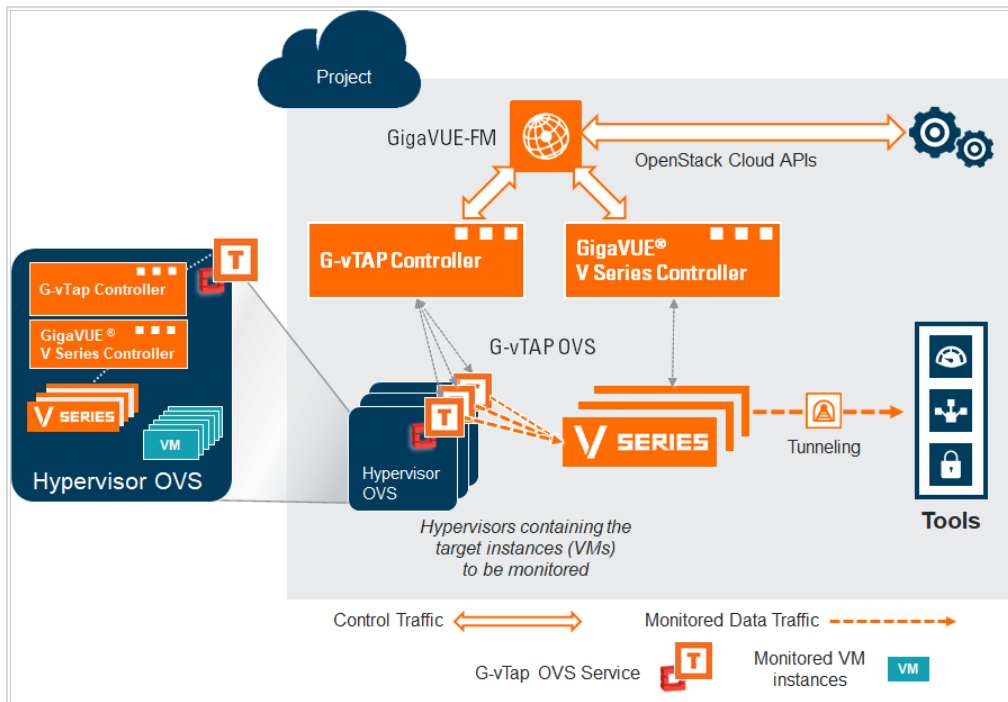


Figure 2: GigaVUE Cloud Components for OpenStack using OVS Mirroring or OVS Mirroring + DPDK

## OVS Mirroring Prerequisites

The following items are required to deploy a G-vTAP OVS agent:

- An existing OpenStack cloud environment should be available with admin login credentials
- A user with OVS access is required to enable OVS-Mirror. The user can be an admin or can be a user with a custom role that has the permissions and the ability to list projects. Refer to [OpenStack Role Privileges Required to Enable OVS Mirroring](#) for the elevated privileges required.
- A working GigaVUE-FM with latest build.

- OpenStack Cloud Environment Requirements:
  - OpenStack Version: Rocky and above.
  - Ubuntu Version 16.04 and above or RedHat version 7.6 and above.
  - ML2 mechanism driver: OpenVSwitch

**Tip:** If the OpenStack CLI does not return a reachable IP for the hypervisors that are being monitored, you must manually enter a reachable IP for each hypervisor in OpenStack using project properties. For each hypervisor you will need to add a key value pair property in the following format:

- key: value
- key: must be in the form gigamon-hv-<hypervisorID>
- value: reachable IP for hypervisor

For example: gigamon-hv-1 : 10.120.10.2

## OpenStack Role Privileges Required to Enable OVS Mirroring

OpenStack CLI command	Supported API/Action	Description
openstack hypervisor list	GET /os-hypervisors	Should list all hypervisors in the domain
openstack server list --all --host <hostname>	GET /servers	Should list all the servers on a specified host
openstack server list -all	GET /servers	Should list servers of all projects in the domain
openstack project list	GET /v3/projects	Should list all projects in the domain
openstack project list --user <user with custom role>	GET /v3/projects	Should list all projects that a specified user (user specified in FM config) is associated with
openstack user list	GET /v3/users	Should list all users in the domain
openstack subnet list	GET /subnets	Should list subnets for all projects in the domain
openstack network list	GET /network	Should list networks for all projects in the domain
openstack floating ip list	GET /floatingips	Should list floating ips for all projects in the domain
openstack floating ip set --port <portId> <floating ip>	PUT /floatingips/{floatingip_id}	Used to attach floating ip to fabric nodes
openstack security group list	GET /security-groups	Should list security groups for all projects in the domain
openstack security group show <security group id>	GET /security-groups/{security_group_id}	Should list details of specified security group
openstack port list	GET /ports	Should list ports for all projects in the domain

## OpenVSwitch (OVS) Mirroring + DPDK



**Note:** OVS Mirroring also supports OpenVSwitch with DPDK as a preview. The configuration steps for OVS Mirroring and OVS Mirroring with DPDK are the same. See instructions for OVS Mirroring throughout this guide when testing OVS Mirroring + DPDK.

# 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:

## MiniTOC Proxy

### Short list. Compare to MiniTOC

- [Before You Begin on page 13](#)
- [Uploading the Images on page 22](#)
- [Launching the GigaVUE-FM Instance on page 23](#)
- [Installing the G-vTAP Agents on page 34](#)
- [Configuring the GigaVUE Cloud in OpenStack on page 38](#)

## Before You Begin

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

- [Supported Hypervisor on page 13](#)
- [Network Requirements on page 15](#)
- [Virtual Network Interface Cards \(vNICs\) on page 16](#)
- [Security Group on page 16](#)
- [Key Pairs on page 21](#)

## Supported Hypervisor

Table 2: Hypervisor for OpenStack lists the hypervisor with the supported versions for G-vTAP.

Table 2: Hypervisor for OpenStack

Hypervisor	Version
KVM	<b>G-vTAP</b> cPike, Queens, Ocata, Newton, Mitaka, and Liberty <b>OVS Mirroring</b> —Rocky and above

## 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 3: Minimum Compute Requirement](#).

Table 3: 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 OVS Agent	N/A	N/A	N/A	Available as rpm or debian package.
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)

Compute Instances	vCPU	Memory	Disk Space	Description
				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 4: Types of Networks lists the recommended requirements to setup the network topology.

Table 4: Types of Networks

Network	Purpose
<b>Management</b>	Identify the Network Interface Card (NIC) that GigaVUE-FM uses to communicate with the GigaVUE V Series nodes and controllers.
<b>Data</b>	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 with GvTAP Agents 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.

**NOTE:** vNICs are only applicable if the GvTap Agent is installed on the instances being monitored. It is not applicable for OVS Mirroring or OVS Mirroring +DPDK.

## 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 the "[Security Group Rules](#)" on the facing page table.

The "[Security Group Rules](#)" on the facing page table lists the rules and port numbers for each component.



Table 5: Security Group Rules

Direction	Ether Type	Protocol	Port	CIDR	Purpose
<b>GigaVUE-FM</b>					
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
<b>G-vTAP Controller</b>					
Inbound	IPv4	TCP	9900	GigaVUE-FM IP address	Allows GigaVUE-FM to communicate with G-vTAP Controllers
<b>G-vTAP Agent</b>					
Inbound	IPv4	TCP	9901	G-vTAP Controller IP address	Allows G-vTAP Controllers to communicate with G-vTAP agents
Inbound	IPv4	TCP	9902	GigaVUE-FM IP address	Allows GigaVUE-FM to communicate with GigaVUE V Series Controllers

Direction	Ether Type	Protocol	Port	CIDR	Purpose
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
---------	----------------------	----------	----	----------------	--

#### VXLAN Traffic

Inbound	Custom UDPRule	UDP	4789	Any IP address	Allows mirrored traffic from G-vTAP agents to be sent to GigaVUE V Series nodes using the VXLAN tunnel  Allows monitored traffic from GigaVUE V Series nodes to be sent to the monitoring tools using the VXLAN tunnel
---------	----------------	-----	------	----------------	--

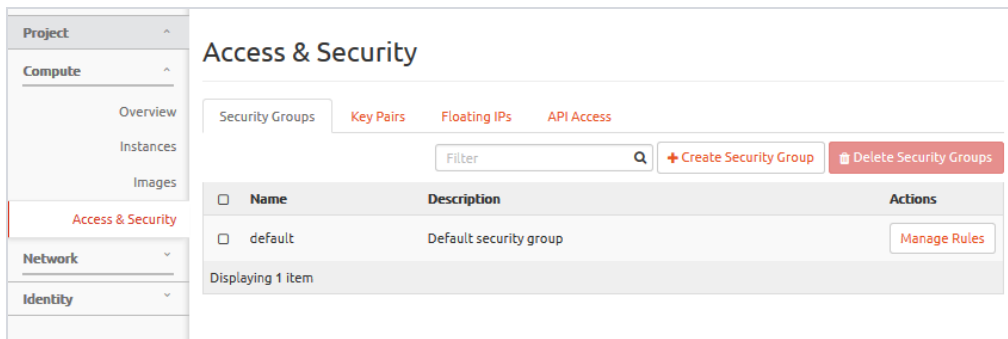
**NOTE:** The "Security Group Rules" on the previous page table lists only the ingress rules. Make sure the egress ports are open for communication.

Along with the ports listed in the "Security Group Rules" on the previous page table, 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.



3. Click **Create Security Group** .
4. Enter a name and description in the respective fields and click **Create Security Group**.

### Create Security Group ✕

**Name \***

**Description:**

Security groups are sets of IP filter rules that are applied to the network settings for the VM. After the security group is created, you can add rules to the security group.

**Description**

Security Group for G-VTAP agents

Cancel
Create Security Group

The security group is created and added to the Access & Security page.

### Access & Security

Security Groups
Key Pairs
Floating IPs
API Access

+ Create Security Group
Delete Security Groups

<input type="checkbox"/> Name	Description	Actions
<input type="checkbox"/> default	Default security group	<span style="border: 1px solid #ccc; padding: 2px 5px;">Manage Rules</span>
<input type="checkbox"/> sg_gvtap-agent		<span style="border: 1px solid #ccc; padding: 2px 5px;">Manage Rules</span> ▼

Displaying 2 items

5. For the new security group added, click **Manage Rules**. The Manage Security Group Rules page is displayed.
6. Click **Add Rule**. The Add Rule page is displayed.
7. Enter the appropriate values in the respective fields.
8. Click **Add**. The Manage Rules page is displayed with the newly added rule.

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

### Manage Security Group Rules: sg\_gvtap-agent (5e2c05fb-2cd3-42f5-9333-18f9e8beb7e4)

[+ 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	-	<a href="#">Delete Rule</a>
<input type="checkbox"/>	Egress	IPv4	Any	Any	0.0.0.0/0	-	<a href="#">Delete Rule</a>
<input type="checkbox"/>	Ingress	IPv4	TCP	9901	10.115.46.131/32	-	<a href="#">Delete Rule</a>

Displaying 3 items

9. Repeat steps 2 to 8 to create security groups for all the components.

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

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

Displaying 7 items

## 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

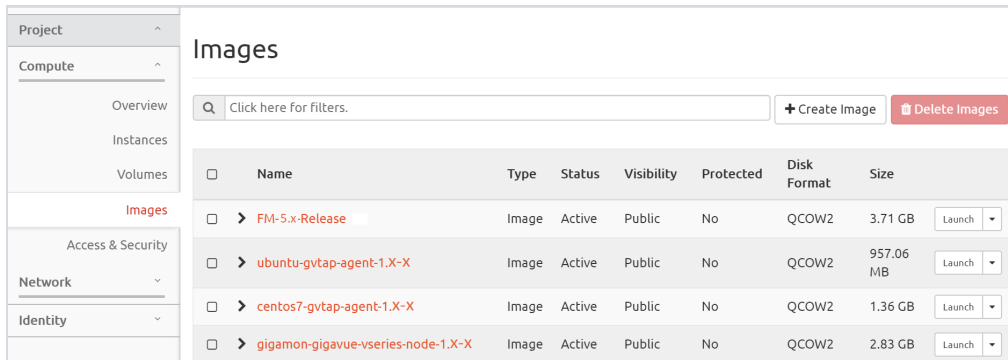
```
openstack image create --disk-format qcow2 -public --container-format  
bare --file gigamon-gigavue-vseries-cntlr-1.x-x gigamon-gigavue-vser-  
ies-cntlr-1.x-x.qcow2
```

**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
- gigamon-gigavue-gvtap-ovs-cntlr-1.x-x

Once the images are uploaded, they are displayed under **Compute** > **Images**. Refer to [Figure 3: Instances Uploaded in Images Page](#) on page 23.



The screenshot shows the 'Images' page in the OpenStack Horizon GUI. The left sidebar contains navigation options: Project, Compute (selected), Overview, Instances, Volumes, Images (highlighted), Access & Security, Network, and Identity. The main content area is titled 'Images' and includes a search bar with the placeholder 'Click here for filters.', a '+ Create Image' button, and a 'Delete Images' button. Below this is a table listing four images:

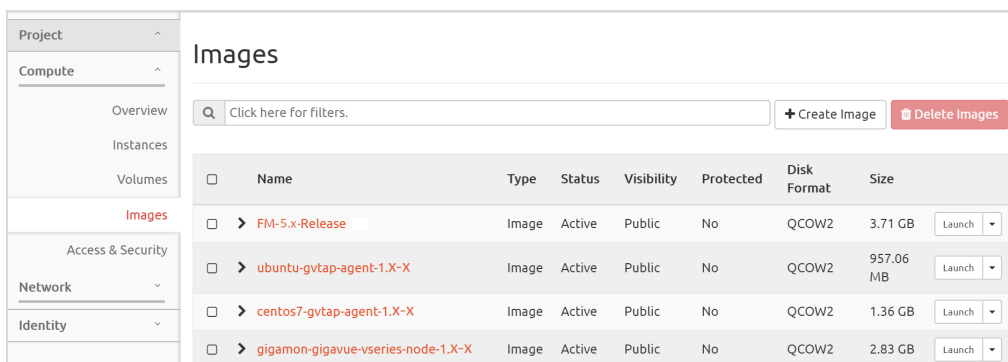
<input type="checkbox"/>	Name	Type	Status	Visibility	Protected	Disk Format	Size	
<input type="checkbox"/>	> FM-5.x-Release	Image	Active	Public	No	QCOW2	3.71 GB	Launch ▾
<input type="checkbox"/>	> ubuntu-gvtap-agent-1.X-X	Image	Active	Public	No	QCOW2	957.06 MB	Launch ▾
<input type="checkbox"/>	> centos7-gvtap-agent-1.X-X	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 3: 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 4: List of Available Images](#) on page 24.



This screenshot is identical to the one above, showing the 'Images' page in the OpenStack Horizon GUI with the same list of four images: FM-5.x-Release, ubuntu-gvtap-agent-1.X-X, centos7-gvtap-agent-1.X-X, and gigamon-gigavue-vseries-node-1.X-X.

Figure 4: List of Available Images

3. Select the GigaVUE-FM image and click **Launch**. The Launch Instance dialog box is displayed. Refer to [Figure 5: Launch Instance Page on page 24](#).

Launch Instance

Details

Source

Flavor \*

Networks \*

Network Ports

Security Groups

Key Pair

Configuration

Server Groups

Scheduler Hints

Metadata

Please provide the initial hostname for the instance, the availability zone where it will be deployed, and the instance count. Increase the Count to create multiple instances with the same settings.

Instance Name \*

GigaVUE-FM-Gigamon

Availability Zone

nova

Count \*

1

Total Instances (10 Max)

70%

6 Current Usage

1 Added

3 Remaining

Cancel

< Back

Next >

Launch Instance

Figure 5: Launch Instance Page

1In the **Details** tab, enter the information as shown in [Figure 5: Launch Instance Page on page 24](#).

- **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.

4. Click **Next**.



- 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 6: Source tab on page 25](#).

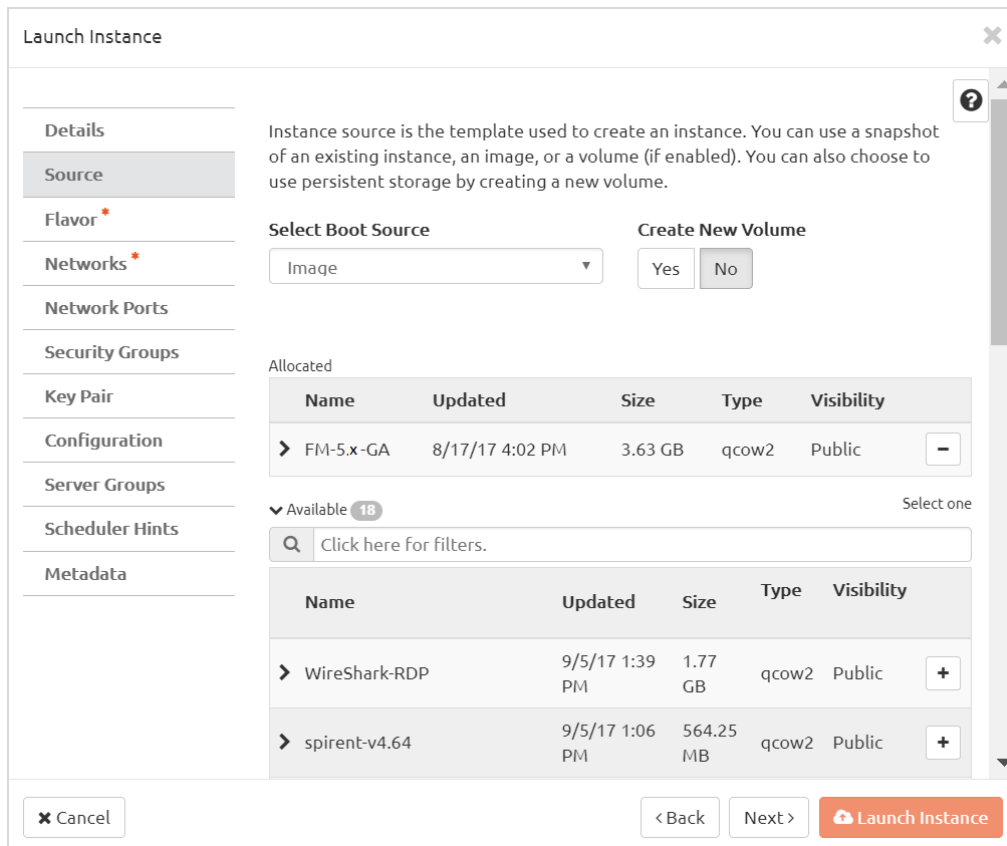


Figure 6: Source tab

- 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 7: Flavor tab on page 26](#).

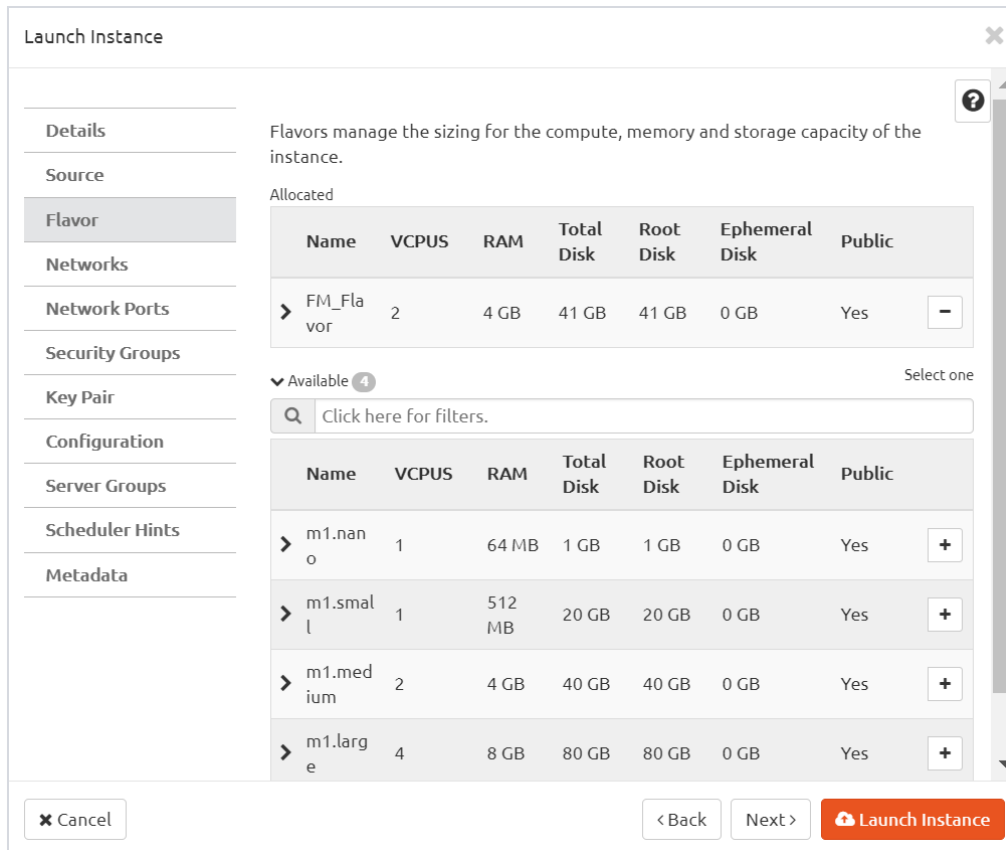


Figure 7: Flavor tab

- 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 15](#). The selected network is displayed under Allocated. Click **Next**. Refer to [Figure 8: Networks tab on page 27](#).

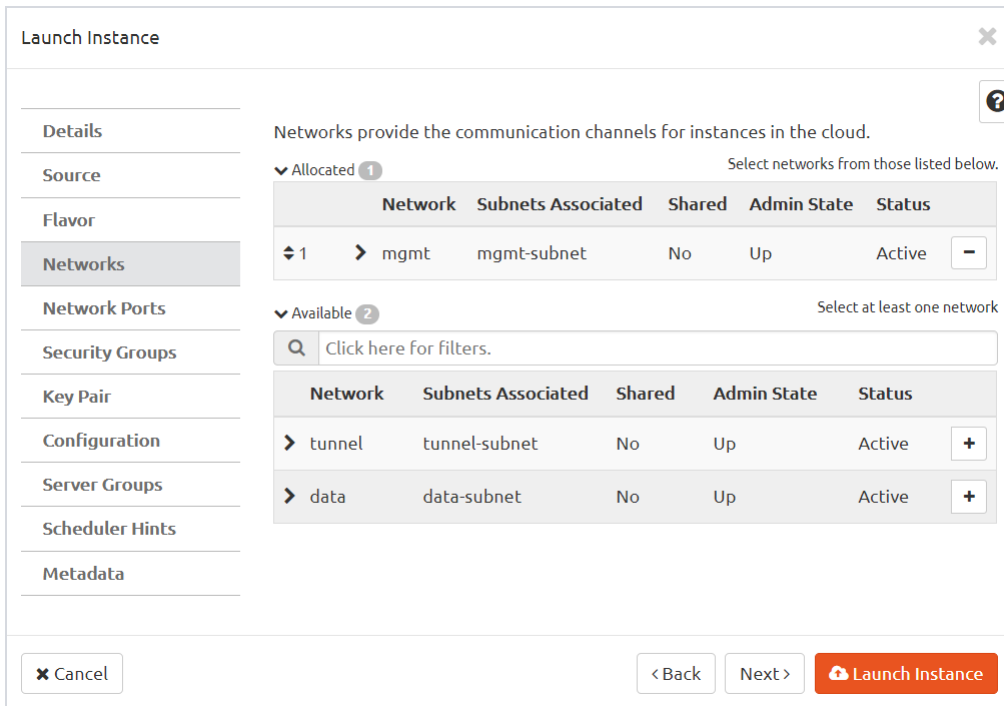


Figure 8: Networks tab

8. In the **Network Ports** tab, click **Next** again.
9. 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 16](#). The selected security group is displayed under Allocated. Click **Next**. Refer to [Figure 9: Security Groups tab on page 28](#).

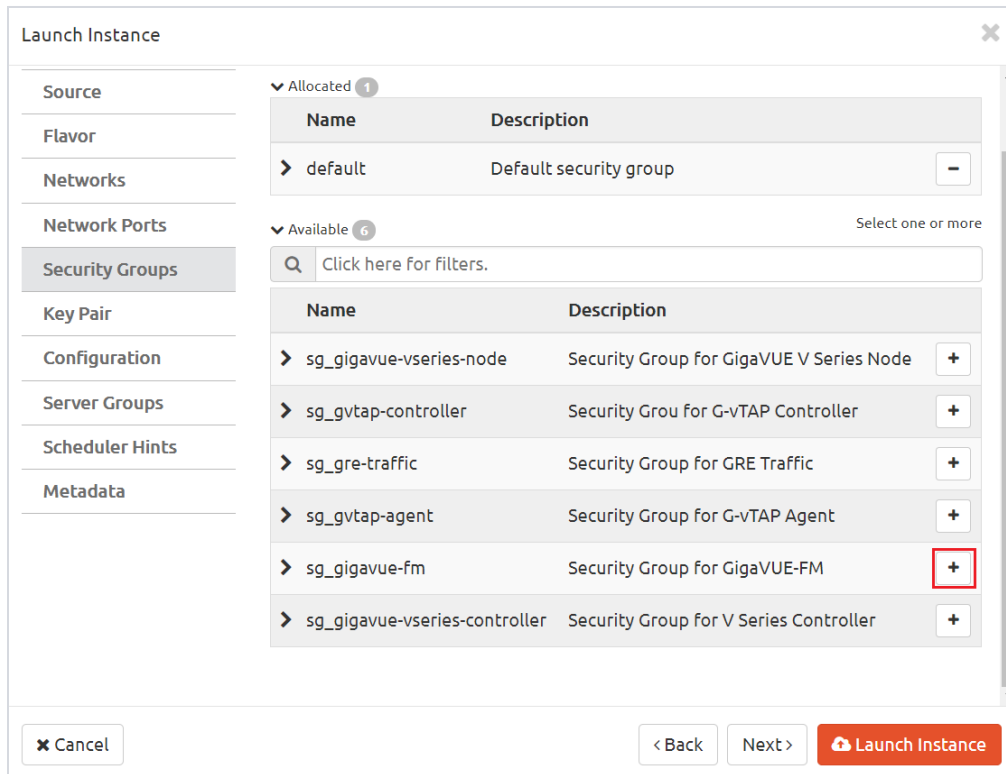


Figure 9: Security Groups tab

- 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 21](#). The selected key pair is displayed under Allocated. Click **Next**. Refer to [Figure 10: Key Pair tab on page 29](#).

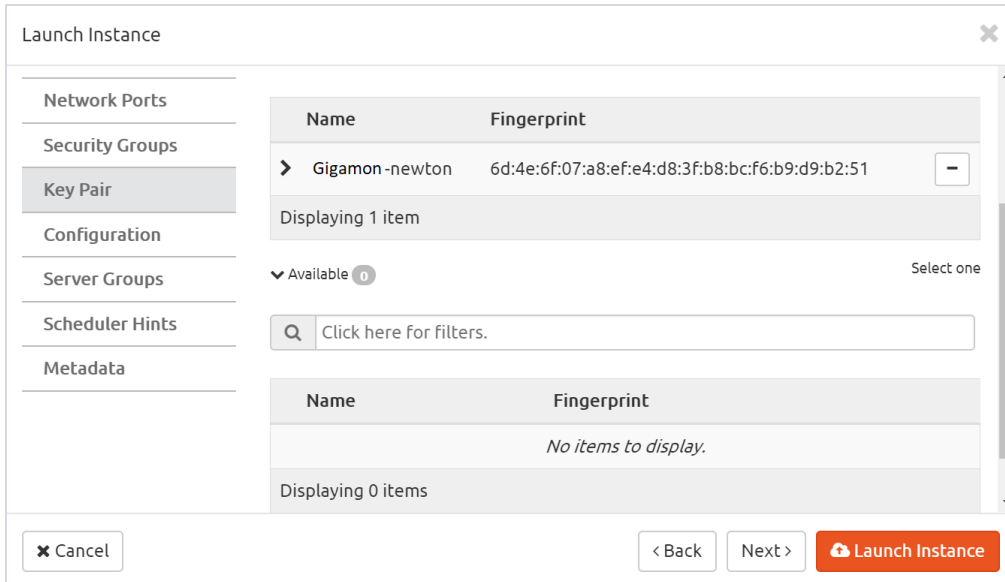


Figure 10: Key Pair tab

- 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 11: Instance Launched](#) on page 29.

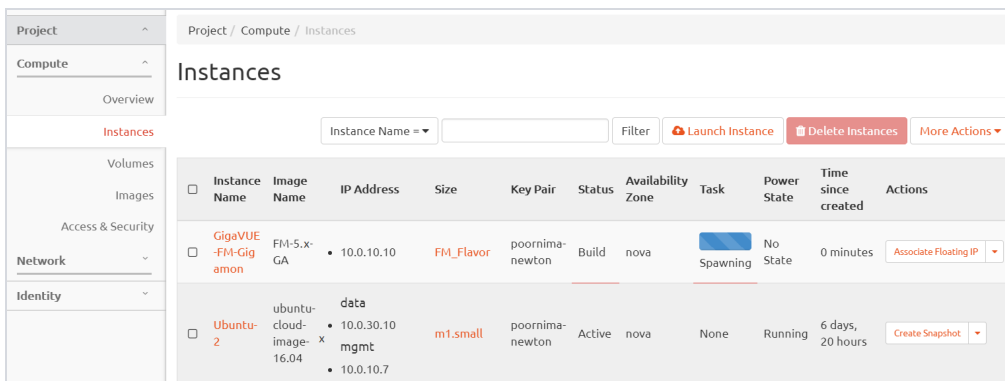


Figure 11: Instance Launched

- During the initial boot-up sequence, click **Allocate Floating IP**. The Manage Floating IP Associations dialog box is displayed. Refer to [Figure 12: Assigning Floating Address on page 30](#).

*Figure 12: Assigning Floating Address*

- In the Manage Floating IP Associations dialog box, enter the following information as shown in [Figure 12: Assigning Floating Address on page 30](#):
  - 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 13: Floating IP of GigaVUE-FM Instance on page 30](#).

<input type="checkbox"/>	GigaVUE-F M-Gigamo n	FM-5.x-GA	<ul style="list-style-type: none"> <li>10.0.10.10</li> <li>10.210.219.144</li> </ul>	Floating IPs:	FM_Flavor	poornima-newton	Active	nova	None	Running
--------------------------	----------------------------	-----------	--	---------------	-----------	-----------------	--------	------	------	---------

*Figure 13: 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.

1. 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 14: GigaVUE-FM Instance Console on page 31](#).

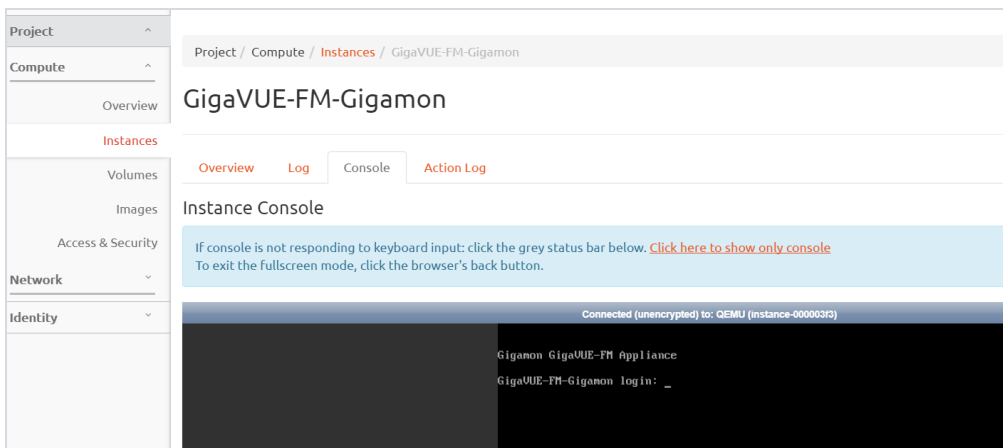


Figure 14: GigaVUE-FM Instance Console

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 15: GigaVUE-FM Jump Start Configuration on page 32](#).

```
Gigamon GigaVUE-FM Appliance
gigavue-fm-199d16 login: admin
Password:

Gigamon GigaVUE-FM Appliance

GigaVUE-FM configuration wizard

Step 1: Hostname? [gigavue-fm-199d16] GigaVUE-FM-Gigamon
Step 2: Use DHCP on eth0 interface? [no] yes
Step 3: Admin password (Enter to leave unchanged)?
Step 4: Additional Domain Name Server IP addresses? 10.210.19.120
Step 5: Additional DNS Domains? gigaga.com
Step 6: Enable NTP? [yes] yes
Step 7: NTP Server IP address? 10.115.1.55
Step 8: NTP Server version? [4] 4_
```

Figure 15: 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 `eth1`; 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 the G-vTAP OVS Agents for OVS Mirroring

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

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

- [Installing the G-vTAP OVS Agent from an Ubuntu/Debian Package](#)
- [Installing the G-vTAP OVS Agents for OVS Mirroring](#)

### Installing the G-vTAP OVS Agent from an Ubuntu/Debian Package

To install from a Debian package:

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

```
$ ls gvtap-ovs-agent_1.x-x_amd64.deb
ubuntu@ip-10-0-0-246:~$ sudo dpkg -i gvtap-ovs-agent_1.x-x_
amd64.deb
```

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

3. Once the G-vTAP OVS package is installed, start the agent:

```
$ sudo service gvtap-agent start
```

The G-vTAP OVS agent status will be displayed as running.

4. Check the status using the following command:

```
$ sudo service gvtap-agent status
G-vTAP Agent is running
```

### Installing the G-vTAP OVS Agent 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 OVS Agent RPM (.rpm) package from the [Gigamon Customer Portal](#).
2. Copy this package to OpenStack compute nodes. Install the package with root privileges, for example:

```
$ ls
gvtap-ovs-agent_1.x-x_x86_64.rpm
$ sudo rpm -ivh gvtap-ovs-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. After the installation completes, start the G-vTAP OVS agent service and verify its status.

```
$ systemctl start gvtap-agent.service
$ sudo service gvtap-agent status
G-vTAP Agent is running
```

## Configuring the GigaVUE 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 6: Pre-configuration Checklist on page 39](#) 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 GigaVUE Cloud Suite for OpenStack:

*Table 6: Pre-configuration Checklist*

Required Information	
<input type="checkbox"/>	Authentication URL
<input type="checkbox"/>	Project Name
<input type="checkbox"/>	Peering
	<b>NOTE:</b> 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 [Figure 16: GigaVUE-FM Login Page on page 40](#).



Figure 16: 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.

1. 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 31](#), enter the changed password.
2. Click **Login**. The GigaVUE-FM Dashboard page is displayed. Refer to [Figure 17: GigaVUE-FM on page 41](#).



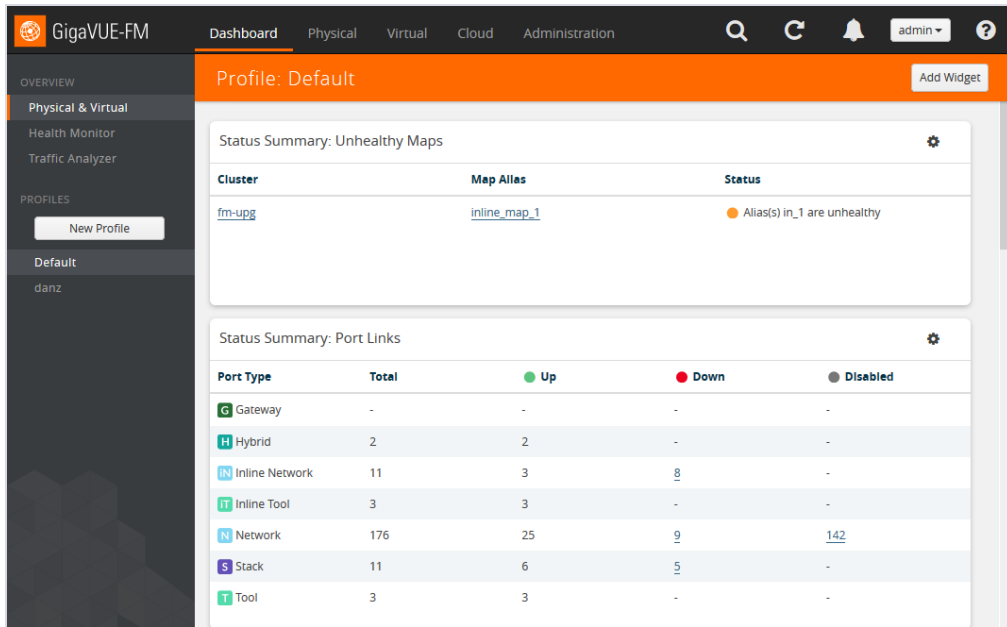


Figure 17: GigaVUE-FM

## Connecting to OpenStack

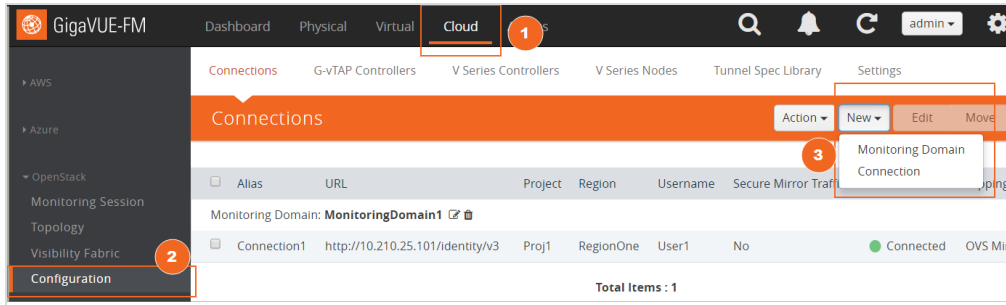
You can log in to GigaVUE-FM and use the CLI command: `ip host <controller-hostname> <ip-address of the controller>`. (For example: `ip host os-controller1 192.168.2.3`.) Then, add the connection to the OpenStack tenant.

**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.

To create a new connection:

1. Click **Cloud** from the GigaVUE-FM top navigation.

2. Go to **OpenStack > Configuration > Connection** to view the OpenStack connections.



3. 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.
4. Enter or select the appropriate information to set up the connection. Refer to [Table 7: OpenStack Connection](#) for field-level details.

Connection
Save Cancel

**Alias**

**Monitoring Domain**

**URL**

**Domain Name**

**Project Name**

**Region**

**Username**

**Password**

**Tap Method**

**Projects to Monitor**

OVS Mirroring only

**Secure Mirror Traffic**

Table 7: OpenStack Connection

Field	Description
<b>Alias</b>	An alias used to identify the connection to OpenStack.
<b>Monitoring Domain</b>	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. <b>NOTE:</b> Monitoring domain consists of set of connections.

Field	Description
<b>URL</b>	<p>The authentication URL is the Keystone URL of the OpenStack cloud. This IP address must be DNS resolvable.</p> <p>To get the authentication URL from the OpenStack dashboard:</p> <ol style="list-style-type: none"> <li>Login to OpenStack Horizon.</li> <li>Go to <b>Compute &gt; Access &amp; Security</b>.</li> <li>Click the <b>API Access</b> tab and copy the Identity URL.</li> </ol>  <p>Paste the Identity URL into the URL field.</p>
<b>Domain Name</b>	The DNS domain name of the project.
<b>Project Name</b>	<p>For GvTAP, this is the name of the project to monitor.</p> <p>For GvTAP-OVS (OVS Mirroring), this is the project name that will be used for authentication.</p> <p>This is a required field.</p>
<b>Region</b>	<p>The region where the Project resides. You can find your region by running one of these commands, depending on your OpenStack version.</p> <pre>keystone endpoint-list or openstack endpoint list</pre>
<b>Username</b>	<p>The user name used to connect to the OpenStack cloud.</p> <p><b>NOTE:</b> The user must belong to a role that meets the OpenStack minimum requirements for OVS Mirroring. Refer to <a href="#">OVS Mirroring Prerequisites</a>.</p>

Field	Description
<b>Password</b>	The password for the OpenStack cloud.
<b>Tap Method</b>	<p>Select the type of agent used to capture traffic for monitoring:</p> <ul style="list-style-type: none"><li>• TaaS</li><li>• G-vTAP</li><li>• OVS Mirroring</li><li>• OVS Mirroring + DPDK</li><li>• None</li></ul> <p><b>NOTE:</b> None is used if you are not using the connection for tapping and are only launching the V Series nodes for processing traffic from other connection, such as Kubernetes.</p>

Field	Description
<b>Projects to Monitor</b>	<p>This field only appears, and is required, for OVS Mirroring or OVS Mirroring + DPDK.</p> <ul style="list-style-type: none"> <li>Click the <b>Get Project List</b> button to view the list of projects.</li> </ul> <p><b>NOTE:</b> The <b>Get Project List</b> button will only work if all the OpenStack credentials have been provided. Refer to <a href="#">OVS Mirroring Prerequisites</a>.</p> <ul style="list-style-type: none"> <li>Select the name of the project or projects you want to monitor from the list. (There is a limit of 128 projects.)</li> <li>You can click <b>Select None</b> to clear existing selections or <b>Select All</b> to add all available projects to the connection configuration.</li> </ul> <div data-bbox="516 688 1317 1108" style="border: 1px solid #ccc; padding: 10px;"> <p><b>Projects to Monitor</b> <span>▼ Select All</span> <span>✕ Select None</span></p> <div style="display: flex; align-items: center;"> <div style="border: 1px solid #ccc; padding: 2px; margin-right: 5px;">Proj1 ✕</div> <div style="border: 1px solid #ccc; padding: 2px; margin-right: 5px;">Proj2 ✕</div> <div style="border: 1px solid #ccc; padding: 2px; margin-left: 5px;">▼</div> <div style="margin-left: 20px; border: 1px solid #ccc; padding: 5px;">Get Project List</div> </div> <ul style="list-style-type: none"> <li style="background-color: #0070c0; color: white; padding: 2px;">Proj3</li> <li style="padding: 2px;">automation</li> <li style="padding: 2px;">Proj4</li> <li style="padding: 2px;">Proj5</li> <li style="padding: 2px;">TestProject1</li> <li style="padding: 2px;">admin</li> </ul> </div> <p><b>NOTE:</b> OVS Mirroring or OVS Mirroring + DPDK supports multiple projects, without limit. As stated in the prerequisites for OVS Mirroring, you need permission to list all the projects in your OpenStack cloud environment.</p>
<b>Secure Mirror Traffic</b>	<p>Check box to establish secure tunnel between G-vTAP agents and GigaVUE V Series nodes (especially in a shared controller and GigaVUE V Series node configuration)</p> <p><b>NOTE:</b> Must be deselected for OVS Mirroring or OVS Mirroring + DPDK.</p>

5. Click **Save**.

## Viewing Connections

If GigaVUE-FM connects to OpenStack successfully, the status is displayed as "Connected" in the **Status** column on the Connections page. GigaVUE-FM discovers the

inventory of the cloud in the background.




Figure 18: OpenStack Connection in GigaVUE-FM

The Connections page has the following controls:

Control	Description
<b>Action</b>	Allows to refresh inventory.
<b>New</b>	Opens the page for specifying the connection details for a new connection.
<b>Edit</b>	Allows to make changes to a connection.
<b>Delete</b>	Deletes the connection. <b>NOTE:</b> Deleting a connection destroys all GigaVUE V Series Nodes, G-vTAP Controllers, and the virtual maps on the project.

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 19: Audit Logs on page 48](#).

Time	User	Operation Type	Source	Status	Description
2017-09-05 16:59:14	admin	create openstackCo nnection	VM	SUCCESS	
2017-09-05 16:58:31	admin	delete openstackCo nnection 695bc7e0- a46c-4567-9b01	VM	SUCCESS	
2017-09-05 16:52:36	admin	create openstackCo nnection	VM	SUCCESS	
2017-09-05 16:51:22	admin	create openstackCo nnection	VM	FAILURE	OpenstackConnectionException: Region: nova does not exist
2017-09-05 16:49:19	admin	create openstackCo nnection	VM	FAILURE	OpenstackConnectionException: Region: Region1 does not exist
2017-09-05 16:49:05	admin	create openstackCo nnection	VM	FAILURE	OpenstackConnectionException: Unauth orized. Verify username, password, doma

Figure 19: Audit Logs

## 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 56](#).

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 have 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** from the GigaVUE-FM top navigation.



2. Go to **OpenStack > Configuration > G-vTAP Controllers**. The G-vTAP Configuration page is displayed.



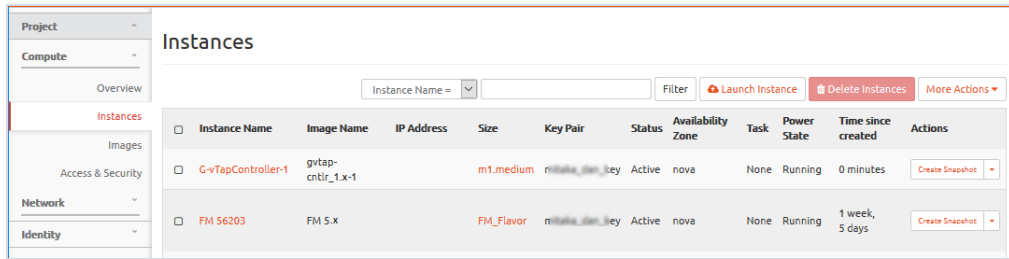
Options:

- a. Click **New** to create a new G-vTAP Controller configuration.
  - b. Click the check box of a configuration and click **Edit** to edit an existing configuration.
3. Complete the fields in the G-vTAP Configuration page to create or edit the controller's configuration settings. Refer to [The OpenStack G-vTAP Configuration Page on page 51](#) for details.
  4. Click **Save** to save the configuration.
  5. To verify your instance is running, refer to "[Verify the G-vTAP Controller Instance](#)" below.

## Verify the G-vTAP Controller Instance

After creating or editing the G-vTAP Controller, verify the instance is running:

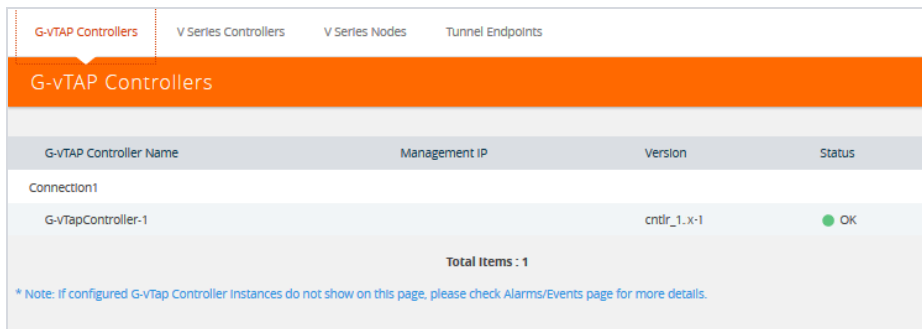
1. Log in to OpenStack Horizon and view the compute instances for the project to verify the launch of the G-vTAP Controller.



**NOTE:** The above screen is an example only. It is not a GigaVUE-FM screen and is subject to change without notice. Refer to OpenStack documentation for how to use the OpenStack interface.

- 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.




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

## Delete a G-vTAP Controller

To delete a G-vTAP Controllers:

- Click **Cloud** from the GigaVUE-FM top navigation.

2. Go to **OpenStack > Configuration > G-vTAP Controllers**. The G-vTAP Configuration page is displayed.



3. Click the check box to select an existing controller from the list.
4. Click **Delete**.

## The OpenStack G-vTAP Configuration Page

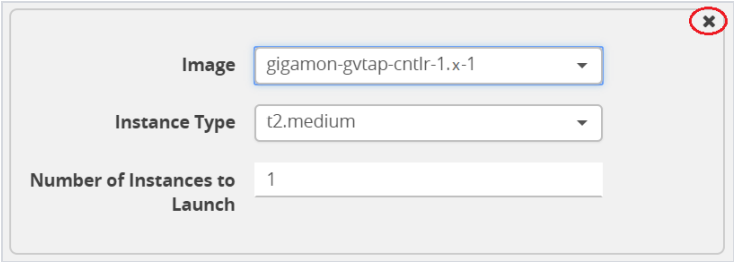
Use the G-vTAP Configuration page to the configuration settings of a G-vTAP Controller.

Figure 20: OpenStack G-vTAP Controller Configuration profile page

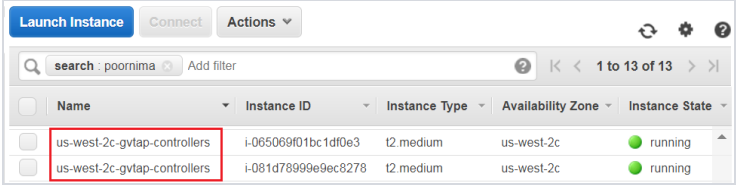
The following table provides a description of the fields in the G-vTAP Configuration page. The specific fields and options may vary based on your selection.

Table 8: Fields for G-vTAP Configuration

Fields	Description
<b>Connection</b>	<p>The name of the connection.</p> <p><b>NOTE:</b> 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.</p>
<b>SSH KeyPair</b>	<p>The SSH key pair for the G-vTAP Controller.</p> <p>For more information about SSH key pair, refer to <a href="#">Key Pairs on page 21</a>.</p>
<b>Project</b>	<p>Only available for OVS Mirroring or OVS Mirroring + DPDK. Select the OpenStack project where you want to deploy the fabric.</p> <p><b>NOTE:</b> When setting up the V Series Controller, this field is populated with the project selected in the G-vTAP Controller setup.</p> <p><b>NOTE:</b> If using with OVS Mirroring or OVS Mirroring + DPDK, a minimum of one project is required for OVS authentication.</p>
<b>Security Groups</b>	<p>The security group created for the G-vTAP Controller. For example, sg_gvtap-controller. For more information, refer to <a href="#">Security Group on page 16</a>.</p>

Fields	Description
<p><b>Controller Version(s)</b></p>	<p>The G-vTAP Controller version that 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><b>NOTE:</b> 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"> <li>Under <b>Controller Versions</b>, click <b>Add</b>.</li> <li>From the <b>Image</b> drop-down list, select a G-vTAP Controller image that matches with the version number of G-vTAP agents installed in the instances.</li> <li>From the <b>Flavor</b> down-down list, select a flavor for the G-vTAP Controller.</li> <li>In <b>Number of Instances to Launch</b>, specify the number of G-vTAP Controllers to launch. The minimum number you can specify is 1.</li> </ol> <p>To delete a G-vTAP Controller version:</p> <p>An older version of G-vTAP Controller can be deleted once all the G-vTAP agents are upgraded to the latest version.</p> <ol style="list-style-type: none"> <li>Click <b>x</b> (delete) next to the G-vTAP Controller image to delete that version.</li> </ol> <div data-bbox="651 1262 1380 1522" style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;">  </div> <ol style="list-style-type: none"> <li>When you delete a G-vTAP Controller image from the G-vTAP Configuration page, all the G-vTAP Controller instances of that version are also deleted.</li> </ol>

Fields	Description
<b>Management Network</b>	<p>This segment defines the management network that GigaVUE-FM uses to communicate with G-vTAP Controllers, GigaVUE V Series Controllers, and GigaVUE V Series Nodes.</p> <p><b>Network</b> - Select the management network ID.</p> <p><b>IP Address Type</b></p> <p>The type of IP address GigaVUE-FM needs to communicate with G-vTAP controllers:</p> <ul style="list-style-type: none"> <li>• <b>Private</b>—A private IP can be used when GigaVUE-FM, the G-vTAP Controller, or the GigaVUE V Series Controller reside inside the same project.</li> <li>• <b>Floating</b>—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.</li> </ul> <p><b>NOTE:</b> If GigaVUE-FM resides inside the same project, no floating IPs are necessary for the controllers.</p> <p><b>NOTE:</b> For V Series Nodes data network deployments, select <b>Floating</b> and then specify the IPs in the <b>Floating IPs</b> field.</p>
<b>Additional Network(s)</b>	<p>(Optional) If there are G-vTAP agents on networks that are not IP routable from the management network, additional networks or subnets must be specified so that the G-vTAP Controller can communicate with all the G-vTAP agents.</p> <p>Click <b>Add</b> to specify additional networks (subnets), if needed. Also, make sure that you specify a list of security groups for each additional network.</p>

Fields	Description															
<b>Tag(s)</b>	<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 (also known as a tag) that is easy to identify such as us-west-2-gvtap-controllers. There is a specific GvTAP Controller Version for OVS Mirroring and OVS Mirroring + DPDK.</p> <p>To add a tag:</p> <ol style="list-style-type: none"> <li>Click <b>Add</b>.</li> <li>In the <b>Key</b> field, enter the key. For example, enter Name.</li> <li>In the <b>Value</b> field, enter the key value. For example, us-west-2-gvtap-controllers.</li> </ol> <p>When the G-vTAP Controllers are launched in the VPC, they will appear with the custom tag:</p>  <table border="1" data-bbox="662 890 1393 1073"> <thead> <tr> <th>Name</th> <th>Instance ID</th> <th>Instance Type</th> <th>Availability Zone</th> <th>Instance State</th> </tr> </thead> <tbody> <tr> <td>us-west-2c-gvtap-controllers</td> <td>i-065069f01bc1df0e3</td> <td>t2.medium</td> <td>us-west-2c</td> <td>running</td> </tr> <tr> <td>us-west-2c-gvtap-controllers</td> <td>i-081d78999e9ec8278</td> <td>t2.medium</td> <td>us-west-2c</td> <td>running</td> </tr> </tbody> </table>	Name	Instance ID	Instance Type	Availability Zone	Instance State	us-west-2c-gvtap-controllers	i-065069f01bc1df0e3	t2.medium	us-west-2c	running	us-west-2c-gvtap-controllers	i-081d78999e9ec8278	t2.medium	us-west-2c	running
Name	Instance ID	Instance Type	Availability Zone	Instance State												
us-west-2c-gvtap-controllers	i-065069f01bc1df0e3	t2.medium	us-west-2c	running												
us-west-2c-gvtap-controllers	i-081d78999e9ec8278	t2.medium	us-west-2c	running												
<b>Agent Tunnel Type</b>	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.															
<b>G-vTAP Agent MTU (Maximum Transmission Unit)</b>	<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> <ul style="list-style-type: none"> <li>For GRE, the default value is 1450.</li> <li>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.</li> </ul>															

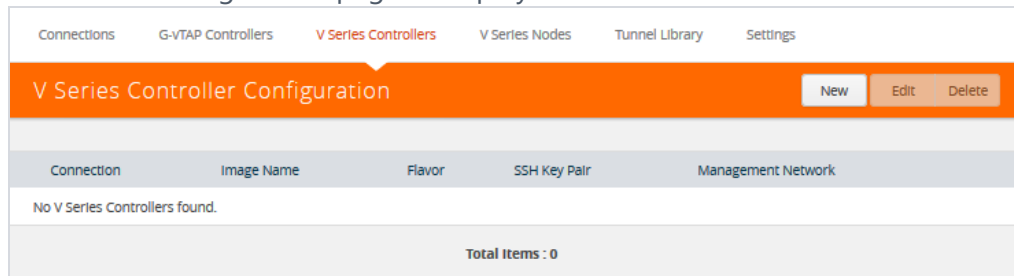
## Configuring the GigaVUE V Series Controllers

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** from the GigaVUE-FM top navigation.
2. Go to **OpenStack > Configuration > V Series Controllers**. The V Series Controller Configuration page is displayed.



Options:

- a. Click **New** to create a new V Series Controller configuration.
  - b. Click the check box of a configuration and click **Edit** to edit an existing configuration.
3. Complete the fields in the V Series Controller Configuration page to create or edit the controller's configuration settings. The fields in this page are the same as those on the G-vTAP Configuration page. Refer to "[The OpenStack G-vTAP Configuration Page](#)" on page 51 for the common fields.

Save    Cancel

### V Series Controller Configuration

**Connection** Connection1 ▾

**SSH Key Pair** key-pair1 ▾

**Project** Proj1 ▾

**Image** 48374f2b-743a-46c3-a754-1684a713e9e1 (gigamon-gigavue-vseries-cntrlr-1.7-1) ▾

**Flavor** m1.medium ▾

**Security Groups** default ▾

**Number of Instances**

**Management Network**

**IP Address Type**  Private  Floating

**Network** f7db4aa5-4c8a-475a-a908-2b50668a6a2b (mgmt-...) ▾

**Floating IPs** 10.218.17.148 × ▾

**Additional Network(s)** Add

**Tag(s)** Add

**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.

4. Click **Save** to save the configuration.
5. To verify your instance is running, refer to "Verify the V Series Controller Instance" on the facing page.

## Verify the V Series Controller Instance

After creating or editing the V Series Controller, verify the instance is running:

1. Log in to OpenStack Horizon and view the compute instances for the project to verify the launch of the V Series Controller.

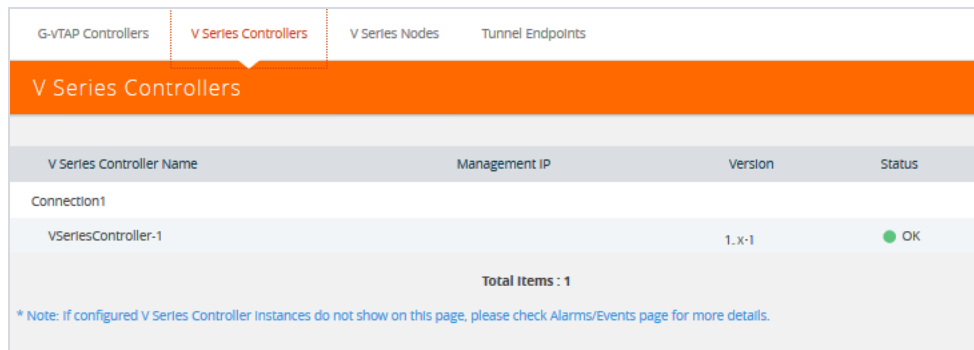


Instance Name	Image Name	IP Address	Size	Key Pair	Status	Availability Zone	Task	Power State	Time since created	Actions
<input type="checkbox"/> VSeriesController-1	gigamon-gigavue-vseries-cntlr-1.x-1		m1.medium	mitaka_dan_key	Active	nova	None	Running	1 minute	<a href="#">Create Snapshot</a>
<input type="checkbox"/> G-vTapController-1	gvtap-cntlr_1.x-1		m1.medium	mitaka_dan_key	Active	nova	None	Running	19 minutes	<a href="#">Create Snapshot</a>

**NOTE:** The above screen is an example only. It is not a GigaVUE-FM screen and is subject to change without notice. Refer to OpenStack documentation for how to use the OpenStack interface.

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

The V Series Controller should appear with the status as OK.



V Series Controller Name	Management IP	Version	Status
Connection1 VSeriesController-1		1.x-1	OK

Total Items : 1

\* Note: If configured V Series Controller instances do not show on this page, please check Alarms/Events page for more details.

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

## Configuring the GigaVUE V Series Node

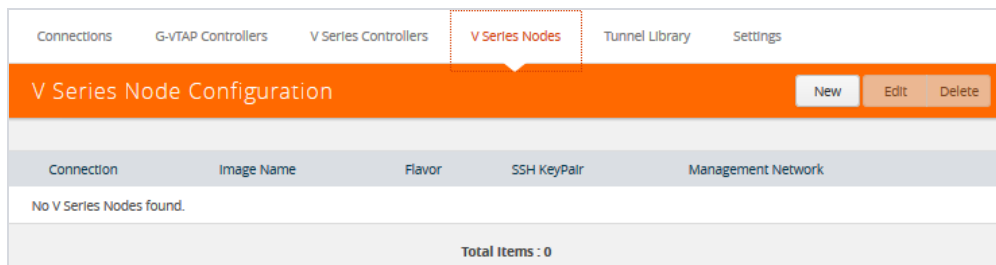
GigaVUE® V Series node is a visibility node that aggregates mirrored traffic from multiple G-vTAP agents. It applies filters, manipulates the packets using GigaSMART applications, and distributes the optimized traffic to cloud-based tools or backhaul to GigaVUE Cloud using the standard IP GRE tunnels.

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** from the GigaVUE-FM top navigation.
2. Go to **OpenStack > Configuration > V Series Nodes**.



Options:

- a. Click **New** to create a new V Series node configuration.
- b. Click the check box of a configuration and click **Edit** to edit an existing configuration.

3. Complete the fields in the V Series V SeriesNode Configuration page to create or edit the node's configuration settings. The fields in this page are the same as those on the G-vTAP Configuration page.

Refer to "The OpenStack G-vTAP Configuration Page" on page 51 for the common fields. For additional configuration settings, refer to [Table 9: Fields for GigaVUE V Series Node Launch Configuration](#).

V Series Node Configuration

Save Cancel

Connection: Connection1

SSH Key Pair: key-pair1

Project: Proj1

Image: 83d10001-7d71-47bb-876b-e1de0f582b5 (gigamon-gigavue-vseries-node-1.7-1)

Flavor: m1.medium

Security Groups: default

Management Network: f7db4a05-4c3e-475e-a908-2b50663afe2b (mgmt-network)

Data Network(s): Add

IP Address Type:  Private  Floating

Network 1: 5bb9b2c2-86ee-489b-912e-a4e92bb5146 (data-network)

Floating IPs: 10.210.17.135

Tag(s): Add

Min Instances: 1 ⓘ

Max Instances: 1

Tunnel MTU: 1450

Shared:  ⓘ

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

Table 9: Fields for GigaVUE V Series Node Launch Configuration

Parameter	Description
<b>Management Network</b>	<p>For the V Series Node, the Management Network is what is used by the V Series Controller to communicate with the V Series Nodes. Select the management network ID.</p>
<b>Data Network</b>	<p><b>Network 1</b></p> <p>Click <b>Add</b> to add additional networks . This is the network that the GigaVUE V Series node uses to communicate with the monitoring tools. Multiple networks are supported.</p> <p><b>IP Address Type</b></p> <p>The type of IP address GigaVUE-FM needs to communicate with G-vTAP or G-vTAP-OVS controllers:</p> <ul style="list-style-type: none"> <li>• <b>Private</b>—A private IP can be used when GigaVUE-FM, the G-vTAP Controller, or the GigaVUE V Series Controller reside inside the same project.</li> <li>• <b>Floating</b>—A floating IP address specified here will be where the monitored traffic is tunneled to. The monitored traffic must be able to reach the V Series Node.</li> </ul> <p><b>NOTE:</b> For OVS Mirroring or OVS Mirroring + DPDK deployments, must select <b>Floating</b> in the Data Network section and then specify the IPs in the <b>Floating IPs</b> field. You can have multiple Floating IPs.</p> <p><b>NOTE:</b> A provider network that is able to receive the monitored traffic may also be used here for OVS Mirroring and OVS Mirroring + DPDK. In this case, you would not need to provide a floating IP; but could select "private" and choose the provider network.</p>
<b>Min Instances to Launch</b>	<p>The minimum number of GigaVUE V Series nodes to be launched in OpenStack. The minimum number can be 0.</p> <ul style="list-style-type: none"> <li>• When you deploy an OVS Mirroring or OVS Mirroring + DPDK monitoring session, the V Series nodes will automatically be deployed based on the # of hypervisors being monitored.</li> <li>• When you deploy a G-vTAP monitoring session, the V Series nodes will automatically be deployed based on the # of VMs being monitored.</li> </ul>

Parameter	Description
	<b>NOTE:</b> GigaVUE-FM will delete the nodes if they are idle for over 15 minutes.
<b>Max Instances to Launch</b>	The maximum number of GigaVUE V Series nodes that can be launched in OpenStack.
<b>Tunnel MTU (Maximum Transmission Unit)</b>	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.
<b>Shared</b>	Only one V Series node configuration can be shared in a monitoring domain. <b>NOTE:</b> Not used for OpenStack configurations.

## Verify the V Series Node Instance

After creating or editing the V Series Node Controller, verify the instance is running:

1. Log in to OpenStack Horizon and view the compute instances for the project to verify the launch of the V Series Node .

The screenshot shows the OpenStack Horizon 'Instances' page. The table lists the following instances:

Instance Name	Image Name	IP Address	Size	Key Pair	Status	Availability Zone	Task	Power State	Time since created	Actions
		danz-mgmt 10.0.0.244								
<input type="checkbox"/> VSeriesNode-1	gigamon-gigavue-vseries-node-1.x-1	danz-mirror 10.0.1.77 danz-tunnel 10.0.2.20	m1.medium	mitaka_dan_key	Active	nova	None	Running	2 minutes	Create Snapshot
<input type="checkbox"/> VSeriesController-1	gigamon-gigavue-vseries-ctrlr-1.x-1	10.0.0.243	m1.medium	mitaka_dan_key	Active	nova	None	Running	21 minutes	Create Snapshot
<input type="checkbox"/> GvTapController-1	gvtap-ctrlr_1.x-1	10.0.0.242	m1.medium	mitaka_dan_key	Active	nova	None	Running	39 minutes	Create Snapshot

**NOTE:** The above screen is an example only. It is not a GigaVUE-FM screen and is subject to change without notice. Refer to OpenStack documentation for how to use the OpenStack interface.

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

The V Series Node Controller is displayed with the status as OK.

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.			

**NOTE:** The launch of the V Series Controller is also displayed in **Cloud > Audit Logs**.



# 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 65](#)
- [Creating Tunnel Endpoints on page 69](#)
- [Creating a Monitoring Session on page 71](#)
- [Configuring the OpenStack Settings on page 134](#)

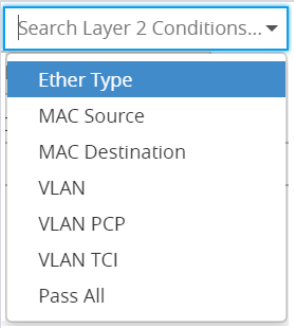
## Overview of Visibility Components

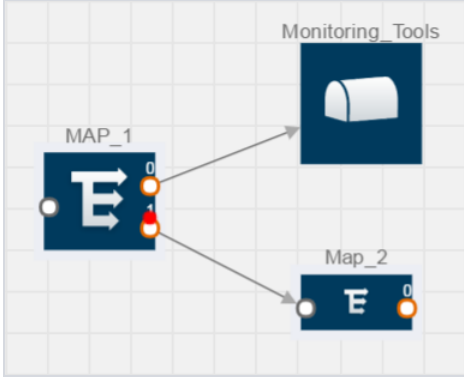
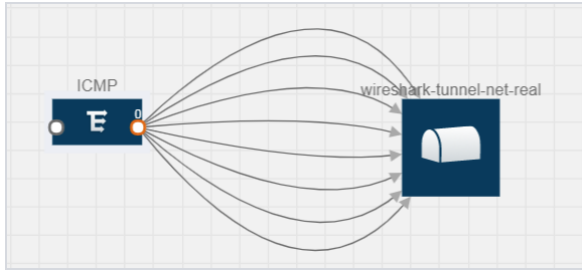
The GigaVUE V Series node aggregates the traffic from multiple G-vTAP agents or TaaS and filters them using maps. It applies intelligence and optimization to the aggregated traffic using GigaSMART applications such as Flow Mapping™, sampling, slicing, and masking, and distributes them to the tunnel endpoints.

[Table 10: Components of Traffic Visibility Sessions on page 65](#) lists the components of the monitoring session:

*Table 10: Components of Traffic Visibility Sessions*

Parameter	Description
<b>Map</b>	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.
<b>Rule</b>	A rule (R) contains specific filtering criteria that the packets must match.

Parameter	Description
	<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"> <li>• Layer 2—Ethertype IPv4</li> <li>• Layer 3—Protocol TCP</li> <li>• Layer 4—Port Destination 80</li> </ul> <p>By default, a rule always displays conditions based on the attributes of L2. Refer to <a href="#">Figure 21: Default Rule Conditions on page 66</a>.</p>  <p style="text-align: center;"><i>Figure 21: Default Rule Conditions</i></p> <p>A rule is also associated with priority and action set.</p>
<b>Priority</b>	<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>
<b>Action Set</b>	<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>

Parameter	Description
	<p>In the following example (refer to <a href="#">Figure 22: Action Set</a>), 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>  <p><i>Figure 22: Action Set</i></p> <p>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 <a href="#">Figure 23: Action Set with Multiple Links on page 67</a>.</p>  <p><i>Figure 23: Action Set with Multiple Links</i></p>
<b>Link</b>	<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 <a href="#">Figure 22: Action Set on page 67</a>, the link originating from action set 0 is</p>


Parameter	Description
	<p>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 <a href="#">Adding Header Transformations on page 124</a>.</p>
<b>Group</b>	A group is a collection of maps that are pre-defined and saved in the map library for reuse.
<b>Application</b>	An application performs operations such as sampling, slicing, and masking on the traffic.
<b>Inclusion Map</b>	An inclusion map determines the instances to be included for monitoring. This map is used only for target selection.
<b>Exclusion Map</b>	An exclusion map determines the instances to be excluded from monitoring. This map is used only for target selection.
<b>Target</b>	<p>A target determines the instances that are to be monitored.</p> <p>Targets are determined based on the following formula:</p> $\text{Target} = (\text{Maps} \cap \text{Inclusion map}) - \text{Exclusion map}$
<b>Automatic Target Selection (ATS)</b>	<p>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.</p> <p>For OVS Mirroring and OVS Mirroring + DPDK, the hypervisors will be selected as targets although this will not be displayed in the topology graph. It will be shown in the monitoring session deployment report.</p> <p>The instance targets will be selected based on two additional criteria:</p> <ul style="list-style-type: none"> <li>• they must be in one of the projects selected in the monitored projects list from the connection page, and</li> <li>• they must reside on a hypervisor that has the GvTAP Agent installed on it.</li> </ul>
<b>Tunnel</b>	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.



4. Click **New**. The Edit Tunnel page appears.

5. On the Edit Tunnel page, select or enter the appropriate information in the fields.  
Refer to [Table 11: Field Descriptions for Tunnel Endpoint on page 70](#).

Table 11: Field Descriptions for Tunnel Endpoint

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

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	22e2878c-553f-92a5-7966a0a00a7	VSeriesNode-1	tep0	Connection1_Monitoring

Total Items : 1

# 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 71](#)
- [Cloning a Monitoring Session on page 74](#)
- [Splitting a Monitoring Session on page 74](#)
- [Creating a Map on page 75](#)
- [Adding Applications to the Monitoring Session on page 85](#)
- [Deploying the Monitoring Session on page 120](#)
- [Viewing the Statistics on page 127](#)
- [Viewing the Topology on page 131](#)

## 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**.

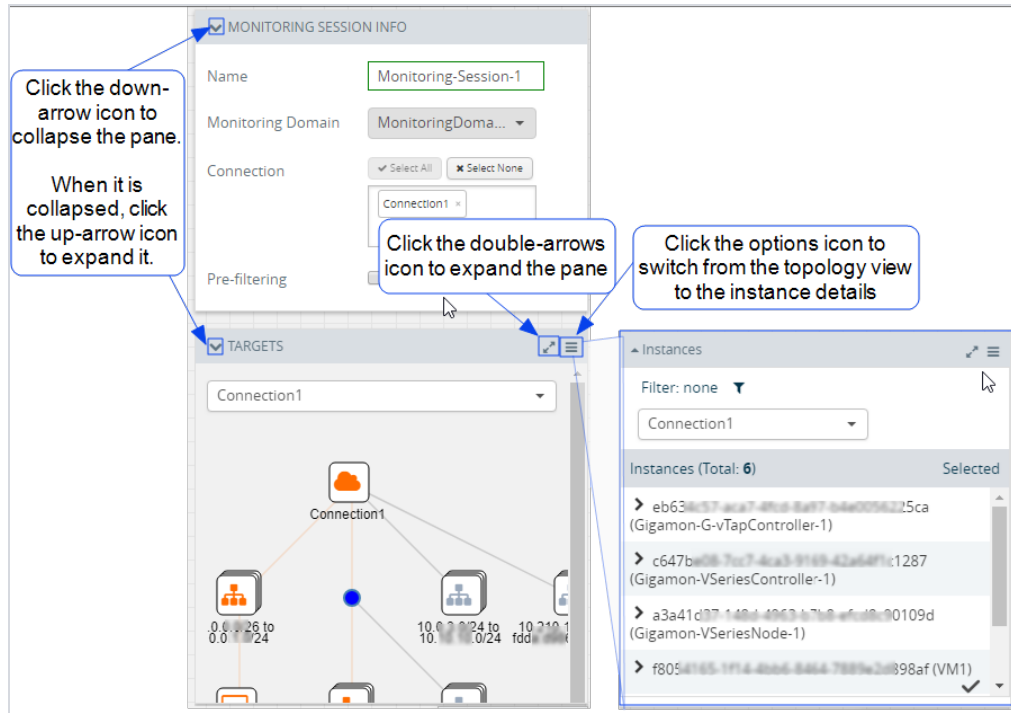
2. Select **OpenStack > Monitoring Session**. The Monitoring Sessions page appears.
3. Click **New** to open the Create a New Monitoring Session page.

The screenshot shows a web form titled "Create A New Monitoring Session". The form contains the following fields and controls:

- Alias:** A text input field containing "Monitoring-Session-1".
- Monitoring Domain:** A dropdown menu showing "MonitoringDoma..." with a downward arrow.
- Connection:** Two buttons, "Select All" and "Select None", are positioned above a list box. The list box contains one item, "Connection1", with a close button (x) to its right.
- Pre-filtering:** A checkbox that is currently unchecked.
- Buttons:** At the bottom right of the form, there are two buttons: a green "Create" button and a red "Cancel" button.

4. Enter the appropriate information for the monitoring session. Refer to [Table 12: Fields for Monitoring Session Info](#) for more information about the fields.
5. Click **Create**. The Monitoring Session details page appears displaying the specified session information and target VMs.





- If multiple projects had been selected in the connections page, the topology view will show instances in all of the selected projects.

Table 12: Fields for Monitoring 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.  <b>NOTE:</b> Agent Pre-filtering must be deselected for OVS Mirroring or OVS Mirroring + DPDK.

## 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 12: Fields for Monitoring Session Info on page 73](#).

*Table 13: Fields for Cloning the Monitoring Session.*

Field	Description
<b>Alias</b>	The name of the monitoring session.
<b>Monitoring Domain</b>	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 12: Fields for Monitoring Session Info](#) on page 73.

Table 14: Fields for Splitting the Monitoring Session.

Field	Description
<b>Original Monitoring Session</b>	<p><b>Alias:</b> The name of the original monitoring session from which a split monitoring session is to be created.</p> <p><b>Connections:</b> Connections that belong to the original monitoring session.</p>
<b>New Monitoring Session</b>	<p><b>Alias:</b> The name of the new monitoring session that is to be created.</p> <p><b>Connections:</b> Connections that have been added to the new monitoring session.</p>

**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 16: Fields for Creating a New Map](#) lists the various rule conditions that you can select for creating a map, inclusion map, and exclusion map.

Table 15: Conditions for the Rules

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

Conditions	Description
<b>VLAN Tag Control Information (TCI)</b>	All the traffic matching the specified VLAN TCI value is filtered. Specify the exact TCI value.
<b>Pass All</b>	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 71](#).
4. From **Maps**, drag and drop a new map template to the workspace.
5. Click on the map, then click details. Refer to [Figure 25: Map Details on page 78](#).




Figure 25: Map Details

The map rules quick view is displayed as shown in [Figure 26: Creating a New Map](#).

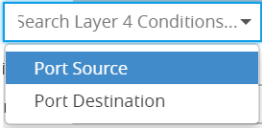
A screenshot of the 'Map\_1' configuration form. The form has an orange header bar with 'Map\_1' on the left and 'Save' and 'Add to Library' buttons on the right. Below the header, there are several sections: 'Alias' with a text input field containing 'Map\_1'; 'Comments' with a text input field containing 'Comments'; 'Map Rules' with an 'Add a Rule' button; 'Rule 1' with three dropdown menus for 'Search Layer 2 Conditions...', 'Search Layer 3 Conditions...', and 'Search Layer 4 Conditions...'; 'Priority' with a text input field containing '0' and 'ActionSet' with a text input field containing '0'; and 'Rule Comment' with a text input field containing 'Comment'. Below the 'Rule Comment' field, there are two expandable rule entries. The first is 'Ether Type' with a value of 'IPv4' and '0x0800'. The second is 'Protocol' with a value of 'TCP' and '6'. Each entry has a close button (x) on the right.

Figure 26: Creating a New Map

6. Enter the appropriate information for creating a new map as shown in [Table 16: Fields for Creating a New Map](#).

Table 16: Fields for Creating a New Map

Parameter	Description
Alias	<p>The name of the new map.</p> <p><b>NOTE:</b> The name can contain alphanumeric characters with no spaces.</p>
Comments	The description of the map.
<b>Rule Conditions</b> <b>Map Rules</b>	<p>The rules for filtering the traffic in the map.</p> <p>To add a map rule:</p> <ol style="list-style-type: none"> <li>Click <b>Add a Rule</b>.</li> <li>Select a condition from the <b>Search L2 Conditions</b> drop-down list and specify a value. Based on this selection, the Search L3 Conditions drop-down list is automatically updated.           <div data-bbox="607 823 870 1115" data-label="Image"> </div> </li> <li>Select a condition from the <b>Search L3 Conditions</b> drop-down list and specify a value.           <div data-bbox="607 1220 870 1512" data-label="Image"> </div> </li> <li>(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</li> </ol>

Parameter	Description
	<p>conditions other than TCP or UDP, then the Search L4 Conditions drop-down list is disabled.</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><b>NOTE:</b> Repeat steps <b>b</b> through <b>f</b> to add more conditions.</p> <p><b>NOTE:</b> Repeat steps <b>a</b> through <b>f</b> to add nested rules.</p>

**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**.

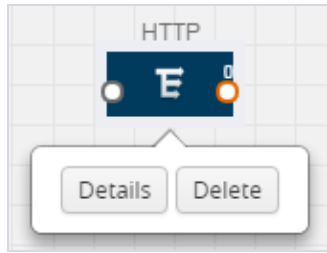
**NOTE:** The maps saved in the Map Library can be reused in any monitoring session created in the project.

8. Click **Save**.

Options:



- To edit or delete a map, click a map and select **Details** to edit the map or **Delete** to delete the map.






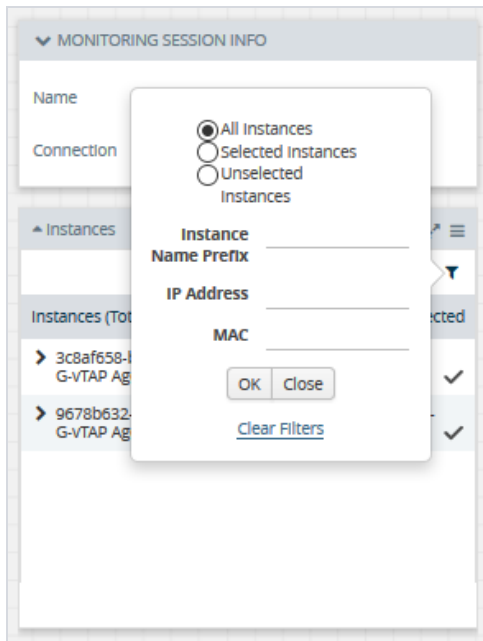
- Click the **Show Targets** button to view the monitoring targets highlighted in blue.



- 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.



- Click on the Filter icon to filter Instances based on the Instance Name Prefix, IP address, or MAC address.



## 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.

**NOTE:** Agent pre-filtering is not supported for OVS Mirroring and OVS Mirroring + DPDK.

### 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:

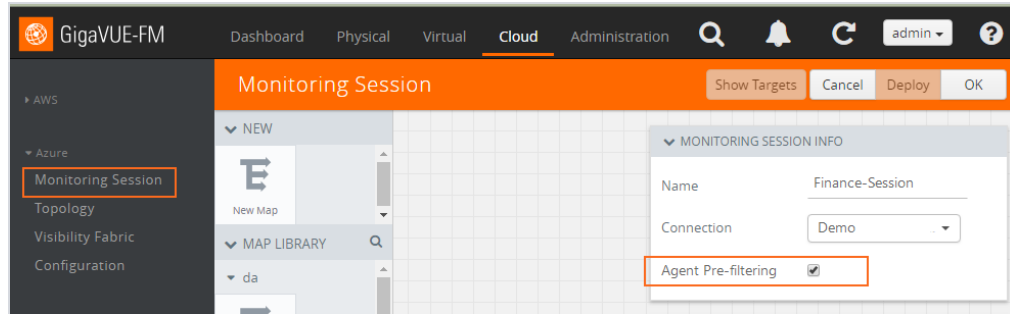
- The agent pre-filtering option can be enabled or disabled at the monitoring-session level and is enabled by default.
- When enabled, traffic is filtered at the G-vTAP agent-level, before mirroring to the V Series Nodes. Consequently, traffic flow to the V Series Nodes is reduced, which reduces the 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						
Name	Connection	# of Targets	Status	Statistics	Pre-capture Filtering	
<input type="checkbox"/> Finance-Session	Demo	4	<span style="color: green;">●</span> Success	<a href="#">View</a>	Yes	
<input type="checkbox"/> HR-Session	Demo	4	<span style="color: green;">●</span> Success	<a href="#">View</a>	No	

## Adding Applications to the Monitoring Session

Gigamon supports the following GigaSMART applications:

- [Sampling on page 86](#)
- [Slicing on page 88](#)
- [Masking on page 90](#)
- [NetFlow on page 92](#)

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 27: Dragging the Sample Application

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



Figure 28: Selecting Details

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

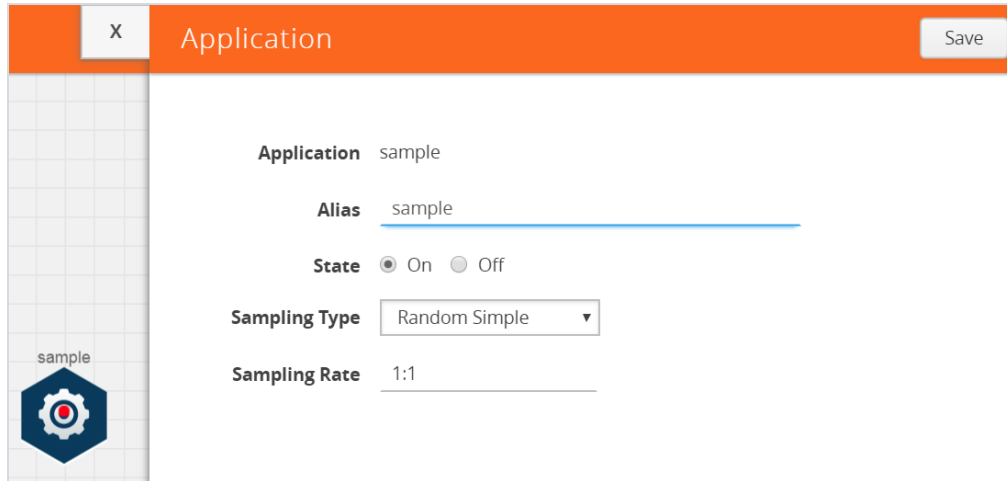


Figure 29: Viewing Sample Application Quick View

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.

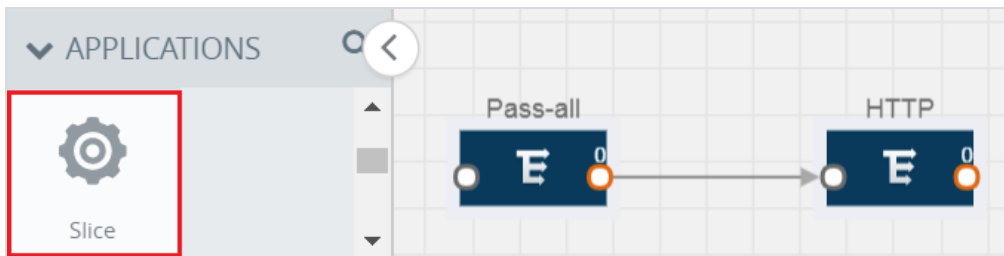


Figure 30: Dragging the Slice Application

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

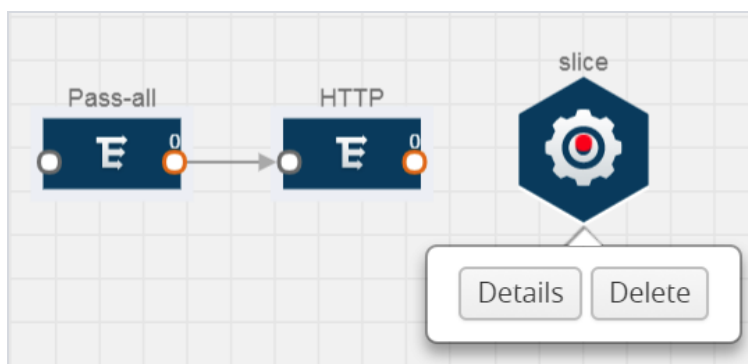


Figure 31: Selecting Details

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



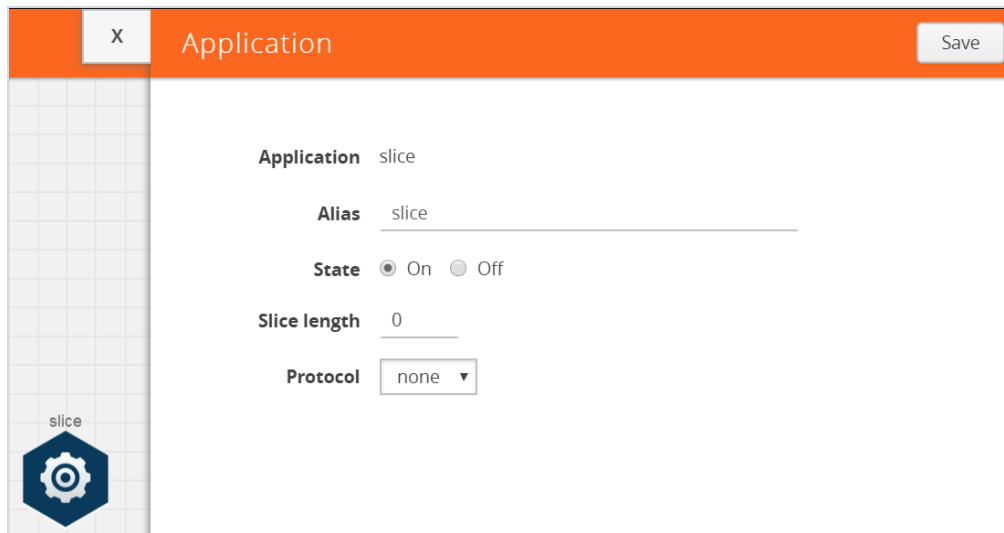


Figure 32: Viewing Slice Application Quick View

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 33: Dragging the Mask Application

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

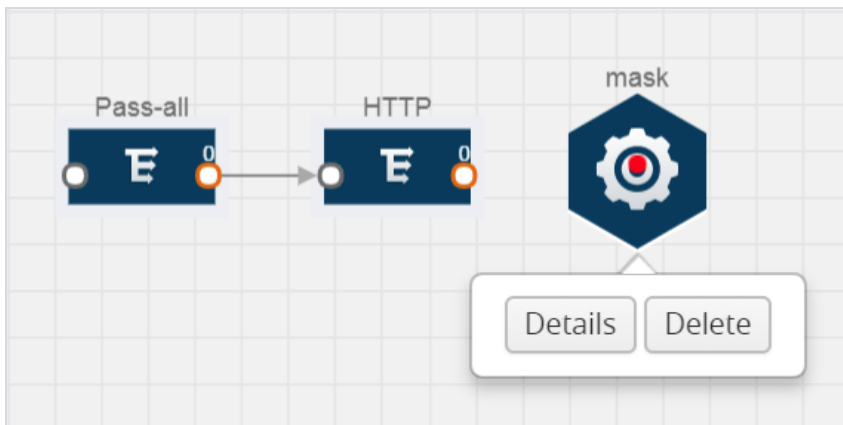


Figure 34: Selecting Details

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

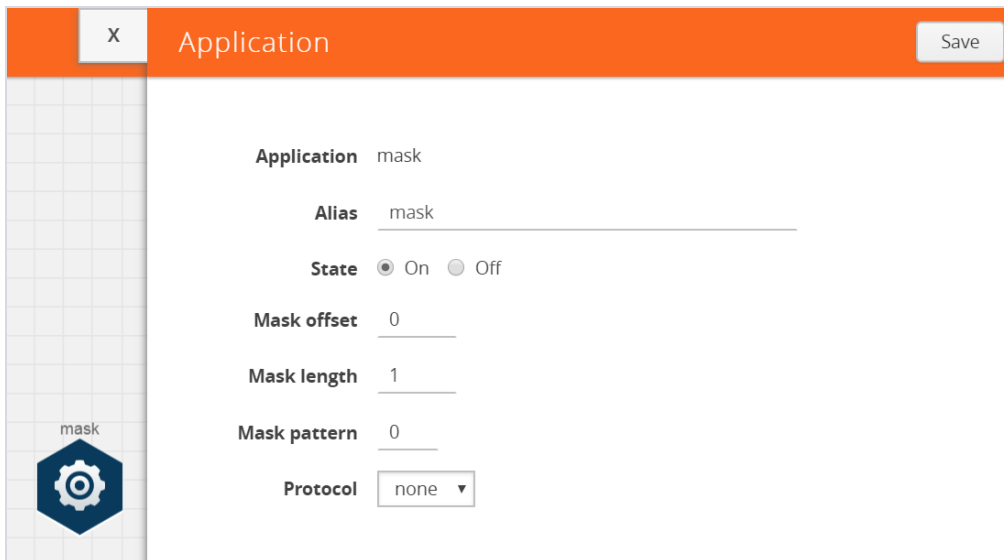


Figure 35: Viewing Mask Application Quick View

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.

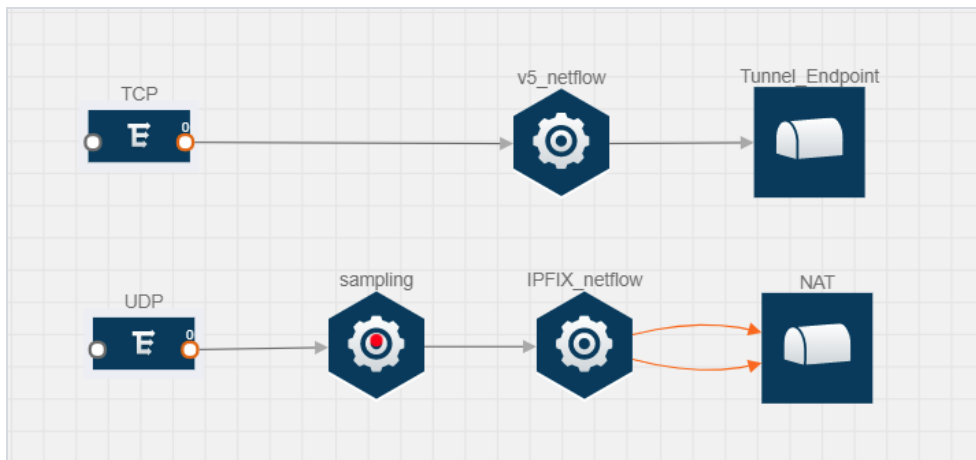
The following are the key benefits of NetFlow application:

- 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 94](#). 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 97](#).

[Figure 36: NetFlow on GigaVUE V Series Node on page 93](#) shows an example of a NetFlow application created on a GigaVUE V Series node in the monitoring session.



*Figure 36: NetFlow on GigaVUE V Series Node*

The NetFlow record generation is performed on GigaVUE V Series node running the NetFlow application. In [Figure 36: NetFlow on GigaVUE V Series Node on page 93](#), 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 105](#).

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 17: Match/Key Elements*

	Description	Supported NetFlow Versions
<b>Data Link</b>		
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

	Description	Supported NetFlow Versions
Source MAC	Configures the IEEE 802 source MAC address as a key field.	v9 and IPFIX
<b>IPv4</b>		
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
<b>Network</b>		
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

	Description	Supported NetFlow Versions
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
<b>IPv6</b>		
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
<b>Transport</b>		
L4 Dest Port	Configures the destination port identifier in the transport header as a key field.	v9 and IPFIX
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



	Description	Supported NetFlow Versions
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 18: Collect/Non-Key Elements*

	Description	Supported NetFlow Versions
<b>Counter</b>		
Byte Count	Configures the number of octets since the	v9 and IPFIX

	Description	Supported NetFlow Versions
	previous report in incoming packets for the current flow as a non-key field.	
Packet Count	Configures the number of incoming packets since the previous report for this flow as a non-key field.	v9 and IPFIX
<b>Data Link</b>		
Destination MAC	Configures the destination MAC address as a non-key field.	v9 and IPFIX
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
<b>Timestamp</b>		
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
<b>Flow</b>		

	Description	Supported NetFlow Versions
Flow End Reason	Configures the reason for Flow termination as a non-key field.	IPFIX
<b>IPv4</b>		
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
<b>Network</b>		
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
<b>IPv6</b>		
IPv6 Dest IP	Configures the IPv6 destination address in the IP packet header as a key field.	v9
<b>IPv6 Flow Label</b>	Configures the value of the IPv6 flow label field in the IP packet header as a key field.	v9

	Description	Supported NetFlow Versions
<b>IPv6 Src IP</b>	Configures the IPv6 source address in the IP packet header as a key field.	v9
<b>Transport</b>		
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.

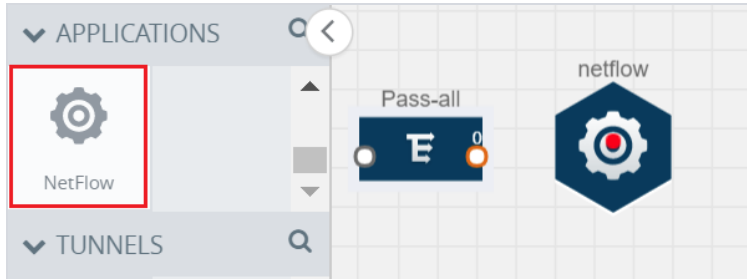


Figure 37: Dragging the NetFlow Application

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

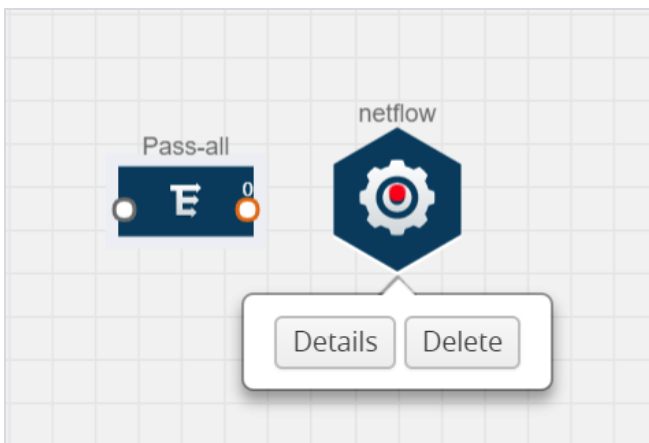


Figure 38: Selecting Details

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

Application Save

**Application** netflow

**Alias** Netflow\_V5

**State**  On  Off

**NetFlow version** v5

**Active cache timeout** 1800

**Inactive cache timeout** 15

Figure 39: 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 109](#).

### 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.

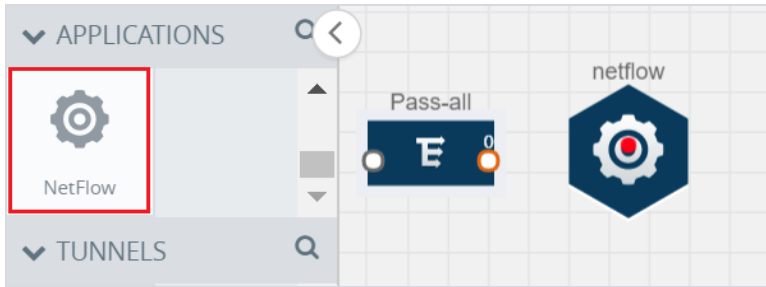


Figure 40: Dragging the NetFlow Application

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

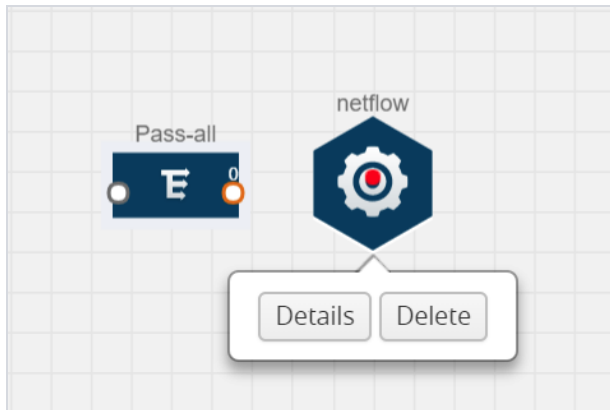


Figure 41: Selecting NetFlow Details

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

**Application** netflow

**Alias** Netflow\_IPFIX

**State**  On  Off

**NetFlow version** IPFIX

**Source Id** 1

**Match fields** L4 Src Port × IPv4 Src IP ×

**Collect 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 × IP Total Length ×

**Active cache timeout** 1800

**Inactive cache timeout** 15

**Template refresh interval** 1800

Figure 42: 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 94](#).
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 97](#).
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 109](#).

### **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.

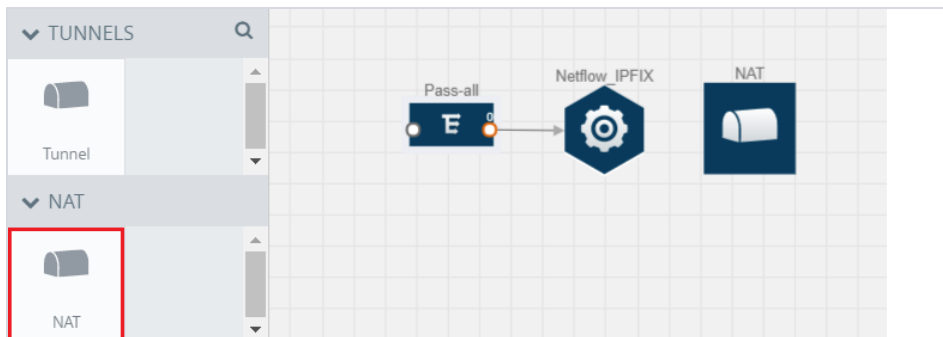


Figure 43: Adding NAT

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



Figure 44: Selecting Details

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

Save

Alias: NAT-to-Splunk

Local Subnet: 192.168.50.0/24

	Destination IP	Node Interface Subnet Cidr
x	10.0.2.189	10.0.1.0/24

Figure 45: 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.

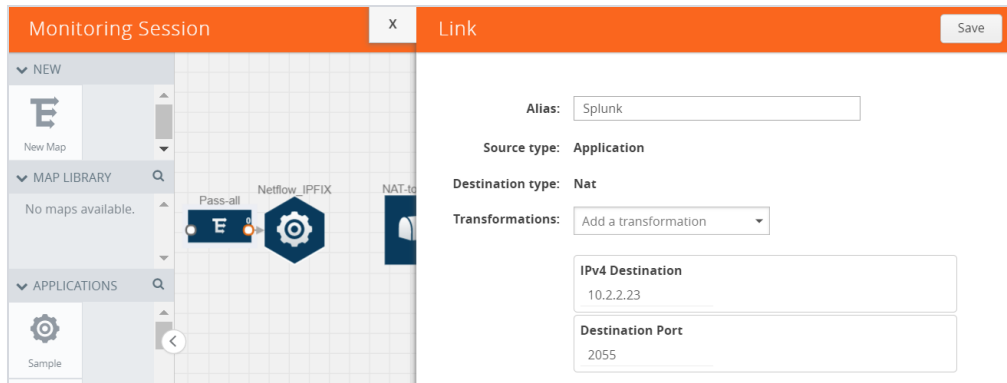


Figure 46: Creating a Link from NetFlow to NAT

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.

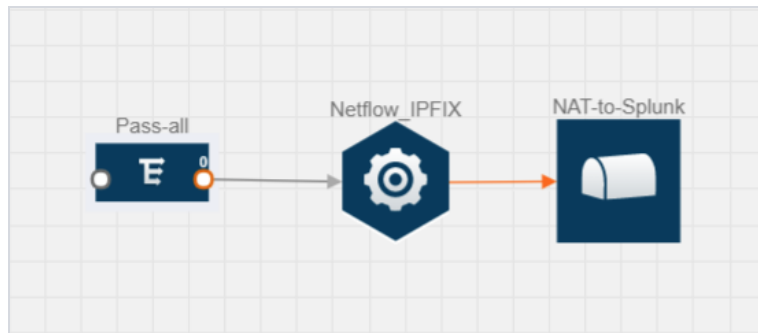


Figure 47: Linking NetFlow to NAT

7. 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 109](#)
- [Example 2 on page 114](#)

#### 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.

1. Create a monitoring session. For steps, refer to [Creating a Monitoring Session on page 71](#).

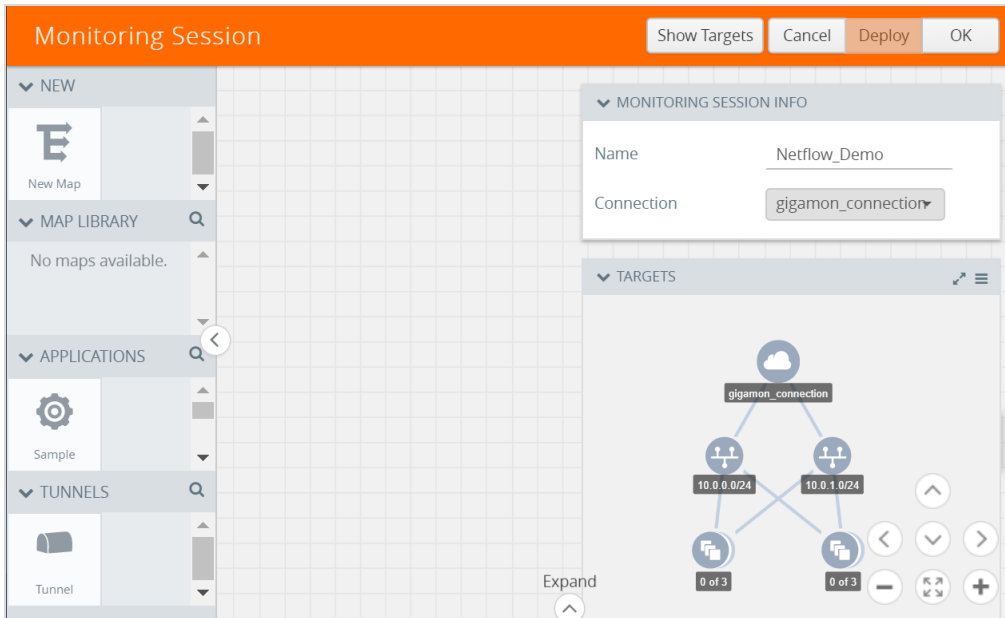


Figure 48: Creating a Monitoring Session

2. 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 75](#).

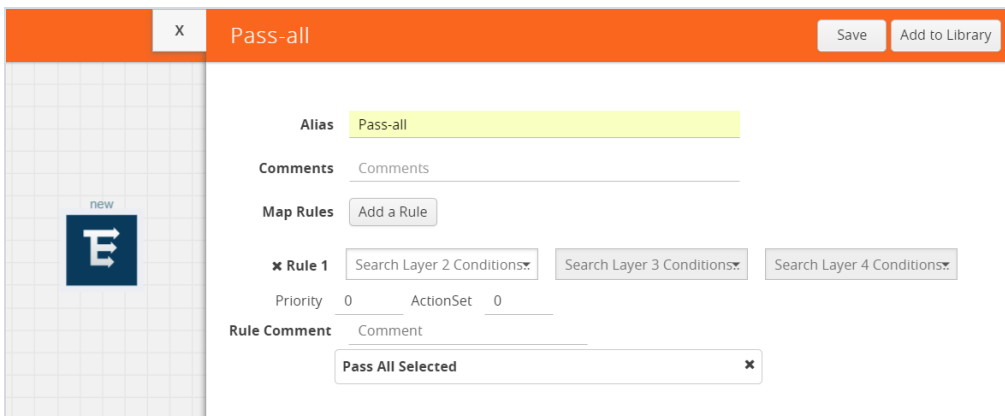
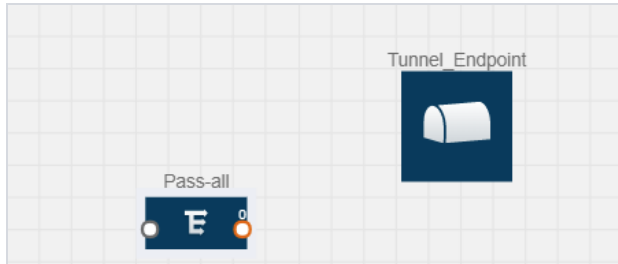


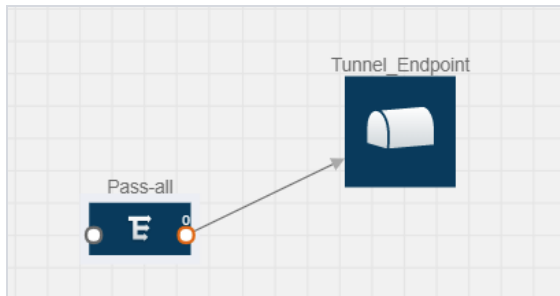
Figure 49: Creating a Pass All Map

3. 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 50: Adding a Tunnel*

4. 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 51: Creating a Link from Pass-all Map to Tunnel\_Endpoint*

5. Drag and drop a v5 NetFlow application.




Figure 52: Adding a link from Pass-all Map to Tunnel\_Endpoint

- 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 100](#).

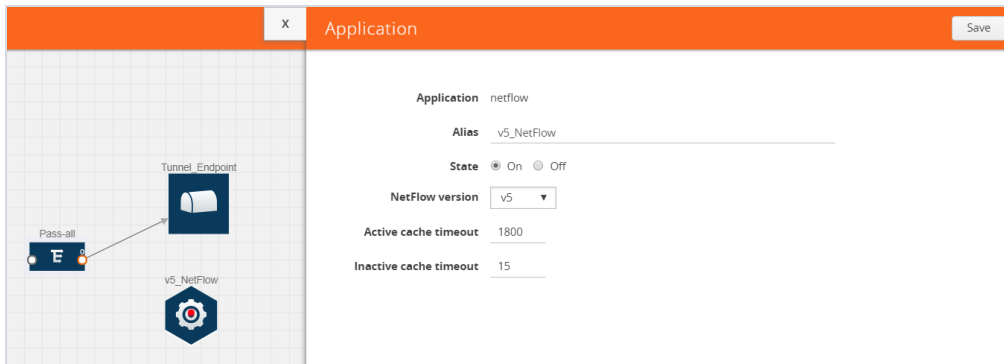


Figure 53: Configuring the NetFlow Application

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






Figure 54: 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 106](#).

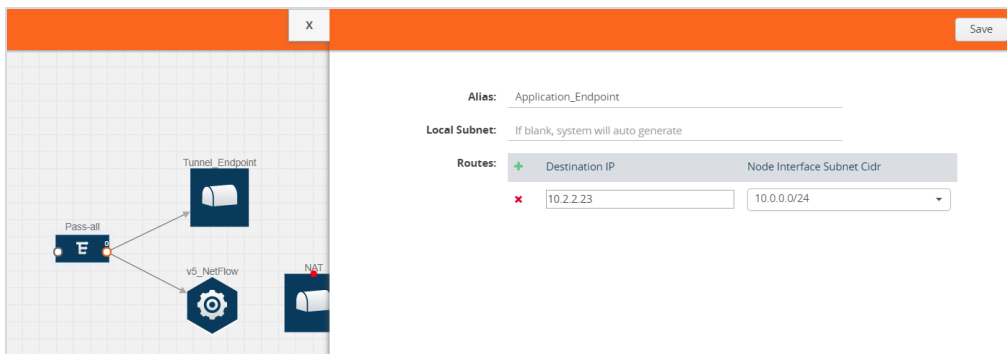


Figure 55: Adding a NAT Device

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 107](#).

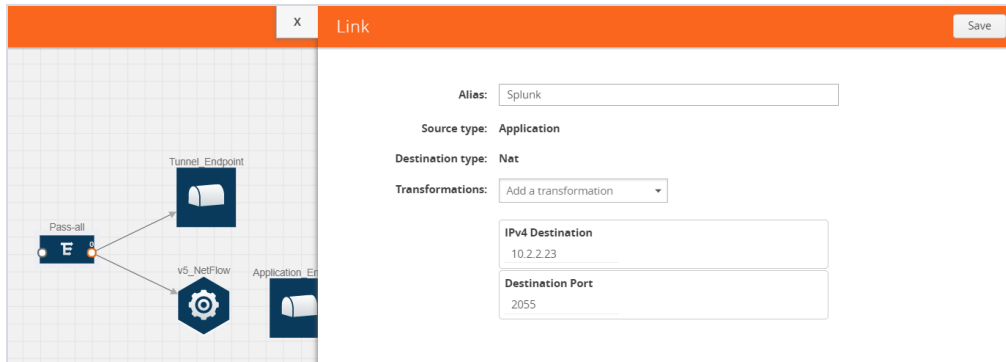


Figure 56: Adding a Link from v5 NetFlow Application to NAT

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

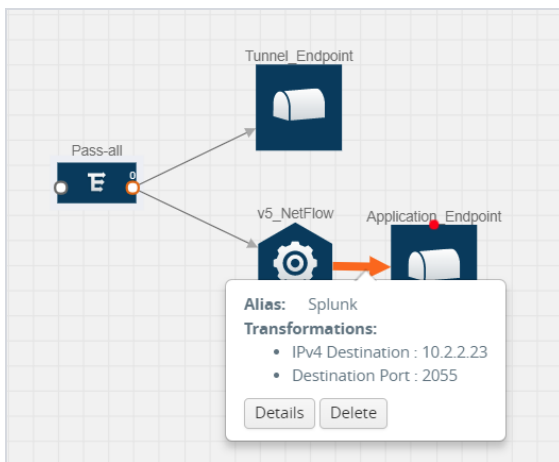


Figure 57: Viewing the Transformation Dialog Box

## 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.

1. Create a monitoring session. For steps, refer to [Creating a Monitoring Session on page 71](#).

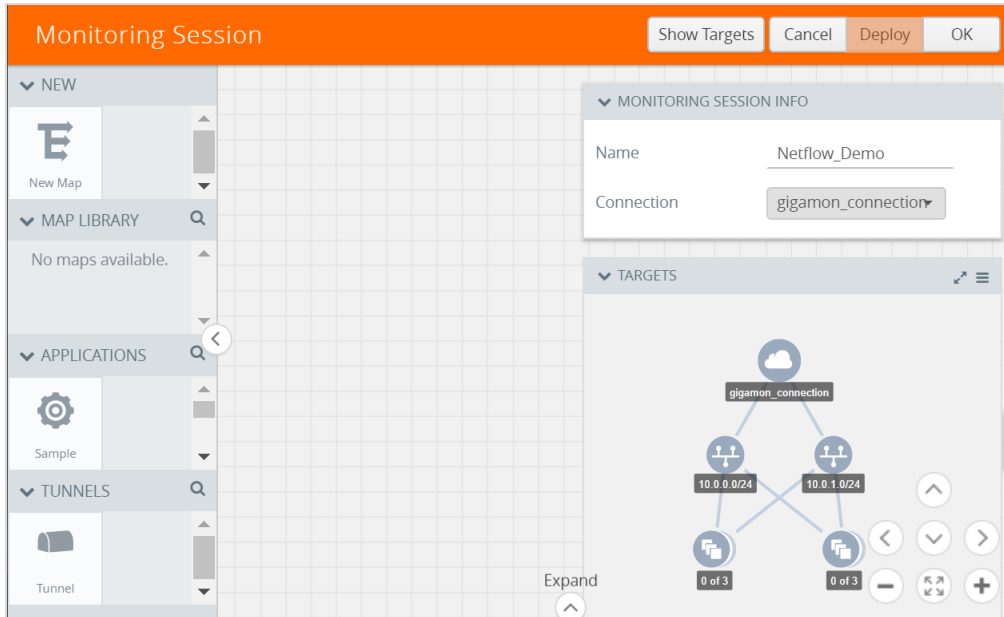


Figure 58: Creating a Monitoring Session

2. Create a map rule to filter the TCP packets. For steps on creating a map, refer to [Creating a Map on page 75](#).

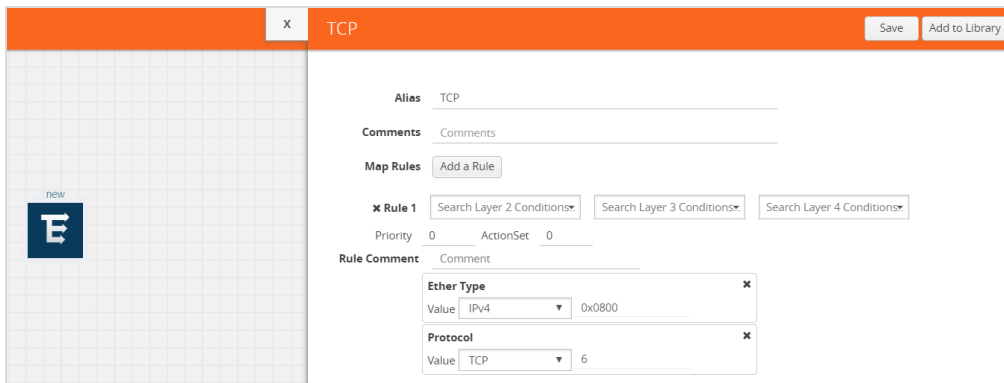


Figure 59: Creating a TCP Map

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

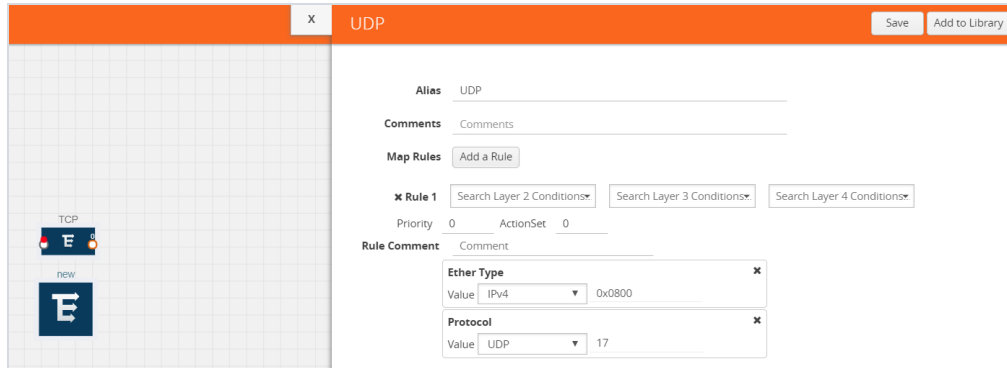


Figure 60: Creating a UDP Map

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

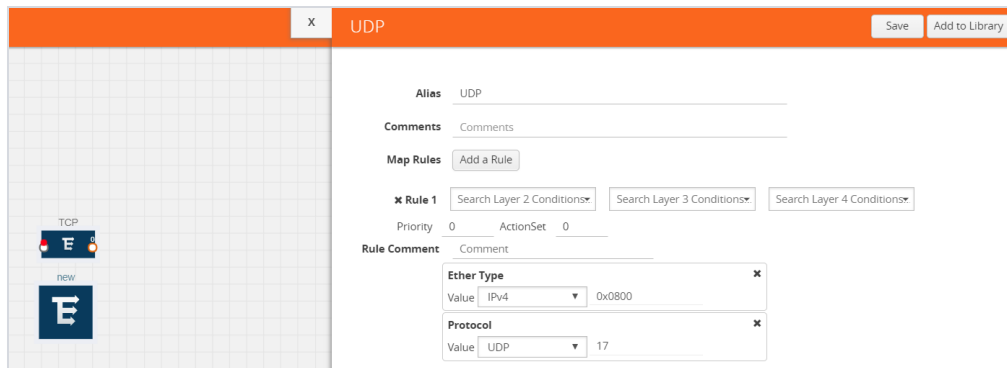


Figure 61: Creating a Map to Filter UDP Packets

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 102](#).

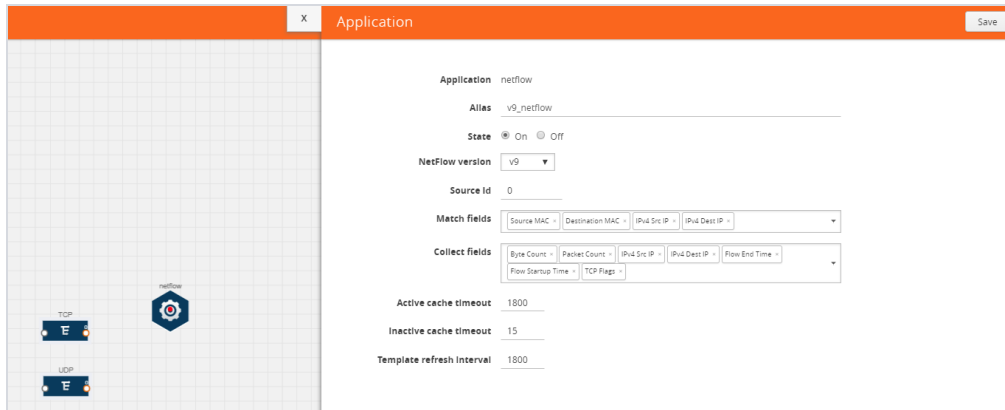


Figure 62: Adding a v9 NetFlow Application

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




Figure 63: 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 102.

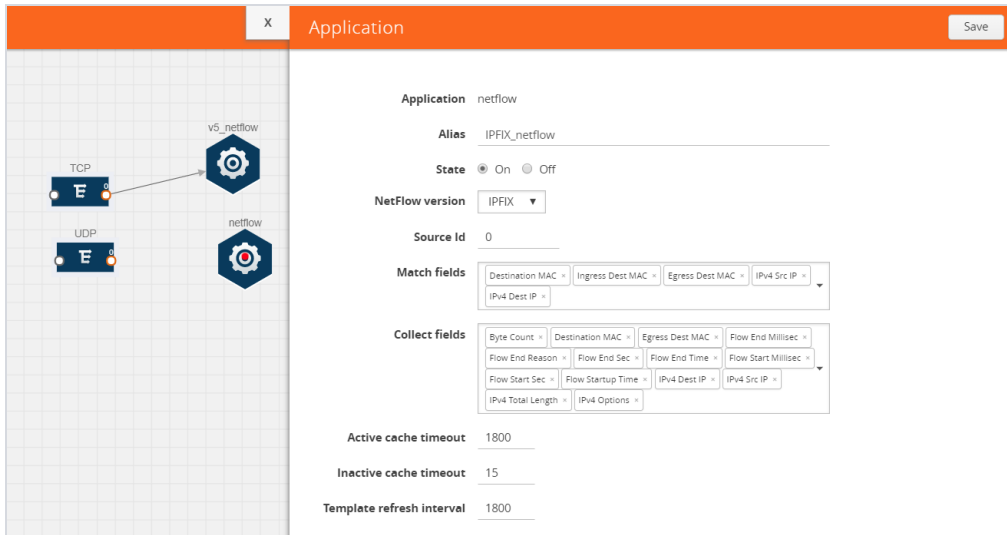


Figure 64: Adding a IPFIX NetFlow Application

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

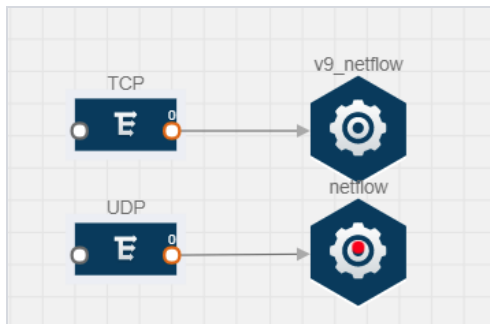


Figure 65: Adding a Link from the UDP Map to the IPFIX NetFlow Application

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 106](#).

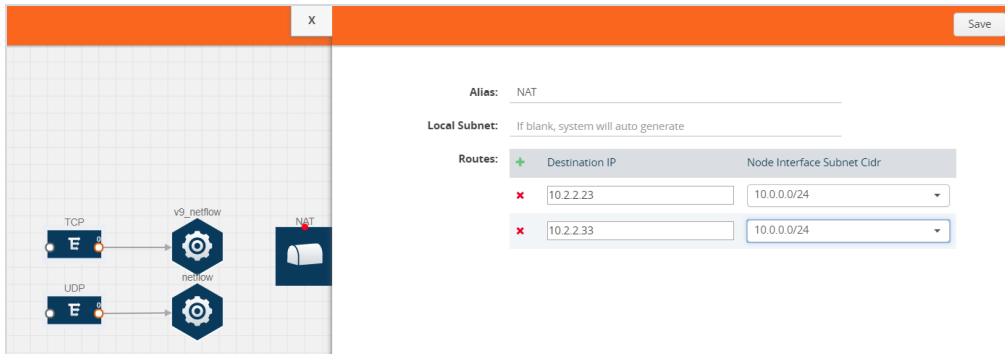


Figure 66: Adding a NAT

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

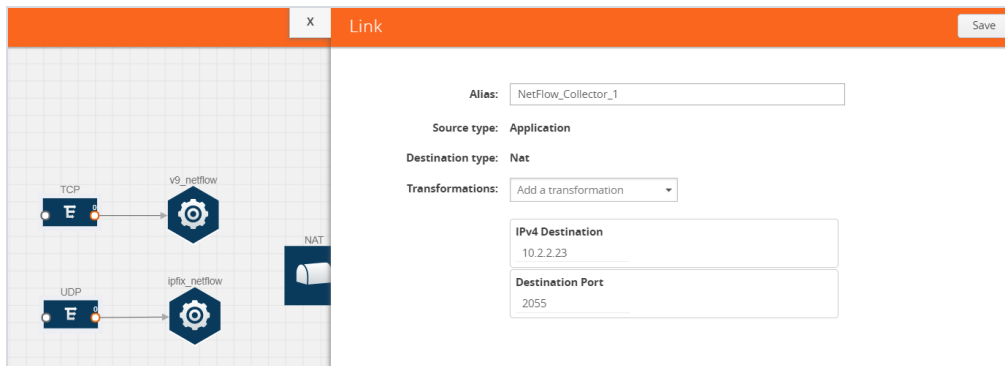


Figure 67: Adding a Link from NetFlow Application to NAT

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

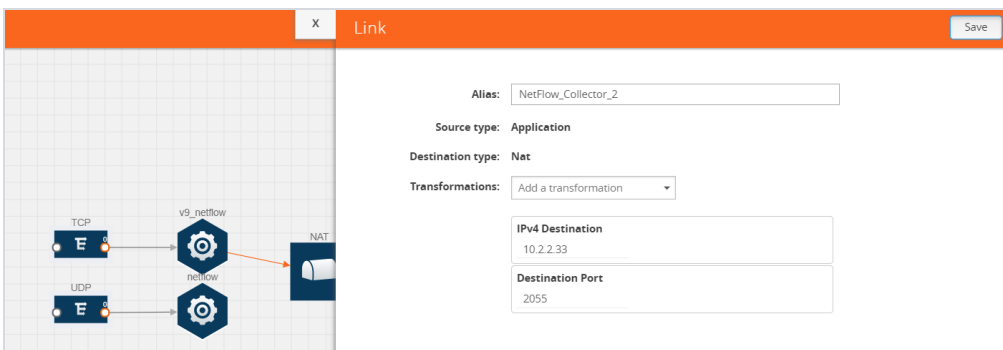


Figure 68: Adding a Link from NetFlow Application to NAT

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

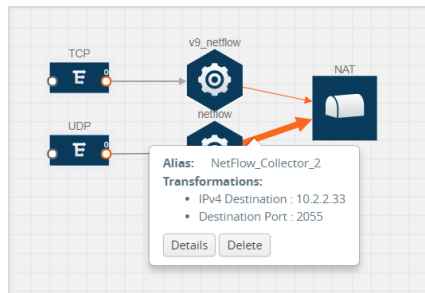


Figure 69: Viewing the Header Transformation

## 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 85](#).

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

Figure 70: Dragging and Dropping the Maps, Applications, and Monitoring Tools 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.



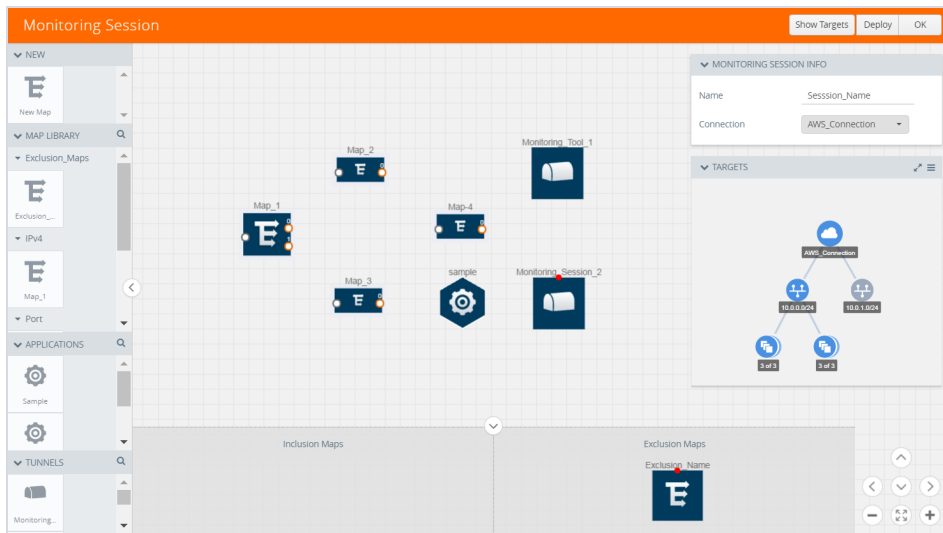


Figure 70: Dragging and Dropping the Maps, Applications, and Monitoring Tools

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 71: Connecting the Maps, Applications, and Monitoring Tools](#) on page 122.

**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 71: Connecting the Maps, Applications, and Monitoring Tools](#), the traffic matching the rules in each action set is routed to maps, applications, or monitoring tools.

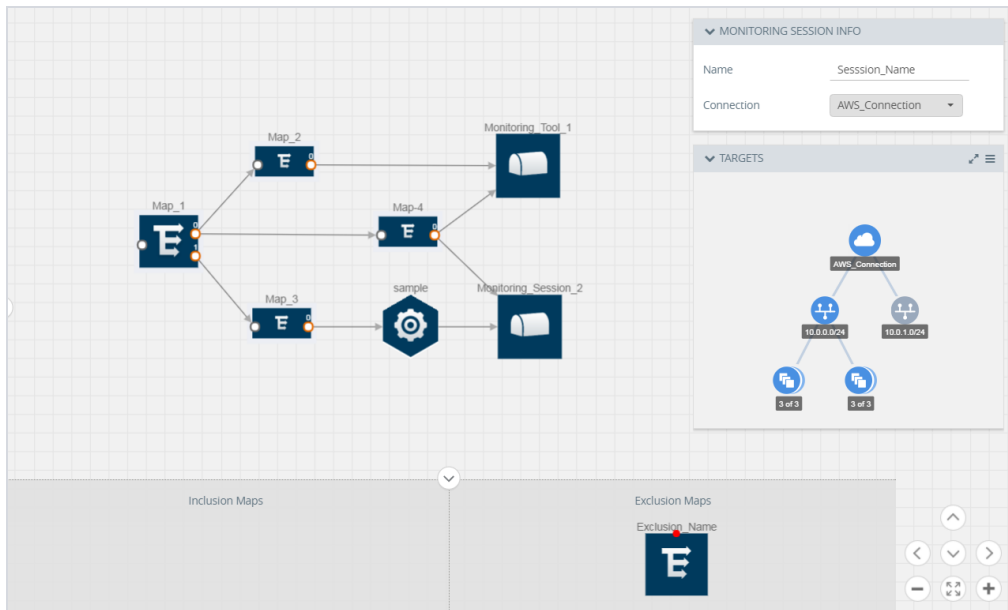


Figure 71: Connecting the Maps, Applications, and Monitoring Tools

1 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.

7. 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				
<a href="#">Redeploy</a> <a href="#">Undeploy</a> <a href="#">New</a> <a href="#">Clone</a> <a href="#">Edit</a> <a href="#">Delete</a>				
<input type="checkbox"/> Name	Connection	# of Targets	Status	Statistics
<input type="checkbox"/> Connection1_Monitoring	Connection1	2	<span style="color: green;">●</span> <a href="#">Success</a>	<a href="#">View</a>
Total Items : 1				

When you click on the Status link, the Deployment Report is displayed. Refer to [Figure 72: Monitoring Session Deployment Report](#).

Monitoring Sessions		Deployment Report	
<input type="checkbox"/> Name	Connection	Monitoring Session Alias :	Example1_Monitor
<input type="checkbox"/> Example1_Monitor	Example1	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 72: 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 73: Action Set with Multiple Links on page 125](#), 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 73: Action Set with Multiple Links

GigaVUE V Series node supports the following header transformations:

Table 19: Header Transformations

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

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.

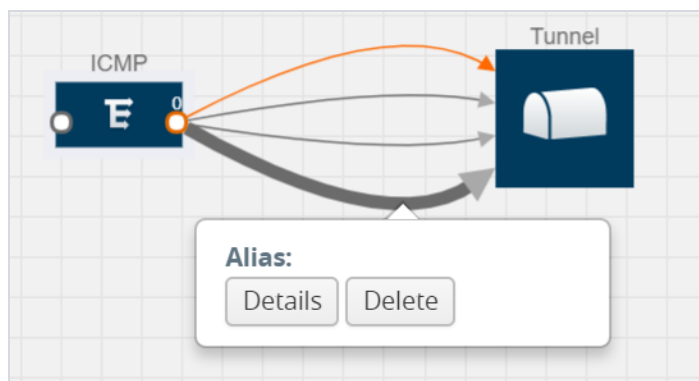


Figure 74: Opening the Link Quick View

1 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.

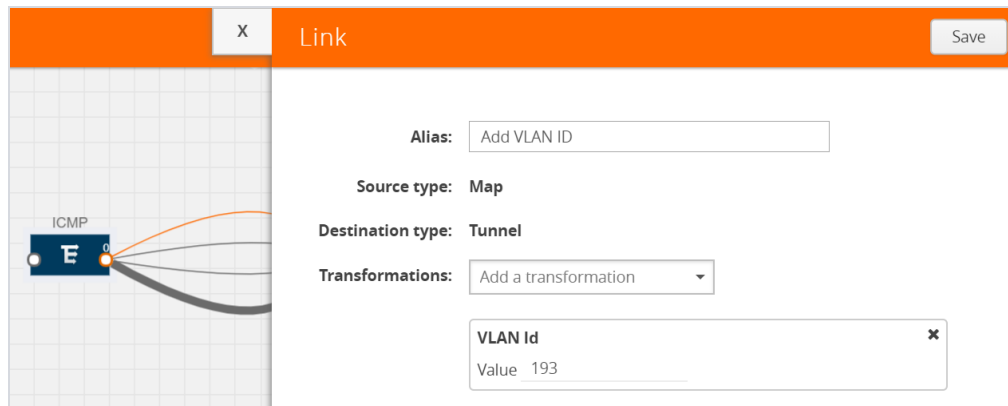


Figure 75: Adding Transformation

2. Click **Save**. The selected transformation is applied to the packets passing through the link.
3. 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					
<a href="#">Redeploy</a> <a href="#">Undeploy</a> <a href="#">New</a> <a href="#">Clone</a> <a href="#">Edit</a> <a href="#">Delete</a>					
<input type="checkbox"/>	Name	Connection	# of Targets	Status	Statistics
<input type="checkbox"/>	Connection1_Monitoring	Connection1	2	<span style="color: green;">●</span> <a href="#">Success</a>	<a href="#">View</a>
Total Items : 1					

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




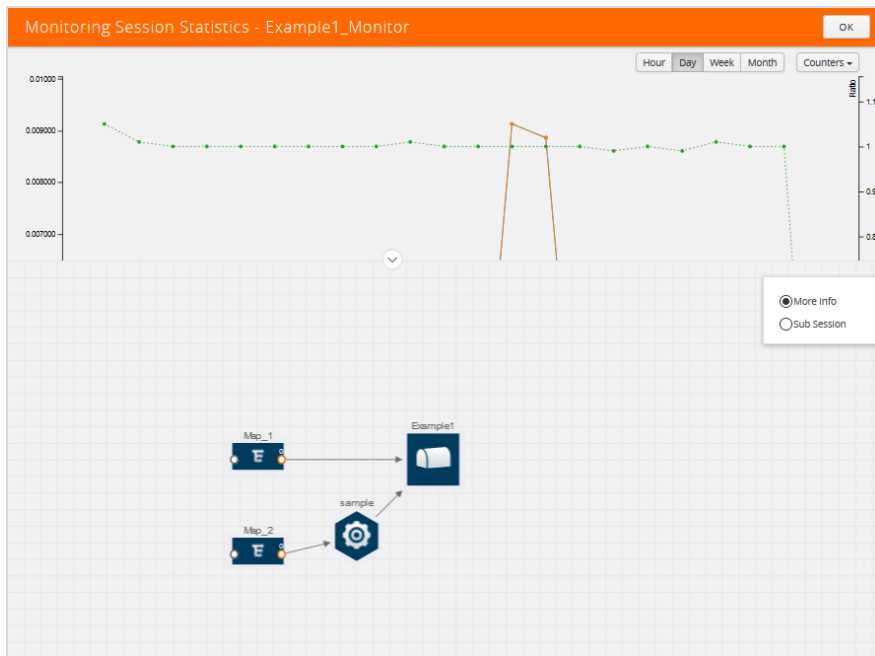
Figure 76: Monitoring Session Statistics View



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 77: Viewing the Map Statistics](#).

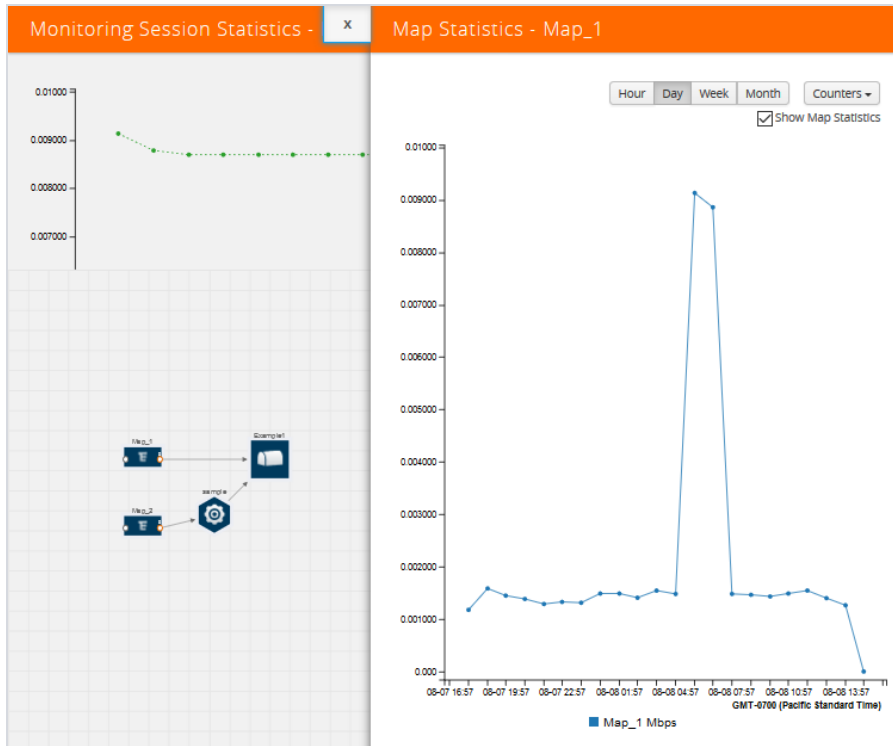


Figure 77: Viewing the Map Statistics

Figure 77: Viewing the Map Statistics 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			
<input type="checkbox"/> Action Set 0			
Map Info			
<b>Map Alias</b> Map_2			
<b>Comment</b>			

## 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 78: 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.

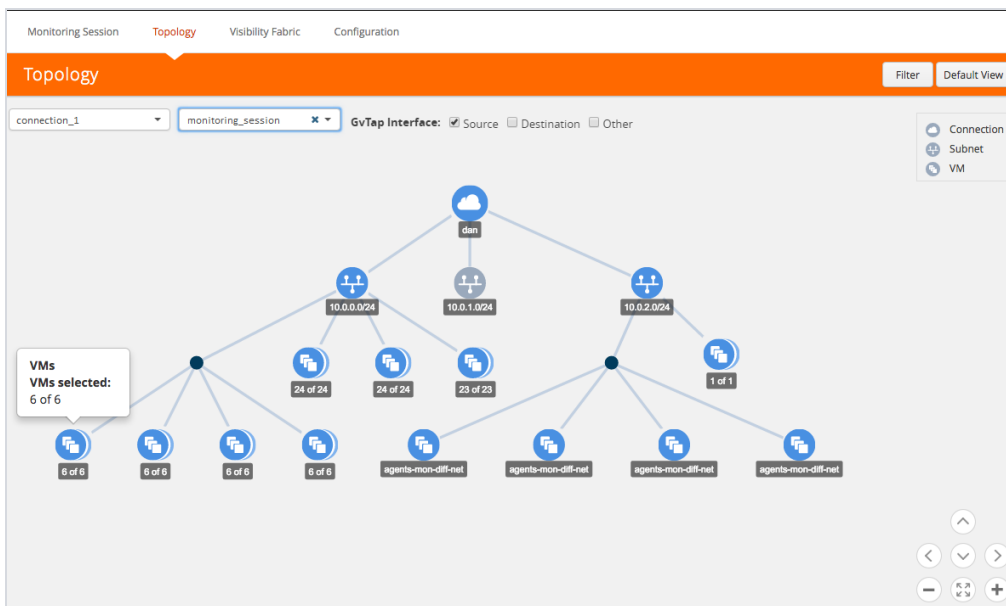


Figure 79: Viewing the Topology

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 80: Filtering in Topology View on page 133](#).

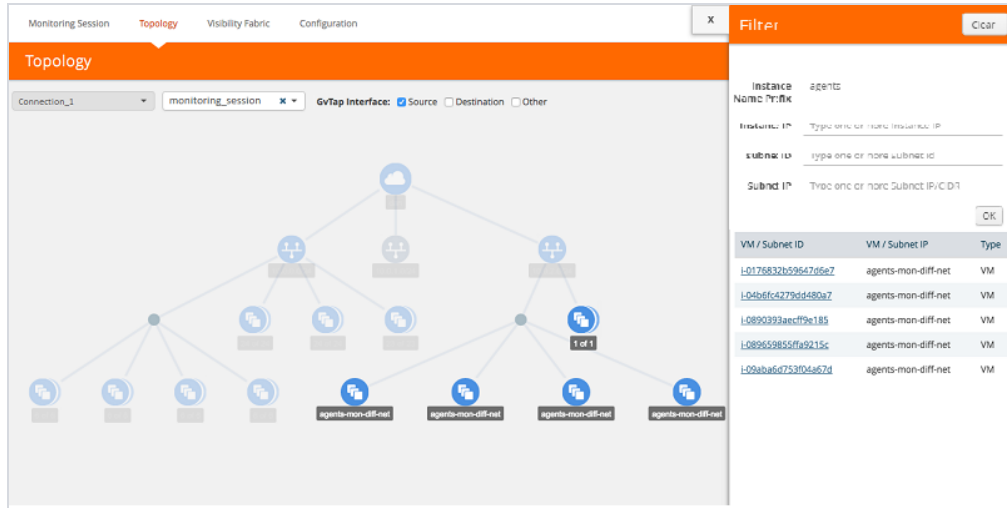


Figure 80: Filtering in Topology View

- 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 80: Filtering in Topology View on page 133](#).

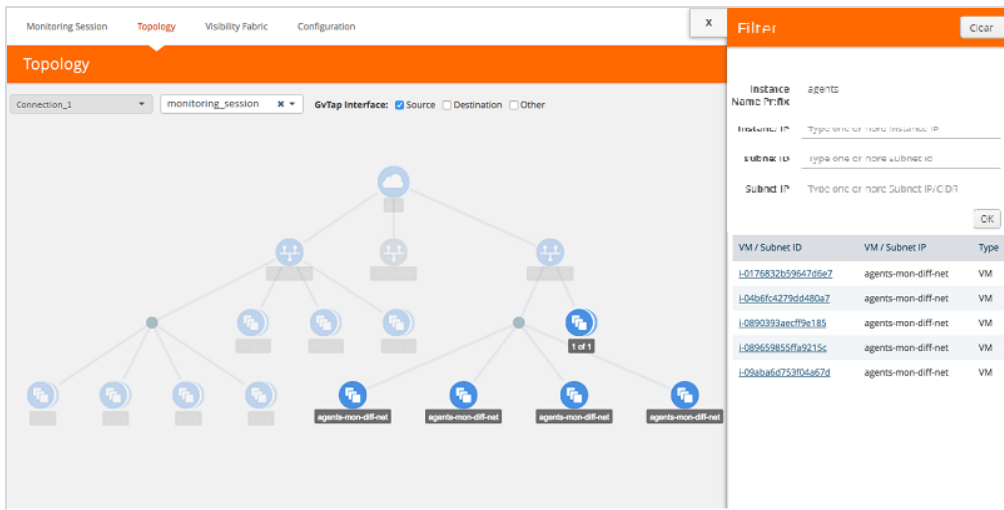



Figure 81: Filtering the Topology page

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**.

1 Select **Settings** to edit the OpenStack settings. The **Settings** page appears. Refer to [Figure 82: Cloud > OpenStack > Connections > Settings](#).

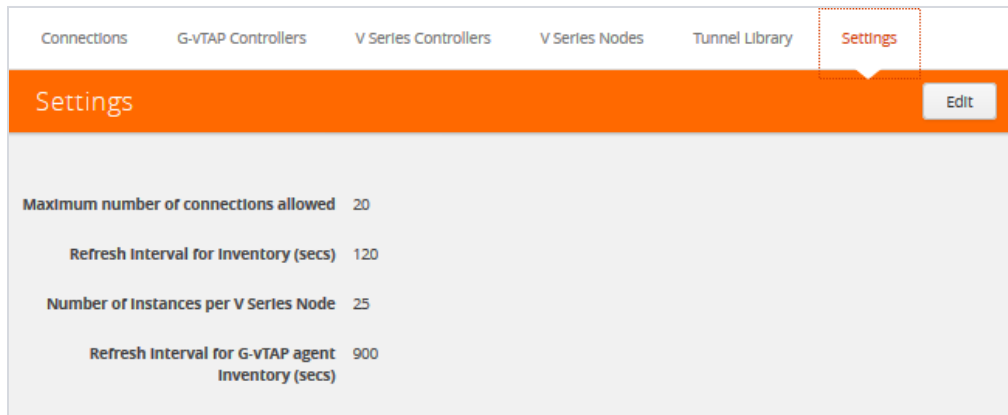


Figure 82: Cloud > OpenStack > Connections > Settings

3. Click **Edit** to edit the Settings fields. Refer to [Table 20: OpenStack Settings on page 135](#) for descriptions of the Settings fields:

Table 20: OpenStack Settings

Settings	Description
<b>Maximum number of connections allowed</b>	Specifies the maximum number of project connections you can establish in GigaVUE-FM.
<b>Refresh interval for instance inventory (secs)</b>	Specifies the frequency for updating the state of cloud instances in OpenStack.
<b>Number of instances per V Series Node</b>	Specifies the maximum number of instances that can be assigned to the V Series node.
<b>Refresh interval for G-vTAP agent inventory (secs)</b>	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 136](#)
- [Supported Features in GigaVUE V Series Nodes on page 136](#)
- [Supported Features in G-vTAP Agents on page 137](#)

## GigaVUE-FM Version Compatibility

The following table lists the different versions of GigaVUE 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



Features	GigaVUE V Series v1.4-x
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 138.

## 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 140](#)
- [Documentation Feedback on page 140](#)
- [Contacting Technical Support on page 141](#)
- [Contacting Sales on page 141](#)

## Documentation

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

Document	Summary
<b>GigaVUE-FM User's Guide</b>	Describes how to install, deploy, and operate the GigaVUE® Fabric Manager (GigaVUE-FM) from Gigamon® Inc.
<b>GigaVUE VM User's Guide</b>	Describes how to install, deploy, and operate GigaVUE® Virtual Manager (GigaVUE-VM) from Gigamon® Inc.

## 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](mailto:support@gigamon.com).

## Contacting Sales

**Table i: Sales Contact Information** shows how to reach the Sales Department at Gigamon.

*Table i: Sales Contact Information*

<b>Telephone</b>	+1 408.831.4025
<b>Sales</b>	<a href="mailto:inside.sales@gigamon.com">inside.sales@gigamon.com</a>

## 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](https://community.gigamon.com)

Questions? Contact our Community team at [community.gigamon.com](https://community.gigamon.com)