



GigaVUE Cloud for OpenStack Configuration Guide

GigaVUE Cloud Suite

Product Version: 5.8.00

Document Version: 1.0

Copyright 2019 Gigamon Inc.. All rights reserved.

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

Trademark Attributions

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

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

CONTENTS

GigaVUE Cloud Suite for OpenStack	5
Gigamon GigaVUE Cloud Components	6
Traffic Capturing Mechanism	7
G-vTAP Agent	7
OpenVSwitch (OVS) Mirroring	8
OpenVSwitch (OVS) Mirroring + DPDK	10
Configuring the Components in OpenStack	11
Before You Begin	11
Uploading the Images	17
Launching the GigaVUE-FM Instance	19
G-vTAP Agents	27
Configuring the GigaVUE Cloud in OpenStack	32
Configuring Monitoring Sessions	53
Overview of Visibility Components	53
Creating Tunnel Endpoints	56
Creating a Monitoring Session	58
Configuring the OpenStack Settings	110
Compatibility Matrix	112
GigaVUE-FM Version Compatibility	112
Supported Features in GigaVUE V Series Nodes	112
Supported Features in G-vTAP Agents	113
Troubleshooting	114
OpenStack Connection Failed	114
Handshake Alert: unrecognized_name	114
GigaVUE V Series Node or G-vTAP Controller is Unreachable	114
Additional Sources of Information	115
Documentation Downloads	115

Documentation Feedback118
Contact Technical Support118
Contact Sales 119
The Gigamon Community 119

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.

Topics:

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

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

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 7](#) 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:

- [Before You Begin on page 11](#)
- [Uploading the Images on page 17](#)
- [Launching the GigaVUE-FM Instance on page 19](#)
- [Installing the G-vTAP Agents on page 28](#)
- [Configuring the GigaVUE Cloud in OpenStack on page 32](#)

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

Supported Hypervisor

[Table 1: Hypervisor for OpenStack](#) lists the hypervisor with the supported versions for G-vTAP.

Table 1: Hypervisor for OpenStack

Hypervisor	Version
KVM	G-vTAP cPike, Queens, Ocata, Newton, Mitaka, and Liberty OVS Mirroring —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 2: Minimum Compute Requirement](#).

Table 2: Minimum Compute Requirement

Compute Instances	vCPU	Memory	Disk Space	Description
G-vTAP Agent	2 vCPU	4GB	N/A	Available as rpm or debian package. Instances can have a single vNIC or dual vNICs configured for monitoring the traffic.
G-vTAP 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) 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 3: Types of Networks](#) lists the recommended requirements to setup the network topology.

Table 3: Types of Networks

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

Virtual Network Interface Cards (vNICs)

OpenStack Cloud Instances 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 next page table.

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

Table 4: Security Group Rules

Direction	Ether Type	Protocol	Port	CIDR	Purpose
GigaVUE-FM					
Inbound	HTTPS	TCP	443	Any IP address	Allows G-vTAP Controllers, GigaVUE V Series Controllers, and GigaVUE-FM administrators to communicate with GigaVUE-FM
Inbound	IPv4	UDP	68	Any IP address	Allows GigaVUE-FM to communicate with DHCP server for assigning IP addresses and other related configuration information such as the subnet mask and default gateway
Inbound	IPv4	UDP	53	Any IP address	Allows GigaVUE-FM to communicate with DNS server for resolving the host name of the cloud controller for OpenStack
G-vTAP Controller					
Inbound	IPv4	TCP	9900	GigaVUE-FM IP address	Allows GigaVUE-FM to communicate with G-vTAP Controllers
G-vTAP Agent					
Inbound	IPv4	TCP	9901	G-vTAP Controller IP address	Allows G-vTAP Controllers to communicate with G-vTAP agents
Inbound	IPv4	TCP	9902	GigaVUE-FM IP address	Allows GigaVUE-FM to communicate with GigaVUE V Series Controllers
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					

Direction	Ether Type	Protocol	Port	CIDR	Purpose
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 UDP Rule	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

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

Access & Security

Security Groups
Key Pairs
Floating IPs
API Access

	Name	Description	Actions
<input type="checkbox"/>	default	Default security group	<input style="background-color: #e67e22; color: white; padding: 2px 10px;" type="button" value="Manage Rules"/>
<input type="checkbox"/>	sg_gvtap-agent		<input style="background-color: #e67e22; color: white; padding: 2px 10px;" type="button" value="Manage Rules"/> ▾

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

Displaying 3 items

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

Displaying 7 items

Key Pairs

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

Uploading the Images

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

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

```
$ source admin_openrc.sh
```

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

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

OR

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

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

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

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

Once the images are uploaded, they are displayed under **Compute > Images**. Refer to [Figure 1: Instances Uploaded in Images Page on page 18](#).




Figure 1: Instances Uploaded in Images Page

Launching the GigaVUE-FM Instance

To launch the GigaVUE-FM instance inside the cloud:

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




Figure 2: List of Available Images

3. Select the GigaVUE-FM image and click **Launch**. The Launch Instance dialog box is displayed. Refer to [Figure 3: Launch Instance Page on page 20](#).




Figure 3: Launch Instance Page

1 In the **Details** tab, enter the information as shown in [Figure 3: Launch Instance Page on page 20](#).

- **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**.

5. 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 4: Source tab on page 20](#).

Launch Instance

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

Select Boot Source Image **Create New Volume** Yes No

Allocated

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

Available 18 Select one

Click here for filters.

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

Cancel < Back Next > Launch Instance

Figure 4: Source tab

6. 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 5: Flavor tab on page 21](#).




Figure 5: 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 13](#). The selected network is displayed under Allocated. Click **Next**. Refer to [Figure 6: Networks tab on page 22](#).




Figure 6: 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 13](#). The selected security group is displayed under Allocated. Click **Next**. Refer to [Figure 7: Security Groups tab on page 23](#).




Figure 7: 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 17](#). The selected key pair is displayed under Allocated. Click **Next**. Refer to [Figure 8: Key Pair tab on page 24](#).




Figure 8: Key Pair tab

11. In the Configuration, Server Groups, Scheduler Hints, and Metadata tabs, click **Next**.
12. Click **Launch Instance**. The GigaVUE-FM instance takes a few minutes to fully initialize. Refer to [Figure 9: Instance Launched on page 24](#).




Figure 9: Instance Launched

13. During the initial boot-up sequence, click **Allocate Floating IP**. The Manage Floating IP Associations dialog box is displayed. Refer to [Figure 10: Assigning Floating Address on page 25](#).

Figure 10: Assigning Floating Address

14. In the Manage Floating IP Associations dialog box, enter the following information as shown in [Figure 10: Assigning Floating Address on page 25](#):
 - From the IP Address drop-down list, select an IP address.
 - From the Port to be associated drop-down list, select an appropriate port for the GigaVUE-FM instance.
15. Click **Associate**. The Floating IP is displayed in the IP Address column. Refer to [Figure 11: Floating IP of GigaVUE-FM Instance on page 25](#).

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

Figure 11: 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 12: GigaVUE-FM Instance Console on page 26](#).




Figure 12: 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 13: GigaVUE-FM Jump Start Configuration on page 26](#).




Figure 13: GigaVUE-FM Jump Start Configuration

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

NOTE: DHCP must be enabled for GigaVUE-FM.

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

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

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

G-vTAP Agents

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

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

Single vNIC Configuration

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

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

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

Multiple vNICs Configuration

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

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

Installing the G-vTAP Agents

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

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

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

Installing from an Ubuntu/Debian Package

To install from a Debian package:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Installing from an RPM package

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

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

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

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

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

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

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

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

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

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

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

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

Check the status with the following command:

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

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

Installing 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 5: Pre-configuration Checklist on page 32](#) 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 5: Pre-configuration Checklist

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

Logging in to GigaVUE-FM

To login to GigaVUE-FM, do the following:

- 1 Enter the Floating IP address of GigaVUE-FM into a browser. The GigaVUE-FM login page is displayed. Refer to [Figure 14: GigaVUE-FM Login Page on page 33](#).



Figure 14: 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 25](#), enter the changed password.
2. Click **Log in**. The GigaVUE-FM Dashboard page is displayed. Refer to [Figure 15: GigaVUE-FM on page 33](#).




Figure 15: 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 > Connections** 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 6: 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

Secure Mirror Traffic

OVS Mirroring only

Table 6: OpenStack Connection

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

Field	Description
URL	<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"> Login to OpenStack Horizon. Go to Compute > Access & Security. Click the API Access tab and copy the Identity URL.  <p>Paste the Identity URL into the URL field.</p>
Domain Name	The DNS domain name of the project.
Project Name	<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>
Region	<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>
Username	<p>The user name used to connect to the OpenStack cloud.</p> <p>NOTE: The user must belong to a role that meets the OpenStack minimum requirements for OVS Mirroring. Refer to OVS Mirroring Prerequisites.</p>
Password	The password for the OpenStack cloud.

Field	Description
Tap Method	<p>Select the type of agent used to capture traffic for monitoring:</p> <ul style="list-style-type: none"> • TaaS • G-vTAP • OVS Mirroring • OVS Mirroring + DPDK • None <p>NOTE: 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>
Projects to Monitor	<p>This field only appears, and is required, for OVS Mirroring or OVS Mirroring + DPDK.</p> <ul style="list-style-type: none"> • Click the Get Project List button to view the list of projects. <p>NOTE: The Get Project List button will only work if all the OpenStack credentials have been provided. Refer to OVS Mirroring Prerequisites.</p> <ul style="list-style-type: none"> • Select the name of the project or projects you want to monitor from the list. (There is a limit of 128 projects.) • You can click Select None to clear existing selections or Select All to add all available projects to the connection configuration. <div data-bbox="380 909 1180 1325" data-label="Image"> </div> <p>NOTE: 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>
Secure Mirror Traffic	<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>NOTE: 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 16: OpenStack Connection in GigaVUE-FM

The Connections page has the following controls:

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

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 17: Audit Logs on page 39](#).

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 17: 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 46](#).

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.

G-vTAP Configuration			New	Edit	Delete
<input type="checkbox"/>	Connection	SSH Key Pair	Management Network		
<input type="checkbox"/>	Connection1	key-pair1	f78b4aa5-4c8e-475e-a908-2b50668a8fe2b (mgmt-network)		

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 41](#) 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.

Project	Instances											
Compute	<input type="text" value="Instance Name ="/> <input type="button" value="Filter"/> <input type="button" value="Launch Instance"/> <input type="button" value="Delete Instances"/> <input type="button" value="More Actions"/>											
Overview	<input type="checkbox"/>	Instance Name	Image Name	IP Address	Size	Key Pair	Status	Availability Zone	Task	Power State	Time since created	Actions
Instances	<input type="checkbox"/>	G-vTapController-1	gvtap-cntrl_1.x-1		m1.medium	m1key_pair	Active	nova	None	Running	0 minutes	Create Snapshot
Images	<input type="checkbox"/>	FM 56203	FM 5.x		FM_Flavor	m1key_pair	Active	nova	None	Running	1 week, 5 days	Create Snapshot
Access & Security												
Network												
Identity												

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 > G-vTAP Controllers** to verify the launch of the G-vTAP Controller.

The G-vTAP Controller is displayed with the status as OK.

G-vTAP Controllers			
G-vTAP Controller Name	Management IP	Version	Status
Connection1			
G-TapController-1		cntr_1.x-1	OK
Total Items : 1			
* Note: If configured G-vTap Controller instances do not show on this page, please check Alarms/Events page for more details.			

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:

1. Click **Cloud** from the GigaVUE-FM top navigation.
2. Go to **OpenStack > Configuration > G-vTAP Controllers**. The G-vTAP Configuration page is displayed.

G-vTAP Configuration			
Connection	SSH Key Pair	Management Network	
<input type="checkbox"/>	Connection1	key-pair1	770b4aa5-4c8e-475e-a908-2b50668a8e2b (mgmt-network)

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.

G-vTAP Configuration
Save Cancel

Connection Connection1 ▼

SSH Key Pair key-pair1 ▼

Project Proj1 ▼

Security Groups default ▼

Controller Version(s) Add

Select an appropriate version

Image gigamon-gvtap-ovs-cntrlr-1.7-1 ▼

Flavor m1.medium ▼

Number of Instances

Management Network

IP Address Type Private Floating

Network f7db4aa5-4c8e-475e-a908-2b50668a8a2b (mgmt-...) ▼

Floating IPS 10.210.17.146 ▼

Additional Network(s) Add

Tag(s) Add

Agent Tunnel Type VXLAN ▼ ⓘ

G-vTAP Agent Tunnel MTU ⓘ

Select the project where the GvTAP Controller will be launched. This field only appears for OVS Mirroring and OVS Mirroring+DPDK tapping methods.

Select an appropriate version

Tunnel type used to send monitored traffic to the V Series nodes


Figure 18: 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 7: Fields for G-vTAP Configuration

Fields	Description
Connection	<p>The name of the connection.</p> <p>NOTE: For shared controller configuration, you must select the required connection for configuring the G-vTAP Controller. Peering must be active in the selected connection to allow the rest of the connections containing the V-series nodes to be monitored.</p>
SSH KeyPair	<p>The SSH key pair for the G-vTAP Controller.</p> <p>For more information about SSH key pair, refer to Key Pairs on page 17.</p>
Project	<p>Only available for OVS Mirroring or OVS Mirroring + DPDK. Select the OpenStack project where you want to deploy the fabric.</p> <p>NOTE: When setting up the V Series Controller, this field is populated with the project selected in the G-vTAP Controller setup.</p> <p>NOTE: If using with OVS Mirroring or OVS Mirroring + DPDK, a minimum of one project is required for OVS authentication.</p>
Security Groups	<p>The security group created for the G-vTAP Controller. For example, sg_gvtap-controller. For more information, refer to Security Group on page 13.</p>

Fields	Description
Controller Version(s)	<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>NOTE: If there is a version mismatch between the G-vTAP controllers and G-vTAP agents, GigaVUE-FM cannot detect the agents in the instances.</p> <p>To add multiple versions of G-vTAP Controllers:</p> <ol style="list-style-type: none"> Under Controller Versions, click Add. From the Image drop-down list, select a G-vTAP Controller image that matches with the version number of G-vTAP agents installed in the instances. From the Flavor down-down list, select a flavor for the G-vTAP Controller. In Number of Instances to Launch, specify the number of G-vTAP Controllers to launch. The minimum number you can specify is 1. <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"> Click x (delete) next to the G-vTAP Controller image to delete that version. <div data-bbox="527 1014 1333 1299" style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;">  </div> <ol style="list-style-type: none"> 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.

Fields	Description															
Management Network	<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>Network - Select the management network ID.</p> <p>IP Address Type</p> <p>The type of IP address GigaVUE-FM needs to communicate with G-vTAP controllers:</p> <ul style="list-style-type: none"> • Private—A private IP can be used when GigaVUE-FM, the G-vTAP Controller, or the GigaVUE V Series Controller reside inside the same project. • Floating—A floating IP is needed only if GigaVUE-FM is not in the same project in the cloud or is outside the cloud. GigaVUE-FM needs a floating IP to communicate with the controllers from an external network. Make sure that this floating IP will not be used by other instances in the cloud. <p>NOTE: If GigaVUE-FM resides inside the same project, no floating IPs are necessary for the controllers.</p> <p>NOTE: For V Series Nodes data network deployments, select Floating and then specify the IPs in the Floating IPs field.</p>															
Additional Network(s)	<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 Add to specify additional networks (subnets), if needed. Also, make sure that you specify a list of security groups for each additional network.</p>															
Tag(s)	<p>(Optional) The key name and value that helps to identify the G-vTAP Controller instances in your environment. For example, you might have G-vTAP Controllers deployed in many regions. To distinguish these G-vTAP Controllers based on the regions, you can provide a name (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"> Click Add. In the Key field, enter the key. For example, enter Name. In the Value field, enter the key value. For example, us-west-2-gvtap-controllers. <p>When the G-vTAP Controllers are launched in the VPC, they will appear with the custom tag:</p>  <table border="1" data-bbox="537 1528 1333 1724"> <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-2-gvtap-controllers</td> <td>i-065069f01bc1df0e3</td> <td>t2.medium</td> <td>us-west-2c</td> <td>running</td> </tr> <tr> <td>us-west-2-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-2-gvtap-controllers	i-065069f01bc1df0e3	t2.medium	us-west-2c	running	us-west-2-gvtap-controllers	i-081d78999e9ec8278	t2.medium	us-west-2c	running
Name	Instance ID	Instance Type	Availability Zone	Instance State												
us-west-2-gvtap-controllers	i-065069f01bc1df0e3	t2.medium	us-west-2c	running												
us-west-2-gvtap-controllers	i-081d78999e9ec8278	t2.medium	us-west-2c	running												


Fields	Description
Agent Tunnel Type	The type of tunnel used for sending the traffic from G-vTAP agents to GigaVUE V Series nodes. The options are GRE or VXLAN tunnels.
G-vTAP Agent MTU (Maximum Transmission Unit)	<p>The Maximum Transmission Unit (MTU) is the maximum size of each packet that the tunnel endpoint can carry from the G-vTAP agent to the GigaVUE V Series node.</p> <ul style="list-style-type: none"> For GRE, the default value is 1450. 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.

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 41 for the common fields.

V Series Controller Configuration

Connection Connection1 ▼

SSH Key Pair key-pair1 ▼

Project Proj1 ▼

Image 48374f2b-743a-46c3-af54-1684af13e3e1 (gigamon-gigavue-vseries-cntlr-1.7-1) ▼

Flavor m1.medium ▼

Security Groups default ▼

Number of Instances

Management Network

IP Address Type Private Floating

Network f7db-4aa5-4c8e-475e-a908-2b50668a8e2b (projnet-...) ▼

Floating IPs 10.218.17.148 x ▼

Additional Network(s)

Tag(s)

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](#)" below.

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.



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.



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 41 for the common fields. For additional configuration settings, refer to [Table 8: Fields for GigaVUE V Series Node Launch Configuration](#).

V Series Node Configuration
Save Cancel

Connection Connection1 ▾

SSH Key Pair key-pair1 ▾

Project Proj1 ▾

Image 83d10001-7d71-479b-879b-e1deb956215 (gigamon-gigavue-vseries-node-1.7-1) ▾

Flavor m1.medium ▾

Security Groups default ▾

Management Network f7db4a05-4c8e-475e-a908-2b50668a8e2b (mgmt-network) ▾

Data Network(s) Add

IP Address Type Private Floating

Network 1 5bb9b212-85ee-4895-912e-a4e92bb5146 (data-network) ▾

Floating IPs 10.210.17.135 × ▾

Tag(s) Add

Min Instances ⓘ

Max Instances

Tunnel MTU

Shared ⓘ

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

Table 8: Fields for GigaVUE V Series Node Launch Configuration

Parameter	Description
Management Network	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.

Parameter	Description
Data Network	<p>Network 1</p> <p>Click Add 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>IP Address Type</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"> • Private—A private IP can be used when GigaVUE-FM, the G-vTAP Controller, or the GigaVUE V Series Controller reside inside the same project. • Floating—A floating IP address specified here will be where the monitored traffic is tunneled to. The monitored traffic must be able to reach the V Series Node. <p>NOTE: For OVS Mirroring or OVS Mirroring + DPDK deployments, must select Floating in the Data Network section and then specify the IPs in the Floating IPs field. You can have multiple Floating IPs.</p> <p>NOTE: 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>
Min Instances to Launch	<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"> • 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. • When you deploy a G-vTAP monitoring session, the V Series nodes will automatically be deployed based on the # of VMs being monitored. <p>NOTE: GigaVUE-FM will delete the nodes if they are idle for over 15 minutes.</p>
Max Instances to Launch	<p>The maximum number of GigaVUE V Series nodes that can be launched in OpenStack.</p>
Tunnel MTU (Maximum Transmission Unit)	<p>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.</p>
Shared	<p>Only one V Series node configuration can be shared in a monitoring domain.</p> <p>NOTE: Not used for OpenStack configurations.</p>

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. On the left is a navigation sidebar with categories like Project, Compute, Overview, Instances, Images, Access & Security, 1.x-1 network, and Identity. The main content area is titled 'Instances' and features a search bar for 'Instance Name', a 'Filter' button, and three action buttons: 'Launch Instance', 'Delete Instances', and 'More Actions'. Below this is a table with the following columns: Instance Name, Image Name, IP Address, Size, Key Pair, Status, Availability Zone, Task, Power State, Time since created, and Actions.

Instance Name	Image Name	IP Address	Size	Key Pair	Status	Availability Zone	Task	Power State	Time since created	Actions
VSeriesNode-1	gigamon-gigavue-vseries-node-1.x-1	danz-mgmt 10.0.0.244 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
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
G-vTapController-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.

2. 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.

The screenshot shows the 'V Series Nodes' page in GigaVUE-FM. At the top, there are four tabs: 'G-VTAP Controllers', 'V Series Controllers', 'V Series Nodes' (which is selected and highlighted with an orange border), and 'Tunnel Endpoints'. Below the tabs is a header 'V Series Nodes' on an orange background. Underneath is a table with columns: 'V Series Node Name', 'Management IP', 'Version', and 'Status'. There is one entry: 'VSeriesNode-1' with a status of 'OK' (indicated by a green dot).

V Series Node Name	Management IP	Version	Status
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 53](#)
- [Creating Tunnel Endpoints on page 56](#)
- [Creating a Monitoring Session on page 58](#)
- [Configuring the OpenStack Settings on page 110](#)

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 1: Components of Traffic Visibility Sessions on page 53](#) lists the components of the monitoring session:

Table 1: Components of Traffic Visibility Sessions

Parameter	Description
Map	A map (M) is used to filter the traffic flowing through the V Series node. It is a collection of one or more rules (R). The traffic passing through a map can match one or more rules defined in the map.

Parameter	Description
<p>Rule</p>	<p>A rule (R) contains specific filtering criteria that the packets must match.</p> <p>The filtering criteria lets you determine the target instances and the (egress or ingress) direction of tapping the network traffic.</p> <p>The rules must contain the appropriate Layer 2 (L2) to Layer 4 (L4) filters defined in them. For example, if you want to filter the traffic for HTTP Port 80, you must select the following criteria:</p> <ul style="list-style-type: none"> • Layer 2—Ethertype IPv4 • Layer 3—Protocol TCP • Layer 4—Port Destination 80 <p>By default, a rule always displays conditions based on the attributes of L2. Refer to Figure 1: Default Rule Conditions on page 54.</p> <div data-bbox="548 653 846 978" style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <div style="border-bottom: 1px solid #ccc; padding: 2px 5px;">Search Layer 2 Conditions... ▾</div> <div style="padding: 2px 5px;"> <p style="background-color: #0070c0; color: white; padding: 2px 5px; margin: 0;">Ether Type</p> <p style="padding: 2px 5px; margin: 0;">MAC Source</p> <p style="padding: 2px 5px; margin: 0;">MAC Destination</p> <p style="padding: 2px 5px; margin: 0;">VLAN</p> <p style="padding: 2px 5px; margin: 0;">VLAN PCP</p> <p style="padding: 2px 5px; margin: 0;">VLAN TCI</p> <p style="padding: 2px 5px; margin: 0;">Pass All</p> </div> </div> <p style="text-align: center;"><i>Figure 1: Default Rule Conditions</i></p> <p>A rule is also associated with priority and action set.</p>
<p>Priority</p>	<p>A priority determines the order in which the rules are executed. The greater the value, the higher the priority.</p> <p>The priority value can range from 0 to 99.</p>

Parameter	Description
<p>Action Set</p>	<p>An Action Set is an exit point in a map that you can drag and create links to the other maps, applications, and monitoring tools. A single map can have multiple action sets. A single action set can have multiple links connecting to maps and applications. You can create an Action Set when you create a rule for a map.</p> <p>In the following example (refer to Figure 2: Action Set), 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 2: 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 Figure 3: Action Set with Multiple Links on page 55.</p>  <p><i>Figure 3: Action Set with Multiple Links</i></p>
<p>Link</p>	<p>A link directs the packets to flow from a map to the destination. The destination could be the other maps, applications, and the monitoring tools. In Figure 2: Action Set on page 55, the link originating from action set 0 is moving the traffic from MAP_1 to Monitoring_Tools.</p> <p>A link lets you add header transformation to the packets passing through it before they are sent to the destination. This transformation is supported only with GigaVUE V Series node v1.2-1 and above. For more information about Header Transformation, refer to Adding Header Transformations on page 102.</p>
<p>Group</p>	<p>A group is a collection of maps that are pre-defined and saved in the map library for reuse.</p>
<p>Application</p>	<p>An application performs operations such as sampling, slicing, and masking on the traffic.</p>


Parameter	Description
Inclusion Map	An inclusion map determines the instances to be included for monitoring. This map is used only for target selection.
Exclusion Map	An exclusion map determines the instances to be excluded from monitoring. This map is used only for target selection.
Target	A target determines the instances that are to be monitored. Targets are determined based on the following formula: <div style="border: 1px solid black; padding: 5px; width: fit-content; margin: 10px auto;"> $\text{Target} = (\text{Maps} \cap \text{Inclusion map}) - \text{Exclusion map}$ </div>
Automatic Target Selection (ATS)	A built-in feature that automatically selects the cloud instances based on the rules defined in the maps, inclusion maps, and exclusion maps in the monitoring session. 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. The instance targets will be selected based on two additional criteria: <ul style="list-style-type: none"> • they must be in one of the projects selected in the monitored projects list from the connection page, and • they must reside on a hypervisor that has the GvTAP Agent installed on it.
Tunnel	A tunnel lists the monitoring tools to which the traffic matching the filtered criteria is routed.

Creating Tunnel Endpoints

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

To create a tunnel endpoint:

1. In GigaVUE-FM, on the top navigation pane, select **Cloud**.
2. On the left navigation pane, select **OpenStack > Configuration**.
3. Select the **Tunnel Spec Library** tab. The Tunnel Library page appears.



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

Add Tunnel Spec
Save Cancel

Alias

Description

Type

Traffic Direction

Remote Tunnel IP

- On the Edit Tunnel page, select or enter the appropriate information in the fields. Refer to [Table 2: Field Descriptions for Tunnel Endpoint on page 57](#).

Table 2: Field Descriptions for Tunnel Endpoint

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

- Click **Save**.
- Select **OpenStack > Visibility Fabric > Tunnel Endpoints** and verify the tunnel endpoint added to GigaVUE-FM.

Create A New Monitoring Session


Alias

Monitoring Domain

Connection

Pre-filtering

4. Enter the appropriate information for the monitoring session. Refer to [Table 3: 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 3: 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. NOTE: 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 3: Fields for Monitoring Session Info on page 60](#).

Table 4: Fields for Cloning the Monitoring Session.

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

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

Splitting a Monitoring Session

You can split a monitoring session.

To split a monitoring session:

1. Select the monitoring session that you need to split from the **Monitoring Sessions** page.
2. Click **Split**.
3. Enter the appropriate information in the **Split A Monitoring Session** dialog box as shown in [Table 3: Fields for Monitoring Session Info on page 60](#).

Table 5: Fields for Splitting the Monitoring Session.

Field	Description
Original Monitoring Session	Alias: The name of the original monitoring session from which a split monitoring session is to be created. Connections: Connections that belong to the original monitoring session.

Field	Description
New Monitoring Session	<p>Alias: The name of the new monitoring session that is to be created.</p> <p>Connections: 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 7: 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 6: Conditions for the Rules

Conditions	Description
L2, L3, and L4 Filters	

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

When you select a condition without source or destination specified, then both egress and ingress traffic is selected for tapping the traffic. For example, if you select Ether Type as IPv4, TCP as the protocol, and do not specify IPv4 source or destination, then both egress and ingress traffic is selected for monitoring purpose.

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

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

To create a new map:

1. Select **OpenStack > Monitoring Session**.
2. Click **New**. The Monitoring Sessions page is displayed.
3. Create a new session. Refer to [Creating a New Session on page 58](#).
4. From **Maps**, drag and drop a new map template to the workspace.
5. Click on the map, then click details. Refer to [Figure 5: Map Details on page 64](#).




Figure 5: Map Details

The map rules quick view is displayed as shown in [Figure 6: Creating a New Map](#).




Figure 6: Creating a New Map

6. Enter the appropriate information for creating a new map as shown in [Table 7: Fields for Creating a New Map](#).

Table 7: Fields for Creating a New Map

Parameter	Description
Alias	The name of the new map. NOTE: The name can contain alphanumeric characters with no spaces.
Comments	The description of the map.

Parameter	Description
Rule Conditions Map Rules	<p>The rules for filtering the traffic in the map.</p> <p>To add a map rule:</p> <ol style="list-style-type: none"> Click Add a Rule. Select a condition from the Search L2 Conditions drop-down list and specify a value. Based on this selection, the Search L3 Conditions drop-down list is automatically updated. <div data-bbox="467 449 735 743" style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <div style="border: 1px solid #ccc; padding: 2px;">Search Layer 2 Conditions... ▾</div> <div style="background-color: #f0f0f0; padding: 2px;">Ether Type</div> <div style="padding: 2px;">MAC Source</div> <div style="padding: 2px;">MAC Destination</div> <div style="padding: 2px;">VLAN</div> <div style="padding: 2px;">VLAN PCP</div> <div style="padding: 2px;">VLAN TCI</div> <div style="padding: 2px;">Pass All</div> </div> Select a condition from the Search L3 Conditions drop-down list and specify a value. <div data-bbox="467 846 735 1140" style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <div style="border: 1px solid #ccc; padding: 2px;">Search Layer 3 Conditions... ▾</div> <div style="padding: 2px;">Protocol</div> <div style="padding: 2px;">IP Fragmentation</div> <div style="padding: 2px;">IP TTL</div> <div style="padding: 2px;">IP TOS</div> <div style="padding: 2px;">IP ECN</div> <div style="padding: 2px;">IP Source</div> <div style="background-color: #f0f0f0; padding: 2px;">IP Destination</div> </div> (Optional) If you have selected TCP or UDP as the protocol in the L3 conditions, then select Port Source or Port Destination from the Search L4 Conditions drop-down list and specify a value. If you have selected conditions other than TCP or UDP, then the Search L4 Conditions drop-down list is disabled. <div data-bbox="467 1339 735 1467" style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <div style="border: 1px solid #ccc; padding: 2px;">Search Layer 4 Conditions... ▾</div> <div style="background-color: #f0f0f0; padding: 2px;">Port Source</div> <div style="padding: 2px;">Port Destination</div> </div> (Optional) In the Priority and Action Set box, assign a priority and action set. (Optional) In the Rule Comment box, enter a comment for the rule. <p>NOTE: Repeat steps b through f to add more conditions.</p> <p>NOTE: Repeat steps a through f to add nested rules.</p>

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


- 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.

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

Adding Applications to the Monitoring Session

Gigamon supports the following GigaSMART applications:

- [Sampling on page 71](#)
- [Slicing on page 73](#)
- [Masking on page 74](#)
- [NetFlow on page 76](#)

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 7: Dragging the Sample Application

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




Figure 8: Selecting Details

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




Figure 9: 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 10: Dragging the Slice Application

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




Figure 11: Selecting Details

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




Figure 12: 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 13: Dragging the Mask Application

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




Figure 14: Selecting Details

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




Figure 15: 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 78](#). 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 80](#).

[Figure 16: NetFlow on GigaVUE V Series Node on page 77](#) shows an example of a NetFlow application created on a GigaVUE V Series node in the monitoring session.




Figure 16: NetFlow on GigaVUE V Series Node

The NetFlow record generation is performed on GigaVUE V Series node running the NetFlow application. In [Figure 16: NetFlow on GigaVUE V Series Node on page 77](#), 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 87](#).

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 8: Match/Key Elements

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

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

	Description	Supported NetFlow Versions
Transport		
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
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 9: Collect/Non-Key Elements

	Description	Supported NetFlow Versions
Counter		

	Description	Supported NetFlow Versions
Byte Count	Configures the number of octets since the previous report in incoming packets for the current flow as a non-key field.	v9 and IPFIX
Packet Count	Configures the number of incoming packets since the previous report for this flow as a non-key field.	v9 and IPFIX
Data Link		
Destination MAC	Configures the destination MAC address as a non-key field.	v9 and IPFIX
Egress Des MAC	Configures the post source MAC address as a non-key field.	IPFIX
Ingress Des MAC	Configures the IEEE 802 destination MAC address as a non-key field.	IPFIX
Source MAC	Configures the IEEE 802 source MAC address as a non-key field.	v9 and IPFIX
Timestamp		
Flow End Millisec	Configures the absolute timestamp of the last packet of current flow in milliseconds as a non-key field.	IPFIX
Flow End Sec	Configures the flow start SysUp time as a non-key field.	IPFIX
Flow End Time	Configures the flow end SysUp time as a non-key field.	v9 and IPFIX
Flow Start Millisec	Configures the value of the IP Precedence as a non-key field.	IPFIX
Flow Start Sec	Configures the absolute timestamp of the first packet of this flow as a non-key field.	IPFIX
Flow Startup Time	Configures the flow start SysUp time as a non-key field.	v9 and IPFIX
Flow		
Flow End Reason	Configures the reason for Flow termination as a non-key field.	IPFIX
IPv4		
ICMP Type Code	Configures the type and code of the IPv4 ICMP message as a non-key field.	v9 and IPFIX
IPv4 Dest IP	Configures the IPv4 destination address in the IP packet header as a non-key field.	v9 and IPFIX

	Description	Supported NetFlow Versions
IPv4 ICMP Code	Configures the code of the IPv4 ICMP message as a non-key field.	IPFIX
IPv4 ICMP Type	Configures the type of the IPv4 ICMP message as a non-key field.	IPFIX
IPv4 Options	Configures the IPv4 options in the packets of the current flow as a non-key field.	IPFIX
IPv4 Src IP	Configures the IPv6 source address in the IP packet header as a non-key field.	v9 and IPFIX
IPv4 Total Length	Configures the total length of the IPv4 packet as a non-key field.	IPFIX
Network		
IP CoS	Configures the IP Class Of Service (CoS) as a key field.	v9
IP Protocol	Configures the value of the protocol number in the IP packet header as a key field.	v9
IP Version	Configures the IP version field in the IP packet header as a key field.	v9
IPv6		
IPv6 Dest IP	Configures the IPv6 destination address in the IP packet header as a key field.	v9
IPv6 Flow Label	Configures the value of the IPv6 flow label field in the IP packet header as a key field.	v9
IPv6 Src IP	Configures the IPv6 source address in the IP packet header as a key field.	v9
Transport		
L4 Dest Port	Configures the destination port identifier in the transport header as a non-key field.	v9 and IPFIX
L4 Src Port	Configures the source port identifier in the transport header as a non-key field.	v9 and IPFIX
TCP Ack Number	Configures the acknowledgment number in the TCP header as a non-key field.	IPFIX
TCP Dest Port	Configures the destination port identifier in the TCP header as a non-key field.	IPFIX
TCP Flags	Configures the TCP control bits observed for the packets of this flow as a non-key field.	v9 and IPFIX
TCP Header Length	Configures the length of the TCP header as a non-key field.	IPFIX

	Description	Supported NetFlow Versions
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 17: Dragging the NetFlow Application

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




Figure 18: Selecting Details

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

 A screenshot of a configuration form titled 'Application'. The form has an orange header bar with the title and a 'Save' button. The form content includes:

- Application**: netflow
- Alias**: Netflow_V5
- State**: On Off
- NetFlow version**: v5 (dropdown menu)
- Active cache timeout**: 1800
- Inactive cache timeout**: 15

Figure 19: 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 90](#).

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 20: Dragging the NetFlow Application

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




Figure 21: 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 22: 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 78](#).
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 80](#).

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 90](#).

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 23: Adding NAT

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




Figure 24: Selecting Details

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

The screenshot shows a configuration form for a NAT device. At the top right, there is an orange bar with a 'Save' button. The form contains the following fields:

- Alias:** NAT-to-Splunk
- Local Subnet:** 192.168.50.0/24
- Routes:** A table with columns 'Destination IP' and 'Node Interface Subnet Cidr'. The first row has a red 'x' in the 'Destination IP' column with the value '10.0.2.189' and '10.0.1.0/24' in the 'Node Interface Subnet Cidr' column. A green '+' icon is visible to the left of the table header.

Figure 25: 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 26: 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 27: Linking NetFlow to NAT

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

NetFlow Examples

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

- [Example 1 on page 90](#)
- [Example 2 on page 94](#)

Example 1

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

- Create a monitoring session. For steps, refer to [Creating a Monitoring Session on page 58](#).




Figure 28: Creating a Monitoring Session

- In the monitoring session, create a Pass all map. A pass all map sends all the traffic received from the G-vTAP agents to the tunnel endpoint or NAT. For steps, refer to [Creating a Map on page 62](#).




Figure 29: 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 30: 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 31: Creating a Link from Pass-all Map to Tunnel_Endpoint

5. Drag and drop a v5 NetFlow application.




Figure 32: 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 83](#).




Figure 33: Configuring the NetFlow Application

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




Figure 34: 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 87](#).




Figure 35: 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 88](#).




Figure 36: Adding a Link from v5 NetFlow Application to NAT

10. 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 37: 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 58](#).




Figure 38: 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 62](#).




Figure 39: Creating a TCP Map

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




Figure 40: Creating a UDP Map

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




Figure 41: 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 85](#).




Figure 42: Adding a v9 NetFlow Application

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




Figure 43: 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 85](#).




Figure 44: Adding a IPFIX NetFlow Application

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




Figure 45: 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 87](#).




Figure 46: Adding a NAT

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




Figure 47: Adding a Link from NetFlow Application to NAT

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




Figure 48: Adding a Link from NetFlow Application to NAT

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




Figure 49: 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 71](#).

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

Figure 50: 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 50: 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 51: Connecting the Maps, Applications, and Monitoring Tools on page 100](#).

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 51: 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 51: 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					
<input type="checkbox"/>	Name	Connection	# of Targets	Status	Statistics
<input type="checkbox"/>	Connection1_Monitoring	Connection1	2	● Success	View
Total Items : 1					

When you click on the Status link, the Deployment Report is displayed. Refer to [Figure 52: Monitoring Session Deployment Report](#).

Monitoring Sessions		Deployment Report	
<input type="checkbox"/>	Name	Connection	
<input type="checkbox"/>	Example1_Monitor	Example1	
		Monitoring Session Alias :	Example1_Monitor
		Deployment Status :	Success
		Operation :	deploy
		Start Time :	2017-08-08 15:14:58
		End Time :	2017-08-08 15:14:58
		General Failure Messages :	NONE
		Selected Targets :	2
		Target Deployment Successes :	2
		Target Deployment Failures :	0
		NIC License Failures :	0
		V-Series Node Deployment Successes :	1
		V-Series Node Deployment Failures :	0
		Unselected Targets :	0
		Target Undeployment Successes :	0
		Target Undeployment Failures :	0
		V-Series Node Undeployment Successes :	0
		V-Series Node Undeployment Failures :	0

Figure 52: 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 53: Action Set with Multiple Links on page 103](#), 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 53: Action Set with Multiple Links

GigaVUE V Series node supports the following header transformations:

Table 10: Header Transformations

Option	Description
MAC Source	Modify the Ethernet source address.
MAC Destination	Modify the Ethernet destination address.
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 54: 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 55: 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					
<input type="checkbox"/>	Name	Connection	# of Targets	Status	Statistics
<input type="checkbox"/>	Connection1_Monitoring	Connection1	2	● Success	View
Total Items : 1					

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





Figure 56: 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 57: Viewing the Map Statistics](#).




Figure 57: Viewing the Map Statistics

Figure 57: 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.



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 58: 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 59: 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 60: Filtering in Topology View on page 109](#).




Figure 60: 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 60: Filtering in Topology View on page 109](#).





Figure 61: 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**.
1Select **Settings** to edit the OpenStack settings. The **Settings** page appears.
Refer to [Figure 62: Cloud > OpenStack > Connections > Settings](#).




Figure 62: Cloud > OpenStack > Connections > Settings

3. Click **Edit** to edit the Settings fields. Refer to [Table 11: OpenStack Settings on page 111](#) for descriptions of the Settings fields:

Table 11: OpenStack Settings

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

Compatibility Matrix

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

Refer to the following sections for details:

- [GigaVUE-FM Version Compatibility on page 112](#)
- [Supported Features in GigaVUE V Series Nodes on page 112](#)
- [Supported Features in G-vTAP Agents on page 113](#)

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
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 114](#).

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 Downloads](#)
- [Documentation Feedback](#)
- [Contact Technical Support](#)
- [Contact Sales](#)
- [The Gigamon Community](#)

Documentation Downloads

The following table provides a list of the additional documentation provided for GigaVUE H Series and TA Series nodes. "*" indicates new documents in this release. Click the "Get" link, below, to download the latest PDF for that release.



TIP: If you keep all PDFs for a particular release in common folder, you can easily search across the doc set by opening one of the files in Acrobat and choosing **Edit > Advanced Search** from the menu. This opens an interface that allows you to select a directory and search across all PDFs in a folder.

Table 1: Documentation Suite for Gigamon Products

Summary	Document	Downloads ↓	
		5.7	5.6
<ul style="list-style-type: none"> • complete doc set for the respective release, minus Release Notes, in a zip file 	All-Documents Zip	Get-All-5700	Get-All-5600
Release Notes - Read First			
<ul style="list-style-type: none"> • new features, resolved issues, and known issues in this release • important notes regarding installing and upgrading to this release <p>NOTE: In 5.7.00, the Release Notes documents combines GigaVUE-OS, GigaVUE-FM, and GigaVUE Cloud Suite into one document.</p>	GigaVUE-OS, GigaVUE-FM, GigaVUE-VM, and GigaVUE Cloud Suite Release Notes <p>NOTE: Registered Customers can download the Release Notes from the Software & Docs page on to My Gigamon. Refer to How to Download PDFs from My Gigamon.</p>	Get-All-RN	Get-FM Get-OS Get-VM Get-Cloud

Summary	Document	Downloads ↓	
		5.7	5.6

Hardware Installation Guides

<ul style="list-style-type: none"> how to unpack, assemble, rack-mount, connect, and initially configure the respective GigaVUE devices reference information and specifications for the respective GigaVUE devices 	GigaVUE-HC1 Hardware Installation Guide	Get	Get
	GigaVUE-HC2 Hardware Installation Guide	Get	Get
	GigaVUE-HC3 Hardware Installation Guide	Get	Get
	GigaVUE TA Series Hardware Installation Guide	Get	Get

Software Installation and Upgrade Guides

<ul style="list-style-type: none"> how to migrate GigaVUE-FM on VMware ESXi, Hardware Appliance, and AWS. 	*GigaVUE-FM Migration Guide		
<ul style="list-style-type: none"> how to install and upgrade GigaVUE-FM on VMware ESXi, MS Hyper-V, and KVM 	GigaVUE-FM Installation and Upgrade Guide	*Get	in FM User's Guide
<ul style="list-style-type: none"> how to upgrade the embedded GigaVUE-OS on GigaVUE H Series and GigaVUE TA Series nodes 	GigaVUE-OS Upgrade Guide	Get	Get
Administration Guide			
<ul style="list-style-type: none"> how to administer the GigaVUE-OS and GigaVUE-FM software 	GigaVUE-OS and GigaVUE-FM Administration Guide	Get	Get

Configuration and Monitoring Guides

<ul style="list-style-type: none"> how to install, deploy, and operate GigaVUE-FM how to configure GigaSMART operations 	GigaVUE-FM User's Guide	Get	Get
<ul style="list-style-type: none"> how to deploy the GigaVUE Cloud Suite solution in any cloud platform 	GigaVUE Cloud Suite for AnyCloud Configuration Guide	*Get	

Summary	Document	Downloads ↓	
		5.7	5.6
<ul style="list-style-type: none"> how to configure the GigaVUE Cloud Suite components and set up traffic monitoring sessions for the respective cloud platform 	GigaVUE Cloud Suite for AWS Configuration Guide	Get	Get
	GigaVUE Cloud Suite for AWS QuickStart Guide		
	*GigaVUE Cloud Suite for AWS Secret Regions Configuration Guide		
	GigaVUE Cloud Suite for Azure Configuration Guide	Get	Get
	GigaVUE Cloud Suite for Kubernetes Configuration Guide	*Get	
	*GigaVUE Cloud Suite for Nutanix Configuration Guide		
	GigaVUE Cloud Suite for OpenStack Configuration Guide	Get	Get
	GigaVUE Cloud Suite for VMware Configuration Guide	Get	Get

Reference Guides

<ul style="list-style-type: none"> library of GigaVUE-OS CLI (Command Line Interface) commands used to configure and operate GigaVUE H Series and TA Series devices 	GigaVUE-OS CLI Reference Guide	Get	Get
<ul style="list-style-type: none"> guidelines for the different types of cables used to connect Gigamon devices 	GigaVUE-OS Cabling Quick Reference Guide	*Get	
<ul style="list-style-type: none"> compatibility information and interoperability requirements for Gigamon devices 	GigaVUE-OS Compatibility and Interoperability Matrix	*Get	
<ul style="list-style-type: none"> samples uses of the GigaVUE-FM Application Program Interfaces (APIs) <p>NOTE: Content will be merged into the GigaVUE-FM User's Guide in a future release.</p>	GigaVUE-FM REST API Getting Started Guide	Get	Get

In-Product Help

Summary	Document	Downloads ↓	
		5.7	5.6
<ul style="list-style-type: none"> how to install, deploy, and operate GigaVUE-FM. Provided from the GigaVUE-FM interface. 	GigaVUE-FM Online Help	Click the  icon in GigaVUE-FM	
<ul style="list-style-type: none"> the web-based GUI for the GigaVUE-OS. Provided from the GigaVUE-OS H-VUE interface. 	GigaVUE-OS H-VUE Online Help	Click the  icon in GigaVUE-OS H-VUE	

NOTE: Registered customers can log in to [My Gigamon](#) to download documentation for specific releases under Software & Documentation Downloads. Refer to [How to Download PDFs from My Gigamon](#).

How to Download PDFs from My Gigamon

To download release-specific PDFs:

1. Log in to [My Gigamon](#)
2. Click on the **Software & Documentation** link.
3. Use the **Product** and **Release** filters to find documentation for the current release. For example, select Product: "GigaVUE-FM" and Release: "5.7," enter "pdf" in the search box, and then click **GO** to view all PDF documentation for GigaVUE-FM 5.7.xx.

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>

Contact Technical Support

See <https://www.gigamon.com/support-and-services/contact-support> for Technical Support hours and contact information. You can also email Technical Support at support@gigamon.com.

Contact Sales

Use the following information to Gigamon channel partner or Gigamon sales representatives.

Telephone: +1.408.831.4025

Sales: inside.sales@gigamon.com

Partners: www.gigamon.com/partners.html

Premium Support

Email Gigamon at inside.sales@gigamon.com for information on purchasing 24x7 Premium Support. Premium Support entitles you to round-the-clock phone support with a dedicated Support Engineer every day of the week.

The Gigamon Community

The [Gigamon Community](#) is a technical site where Gigamon users, partners, security and network professionals and Gigamon employees come together to share knowledge and expertise, ask questions, build their network and learn about best practices for Gigamon products.

Visit the Gigamon Community site to:

- Find knowledge base articles and documentation
- Ask and answer questions and learn best practices from other members.
- Join special-interest groups to have focused collaboration around a technology, use-case, vertical market or beta release
- Take online learning lessons and tutorials to broaden your knowledge of Gigamon products.
- Submit and vote on feature enhancements and share product feedback. (Customers only)
- Open support tickets (Customers only)
- Download the latest product updates and documentation (Customers only)

The Gigamon Community is a great way to get answers fast, learn from experts and collaborate directly with other members around your areas of interest.

Register today at community.gigamon.com

Questions? Contact our Community team at community.gigamon.com