



GigaVUE Cloud Suite for AnyCloud 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 AnyCloud	2
Audience	2
License Information	2
Bring Your Own License (BYOL)	2
Overview of GigaVUE Cloud Suite for AnyCloud	3
About GigaVUE Cloud Suite for AnyCloud	3
Overview of Components	4
Connect Components	6
Obtain Images	6
Launch GigaVUE-FM	6
G-vTAP Agents	6
Configure Monitoring Sessions	16
Overview of Visibility Components	16
Create Tunnel Endpoints	18
Create Monitoring Session	19
Configure AnyCloud Settings	51
Additional Sources of Information	52
Documentation	52
How to Download PDFs from My Gigamon	54
Documentation Feedback	55
Contact Technical Support	55
Contact Sales	55
Premium Support	55
The Gigamon Community	55

GigaVUE Cloud Suite for AnyCloud

This guide describes how to deploy the GigaVUE Cloud solution in any of the cloud platforms available in the market.

Topics:

Audience

This guide is intended for the users who want to deploy the GigaVUE Cloud solution in any of the cloud platforms such as Google Cloud, Nutanix, and others.

License Information

The GigaVUE Cloud Suite for AnyCloud supports Bring Your Own License (BYOL) model.

Bring Your Own License (BYOL)

The BYOL option can be purchased based on the number of TAP points and the term of the license. Gigamon offers the following options for purchasing the license:

- Traffic visibility for up to 100 virtual TAP points (VMs)
- Traffic visibility for up to 1000 virtual TAP points (VMs)

NOTE: Make sure you purchase a licensing option that can provide traffic visibility to all the TAP points in your VM. If the licensing option cannot support all the TAP points, then the VMs are selected randomly for monitoring the traffic.

The minimum term for the license is 3 months. For purchasing licenses with the BYOL option, contact our Gigamon Sales. Refer to [Contacting Sales](#). To generate and apply license, refer to the "Licensing" section in the *GigaVUE-OS and GigaVUE-FM Administration Guide*.

Overview of GigaVUE Cloud Suite for AnyCloud

This chapter introduces the components of the GigaVUE Cloud Suite solution and the supported architecture for deploying the cloud solution in any of the available cloud platforms. Refer to the following sections for details:

- [About GigaVUE Cloud Suite for AnyCloud](#)
- [Overview of Components](#)

This guide provides instructions for connecting GigaVUE-FM in any of the cloud platforms available in the market. For information about installing GigaVUE-FM in your enterprise data center, refer to the *GigaVUE-FM Installation and Upgrade Guide* available in the Customer Portal.

About GigaVUE Cloud Suite for AnyCloud

The GigaVUE Cloud Suite solution for the following platforms are mature solutions with complete support for configuration and upgrade from GigaVUE-FM:

- GigaVUE Cloud Suite for AWS
- GigaVUE Cloud Suite for Azure
- GigaVUE Cloud Suite for OpenStack

For the other cloud platforms (public or private) available in the market, the GigaVUE Cloud Suite for AnyCloud deployment option provides traffic visibility.

The GigaVUE Cloud Suite for AnyCloud option consists of the following components:

- GigaVUE® Fabric Manager (GigaVUE-FM)
- G-vTAP Agents
- G-vTAP Controllers
- GigaVUE V Series Controllers
- GigaVUE V Series Nodes

GigaVUE-FM is a key component of the GigaVUE Cloud solution. GigaVUE® Fabric Manager (GigaVUE-FM) is a web-based fabric management interface that provides a single-pane-of-glass visibility and management of both the physical and virtual traffic.

In the Any Cloud deployment option, you are responsible for the following:

- Installing and launching GigaVUE-FM from the supported cloud or enterprise platform.
- Launching the fabric components in your platform.
- Sharing the IP addresses and subnet CIDR of the fabric components with GigaVUE-FM.

The images of the components are available in the customer portal.

NOTE: Contact Gigamon Technical Support team if the existing Gigamon images for a specific cloud platform is not compatible.

GigaVUE-FM uses the IP addresses of the fabric components to:

- Identify the traffic
- Monitor the traffic flow
- Forward the traffic to the destination

NOTE: You are responsible for deleting the fabric nodes from the platform when visibility for the platform is no longer required.

Overview of Components

The following table provides a brief description of the components in the AnyCloud deployment option:

Component	Description
GigaVUE® Fabric Manager (GigaVUE-FM)	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. You are responsible for launching GigaVUE-FM from your end on the supported cloud or enterprise platforms.
G-vTAP agent	G-vTAP agent is an agent that is deployed in the target VM. This agent mirrors the selected traffic from the VMs to the GigaVUE® V Series node. The G-vTAP agent is offered as a Debian (.deb), Redhat Package Manager (.rpm) or windows package. Refer to Install G-vTAP Agents .
G-vTAP Controller	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.

Component	Description
GigaVUE® V Series Controller	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	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, provided the cloud platform supports

This guide provides instructions on how to use GigaVUE-FM for configuration in any cloud platform available in the market.

High-Level Architecture

The following diagram shows a high-level architecture of the GigaVUE Any Cloud configuration:

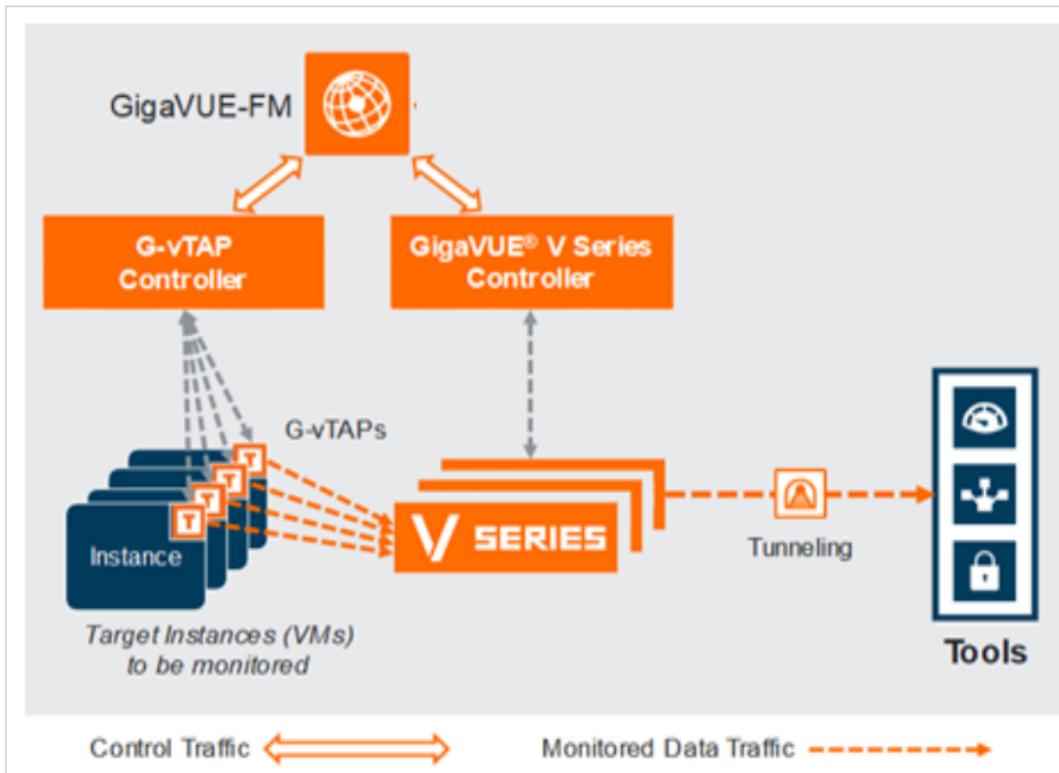


Figure 1: High Level Architecture

Connect Components

This chapter describes how to connect to the GigaVUE Cloud Suite for the AnyCloud solution.

Refer to the following sections for details:

- [Obtain Images](#)
- [Launch GigaVUE-FM](#)
- [G-vTAP Agents](#)
- [Connect to Cloud Platform](#)

Obtain Images

You must have obtained the images for the components of the GigaVUE Cloud Suite for AnyCloud.

NOTE: You are responsible for downloading the image and launching it in your environment. The components can be deployed using the compatible disk formats in respective cloud platform.

The number of GigaVUE V Series node required depends on the number of G-vTAP agents in the environment. The number of G-vTAP agent instances per GigaVUE V Series nodes is based on the metrics available in the Settings tab.

Launch GigaVUE-FM

The GigaVUE-FM software package is available in multiple formats such as OVA, QCOW2, ISO. Use the appropriate media format to deploy GigaVUE-FM.

After you deploy GigaVUE-FM you must perform an initial configuration before you start using GigaVUE-FM. Refer to the *GigaVUE-FM User's Guide* for details.

G-vTAP Agents

A **G-vTAP agent** is an agent that is deployed in the VMs. This agent mirrors the selected traffic from the VMs (virtual machines), encapsulates it using GRE or VXLAN tunneling, and forwards it to the GigaVUE® V Series node.

A G-vTAP agent consists of a source interface and a destination interface. The network packets collected from the source interface are sent to the destination interface. From the destination interface, the packets traverse through the L2 GRE or VXLAN tunnel interface to the GigaVUE V Series node.

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

Linux Agent Installation

Refer to the following sections for the Linux agent installation:

- [Single NIC Configuration](#)
- [Install G-vTAP Agents](#)
- [Install G-vTAP from Debian Package](#)
- [Install G-vTAP from RPM package](#)

Single NIC Configuration

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

For example, assume that there is only one interface eth0 in the monitoring VM. 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 NIC as the source and the destination interface can sometimes cause increased latency in sending the traffic out from the VM.

Example of the G-vTAP config file for a single NIC configuration:

Example—Grant permission to monitor ingress and egress traffic at iface

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

NOTE: The GigaVUE Cloud Suite for AnyCloud supports only single NIC G-vTAP agent configuration.

Install G-vTAP Agents

You must have sudo/root access to edit the G-vTAP agent configuration file.

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

- [Install G-vTAP from Debian Package](#)
- [Install G-vTAP from RPM package](#)
- [Windows Agent Installation](#)

Install G-vTAP from Debian Package

To install from a Debian package:

1. [Download the G-vTAP Agent Debian \(.deb\) package.](#)
2. Copy this package to your VM. Install the package with root privileges, for example:

```
ubuntu@ip-10-0-0-246:~$ ls gvtap-agent_1.7-1_amd64.deb
ubuntu@ip-10-0-0-246:~$ sudo dpkg -i gvtap-agent_1.7-1_amd64.deb
```

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.

[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
```

4. Save the file.
5. Reboot the VM.

The G-vTAP agent status will be displayed as running. Check the status using the following command:

```
ubuntu@ip-10-0-0-246:~$ sudo /etc/init.d/gvtap-agent status
G-vTAP Agent is running
```

Install G-vTAP from 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.](#)
2. Copy this package to your VM. Install the package with root privileges, for example:

```
[VM-user@ip-10-0-0-214 ~]$ ls
gvtap-agent_1.7-1_x86_64.rpm
[VM-user@ip-10-0-0-214 ~]$ sudo rpm -i

gvtap-agent_1.7-1_x86_64.rpm
```

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

[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
```

4. Save the file.
5. Reboot the VM.

Check the status with the following command:

```
[VM-user@ip-10-0-0-214 ~]$ sudo /etc/init.d/gvtap-agent status
G-vTAP Agent is running
```

Windows Agent Installation

To install the Windows agent:

1. [Download the Windows agent package.](#)
2. Extract the contents of the .zip file into a convenient location.
3. Right-click 'WinPcap_4_1_3.exe' (located in the 'winpcap' folder) and select and select **Run as Administrator**.
4. Right-click 'install.bat' and select **Run as Administrator**.
5. If you want to start the Windows G-vTAP agent, you may do one of the following:
 - Reboot the VM.
 - Run 'sc start gvtap' from the command prompt.
 - Start the G-vTAP Agent from the Task Manager.

NOTE: You may need to edit the Windows Firewall settings to grant access to the gvtap process. To do this, access the Windows Firewall settings and find "gvtapd" in the list of apps and features. Select it to grant access. Be sure to select both Private and Public check boxes. If "gvtapd" does not appear in the list, click **Add another app...** Browse your program

files for the gvtap-agent application (gvtapd.exe) and then click **Add**. (Disclaimer: These are general guidelines for changing Windows Firewall settings. See Microsoft Windows help for official instructions on Windows functionality.)

Install IPsec on G-vTAP Agent

If IPsec is used to establish secure connection between G-vTAP agents and GigaVUE V Series nodes, then you must install IPsec on G-vTAP agent instances. To install IPsec on G-vTAP agent you need the following files:

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

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

- "Install from Ubuntu/Debian Package" below
- "Install from Red Hat Enterprise Linux and Centos" on the next page
- "Install from Red Hat Enterprise Linux and Centos with Selinux Enabled" on the next page

Install from Ubuntu/Debian Package

1. Launch the G-vTAP agent in your environment.
2. Copy the G-vTAP package files and strongSwan TAR file to the G-vTAP agent:
 - [strongswan5.3.5-1ubuntu3.8_amd64-deb.tar.gz](#)
 - [gvtap-agent_1.7-1_amd64.deb](#)
 - [gvtap-ipsec_1.7-1_amd64.deb](#)
3. Install the G-vTAP agent package file:

```
sudo dpkg -i gvtap-agent_1.6-1_amd64.deb
```
4. Edit gvtap-agent.conf file to configure the required interface as source/destination for mirror:

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

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

6. Install IPsec package:

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

Install from Red Hat Enterprise Linux and Centos

1. Launch RHEL/Centos agent image.
2. Copy the following package files and strongSwan TAR files to the G-vTAP agent:

- [strongswan-5.7.1-1.el7.x86_64.tar.gz](#) for rhel7/centos7
- [strongswan-5.4.0-2.el6.x86_64.tar.gz](#) for rhel6/centos6
- [gvtap-agent_1.7-1_x86_64.rpm](#)
- [gvtap-ipsec_1.7-1_x86_64.rpm](#)

3. Install G-vTAP agent package:

```
sudo rpm -ivh gvtap-agent_1.7-1_x86_64.rpm
```

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

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

5. Install strongSwan:

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

6. Install IPsec package:

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

NOTE: You must install IPsec package after installing StrongSwan.

Install from Red Hat Enterprise Linux and Centos with Selinux Enabled

1. Launch the RHEL/Centos agent image.
2. Copy package files and strongSwan TAR file to G-vTAP agent.
 - [strongswan-5.7.1-1.el7.x86_64.tar.gz](#) for rhel7/centos7
 - [strongswan-5.4.0-2.el6.x86_64.tar.gz](#) for rhel6/centos6
 - [gvtap-agent_1.7-1_x86_64.rpm](#)
 - [gvtap-ipsec_1.7-1_x86_64.rpm](#)
 - gvtap.te and gvtap_ipsec.te files (type enforcement files) .

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

```
semodule_package -o gvtap.pp -m gvtap.mod  
sudo semodule -i gvtap.pp
```

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

```
semodule_package -o gvtap_ipsec.pp -m gvtap_ipsec.mod  
sudo semodule -i gvtap_ipsec.pp
```

5. Install G-vTAP agent package:

```
sudo rpm -ivh gvtap-agent_1.7-1_x86_64.rpm
```

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

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

7. Install strongSwan:

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

8. Install IPsec package:

```
sudo rpm -i gvtap-ipsec_1.7-1_x86_64.rpm
```

Connect to Cloud Platform

Prior to connecting to the cloud platform, you must ensure the following:

- GigaVUE-FM must have been launched
- G-vTAP Agent, G-vTAP Controller, GigaVUE V Series Node and GigaVUE V Series Controller must be running and accessible from GigaVUE-FM

To connect your platform using GigaVUE-FM:

1. Click **Cloud** in the top navigation bar.
2. Under **AnyCloud**, select **Monitoring Domain**, and click **New**. The AnyCloud Connection Configuration page appears.

3. Select or enter appropriate information as shown in the following table:

Field	Description
Monitoring Domain	An alias used to identify the monitoring domain. A monitoring domain consists of set of connections.
Connection Alias	An alias used to identify the connection.
G-vTAP Agent	
Subnet CIDRs	The CIDR range for the G-vTAP controller to scan the range of G-vTAP agents.
IP Ranges (optional)	IP address range to scan with in the CIDR specified. This is optional.
Tunnel Type	Tunnel type for carrying the mirrored traffic from the G-vTAP agents to the GigaVUE V Series nodes/physical nodes. The tunnel type can be: <ul style="list-style-type: none"> L2GRE VXLAN
Tunnel MTU	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. For GRE, the default value is 9001. For VXLAN, the default value is 8951. However, the G-vTAP agent tunnel MTU should be 50 bytes less than the agent's destination interface MTU size. If Secure Mirror Traffic option is enabled, then to minimize fragmentation you must configuring MTU value for G-vTAP agent as follows: With agent tunnel type L2GRE: <ul style="list-style-type: none"> If secure tunnel is enabled, MTU must be set as $(9001-42-53) = 8906$. If secure tunnel is not enabled, MTU must be set as $(9001-42) = 8959$ With agent tunnel type VXLAN <ul style="list-style-type: none"> If secure tunnel is enabled, MTU must be set as $(9001-50-53) = 8898$ If secure tunnel is not enabled, MTU must be set as must be set as 8951.
Secure Mirror Traffic	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)
G-vTAP Controller	
IP Addresses	IP address of the G-vTAP controller
V Series Controller	
IP Addresses	IP address of the GigaVUE V series controller
V Series Node	
IP Addresses	IP address of the GigaVUE V series node

Field	Description
Tunnel MTU	The Maximum Transmission Unit (MTU) is the maximum size of each packet that the tunnel endpoint can carry from the GigaVUE V Series node to the destination.
Gateway IP (optional)	IP address of the subnet. This field is optional.
Tool Subnet CIDR	Subnet CIDR of the tools
Additional Subnet CIDRs (optional)	Additional subnet CIDR

4. Click **Save**. The new monitoring domain is created and added to the Monitoring Domain list view page.

Configure Monitoring Sessions

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

Refer to the following sections for details:

- [Overview of Visibility Components](#)
- [Create Tunnel Endpoints](#)
- [Create Monitoring Session](#)
- [Configure AnyCloud Settings](#)

Overview of Visibility Components

The GigaVUE V Series node aggregates the traffic from multiple G-vTAP agents 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.

The following table lists the components of the monitoring session:

Parameter	Description
Map	A map (M) is used to filter the traffic flowing through the GigaVUE V Series node. It is a collection of one or more rules (R). The traffic passing through a map can match one or more rules defined in the map.
Rule	<p>A rule (R) contains specific filtering criteria that the packets must match.</p> <p>The filtering criteria lets you determine the target instances and the (egress or ingress) direction of tapping the network traffic.</p> <p>The rules must contain the appropriate Layer 2 (L2) to Layer 4 (L4) filters defined in them. For example, if you want to filter the traffic for HTTP Port 80, you must select the following criteria:</p> <ul style="list-style-type: none">• Layer 2—Ethertype IPv4 or IPv6• Layer 3—Protocol TCP• Layer 4—Port Destination 80 <p>By default, a rule always displays conditions based on the attributes of L2.</p> <p>A rule is also associated with priority and action set.</p>

Parameter	Description
-----------	-------------

Priority

A priority determines the order in which the rules are executed. The greater the value, the higher the priority.
 The priority value can range from 0 to 99.

Action Set

An action set is an exit point in a map that you can drag and create links to the other maps, applications, and the monitoring tools. A single map can have multiple action sets. A single action set can have multiple links connecting to maps and applications.

In the following example, the packets that match the rules in Action Set 0 are forwarded to a tunnel endpoint. The packets that match the rules in Action Set 1 are forwarded to another map.

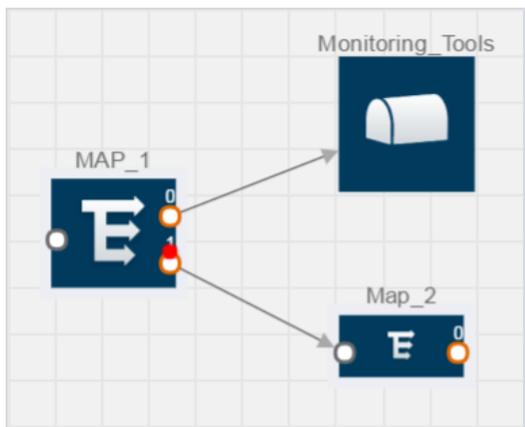


Figure 1: Action Set

A single action set can have up to 8 links connecting the same destination point. The same packets from the map are replicated in 8 different links.

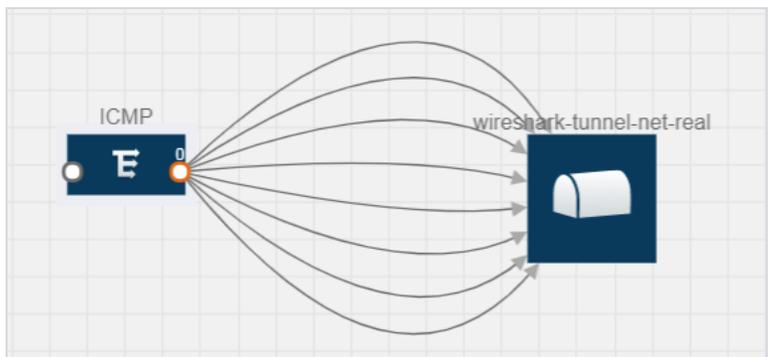


Figure 2: Action Set with Multiple Links

Parameter	Description
Link	<p>A link directs the packets to flow from a map to the destination. The destination could be the other maps, applications, and the monitoring tools. In Figure 2: Action Set with Multiple Links, 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 "Add Header Transformations" on page 46.</p>
Group	A group is a collection of maps that are pre-defined and saved in the map library for reuse.
Application	An application performs operations such as sampling, slicing, and masking on the traffic.
Inclusion Map	An inclusion map determines the VMs to be included for monitoring. This map is used only for target selection.
Exclusion Map	An exclusion map determines the VMs to be excluded from monitoring. This map is used only for target selection.
Target	<p>A target determines the instances that are to be monitored.</p> <p>Targets are determined based on the following formula:</p>
Automatic Target Selection (ATS)	<p>A built-in feature that automatically selects the Virtual Machines based on the rules defined in the maps, inclusion maps, and exclusion maps in the monitoring session.</p> <p>For example, if you create a rule determining the MAC source address in a map and a subnet in the inclusion map, the egress traffic from all VMs matching the MAC address in the specified subnet is selected for tapping the traffic.</p>
Tunnel	A tunnel lists the monitoring tools to which the traffic matching the filtered criteria is routed.

Create Tunnel Endpoints

The customized traffic from the GigaVUE V Series node is distributed to the tunnel endpoints using a standard L2 Generic Routing Encapsulation (GRE) or Virtual Extensible LAN (VXLAN) tunnel.

To create the tunnel endpoints:

1. In GigaVUE-FM, on the top navigation bar, select **Cloud**.
2. On the left navigation pane, select **AnyCloud > Settings > Tunnel Spec Library**.
3. Click **New**. The Add Tunnel page appears.

4. Select or enter the appropriate information as described in the following table.

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 tunnel port.
Traffic Direction	The direction of the traffic flowing through the GigaVUE V Series node. Choose Out for creating a tunnel from the GigaVUE 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 GigaVUE V Series node to the same IP address.
Remote Tunnel Port	Port number for the tunnel end point.

5. Click **Save**. The tunnel endpoints are added successfully.

Create Monitoring Session

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

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

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

- [Create New Session](#)
- [Clone Monitoring Session](#)
- [Create Map](#)
- [Add Applications to Monitoring Session](#)
- [Deploy Monitoring Session](#)

- [Add Header Transformations](#)
- [View Statistics](#)
- [View Topology](#)

Create New Session

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

To create a new session:

1. Select **AnyCloud > Monitoring Session**.
2. Click **New**. The Monitoring Sessions page is displayed.
3. Enter the appropriate information in the Monitoring Session Info as described in the following table.

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

4. Click **Create**.

Clone 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 described in the following table.

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. Click **Edit** to add the connections to the cloned monitoring session.

Create Map

Each map can have up to 32 rules associated with it. The following table lists the various rule conditions that you can select for creating a map, inclusion map, and exclusion map.

Conditions	Description
L2, L3, and L4 Filters	
Ether Type	<p>The packets are filtered based on the selected ethertype. The following conditions are displayed:</p> <ul style="list-style-type: none"> • IPv4 • IPv6 • ARP • RARP • Other <p>L3 Filters</p> <p>If you choose IPv4 or IPv6, the following L3 filter conditions are displayed:</p> <ul style="list-style-type: none"> • Protocol • IP Fragmentation • IP Time to live (TTL) • IP Type of Service (TOS) • IP Explicit Congestion Notification (ECN) • IP Source • IP Destination <p>L4 Filters</p> <p>If you select TCP or UDP protocol, the following L4 filter conditions are displayed:</p> <ul style="list-style-type: none"> • Port Source • Port Destination
MAC Source	The egress traffic from the VMs matching the specified source MAC address is selected.
MAC Destination	The ingress traffic from the instances or VMs 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.

Conditions	Description
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 monitoring the traffic. For example, if you select IPv4 as the Ether Type, TCP as the protocol, and do not specify IP 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. For example, if only IP source is selected as shown in the figure [Creating a Map for Tapping Egress Traffic](#), then the egress traffic from the instances in the subnet 10.0.1.0/24 is selected for monitoring the traffic.

East-zone-1737
Save Add to Library

Alias

Comments

Map Rules

× Rule 1

Priority ActionSet

Rule Comment

Ether Type ×

Value

Protocol ×

Value

IP Source ×

or

Figure 3: Creating a Map for Tapping Egress Traffic

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 **AnyCloud > Monitoring Session**.
2. Click **New**. The Monitoring Sessions page is displayed.
3. Create a new session. Refer to "[Create New Session](#)" on page 20.
4. From **Maps**, drag and drop a new map template to the workspace. If you are creating an exclusion or inclusion map, drag and drop a new map template to their respective section at the bottom of the workspace. The new map page is displayed.

5. Enter the appropriate information for creating a new map as described in the following table.

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.
Map Rules	The rules for filtering the traffic in the map. To add a map rule: <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. Select a condition from the Search L3 Conditions drop-down list and specify a value. (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.
Map Rules	<ol style="list-style-type: none"> (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. NOTE: Repeat steps b through f to add more conditions. NOTE: Repeat steps a through f to add nested rules.

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

6. To reuse the map, click **Add to Library**. Save the map using one of the following options:
- 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 present in the VNet.

7. Click **Save**.

To edit or delete a map, click a map and select **Details** to edit the map or **Delete** to delete the map as shown in the following figure.

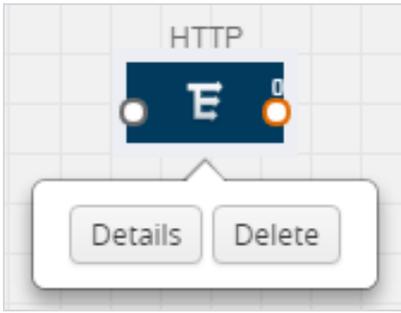


Figure 4: Editing or Deleting a Map

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

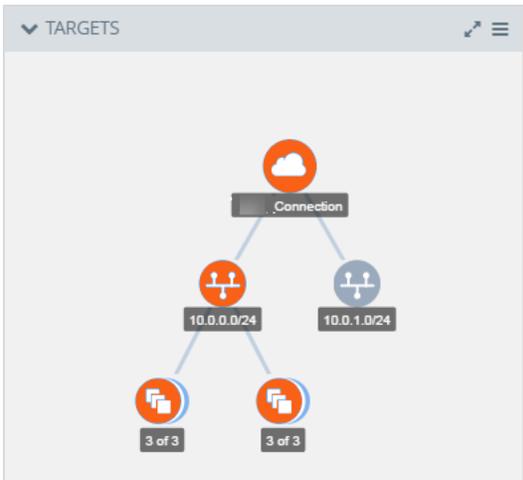


Figure 5: Viewing the Topology

Click  to expand the **Targets** dialog box. Click  to change the view from topology to viewing the target VM names. To view more details about the instance tag name, direction of tapping, and so on, click the arrow next to the instance name.

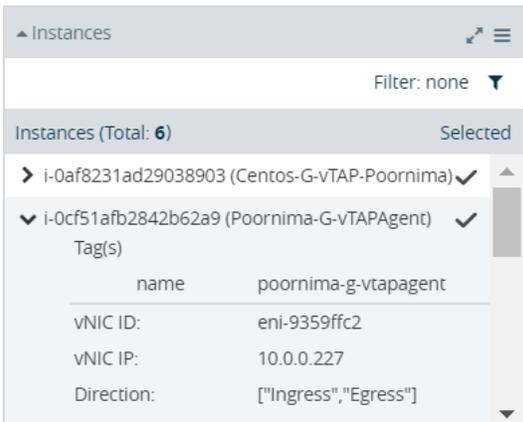


Figure 6: Viewing the instance Details

Filter the instances based on the Instance Name Prefix, IP address, or the 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.

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 > AnyCloud > Monitoring Session**
2. Open a monitoring session by doing one of the following:
3. Click **New** to create a new session.
4. Click the check box next to a session and then click **Edit** to edit an existing session.
5. Select or deselect the **Agent Pre-filtering** check box in the Monitoring Session info box to change the setting. It is enabled by default.
6. Click **OK**.
7. The Monitoring Session view displays the setting in the Agent Pre-filtering column.

Add Applications to Monitoring Session

Gigamon supports the following GigaSMART applications with the GigaVUE Cloud Suite for AnyCloud platform:

- ["Sampling" below](#)
- ["Slicing" on page 29](#)
- ["Masking" on page 30](#)
- ["NetFlow" on page 31](#)

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.

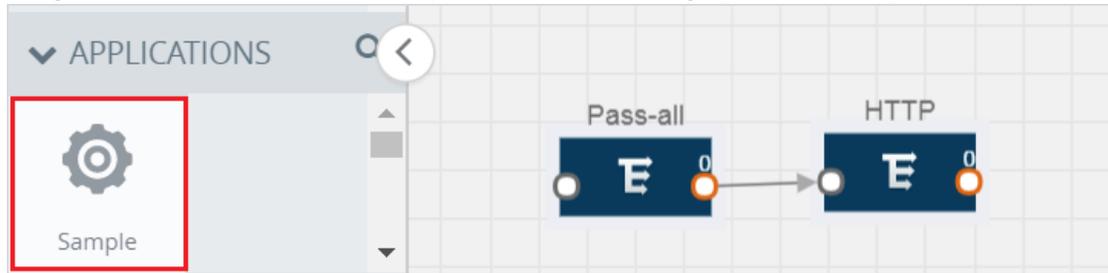


Figure 7: Dragging the Sample Application

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

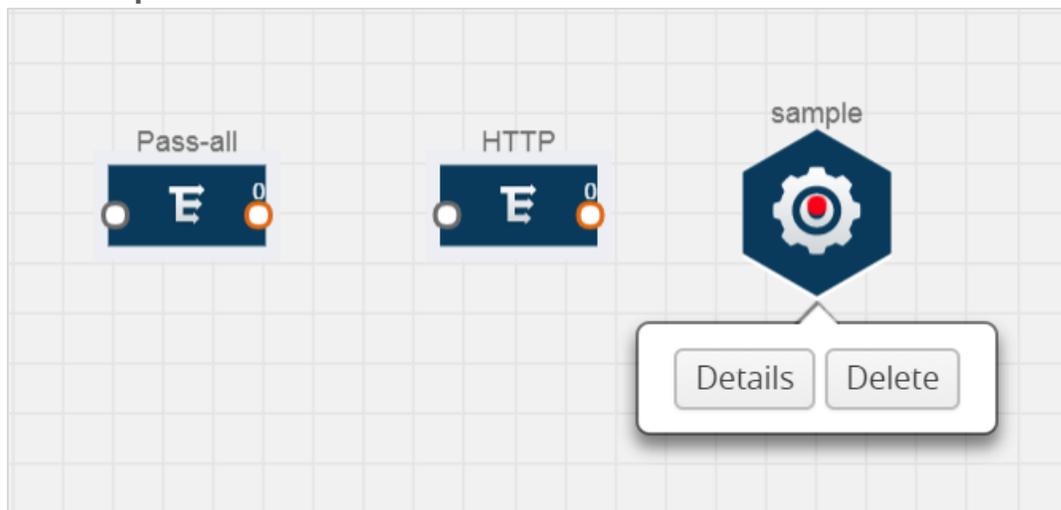


Figure 8: Selecting Details

3. In the **Alias** field, enter a name for the sample.
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 9: Dragging the Slice Application

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

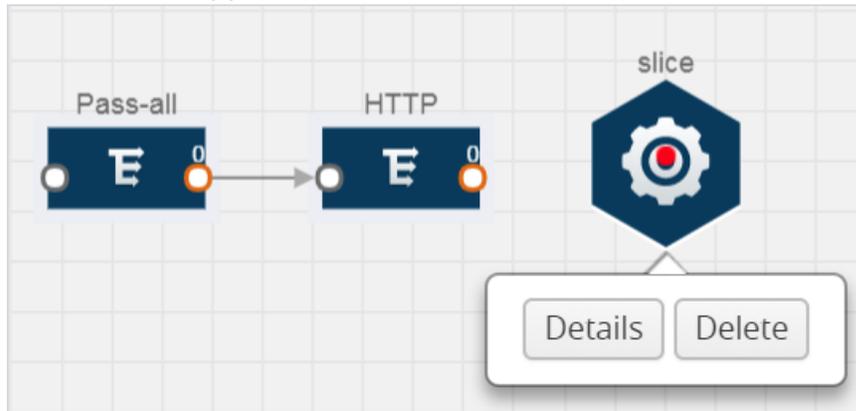


Figure 10: Selecting Details

3. In the **Alias** field, enter a name for the slice.
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 11: Dragging the Mask Application

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

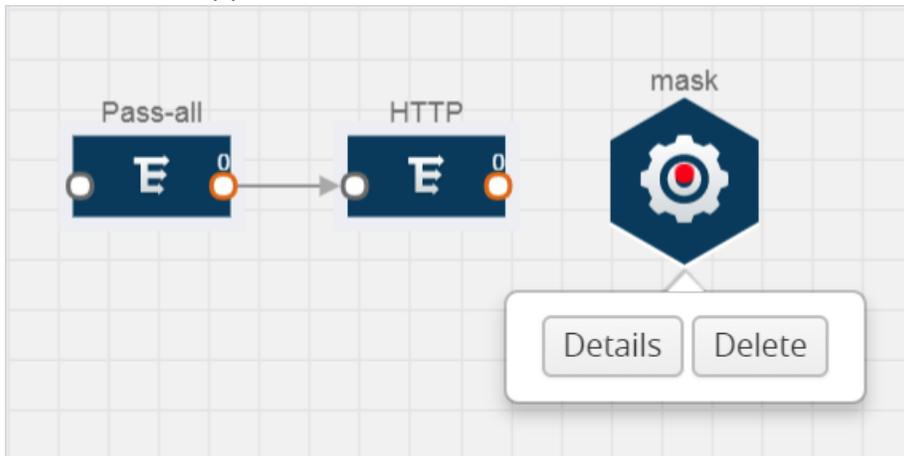


Figure 12: Selecting Details

3. In the **Alias** field, enter a name for the mask.
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 to your cloud environment.
- 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 VNets.

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 33](#). 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 35](#).

The following figure shows an example of a NetFlow application created on a GigaVUE V Series node in the monitoring session.

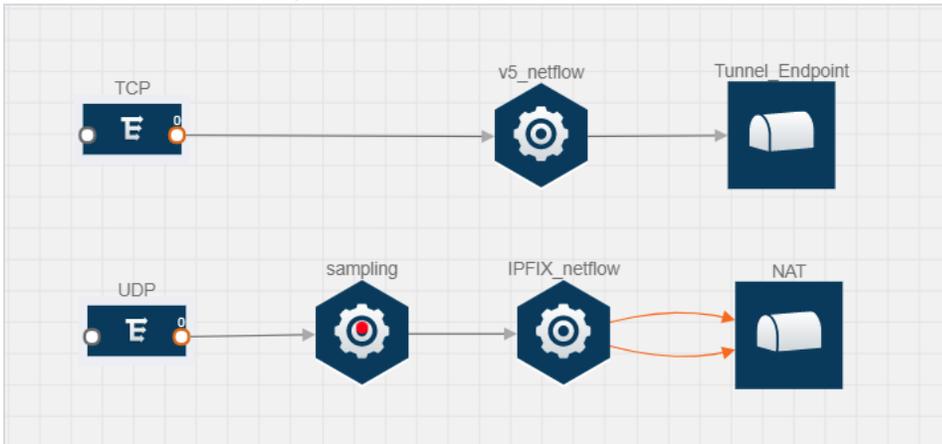


Figure 13: NetFlow on GigaVUE V Series Node

The NetFlow record generation is performed on GigaVUE V Series node running the NetFlow application. In [Figure 13: NetFlow on GigaVUE V Series Node](#), 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 for full packet inspection 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 41.

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 1: Match/Key Elements

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

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

Element	Description	Supported NetFlow Versions
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 2: Collect/Non-Key Elements

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

Element	Description	Supported NetFlow Versions
ICMP Type Code	Configures the type and code of the IPv4 ICMP message as a non-key field.	v9 and IPFIX
IPv4 Dest IP	Configures the IPv4 destination address in the IP packet header as a non-key field.	v9 and IPFIX
IPv4 ICMP Code	Configures the code of the IPv4 ICMP message as a non-key field.	IPFIX
IPv4 ICMP Type	Configures the type of the IPv4 ICMP message as a non-key field.	IPFIX
IPv4 Options	Configures the IPv4 options in the packets of the current flow as a non-key field.	IPFIX
IPv4 Src IP	Configures the IPv6 source address in the IP packet header as a non-key field.	v9 and IPFIX
IPv4 Total Length	Configures the total length of the IPv4 packet as a non-key field.	IPFIX
Network		
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

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

Add Version 5 NetFlow Application

To add a version 5 NetFlow application:

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

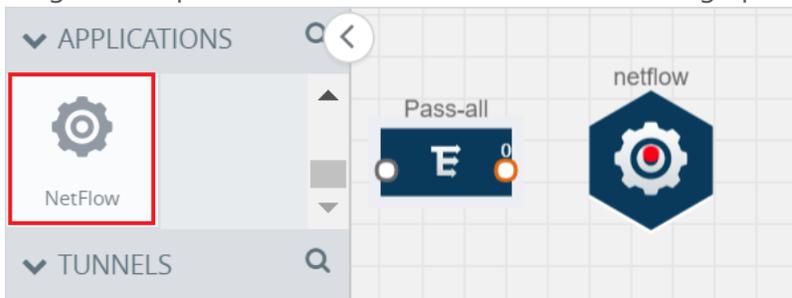


Figure 14: Dragging the NetFlow Application

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

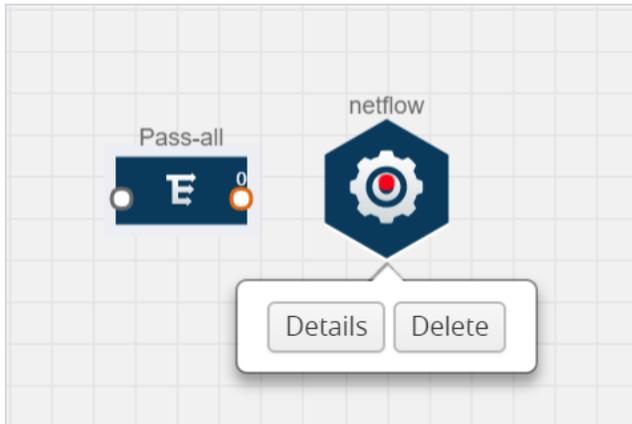


Figure 15: Selecting Details

3. In the **Alias** field, enter a name for the v5 NetFlow application.
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](#).

Add 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.

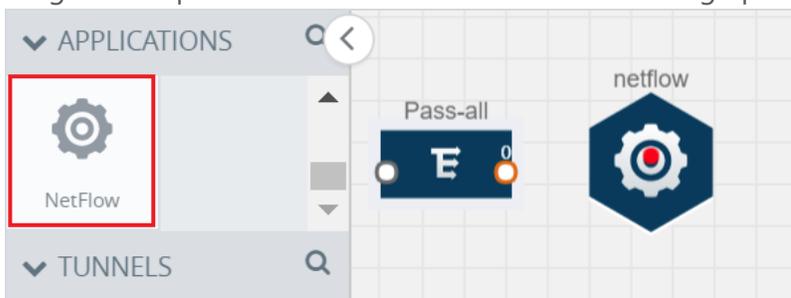


Figure 16: Dragging the NetFlow Application

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

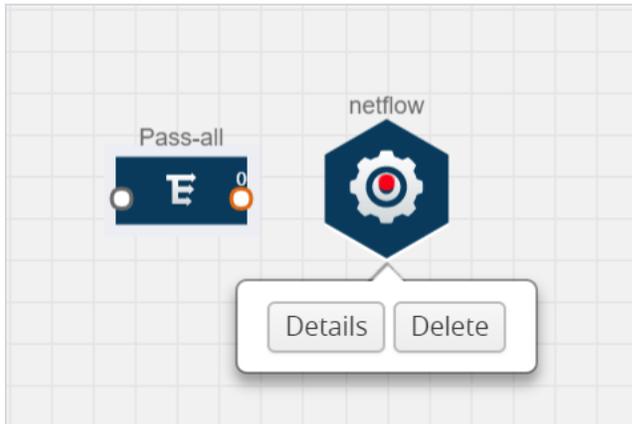


Figure 17: Selecting NetFlow Details

3. In the **Alias** field, enter a name for the NetFlow application.
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](#).
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](#).
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](#).

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.

Add NAT

To add a NAT device:

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

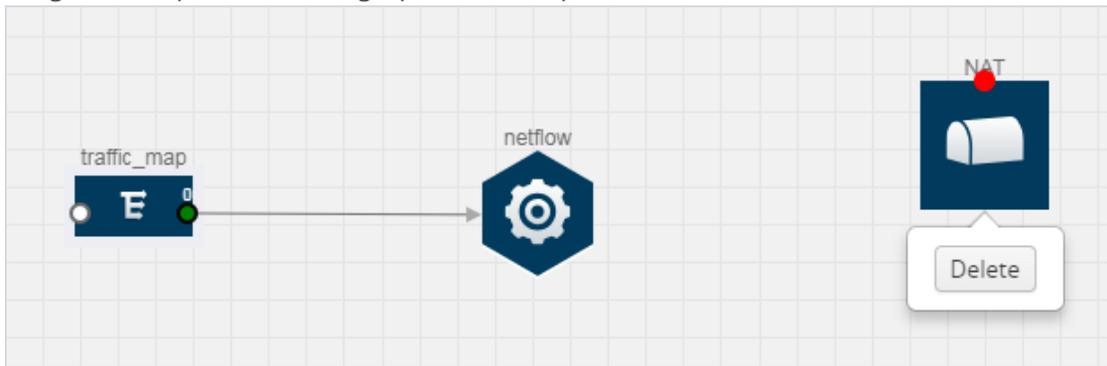


Figure 18: Adding NAT

Link 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.

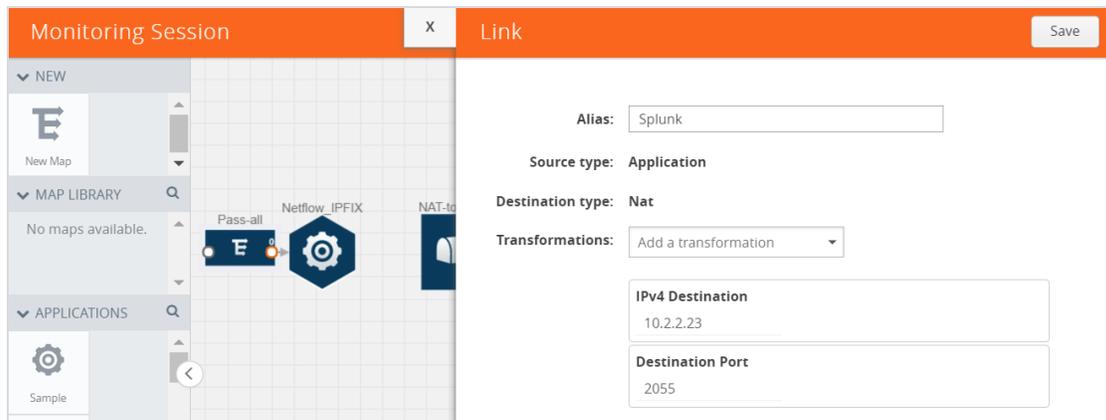


Figure 19: Creating a Link from NetFlow to NAT

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

NetFlow Examples

This section provides an example to demonstrate the NetFlow application configuration in the GigaVUE V Series nodes.

Example 1

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

1. Create a monitoring session. For steps, refer to [Create Monitoring Session](#).
2. In the monitoring session, create a Pass all map. A pass all map sends all the traffic received from the G-vTAP agents to the tunnel endpoint or NAT. For steps, refer to [Create 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.
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.
5. Drag and drop a v5 NetFlow application.

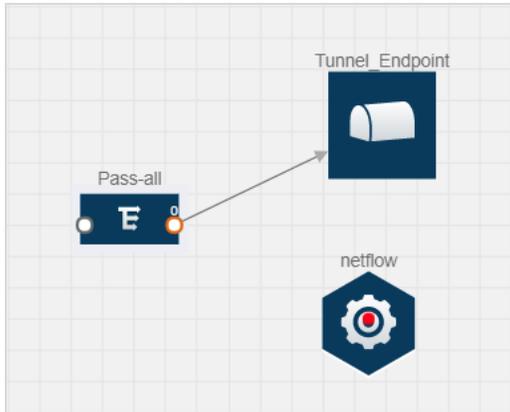


Figure 20: Adding a link from Pass-all Map to Tunnel_Endpoint

6. Click the NetFlow application and select **Details**. The Application quick view is displayed. For steps to configure the v5 NetFlow application, refer to [Add Version 5 NetFlow Application](#).
7. Create a link from the Pass all map to the v5 NetFlow application.
8. Drag and drop **NAT** to the graphical workspace.

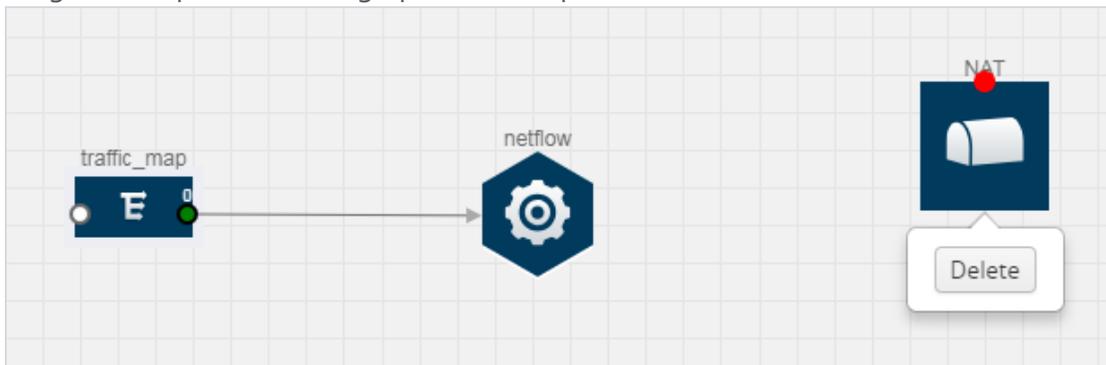


Figure 21: Adding a NAT Device

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 [Link NetFlow Application to NAT](#).

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

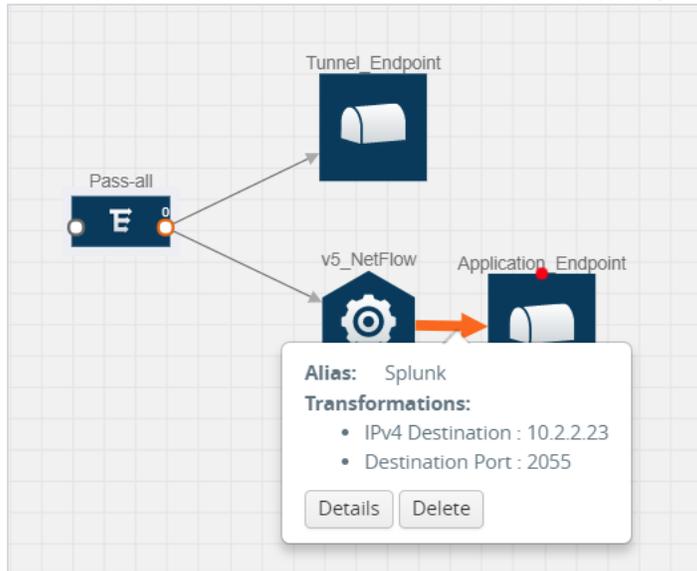


Figure 22: Viewing the Transformation Dialog Box

Deploy Monitoring Session

To deploy the monitoring session:

- Drag and drop one or more maps from the **MAP Library** to the workspace.
- (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.
- (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 [Add Applications to Monitoring Session](#).

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

The following figure illustrates three maps, one exclusion map, one application, and two tunnel endpoints dragged and dropped to the workspace.

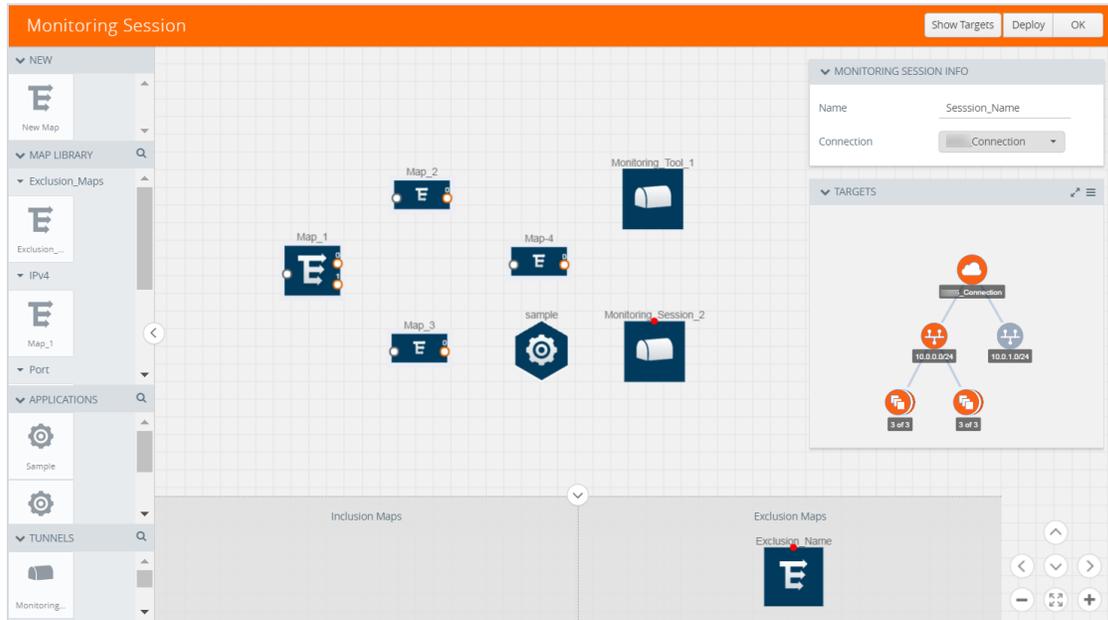


Figure 23: Dragging and Dropping the Maps, Applications, and Monitoring Tools

NOTE: You can add up to 8 links from a single map to different maps, applications, or monitoring tools.

5. Hover your mouse on the map, click the red dot, and drag the link over to another map, application, or tunnel. You can drag more than one link from a map to the destination. On these links, you can apply link transformation to alter the packets. For information about adding link transformation, refer to [Add Header Transformations](#).
6. Hover your mouse on the application, click the red dot, and drag the link (arrow) over to the tunnel endpoints. The traffic matching the rules in each action set is routed to maps, applications, or monitoring tools.
7. Click **Show Targets** to view details about the subnets and monitoring instances. The instances and the subnets that are being monitored are highlighted in orange.

8. Click **Deploy** to deploy the monitoring session. The status is displayed as **Success** in the Monitoring Sessions page. The session is successfully deployed on all GigaVUE V Series nodes and G-vTAP agents.

If the monitoring session is not deployed properly, then one of the following errors is displayed:

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

Click on the status link to view the reason for the partial success or failure.

9. Click **View** under Statistics to view and analyze the incoming and outgoing traffic.

You can also do the following in the Monitoring Session page:

- Use the **Redeploy** button to redeploy a monitoring session that is not deployed or partially successful.
- Use the **Undeploy** button to undeploy the selected monitoring session.
- Use the **Clone** button to duplicate the selected monitoring session.
- Use the **Edit** button to edit the selected monitoring session.
- Use the **Delete** button to delete the selected monitoring session.

Add 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 VNets 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 VNets 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 the following figure, 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.

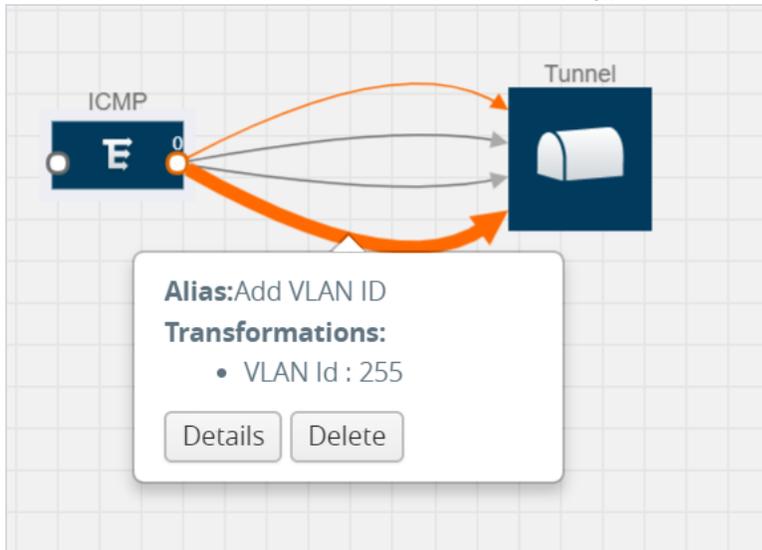


Figure 24: Action Set with Multiple Links

GigaVUE V Series node supports the following header transformations:

Table 3: 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.

Option	Description
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 25: Opening the Link Quick View

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

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

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

View 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.

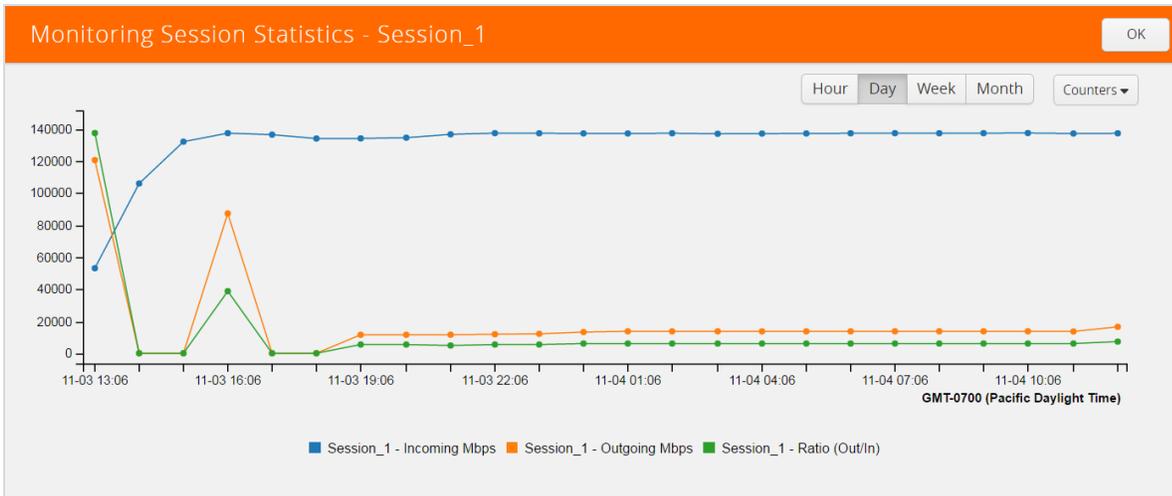


Figure 26: Viewing the Monitoring Session Statistics

You can click on Incoming Maps, Outgoing Maps, and Ratio at the bottom of the graph to view the statistics individually.

You can expand the **View Monitoring Session Diagram** and click on each individual map, application, and tunnel to view more details about the incoming and outgoing traffic on the selected statistics page. The Map Statistics page lets you choose the map rules to view the traffic matching the selected rule.

View Topology

You can have multiple VNet connections in GigaVUE-FM. Each VNet 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 **AnyCloud > Topology**.
2. Select a connection from the **Select connection...** list. The topology view of the subnets and instances is displayed.
3. (Optional) Select a monitoring session from the **Select Monitoring Session...** list. The topology view of the monitored subnets and instances in the selected session are displayed.

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 VMs installed with G-vTAP agents within the subnets being monitored.

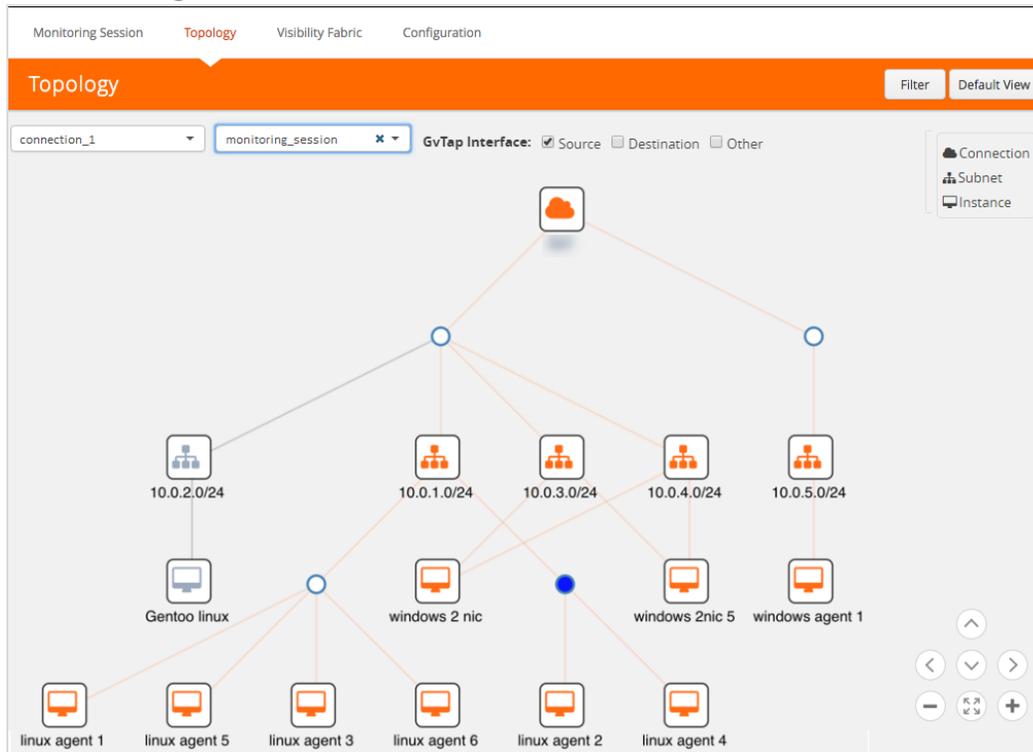


Figure 27: Viewing the Topology

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.
- 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.

Configure AnyCloud Settings

This section provides information on how to configure the maximum number of connections, refresh intervals for instance and non-instance inventory, and maximum batch size for monitoring session updates.

Use the **AnyCloud > Settings > Advanced** to edit the settings. Refer to the following table for more information about the settings:

Settings	Description
Maximum number of connections allowed	Specifies the maximum number of connections you can establish in GigaVUE-FM.
Refresh interval for instance target selection inventory (secs)	Specifies the frequency for updating the state of Virtual Machines.
Refresh interval for fabric deployment inventory (secs)	Specifies the frequency for updating the state of non-instance information such as subnets, security groups, images, and VMs.
Number of instances per GigaVUE V Series Node	Specifies the maximum number of instances that can be assigned to the GigaVUE V Series node.
Refresh interval for G-vTAP agent inventory (secs)	Specifies the frequency for discovering the G-vTAP agents available.
G-vTAP Agent Tunnel Type	Specifies the tunnel type of the G-vTAP agent.

Additional Sources of Information

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

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

Documentation

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.

 **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 4: Documentation Suite for Gigamon Products

Summary	Document
<ul style="list-style-type: none">• complete doc set for the respective release, minus Release Notes, in a zip file	All-Documents Zip
<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 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 .

Hardware Installation Guides

Summary	Document
<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
	GigaVUE-HC2 Hardware Installation Guide
	GigaVUE-HC3 Hardware Installation Guide
	GigaVUE TA Series Hardware Installation Guide
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
<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
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
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
<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
<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
	GigaVUE Cloud Suite for AWS QuickStart Guide
	*GigaVUE Cloud Suite for AWS Secret Regions Configuration Guide
	GigaVUE Cloud Suite for Azure Configuration Guide
	GigaVUE Cloud Suite for Kubernetes Configuration Guide
	*GigaVUE Cloud Suite for Nutanix Configuration Guide
	GigaVUE Cloud Suite for OpenStack Configuration Guide
GigaVUE Cloud Suite for VMware Configuration Guide	

Summary	Document
---------	----------

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
<ul style="list-style-type: none"> guidelines for the different types of cables used to connect Gigamon devices 	GigaVUE-OS Cabling Quick Reference Guide
<ul style="list-style-type: none"> compatibility information and interoperability requirements for Gigamon devices 	GigaVUE-OS Compatibility and Interoperability Matrix
<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

In-Product Help

<ul style="list-style-type: none"> how to install, deploy, and operate GigaVUE-FM. Provided from the GigaVUE-FM interface. 	GigaVUE-FM Online Help
<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

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