



GigaVUE Cloud Suite for Kubernetes 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 Kubernetes	5
Audience	5
Overview of GigaVUE Cloud Suite for Kubernetes	6
About GigaVUE Cloud Suite for Kubernetes Container Visibility	6
Configure Components in Kubernetes	9
Before You Begin	9
Launch GigaVUE-FM Instance	12
Configure GigaVUE Cloud Suite for OpenStack	12
Configure Kubernetes Containers	13
Configure Monitoring Sessions in Kubernetes	17
Overview of Visibility Components	18
Create Tunnel Endpoints	21
Create Monitoring Session	22
Configure Kubernetes Settings	59
Additional Sources of Information	60
Documentation	60
Documentation Feedback	63
Contact Technical Support	63
Contact Sales	63
The Gigamon Community	63

This page intentionally left blank to ensure new

GigaVUE Cloud Suite for Kubernetes

This guide describes how to install, configure, and deploy the GigaVUE Cloud Suite for Kubernetes Container Visibility in the native Kubernetes platform which can run on any of the public and private cloud platforms. Use this document for instructions on configuring the GigaVUE Cloud components and setting up the traffic monitoring sessions for the Kubernetes Containers visibility.

NOTE: GigaVUE Cloud Suite for Kubernetes Container Visibility is qualified to be operational on Kubernetes version 1.14.

Topics:

Audience

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

- **Node:** A node is a working machine in Kubernetes cluster. They are working units which can be physical, VM, or a cloud instance.
- **Cluster:** A group of nodes.
- **Image:** Containers are launched from these images (similar to images in cloud environments).
- **Pod:** A pod is a collection of containers and its storage inside a node of Kubernetes cluster. It is possible to create a pod with multiple containers inside it.
- **Service:** A service is a logical set of pods. It can be defined as an abstraction on the top of the pod which provides a single IP address and DNS name by which pods can be accessed.
- **Namespace:** Provides an additional qualification to a resource name. This is helpful when multiple teams are using the same cluster and there is a potential of name collision.
- **Replication controller:** Responsible for managing the pod life cycle.
- **DaemonSet:** A DaemonSet ensures that all (or some) nodes run a copy of a pod. As nodes are added to the cluster, pods are added to them. As nodes are removed from the cluster, those pods are garbage collected. Deleting a DaemonSet will clean up the pods it created.
- **Job:** Main function of a job is to create one or more pod and tracks about the success of pods.

For a detailed list of terms and definitions, refer to the Kubernetes Glossary:

<https://kubernetes.io/docs/reference/glossary/?fundamental=true>

Overview of GigaVUE Cloud Suite for Kubernetes

This chapter introduces the GigaVUE Cloud Suite for Kubernetes containers Visibility and the supported architecture. Refer to the following sections for details:

- "About GigaVUE Cloud Suite for Kubernetes Container Visibility" below
- "GigaVUE Cloud Components" on page 8
- "Traffic Capturing Mechanism" below

About GigaVUE Cloud Suite for Kubernetes Container Visibility

Kubernetes is an open source system that manages containers for application deployment, scaling and management. Containers are similar to virtual machines but are light weighted as they share the operating system among the applications running on them. Kubernetes is a container management platform with a framework to run distributed systems resiliently. Refer to the <https://kubernetes.io/docs/concepts/> to understand the Kubernetes concepts.

The GigaVUE Cloud Suite for Kubernetes Container Visibility supports the following network options:

- **Flannel:** A simple overlay networking solution for Kubernetes pods using Linux bridge. G-vTAP containers support Flannel networking for Kubernetes pods using VXLAN based encapsulation.
- **Calico:** A highly scalable networking and network policy solution for connecting Kubernetes pods based on the same IP networking principles as the Internet. G-vTAP containers support CALICO with IP-in-IP Encapsulation.

Traffic Capturing Mechanism

GigaVUE Cloud Suite for Kubernetes Containers Visibility captures traffic using the G-vTAP container. Refer to the following section for details.

G-vTAP Container

Gigamon Traffic Acquisition Container or the G-vTAP container is an application container that is launched on each of the worker nodes in the Kubernetes cluster. G-vTAP container acquires and mirrors traffic from other containers in the same worker node. Traffic from various Kubernetes

Networking Infrastructure are encapsulated using VXLAN and sent to one of the following components, where it is further processed with maps and GigaSMART applications:

- GigaVUE® V Series node
- Physical node such as a GigaVUE H Series or a GigaVUE-TA Series node

Traffic from the GigaVUE V Series nodes/physical node is then sent to the required tools.

Figure 1: GigaVUE Cloud Components for Kubernetes using G-vTAP Containers shows a high level architecture of GigaVUE Cloud Suite for Kubernetes using G-vTAP containers as the source for acquiring the traffic.

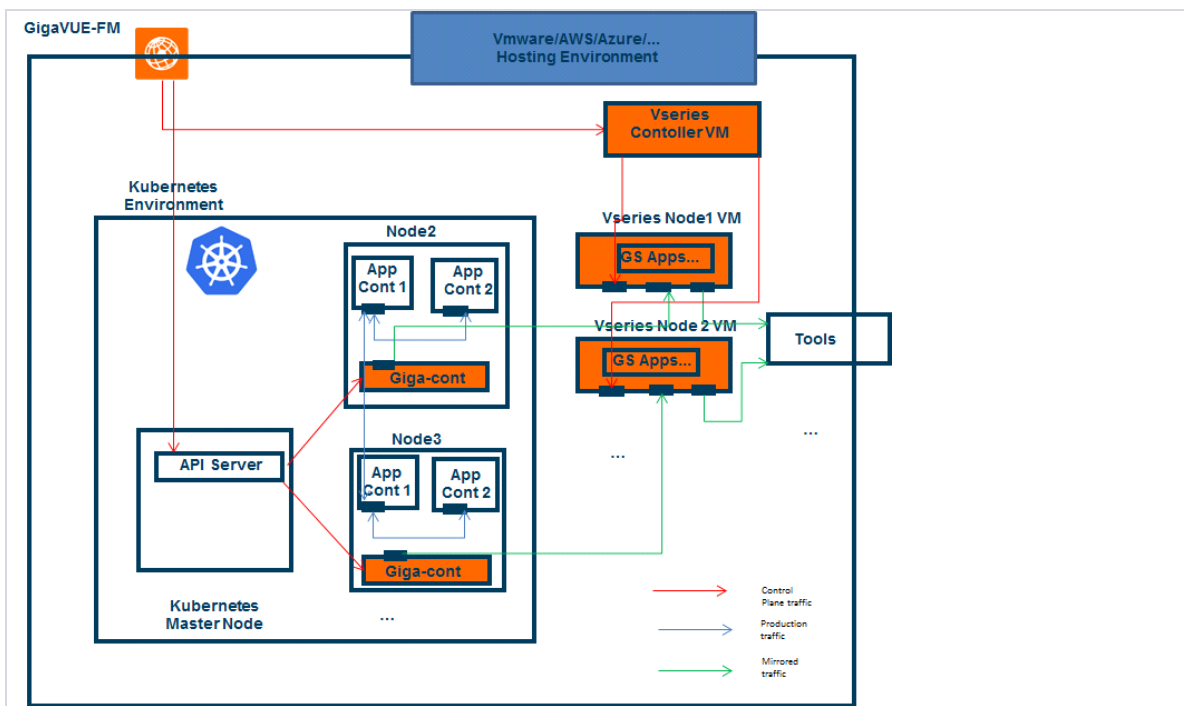


Figure 1: GigaVUE Cloud Components for Kubernetes using G-vTAP Containers

GigaVUE Cloud Components

The GigaVUE Cloud Suite for Kubernetes Containers 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.

You must have GigaVUE-FM installed either on-premises or launched from any of the supported cloud platforms. Refer to the *"GigaVUE-FM Installation Guide"* for details on installing and launching GigaVUE-FM.

GigaVUE-FM manages the configuration of the following visibility components, if you use GigaVUE® V Series node in OpenStack:

- G-vTAP Container which is the Gigamon traffic acquisition container
 - GigaVUE® V Series Nodes
 - GigaVUE® V Series Controllers
- **G-vTAP Container** is the Traffic Acquisition Component of Gigamon's Kubernetes Network Visibility Offering. It receives mirrored traffic from various Kubernetes Networking Infrastructure such as Flannel and Calico and overlays (VXLAN) them to GigaVUE V Series nodes or the physical nodes for further processing.
 - **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 containers. It applies filters, manipulates the packets using GigaSMART applications, and distributes the optimized traffic to:
 - Cloud-based tools
 - Backhaul to GigaVUE Cloud using GRE or VXLAN tunnels
 - On-premise tools, such as a GigaVUE H Series device

Configure Components in Kubernetes

This chapter describes how to configure G-vTAP containers, GigaVUE® V Series Controllers, and GigaVUE® V Series nodes in your environment. Refer to the following sections for details:---

- ["Before You Begin" below](#)
- ["Launch GigaVUE-FM Instance" on page 12](#)
- ["Deploy G-vTAP Containers" on page 15](#)
- ["Configure GigaVUE Cloud Suite for OpenStack" on page 12](#)

NOTE: GigaVUE® Fabric Manager (GigaVUE-FM) must be already installed and configured in your environment.

Before You Begin

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

- ["Security Group " on the next page](#)
- ["Key Pairs" on page 12](#)

Minimum Compute Requirements

The minimum recommended computing requirements are listed in [Table 1: Minimum Compute Requirement](#).

Table 1: Minimum Compute Requirement

Compute Instances	vCPU	Memory	Disk Space	Description
GigaVUE® 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

Compute Instances	vCPU	Memory	Disk Space	Description
GigaVUE® 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.

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 Containers in your project, you add rules that control the inbound traffic to instances, and a separate set of rules that control the outbound traffic.

It is recommended to create a separate security group for each component using the rules and port numbers listed in [Table 2: Security Group Rules](#).

[Table 2: Security Group Rules](#) lists the rules and port numbers for each component.

Table 2: Security Group Rules

Direction		Protocol	Port	CIDR	Purpose
GigaVUE-FM					
Inbound	HTTPS	TCP	443	Any IP address	Allows G-vTAP Container Managers, 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.

Direction		Protocol	Port	CIDR	Purpose
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.

GigaVUE V Series Controller

Inbound	IPv4	TCP	9902	GigaVUE-FM IP address	Allows GigaVUE-FM to communicate with GigaVUE V Series Controllers.
---------	------	-----	------	-----------------------	---

GigaVUE V Series node

Inbound	Custom TCP Rule	TCP(6)	9903	GigaVUE V Series Controller IP address	Allows GigaVUE V Series Controllers to communicate with GigaVUE V Series nodes.
---------	-----------------	--------	------	--	---

GRE Traffic

Inbound	Custom Protocol Rule	GRE (47)	All	Any IP address	Allows monitored traffic from GigaVUE V Series nodes to be sent to the monitoring tools using the L2 GRE tunnel.
---------	----------------------	----------	-----	----------------	--

VXLAN Traffic

Inbound	Custom UDP Rule	VXLAN	4789	Any IP address	Allows mirrored traffic from G-vTAP Containers 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:

- [Table 2: Security Group Rules](#) lists only the ingress rules. Make sure the egress ports are open for communication.
- Along with the ports listed in [Table 2: Security Group Rules](#), make sure the suitable ports required to communicate with Service Endpoints such as Identity, Compute, and Cloud Metadata are also open.

For information about creating security groups refer to the website of the corresponding cloud platform.

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 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.

Launch GigaVUE-FM Instance

You must have a GigaVUE-FM instance launched in your environment. Refer to the *GigaVUE-FM Installation Guide* and *GigaVUE-FM User's Guide* for details on installing and launching GigaVUE-FM.

Configure GigaVUE Cloud Suite for OpenStack

To configure the GigaVUE Cloud Suite for OpenStack using GigaVUE V Series node and GigaVUE V Series controller, you must perform the tasks listed in the following table:

	Task	Chapters in GigaVUE Cloud Suite for OpenStack Configuration Guide
1	Establish a connection between GigaVUE-FM and the environment. NOTE: You must ensure to set the Tap Method field to 'None' for the Kubernetes Container visibility option.	<i>Connect to OpenStack</i>
2	Configure the GigaVUE V Series Controller.	<i>Configure GigaVUE V Series Controllers</i>
3	Configure the GigaVUE V Series Node.	<i>Configure V Series Nodes</i>

Configure Kubernetes Containers

To configure the Kubernetes container you must connect to Kubernetes and configure the G-vTAP Container in the Kubernetes environment. Refer to the following sections for details:

- ["Create Service Account Token" below](#)
- ["Configure Kubernetes" on the next page](#)
- ["Configure G-vTAP Containers" on page 15](#)

Create Service Account Token

Prior to configuring the Kubernetes container, you must establish a connection to the Kubernetes API server and generate the service account token.

The following are the steps to create a service account token:

1. Create a service account token for GigaVUE-FM to use:

```
kubectl create serviceaccount <serviceaccountname >
```

2. View the service account information and token name:

```
kubectl describe serviceaccounts <serviceaccountname >
```

3. Get the token to be input to GigaVUE-FM:

```
kubectl describe secret <secret value in the above output>
```

1>Create the cluster role binding with a role that has proper permissions, in this example an admin role:

```
kubectl create clusterrolebinding <clusterrolebindingname> --  
clusterrole=cluster-admin --serviceaccount=default:<serviceaccountname>
```

2>Create a docker registry key which in input to FM in connection:

```
kubectl create secret docker-registry <docker-key-name> --docker-  
server=<your-registry-server> --docker-username=<your-name> --docker-  
password=<your-pword>
```

4. Patch the service account with the key:

```
kubectl patch serviceaccount <serviceaccountname> -p '{"imagePullSecrets":  
[{"name": "regcred "}]}'
```

5. Display the API endpoint IP and port for GigaVUE-FM to connect to:

```
kubectl cluster-info
```

Configure Kubernetes

To configure Kubernetes:

1. In GigaVUE-FM, on the top navigation pane, select **Cloud**.
2. On the left navigation pane, select **Kubernetes > Monitoring Domain**.
3. Click **New**. The Kubernetes Configuration page appears.

Figure 1: Kubernetes Configuration Page.

4. Enter or select the following details:

Table 1: Fields for Kubernetes Configuration

Field	Description
Monitoring Domain	Name of the monitoring domain.
Alias	Name of the connection.
Authentication Type	Authentication type to connect to the Kubernetes API Server. Service Account Token is the only supported method for this release. To create a Service Account Token, refer to .

Field	Description
Token	Service Account Token details.
API Server URL	API server URL of the Kubernetes Container
V Series Environment	GigaVUE V Series environment configured during OpenStack Configuration. NOTE: If you intend to use a GigaVUE H Series box, then you must not configure the V Series environment.

Configure G-vTAP Containers

The G-vTAP Container is a Gigamon specific container that is deployed on the container worker nodes where visibility is required.

Deploy G-vTAP Containers

GigaVUE-FM launches a DaemonSet with the G-vTAP Containers on the Kubernetes worker nodes where visibility is required and the G-vTAP Container Managers in the Kubernetes master node. The container images are provided through a registry.

NOTE: You can also fetch the image of the Gigamon Traffic Acquisition container from the customer portal using FTP, TFTP, SCP or any other desired method.

When a new worker node comes up, this will automatically have the G-vTAP containers. GigaVUE-FM launches a job that sets up mirroring on the networking layer from the required application containers to the G-vTAP container. This job is specific for the type of networking mechanism used in the Kubernetes environment. GigaVUE-FM passes the relevant parameters to the job, but you are responsible for providing the input about the type of networking present in your environment.

You can configure the traffic to move from the G-vTAP Container to the GigaVUE V Series node where further processing with maps and GigaSMART applications take place. Traffic from the GigaVUE V Series node is then sent to the required tools.

NOTE: If a GigaVUE H Series or a GigaVUE TA Series device is configured, then the traffic is sent from the G-vTAP container to the physical devices where further processing with GigaSMART operations occur.

To configure the G-vTAP containers, enter or select the following details as shown in the following table.

Table 2: Fields for G-vTAP Container Configuration

Field	Description
Name Space	Namespace to separate the cluster resources for users. The default name space is <i>Default</i> .
Tap Manager Image	G-vTAP Container Manager Image
Container Image	G-vTAP Container Image
Container Image Pull Policy	Policy to pull the G-vTAP Container image from the repository. Choose the options as per your requirement: <ul style="list-style-type: none"> • Never: Policy is available locally and therefore need not be pulled from the repository • Always: Policy is not available locally and therefore must always be pulled from the repository • IfNotPresent: Policy is not present
Labels	Key/value pairs attached to objects such as pods
Networking Type	Networking type. Options are <ul style="list-style-type: none"> • Flannel • Calico
Tunnel Type	The type of tunnel for sending the traffic from G-vTAP container to GigaVUE V Series nodes. Only the VXLAN tunnel type is supported.
VXLAN Tunnel ID	VXLAN tunnel ID.
VXLAN Port	VXLAN port.
CPU Request	CPU requested for the G-vTAP container.
CPU Limit	CPU limit for the G-vTAP container.
Memory Request	Memory requested for the container.
Memory Limit	Memory limit for the containers.

Configure Monitoring Sessions in Kubernetes

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 the various monitoring tools.

Refer to the following sections for details:

- ["Overview of Visibility Components" on the next page](#)
- ["Create Tunnel Endpoints" on page 21](#)
- ["Create Monitoring Session" on page 22](#)
- ["Configure Kubernetes Settings" on page 59](#)

NOTE: For a GigaVUE H Series node or a GigaVUE TA Series node, you can only create a passall map from the G-vTAP Containers to the physical nodes. The applicable ATS filters for this configuration include the following:

- Ip4Src
- Ip6Src
- Ip4Dst
- Ip4Dst
- SrcVmPrefix
- SrcVmTag
- DstVmPrefix
- DstVmTag

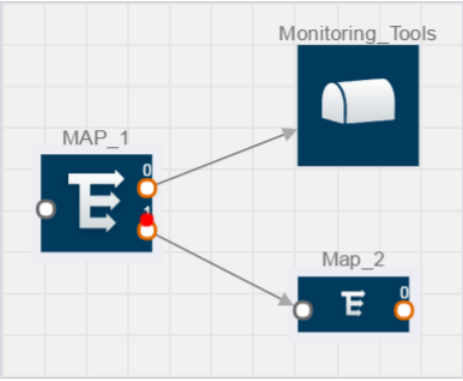
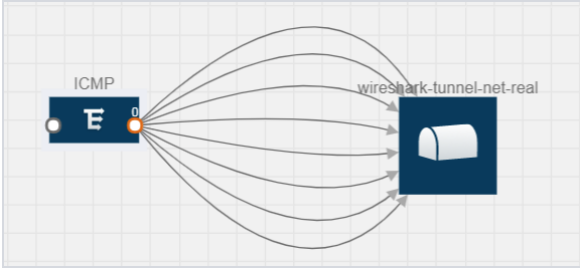
Overview of Visibility Components

The GigaVUE V Series node aggregates the traffic from G-vTAP containers 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 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.
Rule	A rule (R) contains specific filtering criteria that the packets must match. The filtering criteria lets you determine the targets and the (egress or ingress) direction of tapping the network traffic. A rule is also associated with priority and action set.
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.

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 1: 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 1: 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 2: Action Set with Multiple Links.</p>  <p><i>Figure 2: 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 1: Action Set, 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 52.</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>

Parameter	Description
Application	An application performs operations such as sampling, slicing, and masking on the traffic.
Inclusion Map	An inclusion map determines the instances to be included for monitoring. This map is used only for target selection.
Exclusion Map	An exclusion map determines the instances to be excluded from monitoring. This map is used only for target selection.
Target	A target determines the instances that are to be monitored. Targets are determined based on the following formula:
Automatic Target Selection (ATS)	A built-in feature that automatically selects the cloud instances based on the rules defined in the maps, inclusion maps, and exclusion maps in the monitoring session.
Tunnel	A tunnel lists the monitoring tools to which the traffic matching the filtered criteria is routed.

Create 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 **Kubernetes > Settings**.
3. Select the **Tunnel Spec Library** tab. The Tunnel Library page appears. Refer to [Figure 3: Tunnel Library](#).

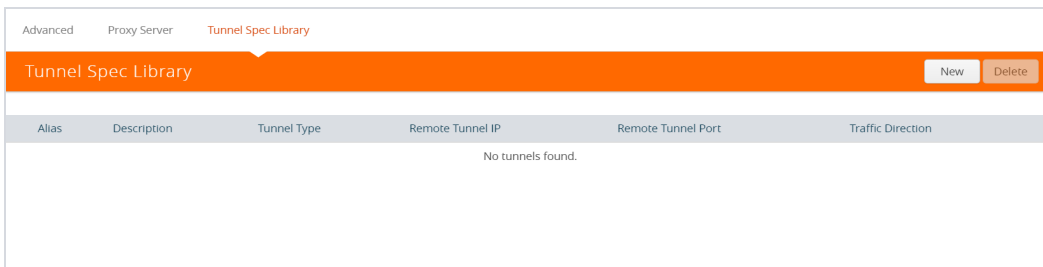


Figure 3: Tunnel Library

4. Click **New**. The Add Tunnel Spec page appears. Refer to [5. On the Add Tunnel Spec page, select or enter the appropriate information in the fields. Refer to Table 2: Field Descriptions for Tunnel Endpoint.](#)
5. On the Add Tunnel Spec page, select or enter the appropriate information in the fields. Refer to [Table 2: Field Descriptions for Tunnel Endpoint.](#)

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.

Field	Description
Traffic Direction	<p>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.</p> <p>NOTE: Traffic Direction In is not supported in the current release.</p>
Remote Tunnel IP	<p>The IP address of the tunnel destination endpoint.</p> <p>NOTE: You cannot create two tunnels from a V Series node to the same IP address.</p>

6. Click **Save**.
7. Select **Kubernetes > Settings > TunnelSpecLibrary** and verify the tunnel endpoint added to GigaVUE-FM.

Create Monitoring Session

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

When a new target VM is added to your environment, GigaVUE-FM automatically detects and adds the VM into your monitoring session. Similarly, when a VM 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" on the next page](#)
- ["Clone Monitoring Session" on page 25](#)
- ["Create Map" on page 25](#)
- ["Add Applications to Monitoring Session" on page 31](#)
- ["Deploy Monitoring Session" on page 50](#)
- ["View Statistics" on page 54](#)
- ["View Topology" on page 57](#)

Create New Session

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

To create a new session:

1. In GigaVUE-FM, on the top navigation pane, select **Cloud**.
2. Select **Kubernetes > Monitoring Session**. The Monitoring Sessions page appears. Refer to [Figure 4: Monitoring Sessions](#).




Figure 4: Monitoring Sessions

3. Click **New**.

The Monitoring Session configuration page appears. Refer to [Figure 5: Creating Monitoring Session](#).

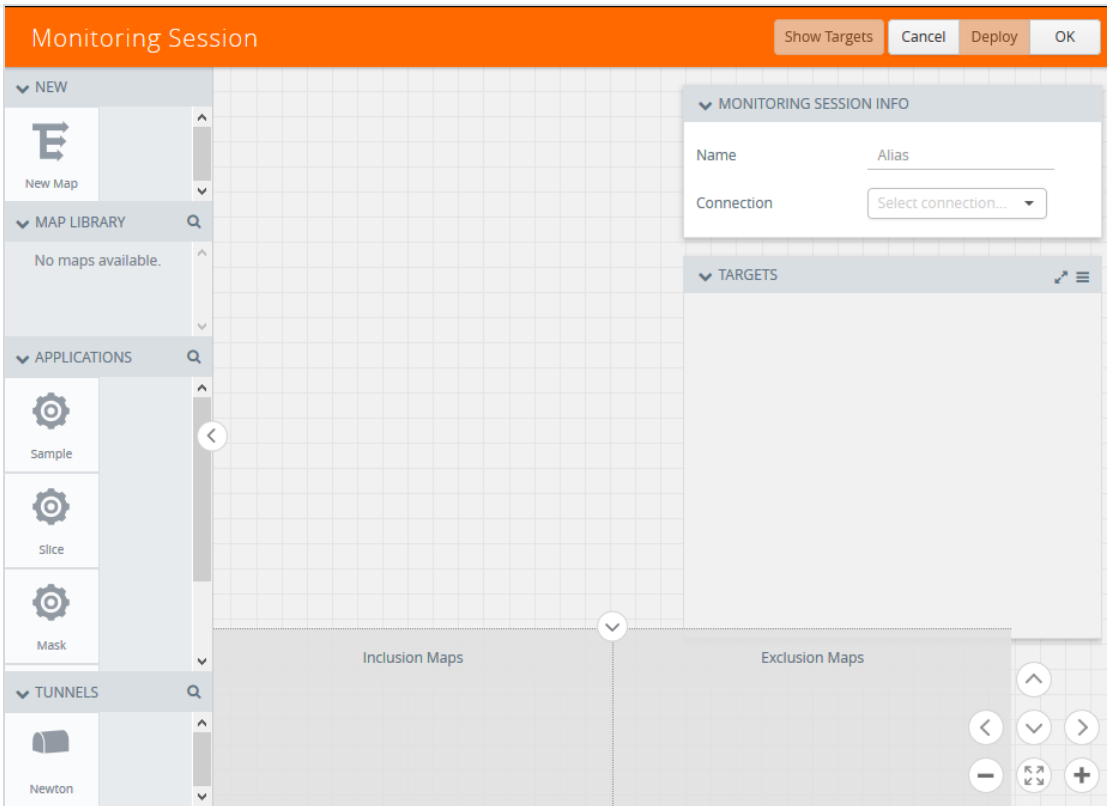


Figure 5: Creating Monitoring Session

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

Table 3: Fields for Session Info

Field	Description
Alias	The name of the monitoring session.
Monitoring Domain	The name of the monitoring domain.
Connection	The OpenStack connection(s) that are to be included as part of the monitoring domain. You can select the required connections that need to be part of the monitoring domain.

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 shown in [Table 3: Fields for Session Info on page 24](#).

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.

Create Map

Each map can have up to 32 rules associated with it. [Table 6: 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 5: 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 matching the specified source MAC address is selected.
MAC Destination	The ingress traffic 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 VMs 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 **Kubernetes > Monitoring Session**.
2. Click **New**. The Monitoring Sessions page is displayed.
3. Create a new session. Refer to "[Create New Session](#)" on page 23.
4. From **Maps**, drag and drop a new map template to the workspace.
5. Click on the map, then click details.

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

The screenshot shows a web interface for configuring a map named 'Map_1'. At the top, there are 'Save' and 'Add to Library' buttons. Below the title, there are input fields for 'Alias' (Map_1) and 'Comments'. A 'Map Rules' section contains an 'Add a Rule' button. Underneath, 'Rule 1' is defined with three search dropdowns: 'Search Layer 2 Conditions...', 'Search Layer 3 Conditions...', and 'Search Layer 4 Conditions...'. The 'Priority' is set to 0 and 'ActionSet' is set to 0. A 'Rule Comment' field is present. Two conditions are listed: 'Ether Type' with a value of 'IPv4' and '0x0800', and 'Protocol' with a value of 'TCP' and '6'. Each condition has a close button (x).

Figure 6: Creating a New Map

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

Table 6: Fields for Creating a New Map

Parameter	Description
Alias	The name of the new map. NOTE: The name can contain alphanumeric characters with no spaces.
Comments	The description of the map.
Rule Conditions 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. Refer to Figure 7: L2 Conditions. <i>Figure 7: L2 Conditions</i> Select a condition from the Search L3 Conditions drop-down list and specify a value. Refer to Figure 8: L3 Conditions. <i>Figure 8: L3 Conditions</i> (Optional) If you have selected TCP or UDP as the protocol in the L3 conditions, then select Port Source or Port Destination from the Search L4 Conditions drop-down list and specify a value. If you have selected conditions other than TCP or UDP, then the Search L4 Conditions drop-down list is disabled. Refer to Figure 9: L4 Conditions. <i>Figure 9: L4 Conditions</i> (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. 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.

8. Click **Save**.

To edit or delete a map, click a map and select **Details** to edit the map or **Delete** to delete the map as shown in [Figure 10: Editing or Deleting a Map](#).

Figure 10: Editing or Deleting a Map

Click the **Show Targets** button to view the monitoring targets highlighted in orange. Refer to [Figure 11: Viewing the Topology](#).

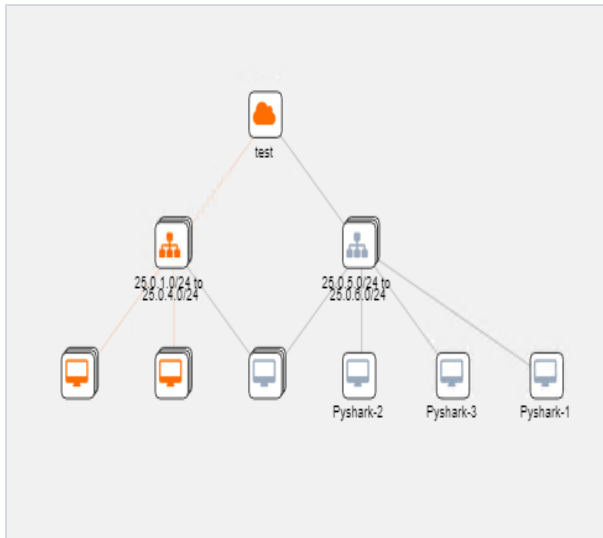




Figure 11: Viewing the Topology

Click on  to expand the **Targets** dialog box. Click on  to change the view from the Topology view to the targets view. To view details about a VM, click the arrow next to the VM. Refer to [Figure 12: Viewing the VMs](#).

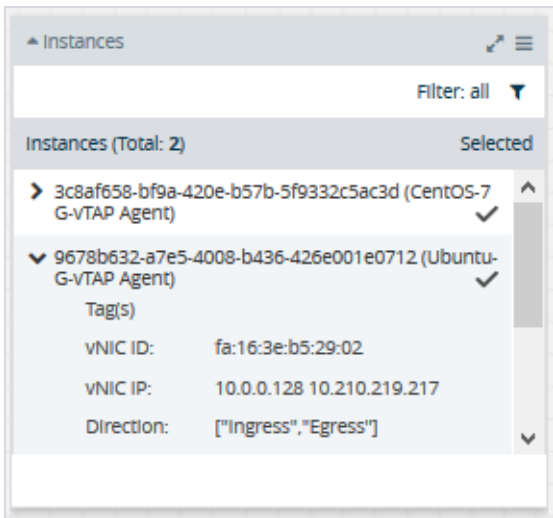


Figure 12: Viewing the VMs

Click on the Filter icon to filter Instances based on the Instance Name Prefix, IP address, or MAC address. Refer to [Figure 13: Filtering Instances](#).

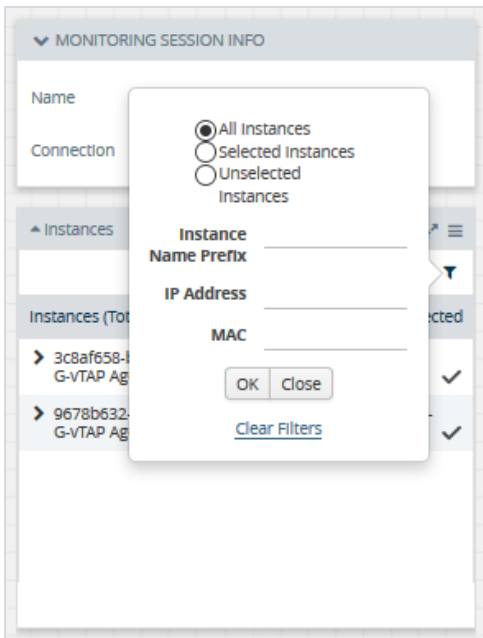


Figure 13: Filtering Instances

Add Applications to Monitoring Session

Gigamon supports the following GigaSMART applications:

- "Sampling" below
- "Slicing" on page 33
- "Masking" on page 34
- "NetFlow" on page 36

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

NOTE: If you have opted to send the traffic to a physical node such as a GigaVUE H Series box, then these GigaSMART operations will not be visible in the monitoring session page.

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.

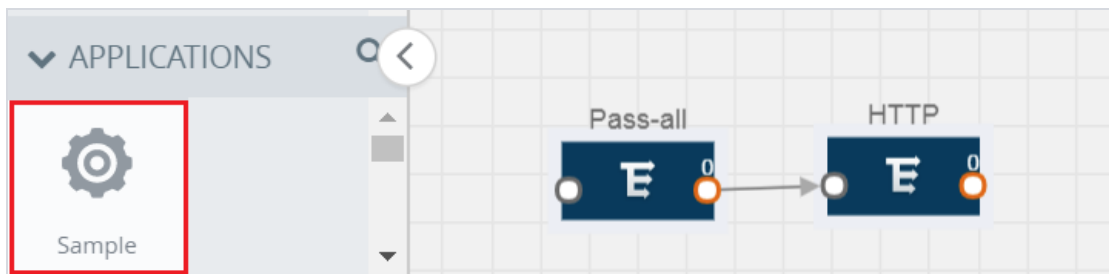


Figure 14: Dragging the Sample Application

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

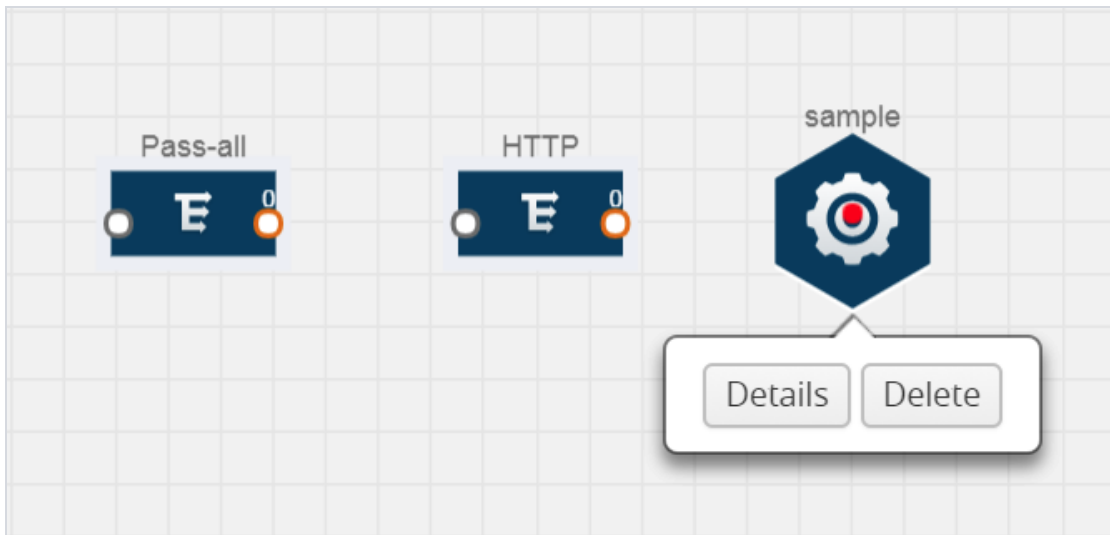


Figure 15: Selecting Details

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

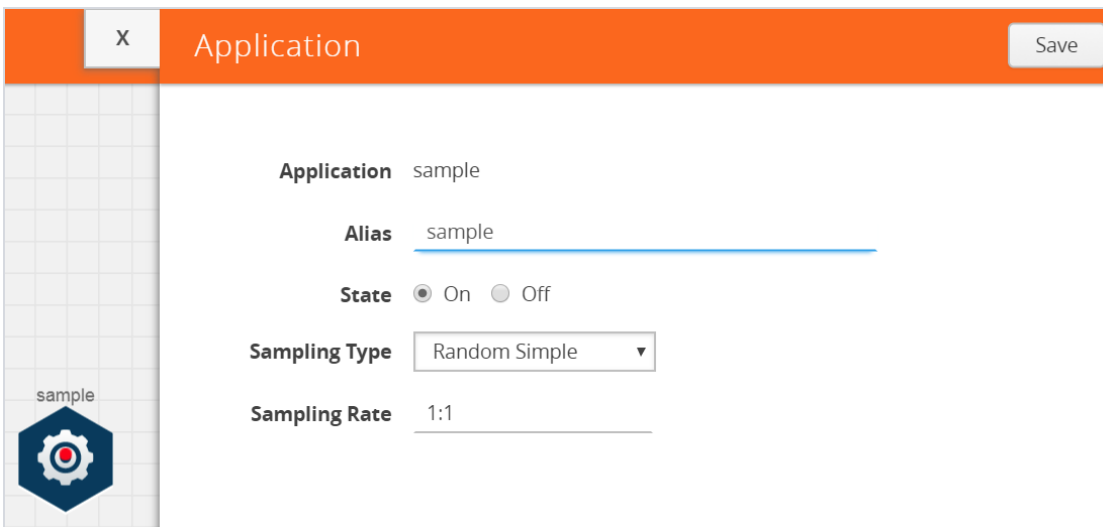


Figure 16: Viewing Sample Application Quick View

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.

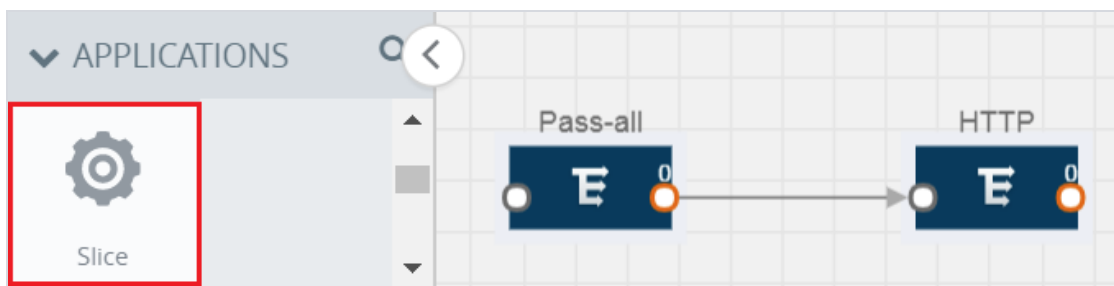


Figure 17: Dragging the Slice Application

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

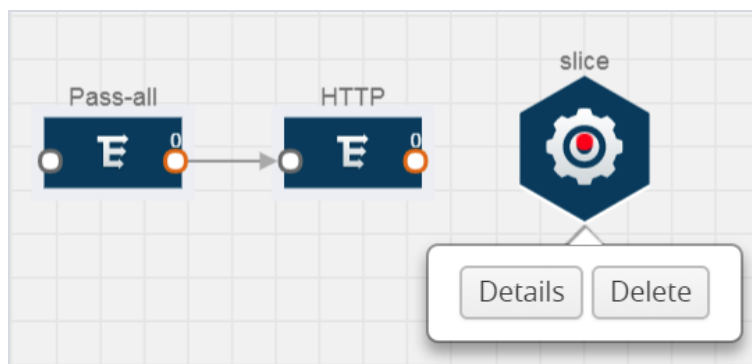


Figure 18: Selecting Details

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

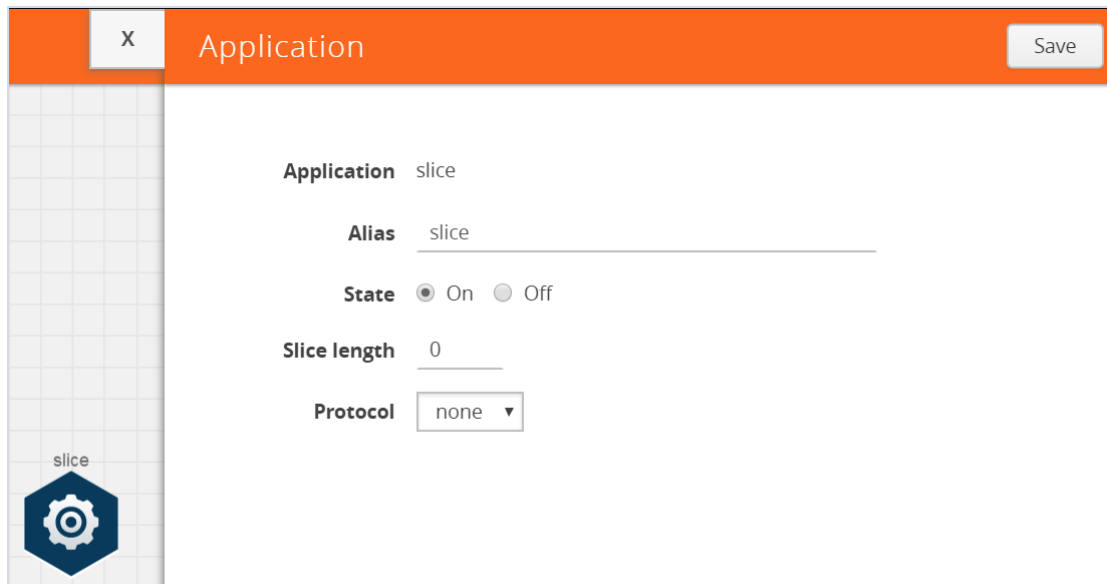


Figure 19: Viewing Slice Application Quick View

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 20: Dragging the Mask Application

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

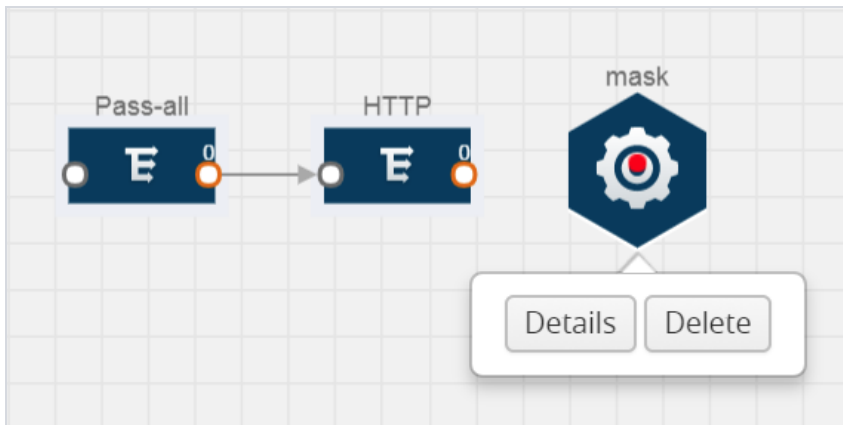


Figure 21: Selecting Details

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

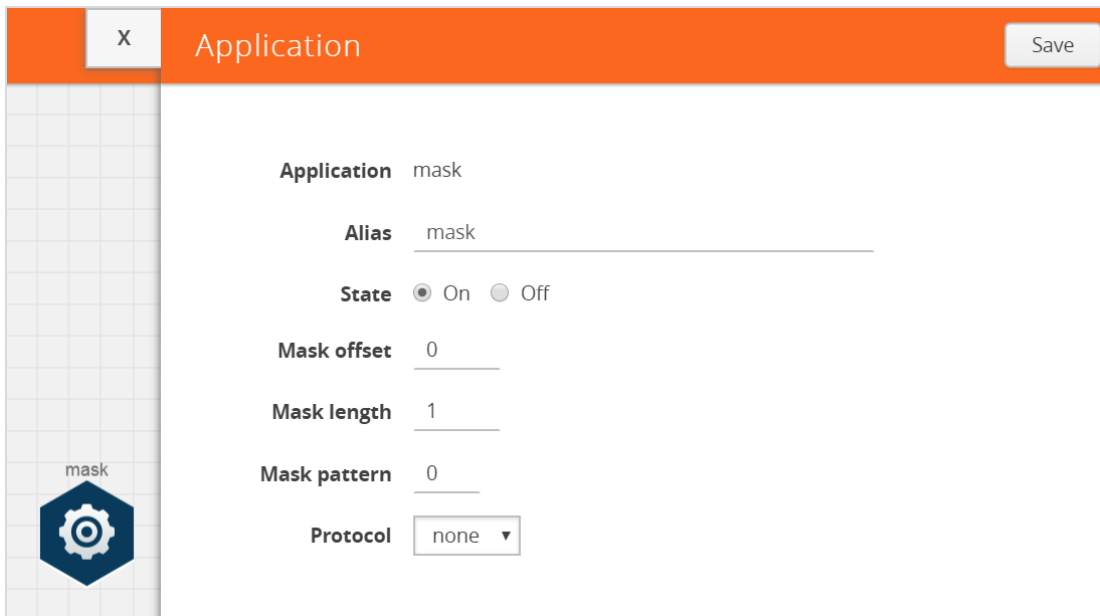


Figure 22: Viewing Mask Application Quick View

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 the next page](#). 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 41](#).

Figure 23: NetFlow on GigaVUE V Series Node shows an example of a NetFlow application created on a GigaVUE V Series node in the monitoring session.

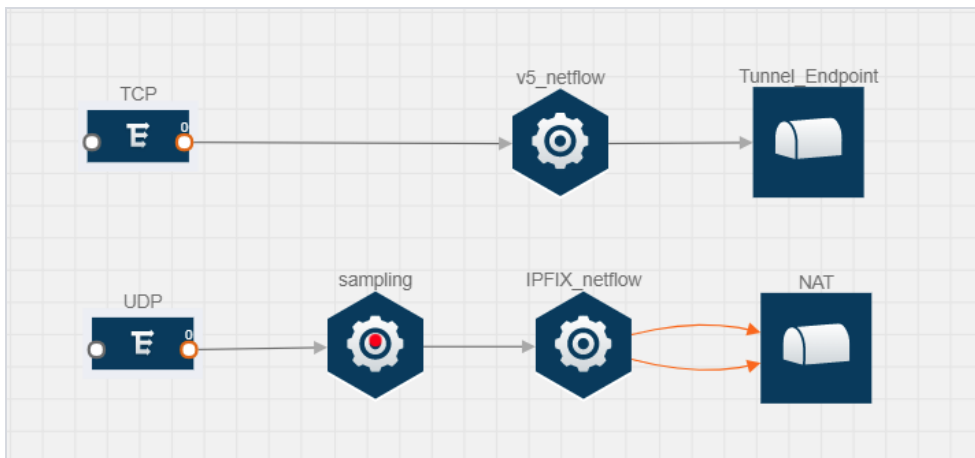


Figure 23: NetFlow on GigaVUE V Series Node

The NetFlow record generation is performed on GigaVUE V Series node running the NetFlow application. In [Figure 23: NetFlow on GigaVUE V Series Node](#), incoming packets from G-vTAP containers 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 46.

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 7: 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

	Description	Supported NetFlow Versions
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
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

	Description	Supported NetFlow Versions
IP TTL	For IPv4, configures the value of Time to Live (TTL) as a key field. For IPv6, configures the value of the Hop Limit field as a key field.	IPFIX
IP Version	Configures the IP version field in the IP packet header as a key field.	v9 and IPFIX
IPv6		
IPv6 Dest IP	Configures the IPv6 destination address in the IP packet header as a key field.	v9 and IPFIX
IPv6 Flow Label	Configures the value of the IPv6 flow label field in the IP packet header as a key field.	v9 and IPFIX
IPv6 ICMP Code	Configures the code of the IPv6 ICMP message as a key field.	IPFIX
IPv6 ICMP Type	Configures the type of the IPv6 ICMP message as a key field.	IPFIX
IPv6 ICMP Type Code	Configures the type and code of the IPv6 ICMP message as a key field.	IPFIX
IPv6 Payload Length	Configures the value of the payload length field in the IPv6 header as a key field.	IPFIX
IPv6 Src IP	Configures the IPv6 source address in the IP packet header as a key field.	v9 and IPFIX
Transport		
L4 Dest Port	Configures the destination port identifier in the transport header as a key field.	v9 and IPFIX
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

	Description	Supported NetFlow Versions
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 8: Collect/Non-Key Elements

	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

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

	Description	Supported NetFlow Versions
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
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

	Description	Supported NetFlow Versions
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.

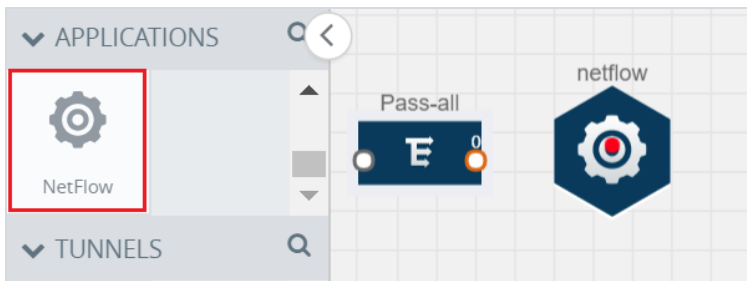


Figure 24: Dragging the NetFlow Application

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

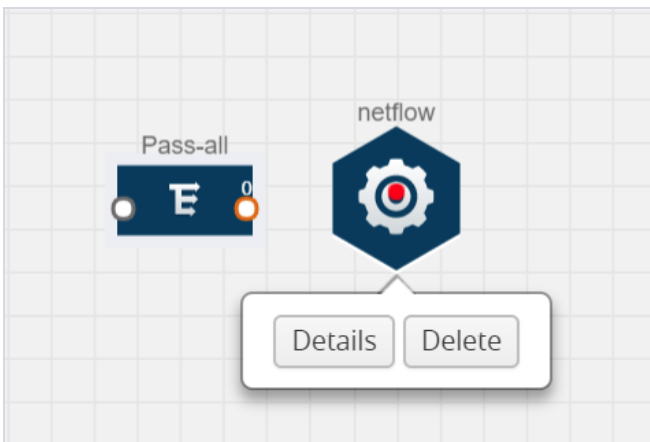


Figure 25: Selecting Details

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

Figure 26: 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 48.

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.

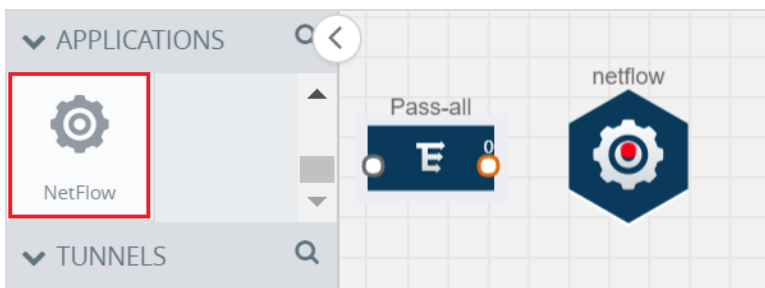


Figure 27: Dragging the NetFlow Application

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

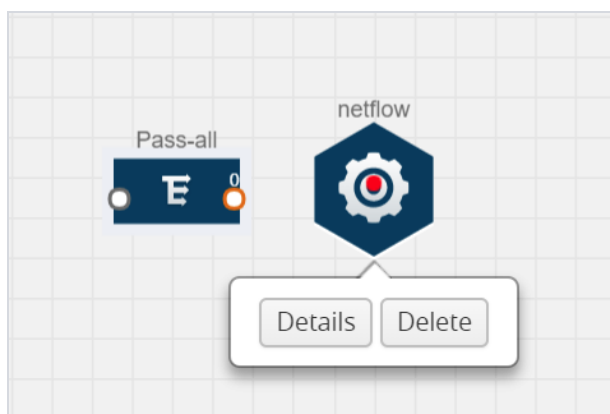


Figure 28: Selecting NetFlow Details

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

Figure 29: 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 containers. 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 38.
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 41.
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 48.

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:

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

Figure 30: Adding NAT

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

Figure 31: Selecting Details

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

Figure 32: 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**.

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 33: 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 34: Linking NetFlow to NAT

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

NetFlow Examples

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

- ["Example 1" below](#)

Example 1

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

1. Create a monitoring session. For steps, refer to ["Create Monitoring Session" on page 22](#).

Figure 35: Creating a Monitoring Session

2. In the monitoring session, create a Pass all map. A pass all map sends all the traffic received from the G-vTAP containers to the tunnel endpoint or NAT. For steps, refer to ["Create Map" on page 25](#).

Figure 36: 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 37: 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 38: Creating a Link from Pass-all Map to Tunnel_Endpoint

5. Drag and drop a v5 NetFlow application.

Figure 39: 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](#)" on page 44.

Figure 40: Configuring the NetFlow Application

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

Figure 41: 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 "[Add NAT](#)" on page 47.

Figure 42: 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](#)" on page 47.

Figure 43: 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 44: Viewing the Transformation Dialog Box

Deploy 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 "[Add Applications to Monitoring Session](#)" on page 31.

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

[Figure 45: 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 45: 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 46: Connecting the Maps, Applications, and Monitoring Tools](#).

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 46: 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 46: Connecting the Maps, Applications, and Monitoring Tools

1Click **Show Targets** to view details about the subnets and monitoring instances.

The instances and the subnets that are being monitored are highlighted in orange.

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 containers. Click on the status link in the Status column on the Monitoring Session page to view the Monitoring Session Deployment Report.

Monitoring Sessions					
Redeploy Undeploy New Clone Edit Delete					
<input type="checkbox"/>	Name	Connection	# of Targets	Status	Statistics
<input type="checkbox"/>	Connection1_Monitoring	Connection1	2	● Success	View
Total Items : 1					

When you click on the Status link, the Deployment Report is displayed. Refer to [Figure 47: Monitoring Session Deployment Report](#).

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

Figure 47: 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 V Series node failure.
- **Failure**—The session is not deployed on any of the V Series nodes and G-vTAP containers.

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.

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 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 48: Action Set with Multiple Links](#), 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 48: Action Set with Multiple Links

GigaVUE V Series node supports the following header transformations:

Table 9: 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 49: Opening the Link Quick View

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

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

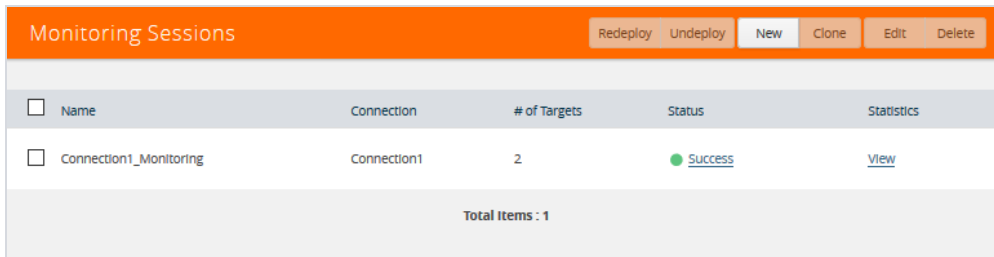
Figure 50: Adding Transformation

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

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.

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



The screenshot shows a web interface for 'Monitoring Sessions'. At the top, there is an orange header bar with the title 'Monitoring Sessions' and several action buttons: 'Redeploy', 'Undeploy', 'New', 'Clone', 'Edit', and 'Delete'. Below the header is a table with the following columns: 'Name', 'Connection', '# of Targets', 'Status', and 'Statistics'. There is one row of data with the following values: 'connection1_Monitoring' in the Name column, 'Connection1' in the Connection column, '2' in the # of Targets column, and 'Success' (with a green dot icon) in the Status column. The Statistics column contains a 'View' link. At the bottom of the table, it says 'Total Items : 1'.

<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.

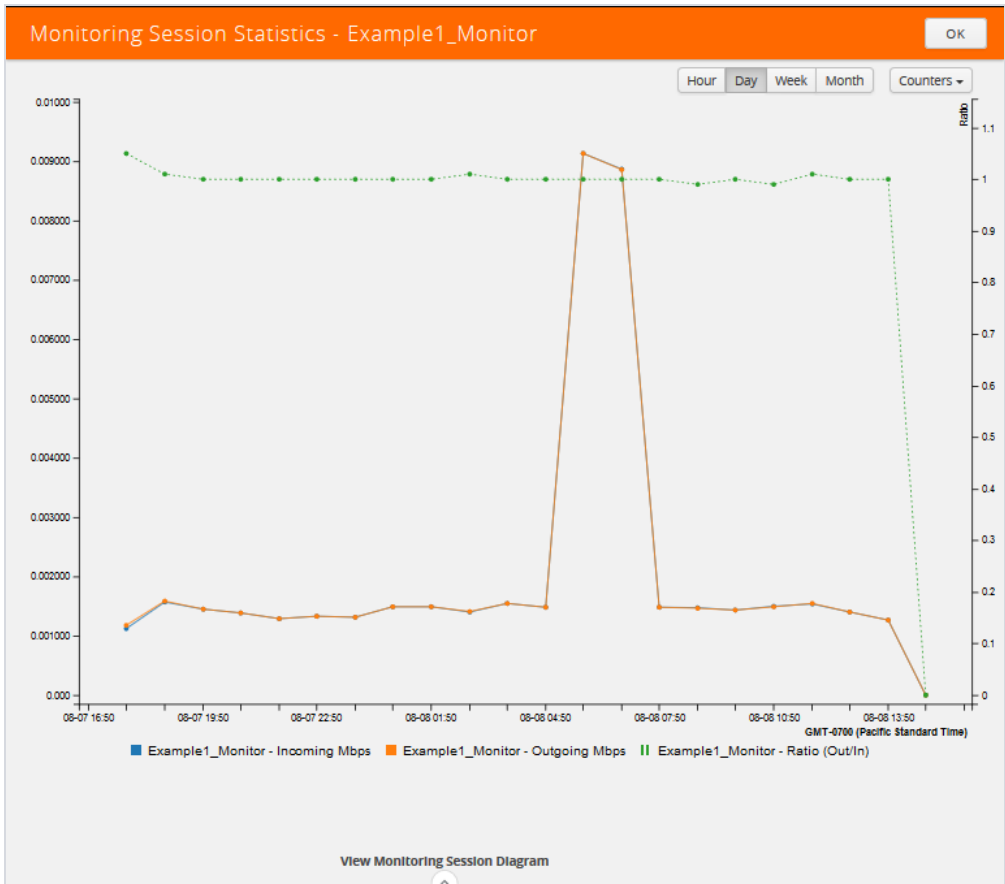
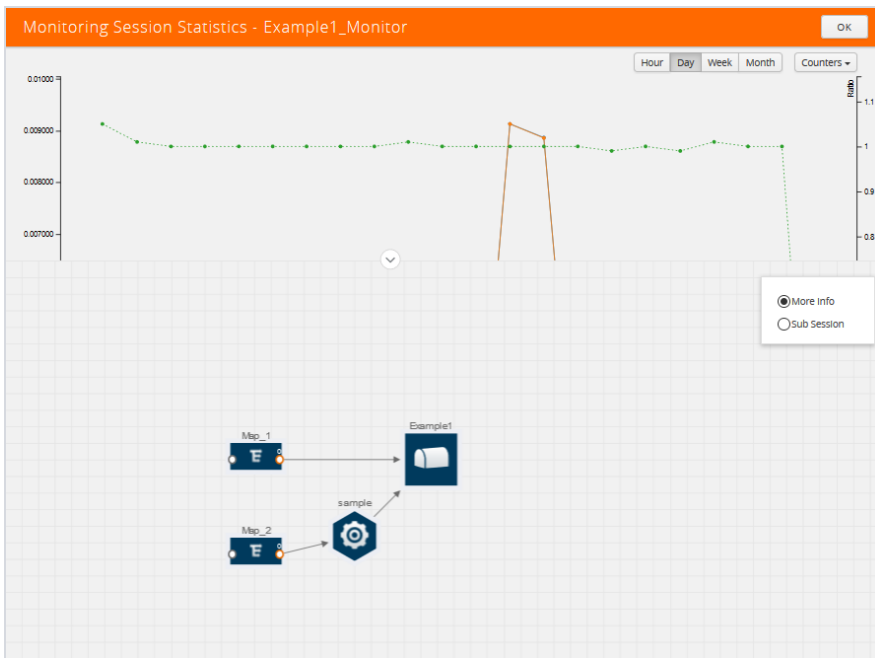


Figure 51: Monitoring Session Statistics View

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 52: Viewing the Map Statistics](#).

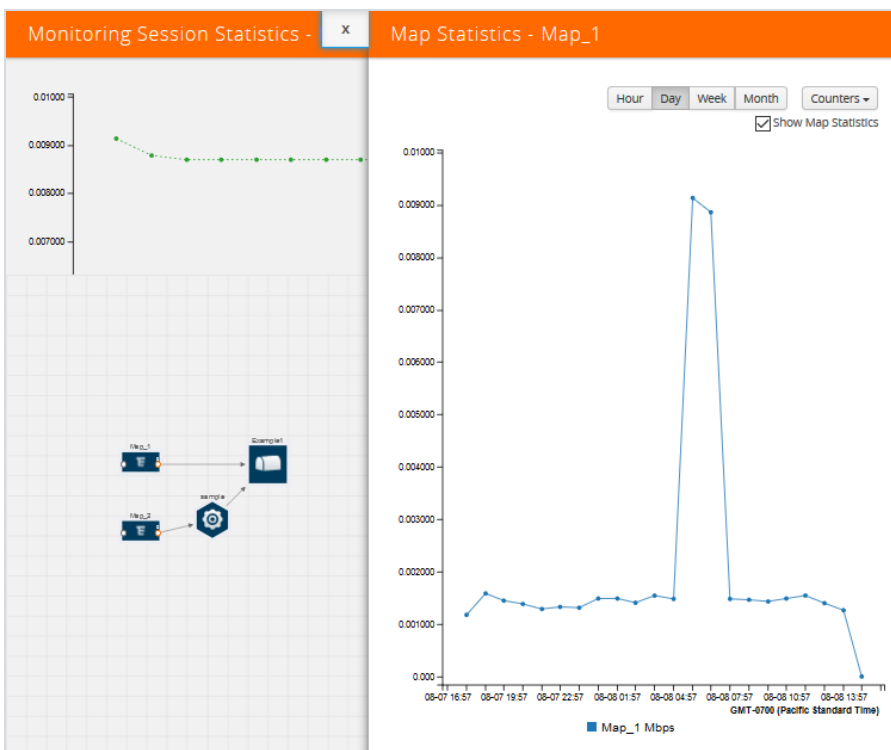


Figure 52: Viewing the Map Statistics

Figure 52: Viewing the [Map Statistics](#) shows the Map Statistics Quick View with a graph of the traffic for Map_1. You can also scroll down the Map Statistics Quick View to see the Map Rules, Action Sets, and Map Info for this map. You can select Map Rules or Action Sets to view the traffic matching the selected rule on the graph in the Quick View.

Map Rules			
RULE	PRIORITY	ACTION SET	CONDITIONS
<input checked="" type="checkbox"/> Rule 0	0	0	etherType 0x0800 IpProto 1
Action Sets			
<input type="checkbox"/> Action Set 0			
Map Info			
Map Alias Map_2			
Comment			

View 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 **Kubernetes > 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 53: 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 GigaVUE V Series Controllers, G-vTAP Controllers, monitoring tools, and targets that are being used in the connection.

Figure 54: 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 55: Filtering in Topology View](#) .

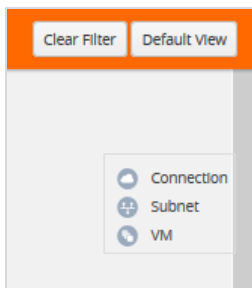
Figure 55: Filtering in Topology View

- 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.

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



Configure Kubernetes Settings

To configure the Kubernetes Settings:

1. In GigaVUE-FM, on the top navigation pane, select **Cloud**.
2. On the left navigation pane, select **Kubernetes > Settings > Advanced**.

Select **Advanced** to edit the Kubernetes settings. The **Settings** page appears. Refer to [Figure 56: Cloud >Kubernetes > Settings > Advanced..](#)




Figure 56: Cloud >Kubernetes > Settings > Advanced.

3. Click **Edit** to edit the Settings fields. Refer to [Table 10: Kubernetes Settings](#) for descriptions of the Settings fields:

Table 10: Kubernetes Settings

Settings	Description
Maximum number of connections allowed	Specifies the maximum number of connections you can establish in GigaVUE-FM.
Refresh interval for VM target selection inventory (secs)	Specifies the frequency for updating the state of Containers in Kubernetes.
Number of Kubernetes nodes per V Series Node	Specifies the number of Kubernetes nodes per GigaVUE V Series nodes.

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 1: 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 <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>

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