



GigaVUE Cloud Suite for Kubernetes Guide

GigaVUE Cloud Suite

Product Version: 5.13

Document Version: 1.0

(See Change Notes for document updates.)

Copyright 2021 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

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

Change Notes

When a document is updated, the document version number on the cover page will indicate a new version and will provide a link to this Change Notes table, which will describe the updates.

Product Version	Document Version	Date Updated	Change Notes
5.13.00	1.0	08/13/2021	Original release of this document with 5.13.00 GA.

Contents

GigaVUE Cloud Suite for Kubernetes Guide	1
Change Notes	3
Contents	4
GigaVUE Cloud Suite for Kubernetes	6
Audience	6
GigaVUE Cloud Suite for Kubernetes	7
About GigaVUE Cloud Suite for Kubernetes Container Visibility	7
Traffic Capturing Mechanism	7
GigaVUE Cloud Suite Cloud Components	9
Role Based Access Control	9
Configure Components in Kubernetes	11
Before You Begin	11
Minimum Compute Requirements	11
Security Group	12
Key Pairs	13
Launch GigaVUE-FM Instance	13
Configure GigaVUE Cloud Suite for Kubernetes	13
Configure Kubernetes Containers	14
Create Service Account Token	14
Configure Kubernetes	15
Configure G-vTAP Containers	16
Configure Monitoring Sessions in Kubernetes	18
Overview of Visibility Components	19
Create Monitoring Session	20
Create a New Session	21
Create Tunnel Endpoints	22
Create Map	22
Add Applications to Monitoring Session	25
Deploy Monitoring Session	39
Add Header Transformations	41
View Statistics	43
View Topology	44
Configure Kubernetes Settings	46

Additional Sources of Information	47
Documentation	47
How to Download Software and Release Notes from My Gigamon	49
Documentation Feedback	50
Contact Technical Support	51
Contact Sales	51
Premium Support	51
The Gigamon Community	51

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 Suite Cloud components and setting up the traffic monitoring sessions for the Kubernetes Containers visibility.

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

Topics:

- [Audience](#)
- [GigaVUE Cloud Suite for Kubernetes](#)
- [Role Based Access Control](#)
- [Configure Components in Kubernetes](#)
- [Configure Monitoring Sessions in Kubernetes](#)

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>

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](#)
- [GigaVUE Cloud Suite Cloud Components](#)
- [Traffic Capturing Mechanism](#)

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 Cloud Suite-TA Series node

Traffic from the GigaVUE V Series nodes/physical node is then sent to the required tools. The following figure shows a high level architecture of GigaVUE Cloud Suite for Kubernetes using G-vTAP containers as the source for acquiring the traffic.



GigaVUE Cloud Suite 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 Suite 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 Kubernetes:
 - 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 Cloud Suite® V Series Controller** manages multiple V Series nodes and orchestrates the flow of traffic from GigaVUE Cloud Suite V Series nodes to the monitoring tools. GigaVUE-FM uses one or more GigaVUE Cloud Suite V Series Controllers to communicate with the GigaVUE Cloud Suite 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 Suite Cloud using GRE or VXLAN tunnels
 - On-premise tools, such as a GigaVUE H Series device

Role Based Access Control

The Role Based Access Control (RBAC) feature controls the access privileges of users and restricts users from either modifying or viewing unauthorized data. Access privileges in GigaVUE Cloud Suite works on the same principles of access privileges in GigaVUE-FM in which the access rights of a user depends on the following:

- **User role:** A user role defines permission for users to perform any task or operation
- **User group:** A user group consists of a set of roles and set of tags associated with that group. When a user is created they can be associated with one or more groups.

To access the resources and to perform a specific operation in GigaVUE Cloud Suite you must be a user with **fm_super_admin** role or a user with write access to the following resource category depending on the task you need to perform.

Resource Category	Cloud Configuration Task
<p>Physical Device Infrastructure Management: This includes the following cloud infrastructure resources:</p> <ul style="list-style-type: none"> • Cloud Connections • Cloud Fabric Deployment • Cloud Configurations • Sys Dump • Syslog • Cloud licenses • Cloud Inventory 	<ul style="list-style-type: none"> • Configure GigaVUE Cloud Components • Create Monitoring Domain and Launch Visibility Fabric
<p>Traffic Control Management: This includes the following traffic control resources:</p> <ul style="list-style-type: none"> • Monitoring session • Stats • Map library • Tunnel library • Tools library • Inclusion/exclusion Maps 	<ul style="list-style-type: none"> • Create, Clone, and Deploy Monitoring Session • Add Applications to Monitoring Session • Create Maps • View Statistics • Create Tunnel End Points

NOTE: Cloud APIs are also RBAC enabled.

Refer to the *GigaVUE Administration Guide* for detailed information about Roles, Tags, User Groups.

Configure Components in Kubernetes

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

- [Before You Begin](#)
- [Launch GigaVUE-FM Instance](#)
- [Deploy G-vTAP Containers](#)
- [Configure GigaVUE Cloud Suite for Kubernetes](#)

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](#)
- [Key Pairs](#)

Minimum Compute Requirements

The minimum recommended computing requirements are listed in the following table.

Table 1: Minimum Compute Requirement

Compute Instances	vCPU	Memory	Disk Space	Description
GigaVUE Cloud Suite® 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
GigaVUE Cloud Suite® 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	2 vCPU	16GB	2 x 40GB	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 Cloud Suite V Series Controllers, GigaVUE Cloud Suite 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

Direction		Protocol	Port	CIDR	Purpose
GigaVUE-FM					
Inbound	HTTPS	TCP	443	Any IP address	Allows G-vTAP Container Managers, GigaVUE Cloud Suite V Series Controllers, and GigaVUE-FM administrators to communicate with GigaVUE-FM.
Inbound	IPv4	UDP	68	Any IP address	Allows GigaVUE-FM to communicate with DHCP server for assigning IP addresses and other related configuration information such as the subnet mask and default gateway.
Inbound	IPv4	UDP	53	Any IP address	Allows GigaVUE-FM to communicate with DNS server for resolving the host name of the cloud controller for Kubernetes.
GigaVUE Cloud Suite V Series Controller					
Inbound	IPv4	TCP	9902	GigaVUE-FM IP address	Allows GigaVUE-FM to communicate with GigaVUE Cloud Suite V Series Controllers.
GigaVUE Cloud Suite V Series <small>node</small>					
Inbound	Custom TCP Rule	TCP(6)	9903	GigaVUE Cloud Suite V Series Controller IP address	Allows GigaVUE Cloud Suite V Series Controllers to communicate with GigaVUE Cloud Suite V Series nodes.
GRE Traffic					

Direction		Protocol	Port	CIDR	Purpose
Inbound	Custom Protocol Rule	GRE (47)	All	Any IP address	Allows monitored traffic from GigaVUE Cloud Suite 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 Cloud Suite V Series nodes using the VXLAN tunnel. Allows monitored traffic from GigaVUE Cloud Suite 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 Kubernetes 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 Fabric Management Guide* for details on installing and launching GigaVUE-FM.

Configure GigaVUE Cloud Suite for Kubernetes

To configure the GigaVUE Cloud Suite for Kubernetes 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 Kubernetes
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 Kubernetes</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](#)
- [Configure Kubernetes](#)
- [Configure G-vTAP Containers](#)

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. `kubectl create serviceaccount <serviceaccountname>`
2. View the service account information and token name. `kubectl describe serviceaccounts <serviceaccountname>`
3. Get the token as an input to GigaVUE-FM. `kubectl describe secret <secret value in the above output>`
4. Create the cluster role binding with a (admin) role that has proper permissions. `kubectl create clusterrolebinding <clusterrolebindingname> --clusterrole=cluster-admin --serviceaccount=default:<serviceaccountname>`
5. Create a docker registry key with the input to the FM connection. `kubectl create secret docker-registry <docker-key-name> --docker-server=<your-registry-server> --docker-username=<your-name> --docker-password=<your-pword>`
6. Patch the service account with the key. `kubectl patch serviceaccount <serviceaccountname> -p '{"imagePullSecrets": [{"name": "<docker-key-name>"}]}'`
7. Display the API endpoint IP and port for GigaVUE-FM. `kubectl cluster-info`

Configure Kubernetes

To configure Kubernetes:

1. From the left navigation pane, select **Inventory > VIRTUAL > Kubernetes > Monitoring Domain**. The Monitoring Domain page appears.
2. Click **New**. The **Kubernetes Configuration** page appears.
3. Enter or select the following details:

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 Create Service Account Token .
Token	Service Account Token details.
API Server URL	API server URL of the Kubernetes Container
V Series Environment	GigaVUE V Series environment configured during Kubernetes Configuration. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>NOTE: If you intend to use a GigaVUE H Series box, then you must not configure the V Series environment.</p> </div>

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 lead node. The container images are provided through a registry.

NOTE: You can also fetch the image of the Gigamon Traffic Acquisition container from the [Gigamon 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 3: 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

Field	Description
	<ul style="list-style-type: none"> ■ 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 Cloud Suite 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](#)
- [Create Monitoring Session](#)
- [Configure Kubernetes Settings](#)



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 Cloud Suite 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.
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 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. In the following example, 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. 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.
Link	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 Overview of Visibility Components , the link originating from action set 0 is moving the traffic from MAP_1 to Monitoring_Tools. 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 Cloud Suite V Series node v1.2-1 and above. For more information about Header Transformation, refer to Add Header Transformations .
Group	A group is a collection of maps that are pre-defined and saved in the map library for reuse.
Application	An application performs operations such as sampling, slicing, and masking on the traffic.
Inclusion Map	An inclusion map determines the instances to be included for monitoring. This map is used only for target selection.
Exclusion Map	An exclusion map determines the instances to be excluded from monitoring. This map is used only for target selection.
Target	A target determines the instances that are to be monitored. Targets are determined based on the following formula:

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

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 a New Session](#)
- [Create Tunnel Endpoints](#)
- [Create Map](#)
- [Add Applications to Monitoring Session](#)
- [Deploy Monitoring Session](#)
- [View Statistics](#)
- [View Topology](#)

Create a New Session

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

To create a new session:

1. From the left navigation pane, select **Traffic > VIRTUAL > Orchestrated Flows > Kubernetes**. The Monitoring Session page appears.
2. Click **New**. The **Create A New Monitoring Session** window appears.

Create A New Monitoring Session

Alias	Alias
Monitoring Domain	Select domain...

3. Enter the appropriate information in the **Create a New Monitoring Session Info** dialog box as shown in the following table.

Table 2: Fields for Session Info

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

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 new tunnel endpoint:

1. After creating a new monitoring session, or click **Edit** on an existing monitoring session, the GigaVUE-FM canvas appears.
2. In the canvas, select **New > New Tunnel**, drag and drop a new tunnel template to the workspace. The **Add Tunnel Spec** quick view appears.
3. On the New Tunnel quick view, enter or select the required 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 IP interface.
Traffic Direction	The direction of the traffic flowing through the V Series node. Choose Out for creating a tunnel from the V Series node to the destination endpoint. NOTE: Traffic Direction In is not supported in the current release.
Remote Tunnel IP	The IP address of the tunnel destination endpoint. NOTE: You cannot create two tunnels from a V Series node to the same IP address.

4. Click **Save**.

To delete a tunnel, select the required tunnel and click **Delete**.

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	
EtherType	The packets are filtered based on the selected ethertype. The following conditions are displayed:

Conditions	Description
	<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 EtherType 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 EtherType and Pass All.

To create a new map:

1. In the Monitoring Session canvas, from **Maps** section, drag and drop a new map template to the workspace.
2. Click on the map, then click details. The map rules quick view is displayed.
3. Enter the appropriate information 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	The rules for filtering the traffic in the map. To add a map rule: <ol style="list-style-type: none"> a. Click Add a Rule. b. Select a condition from the Search L2 Conditions drop-down list and specify a value. Based on this selection, the Search L3 Conditions drop-down list is automatically updated. c. Select a condition from the Search L3 Conditions drop-down list and specify a value.
	<ol style="list-style-type: none"> d. (Optional) If you have selected TCP or UDP as the protocol in the L3 conditions, then select Port Source or Port Destination from the Search L4 Conditions drop-down list and specify a value. If you have selected conditions other than TCP or UDP, then the Search L4 Conditions drop-down list is disabled. e. (Optional) In the Priority and Action Set box, assign a priority and action set. f. (Optional) In the Rule Comment box, enter a comment for the rule. NOTE: Repeat steps b through f to add more conditions. Repeat steps a through f to add nested rules.

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

4. To reuse the map, click **Add to Library** Save the map using one of the following ways:
 - o Select an existing group from the **Select Group** list and click **Save**.
 - o 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.


5. Click **Save**.

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

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

Click on  to expand the **Targets** dialog box. Click on 

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



Add Applications to Monitoring Session

GigaVUE Cloud Suite with V Series 1 node supports the following GigaSMART applications:

- [Sampling](#)
- [Slicing](#)
- [Masking](#)
- [NetFlow](#)

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.



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



2. Click the Slice application and select **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.



2. Click the Mask application and select **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 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](#). 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](#).

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



The NetFlow record generation is performed on GigaVUE Cloud Suite V Series node running the NetFlow application. In [Add Applications to Monitoring Session](#), incoming packets from G-vTAP agents are sent to the GigaVUE Cloud Suite 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 \(NAT\)](#).

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:

	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	IPFIX

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

	Description	Supported NetFlow Versions
	header as a key field.	
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:

	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-	v9 and IPFIX


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

	Description	Supported NetFlow Versions
IPv6 Src IP	Configures the IPv6 source address in the IP packet header as a key field.	v9
Transport		
L4 Dest Port	Configures the destination port identifier in the transport header as a non-key field.	v9 and IPFIX
L4 Src Port	Configures the source port identifier in the transport header as a non-key field.	v9 and IPFIX
TCP Ack Number	Configures the acknowledgment number in the TCP header as a non-key field.	IPFIX
TCP Dest Port	Configures the destination port identifier in the TCP header as a non-key field.	IPFIX
TCP Flags	Configures the TCP control bits observed for the packets of this flow as a non-key field.	v9 and IPFIX
TCP Header Length	Configures the length of the TCP header as a non-key field.	IPFIX
TCP Seq Number	Configures the sequence number in the TCP header as a non-key field.	IPFIX
TCP Src Port	Configures the source port identifier in the TCP header as a non-key field.	IPFIX
TCP Urgent	Configures the urgent pointer in the TCP header as a non-key field.	IPFIX
TCP Window Size	Configures the window field in the TCP header as a non-key field.	IPFIX
UDP Dest Port	Configures the destination port identifier in the UDP header as a non-key field.	IPFIX
UDP Src Port	Configures the source port identifier in the UDP header as a non-key field.	IPFIX


Add Version 5 NetFlow Application

To add a version 5 NetFlow application:

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



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




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 more 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.



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



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


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 and Link NetFlow Application to NAT

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

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



2. 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.

X Link
Save

Alias:

Source type: Application

Destination type: Tunnel

Transformations:

IPv4 Destination ✕

10.2.2.23

Destination Port ✕

0 to 65535

3. Creating a Link from NetFlow to NAT
4. In the **Alias** field, enter a name for the link.

5. 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.

6. In **IPv4 Destination**, enter the IP address of the NetFlow collector.
7. (Optional) By default, the Destination Port is 2055. To change the destination port, enter a port number.
8. Click **Save**. The transformed link is displayed in Orange.
9. 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 Cloud Suite V Series nodes. Refer [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.
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.
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.



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.



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 Cloud Suite V Series node interface. For steps to configure the link, refer to [Add Applications to Monitoring Session](#).
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.

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](#).

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 that have been dragged and dropped to the workspace. The tunnel endpoints are named `Monitoring_Tool_1` and `Monitoring_Session_2`.



5. Hover your mouse on the map, click the red dot, and drag the arrow over to another map, application, or tunnel.

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. 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 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.

When you click on the Status link, the Deployment Report is displayed.

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

- o 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 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.



GigaVUE Cloud Suite 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.
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.



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

Alias:

Source type: **Map**

Destination type: **Tunnel**

Transformations:

VLAN Id ✕
Value


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.

On the Monitoring Sessions page, click **View** in the Statistics column to view the Monitoring Session Statistics page. The Monitoring Session Statistics page appears where you can analyze incoming and outgoing traffic.



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. The Map Statistics Quick View with a graph of the traffic for Map_1 is displayed. 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.


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. From the left navigation pane, select **Traffic > VIRTUAL > Orchestrated Flows > Kubernetes**. The Monitoring Session page appears..
2. In the Monitoring Session page, click the **Topology** tab.

3. Select a connection from the **Select connection...** drop-down list. The topology view of the subnets and instances is displayed.
4. (Optional) Select a monitoring session from the **Select Monitoring Session...** drop-down list. The monitored subnets and instances change to blue.
5. Select one of the following check boxes:
 - o **Source**—Displays the topology view of the source target interfaces that are being monitored.
 - o **Destination**—Displays the topology view of the destination target interfaces where the traffic is being mirrored.
 - o **Other**—Displays the topology view of the GigaVUE Cloud Suite V Series Controllers, G-vTAP Controllers, monitoring tools, and targets that are being used in the connection.



6. (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.

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. From the left navigation pane, select **Inventory > VIRTUAL > Kubernetes > Settings**. The Settings page appears.
2. Click **Edit** to edit the Settings fields. Refer to the following table for descriptions of the Settings fields:

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

This table lists all the guides provided for GigaVUE Cloud Suite software and hardware. The first row provides an All-Documents Zip file that contains all the guides in the set for the release.

NOTE: In the online documentation, view [What's New](#) to access quick links to topics for each of the new features in this Release; view [Documentation Downloads](#) to download all PDFs.

Table 1: Documentation Set for Gigamon Products

GigaVUE Cloud Suite 5.13 Hardware and Software Guides
<p>DID YOU KNOW? If you keep all PDFs for a 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.</p>
<p>Hardware</p> <p>how to unpack, assemble, rack-mount, connect, and initially configure ports the respective GigaVUE Cloud Suite devices; reference information and specifications for the respective GigaVUE Cloud Suite devices</p>
<p>*G-TAP A Series 2 Installation Guide</p>
<p>GigaVUE-HC1 Hardware Installation Guide</p>
<p>GigaVUE-HC2 Hardware Installation Guide</p>
<p>GigaVUE-HC3 Hardware Installation Guide</p>
<p>GigaVUE M Series Hardware Installation Guide</p>
<p>GigaVUE TA Series Hardware Installation Guide</p>
<p>*GigaVUE-OS Installation Guide for DELL S4112F-ON</p>
<p>Software Installation and Upgrade Guides</p>
<p>GigaVUE-FM Installation, Migration, and Upgrade Guide</p>
<p>GigaVUE-OS Upgrade Guide</p>

GigaVUE Cloud Suite 5.13 Hardware and Software Guides

Administration

GigaVUE Administration Guide

covers both GigaVUE-OS and GigaVUE-FM

Fabric Management

GigaVUE Fabric Management Guide

how to install, deploy, and operate GigaVUE-FM; how to configure GigaSMART operations; covers both GigaVUE-FM and GigaVUE-OS features

Cloud Configuration and Monitoring

how to configure the GigaVUE Cloud Suite components and set up traffic monitoring sessions for the cloud platforms

GigaVUE Cloud Suite for AnyCloud Guide

how to deploy the GigaVUE Cloud Suite solution in any cloud platform

GigaVUE Cloud Suite for AWS Guide

GigaVUE Cloud Suite for AWS Quick Start Guide

quick view of AWS deployment

GigaVUE Cloud Suite for AWS SecretRegions Guide

GigaVUE Cloud Suite for Azure Guide

GigaVUE Cloud Suite for Kubernetes Guide

GigaVUE Cloud Suite for Nutanix Guide

GigaVUE Cloud Suite for OpenStack Guide

GigaVUE Cloud Suite for VMware—GigaVUE-VM Guide

GigaVUE Cloud Suite for VMware—GigaVUE V Series Guide

GigaVUE Cloud Suite for VMware—GigaVUE-VM Guide

Reference

GigaVUE-OS CLI Reference Guide

library of GigaVUE-OS CLI (Command Line Interface) commands used to configure and operate GigaVUE H Series and TA Series devices

GigaVUE-OS Cabling Quick Reference Guide

guidelines for the different types of cables used to connect Gigamon devices

GigaVUE-OS Compatibility and Interoperability Matrix

compatibility information and interoperability requirements for Gigamon devices

GigaVUE-FM REST API Reference in GigaVUE-FM User's Guide

samples uses of the GigaVUE-FM Application Program Interfaces (APIs)

GigaVUE Cloud Suite 5.13 Hardware and Software Guides

Release Notes

GigaVUE-OS, GigaVUE-FM, GigaVUE-VM, G-TAP A Series, and GigaVUE Cloud Suite Release Notes

new features, resolved issues, and known issues in this release ;
important notes regarding installing and upgrading to this release

NOTE: Release Notes are not included in the online documentation.

NOTE: Registered Customers can log in to [My Gigamon](#) to download the Software and Release Notes from the Software & Docs page on to [My Gigamon](#). Refer to [How to Download Software and Release Notes from My Gigamon](#).

In-Product Help

GigaVUE-FM Online Help

how to install, deploy, and operate GigaVUE-FM.

GigaVUE-OS H-VUE Online Help

provides links the online documentation.

How to Download Software and Release Notes from My Gigamon

Registered Customers can download software and corresponding Release Notes documents from the **Software & Release Notes** page on to [My Gigamon](#). Use the My Gigamon Software & Docs page to download:

- Gigamon Software installation and upgrade images,
- Release Notes for Gigamon Software, or
- Older versions of PDFs (pre-v5.7).

To download release-specific software, release notes, or older PDFs:

1. Log in to [My Gigamon](#)
2. Click on the **Software & Release Notes** link.
3. Use the **Product** and **Release** filters to find documentation for the current release. For example, select Product: "GigaVUE-FM" and Release: "5.6," enter "pdf" in the search box, and then click **GO** to view all PDF documentation for GigaVUE-FM 5.6.xx.

NOTE: My Gigamon is available to registered customers only. Newer documentation PDFs, with the exception of release notes, are all available through the publicly available online documentation.

Documentation Feedback

We are continuously improving our documentation to make it more accessible while maintaining accuracy and ease of use. Your feedback helps us to improve. To provide feedback and report issues in our documentation, send an email to: documentationfeedback@gigamon.com

Please provide the following information in the email to help us identify and resolve the issue. Copy and paste this form into your email, complete it as able, and send. We will respond as soon as possible.

Documentation Feedback Form		
About You	Your Name	
	Your Role	
	Your Company	
For Online Topics	Online doc link	<i>(URL for where the issue is)</i>
	Topic Heading	<i>(if it's a long topic, please provide the heading of the section where the issue is)</i>
For PDF Topics	Document Title	<i>(shown on the cover page or in page header)</i>
	Product Version	<i>(shown on the cover page)</i>
	Document Version	<i>(shown on the cover page)</i>
	Chapter Heading	<i>(shown in footer)</i>
	PDF page #	<i>(shown in footer)</i>
How can we improve?	Describe the issue	<i>Describe the error or issue in the documentation. (If it helps, attach an image to show the issue.)</i>
	How can we improve the content? Be as specific as possible.	
	Any other comments?	

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.